TRANSFORMING THE MUSEUM-COMMUNITY NEXUS WITH TECHNOLOGY

A VIRTUAL MUSEUM INFRASTRUCTURE FOR PARTICIPATORY ENGAGEMENT AND MANAGEMENT

Adeola Ezekiel Fabola

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

2018

Transforming the Museum-Community Nexus with Technology

A Virtual Museum Infrastructure for Participatory Engagement and Management

Adeola Ezekiel Fabola

University of St Andrews

This thesis is submitted in partial fulfilment for the degree of Doctor of Philosophy (PhD) at the University of St Andrews

March 2018

Abstract

Museums play an important role in society as the custodians of heritage, and advances in technology have brought about opportunities for curating, preserving and disseminating heritage through virtual museums. However, this is not matched by an understanding of how these technologies can support these functions, especially given the varying levels of resources that museums have at their disposal. To address this problem, a hybrid methodology which combines underpinning theory and practice has been adopted. Initial investigation of the problem takes place through a contextualisation of museology and heritage studies, followed by exploratory case studies that yield design objectives for a Virtual Museum Infrastructure (VMI). A design of the VMI is proposed based on these objectives, and the VMI is instantiated, deployed and evaluated in real-world scenarios using a combination of quantitative and qualitative techniques. The findings of this investigation demonstrate that the use of technology provides new opportunities for engagement with heritage, as experts and community members alike can create, curate and preserve content, which can then be disseminated in engaging ways using immersive, yet affordable technologies.

This work therefore demonstrates how technology can be used to: (1) support museums in the creation, curation, preservation and dissemination of heritage, through a VMI that provides support for all the stages of the media life cycle, (2) facilitate active use, so that content that is created once can be reused on multiple platforms (for example on the web, on mobile apps and in on-site installations), and (3) encourage connectivity by linking up local museums using a location-aware interface and facilitates the consumption content using digital literacies available to the public. The aforementioned points, coupled with the system instantiations that demonstrate them, represent the contributions of this thesis.

Acknowledgements

To:

- The School of Computer Science at the University of St Andrews for the tuition-free scholarship and stipends that made this research endeavour possible.
- Dr Alan Miller and Dr Uta Hinrichs for unending guidance, direction and encouragement in their roles as primary and secondary supervisors respectively.
- Sarah Kennedy at the Open Virtual Worlds Group for the modelling of the 14th-century St Andrews Cathedral used to build the mobile virtual reality system in Chapter 4 and the curing yard used to build the museum installation in Chapter 5.
- The staff and pupils at Madras College, St Andrews for hosting the classroom exercises in Chapter 4.
- Dr Jo Clements at Timespan for conducting preliminary research on the content and structure of the 19th-century curing yard featured in Chapter 5.
- Audrey Parsons and Angela Aitchison at Timespan and Highlife Highland respectively for organising the event that facilitated the user evaluation of the real-time virtual tour system featured in Chapter 6.
- Dr Iain Oliver at the Open Virtual Worlds Group for setting up the back end data store of the Virtual Museum Infrastructure (VMI), and – together with Dr Karen Brown, Dr Alan Miller and Catherine Cassidy at the University of St Andrews – facilitating the Three Dimensional (3D) workshops in the EU-LAC project featured in Chapters 7 and 8.
- Mark Simmons and the staff at Perth Museum and Art Gallery for curating and hosting the museum exhibit featured in Chapter 8.
- Other (present and past) members – Dr Bess Rhodes, Hussein Bakri, Dr John McCaffery and Dr C.J. Davies – of the Open Virtual Worlds group for their technical, academic and moral support.
- The staff and students at the School of Computer Science and the University at large for the conducive atmosphere for this research, and everyone outwith the University that provided me with an outlet for extra-curricular activities.
- My family and friends for the love, support and, well, everything.

You are very kind, thank you!

Declaration

Candidate's Declaration

I, Adeola Ezekiel Fabola, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 79,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I was admitted as a research student at the University of St Andrews in January 2015. I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date Signature of candidate

Supervisor's Declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date Signature of supervisor

Permission for Publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Adeola Ezekiel Fabola, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date Signature of candidate

Date Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's Declaration

I, Adeola Ezekiel Fabola, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date Signature of candidate

Permission for publication of underpinning research data or digital outputs

We understand that for any original research data or digital outputs which are deposited, we are giving permission for them to be made available for use in accordance with the requirements of the University and research funders, for the time being in force.

We also understand that the title and the description will be published, and that the underpinning research data or digital outputs will be electronically accessible for use in accordance with the license specified at the point of deposit, unless exempt by award of an embargo as requested below.

The following is an agreed request by candidate and supervisor regarding the publication of underpinning research data or digital outputs:

No embargo on underpinning research data or digital outputs.

Date Signature of candidate

Date Signature of supervisor

A Note on Web Resources

Due to the relatively young age of concepts discussed in this document, some resources are only available on the World Wide Web. Such resources have been cited where peer-reviewed, published material were not available. The Universal Resource Identifiers (URIs) have been shortened for brevity using Google's shortening service (https://goo.gl/) and as at the time of submission, all the URIs were verified. However, their longevity cannot be guaranteed.

A Note on Collaboration

This project has been largely collaborative. The design and implementation contributions of this thesis are mine. I have produced five main systems which are discussed in chapters 4 – 8. Some content which make up these systems were created by members of the Open Virtual Worlds (OVW) research group of which I am part. The 3D reconstructions used in chapters 4, 5 and 8 are the work of Sarah Kennedy. The server and data store used in the EU-LAC project is managed by Dr Iain Oliver, and the 3D workshops featured were facilitated by the EU-LAC project, with involvement from Dr Alan Miller, Dr Karen Brown, Dr Iain Oliver and Catherine Cassidy. The evaluation exercises were conducted by me with support from members of the OVW group, and the staff at the museums and schools involved in the evaluation. The data collection, processing and analysis were done by me.

Publications

The content of the chapters in this thesis have been published in conference proceedings and journals:

1. Chapter 4 – **Adeola Fabola**, Alan Miller, and Richard Fawcett. *Exploring the past with Google Cardboard*. In Digital Heritage, 2015, volume 1, pages 277–284. IEEE, 2015. [1].
2. Chapter 4 – **Adeola Fabola** and Alan Miller. *Virtual Reality for Early Education: A Study*. In International Conference on Immersive Learning, pages 59–72. Springer, 2016. [2].
3. Chapter 5 – **Adeola Fabola**, Sarah Kennedy, Alan Miller, Iain Oliver, John McCaffery, Catherine Cassidy, Jo Clements, and Anna Vermehren. *A Virtual Museum Installation for Time Travel*. In International Conference on Immersive Learning, pages 255–270. Springer, 2017. [3].
4. Chapter 6 – **Adeola Fabola**, Alan Miller, Ishbel Duncan. *Aerial Virtual Reality: Remote Tourism with Drones*, presented at the 7th European Immersive Education Summit, Lucca & Pisa, Italy. 2017. [4].
5. Chapters 7 and 8 – Catherine Cassidy, **Adeola Fabola**, and Alan Miller. *A Digital Museum Infrastructure for Preserving Community Collections from Climate Change*. In Workshop, Long and Short Paper, and Poster Proceedings from the Third Immersive Learning Research Network Conference, pages 170–177. iLRN, 2017. [5].
6. Chapter 8 – Catherine Cassidy, **Adeola Fabola**, Elizabeth Rhodes, and Alan Miller. *The Making and Evaluation of Picts and Pixels: Mixed Exhibiting in the Real and the Unreal*. In Proceedings from the Fourth Immersive Learning Research Network Conference, pages 97–112. iLRN, 2018. [6].

To Bola, Bola, Bola and Bola.

CONTENTS

List of Figures

LIST OF TABLES

LISTINGS

ACRONYMS

Abbreviation	Meaning
2D	Two Dimensional
3D	Three Dimensional
AJAX	Asynchronous Javascript and XML
ANOVA	Analysis of Variance
API	Application Programmable Interface
AR	Augmented Reality
AU	Access Unit
AV	Augmented Virtuality
CAD	Computer Assisted Design
CAVE	Computer Automated Virtual Environment
DAMS	Digital Asset Management System
DBMS	Database Management System
FoV	Field of View
FPV	First Person View
GPS	Global Positioning System
HD	High Definition
HMD	Head Mounted Display
JSON	Javascript Object Notation
MR	Mixed Reality
MTU	Maximum Transmission Unit
NAL Unit	Network Abstraction Layer Unit
NPC	Non-Player Character
ODV	OmniDirectional Video
QoE	Quality of Experience
RDBMS	Relational Database Management System
ReST	Representational State Transfer
SDK	Software Development Kit
SFM	Structure From Motion
UAV	Unmanned Aerial Vehicle
UDP	User Datagram Protocol
VMI	Virtual Museum Infrastructure
VR	Virtual Reality
VRML	Virtual Reality Modelling Language
WebGL	Web Graphics Library
WebVR	Web Virtual Reality
Wi-Fi	Wireless Fidelity
XR	Cross Reality

GLOSSARY

EU-LAC A multinational project involving community museums in Europe, Latin America and the Caribbean, funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 693669

H.264 A video compression standard that specifies the conversion of video content for dissemination purposes using an encoder and the conversion of the content into a playable format using a decoder

Part I

Prolegomenon

Chapter One

Introduction

"The past needs to be remembered. If you do not know where you come from, then you don't know where you are, and if you don't know where you are, then you don't know where you're going." – Terry Pratchett [7, p.187]

1.1 Overview

This thesis investigates how to use emergent technologies to design a novel virtual museum system for participatory heritage engagement. The investigation is carried out through the design and implementation of a Virtual Museum Infrastructure (VMI) – a suite of tools that facilitate curation, preservation and dissemination of heritage. The case for the VMI is made through a combination of underpinning theory, physical-world practice and case studies which feature on-site, remote and real-time exploration of heritage sites. The development of the VMI is facilitated by advances in mobile and immersive technologies. A practice-orientated methodology is adopted for the research, where use cases are investigated in turn to build up to the iterative design, implementation and evaluation of the VMI.

The intersection of museums and technology necessitates the exploration of the concept of virtual museums, a concept which is relatively new and varied in meaning. This discussion therefore leads with the concept of a *museum*, its meaning and functions as perceived by practitioners. A discourse on museums naturally leads to the concept of *heritage* as the focus and subject matter of museums, and highlights the categories and instruments for their representation in museums, hence the concept of *digital heritage*[1] is introduced. The idea of museums as a *community* or participatory platform is also introduced as a key concept, such that the museum as an institution provides a platform for its audience to engage with the heritage management process while the audience consider the museum to be its representative voice [8]. The concepts of *museums*, *heritage*, *digital heritage* and *community* are therefore discussed in the following sections.

[1] *Digital heritage*, in the context of this thesis, refers to the application of technology in the heritage domain, rather than the heritage, or history of technology.

1.1.1 Museums

The International Council of Museums (ICOM) defines museums as:

> *"...a non-profit, permanent institution in the service of society and its development, open to the public, which acquires, conserves, researches, communicates and exhibits the tangible and intangible heritage of humanity and its environment for the purposes of education, study and enjoyment."* [9]

This definition is the current standard amongst museum practitioners hence its components are worth examining:

Non-Profit Museums are not commercial entities. This suggests that their sole purpose is not to make profit and would hence rely on external (usually government or public) funding streams and surpluses generated in-house from fundraising activities. The non-profit status is typically bestowed on museums (and heritage organisations) by governments by virtue of the recognition of museums as entities that provide services that are beneficial to the public. This non-profit status exempts museums from taxes that are paid by other commercial entities, but also places restrictions on the financial activities of the museums. Furthermore, the non-profit status does not preclude museums from generating excess funds (i.e. income which exceeds operating costs) from their day-to-day operations; rather, they are compelled to redirect any surpluses into their operations as opposed to allocating these surpluses to external benefactors (in the form of stocks, shares or dividends) [10]. This examination of the non-profit status of museums is important because it dispels any misconceptions that heritage organisations are not expected to generate financial surpluses, but rather stresses that any surpluses generated should be channeled towards the organisations activities.

Permanent Museums are established in place. This component of the definition is subject to interpretation based on the meaning of "permanence". It may refer to the existence of a building infrastructure within which the organisation operates. It may also refer to the standing of the organisation as an entity whose existence is not bound or limited to a time period.

Public Service Museums serve and develop society. They promote understanding of the past and present states of communities, thus fostering togetherness in society.

Functions Museums *acquire*, *conserve*, *research*, *communicate* and *exhibit* heritage. They execute the stages involved in the life cycle, beginning with acquisition and culminating in the dissemination of heritage content.

Focus Museums' focus should be on tangible and intangible heritage of humanity and its environment. Their activities should revolve around the study of humankind – culture, history, practices – as well as other lifeforms, nature and the environment.

Purpose Museums' purposes revolve around education, study and enjoyment. They are a place where students can learn, where research can be conducted and findings can be made. In addition, they can serve recreational purposes as a place for travellers to visit on holiday or for local families to visit on weekends.

There are variations in the terminology and views on the functions of museums. The ICOM definition [9] suggests that the functions of a museum are to acquire, conserve, research, communicate and exhibit heritage. Noble [11] states that the functions of museums are to collect, conserve, study, interpret and exhibit. Mensch [12] condenses the functions of museums proposed by Noble [11] into three: to preserve (collect and conserve), study and communicate (interpret and exhibit). This thesis adopts the functions of museums proposed by Mensch [12] due to its brevity, with modifications to the terminology. The functions of museums are henceforth referred to as the *preservation*, *curation* and *dissemination* of heritage. *Preservation* is used in line with Mensch's definition [12] (and in lieu of conservation as used by ICOM [9] and Noble [11]) to encompass the acquisition and storage of heritage resources and associated content to ensure availability and continuity in the event that these resources are lost or damaged. *Curation* encompasses the processes of researching [9] and studying [11, 12] heritage content. *Dissemination* encompasses the processes of interpreting [11, 12], communicating [9, 12] and exhibiting [9, 11, 12] tangible and intangible heritage.

1.1.2 Heritage

Heritage is the focus and subject matter of museums [9, 11, 12]. The United Nations Educational, Scientific and Cultural Organisation (UNESCO) defines cultural heritage as

> *"... the legacy of physical artefacts and intangible attributes of a group or society that are inherited from past generations, maintained in the present and bestowed for the benefit of future generations [13]..."*

and considers natural heritage as, amongst others

> *"... features consisting of physical and biological formations or groups of such formations, which are of outstanding universal value from the aesthetic or scientific point of view... [14]."*

Heritage can be categorised based on subject matter to yield cultural heritage and natural heritage. Cultural heritage refers to attributes that relate to the history and way of life of a group, community or society. Examples of cultural heritage include paintings, sculpture, oral histories and folk tales. Natural heritage refers to elements that occur and/or exist endemically in a region without human intervention. Examples include the selection of lifeforms (animals and vegetation), landscapes and geological structures that have been formed as a result of a natural phenomenon. Cultural heritage can be further categorised into tangible and intangible cultural heritage. Tangible cultural heritage refers to physical items such as buildings and landscapes that have historical value, while intangible cultural heritage refers to attributes passed down from one generation to the next, which could take the form of beliefs, knowledge or traditions [13, 15]. Tangible heritage can also be moveable or immovable. Examples of moveable heritage include artefacts, figurines and documents. Examples of immovable heritage include landscapes and cityscapes of a cultural significance. While cultural heritage can exist as either tangible or intangible heritage, natural heritage is predominantly tangible. Examples of tangible natural heritage include fossils or remains of animal specimens (such as dinosaurs) and significant landscapes that have been formed through natural phenomena (such as the Grand Canyon), which represent moveable and immovable tangible heritage respectively. It is pertinent to note that the classification of heritage items as tangible or intangible may depend on the context. For instance, a scroll may be classified as intangible heritage if its importance lies in its content, but may also be classified as tangible if its material or composition is of historic significance. These various forms of heritage require variations in metadata to describe them. For instance, size may be expressed in terms of length, breadth and height for moveable, tangible heritage, but for immovable, tangible heritage and intangible heritage, the concept of size may not apply. These variations should be taken into consideration in designing metadata for the VMI. A diagram illustrating the classifications of heritage is shown in Fig. 1.1.

Heritage plays a major role in the society as it has the potential to enrich the lives of individuals, communities and larger societies. It can foster a sense of belonging of a people; it has the ability to stimulate the local economy by promoting tourism and travel which translate to increased business in the hospitality industry and greatly benefits hotels, cafes and gift shops. It also facilitates the creation and continuity of jobs through the initialisation of site preservation projects, and the establishment of new hotels, cafes, museums and heritage organisations [16, 17]. In a similar vein, museums are important because they are vital in educating the public, promoting unity and togetherness, fostering a sense of belonging in society, and facilitating intergenerational heritage transfer [18, 19]. Museums play an important role in educating society members on history, heritage and a diverse range of topics including science, art and fashion. The preservation and dissemination of heritage as performed by museums, ensures that this knowledge can be

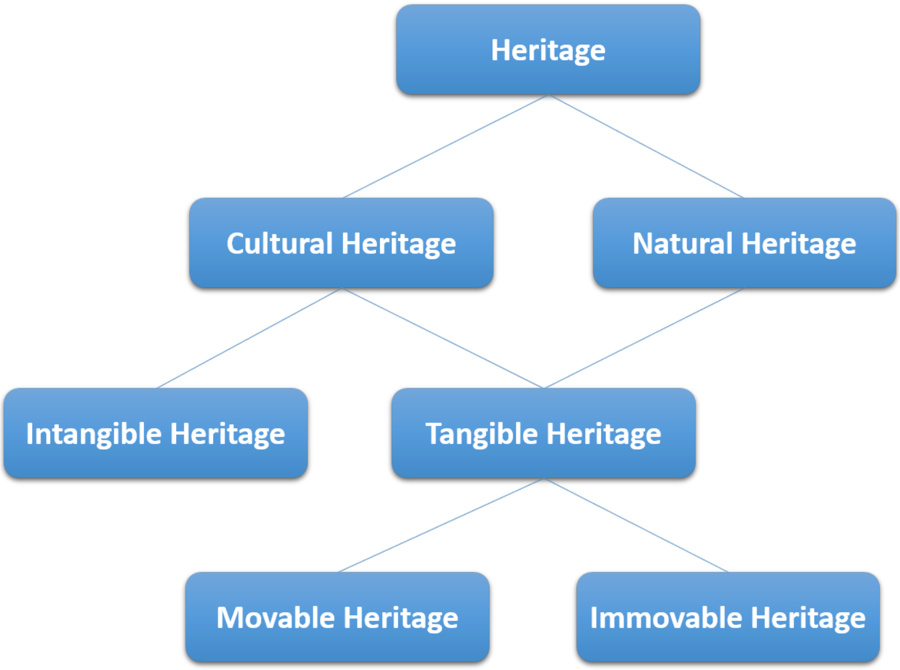

Figure 1.1: Classifications of heritage

passed down from one generation to the next which minimises the extinction of knowledge of the past. Knowledge of the past is important because in the words of Terry Pratchett [7, p.187], *"...if you do not know where you come from...then you don't know where you're going."*

1.1.3 Digital Heritage

In the absence of technology and digital media, intangible heritage can be preserved and presented in the form of written, drawn or carved text, images and symbols on paper, scrolls or stones which may be put on display in museums. The representation and preservation of tangible heritage can also be actualised through writings, drawings and carvings. This includes text and paintings that over time become iconic and display-worthy in museums. In addition, artefacts, documents, preserved animal specimens and other moveable tangible heritage can be displayed on plinths and in show cases, while immovable heritage such as landscapes and cityscapes can be modelled out of wood, cardboard paper or styrofoam and then displayed in museums like moveable heritage.

The concept of digital heritage embodies the use of technology for the representation and interpretation of tangible and intangible cultural and natural heritage. Intangible heritage lends itself to representation using multimedia such as audio and video recordings or in written form using electronic text. Tangible heritage can be represented using a wide range of digital media. Two Dimensional (2D) images and video can be used to capture both immovable and moveable heritage of a cultural or natural kind. For instance, photographs can be taken of landscapes, paintings and animal specimens. Advances in computer graphics facilitate the depiction of remote *scenes* using spherical imagery and the representation of *objects* using Three Dimensional (3D) imagery, as discussed in section 2.2. These have the ability to immerse viewers in a scene by representing that scene as a spherical image, or provide an improved sense of engagement with an artefact by rendering a 3D representation of that artefact using computer graphics. The use of spherical images and 3D objects therefore adds an additional dimension in the digital representation of heritage, because in addition to using text, images, audio and video for capturing intangible heritage, landscapes, cityscapes and geological structures can be captured using spherical images, while artefacts, specimens, documents and sculptures can be captured using 3D objects.

The public domain is rife with discussions and classifications of media types. However, these discussions revolve around the use of media as a tool for mass communication and information dissemination, with platforms such as the television, radio, newspapers, magazines and more recently the Internet. The discussion of digital media in this thesis takes a step back from mass communication media and focuses on electronic components of media that can be leveraged to deliver interactive and engaging experiences to an audience. Attempts have been made to identify and categorise these digital media, notably by Li, Drew and Liu [20, p. 4], who identify multimedia modalities as text, images, drawings, graphics, animations, video, sound and interaction, and Banerji and Ghosh [21, p. 9], who identify multimedia components as digital text, graphics, audio, video, animation. Digital media are therefore categorised as follows:

Text Text is used in line with Li, Drew and Liu [20, p. 4] and Banerji and Ghosh [21, p. 9]. In the context of digital media, text refers to the use of electronic documents that describe subject matter in a natural language. This may take the form of documents generated using word processors, digital scans of hard copy documents, textual components of web pages and so on.

Audio Audio is used in line with Banerji and Ghosh [21, p. 9] and in lieu of sound used by Li, Drew and Liu [20, p. 4]. Audio is used consistently as an alternative to sound in the context of digital media transmission, to refer to sound that is generated and/or transmitted by digital means, typically as electronic recordings of subject matter such as narratives, interviews, performances and natural occurrences. They typically take the form of a seamless sequence of sound waves of varying frequencies that can be perceived by the human ear.

Digital Media	Appeals to audio-visual senses	Leaves things to the imagination	Provides descriptive detail	Allows open-ended interpretation
Text	× Text relies on the ability to read and comprehend subject matter, and thus does not appeal to audio-visual senses.	✓ Text forces the reader to imagine/visualise concepts in the mind.	✓ Text allows for the provision of detail using narrative accounts, anecdotes and vicarious experiences.	✓ Text may be given to multiple (sometimes diverse) meanings depending on interpretation.
Audio	✓ Audio adds an extra dimension (over text) to content in form of intonations, ambience and other forms of expression that may be lost with written text.	×/✓ Audio content may not leave a lot to the imagination because of its descriptive nature. In cases, audio may compel the audience to picture the content in the mind.	✓ Audio excels at providing multi-dimensional detail to augment the core content being communicated.	×/✓ Audio content may not be open to interpretation where unmistakable detail is abundant. In the absence of detail, however, audio content may yield multiple meanings.
2D images (Mono/Stereo)	✓ 2D images add a powerful visual element to content. They may depict environments, events and other forms of expression that may be lost with written text.	×/✓ 2D images may not leave a lot to the imagination because of their descriptive nature. In cases, 2D images may compel the audience to imagine accompanying audio content such as ambient sounds, expressions and conversational content.	✓ 2D images excel at providing multi-dimensional detail to augment the core content being communicated.	× 2D images may not be open to interpretation where unmistakable detail is abundant.
3D images (digital objects)	✓ 3D images add a powerful visual element to content. They typically depict and enable interaction with artefacts and may thus reveal details that may be lost with other media.	× 3D images may not leave a lot to the imagination because of their descriptive and interactive nature.	✓ 3D images excel at depicting content in high levels of detail and often facilitate interaction that can augment the core content being communicated.	× 3D images may not be open to interpretation where unmistakable detail is abundant.
Spherical images	✓ Spherical images add a powerful visual element to content. They typically depict scenes in immersive fashion and may thus reveal details that may be lost with other media.	× Spherical images may not leave a lot to the imagination because of their descriptive and immersive nature.	✓ Spherical images excel at depicting content in high levels of detail and often immerse the audience in the scene which may result in realistic experiences with content.	× Spherical images may not be open to interpretation where unmistakable detail is abundant.
Moving images (Video)	✓ Moving images add a powerful visual element to content. They may depict dynamic environments, events and other forms of expression that may be lost with static media.	× Moving images may not leave a lot to the imagination because of their descriptive and immersive nature.	✓ Moving images excel at depicting content in high levels of detail and they may augment the core content being communicated using dynamic elements that depict actions, events and environments.	× Moving images may not be open to interpretation where unmistakable detail is abundant.
3D Scenes (Environments)	✓ 3D scenes, like moving images, add visual and dynamic elements to content, and take this experience further by giving the audience interactive control over the content being depicted.	× 3D scenes may not leave a lot to the imagination because of their descriptive and immersive nature.	✓ 3D scenes excel at depicting content in high levels of detail and they may augment the core content being communicated using dynamic and interactive elements that depict actions, events and environments.	× 3D scenes may not be open to interpretation where unmistakable detail is abundant.

Table 1.1: An overview of digital modes of heritage interpretation

2D images 2D images refers to all forms of "flat" imagery including electronic images and drawings in line with Li, Drew and Liu [20, p. 4] and graphics in line with Li, Drew and

Liu [20, p. 4] and Banerji and Ghosh [21, p. 9]. 2D in this context is used to refer to the flat component of these images and to distinguish them from other types of interactive visual media. In the context of digital media, 2D images refer to the use of electronic, flat pictorial representations of subject matter such as scenes, events and natural occurrences. 2D images may be monoscopic or stereoscopic, where an image is composed of one or two viewpoints respectively.

3D images 3D images refers to components of interaction (in line with Li, Drew and Liu [20, p. 4]) and animation (in line with Li, Drew and Liu [20, p. 4] and Banerji and Ghosh [21, p. 9]). 3D in this context represents the perceptible depth of images in contrast to the flatness of 2D images. In the context of digital media, 3D images refer to electronic formats for providing representations of physical objects in ways that allow interaction and viewing from multiple perspectives as opposed to the limited perspectives provided by 2D images.

Spherical images Spherical images refers to components of interaction (in line with Li, Drew and Liu [20, p. 4]) and animation (in line with Li, Drew and Liu [20, p. 4] and Banerji and Ghosh [21, p. 9]). Spherical is used to refer to the enclosing nature of images in contrast to the flatness of 2D images. In the context of digital media, spherical images refer to pictorial representations of subject matter that provide the illusion of place by situating the audience at the centre of the scene being represented. Like 2D images, spherical images may also be monoscopic or stereoscopic, but unlike 2D images, and like 3D images, spherical images can provide multiple perspectives of a subject matter, the difference being that spherical images provide an egocentric representation (where the audience is at the centre) while 3D provide an allocentric representation (where the subject matter is at the centre).

Moving images Moving images is used in lieu of video as used by Li, Drew and Liu [20, p. 4] and Banerji and Ghosh [21, p. 9]). "Moving" in this context refers to the fast-changing sequence of (2D, 3D or spherical) images, and thus highlights the relationship between this media component and other image types. In the context of digital media, moving images, also known as videos, refer to the use of rapidly-changing pictorial representations of subject matter. Moving images may be monoscopic or stereoscopic, 2D, 3D or spherical, or some combination of the above (for example stereoscopic, spherical videos or monoscopic, 2D videos).

3D scenes 3D scenes refers to components of interaction (in line with Li, Drew and Liu [20, p. 4]) and animation (in line with Li, Drew and Liu [20, p. 4] and Banerji and Ghosh [21, p. 9]). 3D scenes refer to complex environments that are perceived to have depth as compared to the flat environments depicted by 2D images. 3D scenes may be comprised of 3D images and spherical images, and typically possess both interactive and animated components. In the context of digital media, 3D scenes are a more advanced form of 3D images which can be used to represent larger and more complex environments as opposed to individual objects and artefacts. 3D scenes can provide both an illusion of place (like spherical images) and an illusion of depth (like 3D images) and can also be presented in monoscopic or stereoscopic format.

These digital media can be classified based on how they are intrinsically perceived by audiences, how the audience's perception of them may change over time and space, and how they represent different content types. This yields classification by perception, changeability, and content.

Digital Media – Classification by Perception

The advent of digital media provides new opportunities for *realism*, such that work can be recorded and reproduced at higher fidelity as compared to erstwhile media, and *virtualism*, such that new work can be created that, although may be fictional, are reminiscent of reality [22]. Realism is achieved by providing descriptive, high-fidelity to the user's audio-visual senses, while virtualism is achieved by stimulating the user's imagination. These provide the basis for the classification of digital media based on their abilities to deliver realism and virtualism, as perceived by users. Table 1.1 therefore discusses the characteristics of the aforementioned digital media in terms of their ability to (1) appeal to audio-visual senses by presenting multi-faceted content that can be visualised and heard, (2) leave things to the imagination in the absence of concrete audio-visual representations, (3) provide descriptive detail through rich and varied content, and (4) allow open-ended interpretation in the absence of descriptive detail. The table shows the diversity in the nature of these digital media, and given the strengths and weaknesses of each, it is profitable to combine multiple media for heritage dissemination and interpretation. This is the essence of the multimedia integration and content reuse design goals discussed in section 8.2.

While these media can be classified based on how they are presented to and perceived by audiences, they can also be classified based on their nature. It is pertinent to note that the nature of each media typically informs how the media is perceived. Nonetheless, the characteristics of these media are worth exploring in the context of changeability and content.

Digital Media – Classification by Changeability

Steinmetz and Nahrstedt [23, p. 181], Parekh et al. [24, p. 2], and Raisinghani [25, p. 74] allude to the classification of digital media as either static (such as text and images) or dynamic (such as audio, video and animations). Static media typically remain unchanged to the audience while dynamic media changes from the perspective of the audience. Parekh et al. [24, p. 2] also discuss the concept of linearity (vs non-linearity) and interactivity (vs non-interactivity) as applied to dynamic media. Linearity refers to the property of media that compels the audience to interact with it in a given sequence (for instance a sequence of sound waves or images). By virtue of this, audio and moving images are linear (dynamic) media while spherical images, 3D images and 3D scenes are non-linear (dynamic) media. The difference between these categories is that linear media follow a pre-defined sequence while non-linear media do not. The sequence of non-linear media can be determined by user interaction. For instance, a user can push a button to jump to a portion of content. This makes non-linear (dynamic) media intrinsically interactive, as they give the audience the ability to initiate actions and alter the sequence of events to their satisfaction. The concepts of linearity and interactivity are thus discussed as temporal and spatial dichotomies when applied to dynamic media. For instance, audio is considered temporally dynamic because it changes linearly (i.e. over time, both forwards and backwards in some cases) but, with exceptions, it may not be spatially dynamic because the audience is confined to a pre-defined sequence. In a similar vein, 3D scenes are considered spatially dynamic because the audience determines the order in which they interact with content, but they are typically not temporally dynamic because of their non-linear nature. Spherical moving images can be both temporally dynamic (i.e. linear, given the moving image component) and spatially dynamic (i.e. interactive, given the intrinsic nature that requires audience input). This static vs dynamic (with linearity and/or interactivity) property is referred to as changeability.

Digital media can therefore be classified in terms of their ability to change over time (temporally) and space (spatially). Temporal changeability refers to a variation in the audience's perception of the media content with time. Spatial changeability refers to a variation in the audience's perception of the media content when consumed from different perspectives or positions. Some media remain unchanged temporally and spatially and are therefore static. Examples include text and 2D images. Conversely, some media may change either temporally, spatially or both temporally and spatially; these are dynamic media. Examples include audio, moving images, 3D images and 3D scenes.

Dynamic media may change temporally, spatially or both. Audio content intrinsically changes temporally, given that it is made up of a seamless sequence of sound waves. In addition to temporal changeability, audio may change spatially as well. For instance, stereo sound clips may be perceived differently in the left and right ears, speakers or parts of a sound system. Similarly, an audience may perceive the volume and composition of surround (or 3D) sound to vary over a given space. By virtue of this classification, mono audio possesses temporal changeability while stereo and surround audio possess both temporal and spatial changeability.

Like audio, moving images intrinsically change temporally, given that they are made up of a sequence of (2D, 3D or spherical) images. Moving, 2D images typically possess only temporal changeability, except in the case of moving, stereoscopic 2D images where the audience's perception of the content (in this case the illusion of depth) may vary as a result of the position from which the content is viewed thus resulting in both temporal and spatial changeability. Moving spherical images typically possess both temporal (intrinsically) and spatial changeability given the ability of the media to place the audience at the centre of the scene being depicted.

3D images, or interactive digital objects typically possess spatial changeability given that the audience's perception of the content may vary as interaction takes place through zooming, panning, clicking and dragging. Temporal changeability is not an intrinsic feature of 3D images, but is possible nonetheless. For instance, a 3D image can be transformed (or morphed) into another over a period (of say 30 seconds) on a loop, such that the audience can click, drag, pan and zoom to view the content from multiple perspectives as the transformation takes place; this encompasses both temporal and spatial changeability.

Given that 3D scenes are a more advanced form of 3D images and also lend themselves to zooming, panning, clicking and dragging actions which may alter the audience's perception of the content, they typically possess spatial changeability. In addition to spatial changeability, temporal changeability is possible – albeit not intrinsic – through transformations, animations and/or morphs.

Table 1.2 summarises the changeability of digital media when each is used in isolation. The characteristics may vary when combined with other media. For instance, 2D images by themselves are neither temporally dynamic nor spatially dynamic. However, when multiple, sequential 2D images are combined into a moving 2D image, they become temporally dynamic. Similarly, spherical images are typically spatially dynamic and not temporally dynamic, but when multiple, sequential spherical images are combined into a moving spherical image, they become both spatially and temporally dynamic.

Digital Media – Classification by Content

Kalay, Kvan and Affleck [22] discuss the use of new media to record, represent, reconstruct and/or illustrate already existing objects and events, which refers to the *reproductive* property of new media, but also to create or interpret content that may result in the creation of new work, which refers to the *productive* property of new media. The discussion of the reproductive and productive properties of new media form the basis for a classification based on the content they represent. The productive property, when exercised, usually results in the use of media for artistic

Media	Static	Dynamic	
		Temporal	Spatial
Text	✓	×	×
Audio	×	✓	✓
2D images	✓	×	×
3D images	×	×	✓
Spherical images	×	×	✓
Moving images	×	✓	✓
3D scenes	×	×	✓

Table 1.2: Classification of digital media by changeability

expression, while the reproductive property results in the use of media for concrete or abstract illustrations when exercised.

Artistic expressions are created when digital media are used as a creative means. Text can be used for poetry, audio can be used for poetry, music and melodies, images (whether 2D, 3D or spherical) may be in the form of paintings or photographs, and other types of dynamic imagery such as moving images (videos) and 3D scenes may represent movies, silent films and scripted plays.

Concrete illustrations are used to depict cultural practices, historic events or aspects of daily life that are considered important and worth preserving. Text and audio can be used as forms of commentaries or narratives in visual or auditory form respectively. Static imagery may be used to illustrate aspects of daily life such as dance, communal meals and festivities. Similarly, dynamic imagery may be used to document natural phenomena (such as natural heritage) or re-enact a historical event or practice.

Abstract illustrations are used to depict concepts, beliefs or creatures that are rather intangible. This may take the form of textual descriptions of imaginary creatures, concoctions of the sounds made by these imaginary creatures, or visualisations through static sketches or dynamic representations.

Table 1.3 summarises the classification of digital media based on the content they represent. These media may be used in isolation or combinations to represent these content types. For instance, dynamic imagery (such as moving images) may be used in isolation to make silent

Media	Artistic Expression	Concrete Illustration	Abstract Illustration
Text	Poetry	Commentaries, narratives	Textual descriptions of imaginary creatures
Audio	Music, melodies	Commentaries, narratives	Voicings of imaginary creatures
Static Imagery	Paintings, photographs	Dance representations	Sketches of imaginary creatures
Dynamic Imagery	Movies, films, scripted plays	Re-enactments of historic events	Interactive, digital models of imaginary creatures

Table 1.3: Classification of digital media by content type

films as a means of artistic expression, but may also be combined with audio to make movies and films with sound tracks, also as a means of artistic expression.

Digital Media – Combinations

While each of the aforementioned media can be used in isolation, they can also be combined for the dissemination of heritage. For the purpose of this discussion, the image categories have been combined and will thus be discussed as static images (such as 2D images) and dynamic images (such as moving images and 3D images).

Text: Audio content is perceived through the sense of hearing and is thus impractical for users with limited or no hearing. The addition of textual content can provide an alternative means of content consumption for these users. Furthermore, the addition of text to audio content can provide an alternative (and sometimes preferred) means of sensory input even for users with adequate hearing.

Given that images of both static (such as 2D and spherical images) and dynamic (such as moving images and 3D scenes) nature may not foster open-ended interpretation, textual content can be added to these to provide more information and interpretation, to serve as an educative resource.

Audio: Like the addition of audio to textual content, audio content can be used to augment textual content to provide an alternative for users with visual impairments, and to provide alternative sensory input for information.

In addition, given that both static and dynamic images appeal to visual senses, the addition of descriptive audio content can be used to provide an alternative for users with visual impairments. Audio can also be added to moving images and 3D scenes to provide ambient effects that add an extra dimension to the user experience.

Static and Dynamic Images: Given that textual content may not appeal to visual senses, static and dynamic images can be added to text to visualise content.

Similarly, images can be added to audio content to provide an alternative for users with auditory impairments. Furthermore, dynamic images can be added to add an extra dimension to the user experience.

This discussion has hitherto highlighted the varied nature of digital media and classified them based on their characteristics, their perception by the audience, their changeability over time and space, and the types of content they can represent. It is worth noting that while these media types can be (and have been) used in isolation for heritage dissemination, they can be combined in ways that leverage the strengths and characteristics of each media type to achieve a predefined goal. Table 1.4 highlights the benefits that accrue from combining pairs of digital media, but these combinations can be taken further to include more than two media types, depending on project goals. The next section therefore explores how digital heritage components (and their combinations) can be presented using technology platforms.

Digital Media – Presentation

Advances in networking, processing and graphics capabilities of computers have brought about platforms for heritage presentation characterised by high interactivity and wide reach. These include websites, mobile apps, on-site installations in museums and social archive sites. These platforms have varying degrees of cost, reach and interaction potential, but all excel in their combinations of digital media for the dissemination of heritage content. Foni et al. [26] provide a discussion of cultural heritage visualisation strategies. While a more detailed discussion of this taxonomy is provided in section 2.5, this section, through a cursory examination of the visualisation strategies, identifies common platforms that form the basis for content presentation, as outlined as follows:

Websites Texts are the fundamental building blocks of web pages, which make up websites. Audio, 2D images and moving images can also be embedded traditionally in web pages. In addition, the advent of the Web 2.0 and HTML5 have allowed the embedding of other dynamic images such as spherical images and 3D scenes that users can interact with. Websites are accessible over the Internet which makes them useful for widespread dissemination of heritage content, but this could also be a drawback in cases where network access is unreliable or unavailable.

Mobile Apps Apps are containers that are designed for use on mobile devices and serve a similar function to web pages with the added benefits of: 1) being accessible "on the go" and 2) being accessible offline where pre-packaged content is provided. A limitation of mobile apps is the amount of computing resources available on mobile devices to render content,

Can be added to	Text	Audio	Static Images	Dynamic Images
Text	-------------	To provide an alternative for users with auditory impairments, to provide alternative sensory input for information.	To provide more information and interpretation, to serve as an educative resource or reference.	To provide more information and interpretation, to serve as an educative resource or reference.
Audio	To provide an alternative for users with visual impairments, to provide alternative sensory input for information.	-------------	To provide an alternative for users with visual impairments.	To provide an alternative for users with visual impairments, to provide ambient effects which may improve the experience.
Static Images	To provide interpretive and visual content.	To provide an alternative for users with auditory impairments.	-------------	-------------
Dynamic Images	To provide interpretive and visual content.	To provide an alternative for users with auditory impairments, to augment the experience.	-------------	-------------

Table 1.4: Combinations of digital media

as compared to the typical amount of computing resources available on desktop-based systems.

Desktop-Based Installations Digital media can be disseminated through standalone installations on desktop-based systems. These can be used to provide canned content and typically provide substantial computing resources for rendering high fidelity content such as 3D scenes, 3D images, spherical images and videos. Where network access is available and reliable, desktop-based installations can also be used to consume content that is hosted on remote servers over the Internet.

Social Archives Social archive sites are hosted and accessed over the Internet, like websites, but differ in the sense that they provide services for hosting and storing resources in addition to presenting resources to users. They may also serve as a forum that encourages the discussion of content.

These platforms can be made to deliver more compelling, interactive and immersive experiences using Virtual Reality (VR) technologies. Headsets like the Oculus Rift [27] and HTC Vive [28] can be tethered to desktop systems to turn a screen-based standalone installation into an immersive one. Web applications can also be designed to output content to these desktop-based headsets using the WebVR standards [29]. Similarly, mobile applications can also be designed to use mobile VR headsets for improved immersion and interactivity. In addition to headsets, Computer Automated Virtual Environment (CAVE) and VR domes can be used for large, immersive displays that groups of users can interact with. The distinction between immersive and non-immersive systems for heritage exploration forms the basis of the technology taxonomy discussed in section 2.5.2.

Furthermore, the characteristics of these delivery platforms give rise to varied settings and scenarios in which they can be used. For instance, in the context of location of exploration, these platforms allow for on-site access, remote access or a combination of both. On-site access refers to exploring heritage content at the actual site that the content depicts, while remote access refers to exploring heritage content of a location at which the user is not physically present. Examples of on-site access include exploring the remains of a heritage site or undertaking a guided tour of a site. Examples of remote access include undertaking a virtual tour of a site over the Internet or exploring heritage content in a museum, where the museum's location is not the same as (or at best is loosely-tied) to the location of the content being explored. This is in line with Bruno et al. [30] who identify remote and on-site delivery of experiences in the context of virtual museums. The variations between on-site and remote heritage exploration scenarios forms the basis of the use case taxonomy discussed in section 2.5.3. In addition, the nuances in these exploration scenarios necessitate the investigation of how the aforementioned delivery platforms can be used to deliver the varied digital media components, and part of this investigation is carried out using case studies as introduced in section 1.4.

1.1.4 Community

"Community" may refer to a geographical area, tribe or ethnicity. It may also evoke feelings of identity with a group or affinity for a social practice. Hence in the broadest sense, a community museum is a type of museum which is concerned with the heritage of a tribe, ethnic group, geographical area or a selection of people and/or places which identify or are affiliated with a common theme. While this applies to museums that may or may not identify as community museums, a defining characteristic of a community museum is the relationship between the museum and its community. This refers to a two-way relationship, such that the museum regularly solicits input from the community members on matters concerning projects, exhibitions and research, while the community looks to the museum as a representative voice on its heritage [8]. This definition alludes to the significant variation in the nature and practice of community museums, due to the multiple connotations of the term "community". This variation is made even more pronounced by the varying extents to which heritage organisations consider themselves to be community museums. A heritage organisation's identification as a community museum may be viewed as a question of *scale*, *focus* and *participation*; specifically, the *scale* of the operation, the *focus* of the operation and the extent to which *participation* is embodied in the operation.

Scale The *scale* of operation may be determined by the size of the target audience, or how much of the operation is centred around a geographical area (a village, city or country for instance). The smaller the area of operation, the more likely a heritage organisation will identify as a community museum.

Focus The *focus* of operation may be determined by the subject that the organisation specialises in; for instance, military history, natural history or the heritage of a local vicinity. A heritage organisation whose focus of operation is centred around the heritage of the locals is more likely to identify as a community museum.

Participation The extent to which *participation* is embodied in the operation may be determined by the philosophy of the museum stakeholders with respect to the role of the museum in the society. Cameron [8] identifies two schools of thought in this regard: the museum as a *temple* and the museum as a *forum*. The museum as a *temple* is the traditional school of thought where the museum is ascribed the highest authority as the principal custodian of heritage, such that its visitors are expected to conduct themselves in a solemn manner while passively consuming heritage content. The museum as a *forum*, on the other hand, is an emerging movement which views the museum as a platform for visitors to contribute, discuss, share, learn and actively engage with heritage. Cameron [8] expressed that these two ideas should work together instead of one replacing the other. By so doing, the museum as a *temple* should inspire trust and faith in the community, which enables the members to use it as a *forum* to contribute, express their views, and tackle issues that are pertinent to society. The greater the emphasis on the view of the museum as a *forum*, the greater the emphasis on visitor participation and in turn the greater the identification as a community museum.

By virtue of their *scale*, *focus* and *participation*, community museums tend to be small- to medium-sized, non-profit organisations which are founded, funded and/or run mainly by a group of volunteers who identify as part of the community in focus. This makes community museums important because they perform the functions of a museum (as outlined in section 1.1.1) on a smaller scale, which makes them effective in fostering a sense of togetherness in society. Irrespective of scale, and whether a museum does or does not identify as a community museum, as long as the museum encourages some kind of participation – through contributing, collaborating, co-creating or hosting [31] – with their audience, the concepts discussed in this thesis largely apply.

1.2 Motivation

Museums face challenges such as the shortage (or altogether absence) of operating funds, resources and infrastructure [32, 33, 34]. The paucity of formal infrastructure inhibits proper curation and preservation of heritage content, and the shortage of funds and resources (which may include skilled staff) inhibits the development of interactive exhibits for public engagement [32, 33]. Furthermore, there is a pressing need to preserve heritage content in museums, precipitated by the prospect of loss and degradation of content through climate change and natural disasters [35, 36, 37, 38]. Given the significant variation in the nature of community museums (see section 1.1.4), it follows that there is a significant variation in the level of resources available to organisations that identify as community museums, thus the extent of the paucity of resources available to community museums is unclear. This phenomenon is visualised as a spectrum of resources (see Fig. 1.2), where the leftmost point corresponds to a museum with the barest minimum level of resources while the rightmost point corresponds to a museum with the absolute maximum level of resources. It is difficult to assign quantifiable values to the leftmost (barest minimum) and rightmost (absolute maximum) points on the spectrum, hence these two points are theoretical and are for illustrative purposes. In practice however, each museum will fall somewhere on this spectrum, such that given a finite set of museums, each museum can be assigned a relative position on the spectrum by ranking them based on the level of resources available. The relative positions of museums on the spectrum may change depending on the resources considered, the metrics used to quantify these resources and the weights assigned to each resource category. Nonetheless, the spread of museums on this resource spectrum – as a result of the variation in the level of resources available – poses a major challenge that emergent technologies can address, through design and subsequent implementation that caters to the needs of organisations along the spectrum.

Virtual museums bring about global reach for, new forms of engagement with, and preservation of, heritage content. A discussion of virtual museums is provided in chapter 2, but in this context the term refers to the use of technology to achieve the goals of curating, preserving, interpreting and disseminating heritage, as performed by traditional museums. A prerequisite for

Figure 1.2: Spectrum of resources available to museums

the virtual museums in question is a focus on participatory engagement in small- and medium-sized organisations, such that heritage experts and community members alike can contribute towards the processes of curating, preserving and disseminating heritage. Furthermore, the proliferation of mobile, web and social media technologies at the dawn of the new millennium ushered the world into a new era in which participation, collaboration and engagement can be achieved on a global scale [31]. By using virtual museums to achieve global reach, direct connections between cultures can be strengthened, which in turn fosters inter-cultural acceptance and understanding. In addition, the use of technology for virtual museums leverages digital literacies[2] and readily-available resources for heritage curation, preservation and dissemination, so that at each point on the spectrum, solutions can be provided that utilise the resources available for optimum engagement. For instance, the use of smartphones and game-like systems in heritage dissemination offers a new appeal to the younger generations and allows for inter-generational heritage transfer. Furthermore, the use of virtual museum technologies to digitise heritage content provides a form of preservation in the event that the original artefact is lost or damaged as a result of climate change or natural disasters. Advances in mobile communications and computer graphics provide opportunities for the development of such systems and thus make it possible to now conduct this research, where it was previously infeasible. To demonstrate how these technological advances have now made this research possible, and relevant, it is pertinent to look at computing trends over the last seven decades. Moore's law states that the number of transistors in a single silicon chip is expected to double every year [40] (this was later revised to "double every two years" [41]). This projection has held up for five decades since it was put forward in 1965. Although it is unlikely that the literal implications of the law (in terms of growth in the number of transistors per chip) will hold up for another decade [42], the law still holds in the sense that computing aspects – computer graphics, power, communication and parallelisation –

[2] Knobel [39, pp.166–167] defines digital literacy as *" the awareness, attitude and ability of individuals to appropriately use digital tools and facilities to identify, access, manage, integrate, evaluate, analyse and synthesize digital resources, construct new knowledge, create media expressions, and communicate with others, in the context of specific life situations, in order to enable constructive social action; and to reflect upon this process."* This thesis uses the term to refer to skills and resources – such as gaming proficiency, smartphone ownership, basic knowledge of software, web browsers, the Internet, and so on – that people have come to possess by virtue of the proliferation of technology in aspects of daily life.

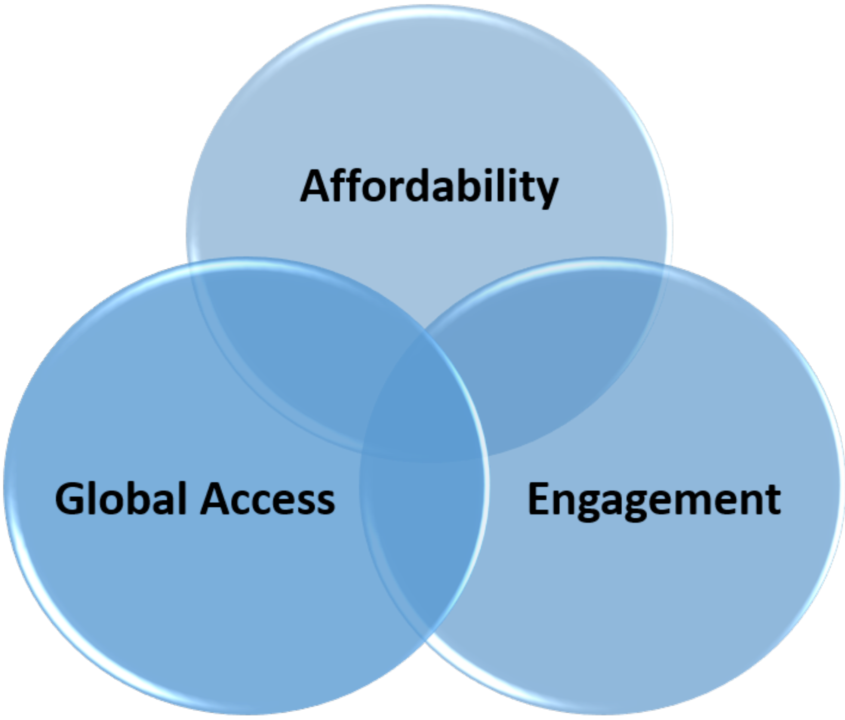

Figure 1.3: Three facets of participatory virtual museums

are advancing at an exponential rate. One implication of these advancements is that computing devices have become smaller, cheaper and faster, and this trend is applicable to mobile, ubiquitous and pervasive computing as well. Mobile devices have higher processing power, higher storage capacities, higher screen resolutions, longer-lasting batteries and miniaturised sizes compared to erstwhile decades. Miniaturisation of computing devices has also brought about the emergence of flight technologies (embodied in portable, affordable drones) made possible by advancements in aerodynamics and availability of affordable components such as portable accelerometers, gyroscopes and GPS technology. Over the last seven decades, attempts have been made to build systems around these technologies. HMD-based VR systems date back to the late 1960s when Ivan Sutherland proposed the idea of the Ultimate Display [43] and created the *Sword of Damocles* [44]. Similar systems (perhaps more appropriately referred to as ancestors of the technology) have been around since before the 1960s.

VR slowly emerged as a result of the advances in media technologies (such as the television) that precede it, hence it is difficult to ascertain exactly when VR was born [45]. Nonetheless, the principles upon which modern VR headsets operate can be traced back to the design of the stereoscope which was invented by Charles Wheatstone in 1838 [46], and was followed by the lenticular stereoscope (invented by David Brewster in 1849 [46]) and the View-Master (invented by William Gruber in 1938 [47]). These systems were followed by precursors to the HMD-style systems of today – the Telesphere Mask which enabled users to view 3D films without any interaction or motion-tracking, a HMD which featured stereoscopy, motion-tracking and connected to input feed from closed circuit cameras, and the Sword of Damocles headset by Ivan Sutherland, all developed in the late 1950s and 1960s [48, 49, 50]. A brief history of HMD-based systems, VR, sensoramas and immersive systems is presented by Sherman and Craig [48] and Craig, Sherman and Will [50], showing how immersive technologies have matured over time. The last five years have seen the release of development versions of Virtual Reality headsets such as the Oculus Rift [27], HTC Vive [28], PlayStation VR (previously known as Project Morpheus [51, 52]) and HoloLens [53], and eventually consumer versions of the Oculus Rift [54, 55] and HTC Vive [56], all of which can be powered by gaming PCs and/or consoles. In the same timespan, consumer versions of the Samsung Gear VR [57] and the Google Cardboard [58, 59] have been released (and more are available [60]), thus enabling the consumption of VR content on mobile phones as opposed to being tethered to computers or game consoles. Web technologies have also evolved from the static web pages of the 1990s to dynamic systems that facilitate the visualisation of geographical content using map-based interfaces as well as interacting with 3D models using WebGL and engaging with immersive content using WebVR [29]. In a similar vein, Unmanned Aerial Vehicles (UAVs), imbued with autonomous navigation and self-stabilising flight capabilities are becoming increasingly popular for recreational and commercial use as exemplified by prototypes, proofs of concept and awarded patents which demonstrate huge potential [61, 62, 63, 64], while portable, affordable computing devices such as the Raspberry Pi have come to the fore [65]. These trends allude to three areas of significant development in web, mobile and immersive technologies that make their potential value realisable: improved *processing*, *graphics* and *communication* capabilities. The improved processing power, graphics and communication capabilities of mobile technologies enable the development of affordable, yet interactive systems for the curation, preservation and dissemination of heritage. This research therefore focuses on the use of affordable, immersive (mobile and web-based) systems for local and global heritage engagement, hence it can be viewed in light of *affordability*, *engagement* and *global access*. *Affordability* is actualised through the use of commodity devices, *engagement* is achieved through the use of immersive devices to provide novel experiences as well as through management interfaces that supports active creation and curation of content, and *global access* – which is facilitated by the proliferation of mobile, ubiquitous computing devices, coupled with an increase in the public's proficiency with these devices – is attained through an architecture that facilitates remote storage, access and dissemination of heritage content. The intersection of these three aspects (Fig. 1.3) provides a unique contribution to the development of virtual museums and to the broader field of digital heritage.

1.3 Thesis Statement

The thesis of this dissertation is that technological advances and concomitant upsurges in digital literacies facilitate the development of a novel Virtual Museum Infrastructure (VMI) that enables museums to actively and continuously curate, preserve and disseminate local heritage on a global scale. The pertinent points in this thesis statement are highlighted below:

Technological Advances As a result of the implication of Moore's law [40], computing devices are becoming more powerful, more portable and more affordable, thus they are more readily-available for use by the public. This has resulted in an increase in digital literacies as demonstrated through widespread smartphone ownership and familiarity with games technologies [66, 67].

Virtual Museums Technology-based virtual museums have been around since the advent of the World Wide Web, but as a consequence of technological advances, specifically in the aspects of improved graphics, mobile devices and the 3DWeb, a novel approach to virtual museums development can be attained by combining multiple traditional media with immersive systems and presenting this content using commodity technology. The provision of immersive and engaging, yet affordable ways for deploying virtual museum systems represents a contribution to the state of the art.

Active Use The Virtual Museum Infrastructure (VMI) encourages heritage experts and community members alike to actively create and update content (and describe these content using metadata). This tailors the system for participatory engagement, such that a collaborative approach can be adopted for the curation, preservation and dissemination of heritage, thus strengthening the nexus, or relationships between museums and their respective communities.

Global and Local Access In addition to supporting on-site functions, remote use cases are supported, such that people can engage with content over the Internet using mobile and web-based systems. This extends local heritage content and makes it available remotely, and is particularly useful if a person is planning a trip to a museum (in anticipation), if they have visited the museum in the past (in reminiscence), or if they are unable to travel to an actual museum (in replacement).

1.4 Research Questions

The combination of the 3DWeb, mobile VR headsets, portable, yet affordable personal computers and consumer drone technology offers new opportunities for heritage exploration, but this is not matched by an understanding of how to leverage the potential value of these systems. Research needs to be carried out to **investigate how these systems can support the functions of traditional museums (Q1, Process)**, to **investigate how to utilise existing resources for heritage dissemination and management (Q2, Infrastructure)**, and to **investigate how these systems bring about change in heritage practice (Q3, Transformation)**. This gives rise to the following questions:

Q1 How can emergent technologies support (and replicate) the functions – heritage curation, preservation and dissemination – of traditional museums?

Q2 How can heritage dissemination systems be designed to maximise value in environments constrained by limited (computing and funding) resources?

Q3 How do technological advances transform museums practice?

The essence of these questions as well as the work done to address them is discussed in more detail below.

Q1 How can emergent, affordable technologies support (and replicate) the functions – heritage curation, preservation and dissemination – of traditional museums?

This question is important because for a virtual museum system to be of value, it should support the functions of a traditional museum as perceived by heritage practitioners. Hence, it should perform the functions of curation, preservation and dissemination of heritage content. This question will be addressed through the design and implementation of a Virtual Museum Infrastructure (VMI) that can be used to curate, preserve and disseminate heritage content, and an evaluation of the system to investigate its potential for fulfilling these functions satisfactorily from the perspective of heritage experts.

Q2 How can heritage dissemination systems be designed to maximise value in environments constrained by limited (computing and funding) resources?

This question is in two parts. The first part focuses on maximising value where computing resources are limited. This is important because an emphasis on low-cost, participatory engagement implies the use of commodity hardware, which may be constrained by limited computing resources (such as processing power, memory and graphics capabilities); for instance, the use of an affordable smartphone or workstation instead of a high-powered workstation, or the use of an affordable VR headset over a more expensive one. This question will be addressed through an iterative design and evaluation of a VR system to result in the deployment of a low-resource system for heritage exploration which leverages digital literacies and existing infrastructure.

The second part focuses on maximising value (in terms of experience for users and satisfaction for heritage practitioners) while minimising cost. This question is important because of the aim of this research, which focuses on the use of affordable technology for participatory heritage engagement. Funding is a major factor in the operations of heritage organisations. Therefore, the demonstration of a value-maximising, cost-minimising system design (and implementation) for participatory heritage exploration is valuable to heritage organisations with limited operating funds. To facilitate participatory engagement while minimising cost, digital literacies that museum visitors (and the general public) possess can be used. Digital literacies in this context refer to skills and resources – such as gaming proficiency and smartphone ownership – that people have come to possess by virtue of the proliferation of technology in aspects of daily life. This question therefore explores how museum experiences can be designed around visitors' smartphones and tablets to preclude the need for museums to purchase and maintain dedicated devices, as well as how museum experiences can be gamified to improve visitors' engagement with heritage content.

Q3 How do technological advances transform museums practice?

This question is important because of the disruptive nature of technology, which has historically spurred and/or catalysed change in industries such as education, media and retail. While the adoption of technology by heritage practitioners may bring about new opportunities for heritage dissemination and management, it may also bring about new challenges for heritage organisations. It is therefore worth investigating the transformations (for better or worse) that the adoption of technology will have on museums practice and heritage practitioners. This question will be addressed through an investigation of the impact of technology on heritage practice from the perspectives of both the practitioners (such as museum staff) and the end consumers (such as museum visitors).

These questions are addressed through a combination of underpinning (museum) theory, physical-world practice and case studies. Underpinning theory on museology and digital heritage is investigated by conducting a context survey that results in the development of taxonomies which identify common technology platforms and use cases for digital heritage and virtual museum systems. The context survey also identifies user groups of virtual museum systems, the needs of these user groups, and the basic forms in which virtual museum systems exist; these form the basis for the case studies. Case study 1 focuses on the use of immersive mobile technologies for heritage learning in a school setting, where secondary school pupils learn about local history using VR (both in the classroom and at the remains of a historic site) as part of a school module. The case study provides a platform for the investigation, design and implementation of an affordable alternative to more expensive VR systems. This system is evaluated in a heritage learning context with school pupils so as to gather comprehensive feedback on participants' perceptions of the system as compared to their more expensive counterparts. Case study 2 focuses on the use of an immersive museum installation to facilitate a comparison of the past and present from equivalent vantage points. The system designed in case study 1 is investigated further in case study 2, in a museum context with visitors and managerial staff, so as to ascertain the efficacy and value of the proposed design for heritage dissemination. Case study 3 focuses on the use of immersive communication technologies for real-time, remote tourism, where museum visitors take part in immersive, real-time telexploration[3] of heritage sites. The findings from the three case studies facilitate the design and implementation of a Virtual Museum Infrastructure (VMI) for active management, preservation and dissemination of heritage. This infrastructure is novel for the following reasons:

1. It supports active use which encourages stakeholders to continuously create, maintain and manage content using an intuitive management interface that allows content description using metadata.

2. It leverages mobile, VR, 3D, map-based and web-based technologies and combines them into a state-of-the-art system. It also uses multiple media types – Wiki, images, audio, video, spherical media and 3D artefacts – for heritage content management and presentation.

3. It encourages content reuse, such that heritage content that is created using the management interface can be reused on multiple platforms such as mobile, web and on-site museum installations. In this manner the VMI functions akin to the traditional museum archive, where objects can be described and stored, and then used (and reused) in multiple exhibitions.

The findings from the aforementioned case studies, as well as the evaluation of the Virtual Museum Infrastructure (VMI), add to the knowledge of how immersive systems can be leveraged

[3]The term "telexploration" (a portmanteau of tele and exploration) in this context is used to refer to the activity of engaging with and/or learning about a heritage site in a geographical location that is different from that of the explorer.

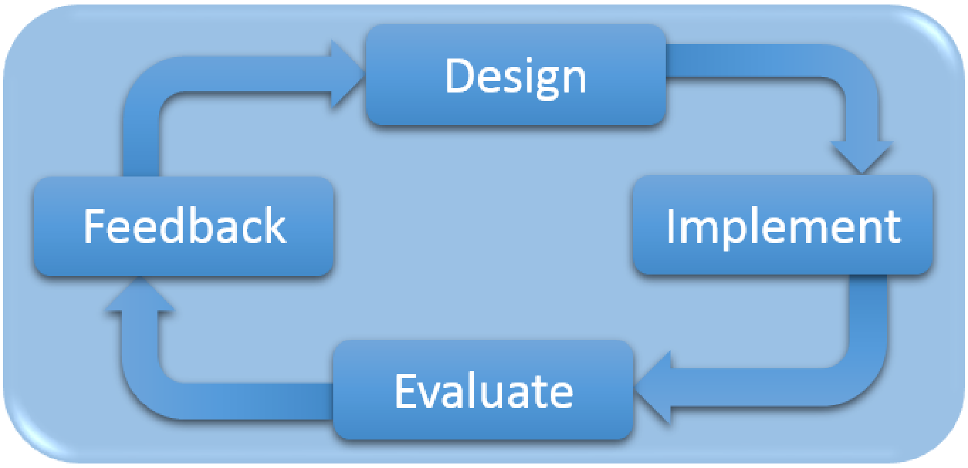

Figure 1.4: Iterative research

for heritage management, preservation and dissemination as outlined by the contributions in section 1.6.

1.5 Methodology

The methodology adopted for this thesis is one of an iterative, practice-orientated research [68], where the primary aim is to further knowledge in the application of affordable, immersive technologies for participatory heritage engagement through a novel Virtual Museum Infrastructure (VMI). The investigation is performed by reviewing literature so as to identify and question existing schools of thought on museum practice, and exploring case studies which feature immersive systems that are deployed in a physical-world heritage learning context and are evaluated using a combination of qualitative and quantitative techniques. Overall the methodology is largely qualitative. This is due to the nature of the topic being investigated, which, by virtue of its relatively young age requires exploration in order to develop theories, the need to present a detailed view of the subject matter being investigated, and the nature of the research questions (which ask *how* as opposed to *what*) [69, chapter 2, p.17]. The experiments take the form of cyclical use cases which entail designing, implementing and deploying systems in physical-world scenarios and gathering feedback from the deployment process and evaluation exercises conducted (see Fig. 1.4). This research incorporates elements of Systems Engineering and Human Computer Interaction, hence experiments comprise of both system evaluations and user studies and feature quantitative and qualitative aspects. System usability is measured using quantitative questionnaires while user experience is investigated using both quantitative and qualitative techniques. The quantitative elements take the form of users' responses to Likert scale-style questionnaire items, while qualitative elements take the form of focus groups, interviews (with heritage experts and non-experts) and observations of users' behaviours. Word frequency analysis and tag clouds are computed on the feedback from user interviews, and observations of subjects' actions are documented and analysed to discern patterns.

1.6 Contributions

As stated in section 1.5, the main aim of this research is to further the understanding of how affordable, immersive systems can promote participatory engagement with heritage. This aim is actualised through the development of a Virtual Museum Infrastructure (VMI) – a suite of tools that facilitate the creation, deployment and management of virtual museums. Specifically, the research outputs will provide a concrete understanding of the optimal design and development of immersive systems, the strengths and limitations of selected immersive technologies as applied to heritage exploration, and users' perceived value of these technologies, as well as propose a set of best practices for applying the technologies in different use cases. This gives rise to the following contributions:

C1 A design, implementation and evaluation of an infrastructure that can support and replicate the functions of traditional museums. This system is referred to as a Virtual Museum Infrastructure (VMI) [chapters 3 and 8].

The design features a conceptual model and architecture which are formulated based on design goals that are identified in the course of the literature review of virtual museums discussed in chapter 2, and the exploratory case studies in chapters 4 – 6. This is followed by an instantiation of the VMI, which is evaluated by means of workshops held with museums and their communities, followed by further design and instantiation based on the findings of the workshops [chapter 8].

This contribution stems from the work done to address **Q1**, **Q2** and **Q3**. The VMI demonstrates the use of emerging technologies for heritage curation (through a web-based platform that encourages content creation and description using metadata), preservation (through a data-store that ensures the continuity of heritage content in digital form if the physical counterparts are lost or damaged) and dissemination (through a suite of web, mobile and on-site systems that interpret and communicate heritage content in engaging ways), thus addressing **Q1**, and further explores how the design of the VMI can cater to the varying needs of museums along a resource spectrum by providing optimal value and experiences that are commensurate with the level of resources available, thus addressing **Q2**. In addition, an instantiation of the VMI espouses the delivery of content via apps that visitors can download on their smartphones and tablets, and thus demonstrates how visitors' devices can be used in lieu of dedicated delivery infrastructure, and a further instantiation of the VMI in the form of an on-site virtual museum incorporates game controllers, game-like menu systems and game-like navigation paradigms to leverage game proficiencies that museum visitors may possess thus fostering engagement with heritage content. This demonstration of how museum visitors' existing resources and skills (in the form of digital literacies) can be utilised further addresses **Q2**. The instantiations of the VMI in chapter 8 demonstrate that advances in technology can bring about a shift in the responsibility of heritage management from a few heritage experts to the greater heritage community, as demonstrated through a suite of tools for heritage content management and dissemination, and that advances in technology can facilitate a change in the perception of, and relationship with content in museums, as demonstrated throuh an evaluation of a VMI instantiation, where the findings suggest an appreciation for the coexistence of physical and digital heritage content in the same museum space. These represent significant change that are brought on by advances in, and consequent adoption of technology in heritage practice, thus addressing **Q3**.

C2 A discussion, design and implementation of an affordable Virtual Reality system for heritage education of school pupils [chapter 4]. The discussion features a comparative analysis of immersive technologies used for heritage exploration in the classroom which provides insights into the strengths and limitations of these technologies in different scenarios. The discussion reveals how a low-end, mobile VR system can be designed to enable an on-site comparison of the past and present states of a historic site, and how its design, featuring distinct view points, can be appropriate for task-based learning, while the graphics and processing resources of desktop VR can be appropriate for exploratory learning [chapter 4]. The discussion also highlights the research conducted, considerations, and best practices that are useful for system developers and heritage organisations who would like to adopt such a system.

An iterative evaluation of the system is also discussed, with an emphasis on how the research was conducted to understand and push the limits of the systems in an environment constrained by limited resources. In addition to the system evaluation, user evaluation exercises are conducted using likert-scale and free-form questionnaires, which add to the knowledge currently in the public domain on how people engage with immersive technologies and provide a starting point on the roles of immersive, headset-based technologies in fostering (or hampering) group heritage exploration in a classroom.

This contribution stems from the work done to address **Q1** and **Q2**. The designed VR system explores the use of spherical imagery for the representation, interpretation and dissemination of scenes (as landscapes, cityscapes and other forms of natural and cultural heritage), all of which support the functions that traditional museums perform with respect to the aforementioned heritage entities, thus partly addressing **Q1**. Furthermore, the discussion and activities that are aimed at pushing the boundaries of technology are geared towards addressing **Q2**, by investigating mechanisms for providing compelling experiences with minimal and already-available resources. This manifests as a comparative evaluation of an affordable VR system against a more established system, which provides insights into the feasibility of such affordable systems for facilitating engaging experiences.

C3 A design and implementation of an immersive, VR museum installation that facilitates a comparison of the past and present from equivalent vantage points [chapter 5]. The discussion shows how museums can maximise the user experience for visitors while minimising cost, as well as how immersive exhibits can be designed to ensure easy management (for museum staff) while facilitating ease of use by means of a controller-free design that provides snackable content to museum visitors. The discussion also features the findings of user evaluation exercises, which are discussed in terms of the content, technology and users' perceptions of immersive systems deployed in a museum context.

This contribution stems from the work done to address **Q1** and **Q2**. The idea of using spherical imagery (which is introduced in chapter 4) to facilitate a comparison of the past and present is further explored in chapter 5, thus partly addressing the dissemination function of museums in **Q1**. Furthermore, the deployment of a mid- to high-end VR system as an on-site museum installation is explored, which makes a case for the delivery of engaging experiences to visitors of museums that are situated towards the right end of the resource spectrum thus partly addressing **Q2**.

C4 A design and implementation of a system for immersive, real-time remote tourism [chapter 6]. The discussion shows how mobile VR systems can be designed to support real-time communication in remote areas constrained by a lack of electrical power and Internet access, and how these systems can support immersive group exploration.

The discussion also features incremental and iterative design and implementation procedures involved in supporting a VR streaming service which sources real-time video from a drone and transmits footage over a portable, affordable Raspberry Pi server. User evaluation exercises are also conducted to investigate the feasibility of the system for real-time virtual tours of remote sites, as well as the roles of immersive, headset-based technologies in fostering (or hampering) group heritage exploration of remote sites.

C5 An examination of the concept known as "Virtual Museum" and what it means in the modern age characterised by rapidly-evolving technology [chapters 1 and 2]. This discussion highlights common technological approaches adopted for the development of virtual museums through a taxonomy of virtual museums which identifies related work on digital heritage dissemination and categorises the work in different dimensions [chapter 2]. By virtue of the proposed virtual museum taxonomy, technological challenges are highlighted, user groups and their needs are identified and design goals are formulated based on the technological challenges and users' needs. In addition, an overview of the state of the art in the area of digital heritage is provided from a technological perspective – through a technology taxonomy of digital heritage applications, categorised based on the technological approach with sample work identified for each category – and a use case perspective – through a use case taxonomy of digital heritage applications identified in the aforementioned technology taxonomy, categorised based on the scenarios in which the work were deployed, and with sample work identified for each category [chapter 2].

1.7 Thesis Outline

Chapter 1 introduces the dissertation topic and discusses the motivation, methodology, thesis statement, research questions and contributions of the research.

Chapter 2 provides a background discussion on digital heritage, virtual museums and immersive systems, identifies the scope of the research, provides a discussion of related work and taxonomies for classifying work done in the public domain within the scope of this research.

Chapter 3 provides a discussion of the research methodology, the research questions, the motivation and methodology for addressing these questions, a discussion of the contributions of the research and an introduction to the Virtual Museum Infrastructure (VMI).

Chapter 4 discusses case study 1, which investigates the use of immersive technologies for heritage learning by early secondary school pupils. The design and implementation of an affordable VR system for on-site and classroom heritage learning are discussed in detail, and the findings of evaluation exercises which are conducted with school pupils are discussed. A discussion of the case study has also been published in conference proceedings – *Exploring the past with Google Cardboard* [1], and *Virtual Reality for Early Education: A Study* [2].

Chapter 5 discusses case study 2, which investigates the use of immersive technologies for on-site exploration in a museum setting. The design and implementation of an immersive, virtual museum installation which facilitates an on-site comparison of the past and present from equivalent vantage points are discussed in detail, and the findings of evaluation exercises which are conducted with museum visitors are discussed. A discussion of the case study has also been published in conference proceedings – *A Virtual Museum Installation for Time Travel* [3].

Chapter 6 discusses case study 3, which investigates the use of immersive technologies for remote tourism. The design and implementation of a real-time, drone-based, mobile VR system for remote, aerial tourism are discussed in detail, and the findings of evaluation exercises which are conducted with community members are discussed. A discussion of the case study has also been published in conference proceedings – *Aerial Virtual Reality: Remote Tourism with Drones* [4].

Chapter 7 introduces the exploratory workshops held in museums in Europe, Latin America and the Caribbean to investigate existing museum infrastructure and empower community members with the skills required to digitally disseminate heritage using commodity hardware. A discussion of these museum workshops has also been published in conference proceedings – *A Digital Museum Infrastructure for Preserving Community Collections from Climate Change* [5].

Chapter 8 discusses the design and implementation of the Virtual Museum Infrastructure (VMI) introduced in chapter 3, as well as its deployment and evaluation with over 20 museums worldwide. A discussion of the VMI design has also been published in conference proceedings – *A Digital Museum Infrastructure for Preserving Community Collections from Climate Change* [5] and *The Making and Evaluation of Picts and Pixels: Mixed Exhibiting in the Real and the Unreal* [6].

Chapter 9 provides a concluding discussion of how the exploratory case studies feed into the VMI and how the work in chapters 4 – 8 collectively address the research questions and result in the contributions of the thesis.

Chapter 10 reiterates the thesis statement, summarises the contributions and activities that address the research questions, highlights challenges associated with the research and outlines directions for future work.

1.8 Chapter Summary

The motivation and methodology for investigating the use of emergent technologies for participatory heritage engagement have been discussed. Research questions have been identified and contributions which stem from the investigation of these questions have been highlighted.

The collaborations and publications that have taken place in the course of this research have been outlined, and a thesis statement has been put forward.

Part II

Background, Theory & Methodology

Chapter Two

Literature Review

This chapter begins with a general discussion of the use of emergent technologies for interactive learning, remote tourism and time travel, as relevant to the case studies investigated in this thesis. A discussion of the concepts of ***digital heritage***[1]*,* ***virtual museums*** *and* ***immersive technologies*** *and the relationships between them is provided. An existing taxonomy for heritage visualisation strategies is discussed, and two taxonomies for classifying digital heritage work based on technological approach and use cases, as well as a taxonomy for classifying virtual museum systems along five dimensions are proposed. The proposed taxonomies of digital heritage and virtual museums, along with a discussion of their featured approaches and classification criteria, characterise the research area at the intersection of digital heritage and virtual museums.*

2.1 Research Background

This research features elements of software engineering, digital heritage and museum studies, and investigates how advances in technology and concomitant development of digital literacies can be leveraged for the curation, preservation and dissemination of heritage both within and outwith museum contexts. Much of this investigation is performed through exploratory work, hence a general discussion which is themed around the case studies investigated in chapters 4 – 6 is provided. Related work on the use of immersive technologies for classroom-based heritage learning, for on-site exploration of heritage in museum spaces and for remote exploration of inaccessible heritage sites are identified, and they provide a backdrop for a closer examination of the fields of digital heritage and virtual museums in the discussions that follow.

2.1.1 Virtual Reality for Learning

Virtual Reality (VR) refers to systems characterised by environments which are comprised solely of artificial elements courtesy of computer-generated imagery, and represent one extreme

[1] *Digital heritage*, in the context of this thesis, refers to the application of technology in the heritage domain, rather than the heritage, or history of technology.

in the Reality-Virtuality Continuum proposed by Milgram et al. [70], with the other extreme representing the real environment which is the world governed by the laws of physics such as gravity, time and space. Between the Virtual Environment and Real Environment exists the concept of Mixed Reality (MR), which represents systems made up of physical and virtual environments [70]. Augmented Reality (AR) refers to a class of MR systems where a physical environment is annotated with virtual (or synthetic) elements, in order to add information and context, or otherwise improve the physical environment. This thesis focuses on VR because it facilitates the distinct exploration of two realities (the virtual and the real) in parallel (as relevant to chapters 4 and 5), whereas AR merges both realities into one.

VR environments are equipped with the ability to facilitate shared experiences, to encourage natural user interaction and to enable unique learning experiences to meet individuals' needs; these features make them suitable for learning [71]. However, the cost, usability and fear of technology have posed significant drawbacks to the adoption of VR since the early 1990s [71]. These challenges notwithstanding, the adoption rate of VR is on the rise, as reductions in cost, increases in the power of microprocessors and improved user interfaces have led to new paradigms that enhance usability and accessibility of VR systems. This growing popularity of VR technologies for learning is investigated by Mikropoulos and Natsis [72], who review empirical studies conducted on the application of VR for learning from 1999 to 2009 and confirm that in the last decade VR has been leveraged for learning in a broad range of disciplines from the sciences to the social sciences; VR is also found to be appropriate for teaching and learning, as well as for training and entertainment purposes [73]. The efficacy of VR as a tool for learning has also been investigated. Studies conducted by Jacobson [74] suggest that VR technologies can serve as an invaluable resource and an effective tool for learning and teaching. Other work that underscores this finding includes the work of McCaffery, Miller and Allison [75], which introduces a VR system for teaching routing algorithms by utilising the graphical capabilities of an open-source VR platform to visualise network packets on a routing island, Perera et al. [76], which reports on a scenario for teaching a university module using a shared, 3D virtual environment, and Kennedy et al. [77], which discusses the deployment of the 3D model in scenarios such as museum exhibitions, science centres, schools and festivals.

Collaborative learning systems have also been developed using VR technology. The importance of social connectedness amongst learners is recognised by the work of Monahan, McArdle and Bertolotto [78]; which facilitates collaboration through features such as dedicated meeting spaces (such as coffee areas) that mimic physical environments. The authors conducted an initial study to ascertain the acceptance of the system as a collaborative learning tool and the results suggest that users are receptive to the use of such systems for learning, albeit with exceptions.

On-site exploration of heritage using mobile technology has also been extensively discussed in literature. A Virtual Time Window (VTW) which facilitates on-site exploration of heritage using an open-source VR server and a mobile client viewer is introduced by Davies, Miller and Allison [79], a location-aware, AR mobile application for exploring a historical street is introduced by Haugstvedt and Krogstie [80], the use of a serious game for heritage learning during museum visits is discussed by Coenen, Mostmans and Naessens [81], and the use of an affordable VR headset system featuring location-awareness and audio narratives to facilitate an on-site comparison of the past and present states of the remains of a monumental cathedral is discussed by Fabola, Miller and Fawcett [1].

The concept of experiential learning has also been facilitated by VR systems. For example, the experience of archaeological excavations is simulated using immersive, game-based VR systems so as to foster a better understanding of the process [82], and the application of game-like methodologies to the learning process in the aforedescribed system is presented by Getchell et al. [83]. The application of VR for learning has also been investigated from a pedagogical standpoint. The attitude of learners towards VR-based educational environments is investigated by Huang, Rauch and Liaw [84], who conclude that appropriate teaching practices should be taken into consideration in the development of VR learning environments, and suggests that the effectiveness of using VR for learning should be further explored. This is investigated in chapter 4.

2.1.2 Virtual Museums for Bridging Time and Space Barriers

The use of 3D technologies to recreate and visualise history has been documented extensively in literature. The reconstruction of an important but now derelict cathedral as it stood in the 14th century is documented by Kennedy et al. [77]. The work describes the reconstruction process including interdisciplinary research, building the landscape based on Ordinance Survey data, establishing the architecture, and embedding sound and scripted Non-Player Characters (NPCs) such as canons and historic figures. The resulting model and associated content have been deployed in learning contexts such as schools (for primary and secondary education), festivals (for community engagement) and on the web (for remote access through a browser). A framework for building interactive virtual museum content and exhibitions is proposed by Kiourt, Koutsoudis and Pavlidis [85]. The system leverages a popular game engine in addition to web frameworks to provide a distributed service (based on data pulled from popular online repositories) that enables users to easily create, manage and share virtual exhibits which are not limited to a specific application domain, but rather support a broad range of applications. A system for visualising 3D models of Mediterranean sculptures, optimised for both small (mobile) and large (desktop) screens is presented by Rodriguez et al. [86]. The system is capable of streaming content over a network and displaying content at multiple resolution levels so as to improve performance and facilitate the inspection of models in high levels of detail. In contrast to the work in Rodriguez et al. [86] which combines a mobile-based and desktop-based approach, Kostadinov and Vassilev [87] adopt a purely mobile-based approach to heritage visualisation. The system, which features cross-platform support (facilitated by web technologies), enables users to visualise a 3D model of a mediaeval town with minimal resources. The impact of 3D technologies in the domain of cultural heritage is investigated by Tait et al. [88], and the findings affirm the case for using digital capture approaches such as photogrammetry and laser scanning to foster community engagement with cultural heritage, as well as the use of immersive technologies such as VR headsets to provide compelling experiences thus increasing users' engagement with heritage content. These studies demonstrate the use of 3D technologies to recreate and visualise the past in different ways. The reconstruction process described by Kennedy et al. [77], the user-centric approach adopted by Kiourt, Koutsoudis and Pavlidis [85], the multi-platform approach adopted by Rodriguez et al. [86] and Kostadinov and Vassilev [87], and the findings of Tait et al. [88] make them relevant to chapter 5.

Technological advances, the proliferation of smart devices and 3D media, and the corresponding reduction in cost have contributed to the popularity of virtual museums in recent decades. The application of these technologies to create virtual museums solves some challenges associated with physical museums by requiring relatively less real estate for exhibitions and mitigating the risk of damage and/or degradation of exhibits [89]. The concept of the "Museum of Pure Form", proposed by Frisoli et al. [90], involves the use of VR headsets and haptic devices to provide immersive experiences to visitors. It mitigates a limitation of traditional museums (in which visitors cannot go close to or touch exhibits) by providing haptic feedback to simulate the sense of touch while interacting with digital objects in the exhibit, usually coupled with a stereoscopic visual display of the digital objects for improved immersion. In addition to bodily-immersive technologies (such as Head Mounted Displays (HMDs) and haptic touch devices), spatially-immersive technologies have also been deployed in museum contexts. The use of a panoramic stereo screen to present artwork to museum visitors is presented by Carrozzino et al. [91], who document a high degree of interaction with, and immersion in the virtual environment. In a similar vein, McCaffery and Miller [92] report on the recreation of a 19^{th}-century township and its deployment in a museum, which enables visitors to explore the township through body gestures. The installation uses the Microsoft Kinect for motion detection and three large projection screens arranged as a semi-hexagon to produce a wide field of view of 150°. A common denominator of the aforementioned work [89, 92, 91, 90] is their use of technologies in museum contexts, either to create digital replicas of museum exhibits (which can be accessed off-site) or to create virtual

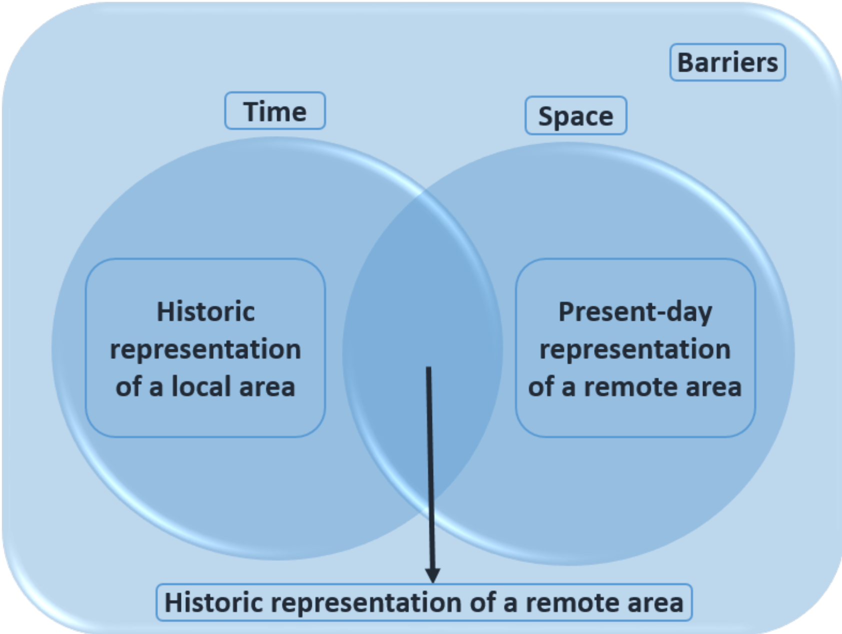

Figure 2.1: The use of technology to conquer time and space barriers

environments which can then be deployed as an on-site installation. This shows the maturity of the concept of virtual museums and the versatility and diversity of immersive technologies in a virtual museums context; for instance, body-based devices are used by Frisoli et al. [90] while space-based devices are used by Carrozzino et al. [91] and McCaffery and Miller [92], all with the common goal of giving users a strong sense of presence in a virtual world, which makes them relevant to chapter 5.

As highlighted in section 1.6, one of the contributions of this research is the development of a system for immersive, real-time remote tourism, which aims to mitigate time and space barriers [50, 93] associated with heritage exploration. Time barriers in this context refer to limitations imposed by being constrained in the present, which result in an inability to naturally re-live the past or experience the future. Space barriers in this context refer to the limitations imposed by geographical distance which result in an inability to be instantly transported to another place. In the physical world constrained by the laws of physics, humans can neither travel through time (physically go back to the past or go forward to the future) nor space (instantly teleport from one place to another) as of this writing. However, immersive technologies can provide high-fidelity visualisations of the past or representations of the future, and can situate the audience at the centre of a geographically-distant area, thus mitigating the barriers of time and space [50, 93]. A conceptual diagram which shows how immersive technologies can mitigate the physical barriers of time and space in the context of heritage exploration is proposed in Fig. 2.1.

Succinctly put, computer graphics and 3D technologies have facilitated the recreation of historical scenes and artefacts, which have been used in the development of exhibits and installations that offer interactive and immersive experiences to museum visitors. These developments have been leveraged in chapter 5, with the contribution of enabling museum visitors to visualise a historical scene from an equivalent vantage point thus fostering a comparison of the past and present.

2.1.3 Virtual Remote Tourism

The role of video in mediating tourist experiences has been investigated by Tussyadiah and Fesenmaier [94], whose findings suggest that such media facilitate imagination and reminiscence; two activities which are usually undertaken before and after touring a remote location. There have also been a proliferation of applications in the form of guides, recommendation systems and content dissemination systems for tourism and cultural heritage. These systems include web-based systems [95, 96], mobile trail apps [97], mobile, immersive systems [1, 98] and fixed, immersive installations [99, 100]. These studies demonstrate the use of 3D and multimedia technologies for remote tourism and cultural heritage dissemination, as well as the incorporation of remote video footage (as seen in the work of Tussyadiah and Fesenmaier [94]) and VR headsets (as seen in the work of Davies, Miller and Fawcett [98] and Fabola, Miller, Fawcett [1]) which are particularly relevant to chapter 6.

User Quality of Experience (QoE) for video streaming applications and stereoscopic image applications has been extensively studied and documented in literature. QoE in this context, refers to how well a video streaming application performs, as perceived by users [101]. Techniques have been proposed for QoE estimation and these can be split broadly into subjective [102], objective [103], and hybrid measures [104]. The determinants of QoE while viewing stereoscopic images is investigated by Xing et al. [105], who found that factors such as the baseline of the camera, the contents of the scene, the size of the screen, and the relationship between these factors all significantly affect QoE. The ability of users to distinguish between levels of stereoscopic video quality is investigated by Goldmann, De Simone and Ebrahimi [102]; the findings suggest that users generally have low sensitivity to changes in stereoscopic video quality, and depending on scene content, changes in camera distance may not impact QoE. In addition, the differences between users' perception of stereoscopic and non-stereoscopic images as viewed on a screen are investigated by Häkkinen et al. [106] and IJsselsteijn et al. [107] and the findings suggest that stereoscopic images tend to promote more feelings of presence than non-stereoscopic images.

VR scenes that incorporate moving components can be depicted using OmniDirectional Videos (ODVs), also known as spherical or equirectangular videos, as they have the ability to place viewers at the centre of a scene. These videos can be viewed on a screen, in a Computer

Automated Virtual Environment (CAVE), VR dome or VR headset. In addition to temporal navigation (rewinding or fast-forwarding) which is typically supported by UniDirectional (i.e. flat) Videos, ODVs also support spatial navigation (looking around the scene), and this is facilitated by their omni-directional nature. Screen-based devices (monitors, phones and tablets) naturally lend themselves to temporal navigation while viewing videos (by clicking or touching the screen), but they are less intuitive for spatial navigation – computer monitors require users to click and drag in order to change the view, and phones/tablets require the viewer to hold up the device and move around. Other (more immersive) means of viewing ODVs such as Computer Automated Virtual Environments (CAVEs), domes and VR headsets naturally support spatial navigation because the viewer (or part of their body) is physically at the centre of the environment in which the ODV is being viewed, hence by moving/turning/looking around, spatial navigation takes place. They are less intuitive for temporal navigation however, due to the absence of peripherals for point-and-click interaction. Given the ability of VR headsets to immerse viewers in a virtual environment and their natural support for spatial navigation of ODVs, the use of gestures, detected by a motion tracker mounted on a VR headset is proposed by Petry and Huber [108] for supporting temporal navigation. ODVs are becoming increasingly more attractive and applicable for various purposes; examples include Anderson et al. [109] who discuss the design, implementation and workflow of a system which captures stereo ODVs – videos in equirectangular projection with left and right pairs of frames for stereoscopic viewing – for consumption and streaming in VR headsets, and Wendrich et al. [110] who present a system that enables geographically-dispersed individuals to collaborate and visualise concepts through the use of HMDs that immerse users in the received video and audio (teleconference) feeds and hence provide the illusion of presence in remote locations. In addition, social media platforms such as Facebook [111] and YouTube [112] now offer support for the delivery of immersive ODVs to multiple users over the Internet. Users can interact with these videos on workstations by dragging with point-and-click devices, and also on mobile phones and tablets by moving the device from side-to-side while holding it up in the air. The real-time exploration system proposed in chapter 6 differs from that provided by these social media platforms in the sense that it can provide footage of inaccessible areas (because a drone can be flown over landscapes that people cannot access), hence it has the ability to provide an aerial view, as opposed to the land view which is usually obtained by people recording and streaming in real-time. Also, the system incorporates VR headsets, and although it relays uni-directional video from a camera to the headsets, it provides the illusion of omni-directional video and its characteristics (such as an encompassing view and the illusion of presence) by translating clients' head movements to the camera gimbal movements thus supporting spatial navigation. Given the real-time nature of the video stream, temporal navigation is irrelevant and is hence not supported.

On another note, Unmanned Aerial Vehicles (UAVs), popularly referred to as drones, are gaining interest amongst scholars, industry practitioners and hobbyists as consumer versions are becoming commercially available and affordable. This has led to various applications in the public domain such as using UAVs to independently detect and track moving targets in the fields of robotics [113], computer vision [114], autonomous tracking [115], thermal imagery and geo-referencing in the environmental sciences [116], disaster management [117], product delivery [62, 118, 64, 63] and of most relevance to this thesis, heritage studies [119, 120, 121, 122]. Of pertinence to chapter 6, is the application of UAVs for virtual tourism, and systems that can stream UAV footage to tourists over a network have been proposed. This streaming could be done over the Internet (as seen in the work of Mirk and Hlavacs [123] and Wu, Wang and Yang [124]) or a Wi-Fi grid (as seen in the work of Wu, Wang and Yang [124]), and the clients typically receive the footage in HMDs such as the Oculus Rift (used by Mirk and Hlavacs [123]) or the HTC Vive (used by Wu, Wang and Yang [124]). The work done in chapter 6 is similar to the work of Mirk and Hlavacs [123] and Wu, Wang and Yang [124] in the sense that clients are imbued with the ability to control the drone's camera using head movements, thus providing the illusion of presence in the remote location. However, it deviates from the aforementioned approaches in the sense that it focuses on the use of smartphones with mobile HMDs (such as the Google Cardboard [58, 59]) so that users can receive the footage without being tethered to a costly workstation thus ensuring mobility, low cost and minimal use of resources. Chapter 6 also incorporates a social element by enabling the video stream from a drone to be received by multiple subscribers. This way, more than one person can consume the footage at a time and a control protocol which dictates how (and when) each individual controls the drone has been implemented.

2.2 Digital Heritage

Heritage refers to the tangible and intangible attributes of society that are passed on from one generation to the next [13]. The concept of *heritage* has been important to societies for centuries who have taken part in the preservation and interpretation of heritage through oral (for example spoken tales and stories), written (for example messages on scrolls) and descriptive (for example engravings on stones) media [125]. The emergence and proliferation of digital technologies have provided additional media for the interpretation and preservation of cultural heritage and these have evolved over time [22].

Societies have always sought to leverage the technology at their disposal to improve their environment [126] and to enhance immersive experiences in virtual ones [127, p. 5]. Kenderdine [125] discusses the use of panoramic technologies for depicting scenes, and provides a brief history of the rise and subsequent fall in popularity of these technologies. A common denominator of these techniques was their circular shape which was meant to enclose the audience to foster a sense of being in the environment being depicted. Advances in mass media technologies led to the development and use of panoramas in the late 18th century. The popularity of the panorama was born out of a need to make sense of the world, propagate war propaganda, and experience historic or inaccessible locations through virtual time or space travel. At a basic level, these panoramic systems involved circular screens that surrounded the audience (and thus encompassed their fields of vision) so as to make them feel that they were at the place being depicted by the panorama. In other words, these systems strove to induce a feeling of immersion in some virtual environment; this was the essence and purpose of their circular shape.

Although these early panoramas were successful in inducing some sense of immersion, they were resource-intensive to maintain and were limited in featured content hence they did not gain traction until the late 19th century when circular screens became more popular, taking on multiple forms like the Mareorama, Pleorama and Cineorama. These panoramic systems were then used significantly to depict scenes, but as the visual realism increased, the viewers' demands for realism in other dimensions increased as well. Viewers longed for the ability to perceive the depictions through sound, smell, and ambient movements. Attempts were made to introduce movement into the scenes but these ultimately failed to produce the verisimilitude desired and in some cases degraded the experience by inducing feelings of nausea, dissonance and what is now referred to as simulation sickness. The shortcomings of the panorama (and its other forms) in failing to meet these demands ultimately led to their decline in the late 20th century.

Of the factors that led to the popularity of the panorama – the need to make sense of the world, the need to propagate war propaganda, and the need to experience historic or inaccessible locations through virtual travel – the first two are not as prevalent as they were in the 18th to 20th centuries when panoramas rose to popularity. However, the third factor – the desire to travel to geographically-distant lands or visualise a place as it was in the past – is still prevalent today. Advances in computer graphics and media technologies have precipitated the resurgence of panoramic technologies and these are now commonplace in cultural heritage interpretations. It is common for museums to host panoramic images of real and virtual environments depicting a significant region, building or monument as it was in the past, or a point of interest that may be inaccessible to interested parties. An example is Timespan [128], which hosts an interactive panoramic view of Caen (a former township in the far north of Scotland) as it was in the 19th century before the Highland Clearances – the coerced, mass emigration of people from the Scottish Highlands. This interactive reconstruction enables the emigrants' descendants and the general public to engage with the town's heritage, and is one of numerous examples of how today's visual media – flat, moving and spherical images – improve on erstwhile panoramic technologies. Moving images (in the form of video and film) introduce ambient movement in scenes and facilitate easy content swaps. However, the flat nature of the media removes the audience from the centre of the scene, unlike spherical media. Now, the advent of spherical media on desktop, mobile and web platforms facilitates highly-realistic ambient movement, scene immersion, seamless content swaps and widespread dissemination, all of which translate to engagement with, and ease of management of content.

Two Dimensional (2D) imagery (in the form of photographs and panoramas) has been available for centuries, but the use of Three Dimensional (3D) imagery is more recent, as the first application of computer graphics for historic reconstructions was in the 1980s [129]. While the use of computers for generating 3D imagery in the form of models and environments has been largely adopted by the heritage community, the adoption is not without resistance as some heritage practitioners view 3D objects as inferior replicas of the original artefact, some believe that creating these digital replicas of objects can undermine the original artefacts [130], and some doubt the efficacy of technology-driven exhibits given the mechanical failures that may arise [131]. Cameron [130] argues that one solution is to adopt an anti-materialist approach to the functions of museums, such that museums should be viewed as establishments that disseminate heritage (among other functions) as opposed to being viewed as central repositories of objects. Furthermore, this change in mindset should dispense with the practice of comparing an artefact – a physical, tangible object – with its digital model – a virtual, intangible one. Rather, the digital model should be treated like the physical artefact – like an entity in its own right with attribution, authorship and provenance. This is akin to the treatment of photographs in museums and art galleries as creative media in their own rights with archival and interpretative roles, not cast in the shadows of the scenes they represent.

This (anti-materialist) trend is becoming more prevalent as established museums continue to digitise artefacts (and entire collections in some cases) to produce 3D objects for online and on-site use. An example is the British Museum [132] who have digitised over 200 artefacts in over 20 collections (as at the time of this writing), and disseminated them on Sketchfab [133], a popular social platform for sharing and consuming 3D objects. This trend is not limited to large museums, as smaller heritage organisations and even non-heritage practitioners are embracing 3D scanning and printing technologies and training their staff and communities on how to use them [134, 135, 136]. Social media sites are also acting as a catalyst for this trend by providing incentives to heritage organisations. These include offering free or discounted account upgrades to heritage and educational organisations [137] and providing support for embedding and interacting with 3D objects on popular social media sites [138, 139].

Now more than ever, there is a need for a change in views concerning the relationship between a physical artefact and a digital representation of that artefact due to the rise in popularity of 3D objects in the heritage community and the increasing availability and affordability of 3D printers and scanners. Using laser scanners, cameras and photogrammetry, a physical object can be digitised to produce a digital object. In a similar vein, a 3D object can be made using Computer Assisted Design (CAD) technology, which can then be printed to produce a physical object. Given the two-way relationship that now exists between a physical, tangible object and a digital, intangible one, both objects should be treated as entities in their own right, as opposed to viewing one as the original and the other as the replica. This raises the question of what is real and what is virtual, or what is the original and what is the replica. The answer(s) to this question notwithstanding, the availability of computer graphics technologies that are capable of representing physical artefacts in a manner that encourages interaction has opened up possibilities for heritage curation, preservation and dissemination through virtual museums, because gone are the days when multimedia were viewed as extras in museums; they are now objects in their own right [140].

In addition to photographic scenes and digital objects, other types of digital interpretation mechanisms have come to the fore. These include the use of Mixed Reality (MR) technologies to combine elements of real and synthetic environments in the exploration of heritage. There are also techniques to foster greater immersion in these virtual worlds. A step up from viewing virtual worlds on computer screens is to view them in Head Mounted Displays (HMDs), Computer Automated Virtual Environments (CAVEs) or domes. A discussion of how these immersive technologies have evolved over time is presented in section 1.2 and further discussion of their different forms is provided in section 2.4.

2.3 Virtual Museums in a Digital Heritage Era

Of the various definitions of the term *virtual museum* in the public domain, there is no widely-accepted one. Lewis [141] defines virtual museums as:

> *"a collection of digitally recorded images, sound files, text documents, and other data of historical, scientific, or cultural interest that are accessed through electronic media. A virtual museum does not house actual objects and therefore lacks the permanence and unique qualities of a museum in the institutional definition of the term."*

Schweibenz [142] adds to Lewis' definition [141] by referring to a virtual museum as a collection of objects of a digital nature, organised in a way that makes sense to the stakeholders. Schweibenz [142] also identifies a virtual museum as one that is not bound by space or confined to a place, in line with the concept of a *"museum without walls"* proposed by Bearman [143]. Irrespective of the context in which the term *virtual museum* is used, the term usually refers to either (1) the like-for-like recreation of a physical museum space, which enables visitors to remotely simulate a virtual visit in the halls and rooms of the museum or (2) the creation (usually from scratch) of artefacts, items, exhibits or experiences which are then presented to users for consumption. This could be done remotely (online for example) or on-site (in a physical museum space) [30].

Three type of virtual museums are identified by Schweibenz [144] as *brochure museums*, *content museums* and *learning museums*. *Brochure museums* are aimed at providing information (such as location and opening hours) about a physical museum thus encouraging potential visitors to actualise the visit. The general public are the primary target audience for this type of virtual museum, because it features informative content that would convince and/or aid users to plan a visit to the physical museum. *Content museums* are aimed at providing information about objects held in a physical museum, catalogued in an object-oriented manner which is usually akin to the way a collection database is arranged. Heritage experts are the primary target audience for this type of virtual museum, because they may contain domain-specific information that are neither useful to, nor understood by the general public. *Learning museums* are aimed at stimulating learning about objects in a context-oriented way based on users' demographics. Like the brochure museum, the general public are the primary target audience because the content is usually tailored to the information needs and the skills of non-experts. Although content museums encourage users to visit the physical museum in order to interact with the real objects, this is not the primary goal (unlike the brochure museum), but rather to provide a substantial learning experience to users who are unable to visit the physical museum. It is this ability of learning (virtual) museums to stimulate users' interests in visiting a physical museum space and to provide an informative experience in the absence of a physical visit, that make them the focus of this thesis.

Three categories of users of virtual museums are identified by Costalli et al. [145] as visitors (or tourists), students (or educators) and experts (or administrators); three groups of virtual visitors are identified by Booth [146] as general visitors, educational visitors, and specialist visitors; and Charitos et al. [147] identify three categorises of (traditional) museum visitors as researchers, students and the general public. Irrespective of the nature of the museum, whether virtual or traditional, by combining [145], [146] and [147], the following three overlapping groups of users are identified based on their interests and needs, as shown in Fig. 2.2.

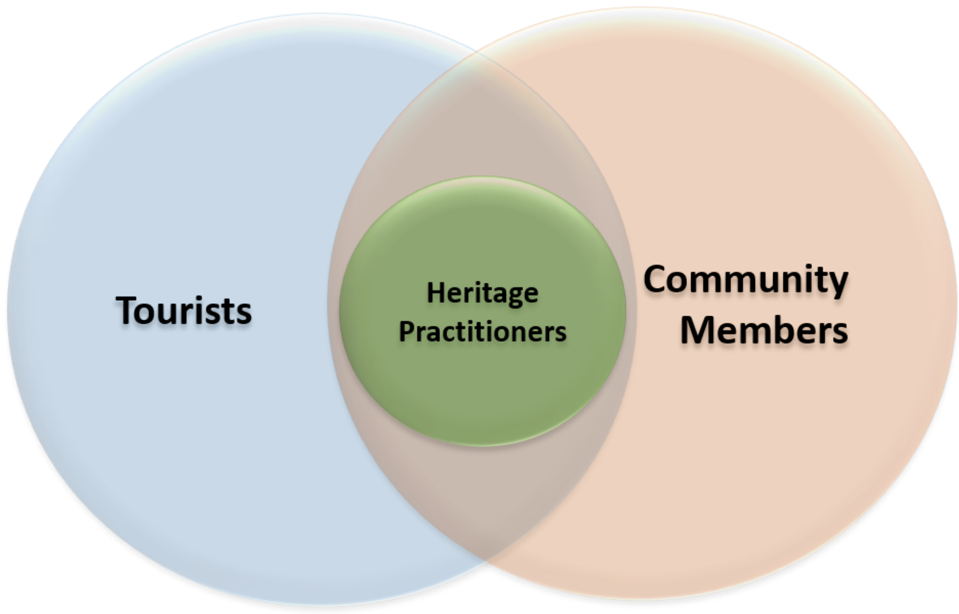

Figure 2.2: Users of heritage content classified by interest and needs

Heritage Practitioners: This is a specialist group of users, and it includes students – individuals who may be studying towards a degree in a heritage-related field (such as art, archaeology or museology), researchers – individuals who may be academics in a heritage-related field, and heritage experts – individuals who may hold jobs or positions in the heritage sector, such as curators, conservators and archivists. These users engage in-depth with a small (often niche) area of heritage, although they may be conversant with a much broader topics. This group are interested in investigating, managing, and preserving heritage content and seek access to original sources through fieldwork, archives and libraries.

Community Members: This group overlaps with heritage practitioners and tourists, and it includes members of the general public with a keen interest in local and community heritage. They may work with (or as part of) teams of heritage practitioners, and are interested in preserving and disseminating the heritage of their local communities. They seek opportunities to get involved with the dissemination process by donating materials, or volunteering with heritage organisations.

Tourists: This group also overlaps with heritage practitioners, and it refers to members of the general public who engage with heritage content as a leisure activity, either in their local area as a pastime or while travelling on holiday. During their travels, the heritage they engage with is usually that of foreign locations as opposed to community members who engage with local heritage. This group are interested in the consumption of heritage content and often seek new learning experiences.

There are overlapping relationships between these groups of users. Heritage practitioners can be community members (who engage with and assist their local heritage organisations) and tourists (who consume local or foreign heritage content as a leisure activity). Some users who consume heritage content as a leisure activity neither identify nor engage with the heritage of a community and it is conceivable that some community members do not engage with heritage content for the sole purpose of leisure. However, a preponderance of users are at the intersection of the community members and tourists groups. It is important for virtual museum systems to cater to these three groups of users, hence the Virtual Museum Infrastructure (VMI) detailed in this thesis adopts a design that caters to the needs of heritage practitioners who seek to investigate and manage heritage, community members who seek to preserve and disseminate heritage and tourists who wish to consume heritage content.

The case for the adoption of digital technologies by traditional museums is extensively documented in literature. A study conducted by Loomis, Elias and Wells [148] revealed that 70% of potential museum visitors are more likely to visit the physical museum after they have had prior engagement with the museum website. Furthermore, the use of digital technologies to represent and disseminate information about objects can be effective in overcoming barriers of space (real estate required to host physical exhibitions) and time (required to initially set-up and rearrange objects based on presentation requirements). In addition, with the advent of Virtual Reality (VR) technologies, digital artefacts can be presented to users in remote locations in an immersive and interactive manner which facilitates an in-depth inspection and engagement with said artefacts by users [149].

While there are benefits that may accrue to museums and heritage organisations that adopt digital technologies, there are drawbacks as well, some of which are discussed by Sylaiou et al. [89]. A major drawback to the adoption of virtual museums is the reluctance of unconvinced heritage experts and sceptical users to the adoption of technology for heritage interpretation. The rationale is often that digital interpretations are based on educated guesses which lead to the misrepresentation of the structure, detail or interpretation of artefacts, hence the end result of (say) a digital reconstruction or visualisation may not be a true representation of a historic artefact. Another drawback is that the use of technology for the dissemination of heritage information may discourage users who are averse to technology from consuming or engaging with the disseminated content. Some other challenges associated with the adoption of digital heritage include technical limitations (such as poor system usability), research limitations (such as the inability to ascertain the veracity of content) and financial limitations (such as the cost of purchasing and maintaining equipment, and staffing skilled personnel) [150, 151, 152, 153]. The design of a Virtual Museum Infrastructure (VMI) – as detailed in this thesis – which encourages

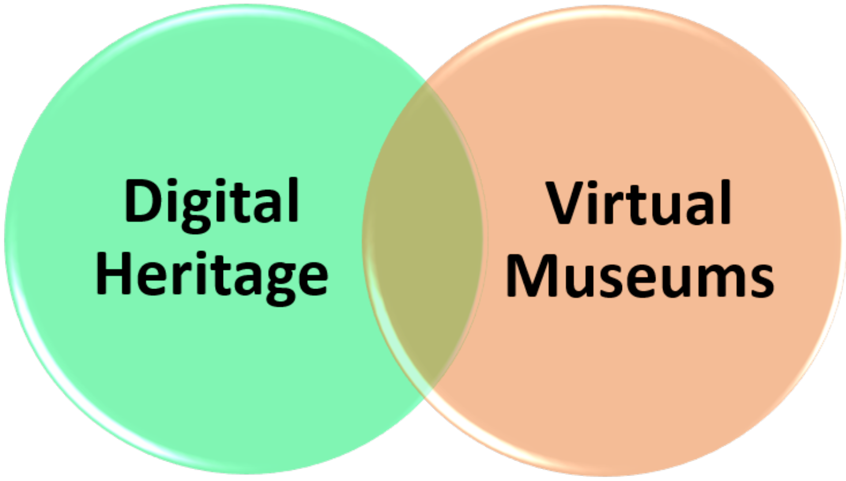

Figure 2.3: The relationship between digital heritage and virtual museums

active use by stakeholders to create, maintain and manage heritage content, can overcome the aforementioned limitations associated with virtual museums by involving heritage experts and community members in the content creation process thus mitigating the misrepresentation of content, as well as eliciting requirements from museum visitors on how to present snackable heritage content using a controller-free design, thus making it more appealing to less tech-savvy individuals.

Virtual Museums and Digital Heritage

Section 2.2, introduces the term *"digital heritage"* as the use of technology to preserve, interpret and/or disseminate heritage, whether tangible or intangible, cultural or natural and section 2.3 introduces the term *"virtual museums"* as the use of technology to perform the functions of a traditional (often physical) museum. *Digital heritage* should not be confused with *digital artefacts*, that is, heritage items (such as pottery, documents, carvings and so on) that have been recorded or represented using technology. Similarly, *virtual museums* should not be confused with *digital exhibits*, that is, collections of museum offerings (such as documents and images) that have been digitally recorded. Therefore, in the context of this discussion, the term *digital heritage* is used to refer to the field, that is, the body of work that is concerned with the use of technology to record, represent and interpret heritage, and the term *virtual museums* is used to refer to the body of work that is concerned with supporting the functions of traditional museums using technology. Where the connotation of digitally-represented heritage items is required, the term *digital artefacts* will be used. Similarly, where collections of these heritage items is implied, the term *digital exhibits* will be used. *Virtual museum system*, or *virtual museum* (note the singular form) in the context of this discussion, is used to refer to a system (or platform, framework or infrastructure) that maintains one or more *digital exhibits* for the purposes of curating, preserving and disseminating its contents. Where more than one *virtual museum* is referred to, the plural form of *virtual museum systems* will be used.

Museum functions include preserving, curating and disseminating heritage, but also include interpreting and disseminating contemporary art, science and fashion – tasks and topics which may be considered unrelated to heritage. In a similar vein, there are aspects of *digital heritage* that are not museum-related; for instance the use of technology to simulate an archaeological excavation furthers the goal of *digital heritage*, but it may or may not be performed in the context of a museum, and when it is not performed under the auspices of a museum, it may not be considered as part of the *virtual museums* field. The implication of this relationship is that the fields of *digital heritage* and *virtual museums* are not mutually-exclusive, and thus overlap substantially. However, this relationship is non-equivalent, such that the two fields do not entirely span the same area and hence should not be used interchangeably, and non-inclusive such that one field is not a subset of the other. In other words, while the fields of *digital heritage* and *virtual museums* overlap substantially, there are aspects of *digital heritage* that have little relevance to *virtual museums* and there are aspects of *virtual museums* that have little relevance to *digital heritage*. This relationship is visualised in Fig. 2.3. The green area represents the field of *digital heritage*, while the orange area represents the field of *virtual museums*. The area of overlap – where *virtual museums* are used for the preservation, interpretation, curation and dissemination of heritage – is investigated in this thesis. This distinction, specifically the non-equivalent and non-inclusive relationship between *digital heritage* and *virtual museums*, provides rationale for the classification of *digital heritage* work and *virtual museums* work separately which results in distinct taxonomies for *digital heritage* and *virtual museums*, as discussed in section 2.5.

2.4 Immersive Technologies

Section 2.2 introduced the concept of using immersive technologies for heritage exploration. This section will explore the nature of immersive technologies, specifically Virtual Reality (VR) systems, how they have evolved over time and work that has been done on using such systems for immersive learning.

2.4.1 Virtual Reality (VR)

Virtual Reality (VR) refers to a class of systems characterised by synthetic (often computer-generated) environments which may be purely fictional or factual representations of the real world at a given point in time, and are composed of three facets (or Is): *Immersion*, *Interaction* and *Imagination* [49]. VR occupies one end of the Reality-Virtuality Continuum described by

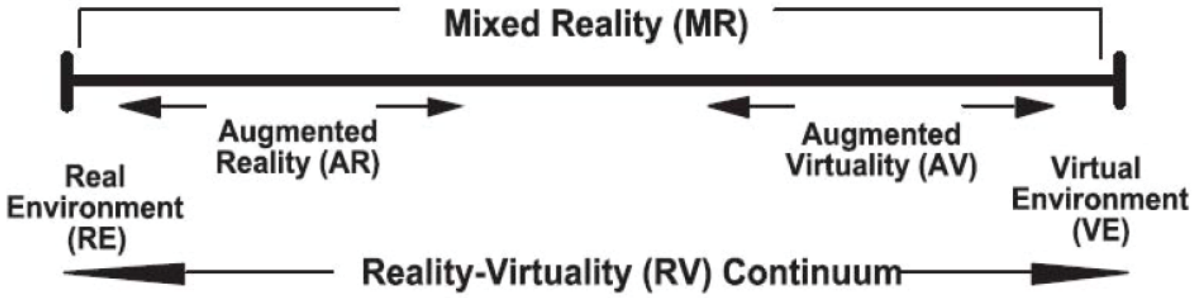

Figure 2.4: Reality-Virtuality Continuum (Source: [154])

Milgram and Colquhoun [154] and Milgram et al. [70]. Examples of VR systems include the interactive model of the St Andrews Cathedral as it stood in 1318, deployed on a computer with an Oculus Rift and an Xbox Controller [77] and the use of a full-dome to recreate cultural scenes with a multiplayer experience [99]. VR systems have been investigated and applied in several fields including classroom education [155], medicine [156, 157], tourism [158] and archaeology [159, 160, 161, 162, 163, 164]. Virtual Reality (VR) should not be confused with Augmented Reality (AR), which is a system made up of a superimposition of synthetic elements on a representation of the real world [154]. Examples include a first person shooting game with positional tracking [165], a system that superimposes layers of content onto real-world objects [166], and more recently, a system that enables users catch, train and fight animated characters (reminiscent of an erstwhile television show) which can be found existing in representations of the real world [167]. Augmented Reality (AR) is a class of Mixed Reality (MR) systems, systems which are made up of two realities (the real environment and the virtual environment). Another class of Mixed Reality (MR) systems is Augmented Virtuality (AV), which are systems that superimpose real world content on virtual environments [70]. Fig. 2.4 highlights the concepts of Virtual Reality (VR), Augmented Reality (AR) and Mixed Reality (MR). The concept of Cross Reality (XR), also known as Parallel Reality or Dual Reality, has also recently emerged [168, 169] as a class of systems where the real world and virtual world exist side by side (with representative mapping), thus facilitating the exploration of both realities in tandem and allowing users to switch between one and the other. The distinction between Virtual Reality (VR), Augmented Reality (AR) and Cross Reality (XR) notwithstanding, it is important that they be viewed as mediums (or concepts) that are implemented using technology in order to fulfil predefined goals or objectives, rather than viewing them as technologies themselves [126, 170, 50]. Augmented Reality (AR) blends the real environment and the virtual environment, Cross Reality (XR) facilitates tandem exploration of the real and virtual environment, and Virtual Reality (VR) enables interaction with a synthetic environment which may or may not be representative of the real environment. It is this feature of Virtual Reality

Figure 2.5: The relationship between presence and immersion (Visually represented from [173])

(VR) – the ability to engulf users in a given representation of reality – that makes it amenable to the exploration of heritage and hence of interest to this thesis.

2.4.2 Presence and Immersion

The terms *"presence"* and *"immersion"* in the context of Virtual Reality (VR) are relevant due to the focus of this thesis on participatory engagement and interactive experiences in museum contexts. *"Presence"* has different (sometimes related) connotations in different circles. Lombard et al. [171] identify six dimensions of presence: presence in terms of social richness, realism, transportation, immersion, social actor within medium, and medium as social actor. These six dimensions can be further grouped into two broad categories: physical presence (the feeling of being in a geographical location which is usually different from the current location of the subject) and social presence (the feeling of being or communicating with another entity) [172].

The definition of presence adopted in this thesis is that proposed by Slater and Wilbur [174] which states that *"presence is a state of consciousness, the (psychological) sense of being in the virtual environment"*. Presence should be distinguished from immersion, which is defined as the ability of a system to induce an *"...illusion of reality to the senses"* of a user [174]. Presence has been suggested to be a function of immersion [173]; the more immersive the technology is, the greater the sense of presence experienced by the user, all other factors being constant. However, literature [173, 174, 175, 176] suggests that immersion does not translate directly into presence, as presence is a function of cognitive processes which take immersive elements of a Virtual Environment as input. Thus, two users who interact with a VR system (and are subject to the same levels of immersion provided by the system) can experience varying levels of presence by virtue of differences in the mental processes that translate the effect of immersion into the feeling of presence. Succinctly put, cognition (imagination) and action (interaction) translate immersion, an objective property of the technology, into presence, a subjective property of a system (recall the three facets of VR identified by Burdea and Coiffet as immersion, interaction and imagination [49]). This relationship is depicted in Fig. 2.5, showing that immersion is the objective property of a given technology which serves as input to a user's cognitive faculty. Cognitive process take place in the user to translate the effect of the immersion (input) provided by the technology into presence (output). Presence is therefore the subjective experience derived as a result of the cognitive processes acting on the immersion provided by the technology, hence the experience of presence in a given virtual environment is different for each user [173]. This is in line with Witmer and Singer [177] , who identify two distinct (but interplaying) components of presence: *psychological immersion* and *involvement*. *Psychological immersion* is a consequence of the construction of the mental model of a virtual environment, while *involvement* is a consequence of the allocation of attention. Given the distinction between presence (as a subjective experience of a user) and immersion (as an objective property of a system), the latter is predominantly used to refer to the ability of systems to place users in virtual environments because of the objective (and somewhat quantifiable) nature of the term.

2.5 Taxonomies and Classifications

The term *"taxonomy"* is used in this thesis to refer to *"A classification of something; a particular system of classification"*, as defined by the Oxford English Dictionary[2]. This section identifies an existing taxonomy proposed by Foni, Papagiannakis and Magnenat-Thalmann [26] which is confined to the approaches used for cultural heritage visualisation strategies. A technology taxonomy for digital heritage is proposed to classify digital heritage systems based on their technology characteristics, beginning with the level of immersion used in the technology. The work identified in the technology taxonomy is further categorised based on the scenarios in which the systems are deployed to produce a use case taxonomy. As discussed in section 2.3, the relationship between digital heritage and virtual museums necessitates a distinction between these fields, hence a separate taxonomy for virtual museums is proposed, to categorise work in five dimensions based on common themes. These taxonomies provide an overview of the research space at the intersection of digital heritage and virtual museums, and identify the commonly-adopted approaches and existing gaps in the research area.

2.5.1 4D Taxonomy for Cultural Heritage Visualisation Strategies

A taxonomy for cultural heritage visualisation strategies is proposed by Foni, Papagiannakis and Magnenat-Thalmann [26]. The taxonomy, which is depicted in the form of a three-dimensional cube (Fig. 2.6), categorises work in four dimensions:

1. Precision and visual consistency: The extent to which the work features photo-realistic content, represented by the Y axis of the cube.

[2] Oxford English Dictionary Online: http://www.oed.com/view/Entry/198305?redirectedFrom=taxonomy#eid (Accessed: 2018-03-09).

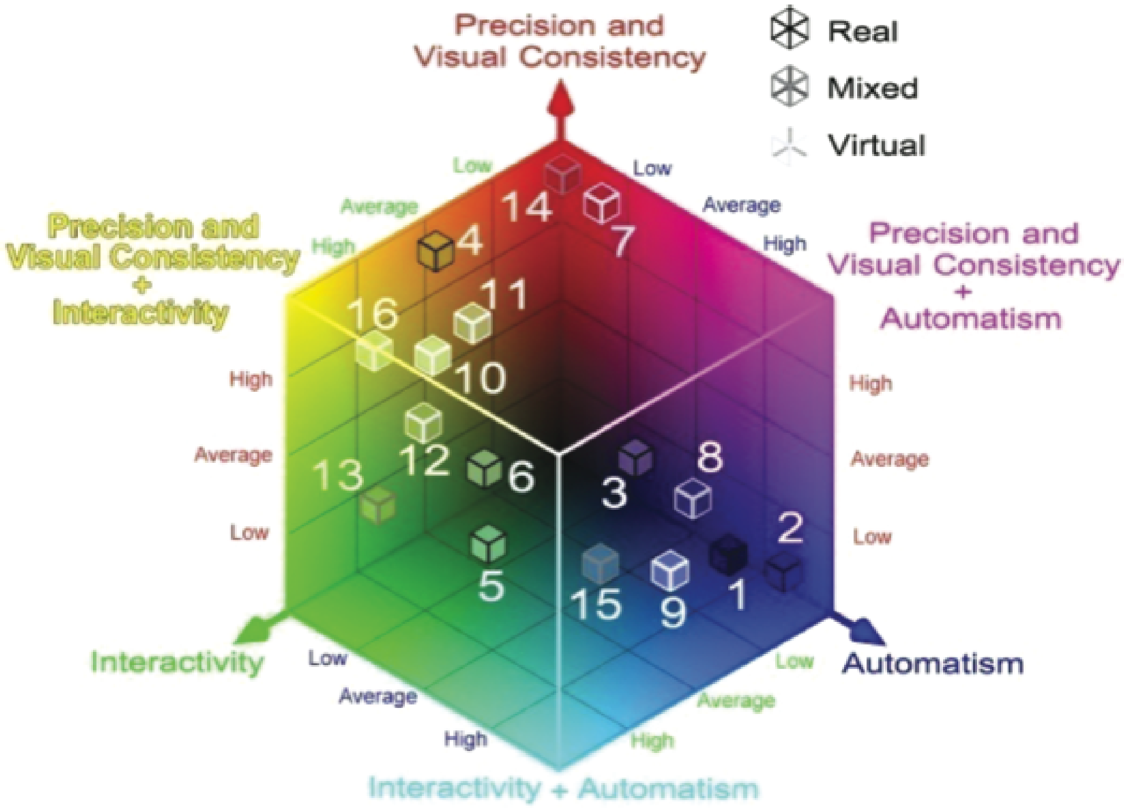

Figure 2.6: Foni, Papagiannakis and Magnenat-Thalmann's taxonomy (Source: [26])

2. Automatism: The extent (in terms of time and effort) to which the work flow of the technique can be easily reproduced, represented by the X axis of the cube.

3. Interactivity: The extent to which users can engage with the content, represented by the Z axis of the cube.

4. Virtuality: The extent to which the work features real, virtual or mixed content, represented by the shade of grey of a point in the three-dimensional cube space.

Foni, Papagiannakis and Magnenat-Thalmann's taxonomy [26] identifies sixteen (16) types of cultural heritage visualisation strategies (see Table 2.1) and categorises them based on the four aforementioned dimensions. It also identifies five categories of players involved in the cultural heritage application process: *Content Holders* (such as heritage organisations, museums, art galleries), *Field Personnel* (such as conservators, restorers, curators), *Cultural Intermediaries* (such as heritage researchers, journalists, writers), *Content Consumers* (such as tourists, museum visitors, art aficionados, the general public), and *Digital Intermediaries* (such as technology experts, digitisation personnel, system developers). The preferences of these players with respect

		Precision	Interactivity	Automatism	Virtuality
		R	G	B	V
1	Restitution drawings	0.1	0	0.7	0
2	Augmented pictures	0.1	0.1	0.8	0
3	Scale models	0.4	0.3	0.6	0
4	Physical reconstructions	0.8	0.6	0	0
5	Interactive scale models	0.4	0.8	0.5	0
6	Live experiments	0.5	0.7	0.4	0
7	Renderings	0.85	0	0.15	1
8	Digital catalogs	0.2	0.15	0.7	1
9	Digital panoramas	0.3	0.4	0.9	1
10	Real time VR simulations	0.8	0.9	0.4	1
11	Stereoscopic visualizations	0.7	0.7	0.3	1
12	Computer games	0.6	0.75	0.25	1
13	Real time AR simulations	0.5	0.85	0.25	0.5
14	Augmented movies	0.9	0	0.1	0.5
15	Semantically supplemented 2D	0.2	0.5	0.65	0.5
16	Semantically supplemented 3D	0.8	0.95	0.25	1

Table 2.1: Cultural heritage visualisation strategies and their 4D coordinate values (Source: [26])

to the four items are also investigated, and the players' perceptions of each of the four dimensions are grouped as either necessary, optional or unnecessary.

Foni, Papagiannakis and Magnenat-Thalmann's taxonomy identifies visualisation strategies, and then proceeds to compare them along the four dimensions of precision, interactivity, automatism and virtuality, which is beneficial in the sense that it enables digital exhibit designers (and stakeholders) to prioritise and make decisions depending on their project objectives. That said, this thesis advocates for the combination of strategies (instead of evaluating them in isolation) to result in an approach that accentuates the strengths and attenuates the weaknesses of each strategy, thus resulting in the design of engaging and informative digital heritage applications. Furthermore, while the taxonomy is a useful starting point for categorising cultural heritage applications, the rapidly-evolving nature and the adoption of technology for heritage exploration necessitate a taxonomy based on technological characteristics. This thesis therefore proposes a technology taxonomy for classifying cultural heritage applications, and in addition, a use case

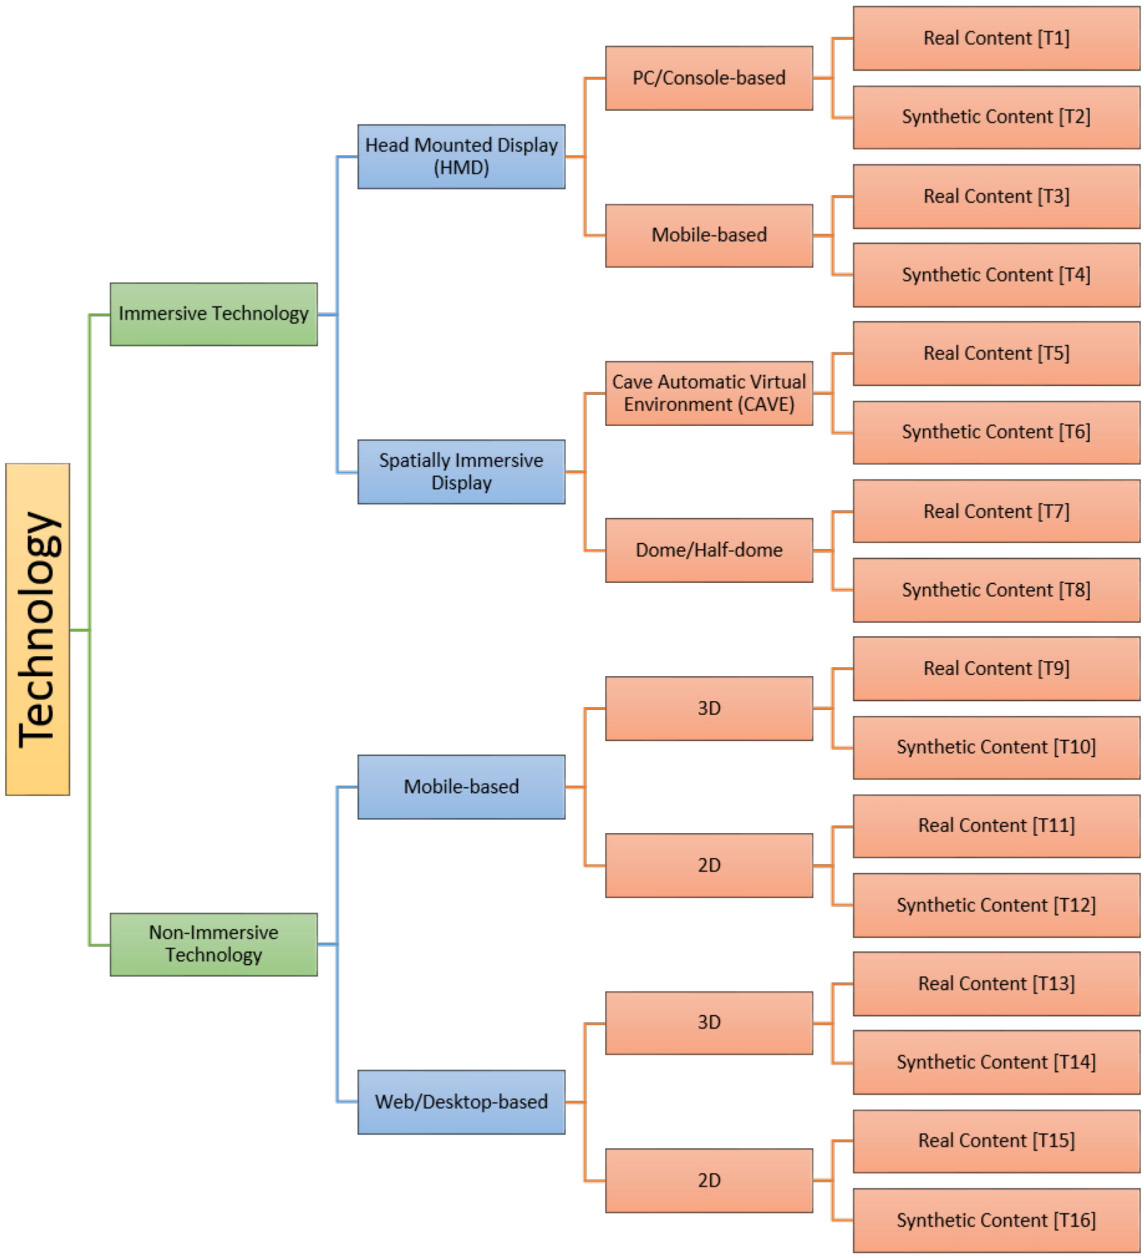

Figure 2.7: A technology taxonomy of heritage exploration systems

taxonomy that classifies cultural heritage applications based on scenarios that are facilitated by the technology identified in the technology taxonomy. It is pertinent to note that the scope of the proposed taxonomies is confined to work that leverage technology (i.e. digital techniques) for heritage exploration, hence both taxonomies exclude work that exclusively feature traditional means of heritage exploration (such as hand-written archives and unrecorded oral stories).

2.5.2 A Proposed Technology Taxonomy of Digital Heritage

Diverse approaches have been adopted for the exploration of heritage, as documented in literature. A discussion of VR technologies as applied for learning is presented by Lee and Wong [178], who identify three categories: partially-immersive, fully-immersive and Augmented Reality (AR), and provide further categorisation into online and offline Virtual Reality (VR) based on interaction and complexity. This thesis adopts the definition of AR proposed by Milgram et al. [70] – a system composed predominantly of real-world elements on which synthetic elements are superimposed – and hence proposes a modification to the categories identified by Lee and Wong [178]. AR systems need not form a stand-alone category, as they could be either fully-immersive (for example a headset-based system for fighter-jet pilots, which provides a see-through display on which data is overlaid) or partially-immersive (for example a pre-recorded, spherical video of a landscape on which an animated dinosaur is superimposed, experienced through a CAVE). Thus, an alternative classification of immersive systems – based on characteristics of the technology – would categorise them as either bodily-immersive (for instance a HMD system) or spatially-immersive (for instance a dome-like, spatially-immersive display). This is the basis for the proposed technology taxonomy shown in Fig. 2.7. Bodily-immersive, HMDs are typically driven by either workstations (PCs) or mobile devices (smartphones), and spatially-immersive systems are typically made up of CAVEs or domes, the difference being in the geometric structure and projection of the surface. Non-immersive systems, on the other hand, typically deliver content on smartphones (via mobile apps) or workstations (via desktop applications or web access), and while web content can be accessed on smartphones, web-based systems are grouped with desktop-based systems in this taxonomy due to the commonness of work that feature workstations with web access in digital heritage applications. It is pertinent to note that the categories in this taxonomy are neither exhaustive nor mutually-exclusive; rather they are based on the approaches that have been identified in the course of the literature review and the distinguishing features between them.

Overview of Technologies

The use of technology for heritage exploration is commonplace, and such technology can be immersive or non-immersive. Immersive technology is defined as technology which provides a compelling illusion of place to users, such that users' awareness of their surrounding environment may be reduced relative to their perception in an alternate reality [174]. Conversely, non-immersive technology does not compel users' senses to provide an illusion of reality, hence users remain fully aware of their surrounding environment while interacting with such technology. The immersive technology category can be further categorised as bodily-immersive (for example Head Mounted Display (HMD)) or spatially-immersive (for example Computer Automated

Key	Features	Sample Work
T1	Immersive Head Mounted Display, PC-based, Real Content	[168, 98]
T2	Immersive Head Mounted Display, PC-based, Synthetic Content	[179, 98]
T3	Immersive Head Mounted Display, Mobile-based, Real Content	[180]
T4	Immersive Head Mounted Display, Mobile-based, Synthetic Content	[1, 181]
T5	Immersive Spatial Display, CAVE-based, Real Content	[182]
T6	Immersive Spatial Display, CAVE-based, Synthetic Content	[183, 184]
T7	Immersive Spatial Display, Dome-based, Real Content	[185]
T8	Immersive Spatial Display, Dome-based, Synthetic Content	[99]
T9	Non-Immersive Mobile, 3D, Real Content	[186]
T10	Non-Immersive Mobile, 3D, Synthetic Content	[79]
T11	Non-Immersive Mobile, 2D, Real Content	[97]
T12	Non-Immersive Mobile, 2D, Synthetic Content	[80]
T13	Non-Immersive Web/Desktop, 3D, Real Content	[187]
T14	Non-Immersive Web/Desktop, 3D, Synthetic Content	[77, 95]
T15	Non-Immersive Web/Desktop, 2D, Real Content	[96]
T16	Non-Immersive Web/Desktop, 2D, Synthetic Content	[188]

Table 2.2: A technology taxonomy of heritage exploration systems

Virtual Environment (CAVE)) technology. In the bodily-immersive sub-category, devices can either be driven by a workstation or game-like console, or by a mobile phone, while common examples of spatially-immersive displays include CAVEs and (semi) domes. The quintessential difference between a CAVE and a Dome – in the context of this thesis – is the geometrical structure. A CAVE has rigid edges and is akin to a cubic environment while a dome has curved edges and is akin to a spherical environment. Each of these technology sub-categories may feature real (for example a photograph of a landscape) content, synthetic content (for example a 3D model of a building) or a mixture of both. The non-immersive technology can be further categorised into mobile-based and web/desktop-based technology. In the mobile sub-category, the content can either be predominantly Two Dimensional (2D) or Three Dimensional (3D). Like the immersive technologies, the 3D sub-category can feature real or synthetic content. Non-immersive, mobile-based, 2D systems usually come in the form of trail apps or AR apps. The web/desktop-based category can also be 3D-based (with real or synthetic content) or 2D-based (for example digital exhibitions and multimedia collections). An overview of the features (and sample work) in this taxonomy is provided in Table 2.2 the categories are highlighted below.

1. **T1 – Immersive Head Mounted Display, PC-based, Real Content**
   The work in this category feature an immersive HMD which is tethered to (and powered by) a workstation and provides captured, real-world content to the user. Examples of work in this

category includes [98], which uses the Oculus Rift [189] headset to provide both real and synthetic content. The real-world content is provided by means of a stereoscopic see-through camera, while the synthetic content is in the form of a 3D reconstruction of the site of interest, served from a laptop to which the Oculus Rift is tethered. The work is a typical example of the concept of Cross Reality (XR) [168] also known as Parallel or Dual Reality, which facilitates simultaneous exploration of two alternate (i.e. the real and the virtual) realities.

2. **T2 – Immersive Head Mounted Display, PC-based, Synthetic Content**
This category is comparable to T1, as it also features a workstation-powered, immersive, HMD. However, it differs from T1 in the sense that the content is created or generated by means of computer graphics, and is thus synthetic rather than real. The Cross Reality (XR) system proposed by Davies, Miller and Fawcett [98] provides a good example of work in this category because it features synthetic content in the form of a 3D reconstruction of a 15^{th}-century chapel.

3. **T3 – Immersive Head Mounted Display, Mobile-based, Real Content**
Like the T1 and T2 categories, this category features an immersive HMD. However, the display is powered by a mobile device rather than a workstation. This approach provides the added benefits of mobility and portability at the expense of computing power and rendering fidelity. Like T1, this approach features real-world content as seen in the work of Rasheed, Onkar and Narula [180], a mobile application designed for use with the Google Cardboard [190], which immerses users in virtual environments represented by spherical images of real-world locations.

4. **T4 – Immersive Head Mounted Display, Mobile-based, Synthetic Content**
This approach features a mobile-powered, immersive HMD like T3, but features synthetic (rather than real) content. An example of this approach is seen in the work of Fabola, Miller and Fawcett [1], which combines the Google Cardboard with mobile phones to facilitate an exploration of the 3D model of a 14^{th}-century cathedral. The system is location-aware, such that it supports automatic navigation to a location in the virtual environment when a user comes in close proximity to the on-site location, and it features the use of panoramas and spherical media to depict virtual locations in a low-end, yet high fidelity approach as compared to the deployment of a resource-intensive 3D model.

5. **T5 – Immersive Spatial Display, CAVE-based, Real Content**
Like the T1 - T4 categories, immersion is an essential feature of the T5 category. However, the immersion is achieved by means of a spatially-encompassing display, which takes the form of a Computer Automated Virtual Environment (CAVE) featuring real-world content.

An example is seen in the work of Smith et al. [182] which features an immersive CAVE environment in which stereoscopic images of historic sites are projected.

6. **T6 – Immersive Spatial Display, CAVE-based, Synthetic Content**
This category is similar to T5 in its use of an immersive CAVE, but features synthetic (rather than real) content. An example is [183], which uses a CAVE-based virtual environment for archaeological research. The system features an 8ft by 8ft by 8ft cube, coupled with shutter glasses for stereo vision and a tracker that controls position and orientation in space, with the aim of immersing users in a virtual environment that represents a historic temple from which excavation data was previously obtained. Another example is [184] which uses a CAVE-like system coupled with stereoscopic goggles for exploring reconstructed heritage sites immersively in a shared public space.

7. **T7 – Immersive Spatial Display, Dome-based, Real Content**
Like the T5 and T6 categories, this category features an immersive, spatial display. However, the virtual environment is rendered in a dome (or half/semi-dome) rather than a CAVE, and features real-world content. An example is seen in the work of Kenderdine [185] which uses an immersive spatial display for the visualisation of high-resolution images.

8. **T8 – Immersive Spatial Display, Dome-based, Synthetic Content**
This category features an immersive, dome environment, like T7. However, it features synthetic (rather than real) content. An example is the work of Tredinnick and Richens [99], which proposes a full dome environment with an equirectangular projection that enables multiple users to simultaneously interact with heritage content. Multi-user support is provided by means of fiducial markers (one for each user) and an infra-red camera which tracks the markers, while the content represents a 3rd-century villa and surrounding landscape.

9. **T9 – Non-Immersive Mobile, 3D, Real Content**
Unlike the T1 - T8 categories, immersion is not a quintessential feature of this category (and subsequent ones). This category features a mobile device on which real-world content is displayed to users in a (sometimes faux) 3D manner. An example is the use of spherical media such as spherical images or spherical videos of real-world locations displayed on a smartphone or tablet. The mobile application proposed by Li, Chuah and Tian [186] provides an example of this category, as it enables users to view 3D content in the form of spherical images on a smartphone or tablet screen.

10. **T10 – Non-Immersive Mobile, 3D, Synthetic Content**
This category is similar to T9 as it displays 3D content on flat, mobile surfaces, but features

synthetic (rather than real) content. An example is [79], which features a tablet and a virtual world viewer that facilitates an on-site exploration of heritage sites such that users can compare the virtual representation of a site provided by the system to the real-world view in the surrounding environment. The system attempts to map the virtual world to the real world using GPS location data, and as the findings suggest, the veracity of the mapping was constrained by the accuracy of the location data.

11. **T11 – Non-Immersive Mobile, 2D, Real Content**
This category features non-immersive, mobile technology which presents real-world, 2D content to users. An example is seen in the work of Fabola et al. [97], which features a location-aware, trail app that enables users undertake a tour of a town. The system features a map interface which shows points of interest, accompanied with multimedia such as images, audio and video clips for each point.

12. **T12 – Non-Immersive Mobile, 2D, Synthetic Content**
This category features non-immersive, mobile technology which presents synthetic (rather than real, unlike T11), 2D content to users. An example is seen in the work of Haugstvedt and Krogstie [80], which features an AR application running on a tablet for the exploration of a historical street. In addition to the AR view, the system superimposes text and images on real-world photographs. The findings of a survey conducted using the system reveal that the use of AR mobile applications is influenced by the usefulness and value of the system as perceived by the users.

13. **T13 – Non-Immersive Web/Desktop, 3D, Real Content**
Like the T9 - T12 categories, immersion is not an essential feature of this category. It features real-world, 3D content displayed on a workstation screen (either on a website or as downloadable content) instead of on a mobile device. Examples include [187], which uses web-based panoramas and map interfaces for visualising heritage objects and real-world locations.

14. **T14 – Non-Immersive Web/Desktop, 3D, Synthetic Content**
This category is similar to T13, as it displays 3D content on flat, desktop screens, but features synthetic (rather than real) content. Examples include [77] which provides a 3D model of a cathedral as it stood in the 14^{th} century, accessible online via a web browser, and as a standalone desktop application and [95] which provides a web-based portal for exploring 3D models of artefacts.

15. **T15 – Non-Immersive Web/Desktop, 2D, Real Content**
    This category features non-immersive, web or desktop-based technology which present real-world, 2D content to users. An example is seen in the work of Albanese et al. [96], which proposes a web-based recommender system for browsing multimedia content such as image galleries.

16. **T16 – Non-Immersive Web/Desktop, 2D, Synthetic Content**
    This category features non-immersive, web or desktop-based technology which present synthetic, 2D content to users. An example is seen in the work of Barak et al. [188], which proposes a web-based system for preserving, curating and disseminating cultural heritage.

The categories that have been identified in the proposed technology taxonomy demonstrate the diversity in the technical approaches that have been adopted for heritage exploration, featuring both immersive and non-immersive systems, mobile, web and standalone applications, as well as real and synthetic imagery. These categories are not mutually-exhaustive, as they can be, and as advocated for in this thesis, should be combined to leverage the strengths of each technique for the effective dissemination of heritage content.

2.5.3 A Proposed Use Case Taxonomy of Digital Heritage

A use case taxonomy is presented in Fig. 2.8 based on work identified in the public domain. This taxonomy categories the use cases that have been enabled by the technological approaches identified in the technology taxonomy in Fig. 2.7. The categories in this taxonomy are not exhaustive and in the majority of the identified cases they are not mutually-exclusive. An exception is the remote, real-time use case, where offline modes cannot exist because the nature of the use case necessitates a network connection. All other identified use cases can exist in both online and offline modes. However, the majority of these use cases incorporate a mixture of online and offline modes, with preferences for the offline mode when possible so as to minimise network and bandwidth resource usage.

Overview of Use Cases

The technologies identified in Fig. 2.8 can be categorised based on the scenarios in which they are deployed, to yield a use case taxonomy. Digital heritage applications may facilitate remote exploration (online, for example) or on-site exploration (in a physical museum space, for example) [30], hence this forms the basis for the top-level classification. The remote use case refers to situations in which the explorer is in a different geographical location (and is thus physically absent) from the site of interest, while the on-site category refers to situations where the explorer is physically present at the site of interest. In the remote category, content can be

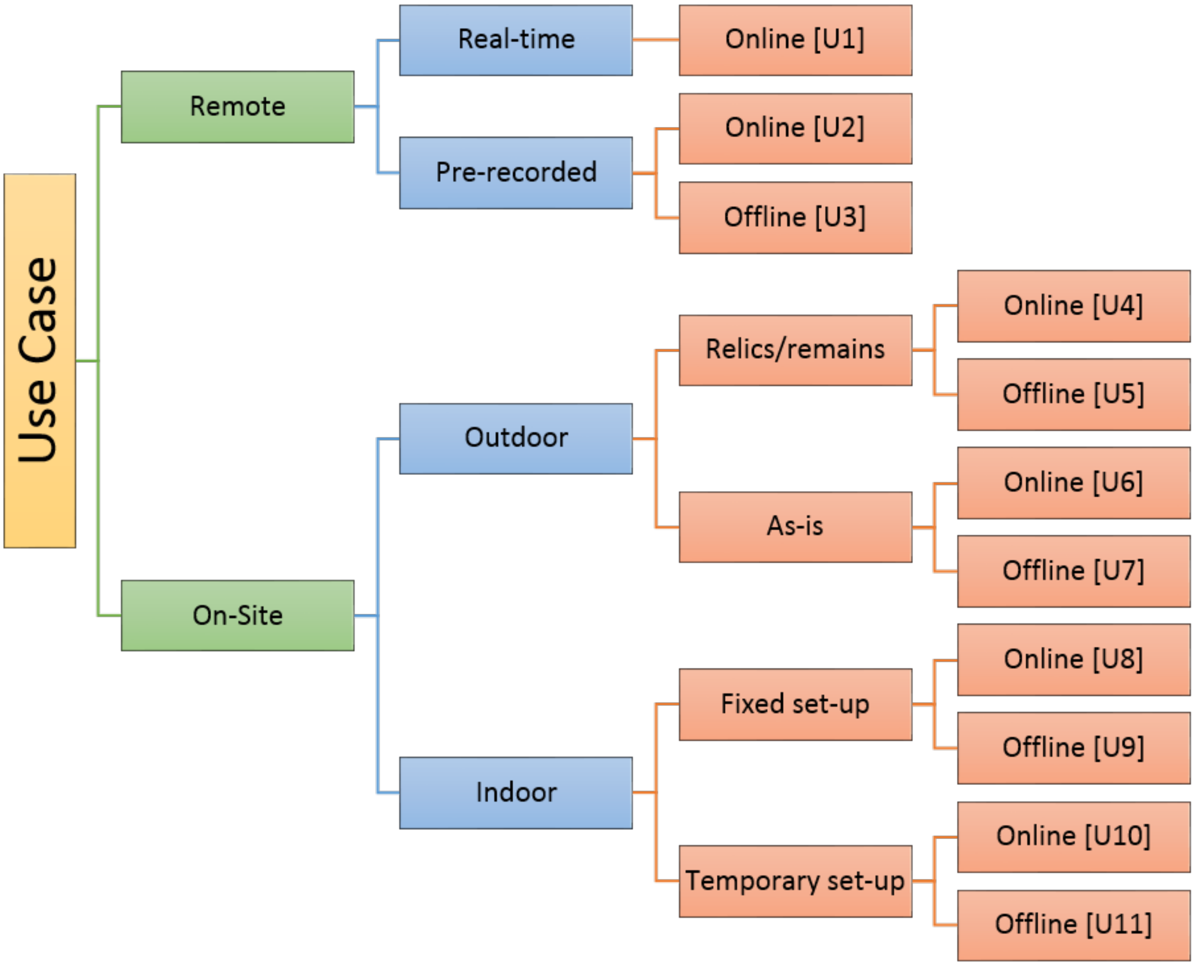

Figure 2.8: A use case taxonomy of heritage exploration systems

served in real-time (live) or pre-recorded. Real-time content is served online (over a network infrastructure) while pre-recorded content can be served online or offline (as canned content). The on-site category is further sub-divided based on whether the area of exploration is indoors or outdoors. In outdoor exploration, the area of interest can represent relics (or remains) of a historic site, or a landscape or structure in its current form. The indoor category is further sub-divided based on the permanence of the environment. The environment could be fixed (for example a museum building) or temporary (for example a portable exhibition). Irrespective of the sub-category, the content can either be served online or offline, with the exception of the remote, real-time category in which content can only be served online. An overview of the features (and sample work) in this taxonomy is provided in Table 2.3 the categories are highlighted below.

Key	Features	Sample Work
U1	Remote, Real-Time, Online	[123, 124]
U2	Remote, Pre-recorded, Online	[95]
U3	Remote, Pre-recorded, Offline	[191]
U4	Outdoors, On-site, Relics, Online	[97]
U5	Outdoors, On-site, Relics, Offline	[92, 192]
U6	Outdoors, On-site, As-Is, Online	[97]
U7	Outdoors, On-site, As-Is, Offline	[98]
U8	Indoors, Fixed Set-up, Online	[193]
U9	Indoors, Fixed Set-up, Offline	[92, 194]
U10	Indoors, Temporary Set-up, Online	[92]
U11	Indoors, Temporary Set-up, Offline	[92, 100]

Table 2.3: A use case taxonomy of heritage exploration systems

1. **U1 – Remote, Real-Time, Online**
   This use case features online (i.e. over a network) exploration of remote (i.e. geographically-distant) locations in real-time. Examples include the work done by Mirk and Hlavacs [123] and Wu, Wang and Yang [124] which feature drones and VR headsets for undertaking live virtual tours of remote locations over a wireless network infrastructure.

2. **U2 – Remote, Pre-recorded, Online**
   This use case is similar to U1 in the sense that it features online exploration of remote locations. However, it differs in the sense that it streams pre-recorded content (rather than content that is generated in real-time) over a network. An example is seen in the work of Guarnieri, Pirotti and Vettore [95], in which the data is gathered, prepared and made available online for remote access.

3. **U3 – Remote, Pre-recorded, Offline**
   This use case features pre-recorded content of remote locations, but rather than streaming the content (like the U2 use case), the content is packaged and provided offline such that it can be accessed without a network. An example is the work of Rua and Alvito [191], in which 3D models were created using Computer Assisted Design (CAD) and Geographic Information Systems GIS) and used to represent an ancient Roman villa, accessible offline in a game engine.

4. **U4 – Outdoors, On-site, Relics, Online**
   Unlike the U1 - U3 use cases, this use case (and subsequent ones) features the on-site (rather than remote) exploration of a location of interest. In this use case, the content is provided

online and the location of interest is represented by remains (or relics) of a historic site. An example is seen in the work of Fabola et al. [97], which features a mobile, location-aware trail app that encourages users to visit points of interest drawn on a map, some of which represent relics of historic sites.

5. **U5 – Outdoors, On-site, Relics, Offline**
   This use case is similar to U4, as it facilitates on-site exploration of the relics of a historic site. However, the content is packaged and provided offline (rather than online), as seen in the work of Gaucci, Garagnani and Manferdini [192], which discusses the exploration of the site of an ancient Etruscan town using mobile, VR headsets, as well as [92], which features an offline trail app as one of the modes exploration of a 19th-century township.

6. **U6 – Outdoors, On-site, As-Is, Online**
   Like U4 and U5, this use case involves the on-site exploration of a site of interest. However, unlike U4 and U5, the exploration is not of the relics of a site but rather of the site in its present state (i.e. as-is). The content is provided online, as seen in the work of Fabola et al. [97] where some of the points of interest on the map interface represent sites in their current state (as opposed to relics or remains).

7. **U7 – Outdoors, On-site, As-Is, Offline**
   This is similar to U6 as it involves the on-site exploration of a site as-is. However, the content is canned and provided offline (rather than online). An example is [98], which facilitates a parallel comparison of the present state of a chapel with a representation of said chapel set in the 15th century.

8. **U8 – Indoors, Fixed Set-up, Online**
   This use case features the exploration of heritage content, provided online in a fixed, indoor environment such as a museum building. An example is the work of Patel et al. [193] which discusses a system for building and presenting web-based digital exhibits that are accessible remotely over the Internet and as well as on-site in museum spaces.

9. **U9 – Indoors, Fixed Set-up, Offline**
   This use case is similar to U8, as it features the exploration of heritage in a fixed, indoor environment. However, the content is canned and provided offline (rather than online), as seen in the work of Pietroni et al. [194] which discusses the recreation and visualisation of Etruscan graves as part of an interactive VR installation, as well as [92], which discusses a fixed, immersive museum installation for exploring a 19th-century township using bodily-gestures such as raising one's arm to move or rotating one's torso to turn while in front of a

semi-hexagonal projection screen.

10. **U10 – Indoors, Temporary Set-up, Online**
    This use case is similar to U8 and U9 in the sense that the exploration of heritage takes place indoors. However, it differs in the sense that the environment is usually temporary i.e. a structure that has been erected for a particular purpose for a period of time such as a moving exhibition or an event. In addition, the content is usually provided online. An example is seen in the work of McCaffery and Miller [92] which discusses a portable, immersive exhibition which can be temporarily set-up for community events.

11. **U11 – Indoors, Temporary Set-up, Offline**
    This use case is similar to U10 as it features the exploration of heritage in a temporary, indoor environment. However, the content is canned and provided offline (rather than online), as seen in the work of Lo et al. [100] which discusses the use of a kiosk equipped with stereoscopic display, 3D graphics and panoramas for heritage exploration, as well as [92] which discusses the use of portable, immersive technology for community festival events.

The categories that have been identified in the proposed use case taxonomy demonstrate the variety in the scenarios in which digital heritage applications have been deployed. While the scenarios are categorised based on whether they facilitate remote or on-site exploration, it is beneficial, where possible, for digital heritage applications to support both remote and on-site use exploration, and this hybrid approach is advocated for in this thesis. For instance, on-site access can enable heritage exploration at the remains of a site or in a museum setting, both of which can serve as social and engaging learning experiences, while remote access can facilitate heritage exploration and learning where a site (such as a museum or an archaeological dig) is inaccessible due to time or space barriers. These benefits are discussed in chapter 4, which investigates the use of technology for heritage learning both on-site (at the remains of a monumental cathedral) and remotely (in the classroom).

2.5.4 A Proposed Taxonomy of Virtual Museums

As discussed in section 2.3, there is a significant overlap between the fields of *digital heritage* and *virtual museums*. However, these terms are neither equivalent nor interchangeable, hence the need for a taxonomy of *virtual museums* in addition to the proposed technology and use case taxonomies for *digital heritage*. The following discussion therefore outlines the rationale for the classification of virtual museums based on related work.

Museums may face constraints associated with a shortage of resources (see section 1.2). Some challenges associated with space and resource constraints are highlighted by Tsichritzis et al. [195]. Firstly, museums are traditionally expected to display physical artefacts. This is sometimes problematic because some artefacts may be too large, incomplete or harmful to put on display when there are space constraints. Secondly, these artefacts are traditionally meant to be displayed in a public, physical setting. This may be problematic because of the associated cost and logistics involved in the construction and maintenance of the physical space. Thirdly, exhibiting artefacts in a public, physical space poses limitations because it requires visitors to be in close proximity to said artefacts in order to engage with them. This may require visitors to travel long distances to where the artefact is hosted, and in some cases it prevents visitors from viewing an artefact in its natural setting (for instance in the artefact's place of origin). Lastly, artefacts may be sensitive, harmful or in peculiar condition, which is problematic because it may exclude the option for visitors to touch, move or closely observe the artefacts while engaging with them. The use of virtual museums can mitigate and (in some cases) eliminate these constraints by displaying digital objects which may or may not be representations of real artefacts, in virtual environments, accessible remotely over the Internet, with interactive paradigms which support touch, rotation, movement and close observation. Tsichritzis et al. [195] discuss how technological developments can be leveraged to mitigate the aforementioned constraints of traditional museums, stating that advances in computer graphics, networking and mobile technologies facilitate the development of virtual museum systems that can deliver engaging on-site and remote exploration of heritage. Although the discussion by Tsichritzis et al. [195] is dated and the technologies are obsolete, the overarching concepts are still relevant, as the technologies discussed can be viewed as precursors to today's network, graphics and mobile systems. Network systems facilitate remote visits, advanced graphics facilitate the digitisation of physical artefacts to create 3D artefacts, and mobile technologies can be used to interact with 3D objects, hence the development of virtual museums can now be actualised.

As discussed in section 2.2, virtual museums represent a divergence from traditional museums, as many still view a visit to a virtual museum differently from a visit to a traditional museum [130]. Paolini et al. [196] identify four ways in which a virtual museum visit may be interpreted: (1) a website that organises virtual content in a gallery that replicates the structure of the content of a traditional museum, (2) a website that does not attempt to replicate the structure of a traditional museum and is hence not bound by physical limitations, (3) an immersive virtual environment that replicates a traditional museum and hence simulates a visit along the halls and rooms of the museum, and (4) an immersive virtual environment that represents an imaginary space which does not replicate any physical space. The varying interpretations of a virtual museum visit [196] and the diverse functions that virtual museums can perform in easing the constraints of traditional museums [195], have resulted in virtual museum systems that differ based on function, features, content and technology. For instance, a virtual museum could support remote (online) visits or on-site visits, it may provide an immersive experience or a non-immersive experience, it may feature a single media type or may combine multiple media types, it may be a permanent installation at a physical site or a temporary system that is commissioned to run for a specified period, and it may or may not provide a management interface for heritage experts to create and update content. A virtual museum system could also be based on one or more technologies such as the 2D web, 3D web, VR, AR, mobile technologies and game engines. By virtue of these differences, a taxonomy that classifies virtual museums based on *accessibility*, *level of immersion*, *media type(s)*, *permanence*, and *manageability* is proposed and the five dimensions of this taxonomy are highlighted as follows.

Accessibility refers to the manner in which the visit to the virtual museum is actualised. The visit could be online (over the Internet, using a web browser or mobile app) or on-site (at a physical location such as a museum or a kiosk). It is possible for a virtual museum to support both online and on-site visits. For instance, a system may facilitate a remote (online) virtual tour of a physical museum, with the option to embark on a location-aware, guided tour while at the physical museum space.

Level of Immersion refers to the amount of engagement with the virtual content experienced by visitors. Immersion is a function of the technology used in deploying the virtual museum. Immersive virtual museums (such as VR headsets or CAVEs) can induce a feeling of presence in the virtual environment, while non-immersive virtual museums (such as a 2D web gallery) maintain a perceptible separation between the visitor's physical environment and the virtual environment. Semi-immersive virtual museums (such as a spherical virtual museum tour) lie between the immersive and non-immersive virtual museums.

Media Type(s) refers to the content featured in the virtual museum system. Media types include 2D images, 3D images, spherical images, audio, 2D video, 3D video, 3D artefacts and 3D scenes. A virtual museum system may feature only one of these media types, or may combine two or more of them to provide content to visitors.

Permanence refers to how fixed the system is in a given physical space. A CAVE VR system mounted in a museum is considered permanent because it requires the installation of large screens and/or other display devices. A mobile AR virtual museum system may be considered semi-permanent if it requires fiducial markers to be placed in a physical space, such that it can only function in said space, and using it elsewhere would require porting the fiducial markers to the new space. A web-based virtual museum is considered non-permanent because it is not tied to any physical space (even if it is a replica of physical museum) and it can be accessed irrespective of location.

Work	Accessibility	Immersion	Media Type(s)	Permanence	Manageability
[197]	Online	Non-Immersive	Multiple	Non-Permanent	Non-Manageable
[198]	Online	Semi-Immersive	Multiple	Non-Permanent	Manageable
[199]	Online	Semi-Immersive	Multiple	Non-Permanent	Non-Manageable
[200]	On-site & Online	Non-Immersive	Multiple	Non-Permanent	Non-Manageable
[201]	On-site	Non-Immersive	Single (Video)	Semi-Permanent	Non-Manageable
[202]	On-site & Offsite	Non-Immersive	Single (3D artefacts)	Non-Permanent	Non-Manageable
[203]	Online	Non-Immersive	Single (3D artefacts)	Non-Permanent	Non-Manageable
[204]	On-site	Non-Immersive	Single (3D artefacts)	Semi-Permanent	Non-Manageable
[205]	Online	Non-Immersive	Multiple	Non-Permanent	Manageable
[206]	On-site	Non-Immersive	Single (3D scene)	Permanent	Non-Manageable
[207]	Online	Non-Immersive	Multiple	Non-Permanent	Manageable
[208]	On-site	Non-Immersive	Multiple	Semi-Permanent	Manageable
[91]	On-site	Immersive	Multiple	Permanent	Non-Manageable
[209]	Online	Non-Immersive	Single (3D artefacts)	Non-Permanent	Non-Manageable
[210]	On-site	Non-Immersive	Single (3D artefacts)	Permanent	Non-Manageable
[211]	On-site	Semi-Immersive	Multiple	Permanent	Non-Manageable
[212]	On-site	Immersive	Single (3D artefacts)	Permanent	Non-Manageable
[213]	On-site & Online	Semi-Immersive	Multiple	Non-Permanent	Manageable
[193]	Online	Semi-Immersive	Multiple	Non-Permanent	Manageable
[214]	On-site & Online	Immersive	Single (3D artefacts)	Permanent	Non-Manageable
[147]	On-site	Immersive	Single (3D artefacts)	Permanent	Non-Manageable
[215]	Online	Non-Immersive	Single (hypertext)	Non-Permanent	Non-Manageable
[196]	Online	Non-Immersive	Single (3D artefacts)	Non-Permanent	Non-Manageable
[216]	Online	Non-Immersive	Single (3D artefacts)	Non-Permanent	Manageable
[195]	On-site	Non-Immersive	Single (3D artefacts)	Non-Permanent	Non-Manageable

Table 2.4: An analysis of sample virtual museums using the proposed taxonomy features

Manageability refers to the availability of features that allow heritage practitioners to curate virtual exhibits. This entails creating, maintaining, updating and presenting content to visitors, often by means of a management interface which stores content and preferences in a data store from which content is retrieved and presented to visitors.

Overview of Related (Virtual Museums) Work

A virtual museum taxonomy that identifies related work in the public domain and classifies them based on the aforementioned criteria is presented in Table 2.4, and the virtual museum systems identified in the taxonomy are discussed below in reverse chronological order. This taxonomy is not exhaustive, but is rather a representative sample of virtual museum systems and approaches spanning three decades.

A virtual museum space dedicated exclusively to tourism is presented by Aurindo and Machado [197]. It features a web-based virtual museum which presents digital material to users, as well as a map interface that displays information based on location. Given the web-based nature of the virtual museum system, it is accessible online in the form of a web gallery that combines multiple media types. It features no immersive elements, nor does it provide a means of managing content tailored to heritage practitioners.

A framework for managing virtual museum exhibitions is discussed by Kiourt et al. [198]. The system uses 3D technologies and the Unity3D game engine to visualise content which is accessible over the web, and the system draws data from online sources including Google and Europeana, which are then stored in a globally-accessible back-end. Given the web-based nature of the virtual museum system, it is accessible online in the form of a 3D environment that combines multiple media types. The interactive 3D environment which users can navigate using an avatar provides for a semi-immersive experience and the ability to draw data from multiple sources provides a measure for content management by heritage experts.

A study conducted to evaluate the experience of older adults while engaging in remote virtual museum visits is discussed by kostoska et al. [199]. The findings suggest that the design of a virtual museum influences the nature of communication (between groups of users), exploration and navigation patterns. In addition, kostoska et al. [199] suggest that a virtual museum system with an interaction-free design, where users can consume content with no active interaction, can result in a positive experience and high usability. This provides rationale for the controller-free design adopted in chapter 5. It is worth noting that the virtual museum presented by kostoska et al. [199] facilitates a remote visit, contrary to the system designed in chapter 5 which allows visitors to explore the past on-site. Nonetheless, the design principles apply because the primary objective is to present snackable content to visitors thus providing an informative and engaging experience in little time. The virtual museum presented by kostoska et al. [199] is accessible online in the form of a web and mobile-based virtual tour with semi-immersive elements (using spherical images). It does not provide a means of managing content tailored to heritage experts.

A virtual museum for cultural heritage education is presented by Fevgas et al. [200]. The virtual museum (dubbed iMuse) features an on-site mobile tour facilitated by Radio Frequency IDentification (RFID) tags, an online virtual presenter facilitated by web technologies and a virtual wing for on-site use (though not necessarily at the actual site) facilitated by a temporary kiosk set-up. The combination of these three features – an on-site mobile tour, a web-based virtual presenter and an on-site virtual wing – enables the iMuse Virtual Museum to support both on-site and online use cases. The iMuse virtual museum does not feature immersive elements, and does not provide a means of managing content tailored to heritage experts.

A video-based virtual museum designed for on-site use, which generates video content in real-time and presents the content to users as they explore a physical museum space is presented by Hayashi et al. [201]. The virtual museum adopts an AR approach, where the video-generated content is superimposed on physical world content such as paintings. The virtual museum does not provide a means of managing content tailored to heritage experts.

A mobile virtual museum system for viewing 3D artefacts is presented by Neale et al. [202]. The system operates independent of location (although it was evaluated in a museum), hence it supports both on-site and offsite (albeit not online) use cases. It features no immersive elements, nor does it provide a means of managing content tailored to heritage experts.

A web-based virtual museum system for presenting 3D representations of artefacts using Web3D technologies (including X3D, X3DOM, SWF and PDF3D) is presented by Flotyński, Dalkowski and Walczak [203]. The system features no immersive elements, nor does it provide a means of managing content tailored to heritage experts.

A virtual museum system that is centred on mobile AR technologies is presented by Mata, Claramunt and Juarez [204]. The system is designed for use on-site in a physical museum setting. The system features no immersive elements, nor does it provide a means of managing content tailored to heritage experts.

A virtual museum that is coupled with a content management system is discussed by Rojas-Sola, Castro-Garcia and del Pilar Carranza-Cañadas [205]. The virtual museum features 3D models, VR and documents which are stored in a Database Management System (DBMS). The content of the database can be queried by user groups including researchers and curators, and it features a framework that enables heritage experts to create, manage and present content. It does not incorporate immersive elements.

A virtual museum that features realistic (human) virtual guides, aimed at increasing students' interest in science is presented by Swartout et al. [206]. The system is deployed in a physical museum to facilitate an on-site use case, and features realistic (albeit non-immersive) virtual environments and virtual guides. It does not provide a means of managing content tailored to heritage experts.

A Web3D-based virtual museum system that features a virtual environment with avatars and artefacts is presented by Zhang and Yang [207]. The system, which is powered by VRML and is accessible via a web browser, features audio clips for narratives, video capabilities for sharing ideas and multi user capabilities. It features no immersive elements, nor does it provide a means of managing content tailored to heritage experts.

A game engine-based system for creating and presenting virtual museums is discussed by Mateevitsi et al. [208]. The system uses the Torque Game Engine and combines multiple media types to provide on-site experiences to users. It also enables heritage experts to manage the content of virtual exhibits. It does not feature immersive elements.

A virtual museum system which extends the concept of the museum of pure form (proposed by Frisoli et al. [90]) is discussed by Carrozzino et al. [91]. The system uses a cylindrical, panoramic screen (albeit not a full cylindrical screen) to provide an immersive experience on-site, where users can interact with virtual sculptures with visual (text and images) and audio feedback. While this takes the form of a physical installation at a site, it is important to note that there need not be a relationship between the physical site at which the installation is hosted and the virtual site represented in the installation. The system features immersive elements and combines multiple multimedia, but it does not provide a means of managing content tailored to heritage experts.

The implementation procedure of a virtual museum of baskets is discussed by Isler, Wilson and Bajcsy [209]. It discusses the use of a laser scanner for data acquisition, as well as data analysis and data visualisation using VRML. The virtual museum features 3D artefacts and uses web-based technologies, but does not integrate other media types. It also does not feature immersive elements, nor does it provide a means of managing content tailored to heritage experts.

A virtual museum system that presents 3D artefacts to users on-site using augmented spherical images is presented by Huang, Chen and Chung [210]. The system uses an augmented panorama, a panoramic image representing a location (such as a museum room) which is juxtaposed with sets of images that depict interactive artefacts. It requires a camera mounted under a screen (hence it can only be used at specific sites and not remotely over the Internet), and it does not integrate multiple media types. It features no immersive elements, nor does it provide a means of managing content tailored to heritage experts.

The use of the Unreal Tournament game engine for the development of a virtual museum is discussed by Lepouras and Vassilakis [211]. The virtual museum combines photos, audio, video, 3D models and text descriptions for artefacts, and simulates the halls and rooms of a physical museum using in virtual environment. It features limited immersive elements, and it does not provide a means of managing content tailored to heritage experts. The findings suggest that game engines can offer viable experiences while engaging with virtual museums, and thus lends credence to the decision to develop and deploy virtual environments in museums using Unreal Engine 4 [217].

An immersive virtual museum installation that enables visitors to interact with 3D artefacts is discussed by Loscos et al. [212]. It uses an on-site, CAVE-like installation with a 3D stereo screen, VR goggles and haptic devices. While the system uses immersive technologies, it does not combine multiple media and it does not provide a means of managing content tailored to heritage experts.

A system for creating and presenting virtual museums is discussed by Wojciechowski et al. [213]

and Patel et al. [193]. The system incorporates VR and AR technologies and provides support for both on-site and online use cases. The system, dubbed Augmented Representation of Cultural Objects (ARCO), combines multiple media types and provides a management interface that allows heritage experts to create and present virtual exhibitions. The ARCO system is based on the (now defunct) Virtual Reality Modelling Language (VRML), and it does not incorporate immersive elements. The discussion by Patel et al. [193] also highlights other features of the virtual museum system designed in the ARCO project, and highlights the combination of multiple media, immersive technologies and a web-based system for online access towards the goal of creating, managing and presenting heritage content. The work done by Wojciechowski et al. [213] and Patel et al. [193] under the auspices of the ARCO project are aimed at supporting the dissemination of tangible cultural heritage artefacts by enabling museums and heritage experts to create and maintain digital exhibitions for both on-site and online (remote) use. The project also utilises Virtual Reality (VR) interfaces and combines multiple media types such as 3D artefacts and virtual galleries. A management interface also enables heritage practitioners to build and maintain exhibits.

The concept of a virtual museum and how the experience can be made substantial by relating a developed virtual museum to a corresponding physical museum is explored by Kwon et al. [214]. The system features an immersive CAVE-like structure, 3D environments and a web explorer to facilitate both on-site and online use cases. While the system uses immersive technologies, it does not combine multiple media and it does not provide a means of managing content tailored to heritage experts.

The development of a virtual museum installation for use on-site in a physical museum is discussed by Charitos et al. [147]. It simulates movement along the halls of a physical museum, as well as interaction with digitised artefacts. While the system uses immersive technologies, it does not combine multiple media and it does not provide a means of managing content tailored to heritage experts.

A web-based virtual museum system that provides personalised experiences to users, which is facilitated through unique login IDs (and passwords) is presented by Bertoletti, Moraes and da Rocha Costa [215]. The system – dubbed SAGRES – caters to the needs of three types of users who are identified as visitors, teachers and students. The system uses web-based media to facilitate online access (over the Internet) and it does not provide a means of managing content tailored to heritage experts.

The development of a virtual museum that enables visitors to simulate an online visit to a physical museum is discussed by Paolini [196]. The system features no immersive elements, nor does it provide a means of managing content tailored to heritage experts.

A virtual museum for building (creating and managing) and browsing (presenting) 3D artefacts online is presented by Ciabatti et al. [216]. It is powered by VRML for web access and does not incorporate other media types. While it provides a means of managing content tailored to heritage experts, it does not feature immersive technologies.

An onsite, desktop-based virtual museum featuring 3D artefacts is discussed by Tsichritzis and Gibbs [195], who also identifies the constraints that traditional museums face and how virtual museums can mitigate these constraints. The discussion also features useful considerations such as the combination of multiple media types, interactive 3D artefacts and remote access over a network. However, given that the work pre-dates the proliferation of personal computing and Internet access, the implemented prototype does not fully actualise the proposed design, and although the discussion is dated, the underlying principles and technologies – such as compelling computer graphics, remote network access, VR interfaces – remain relevant.

The ARCO project highlighted in the discussions of Wojciechowski et al. [213] and Patel et al. [193], part of which is also described by Petridis et al. [218] and White et al. [219], is arguably the most relevant work and may be considered a precursor to the Virtual Museum Infrastructure (VMI) proposed in this thesis, owing to the similarities in the approaches adopted, such as the combination of multiple media types and the provision of a management interface for heritage practitioners. Similarities aside, there are significant differences. Technology is rapidly evolving, thus given the datedness of the ARCO project (which ran in the early 2000s [193]), the technologies adopted are either obsolete or defunct. For instance, VRML is now defunct (replaced by X3D), and web-based 3D content is now driven by WebGL. In addition, web (and other) technologies have evolved since the implementation of the ARCO project, hence the VMI proposed in this thesis adopts state-of-the-art immersive, mobile, web and Virtual Reality (VR) technologies which are designed with affordability in mind and tailored for use by both heritage practitioners and community members. This is underscored by the use of high-fidelity 3D models, created using affordable photogrammetry set-ups, and spherical imagery coupled with affordable VR headsets for improved immersion and engagement. Furthermore, an underlying theme in this thesis is the leveraging of digital literacies of the public for heritage digitisation and dissemination. An example of this is the use of museum visitors' smartphones to provide content to them through mobile VR and trail apps (for example), and game controllers in the design of interactive museum exhibits. This approach has the added benefit that it is suitable to a broad range of museums, irrespective of where they fall on the resource spectrum (see Fig. 1.2, section 1.2) hence it can cater to the requirements of both museums with minimal resources and museums with substantial resources. The combination of affordable yet interactive, immersive technologies, which leverage digital literacies of both heritage practitioners and community members to facilitate the creation, curation, preservation and dissemination of heritage content, represents an improvement on the work in Wojciechowski et al. [213] and Patel et al. [193] and a contribution to the state of the art in virtual museum systems.

2.6 Chapter Summary

The concept of virtual museums has been discussed extensively. Three functions of museums have been identified as curating, preserving and disseminating heritage, and three groups of virtual museum users have been identified as heritage practitioners, community members and tourists. A technology taxonomy for digital heritage has been proposed to categorise existing work based on common technological approaches identified in the public domain, and the identified work have been categorised based on the scenarios in which they have been used to result in a use case taxonomy. A taxonomy has also been proposed to categorise virtual museum work based on accessibility, immersion, media types, permanence and manageability. These three taxonomies collectively provide an overview of the state of the art in the research area at the intersection of digital heritage and virtual museums. Finally, relevant literature has been discussed with particular focus on the use of VR and commodity technologies for heritage learning and exploration both in and out of museums, and these discussions have been situated in the context of the exploratory activities discussed in chapters 4 – 6.

CHAPTER THREE

RESEARCH METHODOLOGY

This chapter provides an overview of the hybrid approach adopted for this research. It highlights the case studies that are investigated and provides rationale for the essence and formulation of each case study, as well as how these case studies are designed to address the research questions and yield research contributions. The chapter concludes with a discussion of the preliminary design of a Virtual Museum Infrastructure (VMI).

3.1 Research Overview

The aim of this research is to further the understanding of how emergent technologies can promote the curation, interpretation and dissemination of heritage. This aim is actualised through the development of a Virtual Museum Infrastructure (VMI) which can support (and replicate) the functions of a traditional museum by facilitating the curation, preservation and dissemination of heritage. The methodology for this research combines underpinning theory and practice in the fields of Museology, Heritage Studies and Computer Science.

Underpinning Theory This research begins with a context survey to understand and scope out the meaning and functions of the traditional museum as an institution, its relationship with heritage, and the evolution of the concept of virtual museums in a digital heritage era. The investigation into underpinning theory is carried out through a review of related literature in the fields of digital heritage and virtual museums. This yields digital heritage taxonomies that identify common technological approaches and use cases, and a virtual museum taxonomy that classifies virtual museum systems along five dimensions. Furthermore, the scenarios in which virtual museum systems can exist, the user groups of virtual museum systems, as well as the needs of these users are identified, and these form the basis for the exploratory case studies and design goals of the VMI.

The research focuses on participatory heritage engagement using affordable, emergent technologies, and part of this investigation is carried out using case studies that entail designing, implementing and deploying quintessential components of the VMI as systems in physical-world scenarios, followed by physical-world evaluation exercises. Each case study features a ***system*** that is used in a specific ***context***, and is evaluated in terms of its ***impact***, i.e. how the ***system*** adds value in that ***context***.

Case Study 1 features 3D visualisation tools for exploring the past using mobile and desktop-based Virtual Reality (VR) devices (collection of VR ***systems***), used by early secondary school pupils (learning ***context***), to show that the use of such systems is a valuable addition to learning and offers new and exciting opportunities that can augment current pedagogical techniques (educational ***impact***).

Case Study 2 features an immersive museum installation that enables visitors to visualise and re-live the building space as it was in the 19^{th} century. It uses a VR headset tethered to a computer with a screen (VR ***system***) to provide interactive experiences to museum visitors (on-site tourism ***context***) and enables them to compare the past and present states of the building from equivalent vantage points, which fosters an appreciation of local history (heritage and organisational ***impact***).

Case Study 3 features a real-time, remote virtual tour system that combines an Unmanned Aerial Vehicle (UAV), a Raspberry Pi and smartphones in VR headsets (video streaming VR ***system***) to facilitate exploration of areas that are inaccessible (remote tourism ***context***) thus breaking geographical (space) barriers. This case study investigates the use of such a system irrespective of power and/or Internet access, making it practical for remote locations, and incorporates a social element in the exploration, such that it allows multiple people to take part and share the experience (collaborative heritage ***impact***).

These case studies have been designed to feature the quintessential components of virtual museum systems as highlighted in chapter 2, hence the systems featured in these case studies collectively represent the state of the art in the development and use of virtual museums. Furthermore, as discussed in section 2.3, a virtual museum can be: (1) a digital representation of artefacts which can be accessed either remotely or on-site, or (2) a simulation of a visit where users find themselves traversing the halls and rooms of an actual museum [30]. By virtue of this, a virtual museum system can support: (1) both remote and on-site exploration, or (2) only on-site exploration, or (3) only remote exploration. This informs the nature of the case studies. Case study 1 (chapter 4) presents a system that enables users to explore the ruins of a monumental cathedral both in the classroom (remotely, using desktop-based VR systems) and at the ruins (on-site, using mobile-based VR systems). Case study 2 (chapter 5) presents an immersive museum installation which enables users to compare the past and present states of a site from equivalent vantage points. Case study 3 (chapter 6) presents a system for exploring remote landscapes without power and/or Internet access.

Museum Workshops A series of workshops is conducted with museums across Europe, Latin America and the Caribbean to investigate the feasibility of a VMI. The workshops are aimed at empowering museum staff and community members with the technical skills required to digitise and disseminate heritage content. The workshops also yield digital output such as 3D artefacts, virtual tours, multimedia and metadata for content curation.

VMI Design The findings of the context survey, case studies and museum workshops are combined to design and implement a Virtual Museum Infrastructure (VMI) – a suite of tools that combines affordable, mobile and web technologies to disseminate and visualise heritage geographically, categorically, descriptively, and exploratorily. The VMI is tailored for participatory heritage engagement (virtual museum ***context***) and it enables museums to preserve, curate and disseminate tangible and intangible heritage in a collaborative and engaging manner using commodity devices (community heritage ***impact***). The VMI design accounts for the variation in museum resources (as demonstrated by the workshops) by providing affordable techniques that leverage existing resources and digital literacies towards the left end of the resource spectrum (see Fig. 1.2) where resources may be scarce, and immersive on-site exhibits that facilitate heritage exploration in engaging and exciting ways towards the right end where resources are available. Entrenched in the design of the VMI is the ability to facilitate active use and management by stakeholders, support the reuse of digital heritage content, provide engaging experiences by combining multiple media and delivery platforms, and maximise the resources that are available to museums. These design goals are based on the functions that museums perform on heritage content, the groups of heritage users, and the needs of these users, all of which are uncovered during the context survey in chapters 1 and 2.

VMI Implementation, Deployment and Evaluation The VMI design is implemented as a multifunctional system. A web-based management portal allows heritage experts and community members to create, curate and manage content. A remotely-accessible data store is implemented to hold this content, which ensures the preservation of content in the event that the original artefacts or objects are lost or damaged. The content held in the data store can also be pushed to social archive sites to expand the reach of the system. To facilitate the interactive dissemination of content remotely, web-based presentation interfaces, mobile apps and immersive on-site exhibits are implemented so that users can be immersed in virtual tours of remote locations, compare the past and present states of a heritage site, and interact with 3D artefacts. The VMI is deployed in over 20 museums worldwide and evaluated in terms of its impact and value for heritage management, and the findings demonstrate the feasibility of digital heritage systems and virtual museums for participatory heritage engagement.

3.2 Research Through Practice

The field of Computer Science is both theoretical and experimental [220]. Experimental Computer Science is concerned with the design of hardware and software systems for the purpose of investigating and or solving problems [221, 222], while Theoretical Computer Science, which is grounded in Mathematics, is concerned with modelling and optimising the underlying processes and components of computer systems [223]. An alternative school of thought views Computer Science as an empirical (rather than experimental) discipline, because unlike physics, chemistry and other natural sciences, some observational techniques that constitute Computer Science research do not conform wholly to the experimental method [224]. Nonetheless, these activities involve experimentation, because (hardware and software) systems that are designed or built to solve problems are experiments in the sense that they impose themselves on their immediate environment, their potential users, and/or the wider society, and by so doing they pose questions that require answers [224].

Theoretical computing research is conducted using formal methods, where research activities entail formulating and proving (or disproving) theories about models, algorithms and processes [223]. While these often form the basis for the implementation of software and hardware, formal methods are of little relevance to this research. Conversely, empirical research techniques are adopted in this research to investigate and further our understanding of computers as applied to museology and heritage studies. These techniques include design, experimentation and observation of both technical phenomena (to optimise systems for tasks) and socio-technical issues (in order to make informed decisions about how to use systems).

The application of computing research paradigms to museology and heritage studies makes this research interdisciplinary. A central concept that facilitates this interdisciplinary research is practice: the practice of curating, preserving and disseminating heritage content, as performed by museums and heritage organisations. Practice-orientated research may be practice-based, where the research results in creative output (such as a sculpture) as the basis for the contribution to knowledge, or practice-led, where the research activities result in a better understanding about practice (such as sculpting techniques) in the field [68]. The interplay between empirical research techniques and museology practice in this research results in a hybrid methodology that demonstrates the design and implementation of affordable virtual museum systems (as in practice-based research), and to contribute to the knowledge of how technology can support museum functions, through the evaluation of these systems in both laboratory conditions and physical-world, museum contexts characterised by heritage learning and management activities (as in practice-led research). This hybrid methodology can be viewed as a form of experimental systems research where the methodology involves designing a system in response to a problem specification, followed by testing, data collection, fine-tuning and deployment. Examples of previous use of such hybrid methodology include computer networks research conducted during the advent of communication protocols and switching-based infrastructure that form the backbone of the Internet [225, 226, 227, 228, 229, 230]. Such systems research methodology may result in the creation of a viable artefact such as a computer network, hardware prototype or software, which is akin to the deployment of a creative artefact like a painting, sculpture or film when practice-based research is conducted in the arts. While research through practice is more popular in the arts and social sciences, where it is used to demonstrate the acquisition of, and contribution to knowledge through creative artefacts, it is also applicable to Computer Science, especially when Computer Science constitutes a crucial part of an interdisciplinary endeavour which seeks to understand how technology can serve society. The adoption of practice-orientated research with community organisations has the added advantage that it allows for tailoring the investigation to the needs of stakeholders [231]. However, it also requires collaboration with these stakeholders, which takes time and effort to build, and there is a tendency for decision making to be hampered by biases or preconceived notions held by stakeholders [231]. Time, effort and resource requirements are also hindrances to the execution of empirical research methods that are conducted in physical-world environments [232]. Conversely, system experiments can be conducted in laboratory environments usually with less time and monetary expense and allow for greater control of system variables. However, the results obtained may not be representative of how similar systems will perform in physical-world scenarios characterised by less control and noisy data [233, 234, 235]. Therefore, the combination of empirical, exploratory and practice research methods in this research allows for the delivery of viable systems that demonstrate pertinent concepts, while simultaneously facilitating the investigation of these systems in close proximity to the user groups that they are meant to serve. Laboratory investigations are conducted while the systems are iteratively developed, strategic partnerships are formed and leveraged to deploy these systems in physical-world scenarios, research questions that allow for a reflection on the outcomes of practice have been formulated, data are gathered and analysed using quantitative and qualitative techniques, and these systems are evaluated empirically with users and stakeholders.

3.3 Research Evaluation Techniques

The hybrid methodology adopted for this research necessitates a hybrid approach to evaluation, which entails a combination of user studies in natural environments, assessments with heritage experts in controlled environments and system measurements in laboratory environments. While user studies conducted in natural environments may require more time and effort, and are susceptible to more data noise as compared to laboratory-based user studies, the natural environments provide more ecological validity, which yields findings that are more representative of physical-world practices [233, 234, 235]. For this reason, user studies in this research are predominantly conducted in the wild – in classroom learning environments (chapter 4), in local community parks (chapters 5 and 6) and in museums (chapters 5 and 8). Where user studies are conducted in museum spaces, multiple data gathering techniques are employed, including collating feedback gathered using questionnaires over several months provided by museum visitors (walk-ins) after using on-site exhibits, conducting open days where events are designed to entice the public to visit the museum and take part in structured evaluation exercises, and conducting studies with focus groups comprised of pre-selected museum visitors. To improve reliability and validity, the questionnaires were piloted before use. For instance, to evaluate usability of the VR system in chapter 4, a System Usability Scale (SUS) questionnaire (proposed by Brooke et. al. [236]) was originally adopted, but after piloting with secondary school teachers, it was significantly modified to make it more appropriate (terminology-wise and length-wise) for first year secondary school pupils who were the intended respondents. While video recordings constitute a valid option for gathering additional data, no recordings are obtained as part of the data gathering process so as to mitigate the observer effect (a phenomenon which may cause participants to talk, act or behave differently due to the presence of recording equipment and/or obtrusive observers), as well as the unsavoury perception and ethical implications associated with recording participants in physical-world settings [237, 238, 239, 240]. As a replacement for video analysis, in-person observations with detailed notes are used during museum open days, as well as detailed notes of participant comments during focus group sessions. These user evaluations are supplemented with expert evaluation exercises, in line with Karoulis, Sylaiou and White [241] who advocate for the combination of domain expert evaluation – which is less resource intensive and provides qualitative insights – with user-based evaluation – which provides reliable data on how systems perform from the perspective of the target audience. The expert evaluation exercises take the form of semi-structured, pre- and post-deployment interviews conducted with museum curators and managers to gauge expectation and satisfaction levels respectively, and focus group-like discussions with stakeholder committees to assess value and impact. Where interviews were used, for instance in chapters 5 and 8 as part of the evaluation of the Curing Yard and the Picts & Pixels exhibits respectively, the interviews were designed and conducted in a semi-structured manner. This approach entails using a series of questions – designed to elicit data on usability, value, acceptance and perceived impact – as the basis for the interview, with flexibility for the discussion to evolve to allow for the discovery of additional information. Where focus groups were used, for instance in chapter 5 as part of the evaluation of the Curing Yard exhibit, the focus groups were organised by recruiting participants through publicity campaigns and word of mouth. The campaigns were designed to publicise the annual conference which takes place at the museum, but also coincided with the planned launch event for the exhibit. This presented an opportunity to invite attendees to express interest in taking part in the focus groups, which were run while the participants, together with a moderator stood around the exhibit. Laboratory-based technical evaluations are also conducted during the iterative design and implementation phases of system components in order to push the limits of technology platforms, compare and contrast two or more platforms and perform quantitative measurements to ascertain system capabilities. This combination of qualitative and quantitative evaluation techniques – known as mixed methods research – has grown in popularity in recent years, and has seen adoption in various fields including medicine [242], education [243], and technology [244]. Mixed methods allow for the combination of data so that the research is not restricted to numbers in a quantitative sense or words in a qualitative sense; rather, the use of mixed methods enables participants feedback to be placed in numbers [245, 246]. Although the use of mixed methods has been adopted to leverage the strengths of quantitative and qualitative data as applied in practice-oriented research, this approach is not without limitations. One limitation of the adopted methodology is the use of a series of exploratory case studies to investigate the feasibility and value of a VMI, at the expense of limited (and focused) treatment of each case study. Each of these studies could have been developed to a greater level of rigour, but a balance was struck between the further exploration of these work and the iterative design and implementation of the VMI. While this is a limitation of the overall research on one hand, on the other hand it yields benefits in terms of validity and reproducibility, as consistent findings were observed across the individual case studies explored. It is pertinent to note that while there were secondary objectives included in the exploratory work discussed in chapters 4– 6, the primary objective, and consequently the findings of these work revolve around developing an understanding of the problems faced by heritage organisations as well as evaluating the feasibility of virtual museums and digital heritage components, which ultimately feed into the design and implementation of the VMI discussed in chapter 8. Another limitation of the work stems from the use of questionnaires in places, which have limited validity and reliability. To overcome these limitations, the questionnaires were designed in line with the work's objectives and were then piloted appropriately. Furthermore, the results of the work do not rely solely on these questionnaires, as a combination of focus groups, expert interviews and observations have been used to gather data; this combined method of data collection ensures the validity of the research instruments and the reproducibility of the findings. A third limitation of the methodology stems from the use of real-world evaluation exercises – which is qualitatively different from a single system evaluation or social science study; this resulted in limited control over the evaluation environment and subsequent introduction of noise into the data. Nonetheless, the use of in-the-wild evaluation exercises has resulted in representative findings with real-world applications.

In designing this methodology, steps have been taken to improve the instruments' *validity* – the extent to which the research instruments measure what they were intended to measure – and *reliability* – the extent to which the research findings are reproducible. In line with Riege [247] who examines validity and reliability in case studies and qualitative research, this methodology has adopted the use of multiple sources of evidence (for instance by combining quantitative Likert-scale data with qualitative interviews, focus groups and observations), as well as an interpretation of the data alongside domain (heritage) experts to improve *construct validity*. In addition, *internal validity* is improved through an examination of the research in the context of theoretical frameworks (such as Foni's taxonomy for cultural heritage visualisation [26]), and through the use of peer debriefing exercises, while *external validity* is improved through a definition of the research scope and boundaries, an examination of evidence to check consistency with literature, and the use of thematic analysis as a specific procedure for coding and analysing qualitatively-generated data. In addition, the deployment, evaluation and consequent large-scale acceptance and approval of the designed systems in real-world settings further makes a case for *validity*. The practice-oriented methodology adopted in this research takes strengths in iteratively designing, deploying and evaluating systems in the environments in which they will eventually be used and thus provides actual, representative findings, but on the other hand it makes *reliability* much harder to demonstrate than simulated, lab-based studies which introduces assumptions and abstraction layers. Nonetheless, to improve *reliability*, the research process was examined and documented for internal consistency through an audit trail that continuously applied checks and bounds to clarify theoretical positions and minimise biases within theoretical frameworks. Furthermore, where available, word-for-word accounts and observations of respondents' actions (and expressions) were examined against data collected through other means (such as questionnaires) to confirm and/or support interpretations, in line with Noble and Smith [248]. The combination of these audit processes, theoretical grounding of evidence, scope definition, and data triangulation establish *validity* and *reliability* of the research.

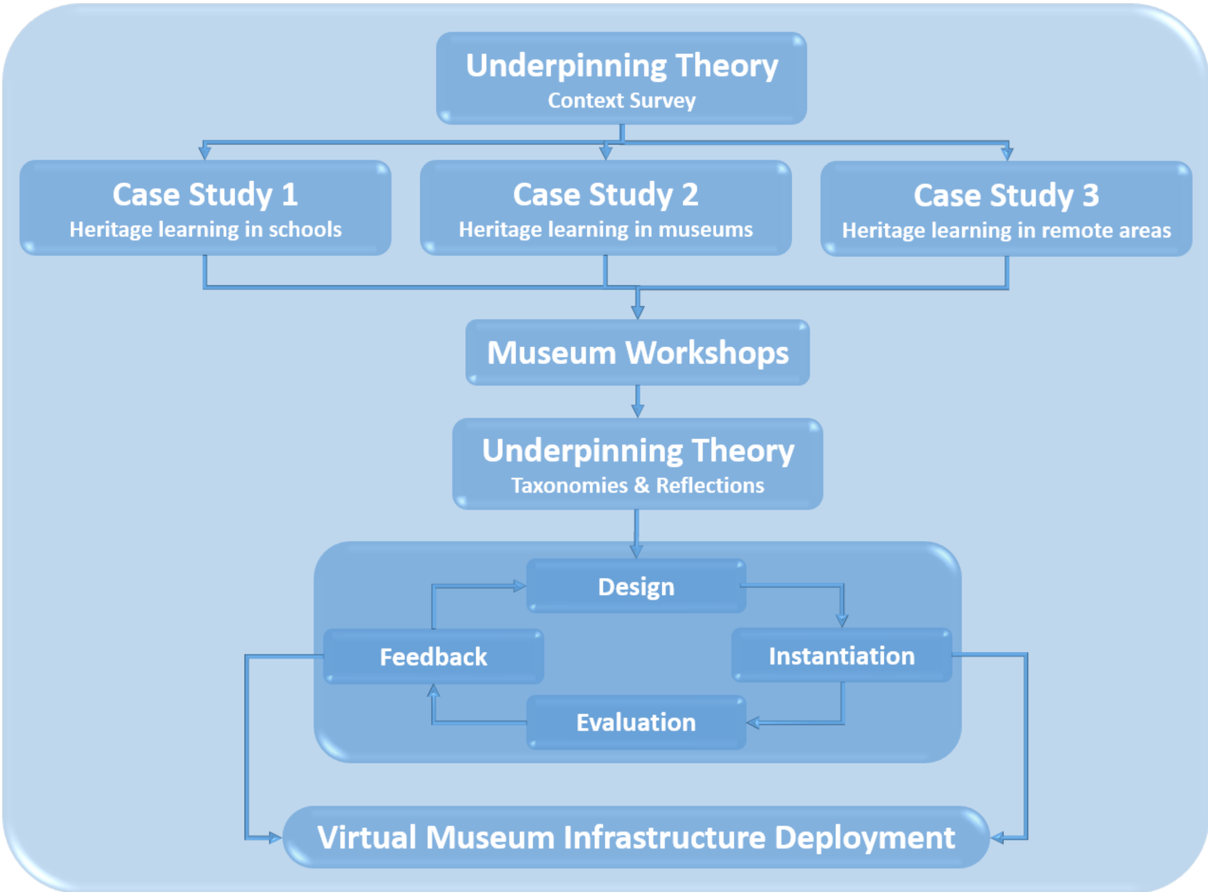

Figure 3.1: Research Overview

3.4 Research Questions

An iterative approach is adopted for the design, implementation and evaluation of the systems in the case studies. This evaluation is conducted in terms of the value and usability of the systems in the identified contexts, such that the system components are representative of the affordable, emergent technologies that constitute the building blocks for the Virtual Museum Infrastructure (VMI). It is therefore worth investigating how efficacious these systems are for heritage management, and this is the motivation for research question Q1.

> *Q1: How can emergent technologies support (and replicate) the functions – heritage curation, preservation and dissemination – of traditional museums?*

To address this question, section 3.5 proposes an architecture and design of a Virtual Museum Infrastructure (VMI) that enables heritage practitioners and community members alike to curate, preserve and disseminate heritage content. An instantiation of this system is implemented and evaluated in chapter 8, in which 3D workshops are held with museums across three continents.

The workshops provide an avenue for the collaborative creation of content (where community members and heritage experts work together to digitise artefacts using affordable and readily-available technology). The created content is described and presented using an instantiation of the VMI ***system***, which is evaluated in the ***context*** of value and usability. The results demonstrate the feasibility and value of the Virtual Museum Infrastructure (VMI) (heritage curation and dissemination ***impact***). Fig. 3.1 shows the relationships between the case studies and how they build towards the development of the VMI.

The design of the Virtual Museum Infrastructure (VMI) is tailored to the needs of museums, which may be constrained by resources (such as operating funds and formal infrastructure), hence an approach to heritage dissemination which utilises existing infrastructure and leverages digital literacies in order to maximise value for users while minimising cost for heritage practitioners is paramount. This is the motivation for research question Q2.

> *Q2: How can heritage dissemination systems be designed to maximise value in environments constrained by limited (computing and funding) resources?*

To address this question, chapter 4 investigates how high-fidelity virtual environments can be depicted using spherical images on a smartphone with an affordable VR headset. This uses affordable technology in people's pockets, coupled with headsets that can be made out of cardboard to deliver immersive experiences. Furthermore, an evaluation is carried out to investigate how the mobile-based VR system compares to a desktop-based system which features a VR headset tethered to a computer, coupled with a game controller for exploring a full 3D model of a historic site. In addition, chapter 5 investigates how a controller-free design can be adopted for a museum installation to deliver snackable content to visitors. This installation uses a VR headset and spherical images to immerse users in a virtual environment that represents an erstwhile view of the same vantage point. This solution represents a viable alternative to deploying a resource-intensive 3D environment on a mid- to high-end workstation, which demonstrates how museums can minimise cost and resources while maximising experience for visitors. The design features introduced in chapters 4 and 5 are then adopted in the Virtual Museum Infrastructure (VMI) which is implemented in chapter 8, with a focus on how the VMI can cater to the varying needs of museums at different points on the resource spectrum (see section 1.2, Fig. 1.2), thus maximising value and resource utilisation in constrained environments. In a similar vein, the use of a VMI for participatory heritage engagement in environments which may be constrained by limited resources necessitates the investigation of how digital literacies can be transformed into additional resources that can improve the experience of museum visitors. This thesis therefore investigates how digital literacies – specifically smartphone ownership and gaming proficiency – can be used for participatory heritage engagement. This is carried out in chapter 8 which shows how visitors' devices can be used in lieu of dedicated infrastructure by delivering content via VR and trail apps that can downloaded on smartphones and tablets, and also explores the use of game controllers, menu systems and navigation paradigms that visitors may be familiar with to foster engagement with heritage content.

The disruptive nature of technology necessitates an investigation of the transformations (for better or worse) that the adoption of technology will have on museums practice and heritage practitioners. This provides the motivation for research question Q3.

Q3: How do technological advances transform museums practice?

To address this question, this thesis investigates the impact of technology on heritage practice from the perspectives of the practitioners (such as museum staff) and the end consumer (such as museum visitors). The impact on the general public is investigated through a combination of questionnaires and semi-structured interviews to elicit feedback on a variety of heritage dissemination systems as discussed in chapters 4 to 8, while the impact on heritage practitioners is investigated through a series of pre- and post-deployment interviews with heritage experts to gauge expectation and satisfaction levels respectively as discussed in chapters 5 and 8. In addition, chapter 8 features an on-site exhibit in which both digital and physical content coexist in the same physical space, and this design is aimed at investigating public perceptions with respect to the differences and relationships between digital heritage content and physical heritage content. The findings demonstrate that advances in technology can bring about a change in mindset regarding: (1) the roles of digital content in museums as entities in their own rights, as opposed to mere supplements or replicas of physical content, and (2) the onus for heritage management, as one that is the collective responsibility of the community as opposed to the responsibility of a few heritage practitioners.

3.5 The Virtual Museum Infrastructure (VMI)

The Virtual Museum Infrastructure (VMI) is a suite of tools designed to facilitate, support and replicate the functions of traditional museums, namely the *curation*, *preservation*, and *dissemination* of heritage. Section 2.3 identifies three groups of heritage users and discusses their characteristics and needs. *Heritage practitioners* seek to curate, preserve and disseminate content, *community members* seek to preserve and disseminate content and *tourists* seek to consume heritage content that is disseminated in engaging and informative ways. The design of the VMI should support the curation (including creation and description), preservation (study and conservation) and dissemination (communication and presentation) of heritage content. It should enable heritage practitioners to easily and intuitively input content into the system and

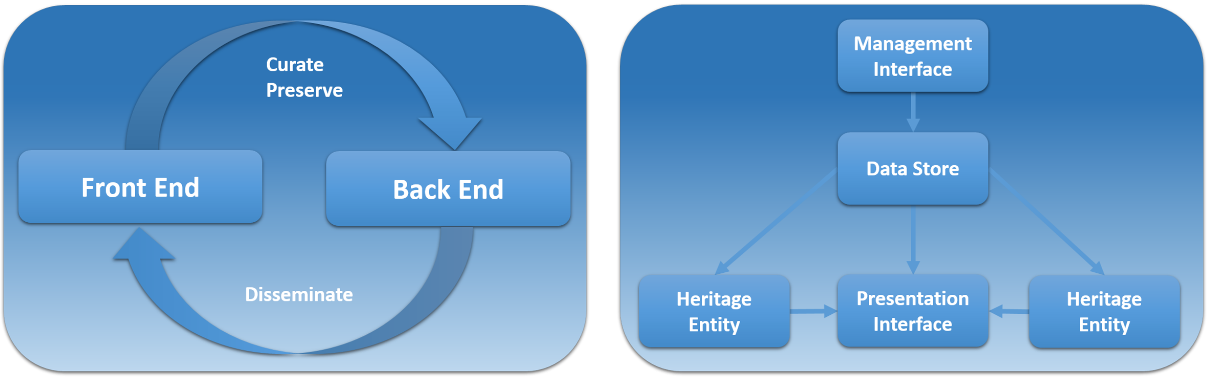

Figure 3.2: VMI preliminary design (left) and components (right)

create multiple representations of the content. During (and after) the creation of content, heritage practitioners and community members should be able to add metadata to describe the content. The content created and associated metadata should be stored and organised for easy retrieval and presentation. These serve as the basis for the design of the VMI.

3.5.1 Design

As stated in section 2.3 (and illustrated in Fig. 2.2), three groups of virtual museum users are identified in this thesis based on their interests and needs – *heritage practitioners*, *community members* and *tourists*. Furthermore, the information-seeking behaviour of virtual museum users is characterised by a desire for visual exploration, serendipitous discovery, detail gathering and sense making [249, 250]. In other words, users need to be able to interact with content (by zooming, clicking, panning and so on), explore content in an unstructured manner, look for specific information, and make sense of the information found. Hence the museum functions – of curation, preservation and dissemination of heritage – coupled with the needs and characteristics of users as they relate to these functions, give rise to the following design goals for the VMI:

1. **Active Use and Management**: The VMI should cater to the interests of different user groups in encouraging active use and facilitating continuous management of content. In addition to creating content, heritage practitioners should be able to continuously update and modify previously-created content, which should be instantly available to the community. This goal is central to the design of the VMI because the creation and management of content should not be limited to heritage practitioners, given the significant benefits that accrue when heritage organisations encourage public participation and community engagement through contribution, collaboration and/or co-creation [31].

2. **Content Reuse**: The VMI should support the reuse of content in different scenarios, so that

users can upload content once which can then be re-purposed for different platforms such as mobile, web and on-site use, without having to explicitly supply data for each platform. This is akin to traditional museum archives which store and catalogue artefacts that can be reused in multiple exhibitions. The reuse of content can help make the financial investments more lucrative and may provide a counterargument to high cost serving as a deterrent to the adoption of technology [71].

3. **Engaging Experiences**: Content consumption using the VMI should be engaging and interactive, to encourage active use and participation in museums. This is central to the theme of this thesis and is achieved through the use of novel, immersive technologies, as technology has been demonstrated to provide opportunities for situated, personalised and collaborative heritage learning, as well as improved engagement with content [251, 252, 253].

4. **Multimedia Integration**: To further enrich the user experience, the VMI leverages multiple media types – Wiki (electronic text), audio, video, images, spherical media, 3D artefacts, 3D scenes and immersive technologies – to provide immersive and interactive experiences to users. As discussed in section 1.1.3, multiple media types can be combined to leverage the strengths and characteristics of each media so as to provide additional or alternative sensory input (in cases of visual or auditory impairment for example) and to improve the user experience.

5. **Resource Maximisation**: To cater to the varying needs of museums along the resource spectrum (Fig 1.2) the VMI design should incorporate both affordable low-end and mid- to high-end systems for heritage dissemination. It should also utilise the digital literacies of potential visitors as an additional source of resources. This recognises that museums possess varying levels of financial and human resources, and may be plagued by a host of challenges that inhibit the successful adoption of technology [32, 33]. The existing skills and resources of the target audience can be leveraged to provide engaging and meaningful experiences both within and outwith museum spaces.

3.5.2 Implementation and Evaluation

The identified design goals give rise to a system architecture which features a back end for data storage, as well as a front end for data curation and presentation as shown in Fig. 3.2. An overview of the system components is presented in Fig. 3.2. The components include a data store to hold and organise data for retrieval, a management interface for creating and updating content that is held in the data store, a presentation interface for delivering content to users and representations of heritage entities in the system. Chapters 4 – 6 explore how these heritage entities can be managed and disseminated using the Virtual Museum Infrastructure (VMI), and a comprehensive discussion of the design and instantiation of the VMI is provided in chapter 8. The components of the VMI are instantiated as both online and on-site exhibits that are collectively used by over 20 museums in 9 countries on 3 continents. The evaluation is carried out to investigate value and impact from the perspectives of both heritage experts and museum visitors. Data are gathered from the museum workshop participants, as well as museum visitors during open-day events, and are analysed using a combination of quantitative and qualitative techniques (see section 3.3).

3.6 Chapter Summary

The hybrid methodology which involves the iterative design, deployment and evaluation of systems in physical-world environments, aimed at investigating the research questions identified has been discussed. A preliminary design of a Virtual Museum Infrastructure (VMI) for participatory heritage engagement has been proposed to fulfil the functions of museums based on the needs of virtual museum user groups identified in chapter 2. The next part of this dissertation discusses three exploratory case studies that have been designed to investigate heritage exploration in learning environments (both in the classroom and at the remains of a heritage site), in museum spaces and in remote environments.

Part III

Exploratory Work

4

CHAPTER FOUR

IMMERSIVE HERITAGE EDUCATION

This case study investigates the use of mobile Virtual Reality (VR) systems as an affordable tool for cultural heritage learning by early secondary school pupils using the St Andrews Cathedral as the subject matter. The findings suggest that, regardless of prior experience, participants find VR systems easy to use for learning, and that mobile headset-based VR systems stimulate interest in learning history more than screen-based VR systems; this makes a case for VR-based heritage learning. The findings also suggest that there is no significant difference in the levels of immersion provided by affordable, mobile VR systems as compared to mid- to high-end VR systems. The implication of these findings is that, in the context of heritage learning, mobile VR systems are comparable to desktop-based VR systems in terms of levels of engagement and immersion, yet mobile VR systems are significantly cheaper, hence they are included in the design of the Virtual Museum Infrastructure (VMI) to further the goals of facilitating engaging experiences for the general public and providing more affordable alternatives to heritage organisations where resources are in short supply.

Relevant publications:

1. **Adeola Fabola**, Alan Miller, and Richard Fawcett. *Exploring the past with Google Cardboard*. In Digital Heritage, 2015, volume 1, pages 277–284. IEEE, 2015. [1].

   Contributions: Design, implementation and preliminary evaluation of an affordable, mobile, VR system for heritage exploration.

2. **Adeola Fabola** and Alan Miller. *Virtual Reality for Early Education: A Study*. In International Conference on Immersive Learning, pages 59–72. Springer, 2016. [2].

   Contributions: Design and implementation of two mobile, VR systems for heritage exploration, and a user study aimed at performing a comparative analysis of mobile-based and desktop-based VR systems in the context of heritage learning.

4.1 Introduction

St Andrews Cathedral represents an iconic monument in Scottish history, as it was the centre of religious activities and domicile of eminent figures in its heyday. Owing to its importance, the Open Virtual Worlds Group at the University of St Andrews embarked on a reconstruction of the Cathedral as it stood in 1318 (see Fig. 4.1), and this reconstruction has been a valuable resource in exploring and teaching cultural heritage through deploying it in schools, museum installations and over the Internet [77]. The reconstruction has served as the basis upon which other resources have been built and has facilitated a series of educational activities such as pilgrimages and excursions to the Cathedral site, reenactments of historic events, and classroom-based writing, drawing and illustration, as part of the Curriculum for Excellence that is espoused by the Scottish Government. The reconstruction of St Andrews Cathedral and its subsequent use to support and facilitate educational activities is part of a larger body of work that leverage digital media and Virtual Reality (VR) systems for exploratory and experiential learning. Getchell et al. [82] discuss the use of a VR-based system for teaching and learning about archaeological excavations with third year university archaeology students, and the findings of evaluation exercises conducted as part of their study validate their premise that situating students in immersive environments where they explore content and perform tasks facilitates a better understanding of the subject matter. In a similar vein, Davies, Miller and Allison [79] discuss the use of a Virtual Time Window (VTW), a VR system which facilitates exploration at the remains of cultural heritage sites by matching the physical-world orientation with the orientation of a virtual environment that is viewed through a tablet device. The idea of matching orientation between a physical and virtual environment is taken further in the design and implementation of a Head Mounted Display (HMD)-based Cross Reality (XR) system for real-time exploration of a 15th-century chapel which has changed over time, and the findings of evaluation exercises conducted reveal an increased appreciation for cultural heritage content when the past and present can be compared through tandem exploration [98, 169]. While the value of immersive VR for exploratory learning has been established [82, 98, 169], there remain issues of cost, logistics and scale [71]. A common denominator of the aforementioned VR-based exploratory work is the relatively-high resource requirements, manifested in the form of workstations with high-end processors and graphics cards, networking infrastructure and tethered VR headsets. This chapter therefore discusses the design, implementation and evaluation of a system for on-site and classroom-based exploration of St Andrews Cathedral which works with the Google Cardboard, an affordable VR headset for mobile devices [58, 59]. A study carried out to investigate the efficacy of the affordable, mobile system – as compared to the more resource-intensive alternatives – for heritage dissemination and exploratory learning is discussed in this chapter. The efficacy of the mobile system is investigated in terms of its value for heritage learning as perceived by students and teachers, its ability to

Figure 4.1: St Andrews Cathedral Cloister: 1318 representation (left) and today (right)

stimulate more interest in local history, and its ability to immerse users and encourage them to engage with virtual content, with the premise that the system constitutes an affordable, yet viable alternative for exploratory learning if it is perceived to be as efficacious as the more expensive VR systems. This study is conducted in the context of a secondary school module on local history, which combines classroom-based learning activities with excursions to the remains of the Cathedral site in the Scottish city of St Andrews. Section 4.2 discusses the motivation for this case study, highlighting how it contributes to the research objectives and fits with the theme of the thesis. Section 4.3 discusses the methodology adopted for this case study. Section 4.4 describes the system design and implementation activities. Section 4.5 discusses the classroom, and on-site exercises, section 4.6 discusses the data analyses and results, section 4.7 reflects on the findings of the work and section 4.8 provides concluding thoughts.

4.2 Motivation

The motivation for this case study stems from the need for an affordable system that can facilitate immersive heritage dissemination. The affordability requirement arises because of the funding and other resource pressures often experienced by the museums that engage in participatory heritage engagement. As discussed in section 1.2, museums may have little or no operating funds and may thus be constrained by limited resources, hence cost is a major determinant of the feasibility of the proposed approach. This approach is actualised through the design of a mobile-based VR system as an alternative to mid- to high-end desktop-based VR systems. The choice of mobile technologies is made because of the proliferation of smartphones and the availability of affordable VR headsets for mobile devices such as the Google Cardboard [58, 59] which can be purchased for £10. The Google Cardboard has grown in popularity since its release and there has been interest in exploring (as well as comparing and contrasting) interaction paradigms for

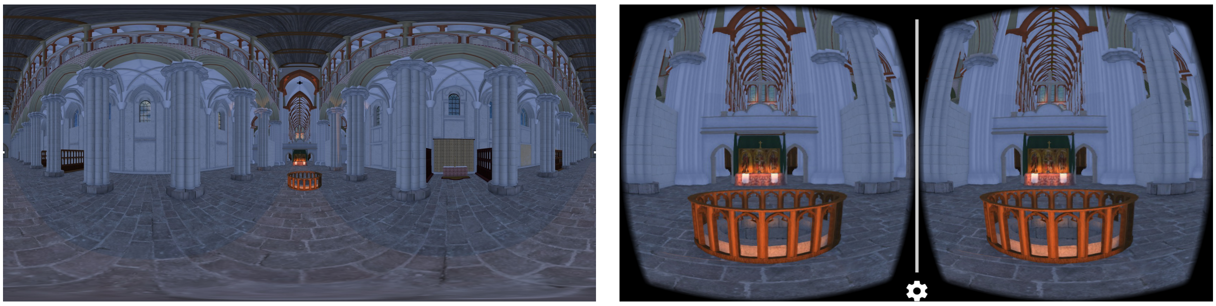

Figure 4.2: Cathedral Nave: Equirectangular projection (left) and view in VR mode (right)

System	Cost	Processing Power	Physical World Mobility	Virtual World Mobility
Google Cardboard	Low	Low	High	Static
Samsung Gear VR	Medium	Medium	High	Static
Oculus Rift + Xbox Controller	High	High	Low	Dynamic
Screen + Xbox Controller	High	High	Low	Dynamic

Table 4.1: Features of VR systems used

such mobile-based headsets; and the findings of a survey conducted by Yoo and Parker [254] suggest that controller-free techniques such as gaze- and hotspot-based navigation are more popular than techniques that rely on external, hand-held controllers. In addition, from 2013 to 2015, smartphone ownership figures increased by over 25% in emerging economies such as Chile and Brazil, and as at 2015, over two-thirds (67%) of the inhabitants of developed countries such as South Korea, Australia and Israel owned a smartphone [66]. This raises the possibility of leveraging existing digital literacies of the general public and significantly minimises cost for heritage organisations, because visitors can walk into a museum with their smartphones, download an app and use a £10 headset to immersively engage with heritage content. The value of the proposed mobile VR system remains to be seen, hence this case study investigates the efficacy of this system in delivering interactive experiences to users, as compared to mid- to high-end VR systems. This case study has therefore been designed to address a research question (see section 1.4) which asks how heritage dissemination systems can be designed to maximise value in environments constrained by limited resources. This investigation is conducted in an educational context – facilitated by the annual visit of the Open Virtual Worlds group to a local secondary school – with first year secondary school pupils aged 11-13 years at the remains of a heritage site and in a classroom setting. These school pupils constitute an appropriate audience for this investigation because young people tend to be inquisitive, enthusiastic and forthcoming with feedback.

System	Development Platform	Running Platform	Content
Google Cardboard	Android OpenGL	Mobile	Photosphere tour
Samsung Gear VR	Unity 3D	Mobile	Photosphere tour
Oculus Rift + Xbox Controller	OpenSim	PC	3D Model
Screen + Xbox Controller	OpenSim	PC	3D Model

Table 4.2: Platforms and content used

System	Task	Activity	Freedom
Google Cardboard	Structured	Guided tour	Low
Samsung Gear VR	Structured	Guided tour	Low
Oculus Rift + Xbox Controller	Unstructured	Free-form exploration	High
Screen + Xbox Controller	Unstructured	Free-form exploration	High

Table 4.3: Nature of activities performed

4.3 Methodology

The content used in this case study was extracted from a 3D model of the 14th-century St Andrews Cathedral, reconstructed by the Open Virtual Worlds research group. A detailed account of the reconstruction process that began with research and culminated in the development of the 3D model is presented by Kennedy et al. [77]. The 3D model was deployed on workstations with computer screens, Xbox controllers and Oculus Rift VR headsets to make up the desktop-based systems, while spherical images (also known as photospheres) were extracted from the model at ten sites of significance to make virtual tours on Android smartphones, coupled with the Google Cardboard and Samsung Gear VR. An iterative approach was adopted for the implementation of the mobile VR system, such that the components were incrementally implemented through several design, development and testing cycles in order to understand and push the system boundaries. These desktop- and mobile-based systems were then used to run a week-long exercise with school pupils to ascertain the value of immersive technologies for heritage exploration and to perform a comparative analysis of mobile- and desktop-based VR systems. The learning activities took place in the classroom (using both the desktop-based and mobile-based systems) as well as at the remains of the St Andrews Cathedral (using the mobile-based systems). The evaluation adopted a combination of quantitative elements including lab-based system studies, Likert-scale questionnaires with Analysis of Variance (ANOVA) and qualitative elements such as observations, interviews and word frequency analysis. The findings of the evaluation exercise provide insights into the feasibility of using affordable, mobile VR systems for delivering immersive experiences as compared to mid- to high-end VR systems.

4.4 Design and Implementation

The existence of the 3D model of St Andrews Cathedral – which was created by the Open Virtual Worlds group graphic designer, and was popular and appropriate for the case study – facilitated the design and implementation of the systems. The content of the model was made available on desktop and mobile platforms, and VR headsets and controllers were integrated for added engagement. Spherical environments – in the form of spherical images and spherical videos – were used to represent content on the mobile devices instead of 3D models so as to conserve processing power and memory resources. The mobile system takes advantage of inbuilt features of commodity smartphones. The multimedia resources (images, audio and video) are stored on the device's storage, the gyroscope is used for head-tracking to provide an encompassing view of the environment, the GPS is used for location-awareness and the screen provides touch input into and visual output from the system. Four system set-ups were implemented: Samsung Gear VR (SG, Fig. 4.4), Google Cardboard (GC, Fig. 4.4), Oculus Rift and Xbox controller (OX) as well as screen, mouse and keyboard with Xbox controller (SX).

The mobile-based systems – the Samsung Gear VR and Google Cardboard – facilitated discrete exploration of the Cathedral from distinct viewpoints. These systems featured ten locations obtained from the 3D reconstruction of St Andrews Cathedral (see Kennedy et al. [77]). Each location was represented using spherical images (see Fig. 4.2) and linked together to form a trail, similar to the practice of combining cylindrical panoramas to form a virtual tour [255]. Audio narratives were associated with each location to serve as tour guides and provide additional content for the user. In contrast, the use of the Xbox controller with the Oculus Rift and the computer screen on the desktop-based systems facilitated free-form exploration of the Cathedral. With these systems, the participants had the freedom to explore the 3D model with an avatar that can walk, run and fly around the 3D space. Participants had the option to see the Cathedral from different perspectives and vantage points, as they could toggle between a first-person view and a third-person view. Although the mobile systems restricted participants' exploration to distinct viewpoints, their mobile nature allowed for on-site exploration, such that participants could walk around the (remains of the) actual site of the Cathedral while viewing the virtual reconstruction from equivalent vantage points. This would not have been possible with the screen and/or Oculus Rift tethered to a computer because the size and power requirements confine their usage to

Figure 4.3: Oculus Rift Headset [189] (left) and Xbox Controller (right)

specific (usually indoor) locations. The distinction between the set-ups made for an interesting comparison of the benefits and limitations of each, as discussed in section 4.5. A comparison of the system set-ups in terms of cost, processing power, physical world mobility and virtual world mobility is shown in Table 4.1. The development platform, running platform and content for each system set-up are highlighted in Table 4.2, and the nature of activities performed while using each system set-up is shown in Table 4.3.

4.4.1 Desktop-Based Set-up

The desktop-based set-up featured a 3D model of the St Andrews Cathedral as it stood in 1318, which was deployed for interaction in two modes. The first mode involved a screen and traditional computer peripherals (keyboard and mouse) to view and interact with content respectively. This enabled users to view avatars in a 3D model on a screen, move the avatars around the model using the direction keys on the keyboard, switch between First and Third Person Views and zoom in and out using the mouse. The second mode involved using the Oculus Rift Development Kit 2 [189] (shown in Fig. 4.3) for viewing the 3D model through the eyes of an avatar (in First Person View only) and an Xbox controller to move the avatar around the model. The deployment of these two modes facilitated a comparison of user experiences across immersive vs non-immersive systems as well as a comparison of desktop-based and mobile-based systems, as discussed in section 4.5. The model of the St Andrews Cathedral was neither built as part of this case study, nor for the activities carried out in this research. Rather, the model was developed by the Open Virtual Worlds Group, and was adopted for this case study owing to the rich amount of detail featured, which makes it serve as valuable tool for heritage education. A detailed account of the development of the St Andrews Cathedral as it stood in 1318 is provided by Kennedy et al. [77].

Figure 4.4: Google Cardboard [58] (left) and Samsung Gear [57] (right) VR Headsets

4.4.2 Mobile-Based Set-up

The mobile-based set-up featured the use of the Google Cardboard [58, 59] shown in Fig. 4.4 and the Samsung Gear VR [57] shown in Fig. 4.4 with smartphones on which spherical images of the Cathedral model were deployed. The spherical environments for the Google Cardboard and Samsung Gear VR were developed using OpenGL and Unity3D respectively.

Literature is rife with studies that investigate the relationship between resolution and user experience [256, 257, 258, 259]. Much of this investigation is done in the context of large displays, streaming and gaming environments. However, more recently, studies that examine these effects in the context of mobile devices have been done, as mobile devices have grown in popularity as a means of consuming multimedia content. For instance, to yield quicker streaming and downloads for mobile users in bandwidth-constrained environments, Oeldorf-Hirsch, Donner and Cutrell [260] suggest that quality (and consequently file size) can be significantly decreased while resulting in relatively less reduction in satisfaction. Furthermore, there exist lower thresholds below which the user experience is unacceptable and upper thresholds beyond which there is little (if any) improvement in user experience [261]. It is therefore important for multimedia systems to deliver acceptable satisfaction levels to users with a judicious expenditure of resources, especially on constrained, mobile platforms. Considerations for screen-based systems such as content quality, resolution, display size, bit rate and frame rate largely apply to headset-based VR systems, depending on the application. There are exceptions, however, first of which is the consideration of distance, as the display size for screen-based systems depends on the viewer's distance to the screen, hence the concept of a *viewing ratio* – which is expressed as the ratio of the user's viewing distance to the screen height – is applicable [261]. For headset-based VR systems however, the viewing distance is largely held constant, as this corresponds to the distance between the user's eyes and the headset lenses. Large displays also vary more significantly in size than mobile devices; for example typical TV screens range from 15 to 60 inches, and projectors can be well in excess of this range, but typical smartphones range

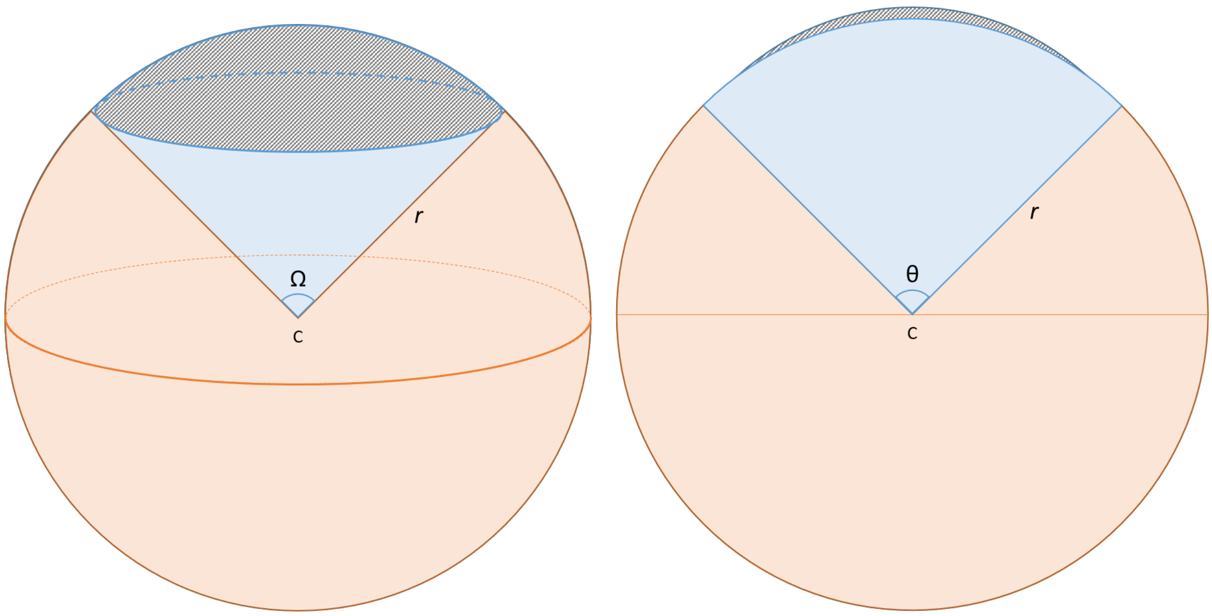

Figure 4.5: 3D sphere (left) and 2D circle (right) demonstrating that spherical images reveal a portion of content at a time, proportional to the FoV of the viewing medium

from 4.5 to 6 inches while tablets are typically 8 to 10 inches. Furthermore, flat images (and videos), when viewed on a screen at the default zoom level, typically allow for the rendering and consumption of the entire pixels in the content, in contrast to spherical images (and videos) which only render a portion of the media at any one time and only reveal other portions of the media when the user initiates interaction through head movement (for VR headsets) or clicking/dragging/panning (for screen-based systems). To illustrate this, Fig. 4.5 shows a 3D sphere with radius r and a spherical sector with a solid angle Ω in steradians represented by the blue area. The surface area of the sphere is $4\Pi r^2$ while the area of the spherical sector on the surface of the sphere, excluding the surface of the cone is Ωr^2. If the surface of the sphere represents the spherical image and the angle Ω represents the Field of View (FoV)[1] of the viewing medium, then the area of the spherical sector on the surface of the sphere (shown in the shaded part of the blue sector) represents the portion of the spherical image that is visible to the user (at the centre C). For simplicity, let this 3D sphere be mapped to a 2D circle shown in Fig. 4.5, with radius r and a circular sector with an angle Θ in degrees represented by the blue area. Since the circular construct is Two Dimensional, the circumference (and not the area) is of the essence and represents the image content. The circumference of the circle is $2\Pi r$ while the perimeter of the circular sector on the circle's circumference (excluding the perimeter of the slice made up of the two radii) is $\Theta/360 * 2\Pi r$. Similarly, if the circumference of the circle represents

[1]Field of View (FoV) in this context, refers to how much of a virtual environment is visible to the user at any time, expressed in angular degrees.

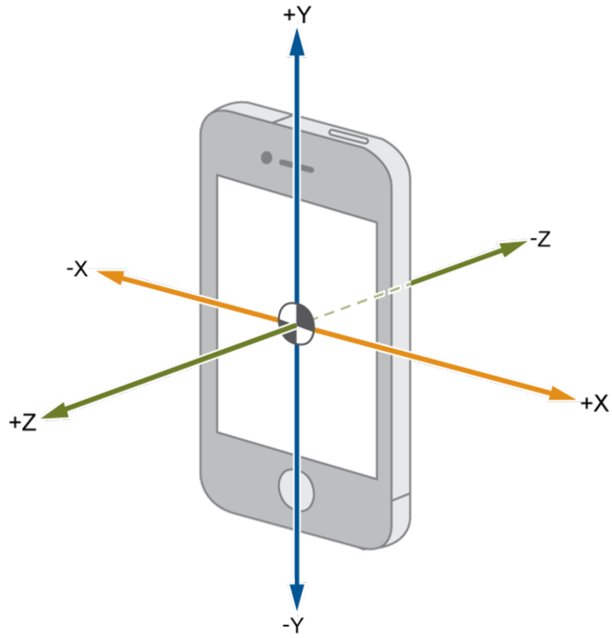

Figure 4.6: Accelerometer values measured about the X, Y and Z axes

the image content and the angle Θ represents the FoV of the viewing medium, then the perimeter of the circular sector on the circumference of the circle (shown in the shaded part of the blue sector) represents the portion of the image that is visible to the user (at the centre C). The portion of content that is visible at a time is the ratio of the shaded part to the rest of the geometric structure and is therefore dependent on the FoV of the viewing medium. This can be represented as $(\Omega r^2)/(4\Pi r^2)$, simplified as $\Omega/4\Pi$ steradians or $(\Theta/360 * 2\Pi r)/(2\Pi r)$, simplified as $\Theta/360$ degrees. For a viewing medium with a 90 degrees FoV, $90/360 = 0.25$, or 25% of the image will be visible at a time and for a viewing medium with a 60 degrees FoV, $60/360 = 0.1667$, or approximately 17% of the image will be visible at a time. A viewing medium with a 45 degrees FoV will reveal approximately 13% of a spherical image and a viewing medium with a 120 degrees FoV will approximately 33% of a spherical image. The portion of a spherical image that is visible at a time lies in this range by virtue of the FoV values of mainstream VR headsets like the Oculus Rift, HTC Vive, Samsung Gear VR and Google Cardboard. Consequently, to match the true resolution levels of flat images, spherical images need to be presented at much higher resolution levels than their flat image counterparts, and this is particularly problematic due to the resource constraints on mobile devices. Also, the consumption of these spherical images using VR devices precludes the zooming functions that are possible on screen-based systems through clicking or pinching. Therefore, techniques for rendering spherical images at higher resolutions by adjoining image tiles, and alternative interaction paradigms for zooming in and out of spherical images have been designed and discussed as follows.

4.4.2.1 Zooming

The zoom feature is facilitated by reading the accelerometer data on the smartphone and translating the values on the Z-plane (or axis) to changes in the Field of View (FoV) of the virtual camera. The logic behind this translation stems from the instinct of a person to walk forward in order to reduce the distance between said person and a point of interest. When a person cannot walk, for instance when there is an obstacle in their path, the natural instinct is to lean forward to get closer to a point of interest. Statically leaning forward can enhance users perception in VR, hence this represents a practical approach for navigation in affordable VR systems [262]. Leaning interaction paradigms have been designed and investigated in the past, but these feature dedicated (and sometimes costly) devices such as floor pads (for standing), leaning stools and gaming chairs (for sitting) [262, 263, 264, 265, 266]. This work proposes a leaning interaction paradigm that can function without additional (and dedicated) devices by using the sensors embedded in smartphones. Commodity smartphones are equipped with a magnetometer for calculating compass readings in relation to the Earth's magnetic field, a gyroscope for measuring orientation about three axes, and an accelerometer for measuring acceleration about three axes (see Fig. 4.6). The readings from these sensors can be combined to provide more accurate data about a device's position, orientation and movement in space. As a user with a mobile headset walks or leans forward, their head (and by extension the smartphone in the headset) accelerates in the direction of movement. Given that a Z-plane reading of a smartphone's accelerometer corresponds to forward acceleration when the device is in landscape (wide rather than long) orientation, an increase in the Z-value of the accelerometer indicates forward movement or leaning, and conversely a decrease indicates backward movement (see Fig. 4.6).

In a similar vein, the Field of View (FoV) of the virtual camera controls how much of the virtual environment the user sees at a time. The larger the FoV, the greater the area that is visible at a time and vice versa. Given the spherical nature of the virtual environment, a smaller FoV results in a closer view of the virtual environment (because less of the environment is seen at a time) and vice versa. For this reason, changes in the Z-value can be continuously translated to changes in the FoV of the virtual camera using the formula shown in Table 4.4 such that forward movements result in negative Z-values which translate to lower FoV values which in turn result in a closer view of the virtual environment. The Z-values have been translated to FoV values using a linear function; the pseudocode featuring the linear function is shown in Table 4.4.

4.4.2.2 Tiling

As discussed earlier, spherical images should be presented at higher resolution levels relative to flat images because only a portion of the image is visible to the user at a time. However,

1. Define *defaultFov, maxFov, fovTransitionMultiplier, currentFov, computedFov, fovTransition, accelerationZ, k1, k2;*

2. Set $k1 = 2 * (maxFov - defaultFov)$, $k2 = defaultFov$

3. On each frame update, while phone is in landscape mode (only):

   a) Get *accelerationZ* as device acceleration on Z-plane.

   b) Compute FoV:
   $$computedFov = (accelerationZ * k1) + k2$$

   c) Linear interpolate between *currentFov* and *computedFov* while applying *computedFov* to camera using the time increment defined as:
   $$fovTransition = Time.deltaTime * fovTransitionMultiplier$$

   d) Interpolate between current and new FoV values:
   $$FoV = (1 - fovTransition) * currentFov + (fovTransition * computedFov)$$

Table 4.4: Pseudocode for translating acceleration to FoV values

convert -crop 4x2@ inputimg.ext outputimg.ext

Table 4.5: ImageMagick command used to split images into tiles

mobile VR environments are often constrained by limitations such as memory, processing power and graphics. The latter is particularly challenging because 3D rendering engines impose limitations on texture size; for instance, Unity3D imposes a maximum texture size of 2,048 pixels, 4,096 pixels or 8,192 pixels depending on a combination of hardware specifications including graphics card mode, amount of graphics memory available and system architecture (32 vs 64 bit). This is lower than desirable for rendering detailed spherical environments because only a portion (around 20%, depending on the FoV) of these pixels are visible at a time. For this reason, a tiling mechanism has been developed to render as much detail as possible in the virtual environment. Although the tiling mechanism was implemented in the Unity3D game engine, the design principles are applicable to other 3D development environments. Furthermore, given the rapid advancements in computer graphics, processing power and technology, it is plausible that newer versions of 3D environments which support automatic tiling have been released in the time between the beginning of this investigation and the writing of this thesis.

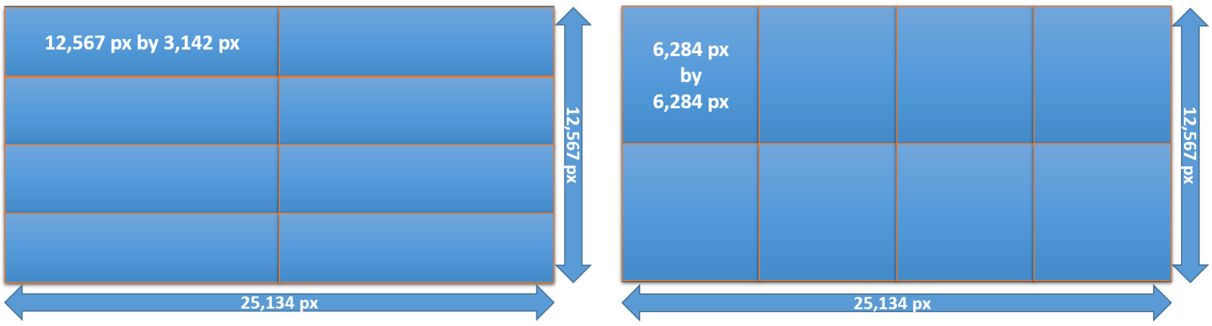

Figure 4.7: 8 rectangular tiles (suboptimal) and 8 square-shaped tiles (implemented)

Nonetheless, an understanding of how this process is carried out enables 3D developers to further investigate, optimise and push the limits of 3D engines, which facilitate the development of immersive systems in environments constrained by limited resources. Tiling works by splitting an equirectangular (spherical) image into smaller, adjoining images (called tiles) which, when put together, make up the original image. These images can then be rendered side by side in the virtual environment. It allows a series of images, each at the maximum texture size (or lower) to be held in memory and manipulated using shaders, such that images collectively yield an overall image of much larger resolution. For instance, in Unity3D, each tile will be treated as a texture (which can have a maximum size of 4,096 px), hence an image of much larger resolution (say 25,134 px by 12,567 px) can be split into smaller images such that each image will be no more than the maximum texture size. This maximises the number of pixels that can be rendered in the development environment and ensures higher fidelity rendering than that which is supported by default. This tiling mechanism consists of the following steps:

1. **Determining the number of tiles:** The number of tiles that can be rendered depends on the resolution of the original image and the development environment. Just as 3D engines impose limitations on texture size, they may also pose limitations on the number of textures that can be rendered per shader or per material, and this may vary based on hardware. For instance, a maximum of 14 textures per shader was observed for Unity3D[2]. Furthermore, a proper alignment of tiles requires dividing pixels evenly across the horizontal and vertical axes, and these divisions may introduce gaps or overlaps between tiles or may produce skewed or stretched images if the results do not produce exact remainders (for instance repeating decimal numbers such as 10/3 = 3.3333 or 10/6 = 1.6667). Hence for simplicity and optimum performance, the ideal number of tiles are powers of 2 (i.e. 2, 4, 8, 16, 32 and so on), because powers of 2 can always be split into equal parts across the horizontal and vertical axes without resulting in imprecise remainders or repeating decimal numbers. The

[2]This was verified by varying the number of textures using step-wise increments until system failure.

chosen power of 2 should be high enough to maximise the number of pixels, yet low enough to compensate for the texture size constraint.

2. **Splitting an equirectangular image into tiles:** A suitable image processing program (such as Photoshop, GIMP or Imagemagick) can be used to split an image into smaller ones. For this exercise, Imagemagick was used with the *convert* command and the *crop* argument (shown in Table 4.5) to split the image into 8 square tiles (each with 1:1 aspect ratio), provided the original equirectangular image has a 2:1 aspect ratio. For fidelity and resource maximisation, it is important to produce square-shaped tiles (tiles with a 1:1 aspect ratio) instead of rectangular tiles (tiles with a 2:1 or 1: 2 aspect ratio), because rectangular tiles may "waste" pixels along the horizontal axis. To illustrate this, consider two sets of 8 tiles – one rectangular set of tiles as shown in Fig. 4.7 and one square-shaped set of tiles as shown in shown in Fig. 4.7 – which are produced from an equirectangular image of 25,134 pixels by 12,567 pixels. If both sets of tiles are used to render a spherical environment on a platform with a maximum texture size of 4,096 pixels, then more pixels will be lost in the compression of 12,567 px to 4,096 px than in the compression of 6,284 px to 4,096 px. Each tile in the rectangular set will compressed to 4,096px by 1,024 px while each tile in the square-shaped set will be compressed to 4,096px by 4,096px (both the horizontal and vertical dimensions are compressed to maintain the aspect ratio of each tile). Thus, the rectangular set will result in a total pixel area of 4,096 px * 1,024 px * 8 tiles = 33,554,432 px, while the square-shaped set will result in a total pixel area of 4,096 px * 4,096 px * 8 tiles = 134,217,728 px. This shows that the 8 square-shaped tiles produce a spherical environment with significantly-more pixels than the spherical environment produced by the 8 rectangular tiles, and subsequent calculations will demonstrate that the spherical environment produced using 8 rectangular tiles is at par with the spherical environment produced using 4 square-shaped tiles in terms of the total pixel area.

   Where possible, to maximise pixels the ideal aspect ratio of each tile should be 1:1 (and not 2:1 or 1:2); in other words, tiles should be squares (and not rectangles). This is because when tiles with a 2:1 aspect ratio are used, a greater number of tiles are required to render the same amount of pixels as could be rendered using tiles with a 1:1 aspect ratio. This is illustrated in Fig. 4.8, which shows how 4 (rectangular) tiles do not result in more pixels than 2 (square-shaped) tiles, yet 4 tiles have more rendering overhead than 2 tiles because twice the number of images need to be transferred, stored and loaded. The rationale behind this stems from the nature of equirectangular images – they have a 2:1 aspect ratio. In order to divide a sample equirectangular image into 4 equal tiles, the tiles could either be wide (i.e. 2:1 tiles, with larger width than height such that there are two rows and two columns), wide

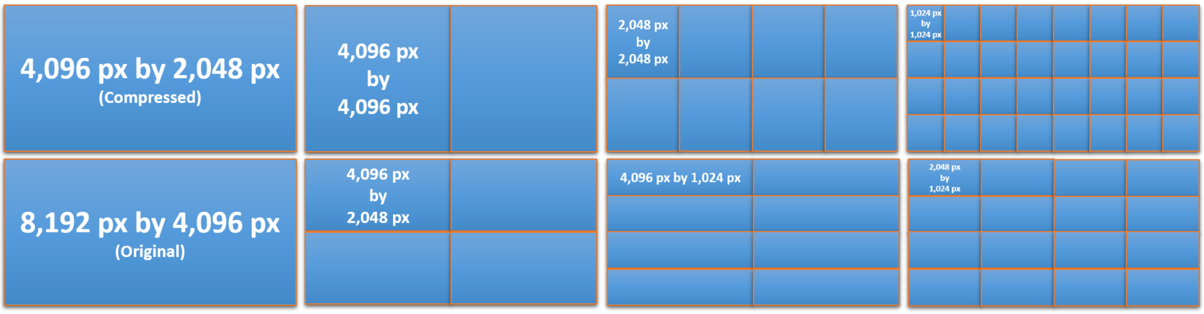

Figure 4.8: Ideal (top row) and suboptimal (bottom row) techniques for splitting an image into ($2^{0\text{-}5}$) tiles

(i.e. 4:1 tiles, with larger width than height such that there are four rows and one column) or long (i.e. 1:2 tiles, with smaller width than height such that there is one row and four columns). These are the three ways to divide a rectangle into four equal quadrilaterals. When dealing with textures, rendering environments have a maximum texture size (4,096 px for the purpose of this illustration) beyond which all textures are compressed. To maximise resources (pixels in this context) it is logical for the larger of the width and height of a texture to have the maximum texture size. If either or both of the width and height are larger than the maximum texture size, the image will be compressed so that the larger dimension represents the maximum texture size. To maintain the aspect ratio of the texture, the smaller dimension is compressed proportionately, hence pixels are lost along that dimension as well as the larger dimension. To render a 8,192 px by 4,096 px image for instance, using 4 wide tiles (each with 2:1 aspect ratio), each tile would be 4,096 px by 2,048 px, and using 4 wide tiles (each with 4:1 aspect ratio), each tile would be 8,192 px by 1,024 px. Conversely, using 4 long tiles (each with 1:2 aspect ratio), each tile would be 2,048 px by 4,096 px. Assuming the maximum texture size is 4,096 px, in the first case (using wide tiles), no compression would take place (because both dimensions are less than or equal to the maximum texture) and no pixels would be lost in the rendering process, as all 8,192 px on the horizontal plane and all 4,096 px on the vertical plane would be rendered. In the second case (using even wider tiles), each tile would be compressed from 8,192 px by 1,024 px to become 4,096 px by 512 px which would result in half the pixels being lost in the rendering process. In the third case (using long tiles), no compression would take place (because both dimensions are less than or equal to the maximum texture size) and no pixels would be lost in the rendering process, as all 8,192 px on the horizontal plane and all 4,096 px on the vertical plane would be rendered. Cases 1 and 3 are preferable to case 2 because all available pixels are rendered and no pixels are lost, while using 4 tiles. However, the same number of pixels (8,192 px on the horizontal plane and 4,096 px on the vertical plane) can be rendered with two square tiles which are both 4,096 px by 4,096 px. Since the goal – of rendering 8,192 px on the horizontal plane and 4,096 px on the vertical plane – can be achieved using 2 tiles, doing so with 4 tiles represents an unnecessary and avoidable resource expense. Hence, the next logical increment from 2 (square) tiles is therefore 8 (square) tiles (4 horizontal by 2 vertical tiles as opposed to 2 horizontal by 4 vertical tiles). In the same vein, 16 tiles do not offer any more pixels than 8 tiles but add to the rendering overhead, hence the next logical increment from 8 (square) tiles is 32 (square) tiles. Fig. 4.8 shows some options for splitting an image into tiles based on powers of two from 2^0 to 2^5. The top row depicts 1, 2, 8 and 32 tiles (from left to right) which all have a 1:1 tile aspect ratio (thus they are all appropriately-square-shaped) to maximise the pixels rendered. However, the bottom row depicts 1, 4, 8 and 16 (rectangular-shaped tiles) tiles from left to right. The image of one (8,192 px by 4,096 px) tile can only be rendered in the environment if it is compressed such that its larger dimension is no greater than the maximum texture size (4,096 px in Unity 3D) to result in the image directly above it in Fig. 4.8. Also note that the images of 4, 8 and 16 (wide-shaped) tiles in the bottom row of Fig. 4.8 do not offer any improvements over the images of 2 and 8 and 32 (square-shaped) tiles in the top row of Fig. 4.8. In fact, the image with 16 wide-shaped tiles in the bottom row does not provide any more pixels than the image with 8 square-shaped tiles in the top row, but rather adds more overhead because an additional 8 tiles (textures) need to be transferred, stored and/or rendered.

In addition, it is counter-intuitive to use image tiles whose larger dimensions exceed the maximum texture size of the development platform (4,096 px for the purpose of this illustration). Doing so adds to the overhead of compressing images at run-time when textures are created from preloaded images which are locally-stored, and may preclude the rendering of textures when created at runtime from remotely-stored images. This implies that even though an image with dimensions 25,134 px by 12,567 px (for example) can be captured, given the implementation of 8 tiles per environment and 4,096 px per texture, it may be worth compressing said image to 4,096 px by 2,048 px for a single tile, 8,192 px by 4,096 px for generating 2 tiles (i.e. each tile will be 4,096 px by 4,096 px) and 16,384 px by 8,192 px for generating 8 tiles (i.e. the image will be split 4 by 2 such that each tile will be 4,096 px by 4,096 px). This ensures that no image tile will have dimensions greater than the maximum and thus mitigates the additional compression overhead (when rendering environments from pre-loaded textures) or the risk of failing to create appropriate textures (when textures are generated at runtime).

3. **Rendering the tiles:** The created tiles can be rendered on the surface of the sphere in the development environment. 3D environments map textures onto materials using graphics shaders. If the default shaders specify a one-to-one mapping between textures and materials

(as is the case with Unity3D), custom shaders can be implemented to map multiple textures to a single material thus specifying a many-to-one mapping. The tiled shader was implemented based on an adaptation of the shader code obtained from Dimitri [267] and Zucconi [268], first using 4 tiles and then using 8, both of which resulted in a significant improvement in the fidelity of the tiled environment, and some lag in head-tracking as compared to the fidelity and head-tracking of the single-image environment. The code was further optimised to store pixel colours and alpha levels when they are first computed and reuse them so as to eliminate redundant computations and minimise calculations.

Given the ability to render spherical environments at multiple tile levels, a responsive rendering approach represents the natural step forward. This involves loading a maximum number of tiles depending on the zoom level initiated by the user. For instance, the virtual environment would initially be represented with one tile, then when a user initiates the zoom feature, the number of tiles will be increased to 2, then 4, then 8 and so on based on the zoom level requested. This is important when transferring and loading resources over a network but it can also be beneficial when resources are stored locally because fewer tiles use up less resources (in the form of processor cycles, memory and battery power) and thus optimise resource utilisation. The multi-level tiling component was implemented by modifying the shaders to support multiple tile levels and programmatically re-render the spherical environment using more tiles as the zoom level is increased. As discussed earlier, some platforms may impose limitations on the number of textures that can be rendered per material. To overcome this limitation, multiple materials can be combined to render a single spherical environment. To illustrate this, recall that a spherical environment is made up of an equirectangular image which may or may not be split into tiles, textured on a sphere mesh (or object). 3D environments typically allocate one material to simple geometric meshes (such as spheres, cubes and cylinders), however, multiple materials can also be layered on these meshes. Furthermore, parts (or the whole) of each material can be made transparent or opaque, thus it is possible to render a sphere with (say) 4 materials, each of which are transparent on 3/4 of their surface and opaque on 1/4 of the surface. These materials can then be aligned so that the 4 opaque quarters seamlessly make up a sphere, and each opaque quarter can be subdivided into 8 parts, for holding 8 tiles . This is similar to dividing an entire material into 8 parts to store 8 tiles (as discussed previously), but the difference here is that only 1/4 of the material – the opaque quarter – is subdivided to store 8 tiles. When these 4 materials are treated as a whole sphere, the resulting sphere is capable of holding 32 tiles (8 tiles per quarter). At 4,096 px per tile, the 32-tile approach allows for rendering spherical environments of up to 32,768 px by 16,384 px. Underlying platform permitting, this approach can be extended to render more than 32 tiles. As discussed earlier, the choice

Figure 4.9: Pupils exploring St Andrews Cathedral in the classroom (left) and on-site (right)

of the number of tiles for representing spherical environments is ideally a power of 2, with twice the number of tiles on the horizontal plane than the vertical plane (i.e. 2 tiles – 2 by 1, 8 tiles – 4 by 2, 32 tiles – 8 by 4), hence the ideal next step from 32 tiles is 128 tiles – 16 by 8, because 64 tiles will not produce more pixels than 32 tiles but will add unnecessary overhead in rendering the additional 32 tiles. To render 128 tiles, a sphere mesh can be layered with 16 transparent materials, each of which is 15/16 parts transparent and 1/16 parts opaque, while the opaque part is subdivided into 8 parts to store 8 tiles. When these 16 materials are treated as a whole sphere, the resulting sphere is capable of holding 128 tiles (8 tiles per 1/16 sector). At 4,096 px per tile, the 128-tile approach allows for rendering spherical environments of up to 65,536 px by 32,768 px. As imaging and graphics technologies continue to develop, there will be increasing availability of cameras capable of taking photographs at higher resolutions and 3D (game) engines capable of building and rendering high-fidelity environments, thus arbitrarily-large spherical environments can be rendered using the aforedescribed approach.

4.5 Evaluation – A Comparative Analysis of VR Systems

A study was conducted with early secondary school pupils to: (1) ascertain the value of VR systems for heritage dissemination, and (2) investigate the efficacy of affordable, mobile VR systems for delivering interactive experiences as compared to mid- to high-end systems. This is the essence for the development of the four systems (shown in Table 4.1) in the form of two mobile-based systems and two desktop-based systems. The exercise was facilitated by the annual visit of the Open Virtual Worlds team to the local secondary school, during which the four systems were set up in the school library for a week (Monday to Friday), and daily cohorts of first year secondary school pupils (aged 11-13 years) were invited to explore St Andrews Cathedral using the systems (see Fig. 4.9). Each cohort comprised of up to 30 pupils and they were further

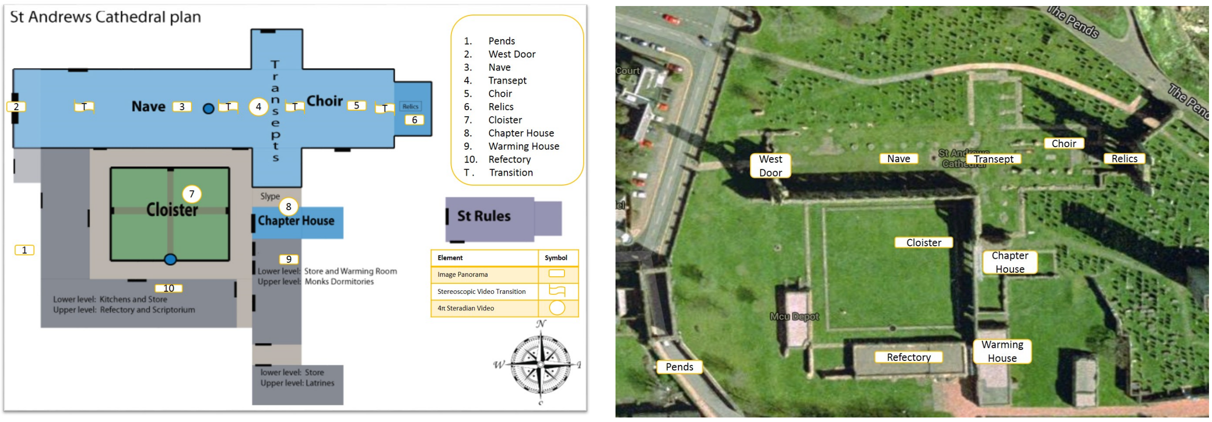

Figure 4.10: Illustration (left) and annotated (right) satellite view of St Andrews Cathedral floor plan

divided into groups of 5, each of which took turns in exploring the systems. Data collection was by means of questionnaires and observation, with both quantitative and qualitative data collected from a total of 132 participants. The school pupils also visited St Andrews Cathedral in groups where they undertook a guided tour of the remains of the site while accompanied by teachers and actors dressed as pilgrims for illustrative activities. The guided tour entailed visiting ten sites around the cathedral (see Fig. 4.10) and at each site, the pupils explored the reconstructed view of the Cathedral at that point using the mobile VR headsets (see Fig. 4.9).

Background information was recorded for each participant using the questionnaire in Table 4.6. Each participant then had up to 10 minutes to explore the cathedral using each system, after which they filled the “Experience” portion of the questionnaire shown in Table 4.7 on a 5-point Likert scale to describe their experiences and facilitate qualitative analysis. The questionnaire shown in Table 4.7 is the result of a piloting exercise on the System Usability Scale (SUS) questionnaire [236], conducted with 2 secondary school teachers and 15 university students. While it would have been ideal to pilot with the target audience themselves (i.e. school pupils), this was not possible due to delays in getting ethical approval for the study. For this reason, the questionnaires were piloted with 2 teachers from the school (these teachers were responsible for the pupils in the class chosen for the study and were going to be involved in the learning exercise, hence they could provide insights on the appropriateness of the questionnaires for the pupils). In addition to piloting with the teachers, piloting was conducted with a sample of 15 university students. The piloting yielded feedback that led to rewording the questionnaire items (to make them fit for first year secondary school pupils as advised by the teachers), as well as the structure and length of the survey to minimise dropouts due to boredom and low attention spans, thus improving response rates. By virtue of the piloting, the resulting questionnaire was streamlined to focus on specific aspects of the experience, worded to be easily understood by school pupils, and overall shortened to be easily completed in a few minutes. Medians of the Likert-scale questionnaire items were calculated as a (non-parametric) measure of central tendency, after which further statistical tests were conducted to examine whether there are significant differences in the mean responses reported by the populations. While non-parametric methods may be used on Likert-scale items for this purpose, parametric tests are also applicable because studies show that the use of parametric vs non-parametric methods do not affect the conclusions drawn from the results [269], and that for relatively-large sample sizes (>15) in Likert-scale data analyses, non-parametric methods provide similar results to parametric methods yet they possess lower statistical power [270, 271]. For these reasons, a parametric method (ANOVA) has been used instead of the non-parametric equivalent (Kruskal Wallis test). For each questionnaire item, a one-way Analysis of Variance (ANOVA) was conducted across the set-ups to determine whether there was a statistically-significant difference in the means. A statistically-significant difference would suggest that the participants felt differently about the set-ups for a particular questionnaire item, while the absence of a statistically-significant difference would suggest otherwise. For cases where a statistically-significant difference was observed, a t-Test (Two-Sample assuming unequal variances) was conducted with set-up pairs (i.e. SG and GC, SG and SX, GC and SX, and so on) to determine which pairs have means with statistically-significant differences.

The questionnaire also provided spaces for the participants to provide three words that describe their experience and to suggest what aspects of the experience could be improved on. In addition to these, observations of the actions and spoken words of participants were noted. These data were collected for qualitative analyses (see section 4.6.1). Specifically, the three words requested from each of the participants were used to form a tag cloud and a frequency analysis was conducted to observe frequently-occurring words in the sequence, which would provide a holistic representation of the user experience. Similarly, the participants' responses to what aspects of the experiences could be improved on, in conjunction with their actions and spoken words while using the system, were analysed to observe recurring patterns. These are discussed in section 4.6.1. The study was aimed at investigating three hypotheses to affirm that VR is suitable for learning and that affordable VR systems can match more costly alternatives in learning environments. There is evidence to suggest that mobile VR headsets are capable of providing similar levels of immersion [272], spatial presence and learning outcomes [273], as compared to the Oculus Rift headset. For this reason, the study focused on comparing the mobile headsets to the desktop screen-based systems. The hypotheses are outlined as follows:

H_1 *VR is easy to use regardless of prior experience*. This hypothesis is important because it is conceivable that some museum visitors may have no prior experience with VR. While it is beneficial to design digital exhibits that leverage digital literacies, it is equally important

Background Information (Please tick the appropriate column)						
		1	2	3	4	5
1	How good are your English skills?	Very Poor	Poor	Fair	Good	Very Good
2	How good are your IT skills?	Very Poor	Poor	Fair	Good	Very Good
3	How interested are you in History?	Very Low	Low	Some	High	Very High
4	Do you have previous experience with Virtual Reality (VR)?	None	Novice	Moderate	A lot	Expert
5	Do you have previous experience with Virtual Reality (VR) Headsets?	None	Novice	Moderate	A lot	Expert

Table 4.6: Participant Background Information (ranked on a 5-point Likert scale)

Experience Questionnaire (Please tick the appropriate column)						
		Strongly disagree	Probably disagree	Neither agree nor disagree	Probably agree	Strongly agree
	Statement	1	2	3	4	5
1	I think that this system is easy to use.					
2	I would recommend this virtual reality system for learning history.					
3	This system has changed how I think about the St Andrews Cathedral.					
4	I am now more interested in learning about local history.					
5	I felt like I was there in the virtual environment.					

Table 4.7: User Experience Items (ranked on a 5-point Likert scale)

to cater to the needs of users with little or no experience. Therefore, this investigation is necessary to ascertain whether special considerations need to be made for users with minimum experience in the design of VR systems for heritage learning. This will be investigated by comparing participants' previous experience with VR (Background Item 4) with participants' perceived ease of use of the system (Item 1). The absence of a significant correlation coefficient between Background Item 4 and Item 1 may suggest that there is little relationship between participants' prior experience with VR and participants' perceived ease of use. In addition, if the measure of central tendency scores for participants' responses to Item 1 (ease of use) are above the point of neutrality, this would suggest that VR is easy to use regardless of prior experience.

H_2 *Mobile headset-based systems stimulate more interest in learning history than desktop screen-based systems*. This is worth investigating due to the novelty of headset-based VR systems, especially on mobile platforms which have lowered the barriers to entry to VR device ownership. This will be investigated by analysing the means of participants' responses to Item 4 (stimulated interest in local history) across the mobile headset-based systems i.e. the Samsung Gear VR (SG) and Google Cardboard (GC) and the desktop screen-based systems i.e. the Screen and Xbox Controller (SX). A one-way Analysis of Variance (ANOVA) will be conducted to determine whether there are statistically-significant differences between the means of the set-ups and if there are, a t-Test (Two-Sample assuming unequal variances) on pairs of the means of SG and SX, and GC and SX will be conducted. A statistically-significant difference between the means would suggest a difference in the levels of stimulation provided by the set-ups and in such a case, a comparison of the means would reveal which systems stimulate more interest in the participants.

H_3 *Mobile headset-based systems and desktop screen-based systems induce similar immersion levels*. This is worth investigating in order to determine whether the experience delivered by affordable mobile systems is comparable to that delivered by mid- to high-end desktop-based systems. The essence of this hypothesis is that the ability of mobile-based VR systems to enable users to explore a virtual environment while moving around in the physical world, coupled with the novelty of these devices, balances out the computing power and graphics fidelity that the desktop-based systems provide. This will be investigated by analysing the means of participants' responses to Item 5 (immersion in the virtual environment) across the Samsung Gear VR, Google Cardboard (mobile headset-based systems) and the Screen and Xbox Controller. A one-way ANOVA will be conducted to determine whether there are statistically-significant differences between the means of the set-ups and if there are, a t-Test (Two-Sample assuming unequal variances) on pairs of the means of SG and SX, and GC and SX will be conducted. The absence of a statistically-significant difference between the means would suggest an indifference in the level of immersion provided by the mobile and desktop systems.

4.6 Results

The data collected was analysed using a combination of qualitative and quantitative techniques. Qualitative techniques took the form of a word analysis on the participants' feedback and observations of their actions while statistical techniques such as a one-way Analysis of Variance (ANOVA) and t-tests were used to analyse the quantitative data.

4.6.1 Qualitative Data Analysis

A tag cloud was created to visualise frequently-occurring words used by the participants to describe their experiences. A frequency analysis of the tag cloud (shown in Fig. 4.11) revealed that the top twelve (12) words were: *fun*, *interesting*, *cool*, *amazing*, *good*, *realistic*, *awesome*,

Figure 4.11: Participants' experiences visualised by a tag cloud

exciting, *educational*, *real*, *weird*, and *different*. Each word was placed in one of three categories depending on whether the experience described by the word was positive, negative or neutral. Negative words (such as weird, scary, uncontrollable) and neutral words (such as different, tricky) combined accounted for ≈6% of the total words while ≈94% of the words were positive. Table 4.8 shows the result of a frequency analysis carried out on the words the participants used to describe their experience. In addition to the tag cloud and word frequency analysis conducted, a thematic analysis was conducted on the qualitative data to identify themes in the feedback and observations during the study. The thematic analysis was conducted in line with the process outlined by Nowell et. al. [274], beginning with familiarising oneself with the data and culminating with production of the report. The process resulted in six themes – device form factor, visual and audio, presence of others, interaction, nature of tasks and users' preferences – which form the basis for the discussion in the following section.

Form Factor The shape, size and physical features of VR headsets can contribute to and/or inhibit user experience. For instance, the Samsung Gear VR is designed to shut off the display and audio when there is no proximity sensor reading i.e. when the headset is not being worn. As the study participants were pre-teen children, the size of the Gear VR meant that it was too big for the heads of some of them (even after fastening the strap as tight as possible) and this caused the headset to lose contact with their faces

Position	Word	Frequency	Percentage	Category
1	fun	44	10.45	Positive
2	interesting	38	9.03	Positive
3	cool	33	7.84	Positive
4	amazing	23	5.46	Positive
5	good	17	4.04	Positive
6	realistic	13	3.09	Positive
7	awesome	10	2.38	Positive
7	exciting	10	2.38	Positive
9	educational	9	2.14	Positive
10	real	8	1.90	Positive
10	weird	8	1.90	Negative
10	different	8	1.90	Neutral

Table 4.8: Frequency analysis of participants' descriptive words

(and by extension the proximity sensor), often resulting in a break in transmission. Some participants were also observed holding the headset up to their face with their hands, which precluded them from performing other tasks (such as gesturing) with their hands.

Visual and Audio Cues Cues (such as eye level) in the virtual environment play a part in the user experience. It was observed that the eye levels in the virtual scenes did not align with the eye levels of the participants. This is because the virtual scenes were taken at the eye level of an average male, which is higher than the eye level of the pre-teen participants. Some of the participants expressed perceived dissonance as a result of this, while others did not report any such symptoms. A participant also reported the strange feeling of being "there [in the virtual world] ...but not there", as they felt like they were in the virtual environment but knew that they were seated in a chair in the physical world.

In addition to visual cues, audio cues can play a big role in the user experience. During the study a participant was observed trying to put the headset close to their ear, and by implication taking the headset away from their face, perhaps to hear the audio narrative. This could suggest that the participant wished for louder audio and needed the narrative to bolster the experience. In attempting to improve the audio experience, the participants inadvertently broke the visual experience. This may not have been the case if the audio was loud enough for them, but as it was not, they were apparently willing to sacrifice the visual experience for the audio. This suggests that audio is an important part of the virtual experience, and this is consistent with findings in literature, as evidence suggests that sound and auditory cues can improve realism, immersion and memory recall in virtual environments [275, 276, 277]. Dinh et. al. [277] also found that an increase in

"Like a big, massive cinema screen"
"Nice shadows" – referring to the realistic nature of the digital model viewed in the display
"I hear people but I don't see them"
"I like this world better than the world we are in"
"It's weird and really sickening"
"Don't look at this if you're scared of heights" – referring to scenes with raised eye levels
"It's a shame you can't walk about yourself"
"Ooh...scary" – said with excitement
"This virtual life is really confusing"

Table 4.9: Selected participants' comments

visual fidelity may not result in a commensurate increase in presence, thus suggesting that the virtual experience is not all about visual cues. This lends further credence to the combination of media in digital heritage applications as introduced in section 1.1.3.

Presence of Others It was observed that the participants expressed themselves more when they used the device with other participants present; and they were quieter when they used the device alone. Participants were also observed asking each other where they were on the tour, as if to suggest that they wanted to be at the same places at the same time. This could affirm the case for social exploration of heritage sites [278].

Interaction Paradigm The interaction paradigm plays a role in the user experience. A participant was observed trying to lean forward and backward to zoom in and out respectively. This suggests that the immersive experience provided by the system caused the participant to assume that by leaning forward in the physical world, their view in the virtual environment would become closer. For this reason, a "lean-to-zoom" feature (as described in section 4.4.2.1 where a user leans forward or backward to zoom in or out respectively) could be valuable (and intuitive) in the use of VR headset systems. The efficacy and intuitiveness of this lean-to-zoom paradigm will be investigated in future work. Also, it was observed that swivel chairs facilitate better exploration of spherical environments while seated, because they provide the ability to completely rotate around a fixed axis as compared to non-swivel chairs. This is not an issue while exploring spherical environments while standing; however, standing poses an increased risk of injury (by falling or bumping into walls or objects for example).

Task and Purpose VR can be a valuable tool for learning (as shown by H_1 in section 4.6.2), but this may not always be the case; it was observed that for some participants, the fun factor of the system detracted from the learning objective. This phenomenon is well-documented in literature [170, 279, 280, 281], and suggests that there are conditions – configurations,

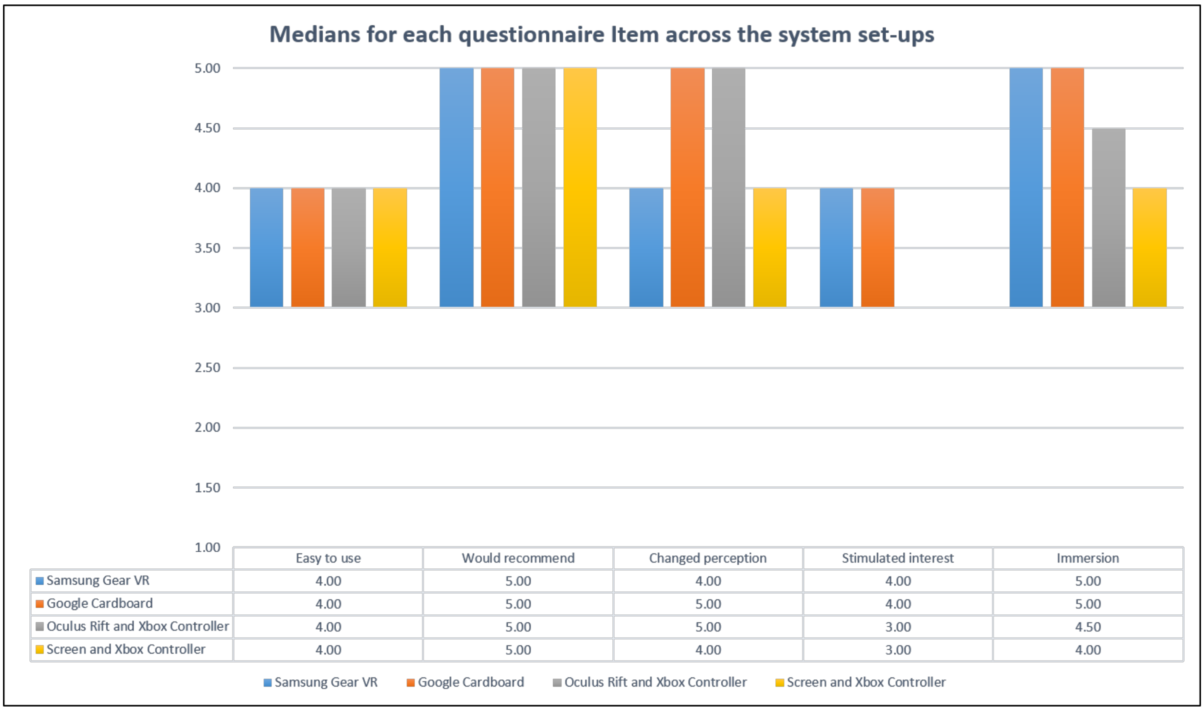

Figure 4.12: Median values for each questionnaire Item across the system set-ups

content, scenarios and so on – under which a VR system could be an effective learning tool, and others under which it would not be ideal for learning. Ultimately, the decision to deploy (or not deploy) VR should be made based on the objectives of the project.

User Preferences User inclinations and preferences can determine the nature of the experience. For instance, it was observed that a few participants refrained from properly fastening the headset strap over their heads. The experiences of said participants were limited, as their hands were occupied with holding up the device on their faces, and this occasionally resulted in shaky movements that led to a loss of contact with the proximity sensor which resulted in a break in transmission. Users may be reluctant to put on VR headsets [282] due to inclinations such as vanity, timidity or fear, and these inclinations may get in the way of effective device usage and consequently inhibit the experience.

Some statements made by the participants while interacting with the system are highlighted in Table 4.9.

4.6.2 Quantitative Data Analysis

The medians and a one-way ANOVA test were computed on the data. A snapshot of the data collected is shown in Fig. 4.12. The feedback to each questionnaire item was ranked on a

5-point Likert scale, where 1 (strongly disagree) and 2 (probably disagree) represent negative feedback, 3 (neither agree nor disagree) represents neutral feedback, and 4 (probably agree) and 5 (strongly agree) represent positive feedback. As shown in Fig. 4.12, the medians of the set-ups for Item 1 represent positive feedback (i.e. above neutral or greater than 3). This implies that the participants found the VR systems easy to use, and hence affirms the prospects of VR for exploratory learning by diminishing usability concerns. A one-way ANOVA of participants' responses to Item 1 for the set-ups revealed that there is no statistically-significant difference in the means (we do not reject the null hypothesis that the means are equal because F_{value} [2.22] < $F_{critical}$ [3.07], P-value [0.11] > alpha level [0.05]). This suggests that participants felt that the four *VR systems are easy to use*, but no one set-up is significantly easier to use than others. In addition, correlation of the participants' previous experience with VR (Background Item 4) with participants' perceived ease of use of the system (responses to Item 1 for all the set-ups) revealed a coefficient of 0.01, which indicates that there is little relationship between prior experience with VR and perceived ease of use of the system, hence the positive median scores recorded by the participants for ease of use suggests that the *VR system is easy to use regardless of prior experience* and thus confirms H_1.

As shown in Fig. 4.12, the medians of the set-ups for Item 2 represent positive feedback. This implies that the participants are likely to recommend VR systems to their peers and colleagues, and hence affirms the prospects of VR adoption for exploratory learning in schools. A one-way ANOVA of participants' responses to Item 2 revealed that there is no statistically-significant difference in the means (we do not reject the null hypothesis that the means are equal because F_{value} [0.08] < $F_{critical}$ [3.07], P-value [0.92] > alpha level [0.05]). This suggests that the *participants would recommend VR as a tool for learning history*, in line with previous work [74, 80, 81], and for learning in general [75, 76]. However, the lack of a statistically-significant difference in the means suggests that the participants would not necessarily recommend one type of the four VR set-ups over another for learning history.

As shown in Fig. 4.12, the medians of the set-ups for Item 3 represent positive feedback. This implies that the VR systems altered the participants' perception of the subject matter, and hence affirms the prospects of VR for broadening users' perspectives in learning environments. A one-way ANOVA of participants' responses to Item 3 revealed that there is no statistically-significant difference in the means (we do not reject the null hypothesis that the means are equal because F_{value} [2.67] < $F_{critical}$ [3.07], P-value [0.07] > alpha level [0.05]). This suggests that the *participants' views about St Andrews Cathedral were changed by the VR systems*, but no one set-up influenced the participants' views of St Andrews Cathedral more than the others. This demonstrates the impact that VR systems can have on users when deployed in a learning context.

For Item 4, it was observed that the medians for the set-ups represent positive feedback for the Samsung Gear VR and the Google Cardboard, and neither positive nor negative feedback for the other two systems. A one-way ANOVA of participants' responses revealed that there is a statistically-significant difference in the means (we reject the null hypothesis that the means are equal because F_{value} [4.71] > $F_{critical}$ [3.07], P-value [0.01] < alpha level [0.05]), which suggests that some set-ups stimulated the participants' interests more than the others. Further tests to determine which pairs of the means have a statistically-significant difference were conducted. A t-Test (Two-Sample assuming unequal variances) on pairs of the means revealed that there is a statistically-significant difference between the Samsung Gear VR and the Screen with Xbox Controller (we reject the null hypothesis that the means are equal because T_{value} [2.24] > $T_{critical}$ [2.03], P-value [0.03] < alpha level [0.05]), and also between the Google Cardboard and the Screen with Xbox Controller (we reject the null hypothesis that the means are equal because T_{value} [2.49] > $T_{critical}$ [2.04], P-value [0.02] < alpha level [0.05]), but there is no statistically-significant difference between the Samsung Gear VR and the Google Cardboard (we do not reject the null hypothesis that the means are equal because T_{value} [-0.22] < $T_{critical}$ [1.99], P-value [0.82] > alpha level [0.05]). This, coupled with the higher median responses for the Samsung Gear VR and the Google Cardboard as compared to the Screen with Xbox Controller, suggests that the Samsung Gear VR and the Google Cardboard individually stimulated the participants' interests more than the Screen with the Xbox Controller. This confirms H_2, that *mobile headset-based systems stimulate more interest in learning history than desktop screen-based systems.*

As shown in Fig. 4.12, the medians of the set-ups for Item 5 represent positive feedback which suggests that the participants felt immersed in the virtual environment. Furthermore, a one-way ANOVA of participants' responses to Item 5 did not find a statistically-significant difference in the means (we do not reject the null hypothesis that the means are equal because F_{value} [2.41] < $F_{critical}$ [3.07], P-value [0.09] > alpha level [0.05]). This suggests that participants reported similar levels of immersion across the system set-ups. This confirms H_3, that *mobile headset-based systems and desktop screen-based systems induce similar immersion levels*.

4.7 Discussion

Mobile VR systems have been pitted against more resource-intensive alternatives and have been found to provide comparable experience to users when interacting with spherical media. This work has found that VR is easy to use regardless of prior experience, thus there is little of a learning curve (if any) to discourage technology-averse users from engaging with content when used for heritage dissemination. Furthermore, these affordable headset-based systems can match mid- to high-end systems in terms of their ability to stimulate interest and induce immersion in the virtual environment. These findings provide the impetus for the inclusion of mobile virtual tours in the Virtual Museum Infrastructure (VMI) for facilitating remote and on-site (outdoors) exploration of heritage sites. The affordability of mobile VR headsets ensures that it can be adopted in areas where resources are limited, and the proliferation of smartphones amongst community members and the general public encourages this means of heritage dissemination.

The main challenge encountered in this work was the shortage of materials for the evaluation exercise. As discussed in section 4.5, approximately 30 participants took part in the evaluation exercise daily, and these were divided into groups of 5, each of which interacted with mobile and desktop VR systems. Only one of each system set-up was available for the study i.e. there were two smartphones, one for the Google Cardboard headset and the other for the Samsung Gear VR headset, and two desktops, one for the Oculus rift with the Xbox controller and the other for the screen and traditional input devices. The shortage of devices introduced some bottleneck as participants had to wait their turn to interact with the devices, and this may have introduced some noise in the data and increased the time required to complete the user studies. There was a silver lining, however, because while participants waited their turns to interact with the devices, they engaged in conversations and behaviours that revealed insights into issues that were worth investigating and provided for additional points of discussion during the post-exercise interviews.

The open questions raised in the course of this case study will be addressed in future work. User studies will be conducted to investigate how to provide the optimum user experience using high fidelity spherical environments, which is important because the constraints on smartphones necessitate the judicious use of resources in terms of transferring, storing and rendering textures in spherical environments.

4.8 Chapter Summary

This chapter has investigated three hypotheses using a combination of quantitative and qualitative techniques, and the results (see section 4.6) confirmed H_1, that *participants find VR systems easy to use for learning regardless of prior experience*, H_2, that *mobile headset-based VR systems stimulate participants' interest in learning history more than desktop screen-based VR systems*, and H_3, that *mobile headset-based systems and desktop screen-based systems induce similar immersion levels*.

The similar immersion levels reported for the affordable, mobile VR and the mid- to high-end desktop VR systems, coupled with the ease-of-use and tendency of the mobile systems to stimulate interest in heritage learning, affirm the case for the adoption of these affordable, mobile VR systems. This suggests that the engagement with heritage content that is achieved when exploration is confined to spherical view points on low-end systems can match that which is achieved using mid- to high-end systems running full 3D models. Furthermore, the use of mobile devices opens up new possibilities for heritage learning, as they facilitate on-site (in addition to off-site) exploration of heritage sites as compared to desktop-based systems that restrict users to off-site (indoor) usage due to their size and power requirements. These findings provide justification for the adoption of affordable, commodity technologies for participatory heritage engagement where resources may be limited, and are thus incorporated into the Virtual Museum Infrastructure (VMI) design in chapter 8. While this chapter has discussed the use of low-end and mid- to high-end systems for heritage engagement, this dichotomy is further explored in chapter 5. with a focus on using spherical media in museums context.

Chapter Five

Immersive On-site Exploration

This case study discusses the methodology for the design, development and deployment of a virtual 19th-century fish curing yard as an immersive museum installation. The museum building now occupies the same space where the curing yard was over 100 years prior, hence the deployment of a virtual reconstruction of the curing yard enables the museum visitors to explore the virtual world from equivalent vantage points in the physical world and thus facilitates a comparison of the past and present. The project methodology achieves the goal of maximising user experience for visitors while minimising cost for the museum, and focus group evaluations of the system revealed the success of the controller-free interaction design with snackable content. A major implication of these findings is that museums can provide compelling and informative experiences that enable visitors to travel back in time with controller-free interaction and relatively low cost systems.

Relevant publication:

1. **Adeola Fabola**, Sarah Kennedy, Alan Miller, Iain Oliver, John McCaffery, Catherine Cassidy, Jo Clements, and Anna Vermehren. *A Virtual Museum Installation for Time Travel*. In International Conference on Immersive Learning, pages 255–270. Springer, 2017. [3].

   Contributions: Design, implementation and evaluation of an immersive, virtual museum exhibit that features a controller-free design with snackable content, aimed at maximising value for heritage organisations.

5.1 Introduction

As discussed in section 2.3, virtual museums can refer to the recreation of a physical museum which enables visitors to remotely simulate a virtual visit in the halls and rooms of the physical museum or the creation of content and/or experiences which are then presented (remotely or on-site) to users for consumption [30]. The second definition is of pertinence to this chapter, as it introduces the design, deployment and evaluation of an immersive museum installation aimed

Figure 5.1: Museum visitors interacting with the virtual curing yard installation

at serving as an educational resource on local history. Virtual museums have gained popularity in recent decades [89] and projects that leverage computer graphics to re-enact historical scenes date back to the early 1980s [129]. Like recent interest in using 3D computer imagery and graphics to deliver engaging experiences to users, previous generations leveraged the techniques at their disposal in order to enhance immersive experiences [283]. Panoramas (and similar film-based techniques such as Cineorama, Diorama, Georama and Stereorama) for cultural heritage applications date back to the 18$^{\text{th}}$ century, and these technologies were leveraged due to their ability to “involve” audiences and place them at the centre of the scene. However, due to content and cost limitations, the popularity of these technologies declined as people sought alternatives [125]. Panoramas (in the form of cubic, cylindrical and spherical images) have seen a resurgence due to advances in photography, computer graphics and digital technologies, and the desire to visualise the past and explore geographically-distant landscapes [125], provides new opportunities for the development of systems that enable us to bridge time (historical) and space (geographic) barriers.

Timespan museum and arts centre is a cultural organisation located in Helmsdale, a small town on the north-east coast of Scotland. The building which Timespan now occupies was a fish curing yard around the late 19$^{\text{th}}$ and early 20$^{\text{th}}$ century. Fishing was a major source of income and livelihood for the inhabitants of 19$^{\text{th}}$-century Helmsdale, hence the curing yard played an important role in the village’s economy as it was a space where the herring and salmon caught by fishermen were gutted, cured and packaged for transportation. The ability to re-live and explore the curing yard as it stood in the 19$^{\text{th}}$ century thus constitutes a great resource for preserving and disseminating the heritage for the locals, and provides a medium for tourists to learn and engage with local history. This provided the impetus for the recreation of the curing yard and a conception of a project to produce an immersive museum installation that would offer novel, interactive and informative experiences to the museum visitors, with a focus on bus (coach) tours carrying over 60s, who may be ambivalent to the adoption of novel technologies [284, 285, 286].

To further evaluate the efficacy of spherical media for content dissemination in museum environments (as discussed in chapter 4), the content of the curing yard exhibit is represented as high-fidelity spherical images instead of more resource-intensive 3D models. A controller-free design which provides snackable content to visitors thus enabling a group of people (a bus tour for example) to explore the past and share their experiences with one another, offers a contribution in the design of immersive virtual museum installations. In addition to deploying the curing yard content as a museum exhibit, the content is also deployed on mobile apps for use outwith of the museum premises. Both the museum exhibit and the mobile app are evaluated using focus groups, questionnaires and interviews, and the findings reaffirm the feasibility of spherical media for heritage content dissemination and demonstrate the acceptance of a controller-free interaction design for museum exhibits.

The remainder of this chapter is organized as follows. Section 5.2 discusses the motivation for this case study, highlighting how the case study contributes to the research objectives and fits with the theme of the thesis. Section 5.3 provides an overview of the methodology adopted for the project, and section 5.4 discusses the design and implementation of the proposed system, highlighting the reconstruction process and the exhibit installation. The results of an evaluation of the system is provided in section 5.5, section 5.6 reflects on the findings of the work and section 5.7 concludes the discussion.

5.2 Motivation

Having investigated the design of an affordable VR system for immersive consumption of heritage content in chapter 4, this case study builds on the findings from chapter 4 and applies them in a museum environment. chapter 4 introduces an affordable system for interactive heritage dissemination, and a comparative evaluation with mid- to high-end VR systems affirms the feasibility of this system. That said, this evaluation was conducted with school pupils in a classroom setting, hence it is pertinent to carry out further investigation in a museum context, with museum visitors (as users) and staff (as heritage experts). The motivation for this case study is therefore twofold: (1) to deploy the proposed design and investigate the feasibility and practicality in a natural, museum environment, and (2) to further investigate how value can be maximised for museum visitors (especially older adults) while minimising cost for the museum. The proposed design has been deployed as an immersive museum installation, and to investigate this value-maximising, cost-minimising approach, the system has been adapted so that museum visitors can come away with a satisfactory experience and an appreciation of the disseminated heritage content with controller-free interaction, such that the briefest of engagements with the system will suffice to make a significant impact on the user. This case study has therefore been designed to further address a research question (see section 1.4) which asks how heritage dissemination systems can be designed to maximise value in environments constrained by limited resources. This investigation is conducted in a museum setting with a target demographic that is focused on (but is not limited to) over 60s who travel through the village on bus tours. The perceived ambivalence (but not necessarily resistance) of older adults towards the adoption of novel technologies [284, 285, 286] makes them appropriate for this investigation because a satisfactory and informative experience as reported by this demographic will demonstrate the success of the museum installation in maximising value for visitors irrespective of prior experience with VR.

5.3 Methodology

During an initial research phase, heritage experts at the museum consulted archival documents from the local and neighbouring museums. Two of the primary sources used for research include Loch [287], which sets out the plan for the first curing yards in Helmsdale and FSBI [288], the annual reports of the British Fisheries Society – the industry regulator. Based on these sources, as well as maps and plans of the area, a preliminary 3D model was created by the Open Virtual Worlds group graphic designer using Unreal Engine 4 [217]. This model then served as the basis for eliciting feedback from the heritage experts and locals, which provided further input for refining the model. An iterative design-implement-feedback cycle (see Fig 1.4) led to the development of a model that was deployed and evaluated with the visitors and heritage experts at the museum.

A prototype of the system was trailed at an outdoor, community event, where members of the general public were invited to explore the reconstruction using mobile VR headsets, and provide feedback in the form of Likert-scale questionnaires and interviews. The first access point (a location of interest in the museum for which data was gathered from the equivalent vantage point in the virtual environment) was installed in the museum shortly afterwards, and visitors were invited to interact with the system during an open day event (Fig. 5.1). This access point was evaluated with focus groups to investigate the impact that this design can have on museum visitors, and an evaluation of the value of the system from the perspective of heritage experts was conducted using semi-structured interviews with the museum director and archive manager. The findings are discussed in section 5.5.

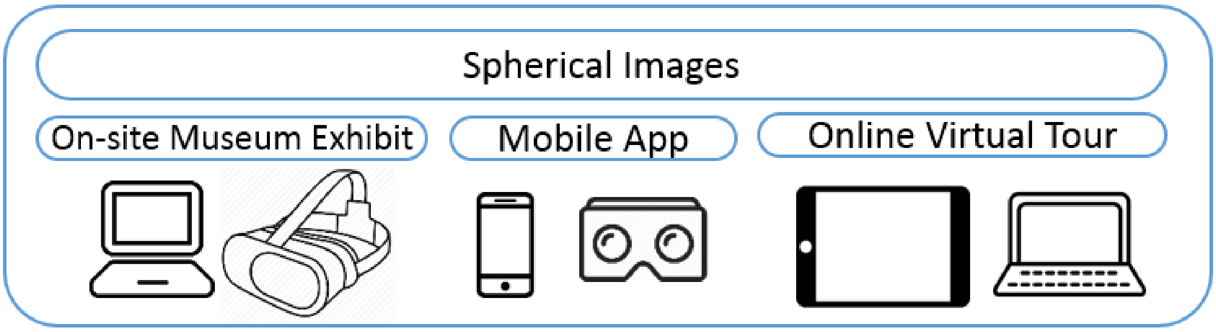

Figure 5.2: Overview of the curing yard system platforms

5.4 Design and Implementation

This work produced a multi-faceted system that depicts the virtual environment using spherical images on three platforms: an on-site museum exhibit, a downloadable mobile app for use with VR headsets and an on online virtual tour which is accessible via web browsers (see Fig. 5.2). The museum exhibit features an Oculus Rift VR headset and a screen, and is void of controllers and peripherals to encourage quick absorption of content via natural interaction through head movements. Furthermore, the exhibit is placed in the exact location that the virtual environment depicts, so that users can compare the past and present from equivalent vantage points. The mobile app can be used with the Google Cardboard headset in VR mode, and without the headset in landscape mode. The app features a string of spherical images that users can navigate through in sequence, and can be used on the go, at home and in the museum. The online virtual tour is designed for use on web browsers with traditional computing devices (screen, keyboard and mouse), by people who are unable or unwilling to use smartphones and/or VR headsets. Design decisions were made on interaction – using fixed-view access points, content – using spherical images, and platform – using VR headsets, after which the system was implemented, beginning with research, followed by 3D modelling, content extraction and system deployment.

5.4.1 Reconstruction Process

The reconstruction process began with gathering visual evidence and culminated with the deployment of stereo equirectangular images which were extracted from an Unreal Engine 3D model.

Historical Research Historical research, which was largely conducted by Timespan Museum's Archive Development Manager, entailed gathering building scale plans to represent the footprint of the curing yard, historical images to observe building transformations overtime, and a village map to serve as the basis for development.

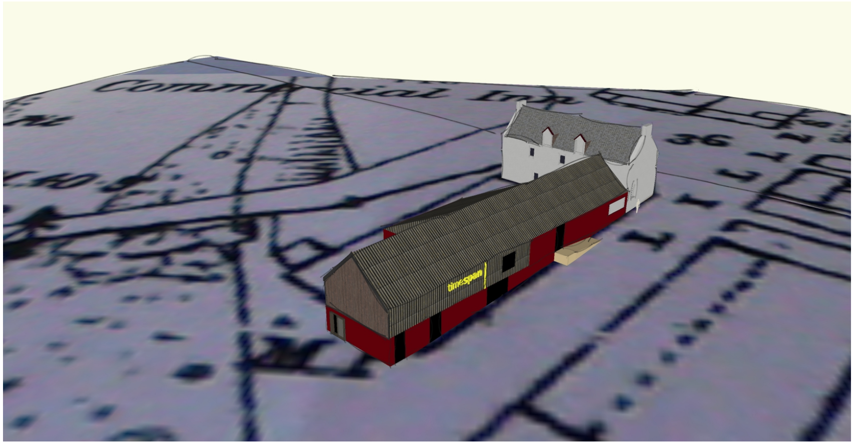

Figure 5.3: Preliminary model in Sketchup

3D Modelling The project team's graphic designer created a 3D model of the building frame using Sketchup [289] (see Fig 5.3), and exported the model as a Collada mesh into Unreal Engine. The model was populated with objects and textured in line with the typical appearance of a 19th-century curing yard (as uncovered by the research). The resulting curing yard model served as the basis for content extraction and deployment.

Content Extraction Stereo equirectangular images, formatted in a top-bottom manner were extracted from the 3D model. The extraction of spherical images from the 3D model opened up the potential for content reuse on multiple platforms – mobile virtual tours, web virtual tours, on-site museum installations, remote access, immersive and non-immersive media – which are represented in the technology and use case taxonomies discussed in section 2.5.

Deployment The images extracted from the model were used to make a virtual tour of the curing yard. The virtual tour was deployed as a mobile app which can be viewed using a mobile VR headset (such as the Google Cardboard [58, 59]) in stereoscopic mode and without a VR headset in wide mode, and as a web-based virtual tour for both online and offline use, which can be navigated using an Xbox controller or with a standard keyboard and mouse. The considerations that influenced these modes of deployment are discussed in sections 5.4.2 and 5.4.3. The spherical images were also deployed as an on-site museum exhibit, where an access point was designed around the museum foyer and a spherical image that corresponds to that space in the 19th-century curing yard was used to create a virtual environment. Historical images (see Fig 5.4) and text snippets were embedded into the virtual environment to provide context, draw attention to the equivalence of the physical and virtual vantage points, and provide heritage interpretation. This was achieved by blending the flat, historical image onto an equirectangular, virtual image such that the

Figure 5.4: Archival image of herring girls (left), embedded in a virtual environment (right)

historical image appeared to be superimposed on the spherical environment represented by the virtual image (see Fig. 5.4).

5.4.2 Design Decisions

Content The choices for content representation lay between deploying the full 3D model or deploying spherical images of vantage points. The 3D model allows users to explore the curing yard in first-person view while spherical images restrict users to exploring the system from distinct vantage points. On the other hand, spherical images enable the deployment of virtual environments with relatively-less resources on a low-end workstation or smartphone, while a 3D model requires mid- to high-end computing power and graphical resources for optimal performance. A controller-free interaction mechanism, where users engage with content in a hands-free fashion without using joysticks, game controllers and pointing devices, was adopted to enable users to explore the virtual environment from spherical viewpoints (representing past states) that are accurately mapped to the present (the physical world view). The interaction takes place via head-tracking, such users' head movements are matched to the orientation of the virtual camera in the virtual environment. Invisible hotspots, in the form of text snippets and image overlays, are triggered when the virtual camera hovers or dwells on a portion of the virtual environment, and these hotspots only remain visible for the duration of the dwell (or hover) time.

Hardware Platform Mobile (smartphone with an enclosing VR headset) and Desktop (computer with a tethered VR headset) platforms were considered for the museum exhibit hardware. Although the mobile platform had relatively lower space (real estate) requirements, the desktop platform was chosen owing to the higher computing power of the cumputer and the lower head-tracking latency of the headset. This also gives the installation a more permanent feel than could be obtained from the use of a mobile set-up. The Oculus

Level	GPU	RAM	CPU
5	Nvidia GeForce GTX 1060	64 GB RAM	Intel i7 processor
4	Nvidia GeForce GTX 980	16 GB RAM	Intel i7 processor
3	Nvidia GeForce GTX 570	8 GB RAM	Intel i7 processor
2	AMD Radeon HD 7700	4GB RAM	Intel Q6600 processor
1	Broadcom VideoCore IV	1GB RAM	Broadcom BCM2837 processor

Table 5.1: Five levels of system specification

Rift [189] and the HTC Vive [290] VR headsets were considered for deployment because they represent the newest generation of high-fidelity, consumer-grade headsets as at the time of deployment of the installation. The Oculus Rift was chosen as it has smaller physical space requirements as compared to the HTC Vive.

To minimise cost without compromising the user experience, it was necessary to determine a suitable price point for the hardware at which optimal software performance was observed. Five (5) levels of system specifications were identified with varying specifications in terms of Graphics Processing Unit (GPU), Random Access Memory (RAM) and Central Processing Unit (CPU). The five levels are outlined in Table 5.1 where level 5 represents the highest system specifications (an ASRock gaming PC tower) while level 1 represents the lowest system specifications (a Raspberry Pi 3 Model B [65]). The virtual environment – as a WebGL-powered spherical image in Chromium browser – was then run on these five system levels at three resolutions – 2K, 4K and 8K – and the number of frames rendered every second was recorded for 60 seconds of activity. The average framerates in Frames per Second (FPS) were computed for each system level and resolution, and the data is summarised in Fig. 5.5. The chart shows that the performance of the top four levels (5, 4, 3 and 2) is practically identical while the performance on level 1 (Raspberry Pi) is significantly lower than, and is insufficient as compared to the other four levels. Furthermore, the chart shows that at each system level, there is no significant difference in the performances recorded at 2K, 4K and 8K resolutions.

Having established satisfactory performance levels for the top four levels when the virtual environment is rendered on a screen, it is important to evaluate performance when rendered with a VR headset, which entailed setting up each system level with the Oculus Rift headset. The recommended and minimum specifications supplied by the headset manufacturers (obtained from Oculus [291]) are shown in Table 5.2. These specifications rule out the level 1 system (Raspberry Pi) because it lacks a USB 3.0 port (it has four USB 2.0 ports instead) and runs a UNIX-based operating system (the Oculus software only runs

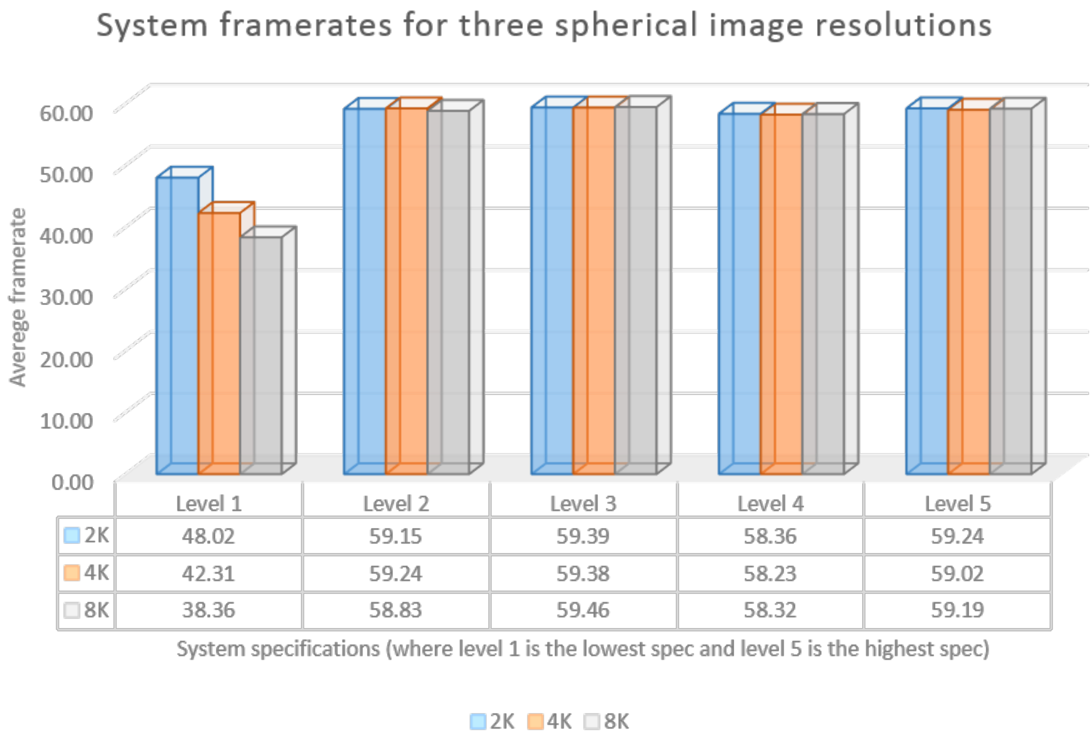

	Level 1	Level 2	Level 3	Level 4	Level 5
2K	48.02	59.15	59.39	58.36	59.24
4K	42.31	59.24	59.38	58.23	59.02
8K	38.36	58.83	59.46	58.32	59.19

Figure 5.5: Average framerates at three resolution levels across five systems

on Windows 7 or higher). Similarly, the level 2 system does not meet the minimum memory and graphics requirements according to the specifications table. However, the Oculus software and headset were successfully set up on the top three system levels, albeit with warning messages about graphics requirements produced on levels 3 and 4. The virtual environment – as a WebGL-powered spherical image in Chromium browser – was then run on the top three system levels at three resolutions – 2K, 4K and 8K – and the number of frames rendered every second was recorded for 60 seconds of activity. The average framerates in Frames per Second (FPS) were computed for each system level and resolution, and the data is summarised in Fig. 5.6. The chart shows that the performance of the top three levels is practically identical, and that there is no significant difference in the performances recorded at 2K, 4K and 8K resolutions. The implication of this is that virtual environments that are rendered using spherical images can be consumed on systems with lower specification GPUs and RAMs than that which is recommended for games and detailed 3D environments, and this yields significant cost savings. For instance by opting for a lower model, older generation GPU (such as the Nvidia GeForce GTX 570) instead of the recommended, high-end GPU (such as the Nvidia GeForce GTX 1070 or

	Recommended Specification	Minimum Specification
Graphics card	Nvidia GTX 1060/AMD Radeon RX 480 or greater	Nvidia GTX 1050 Ti/AMD Radeon RX 470 or greater
Alternative graphics card	Nvidia GTX 970/AMD Radeon R9 290 or greater	Nvidia GTX 960 4GB/AMD Radeon R9 290 or greater
CPU	Intel i5-4590/AMD Ryzen 5 1500X or greater	Intel i3-6100/AMD Ryzen 3 1200, FX4350 or greater
Memory	8 GB+ RAM	8 GB+ RAM
Video output	Compatible HDMI 1.3 video output	Compatible HDMI 1.3 video output
USB ports	3x USB 3.0 ports, plus 1x USB 2.0 port	1x USB 3.0 port, plus 2x USB 2.0 ports
OS	Windows 7 SP1 64 bit or newer	Windows 8.1 or newer

Table 5.2: Recommended and minimum specifications for the Oculus Rift (Source: [291])

1080), cost savings of several hundred pounds[1] can be made.

To complement the in-place museum installation and to leverage the design decision to use lightweight spherical content, a mobile version with similar content was also created using the approach detailed in chapter 4, such that multiple spherical images that were captured from the 3D model of the curing yard were used to make an immersive virtual tour in the form of a VR app that is downloadable onto smartphones. The spherical images were also used to make a web-based virtual tour to facilitate the consumption of content remotely with a web-browser. This furthers the objectives of affordability – because potential museum visitors can explore the virtual curing yard remotely on their smartphones with the option to do so with an affordable VR headset like the Google Cardboard, content reuse – because the developed virtual environment can be redeployed on multiple platforms such as in the museum, on mobile devices and on the web, and global dissemination – because the mobile app can be made available on app stores. The provision of mobile and web options also allow interested parties to consume the content without having to visit the museum. This is particularly useful where people are unable to travel to the museum (due to distance for example) but wish to engage with the content for research or leisure. Furthermore, the use of the mobile- and web-based systems complement the visitor's museum experience when they are used before, during and/or after the museum visit in anticipation and/or reminiscence respectively [94]. Thus, the museum installation, mobile system and web-based system collectively facilitate both on-site and online (remote) exploration of local heritage, in line with the exploration of virtual museums discussed in section 2.3.

[1]Based on the prices of Graphics Processing Units (GPUs) as of the time of this writing.

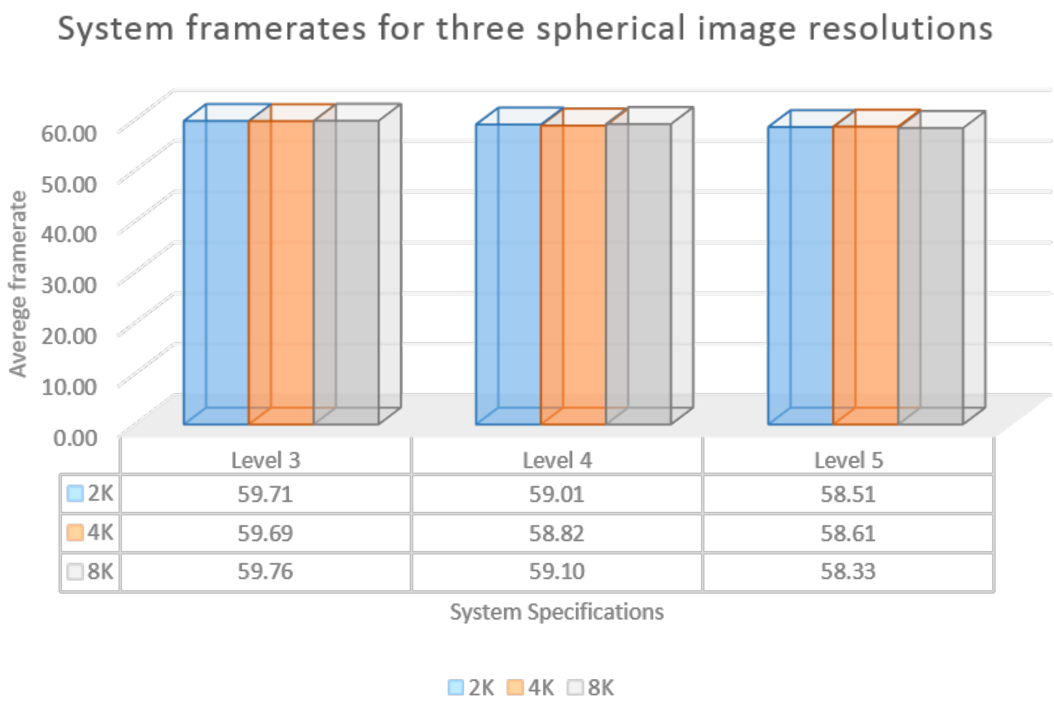

Figure 5.6: Average framerates at three resolution levels across three systems on a VR headset

5.4.3 Exhibit Installation

As discussed in section 5.4.2, a computer with a tethered VR headset was chosen for the installation. The spherical environment was designed as a virtual tour using krpano [292], which retrieves content from a local server running on the computer. To provide input feed into the Oculus Rift headset from the virtual tour, WebVR [29] – an experimental Javascript Application Programmable Interface (API) that facilitates consumption of browser-based content via VR headsets – was used. At the time of developing the installation, the two WebVR-compatible browsers for desktop were Chromium (an open-source version of Google Chrome) and Nightly (a version of Mozilla Firefox). Chromium was chosen (over Nightly) for the installation because it exhibited greater compatibility with the system features when both browsers were tested.

The installation was conducted in conjunction with members of the Open Virtual Worlds team and Timespan Museum staff, who assisted in the logistics of handling hardware, sourcing joiners and printing vinyl labels for the exhibition space. A false wall was built in the museum to accommodate the computer, with a slot for the screen to fit into, outlets for cables and a protruded base for the headset to sit on as shown in Fig. 5.7. Further to the controller-free interaction and easy-to-use requirements, the user interaction was limited to exploring the immediate surroundings of the virtual environment that correspond with each access point, and

Figure 5.7: Curing yard exhibit installation

the exhibit space was stripped of extraneous devices that did not contribute towards exploring the point of interest, thus there were no traditional computer peripherals (keyboard, mouse) or VR controllers (Xbox controller, Oculus touch). This contributed towards the usability of the system, such that the installation sent a clear message to visitors: "*put on the headset and look around*".

Alongside the ease-of-use requirement was the ease-of-management requirement, i.e. the system should run seamlessly and require little or no input from the museum staff. This includes turning on and shutting down automatically at the start and close of business respectively. Automatic startup was achieved by changing the system BIOS setting, and automatic shutdown was achieved using the Task Scheduler feature of the Windows 7 Operating System. In addition to automatically starting up the computer, the spherical environment was automatically launched shortly afterwards. This was done through the use of a batch script (see Listing 5.1) that contained the commands to be executed in sequence. A link to this script was placed in the "Startup" programs folder of the computer so that it is launched after the computer boots up. The VR mode requires fullscreen operation, however most browsers require user input (such as a mouse click or key press) in order to activate the fullscreen mode; this is an inherent security feature which prevents fraudsters from defrauding unsuspecting users by crafting fullscreen websites that look like web environments. To overcome this, a mouse click was simulated using a Java program at the system level, and an "onclick" event was attached to the spherical environment, whose action is to launch the fullscreen VR component. The Java program was packaged as a "jar" archive and executed ten (10) seconds after launcing the spherical environment (see Listing 5.1).

```
START "" "xampp"
REM launch xampp, local server

TIMEOUT /T 5
REM sleep for 5 secs

START "" "chromium_curing_yard"
REM launch chromium, to render the spherical environment

TIMEOUT /T 10
REM sleep for 10 secs

START "" "ExhibitionTools.jar"
REM simulate a click to trigger VR mode
```

Listing 5.1: API call to retrieve all items from the Omeka database

5.5 Evaluation

The system was evaluated with emphasis on *usability* from a user's perspective and *value* from a heritage expert's perspective. *Usability* is defined in terms of user engagement with the system; this is investigated using Likert-scale questionnaires, semi-structured interviews and focus groups. *Value* is defined in terms of the contributions that the system makes towards actualising the heritage organisation's goals as perceived by the heritage experts.

5.5.1 User Evaluation

A user study was conducted at the Helmsdale Highland Games that took place on the 20th of August 2016, where attendees were invited to trial the system, deployed on both smartphones and laptops with VR headsets. Using the questionnaire shown in Appendix Table B.3, feedback was gathered from thirteen (13) participants (6 female), all of whom found the system very engaging and interesting, and expressed that it gives viewers an insight into the past through a captivating visual experience. The ability to explore the virtual environment at one's pace, the ability to focus on areas of interest and the superimposition of archival images on virtual content were cited as positives. In terms of the advisory and negative comments, one participant expressed the desire to dynamically walk around the virtual environment, one mentioned a mild feeling of dizziness while using the headset, another suggested that the inclusion of a head strap would help free the user's hands for more potential interaction, and all participants expressed willingness to remotely explore the curing yard using the mobile and web platforms. Overall, the qualitative feedback suggest that the system is suited for exploring the past in an engaging manner.

Experience Questionnaire (Please tick the appropriate column)						
		Strongly disagree	Probably disagree	Neither agree nor disagree	Probably agree	Strongly agree
	Statement	1	2	3	4	5
1	I think that this system is easy to use.					
2	I would recommend this virtual reality system for learning history.					
3	This system has changed how I think about Helmsdale.					
4	I am now more interested in learning about local history.					
5	I felt like I was there in the virtual environment.					

Table 5.3: System aspects (ranked on a 5-point Likert scale)

Participants also filled in a custom 5-point Likert scale (where 1 represents "strongly disagree" and 5 represents "strongly agree"). The custom questionnaire (shown in Table 5.3) was used instead of a standardised industry questionnaire (such as the System Usability Scale) to ensure a quick turnaround time and to directly elicit quantitative feedback on five aspects of the system: ease of use, recommendation potential, ability to change perception, ability to stimulate interest and level of immersion. The results (summarised in Fig. 5.8) show that median participant responses were all positive (i.e. well above the neutral score of 3), hence participants found the system easy to use, participants would recommend the system for learning history, participants' perceptions of Helmsdale were changed by the system, participants became more interested in learning about local history and participants felt immersed in the virtual environment. Unlike chapter 4 where data was collected on usage of four systems and structured as four respondent populations, all respondents in this section provided feedback on the same system and were thus considered part of the same population, hence no tests were conducted to determine statistically-significant differences between populations.

The curing yard museum exhibit was evaluated with a focus group consisting of nine (9) persons (six female) at the 2017 Timespan Conference. The session was organised by recruiting participants through publicity campaigns and word of mouth aimed at publicising the museum's annual conference and planned exhibit launch. This presented an opportunity to invite attendees to express interest in taking part in the focus groups, which were run while the participants, together with a moderator stood around the exhibit. The session lasted for approximately one (1) hour, in which participants gathered around the curing yard installation and took turns standing in front of the screen and using the VR headset. A discussion ensued during which an observer took notes of the comments and participant behaviour. A thematic analysis was conducted on the data generated by the focus group discussions. The thematic analysis was conducted in line with the process outlined by Nowell et. al. [274], beginning with familiarising oneself with the data and culminating with production of the report. The process resulted in three main themes – system

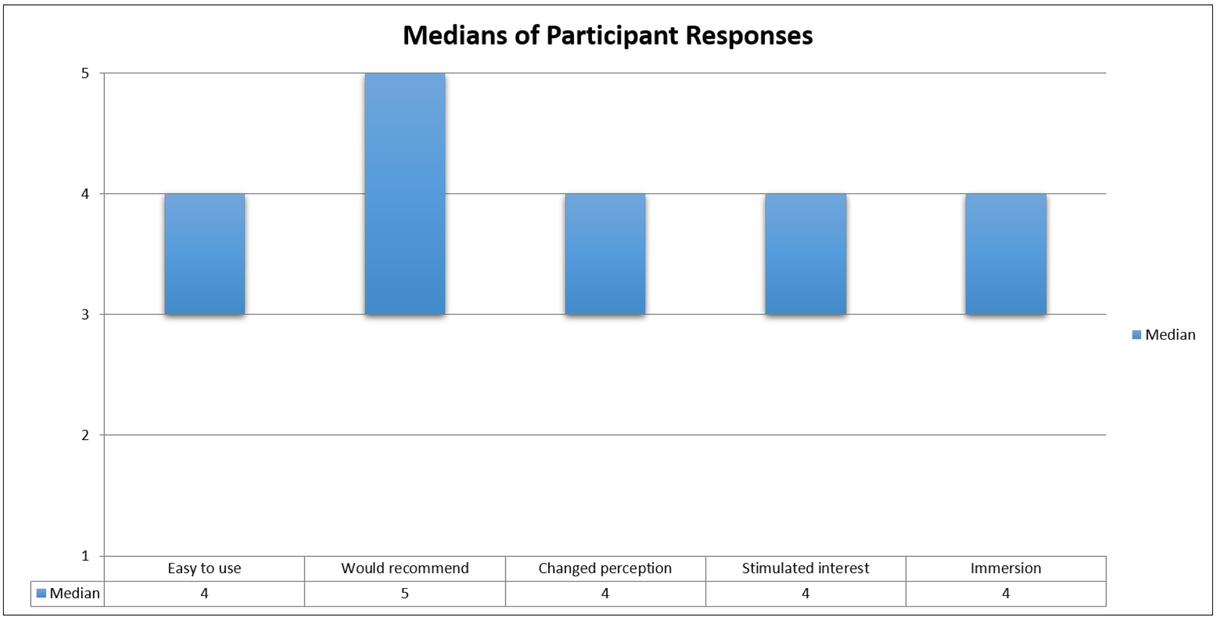

Figure 5.8: Medians of participants' responses to questionnaire items

usability, exhibit content and technology – which form the basis for the following discussion.

Usability A discussion ensued about the ease of use of the system, as one participant remarked about how "straight-forward" the installation was. A participant stated that the position of the installation (in the museum foyer) was ideal, as it made it easily noticeable to new museum visitors and it made a clear statement to visitors to simply "walk up and use" the system. Furthermore, the lack of external controllers and peripherals was cited as particularly beneficial to older-aged visitors, as users simply had to put on a headset and look around. Participants all agreed that the absence of peripherals (besides the headset and the screen) made the installation less intimidating and more inviting to users.

Content Participants expressed curiosity about the single-view nature of the installation, as one participant asked why there were not multiple locations that could be accessed in the virtual environment. A museum staff who was part of the discussion mentioned that the installation serves as the first in a series of three access points around the museum, where each access point provides a single, spherical view of what the 19th-century curing yard would have looked like at that vantage point. This was deemed acceptable by all but one participant, who stated that this decision made the resulting installation seem so "basic", suggesting that a more dynamic environment could have been deployed with the technology used for the installation. Also on the aspect of content, participants seemed to appreciate the blending of an archival image into the virtual environment, together with informative text snippets which fade in and out when an area of interest is looked at. Participants also suggested the addition of audio content to provide more information and context about the 19^{th}-century curing yard. One participant suggested including ambient sound in the form of historical songs that were often sung by herring girls while gutting fish, and another suggested the use of audio narratives instead of (or at least together with) the text snippets for the benefit of children (who may be unable to read) and adults (who may be unwilling to browse around for the snippets).

Technology A participant asked why an Augmented Reality (AR) approach – as opposed to a Virtual Reality (VR) approach – had not been adopted for the installation. The participant suggested that a system where visitors could hold up a device in a direction in the museum foyer and be presented with a synthetic environment with informative text, may have worked better than the current installation. Another participant chimed in to state that such a system would require the use of mobile devices, which would either require purchasing multiple devices or requiring museum visitors to have capable devices with the application downloaded and installed. This would be less practical than the current installation as suggested by the participant. Furthermore, yet another participant made an argument for the use of VR instead of AR by stating that the use of VR instead of AR enables users to immerse themselves in the virtual environment (which represents the past) and compare the experience with the physical environment (which represents the same space in present day). Such a comparison of experiences would be lost if a mixed (augmented) reality approach was adopted.

5.5.2 Expert Evaluation

In addition to the user evaluation, an expert evaluation – in the form of pre- and post-deployment interviews – was conducted with heritage experts at Timespan. This took the form of semi-structured interviews conducted before the system was deployed so as to gauge expectations, and interviews conducted after deployment to review outcomes and elicit feedback on level of satisfaction using the questions in Table 5.4 as a guide, while allowing flexibility for the discussion to evolve to allow for the discovery of additional information. An interview with Anna Vermehren, the Director of Timespan revealed that the museum hoped to offer novel, interactive experiences to visitors, through an installation that is informative, easy-to-use, and leaves a long-lasting impression. When asked about her views on the project output, she said:

> *"...the final product is definitely beyond the expectations that I've had in the beginning both visually and also in terms of the research...I really love the integration of archival material (images) because it demonstrates the research conducted".*

Item 1 (Purpose) - Why did you embark on this project? What did you hope to get out of it?	
Item 2 (Function) - How do you see this being used in a museum? What role(s) do you think it will play?	
Item 3 (Value) - Do you believe this system will add value to the museum visitors' experience? If yes, how? If no, why not?	
Item 4 (Challenges) - What are the main challenges facing the use of such systems in museums?	

Table 5.4: Questions used for expert interviews

An interview was also conducted with Dr Jo Clements, Timespan's Archive Development Manager and the project lead. When asked about how the installation will add value to visitors, she stated that the system offers a unique experience to visitors by bridging time and space barriers, in the sense that it enables to visitors compare their current location (i.e where they are currently standing and using the system) to what it was like in the past, thus fostering a comparison of past and present and enabling them to see how it has evolved. She further stated that the context of location (i.e. comparing the past and present from an equivalent vantage point) is important in achieving this, as context turns the virtual environment into something that is more than just interesting images, thus making the experience more meaningful. In addition to educating visitors, the installation will pique visitors' interests in local history and direct them to the museum archive where they can learn more.

An evaluation also was conducted with a focus group consisting of 3 (female) participants, aged 50-70, who are all members of the Timespan Heritage Group. This exercise was particularly valuable because the participants were all experienced heritage practitioners who represent the target demographic, hence they were able to provide valuable insights. Overall, the participants stated that they were satisfied with the experience. Participants were pleased with the level of detail in the curing yard, and they consider the system as a very informative resource and a good way to learn about heritage. The findings of the exercise are discussed in terms of user experience, technology and content.

User Experience When asked about the user experience, a participant stated that they were pleased. One participant stated that they found the experience *"very smooth"* and another commented *"you're really there, wow, all round you"*. A participant stated that *"the use of old pictures blended into the virtual environment brings the exhibit to life"*. Conversely, another participant was displeased with the embedded historical image because *"it comes in and obscures what is behind it (i.e. the virtual environment) too quickly"*. A participant suggested that a chair might be useful so that when users wish to make a full 360°turn, it can be done safely while seated as opposed to while standing. This might mitigate the risk of falling or getting entangled by the headset wires when turning. At this point, another participant suggested that wireless headsets may improve the user experience.

Technology and Content Participants expressed preference for the use of engaging, VR technology as a means of interpretation over storytelling. In addition, participants thought that the use of VR made the exhibit more exciting, and makes them excited to try new means of heritage interpretation. In terms of content, one participant suggested that the use of audio narratives would be beneficial, while another expressed preference for ambient sound (in the form of voices in the background) instead of a narrative. Overall, all three participants expressed acceptance of the proposed addition of audio to the exhibit.

5.6 Discussion

This work has explored the use of spherical media in the context of on-site, digital museum exhibits. Spherical images have been used to portray the past representation of the physical space in which the exhibit is deployed, hence enabling a comparison of the past and the present from equivalent vantage points. A technical evaluation of five system configurations demonstrates that by using spherical media as opposed to extensive 3D models, comparable performance can be obtained from system configurations that are lower than those recommended by manufacturers, and this offers significant cost savings. Furthermore, by opting for these spherical environments in lieu of model environments, the user experience can be designed for easy and quick absorption of content through natural interaction techniques, and this is particularly appealing to technology-averse museum visitors. In addition, the concept of content reuse (as an integral feature of the Virtual Museum Infrastructure (VMI)) is demonstrated in this work, through the deployment of a virtual environment (which is built in a game engine) on multiple platforms – mobile virtual tours, online web virtual tours and on-site digital exhibit.

A challenge manifested in the form of a trade-off between system performance and cost. As discussed in section 5.4.2, to create an immersive, virtual environment, a 3D model approach – which allows dynamic exploration of virtual environments – was pitted against a spherical image approach – which confines exploration to distinct view points; the latter was adopted by virtue of the findings of chapter 4 that spherical images can provide compelling experiences to users. While these spherical images could have been deployed on mobile platforms, a desktop-based approach with spherical environments was deployed to deliver the ideal trade-off between cost and performance and provide a more permanent feel for the exhibit than that which could be provided with mobile devices.

Another challenge manifested in the form of a trade-off between ease of use (from the perspective of the museum visitors) and convenience of management (from the perspective of the museum staff). For convenience, the exhibit should be relatively easy to turn on/off and debug when issues arise, and for ease of use, the system should be relatively straight-forward to use with minimal (or no) instructions. To this end, the system has been configured to automatically boot-up and shut-down on a schedule to minimise the burden on museum staff, and the interaction has been made as intuitive as possible so that visitors know how to engage with the system without being told what to do.

The next stage of the project involves deploying the model in more access points around the museum so that visitors can explore the curing yard from multiple vantage points. The new access points as well as the existing one will incorporate curing yard-related sound tracks as well as audio narratives to improve on the immersive experience, and further evaluation will be conducted to evaluate whether the site-specific nature of the installation becomes more obvious to users once there are more access points, evaluate how the target demographic (over 65s) will interact with the technology, as well as the extent to which the novelty of the technology contributes to the interaction, and evaluate how site-specific, VR exhibits which are distributed across a museum can function as an educative tool.

5.7 Chapter Summary

The design, implementation and evaluation of a 19^{th}-curing yard as an on-site virtual museum installation discussed in this chapter has explored how participatory heritage engagement can be achieved using mid- to high-end systems for cost-minimisation and profit-maximisation. The use of spherical viewpoints with snackable content has been demonstrated to offer compelling experiences with controller-free interaction, such that both young and old visitors can benefit from the use of interactive technology for heritage dissemination. The use of spherical scenes to depict immersive environments minimises the cost of computing power because the environments can be rendered on commodity hardware that cost significantly less than hardware required to render detailed 3D environments. When this is incorporated into a controller-free design with snackable content, the user experience is not compromised, hence this represents a value-maximising, cost-minimising approach to facilitate a comparison of the past and present from equivalent vantage points. This demonstrates how mid-end systems can facilitate heritage dissemination where resources are available, and thus makes a case for its incorporation into the Virtual Museum Infrastructure (VMI) design which is further explored in chapter 8.

CHAPTER SIX

IMMERSIVE REMOTE TOURISM

This chapter introduces a portable system for exploring remote landscapes in real-time. The system, which has been designed and implemented, incorporates a drone for flight and video capturing, a Raspberry Pi for wireless communication, an Android server for streaming and control, and one or more Android clients for decoding the stream and rendering the footage in mobile VR headsets. The system has been evaluated from a technical perspective to investigate and optimise resource utilisation, as well as from a user perspective with participants to investigate usability and feasibility of the idea. The findings demonstrate that the system is a viable means for group-based, real-time virtual tours of remote areas, and an implementation using affordable system components makes it feasible for participatory heritage engagement.

Relevant publication:

1. **Adeola Fabola**, Alan Miller, Ishbel Duncan. *Aerial Virtual Reality: Remote Tourism with Drones*, presented at the 7th European Immersive Education Summit, Lucca & Pisa, Italy. 2017. [4].

   Contributions: Design and implementation of an immersive system for aerial exploration of remote landscapes, and an evaluation of the system to investigate feasibility and group exploration using headset-based systems.

6.1 Tourism in the New Age

Tourism is a booming industry which contributes to the economic progress of towns and villages. It also fosters an appreciation of the heritage and history of a people. In recent decades the need to explore remote areas have led to the adoption of novel forms of remote heritage exploration which involve the use of multimedia such as audio, video and images.

Video forms of remote tourism are either provided pre-recorded (i.e. where the video is captured offline, stored and then made available on a platform usually through social media) or live (where the video is delivered to users in real time as events unfold). The emergence of digital

Figure 6.1: DJI Phantom 3 Advanced flown over the Strath of Kildonan

technologies, multimedia and social media platforms have given this mode of exploration new light, as users can now immerse themselves in far-away lands through the use of VR and spherical media. Facebook and YouTube platforms now support streaming live spherical images and spherical videos which place the user at the centre of the scene.

The role of video (and other media) in mediating tourist experiences has been investigated and the findings suggest that such media facilitate imagination (before visiting a place) and reminiscence (after visiting), and also provide access to remote landscapes [94]. Furthermore, access to remote landscapes is particularly important when it is not feasible to visit said remote areas. In addition, some places are hard to reach owing to logistic or financial reasons. For some, certain places are also inaccessible due to distance, illness, and/or disability (lack of wheel chair access for instance). The need for remote exploration of inaccessible sites necessitates the development of a novel system; hence a drone-based communication relay system which facilitates the delivery of content from inaccessible locations is designed, implemented and evaluated in this case study. The use of a drone allows users to obtain footage of areas that are inaccessible on land. The system design also enables the footage to be provided irrespective of electricity or Internet access and thus makes it suitable for deployment in remote locations. The remainder of this chapter is organized as follows. Section 6.2 discusses the motivation for this case study, highlighting how the case study contributes to the research objectives and fits with the theme of the thesis. Section 6.3 discusses the methodology adopted for this case study. section 6.4 describes the system design, section 6.5 describes the implementation of the system as well as the tools, platforms and devices used. Section 6.6 outlines an evaluation of the system components and functionality, section 6.7 provides a discussion of the user evaluation and findings, section 6.8 reflects on the findings of the work and section 6.9 concludes the discussion.

6.2 Motivation

The use of virtual museums for remote exploration motivates this case study. Recall that virtual museum systems can be used for either on-site exploration, off-site (remote) exploration or both [30]. Chapter 4 has investigated the use of heritage dissemination systems for both on-site and off-site exploration of heritage using location-aware, mobile- and desktop-based VR systems, while chapter 5 has investigated the on-site exploration of heritage using an immersive virtual museum installation. This chapter therefore investigates a use case for the remote exploration of heritage using affordable system components which are combined to deliver interactive experiences. Remote exploration is important where heritage sites may be inaccessible to the public, may be accessible but difficult to get to, or may be in danger of deterioration when accessed regularly. Furthermore, remote exploration use cases may involve areas without steady power supply or network connectivity, hence the system designed and implemented in this work is aimed at addressing these challenges, and the use of an elevated, open field with a view of some local attractions (including the museum, community centre and remains of the castle) is appropriate for this use case. Furthermore, an integral part of this case study is a focus on digital literacies (in the form of smartphone ownership), intuitive interaction with mobile VR headsets and the design of an infrastructure using affordable components to deliver engaging experiences to museum visitors where (financial and computing) resources are in short supply.

A secondary objective of this case study is a cursory exploration of social aspects associated with mobile VR headsets. The evaluation exercises in chapter 4 uncovered social aspects associated with the use of mobile VR headsets for heritage exploration. For instance, while exploring the virtual environment, some pupils were observed to gesture (by pointing in the air) and speak to their colleagues about their current view in the virtual environment because they assumed that said colleagues shared the same views as them, some pupils made funny faces at their colleagues who had the headsets on because they knew that said colleagues were "absent" from the physical world and thus could not retaliate, and after all the pupils had explored the virtual environment, they often spoke to each other about what they saw during the exploration. Collaborative and group-based interactions in indoor museum settings using mobile technologies, open displays and table-top technologies have been studied extensively and are well-documented in literature [293, 294, 295, 296, 297, 298]. However, the dynamics of how social experiences unfold in outdoor heritage settings – specifically, technology-based, real-time, guided tours of remote landscapes – remain to be seen. It is hoped that this initial exploration will provide insights that serves as the basis for further investigation. This investigation is conducted with school pupils at a community park. Similarly to chapter 4, school pupils constitute an appropriate audience for this investigation because young people tend to be inquisitive, enthusiastic and forthcoming with feedback.

6.3 Methodology

An iterative approach was adopted for this project, whereby the project span entails design-implement-evaluate cycles. The process began with an investigation into the technical capabilities of drones in terms of features such as flight range (distance), multimedia (video cameras, streaming), networking (transmission and communication) and so on. This resulted in the formulation of a conceptual model and system architecture (Fig. 6.2). After designing the architecture, the implementation of the individual system components were then investigated and executed, and after building all components, they were pieced together into a functional system. The next stage involved probing this system so as to understand and stretch its limits, and optimise it for use in live, production environments. The resulting system was then evaluated with participants in a community park to ascertain its feasibility for remote virtual tours of heritage sites and to investigate the social aspects associated with immersive, mobile VR systems.

6.4 Design

The design process began with the requirement to stream live video from a drone to one or more client smartphones configured for mobile VR headsets. In order to relay the footage from the drone to the client devices, an approach in which a server receives the First Person View (FPV) from the drone's controller and then broadcasts the FPV to client devices was designed. The feasibility of this approach is contingent on communication (preferably wireless, to ensure mobility and portability) between the client devices and the server, as well as communication between the server and the drone's controller. The former can be achieved using a network infrastructure that connects the client devices to the source of the video stream, while the latter can be achieved by implementing the server as a mobile app which runs on a device that is directly connected to the drone's controller. This approach was therefore implemented by transferring the raw video stream from the drone's controller into a smartphone (henceforth referred to as the server) connected to the controller via USB, and broadcasting the stream across a wireless network hosted on a Raspberry Pi [65] – an affordable, portable computer that can be powered with a mobile power bank – as a wireless Access Point. Furthermore, use cases which require the system to function in the absence of Internet access can be actualised using this approach,

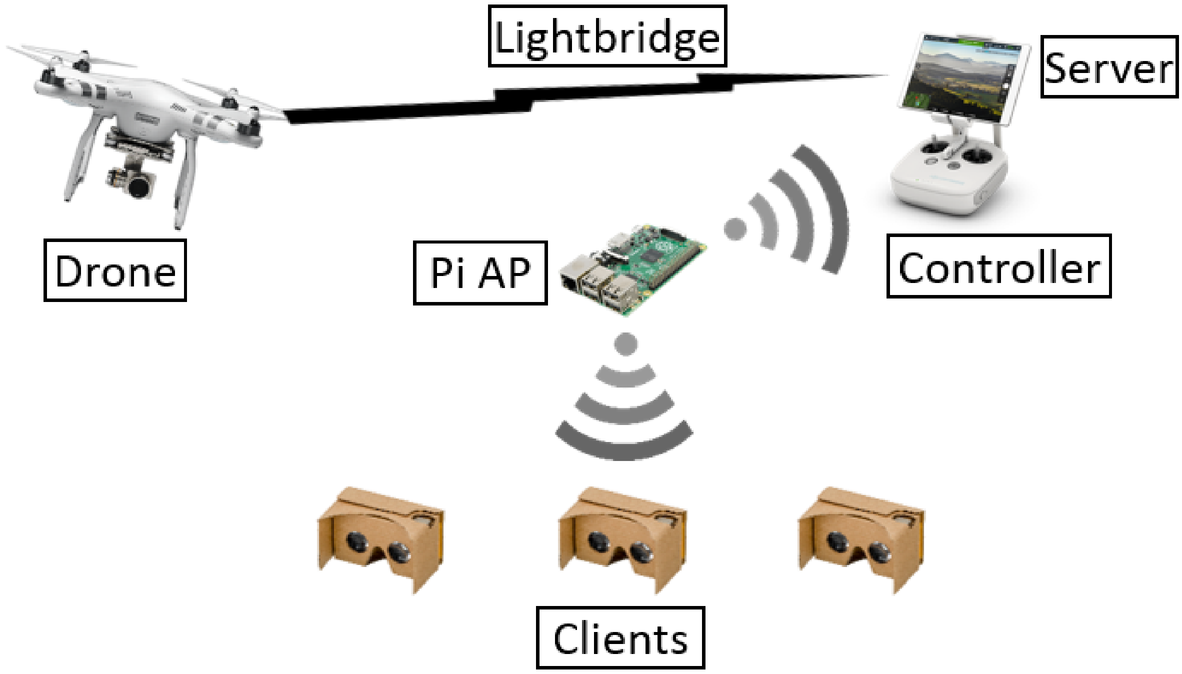

Figure 6.2: System architecture

by hosting a local wireless network on the Raspberry Pi, and broadcasting directly from the server smartphone, while client devices (connected to the Raspberry Pi's wireless network and running the client app) receive the stream. Thus, a drone-relay communication system has been designed to provide a live video stream of inaccessible or distant locations to tourists. This system features a drone (Fig. 6.1) which is equipped with a High Definition (HD) camera mounted on a gimbal. The drone is flown by a pilot – a person equipped with the skill and licence required to operate Unmanned Aerial Vehicles – with a joystick-like controller which communicates with the drone via a wireless infrastructure and has a maximum transmission range which constrains the distance that the drone can be flown away from the controller's location. The controller also provides an Application Programmable Interface (API) and an Software Development Kit (SDK) which enable the development of apps for Android and iOS which can be used to retrieve content from and send instructions to the drone. The system also features a 802.11n wireless network to facilitate communication between a server (an Android smartphone connected to the controller via USB) which transmits the raw video from the drone and subscribing clients (one or more Android phones) which receive and decode the footage and display it in VR mode. The system architecture is presented in Fig. 6.2 while a collaboration diagram that depicts interaction between the system components is provided in Fig. 6.3. The end result is a system that comprises of a drone (with a camera mounted on a gimbal), piloted with a joystick-like remote controller, a server phone which connects to the remote controller via

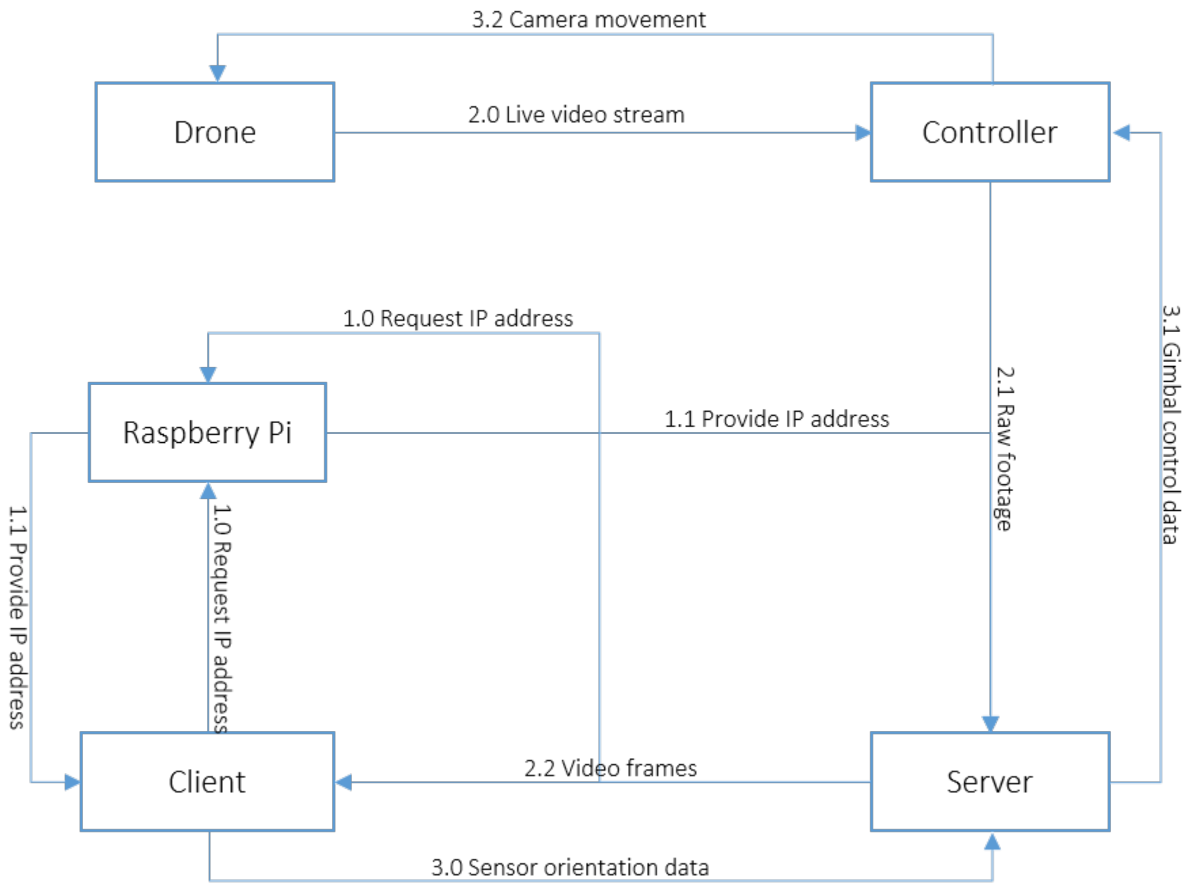

Figure 6.3: System collaboration diagram

USB, one or more client phones in VR headsets, and a Raspberry Pi Access point. The drone can be piloted over distant and inaccessible areas while the camera captures live video which is transmitted to the controller. The server app captures the video stream and sends video frames to the clients over the wireless infrastructure provided by the Raspberry Pi. The clients can also control the orientation of the drone's camera using natural head movements of the user. When head movements occur, the sensor readings of the client device are captured and sent to the server, which instructs the controller to move the drone camera.

6.5 Implementation

The system implementation was carried out in stages, with each stage focused on one compartment of the system. These compartments include the drone for flight and remote access, the wireless infrastructure which is facilitated by the Raspberry Pi, and the server and client which are implemented as Android smartphone apps. The equipment used and the tasks carried out are described in the following sections.

Drone

DJI Phantom 3 Advanced [299] and DJI Inspire 1 [300] drones were used for development and testing. The Phantom 3 Advanced has a maximum flight time of 23 minutes while the Inspire 1 has a maximum flight time of 18 minutes, and for both drones the remote controller has a maximum transmission distance of 5 kilometres or 3.1 miles unobstructed [299, 300]. The camera on the Inspire 1 can record video up to 4096x2160 pixels while the Phantom Advanced can record video up to 2704x1520 pixels. However, on both drones live streaming is capped at 1280x720 pixels at 30 frames per second. The gimbal on the Inspire 1 can pitch from -90°to +30°and pan 330°, while gimbal on the Phantom Advanced can pitch from -90°to +30°.

Wireless Infrastructure

A Raspberry Pi 3 Model B (Fig. 6.4) running hostapd (Host Access Point daemon) [301] was used to create an Access Point which the server and clients connect to. The Raspberry Pi can be powered using a wall socket charger if available, but in order to support a use case in remote areas, it can also be powered with a mobile power bank. The Raspberry Pi has a 1.2GHz 64-bit quad-cire ARMv8 CPU with a 802.11n Wireless LAN interface [65]. It also has 1GB RAM, 4 USB ports 1 HDMI port and a Micro SD card slot. A 32GB Micro SD card was used as secondary memory for the Raspberry Pi.

Server

The server was developed as an Android mobile application, compatible with any smartphone running Android 4.4 or higher. The server app runs on a smartphone connected to the drone controller via USB, interfaces with the DJI Mobile SDK, implements a callback method which receives chunks of the raw video from the drone source and sends the chunks as datagram packets to the subscribing client devices over the wireless network. The server also runs a control protocol by which it signals gimbal control (or lack thereof) to each client. This ensures that only one client has control of the gimbal at each point in time.

Client

In line with the portability requirement of the system, the client was developed as an Android mobile application, compatible with any smartphone running Android 4.4 or higher. The client app runs on a smartphone in a Google Cardboard VR headset and communicates with the server while connected to the wireless network hosted by the Raspberry Pi. It receives datagram packets containing chunks of the raw video from the server and continuously parses the raw video (in real time) using ffmpeg, an open-source, multimedia framework [302]; access to ffmpeg is facilitated

Figure 6.4: Raspberry Pi in a transparent plastic case

by a Java class which calls methods in a shared object library. Each frame that is successfully parsed is then rendered to a surface in VR mode. Each client also listens for control packets (sent by the server) to determine whether it has control of the drone gimbal. When a client receives a packet authorising it to control the gimbal, it begins to send packets containing sensor readings that describe its orientation in space. This enables the client to control the drone's view, such that the movement of the user's head (and consequently the client smartphone) causes the drone gimbal to move, which in turn causes the drone's view to change, thus providing the illusion of an immersive first-person experience.

Three main implementation challenges were investigated and resolved. A discussion of the developed mechanism for decoding raw H.264 streams, the design decisions and implementation procedures that ensure portability, mobility and affordability, and the bespoke protocol implemented to manage the communication between the server and clients, is provided as follows.

6.5.1 Decoding raw H.264 streams

As stated in section 6.5, the server application implements a callback (provided by the SDK) which fires when a new chunk of raw video is available from the source. The raw video chunks put together represent a stream of raw H.264 video which must be parsed and decoded in order to be rendered by the client. Decoding and rendering a raw H.264 video stream on Android is executed using the MediaCodec class (Fig. 6.6). The class provides a set of input buffers which must be fed with H.264 Access Units (AUs) and output buffers which render the decoded

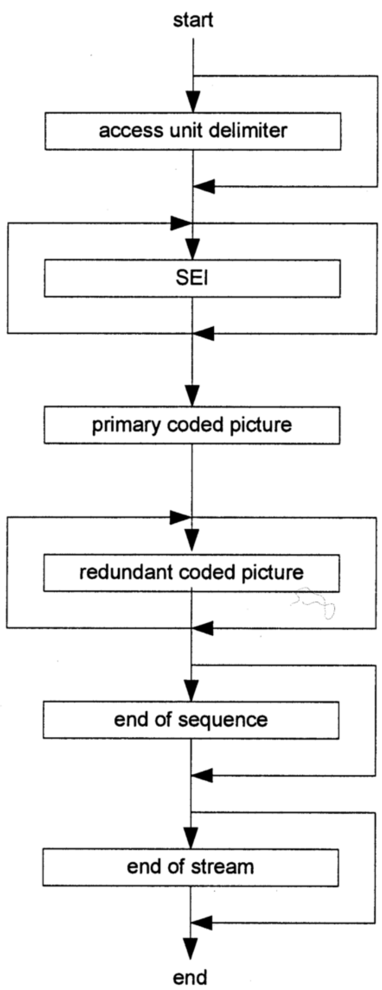

Figure 6.5: Access Unit Structure (Source:[305])

frames onto an Android surface [303, 304]. Decoding this raw stream therefore necessitates an understanding of the structure of raw H.264 streams.

A raw H.264 stream is made up of one or more Network Abstraction Layer Units (NAL Units) which collectively contain both frame payload and key information needed to decode each frame [305]. The raw stream has delimiters that signal the start and end of the NAL Units, and these NAL Units must be parsed into AUs (Fig. 6.5), which can then be fed into an Android MediaCodec decoder that has been initiated and configured. This process was used to implement a parser class that configures a MediaCodec object, continuously parses the chunks of the raw H.264 video into video frames and seamlessly renders them to a VR surface in real-time. The decoding process was designed to run on an independent thread, separate from the thread responsible for receiving the packets so as to ensure that the decoding process is not hampered by the receiving function (which periodically blocks due to the nature of network input/output operations). The decoding process was also optimised to minimise overhead so as to ensure adequate processing power for the other tasks on the client devices.

6.5.2 Portability, mobility and affordability

The system was designed with the requirements that it should be portable – it should work with minimal configurations irrespective of the environment, mobile – it should require minimal fixed infrastructure, and affordable – it should work with inexpensive components. The system portability is fulfilled by the design of the communication paradigm (discussed in section 6.5.3), which enables it to function on any wireless network that supports and permits multicast communication. This implies that the network component of the system can be fulfilled with an off-the-shelf Access Point, Raspberry Pi, smartphone or computer that is capable of hosting a wireless network, albeit subject to sufficient bandwidth and transmission range of the device. Mobility – and by extension portability – is also fulfilled by the fact that the server and client are implemented as mobile apps that can be distributed and downloaded from global App stores, and do not require any pre-configuration owing to the discovery protocol implemented, as discussed in section 6.5.3. In addition, the portability and mobility of the system combine to ensure that it can be actualised with minimal costs. While some costs are inevitable (for example the cost of hiring or purchasing a drone), a smartphone can be procured or re-purposed to act as a server, tourists often travel with their smartphones and can download the client app for free, and some heritage organisations have wireless infrastructure in place or can install one for the cost of a Raspberry Pi ($\approx$ £30). This keeps costs to a minimum and thus ensures affordability.

6.5.3 Communication and control

Communication between the server and client takes place via datagram packets over wireless networks. This uses the User Datagram Protocol (UDP) at the transport layer which is characterised by fast, connectionless, albeit unreliable (in terms of packet ordering and delivery) communication [306]. This makes UDP appropriate for video streaming and other applications that are either not sensitive to (minimal) packet loss or do not benefit from retransmission of lost packets. This informed the decision to use UDP for all communication between client and server, and this communication employs both unicast and multicast datagram packets.

When the server is initiated (i.e. the smartphone is connected to the wireless network, plugged into the drone controller and the server app is launched), it establishes a multicast socket on a port and begins listening for discovery requests from clients. When a client is initiated (i.e the smartphone is connected to the wireless network and the client app is launched or restarted) it sends a discovery request to the multicast group and port on which the server is listening. Upon

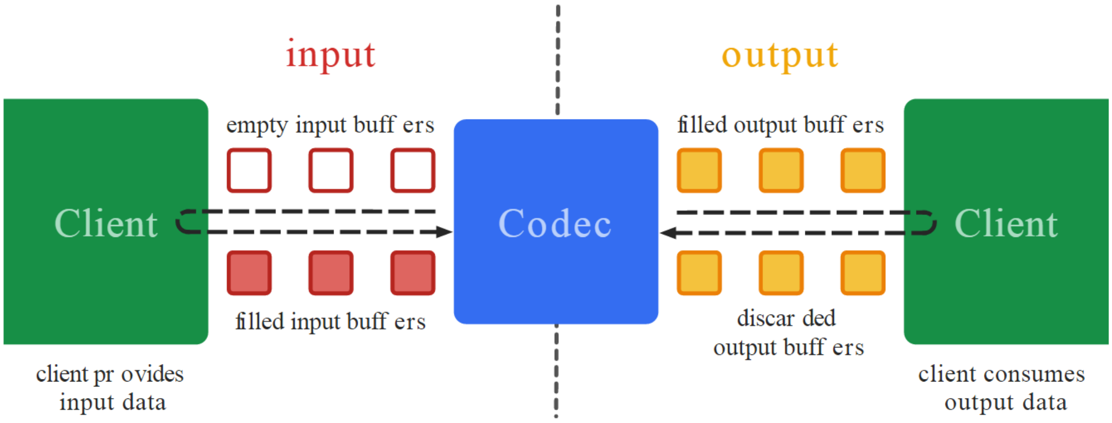

Figure 6.6: Android MediaCodec decoding process (Source:[304])

receiving this request, the server sends an acknowledgement to the client and adds the client (an object represented by an IP address and port number) to a list of connected clients. Clients can also unsubscribe from the video stream by simply closing or minimising the app, which generates and sends a packet, notifying the server that it is leaving the stream. Upon receiving this notification, the server removes the client from the list of connected clients. This ensures that the server can maintain a list of how many clients are subscribed to the stream at any point in time.

By keeping track of connected clients, the server can delegate control of the gimbal to each client in turn and switch control between them. The server implements this control protocol by periodically sending packets which either authenticate clients to (or instructs them to not) send sensor coordinates. When a client has control of the gimbal, it continuously sends sensor coordinates – a reading about three axes (pitch, roll and yaw) that represents the device's orientation in space – to the server. Upon receipt of sensor readings from the controlling client, the server updates the drone gimbal to reflect the orientation of the client, ensuring that the drone's camera orientation aligns with that of the controlling client. Only one client can have control of the gimbal at a time, and the transfer of control between clients depends on which of the two control modes – *tour* and *explore* – is selected. When the *tour* mode is selected, the server always assigns control to the same client. This mode is ideal for giving a guided tour of a landscape, where the tour guide's device is the designated client and thus control's the drone's gimbal. When the *explore* mode is selected, the server rotates control between each client, such that one client has control for a fixed duration, then control passes on to the next client in the list, and so on until control gets to last client and goes back to the first client. This mode is ideal for when a group of people wish to explore a landscape at will, such that each individual benefits from controlling the gimbal for a period of time.

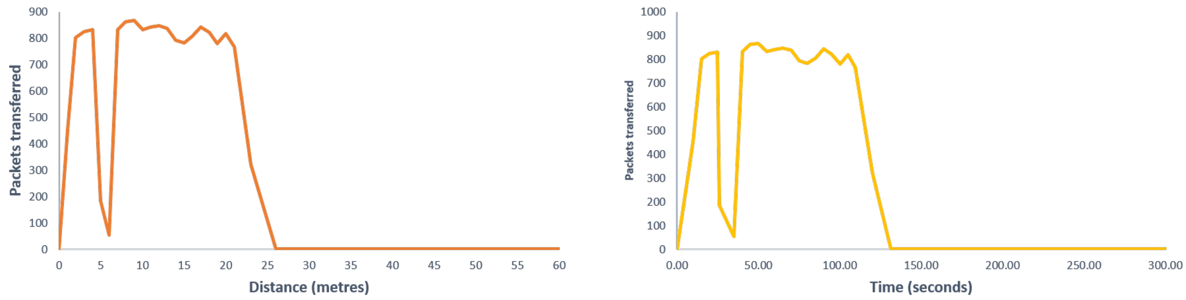

Figure 6.7: Packets transferred vs Distance (left) and Time (right)

6.6 System Evaluation

The relationship between effective transmission speed over the network and the distance between the server and clients was investigated to confirm the transmission speeds that could be attained on the network as well as how far the client devices could be from the server while still receiving footage. In addition, an investigation was carried out to ascertain whether the attainable transmission speeds could support the required bitrates for live video playback.

6.6.1 Transmission Speed and Distance

Given the system architecture shown in Fig. 6.2, there is a question of whether packets from the server are transferred directly to the clients or whether the packets are first sent to the Access Point (Raspberry Pi) which then sends the packets to the clients. Most modern access points act as facilitators that orchestrate communication between devices connected to the network medium without interference, hence it is expected that the Raspberry Pi – acting as an access point – will transparently orchestrate communication without interfering with the process. This distinction is important because if packets are transferred directly from the server to the client, then the maximum transfer distance is determined by the (smaller of the) wireless range of the server and a client that subscribes to the video stream. Furthermore, there is no additional point of failure. On the other hand, if packets are transferred through the Raspberry Pi, then the maximum transfer distance is determined by the transmission range of the Raspberry Pi. Also, this introduces an additional point of failure at which packets can be lost.

tcpdump -i wlan0 -n src host server-ip dst host client-ip

Table 6.1: Using tcpdump to capture packets based on source and destination IP address

Timestamp	Source	Destination	Protocol	Info	Distance
...	...	...	...	...	...
119.967768	192.168.42.39	192.168.42.18	UDP	36136 > 55555 Len=1	23
119.973204	192.168.42.39	192.168.42.18	UDP	36136 > 55555 Len=1	23
131.31901	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	26
156.849807	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	31
165.81583	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	Who has 192.168.42.18? Tell 192.168.42.39	33
181.659965	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	36
206.470458	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	41
225.816298	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	Who has 192.168.42.18? Tell 192.168.42.39	45
231.281178	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	46
256.141488	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	51
280.962092	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	56
305.782807	LgElectr_7a:8d:e2	Wisol_6e:f2:d1	ARP	192.168.42.39 is at c4:9a:02:7a:8d:e2	61

Table 6.2: Snippet of the packets captured between the server and client

To verify this, a stream was set up between the server and a client over the Raspberry Pi Access Point wireless network, and a packet sniffer was set up (in promiscuous mode so as not to interfere with the network) to capture packets being transferred on the network. The packet sniffer was set up on a Linux workstation and the wireless adapter of the workstation was configured so that its channel was set to the same channel on which the Raspberry Pi Access Point was operating. The network traffic was then sniffed using *tcpdump* (using the command shown in Table 6.1). The server and client started next to each other, and were moved one metre (1m) apart at ≈5 second intervals while packets were captured for ≈306 seconds (just over 5 minutes). In Fig. 6.7, the number of packets transferred are plotted against distance. The graph shows that communication failure occurred where the distance between the server and client was ≈25 metres (after ≈120 seconds). As the client and server were taken further apart, the speed and the number of packets transferred reduced, and beyond the communication range, video stream packets were not transferred, but rather Address Resolution Protocol (ARP) requests were transferred instead (as observed in the sniffed packets), where the client and server were attempting to resolve each other's Media Access Control (MAC) addresses. A snippet of the packet sniff in Table. 6.2 shows that after successfully transferring UDP packets (video chunks) for up to ≈120 seconds, the devices began sending ARP requests to reconcile each other's Internet Protocol (IP) addresses. This shows the server (192.168.42.39) and client (192.168.42.18) devices communicating directly while the Raspberry Pi Access Point orchestrates the communication between them and provides the wireless infrastructure over which packets are sent.

```
ffmpeg -i input.mp4 -vcodec copy -bsf h264_mp4toannexb -an -f h264 output.h264
```

Table 6.3: Converting an mp4 video file into a raw H.264 file

Resolution	MP4 size (bytes)	H.264 size (bytes)	H.264 length (Java byte array)
480x270 (Low)	6,148,472	4,438,884	4,438,884
640x360 (Medium)	14,516,753	12,922,727	12,922,727
960x540 (High)	41,441,394	39,206,325	39,206,325

Table 6.4: Resolutions and sizes of videos used for testing

6.6.2 Video Playback and Bitrate

To test the streaming capabilities of mobile devices, a live video stream was simulated by converting an mp4 video file – available in three resolutions – into raw H.264 files. These files – created using the ffmpeg command shown in Table 6.3 – were then read as input streams into a byte array in Java, and chunks of the file were transferred as datagrams over the network. Table 6.4 shows the resolutions and sizes of videos used for testing, Listing 6.1 shows how the byte array input stream was created and Listing 6.2 shows how chunks of the raw video were sent as datagrams. Recall that the bitrate of a video is the ratio of its file size to its duration i.e. $bitrate = filesize/duration$, expressed in Megabits per second (Mbps). The video files in Table 6.4 have the same duration but different file sizes, hence they require different bitrates to be supported for successful playback. Samples of the raw H.264 files were sent from the server to a client device in chunks, and the chunks were reconstructed on the client to simulate a video stream. As expected, the transmission did not produce an intelligible video feed on the client without any time-based mediation of the playback. Playback was mediated by introducing "sleep" intervals between chunks and it was observed that changing the interval affected the intelligibility of the video feed.

```java
//...
try
{
  InputStream inputStream = getResources().openRawResource(R.raw.low);  //low represents a low resolution file stored in the resources folder
    int length = inputStream.available();
    byte [] someBuffer = new byte[length];  //someBuffer represents a byte array which will store the raw h.264 file contents
    inputStream.read(someBuffer);
}
catch (Exception e) {e.printStackTrace();}
//...
```

Listing 6.1: Creating a byte array stream from a raw H.264 file in Java

```java
//...
try
{
  List<byte[]> chunks = BufferUtils.chunkBytes(someBuffer, someBuffer.length, 1348);  //someBuffer contains the raw h.264 file contents, and is divided into chunks of length 1348
    BufferDatagram tempBufferDatagram;
    while (true)
    {
    if(DatagramSender.hasDatagramSenderThreadStarted)
        {
      for (byte[] c : chunks)
            {
        tempBufferDatagram = new BufferDatagram(c, c.length);
        tempPacket.setData(tempBufferDatagram, 0, tempBufferDatagram.length); // store the chunk contents in a packet
        tempSocket.send(tempPacket);  //send the packet using a socket
                Thread.sleep(5);  //sleep for 5 milliseconds
            }
        }
    }
}
catch (Exception e) {Log.e(TAG, "Could not load raw video"); e.printStackTrace();}
//...
```

Listing 6.2: Sending raw video chunks as datagrams in Java

Multiple streams were simulated using the low resolution H.264 file created using the command in Table 6.3 with different values for the interval, while keeping the chunk size constant at 1,348 bytes. The chunk size of 1,348 bytes was chosen to ensure that after the addition of headers at the network and transport layers, the resulting chunk remains less than the Maximum Transmission Unit (MTU = 1,500 bytes) of the transmission medium. At an interval of ≈25 milliseconds (±3 milliseconds), an intelligible video stream was rendered by the client. With an interval of 25 milliseconds, and a chunk size of 1,348 bytes, 53,920 bytes per second ((1,348/25)/0.001) were transferred between the server and client. This amounts to 431,360 bits per second (53,920*8) hence the bitrate of the transmission was calculated to be ≈0.43 megabits per second (0.43 Mbps). In Table 6.4, there are 2.9 times more bytes in the medium resolution file than the low resolution file. Therefore, for the medium resolution video file to produce a stream similar to that obtained for the low resolution file, a bitrate of 1.25 Mbps ($0.43 * 2.9$) must be supported by the transmission. Similarly, a bitrate of 3.83 Mbps must be supported while streaming the high resolution video in order to get a stream akin to that obtained from the low resolution file. As the duration (time) is identical for the three video files but have increasing file sizes

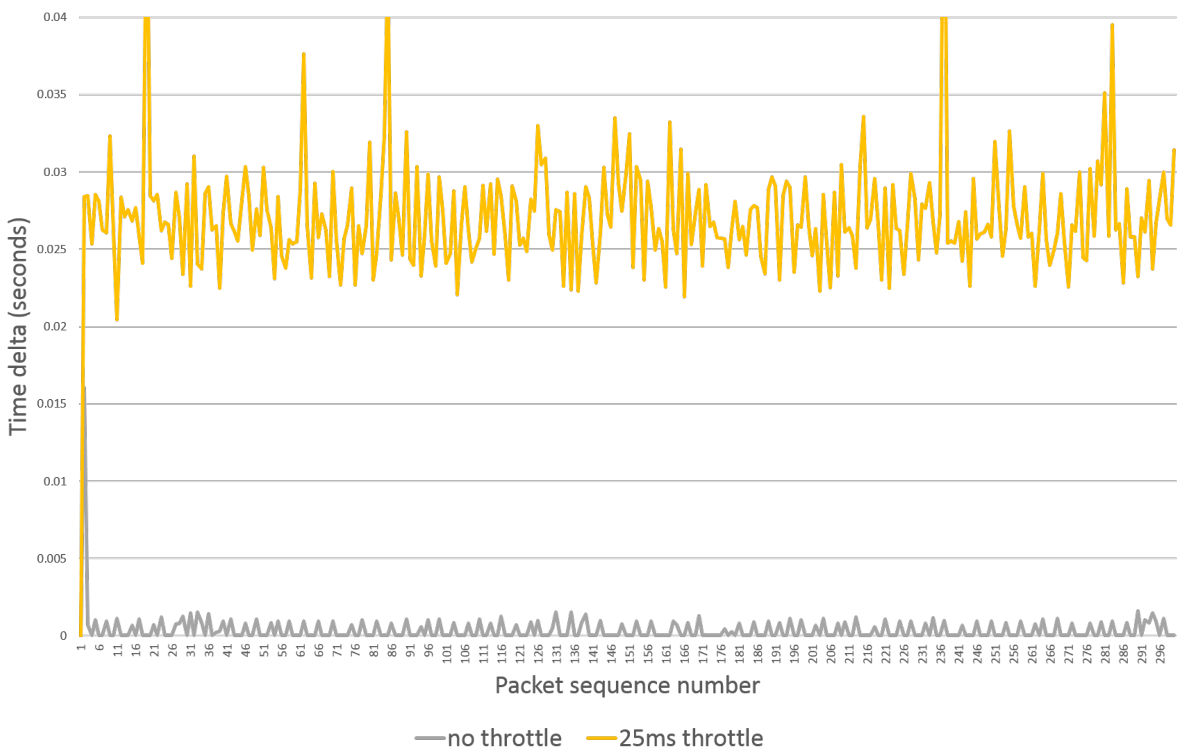

Figure 6.8: Time difference between packet transfers (with and without set intervals)

as the resolution increases, to maintain playback at the required bitrate, more bytes per second must be streamed and rendered. This can be achieved by either increasing the size of each chunk or decreasing the interval between chunks. The chunk length can be increased to 3,919 bytes ($1.25/0.43 * 1348$) and 12,007 bytes ($3.83/0.43 * 1348$) to support the medium and high resolution files respectively. However, this results in individual chunks that are greater than the MTU, causing the chunks to be fragmented and reconstructed before and after transmission respectively. This introduces unnecessary network overhead and may affect the quality of the video stream. A better alternative is to decrease the interval between chunks instead of increasing the chunk size. Thus, the chunk intervals were throttled for the medium and high resolution files, and it was observed that using an interval of $\approx$3 milliseconds for the high resolution video and $\approx$9 milliseconds for the medium resolution video. Intervals outwith of these values by $\pm$3 milliseconds produced jittery streams as clients either received packets faster or slower than required for playback. For comparison, Fig. 6.8 shows the time difference (in seconds) that elapsed between packet transfers (as time deltas), plotted against packet sequence numbers for the first 300 packets, with (yellow) and without (grey) an interval period to throttle the transfer speed. The graph shows that a preponderance of the points on the yellow line lie around 0.025 seconds (or 25 milliseconds, the specified interval) while the points on the grey line lie below 0.001 seconds (or 1 millisecond), indicating that the system is capable of handling lower throttle

Item 1 (System Usability) - Is this system awesome or weird? Why?
Item 2 (System Control) - Did you feel like you were in control of the system? Why?
Item 3 (Social Experience) - Did you feel like this was a good social experience? Why?
Item 4 (Miscellaneous) - Any other comments?

Table 6.5: Questions used for semi-structured interviews

values (of less than 1 millisecond) in order to accommodate higher bit rates.

6.7 User Evaluation

A user evaluation was conducted in a natural setting, where participants in a local community were invited to sign-up to an event involving drone flight over the local landscape. The aim of the user study was to investigate the usability and potential value of the system, as well as how it supports social experiences in the context of heritage exploration. The event took part in the local community park where the participants gathered. A drone pilot flew the drone over the local landscape in a slow, circular motion so as to observe the iconic parts of the community, including a monumental statue, a river and bridge, an old church (with a tower and clock), as well as the local heritage centre and community centre. While the drone was being flown, pairs of participants used the VR headsets to receive the drone footage, where one person actively controlled the gimbal while the other passively received the drone footage. The control switched between the pair every minute, so that the active client became passive and the passive client became active. Each pair had 5 to 10 minutes to engage with the landscape using the headset. Data was collected using an observation technique, questionnaires filled by the participants and semi-structured interviews conducted after the event using the questions in Table 6.5 as the basis for the interviews. The questions in the questionnaire (and consequently the interviews) were chosen to elicit feedback on usability, control (or lack thereof) and social interaction, as applicable to the system, and these formed the basis of the discussion. There were 8 participants (6 female) aged between 12 and 17, thus there were 4 pairs, and participants in each pair were familiar with each other. The exercise took approximately an hour to complete. The findings of the study will be discussed in the following categories: *system usability*, *system control* and *social experience*.

6.7.1 System Usability – How well did the system work?

Overall the feedback from participants was positive. They reported that it was a valuable system for exploring landscapes. This was suggested by comments such as *"it was awesome because you get to see the town from different heights, angles and sides"*, *"it lets you see the world from a* *different perspective"*, *"...you would not see stuff like that when you are on the ground"*, *"...you can ... see what is around you and high up that you can't see"*, and *"you get to see the town from above"*. Participant's feedback also bordered on feelings of immersion in a virtual environment, as evident from comments such as *"it's like you are in a virtual world"*, *"it's like you are inside a game"* and *"you feel like you are flying"* . Overall, participants thought it was an exercise that they would like to do more often.

An interesting observation made while the participants used the system was that some participants did not immediately find the screen on which the real-time video was displayed. The video was displayed on a screen represented as a rectangular plane in a 3D environment and it had a constant position and orientation in 3D space. Participants first instinct when they put on the headset was to look forward i.e. in the direction they were facing when they put on the headset. However, the forward position for some participants did not correspond with the position of the screen; the screen may have been behind or to either side of the participant such that the screen was completely out the participant's field of view. All of the participants eventually found the screen after no more than a few seconds of looking around, but some participants took longer than others and may have missed some footage during the time it took to find and look at the screen. This observation was the basis for updating the behaviour of the system (and subsequent similar systems), so that when a 3D environment features a plane (or screen) on which content (such as videos or flat images) is displayed, the plane should always be placed in front of the user by aligning the plane's position with the forward orientation of the smartphone (obtained from the phone sensors). Placing a screen in front of the user ensures that the user does not miss out on content while trying to find the screen in the virtual environment.

A participant exclaimed about being scared shortly after putting on the headset. During the post-exercise interview the participant stated that they felt like they were flying and this caused mild discomfort because of their fear of heights. No other participants reported a discomfort caused by a fear of heights, but as discussed in the *system usability* section, participants' comments bordered on feelings of presence in an aerial virtual world. This affirms the case for the use of mobile VR to immerse users in virtual environments and thus provide compelling engagement with content.

6.7.2 System Control – How did users find the illusion of control?

As discussed in section 6.5.3, a protocol was implemented to periodically switch control of the gimbal between clients, such that one client's movements determine the orientation of the drone's gimbal at any point in time, and that client changes after a set period of time. Given the system set-up which immersed participants in the virtual environment, it was important to investigate how this illusion of control was perceived by each participant, both at times when each participant had control of the gimbal and at times when they did not have control of the gimbal. Feelings about system control were mixed. Participants were impressed by the ability of the drone's camera to match their head orientation. One participant felt they had sufficient control of the system because *"... you just nod your head and it moves."*. Another participant thought it was unusual at first but after getting used to it, they felt in control, and two others felt they were in control some of the time (when it responded to their head movements) but not at other times when their head movements had no effect (which was probably when the other individual in their pair had control). One participant felt like they were not in control because *"...it kept moving when I was still..."*, while referring to periods when the other individual in their pair had control, and two participants wanted to fly the drone themselves so as to not only control orientation but also location.

6.7.3 Social Experience – How did social interaction unfold?

Participants unanimously felt that the system provided a good social experience. They referred to the ability to *"chat"* to each other while using the headsets and this was important because they could discuss shared views. The provision of shared, common views enabled participants to refer to things in their discussion, as they would have if they were viewing content on a big screen. This improves on the state of the art in headset-based VR systems because these systems often provide personal experiences where a user is totally engaged with a system and is thus cut-off from their physical environment including the people in that environment. However, by providing a shared, immersive view, the system incorporates the social benefits gained when two or more users view content on a screen and the immersive benefits gained when users view content in stereoscopic headsets. Participants also cited the novelty factor, with one stating that *"...it is something different"*, and with another stating that *"...it is something that everybody likes"*. One participant also stated that *"...we were all talking to each other when we were flying it'*. It is pertinent to note that the participant used the phrase *"when we were flying it"*, which suggests that they felt they were in control i.e. flying the drone.

6.8 Discussion

The use of affordable, mobile technologies for the remote exploration of inaccessible, distant or endangered heritage sites has been discussed. Emphasis has been placed on portability – the system can function regardless of existing infrastructure and can thus be used in various settings and locations, mobility – the system works with users' smartphones and thus allows free movement during interaction, and affordability – the system components can be cheaply purchased or freely repurposed. This work also demonstrates how digital means can support the dissemination and preservation functions of traditional museums; dissemination is demonstrated through remote exploration of sites, while preservation is actualised by encouraging remote (rather than on-site) access to endangered sites thus ensuring their longevity.

It is pertinent to note that the server application used in this work is implemented using the DJI mobile SDK and is hence limited to DJI drones. The implication of this is that although DJI are one of the leaders in the market for portable, professional and personal Unnammed Aerial Vehicles and they boast an array of products in their range [307], this decision has introduced a vendor lock-in with switching costs. Nonetheless, the design principles, findings of the system evaluation and recommendations provided in the discussion represent a detailed resource for developing an affordable, portable drone-based VR system for real-time exploration of heritage landscapes.

Limitations of the technology also imposed challenges on the system functionality. The range of the Raspberry Pi's wireless interface limits the distance that the clients can be away from the server and still receive a live video stream. Although this is not an issue in the primary use case, i.e. while using the system outdoors, it poses a challenge in settings where the clients need to be at a considerable distance away from the server. For instance in use cases when the clients wish to receive footage in an indoor setting while the server is located outdoors (in a safe flying zone) with the drone pilot. In the same vein, drones have a maximum transmission distance beyond which the aircraft cannot communicate with the controller or base station and a maximum flying time of twenty-five (25) minutes or less. This imposes further limitations on the size of the area which can be explored and the duration of a session, although the latter can be remedied with the use of multiple batteries.

There are also socio-technical and environmental considerations in the use of drones such as issues with legislation, privacy, public acceptance and the logistics of simultaneously routing multiple drones (akin to road networks for vehicles or rail networks for trains) [308, 64]. Furthermore, a skilled pilot may be required to successfully and safely fly a drone over an area of interest, and considerable time and effort are required in order to become a proficient drone pilot. The lack of skilled personnel may result in the absence of a crucial (human) component of the system and consequent inability to operate altogether, or the risk of harm to people, property and the environment which may have undesirable (and in some cases catastrophic) consequences. That said, autonomous drones which can be programmed for specific routes and trained to avoid collisions are becoming mainstream by virtue of advances in computer vision, machine learning and artificial intelligence [62, 118, 64]. In addition, drones are becoming more affordable in line with trends in other aspects of computing such as processors, mobile and ubiquitous devices. These will facilitate even more possibilities and use cases for virtual tourism and remote exploration of inaccessible heritage sites.

6.9 Chapter Summary

The motivation and methodology for designing, implementing and deploying a system for real-time, remote exploration from an aerial perspective have been discussed in this chapter. An architecture which ensures that it can function in remote areas (where electricity and Internet access are not guaranteed) has been designed and the system has been implemented using affordable system components which make it feasible for participatory heritage engagement. A system evaluation was conducted which provided input into the re-design process, and a user study was carried out to investigate the potential value and usability of the system. The findings suggest that the system is a valuable and cost-effective tool for heritage exploration, as it enables users to observe a landscape from an aerial perspective and it provides an immersive, yet socially-inclusive experience. The system contributes to the state of the art in remote tourism as it combines the engaging benefits of personal, immersive systems with the social benefits of large displays, and it facilitates the exploration of inaccessible areas, distant locations and sites that are at risk of being damaged when accessible to the public, hence it facilitates heritage dissemination and preservation. The discussion affirms the case for the use of mobile immersive systems to deliver engaging experiences and for remote exploration of heritage which constitutes one of the modes in which virtual museums exist as discussed in section 2.3, hence if offers a new dimension for heritage exploration.

CHAPTER SEVEN

EXPLORATION IN MUSEUMS

This chapter discusses a series of workshops conducted in museums in Europe, Latin America and the Caribbean, to investigate existing infrastructure as well as empower heritage experts and community members with the technical skills required to digitally disseminate heritage, while acquiring digital output in the process using a co-creation approach.

7.1 Introduction – Museum Workshops

Over a seven-month period, workshops were held in museums across Europe, Latin America and the Caribbean (EU-LAC), in conjunction with a team of academics at the University of St Andrews. The EU-LAC[1] project is a multinational project aimed exploring how affordable technologies can be leveraged to create digital representations of heritage, while fostering and building relationships between these museums and their respective communities, as well as the relationships between community museums. An overview of the museums involved in the workshops, including their location (country and region) and main governing body is provided in Table 7.1. The workshops entailed working with museum staff, young professionals and community members to explore how affordable technologies can be leveraged for heritage learning and dissemination. Participants learned how to use affordable technologies to create digital representations of their community heritage, as well as how the digital outputs can be combined to provide a comprehensive online presence for the museum and the community. A description of the activities and outputs of the workshops is provided in the following sections.

7.2 Workshop Content

The museum workshops generated digital content in various forms including 3D artefacts, spherical imagery, flat imagery, audio, video and text for curation. These processes involved in creating and curating these content is discussed in the following sections.

[1]The EU-LAC project is funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 693669.

S/N	Name	Governance	Location
1	Moruga Museum	Community	Trinidad, Caribbean
2	CharlesTown Maroons Museum	Community	Jamaica, Caribbean
3	Curre	Community	Costa Rica, Latin America
4	Community museum of Boruca Indians	Community	Costa Rica, Latin America
5	San Vicente Ecomuseo Nicoya Costa Rica	Community	Costa Rica, Latin America
6	Barbados Museum and Historical Society	National	Barbados, Caribbean
7	National Museum and Art Gallery of Trinidad and Tobago	National	Trinidad, Caribbean
8	National Museum Jamaica	National	Jamaica, Caribbean
9	University of the West Indies Museum	Other Body (University)	Jamaica, Caribbean
10	Seixal Eco Museu	Regional Body (Municipal)	Portugal, Europe
11	National Archaeology Museum Lisbon	National	Portugal , Europe
12	Shetland Museum and Archive	Regional Body (Municipal)	Scotland, Europe
13	Unst Heritage Centre	Regional Body (Municipal)	Scotland, Europe
14	Unst Boat Haven	Regional Body (Municipal)	Scotland, Europe
15	Museu Comarcal de l'Horta Sud	Regional Body (Municipal)	Spain, Europe
16	Community Museum of Malalhue	Community	Chile, Latin America
17	Sican Museum	National	Peru, Latin America
18	TUCUME museum	National	Peru, Latin America
19	Museum of the Universidad Austral de Chile	Other Body (University)	Chile, Latin America
20	Timespan Museum and Art Gallery	Community	Scotland, Europe
21	Cortes de Pallas	Regional Body (Municipal)	Spain, Europe
22	Museo del Palmito d'Aldaio	Community	Spain, Europe
23	Taigh Cheasabhagh Museum and Art Gallery	Community	Scotland, Europe

Table 7.1: Overview of museums in the EU-LAC project

7.2.1 3D Artefacts

Chapter 1 introduces the concept of 3D artefacts (or objects) that can be used to digitally represent tangible museum artefacts such as preserved specimens, documents and sculptures. These digital artefacts have the ability to facilitate interaction (such as rotating and zooming) which can provide an improved sense of engagement with heritage content when the physical artefacts are inaccessible or at risk of deterioration. 3D artefacts can be created using scanning technologies, and Structure From Motion (SFM).

Two types of scanners – a Laser Scanner (pictured in Fig. 7.1) and a Structured Light Scanner were used to create digital representations of museum artefacts. Scanners can produce digital representations of artefacts (as 3D objects) relatively quickly; the Laser Scanner pictured in Fig. 7.1 can create a High Definition (HD) model in $\approx$ 60 minutes or a Standard Definition model in about half that time. Scanners are also usually reliable, as they provide instantaneous feedback while automating the digitisation process. However, the technology is usually expensive; the Laser Scanner pictured in Fig. 7.1 costs $\approx$ £4,000 (as of the time of this writing). For this reason, scanning technology may be out of reach of museums and heritage organisations hence Structure From Motion (SFM) – a popular type of photogrammetry [309] – represents a more affordable alternative as it is significantly cheaper than laser scanners [310].

Structure From Motion (SFM) (commonly referred to as photogrammetry) is an affordable

Figure 7.1: Digitising an artefact using a laser scanner (left) and photogrammetry (right)

alternative to scanning, which involves digitally recreating the structure of an entity (an artefact, building or even human body part) from dozens (and sometimes hundreds) of overlapping photos of that entity, captured from various perspectives using a camera [309]. This approach can cost significantly-less than scanning because the minimal requirement is the use of a camera; a mobile phone camera will suffice. A more rigorous photogrammetry set-up (as pictured in Fig. 7.1) may involve a mid- to high-end camera, lights and a soft box for improved control over the environment. These maintain consistency in lighting and improve the quality of the images captured. To create 3D artefacts using photogrammetry, a physical object is placed centrally on a circular base in the light box to ensure uniform background and even distribution of light (see Fig. 7.1). Dozens of images are captured by taking images between overlapping, stepwise rotations of the base, until the object can be seen from all perspectives in the images. This may result in around 70 to 150 images depending on the nature of the physical object. The images are then processed using a SFM toolkit[2] to produce a 3D model. The processing entails pairwise matching of images to identify common points, creating sparse and dense point clouds, generating and trimming meshes, and exporting the meshes as 3D models. Metadata (such as the title, description, size, origin and location of the object) are then gathered and uploaded together with the 3D model using the content management interface.

While photogrammetry is more affordable than scanning [310], the rigorous nature which requires the methodical capture and processing of dozens (sometimes hundreds) of images, and the required computing resources (in terms of processing, memory and graphics) may pose challenges. Nonetheless, given that museums and heritage organisations may be of the small to medium category of enterprises with limited operating costs, photogrammetry is often a more practical alternative to scanning.

[2]Visual SFM (http://ccwu.me/vsfm/) and MeshLab (http://www.meshlab.net/) are free, open-source tools that can be used collectively for SFM. This makes them beneficial and attractive to community museums with limited resources.

Photogrammetry and scanning are two examples of 3D imaging techniques that have recently come to the fore, and the trend towards the reduction in cost of these techniques is beneficial to heritage organisations as well as organisations in other domains [311]. There is a plethora of work that leverage these technologies for content digitisation in diverse application domains including museums, archaeology, geographic surveys, medicine and law enforcement [119, 310, 312, 313, 121, 311, 314]. While these 3D imaging techniques facilitate interactive visualisation and consumption of digitised content which is useful in preserving, recreating and disseminating subject matter, there are challenges that need to be addressed in order to fully actualise their potential [315]. These challenges include the intellectual property issues associated with managing the rights and dissemination of 3D models, the interoperability of 3D model formats and software, and the paucity of infrastructure for indexing and searching 3D repositories [315].

If the trends in the popularity and cost reduction of digitisation technologies hold, issues with intellectual property and format incompatibility will likely become less prevalent as industry standards are developed for metadata and software respectively. This is already evident through the development of the Dublin Core schema [316] which is aimed at providing a universal set of terms for cataloguing, archiving and disseminating heritage content, and the WebGL framework which enables web browsers to render and interact with 3D models [317]. Furthermore, the proposition of a Virtual Museum Infrastructure (VMI) as detailed in this thesis, addresses the issues associated with indexing, searching and interacting with 3D repositories, and thus equips heritage practitioners and community members with the resources required to digitise, archive and disseminate heritage content. By lowering the technical barriers to entry, stakeholders can better engage with (and be more invested in) heritage management processes, which translates to a strong sense of participation and better appreciation of heritage [318].

As discussed in section 1.1.2, the digitisation of physical artefacts to result in 3D artefacts adds an additional dimension in the digital representation of heritage as it facilitates the representation and preservation of tangible cultural and natural heritage such as specimens, documents and sculptures. This constitutes part of a larger body of work that applies digital technologies for heritage preservation [121, 313, 158, 319], and hence informed the incorporation of photogrammetry and scanning into the workshops. Their differences aside, both photogrammetry and scanning result in the creation of 3D artefacts that represent physical objects. These artefacts may be viewed as annexes (representations or replicas) of the physical objects they represent, or as artefacts in their own right. This distinction is important because of the proliferation of 3D printing technology which enables users to create physical representations of 3D models. By virtue of this possibility, an individual may create a digital 3D model of an imaginary entity using computer graphics software and then 3D-print it, which would constitute the reverse of digitising a physical object to result in a 3D model. Given that the link between a 3D object and a corresponding physical object has become bi-directional (because either the physical or the digital object can exist first), it may be appropriate to consider a 3D object and a corresponding physical object as separate entities irrespective of provenance and the relationship between them.

7.2.2 Spherical Images

Chapter 1 introduces the concept of spherical images (also known as equirectangular images or photospheres) that can be used to make virtual tours which immerse users in remote landscapes and cityscapes, or provide visual representations of the past, thus breaking barriers of time and space. The last decade has seen an increase in the availability of cameras meant for capturing spherical images. These include the Ricoh Theta[3], the iris360[4] and the bublcam[5]. Smartphones are also equipped with cameras capable of capturing images, and app stores feature a plethora of apps that are designed for capturing spherical images. The availability of spherical image-capturing apps, coupled with smartphones equipped with mid- to high-end cameras enables the general public (including heritage experts and community members) to freely create virtual tours that can be disseminated over the Internet without purchasing specialised equipment such as spherical cameras. As discussed in section 1.1.2, spherical media facilitates the capture and digitisation of landscapes, cityscapes and geological structures which constitute tangible cultural and natural heritage. The use of spherical photography also facilitates the integration of VR devices such as head-mounted displays, and the findings of chapters 4 and 5 suggest that spherical images, coupled with mobile and/or desktop headsets can provide engaging and immersive experiences to users. Chapter 6 also affirms the case for the use of mobile VR headsets for remote exploration, as users reported feeling immersed in the remote scene. Spherical images thus provide an additional immersive element and this makes a case for its inclusion in the VMI.

7.2.3 Multimedia

As discussed in section 1.1.3, audio and video recordings can be used to capture stories or descriptions associated with entities, in addition to being entities in their own rights. For instance, audio recordings can preserve folk tales and other intangible heritage, or provide additional information by accompanying a virtual tour or video with commentary. In a similar vein, video can preserve and disseminate traditional dances or festival events, or be used to explore a remote area. For these reasons, audio playlists and video galleries are valuable additions to

[3]Online: https://theta360.com/en/ (Accessed: 2018-03-09).
[4]Online: https://nctechimaging.com/iris360-pro/ (Accessed: 2018-03-09).
[5]Online: https://bublcam.com/ (Accessed: 2018-03-09).

Figure 7.2: Capturing spherical images (left) and digitising artefacts (right) during workshops

a VMI, and during the workshops, the participants learnt how to adjust camera settings (such as aperture and shutter speed) for capturing images and videos as well as how to capture high quality audio recordings. Participants exercised their skills by creating virtual tours which constituted recorded video footage of themselves at various locations, coupled with recorded audio narratives that explain the significance of those locations. Aerial footage was also captured using drones to provide a different perspective of landscapes. In addition to the aforementioned multimedia elements, the VMI features editable wiki articles that provide text descriptions. This enables heritage experts and community members to describe content (such as 3D artefacts, spherical images and other multimedia). The wiki articles may refer to the structure, origin or significance of an entity, or pertinent stories associated an entity, and typically feature links to related resources such as galleries, collections and museums. During the workshops, descriptions of artefacts (provided by the museum curators) were used to populate the wiki articles for the artefacts digitised using scanning and photogrammetry, and participants also created wiki articles for other entities including spherical images, videos and museum collections.

7.3 Workshop Format

Workshops were hosted in nine countries across three continents, where participants came from a range of backgrounds including museum professionals, photographers, students, community volunteers, school children and academics. Over 300 people participated in the workshops and where possible, an intergenerational element was included to facilitate the passing down of heritage from community elders to younger generations. A video summary of the workshops is provided[6]. Each workshop examined the six stages involved in the successful creation of digital artefacts:

Selecting artefacts with suitable characteristics in terms of size, shape, rigidity, opacity and

[6] Video summary of 3D workshops: http://openvirtualworlds.org/cadam/eulac_promo_video.mp4

texture.

Preparing affordable equipment such as tripods, lights, soft box, turntable, with participants' smartphones and cameras.

Capturing photos of artefacts from different vantage points.

Creating 3D models using open-source software.

Archiving the source materials (such as the captured images) and the created models.

Curating and uploading content and associated metadata to the VMI and social archive sites.

To better communicate the workshop process, an account of two workshops (conducted in Jamaica and Portugal) is presented below.

7.3.1 Charlestown Maroons Museum, Jamaica

Workshops took place in the Maroons community in the Blue Mountains of Charlestown. The community museum was selected for its remote location and its significance to the country's heritage, as the community forms part of a UNESCO World Heritage Site. The Maroon Museum workshop featured the following activities:

Digitising artefacts: Artefacts held in the community museum were selected and digitised using scanners and photogrammetry to result in 3D representations of the artefacts (see Fig. 7.2). The digital artefacts were combined to produce 3D galleries which are accessible using the VMI[7], as well as on Sketchfab[8], thus facilitating global access and (literally) putting the Maroons community heritage on the map.

Capturing spherical media: Spherical images of both indoor (for example the interior of the community museum) and outdoor (for example the surrounding landscape and river) locations were captured for the purpose of making virtual tours which are accessible using the VMI, thus enabling the general public to remotely visit the Maroons community without having to undertake physical travel (see Fig. 7.2).

Capturing aerial footage: Aerial footage of the community was captured using a drone flown over the Asafu yard around the museum, the Blue Mountains and the community river (see Fig. 7.3). The process of capturing this footage was directed by the locals and the output are accessible using the VMI.

[7] VMI map interface: https://eu-lac.org/map/?page=caribbean
[8] Charlestown: https://sketchfab.com/eu-lac-3D/collections/charles-town-maroons-museum

Figure 7.3: Pre-workshop talks (left) and capturing aerial footage (right)

Traditional dance and storytelling: A dance session took place during which traditional beats were played using drums while the locals put on a display of their ancient festive dance techniques. Audio and video recordings of the dance techniques and traditional beats were captured as a means of digitising and preserving intangible heritage. Also, school pupils participated in storytelling activities with local elders. Flat images and video footage were captured to document these activities, all of which are available online via the VMI.

7.3.2 Seixal Ecomuseum, Portugal

Workshops took place at the Seixal Ecomuseum in Seixal, Portugal. The Seixal Ecomuseum is part of a cultural heritage network with significant community significance in the municipality of Seixal, hence it was selected due to its ties to the community and its importance to the greater area of Lisbon. Attendees of the workshops had varying backgrounds such as photography, archaeology and archiving. The workshops began with an introduction of the team, participants and the technologies that were to be presented (see Fig. 7.3). The participants were asked to specify their backgrounds, interest in 3D technologies and museums, and what skills were they hoping to gain from the workshops. Over 90% of the group had medium to high technical proficiency so the workshop format was adjusted to suit the skill level in the group. Similar to the workshops at the Maroons museum in Jamaica, a group discussion about scanning, photogrammetry, spherical imagery and photography took place, followed by hands-on sessions where the participants digitised artefacts and created virtual tours. Participants also carried out the post-processing tasks required to produce 3D models of artefacts, and the created content (and source material) were uploaded to the VMI[9], archived, and shared on social archive sites such as Sketchfab[10] and Roundme[11].

[9] VMI map interface: https://eu-lac.org/map/?page=europe
[10] Seixal 3D gallery: https://sketchfab.com/eu-lac-3D/collections/seixal-eco-museu
[11] Seixal virtual tour: https://roundme.com/tour/105083/view/264860/

7.4 Discussion

The museum workshops demonstrated that the community members possessed the technical skills required for photogrammetry, as a significant number of community members had smartphones and/or tablets with which they could capture photos and videos. In addition, some community members possessed cameras and had experience with high-end photography equipment and media production. These skills in photography and media production were serendipitous in facilitating the workshop activities, namely the digitisation of artefacts, the preservation of intangible heritage by recording footage of folk dances and the documentation of natural heritage by capturing aerial footage of the community mountains and rivers.

While some variation was observed in the resources available to community museums (in terms of infrastructure), there was also consistency in the level of digital literacies available in the communities (in terms of ownership of, and proficiency with mobile devices). The VMI design therefore taps into these digital literacies to maximise the resources available to community museums. This is exemplified in the incorporation of an automated system that allows museums to upload captured photos of a physical artefact, executes the photogrammetry post-processing on a remote server and produces a 3D artefact, thus enabling museums to digitise artefacts irrespective of their possession of capable hardware. Over 100 3D artefacts[12] have been created using this approach. In addition, the workshops feed into the design of the Virtual Museum Infrastructure (VMI), as the activities demonstrate how museum staff can be empowered with the skills required to leverage emergent technologies for curating, preserving and disseminating heritage (see **Q1**, section 1.4), while the use of affordable technology (such as open-source tools) and existing resources (such as manpower and peoples smartphones) demonstrates value maximisation in environments constrained by limited resources (see **Q2**, section 1.4). Furthermore, the increasing availability, affordability and consequent adoption of technology (such as smartphones, computers and cameras) in heritage contexts allows community members to collectively create media, curate content and contribute to museum offerings, thus shifting the responsibility of heritage management from a few heritage experts to the greater heritage community (see **Q3**, section 1.4).

7.5 Chapter Summary

A discussion of the workshops held as part of the EU-LAC project, with a focus on how museum staff and community members can collaboratively create and digitise heritage content, has been presented in this chapter. The discussion has introduced the use of commodity hardware to

[12]Sample 3D artefacts: http://openvirtualworlds.org/cadam/eulac_objects.mp4

digitise heritage in form of 3D artefacts, spherical imagery and traditional multimedia such as audio, video and flat images. The workshops also uncovered the variations in the amounts of resources available to museums; this is further explored in the VMI design and implementation as discussed in chapter 8.

Part IV

Virtual Museum Infrastructure

CHAPTER EIGHT

VIRTUAL MUSEUM DESIGN AND IMPLEMENTATION

This chapter discusses the design, implementation and evaluation of a Virtual Museum Infrastructure (VMI). The design revolves around the functions of museums – heritage curation, preservation and dissemination – and builds on technology trends that have made digital media more affordable and accessible (chapters 1 – 3). The design also incorporates web, mobile and immersive technologies which have been found to stimulate interest, deliver interactive experiences and add value in heritage dissemination (chapters 4 – 6). As part of the design process, workshops were held with museums across three continents during the EU-LAC project, and an evaluation of the workshops informed the situation of the museums on the resource spectrum – a continuum that represents the level of resources available to museums. The VMI has therefore been implemented to account for this resource variation by leveraging existing resources and digital literacies that are available to the public. Furthermore, the VMI facilitates active, collaborative management and content reuse, such that community members and heritage experts can continuously create and manage heritage content, and the created content can be reused on multiple platforms. The deployment and evaluation of the VMI with over 20 community museums worldwide demonstrates its ability to add value and make positive impact in the heritage community.

Relevant publication:

1. Catherine Cassidy, **Adeola Fabola**, and Alan Miller. *A Digital Museum Infrastructure for Preserving Community Collections from Climate Change*. In Workshop, Long and Short Paper, and Poster Proceedings from the Third Immersive Learning Research Network Conference, pages 170–177. iLRN, 2017. [5].

   Contributions: The proposition of a VMI architecture that can maximise value for museums with varying availability of resources along the resource spectrum proposed in section 1.2.

2. Catherine Cassidy, **Adeola Fabola**, Elizabeth Rhodes, and Alan Miller. *The Making and Evaluation of Picts and Pixels: Mixed Exhibiting in the Real and the Unreal*. In Proceedings from the Fourth Immersive Learning Research Network Conference, pages 97–112. iLRN, 2018. [6].

   Contributions: The implementation, deployment and evaluation of an immersive, on-site digital exhibit, and data analyses to ascertain value and impact from visitors' and staff perspectives.

8.1 Introduction

This chapter discusses the design and implementation of the Virtual Museum Infrastructure (VMI), followed by the deployment and evaluation of the infrastructure components in physical museum contexts. The VMI is a suite of tools that facilitate the curation, preservation and dissemination of heritage by practitioners and the general public. Heritage curation is actualised through a management interface that allows practitioners to create, describe, update and manage content. The created content is held in a back end data store, powered by a Digital Asset Management System (DAMS), which ensures the preservation of heritage in digital form. The content of the data store can be disseminated through multiple platforms – via a web-based portal, mobile apps, on-site installations and social archive sites. The findings of the museum workshops discussed in chapter 7 feed into the design of the VMI. The design of the VMI accounts for the variation in the resources available to museums, a phenomenon which is visualised on a resource spectrum. The design therefore incorporates affordable elements that encourage the use of existing resources such as digital literacies where resources are scarce, as well as elements for deploying immersive on-site installations, where resources are available. These components have been instantiated, deployed and evaluated with over 20 museums across Europe, Latin America and the Caribbean, to investigate the value and feasibility of the VMI for supporting the functions of traditional museums from the perspectives of both heritage experts and the greater community.

The rest of this chapter is organised as follows. Section 8.2 discusses the design of the VMI including the design goals, architecture and relationship between entities, section 8.3 discusses the technical implementation of the components, sections 8.4, 8.5 and 8.6 discuss the deployment and evaluation of the VMI in physical museum contexts, section 8.7 reflects on the findings of the work and section 8.8 provides concluding thoughts.

S/N	Name	Building (1)	Power (2)	Computers (3)	Network (4)	Score (/10)
1	Moruga Museum	Yes	No	No	No	1
2	CharlesTown Maroons Museum	Yes	No	No	No	1
3	Curre	Yes	Yes	No	No	3
4	Community museum of Boruca Indians	Yes	Yes	No	No	3
5	San Vicente Ecomuseo Nicoya Costa Rica	Yes	Yes	No	No	3
6	Barbados Museum and Historical Society	Yes	Yes	Yes	No	6
7	National Museum and Art Gallery of Trinidad and Tobago	Yes	Yes	Yes	No	6
8	National Museum Jamaica	Yes	Yes	Yes	No	6
9	University of the West Indies Museum	Yes	Yes	Yes	Yes	10
10	Seixal Eco Museu	Yes	Yes	Yes	Yes	10
11	National Archaeology Museum Lisbon	Yes	Yes	Yes	Yes	10
12	Shetland Museum and Archive	Yes	Yes	Yes	Yes	10
13	Unst Heritage Centre	Yes	Yes	Yes	Yes	10
14	Unst Boat Haven	Yes	Yes	Yes	Yes	10
15	Museu Comarcal de l'Horta Sud	Yes	Yes	Yes	Yes	10
16	Community Museum of Malalhue	Yes	Yes	Yes	Yes	10
17	Sican Museum	Yes	Yes	Yes	Yes	10
18	TUCUME museum	Yes	Yes	Yes	Yes	10
19	Museum of the Universidad Austral de Chile	Yes	Yes	Yes	Yes	10
20	Timespan Museum and Art Gallery	Yes	Yes	Yes	Yes	10
21	Cortes de Pallas	Yes	Yes	Yes	Yes	10
22	Museo del Palmito d'Aldaio	Yes	Yes	Yes	Yes	10
23	Taigh Cheasabhagh Museum and Art Gallery	Yes	Yes	Yes	Yes	10

Table 8.1: Resources available to museums in the EU-LAC project

8.2 Design

The workshops discussed in section 7.1 was aimed at empowering heritage practitioners and community members with the technical skills required to digitise and capture heritage content. However, it also facilitated the investigation of the feasibility of a VMI, and revealed that there is a significant variation in the level of resources available to museums. This phenomenon can be visualised as a spectrum of resources (see Fig. 1.2), such that the museums at one end of the spectrum are better equipped with power and network facilities to facilitate connectivity than museums at the other end. These two situations represent extremes of the resource spectrum with majority of museums falling in-between. Table 7.1 provides an overview of the museums involved in the workshops, including their location (country and region) and main governing body. Table 8.1 shows the same museums and assigns values to them in four categories: availability of building infrastructure, power supply, computing infrastructure and networking facilities. These four categories – building infrastructure, power supply, computing infrastructure, network facilities – are assigned weights based on the prerequisite resources required to install and operate them in museums.

The installation of network facilities often necessitates the availability of computing infrastructure to utilise the equipment, power supply to operate the equipment and a building to house the

Figure 8.1: Spectrum of resources available to museums in the EU-LAC project

equipment. Consequently, the network category is dependent on the other three categories hence it is assigned the highest weight. In a similar vein, the installation of computers often necessitates the availability of power supply and a building infrastructure, but not necessarily a network infrastructure, because computers can be installed in museums to perform tasks like word processing and label design, which can be carried out offline without Internet access. For this reason, the computing infrastructure category is assigned a weight which is lower than the network category but higher than the power and building categories. Similarly, the building category is assigned the lowest weight because a building can be erected or re-purposed to perform a fundamental function of a museum – to be a place where local heritage is managed and celebrated – without power supply, computers and network facilities, albeit this is rare in practice. The four categories – building infrastructure (shortened as "Building"), power supply (shortened as "Power"), computing infrastructure (shortened as "Computers"), and networking facilities (shortened as "Network") – are therefore assigned weights of 1, 2, 3 and 4 respectively. The values for each of the museums are then totalled to yield a score which represents the museum's position on the spectrum shown in Fig. 8.1. The leftmost position on the spectrum represents a score of 0, which would correspond with a museum that has no building, power, computers or network facilities. Conversely, the rightmost position on the spectrum represents a maximum score of 10, which would correspond with a museum that possesses adequate building, power, computers and network facilities. This spectrum is similar to Fig. 1.2 introduced in chapter 1, but has been modified to situate the EU-LAC museums on the spectrum. The numbers above the spectrum line represent museums using their serial numbers in Table 8.1. By this logic, museums 1 (Moruga Museum) and 2 (CharlesTown Maroons Museum) are closer to the minimum end of the spectrum, museums 9 (The University of the West Indies) to 23 (Taigh Cheasabhagh Museum and Art Gallery) are closer to the maximum end of the spectrum and museums 3 (Curre Museum) to 8 (National Museum of Jamaica) are closer to the middle of the spectrum.

It is pertinent to note that the museums on the spectrum will hold the same positions relative to each other if the categories are assigned equal weights as opposed to weighted values; however, the weighted values are used to demonstrate a spread on the spectrum. Also, while the categories – building infrastructure, power supply, computing infrastructure and network facilities – have been considered for this illustration because they collectively describe a museum's ability to host a digital exhibit on-site with remotely-accessible resources, other categories may be relevant for placing museums on a resource spectrum. For instance, access to public (government) funds (as financial resources) and availability of volunteers (as human resources) may be more appropriate categories (in conjunction with others) for describing museums' abilities to initiate and execute community outreach programs. This variation in resources necessitates a design of the VMI that can provide the appropriate level of functionality based on the available resources.

Towards the left end of the spectrum, the design leverages available resources – community members, their mobile devices and their digital literacies – to support heritage curation, preservation and dissemination by providing a local server with an archive and data store, all of which are accessible over a low-powered, local and affordable infrastructure. This local system allows community members to supply digital heritage content and metadata for these content, which are housed in the data store. Tourists, museum visitors and other community members can also access these content over the local system. Towards the right end of the spectrum, the design shows how mid- to high-end computers, VR devices, smartphones and tablets can be combined to cater to museums with adequate resources through immersive, on-site installations. The design of the on-site exhibit leverages the advances in computer graphics, game engines, and digital literacies. The last decade has seen a proliferation in 3D environments and game engines which have lowered the barriers to entry into the games industry and encouraged individuals and small/medium organisations to design and release games and 3D environments [320]. In a similar timespan, digital literacies have risen, as exemplified by smartphone and tablet ownership [66] and worldwide gaming trends [67]. The rationale behind the leveraging of digital literacies is that the familiarity with games technologies as exemplified by the ownership of games devices (consoles and peripherals) and game playing on computers and mobile devices [67], has resulted in the development of skills that were hitherto unavailable to the public, and these skills provide opportunities in the design of exhibits and installations that are intuitive and engaging. These decisions manifest in the form of replacing traditional computing peripherals (such as a keyboard and mouse) with game controllers (such as joysticks and game pads), and augmenting traditional output devices (such as screens) with immersive displays (such as headsets and domes). The on-site exhibit explores the idea of replicating physical artefacts using digital media, such that museum objects can be digitised using photogrammetry and scanning (as discussed in section 7.1). Multiple 3D artefacts can be arranged by theme or subject into 3D galleries and these can be embedded into virtual environments for interaction using game controllers. In addition to viewing digital objects in virtual environments, scenes (whether historic or present-day, real or fictional) can be depicted using spherical media to facilitate exploration of distinct vantage points. Furthermore, these digital objects and scenes can then be combined by embedding objects in virtual environments depicted by scenes. This can prove useful for providing context and interpretation; for instance, for viewing an object in its location of origin or the location where it was found. The exploration of scenes from distinct vantage points can be useful for task-based learning. However, free-form navigation may be beneficial for exploratory learning (see chapter 4), hence the inclusion of free-form 3D environments that users can navigate through in first- or third-person view can offer an additional dimension to the navigation of scenes from distinct vantage points.

8.2.1 Design Goals

The underpinning theory discussed in chapters 2 and 3, coupled with the exploratory work detailed in chapters 4 – 6 have provided insights into the museum functions, visitors' groups and needs. In order for the Virtual Museum Infrastructure (VMI) to effectively support the functions of traditional museums and meet users' needs, the following design goals have been identified:

Active Use and Management The discussion on museums in section 1.1.1 introduces a definition of and functions of museums, and this discussion leads to an introduction of virtual museums which is further discussed in section 2.3, also highlighting the varied nature of virtual museums in terms of features, components and use cases based on literature (section 2.5.4). Due to this variation, the view of virtual museums adopted in this work is less concerned with these features, components and use cases, but rather focuses on how the combination of these can be used to support the functions of traditional museums identified in section 1.1.1. To achieve this, a Virtual Museum Infrastructure (VMI) should be viewed as an active, live system; a system that is continually used and managed in view of the day to day operations of a museum, and should consequently support this continuous use for heritage curation, preservation and dissemination. To actualise this, the VMI design incorporates a presentation front end that enables the general public to consume heritage content which is held in the back end data store. This is improved upon by pairing the presentation front end with a management front end that heritage experts and community members alike can use to easily populate the data store with content, thus fostering active use and process reproducibility, which are important in the empowerment of museum staff and communities.

Content Reuse The case for the installation of virtual exhibits on-site, where resources are available is made in chapter 5. Where possible, the content of the on-site installation should be reusable, akin to the manner in which physical items in museum collections are catalogued and reused in multiple exhibitions. For instance, spherical media can be used in on-site installations, VR mobile apps and online virtual tours. In this manner, content that is created once can be deployed on multiple platforms thus making a VMI more effective for supporting the curation and dissemination functions of museums. To actualise this, the VMI design incorporates multiple presentation platforms that facilitate the dissemination of content held in the data store. For instance, data that is uploaded using the content management interface can be consumed on a web-based map interface which presents an up-to-date, instantaneous representation of the data store. In addition, the contents of the data store are backed up on local storage servers, hence data can be retrieved from these servers to be used for building immersive, on-site museum exhibitions and mobile apps. The data uploaded to the data store can also be pushed to social archive sites thus ensuring greater reach and dissemination. This allows data to be created once and reused in multiple scenarios and on multiple platforms.

Engaging Experiences The use of VR and immersive technologies for heritage exploration provides opportunities for high levels of engagement with heritage content. This is demonstrated in chapter 6 through the use of headsets, smartphones and drones for aerial first-person exploration of cultural landscapes, as well as the use of both low-end (chapter 4) and mid- to high-end (chapters 4 and 5) VR headsets for immersive experiences, both remotely and on-site. To actualise this, the VMI provides support for VR headsets while undertaking virtual tours or exploring 3D artefacts. This may take the form of affordable headsets (such as the Google Cardboard) for use with mobile apps, or wired headsets (such as the Oculus Rift) for use with immersive, on-site exhibits. The engaging experience is not lost in the absence of VR headsets, as users can still consume the aforementioned content on mobile and desktop screens.

Multimedia Integration Chapter 4 discusses the value in the remote exploration of heritage sites by deploying content in the classroom for learning activities, while chapter 5 discusses the value in the on-site exploration of heritage in a museum setting through an immersive museum installation. In addition, a case for combining media types such as audio, video, flat and spherical images is made while exploring heritage content remotely and on-site. The findings of chapters 4 and 5 demonstrate that the integration of different modes of exploration and multimedia content provides opportunities for use cases such as outreach in schools, group excursions and individual visits to museums. To actualise this, the VMI combines different media types to provide a rich and varied experience to users. Users can listen to oral histories (such as folk tales and ancient stories) or be guided around a site using audio narratives, users can undertake virtual tours of remote sites or compare the present and past states of a local site using spherical media, curators and conservators can create digital representations of artefacts using photogrammetry and scanning technologies, thus enabling users to explore and inspect these 3D artefacts remotely using their mobile devices, text snippets can be added by community members to describe entities using an integrated wiki, and the resulting wiki articles are available to consumers to add context and interpretation to their experiences.

Resource Maximisation Given the variation in the resource availability for museums, a design of virtual museum systems to maximise resource utilisation can increase its value and appeal to heritage practitioners. This design may entail delivering engaging experiences using affordable infrastructure where resources are limited and more advanced installations where resources are available. Chapter 4 demonstrates that low-end, mobile infrastructure can deliver engaging experiences that are comparable to mid- to high-end systems. While there is a need for affordable systems where resources may be scarce, there is also a need for interactive systems that leverage the computing power of mid- to high-end systems where resources are available as investigated in chapter 5. This provides the rationale for designing systems that cater to the varying resource capabilities of heritage organisations as illustrated by the resource spectrum (Fig. 1.2). To actualise this, the VMI design places an emphasis on optimal resource utilisation, which is important due to the varying levels of resources available to museums. This is manifested in the form of the use of low-end systems where resources are scarce and mid- to high-end systems where more resources are available. Furthermore, digital literacies are leveraged as a means of tapping into unused resources, such that museums can take advantage of their visitors' ownership of computing devices (such as smartphones) and proficiency with game technologies to deliver engaging experiences and maximise user satisfaction with the resources available.

The combination of web, mobile and immersive technologies in the VMI achieves the goal of providing engaging experiences to users, as it uses VR headsets to provide immersive virtual tours of remote locations and enables the inspection of and interaction with 3D artefacts. Furthermore, multiple media types – audio, video (flat, spherical and stereoscopic), images (flat, spherical and stereoscopic), 3D artefacts and text – are combined to provide these compelling and informative experiences to users, thus fulfilling the multimedia integration goal. By combining a Digital Asset Management System (DAMS) back end with a management front end, heritage practitioners can create and update content with ease, and a presentation front end provides instant feedback hence heritage practitioners and community members alike can continuously manage and use

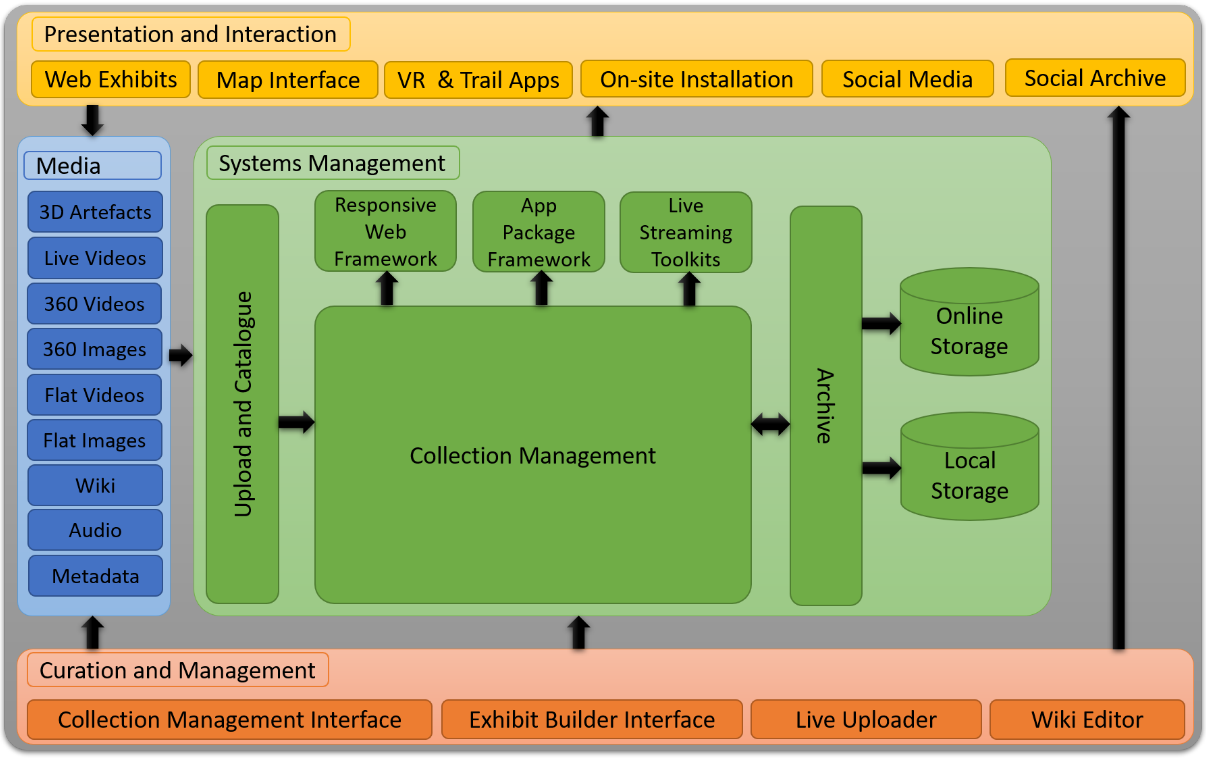

Figure 8.2: Virtual Museum Infrastructure architecture

the virtual museum. The architecture and relationship between entities in the VMI are discussed in the following sections.

8.2.2 Architecture

The architecture of the VMI is shown in Fig. 8.2. A web-based archive form enables heritage practitioners to upload files and supply metadata which, together represent and describe entities that are presented to users using multiple media types which collectively constitute the digital heritage representation mechanisms discussed in section 1.1.2. These include 3D artefacts, spherical images and videos, and traditional multimedia such as flat images, audio and text. 3D artefacts which are created using photogrametry and scanning can be used to depict or represent physical museum artefacts. Spherical videos and images can be used to depict landscapes and cityscapes as scenes. Traditional media such as flat images, videos, and audio can be used to capture both tangible and intangible heritage, and wikis can provide narratives for heritage interpretation. In addition to the aforementioned pre-packaged media, the VMI facilitates the presentation of live media through toolkits and software that are made available to heritage practitioners. For instance, the server and client apps developed in chapter 6 for streaming real-time videos can be downloaded and then configured by heritage practitioners, so that with

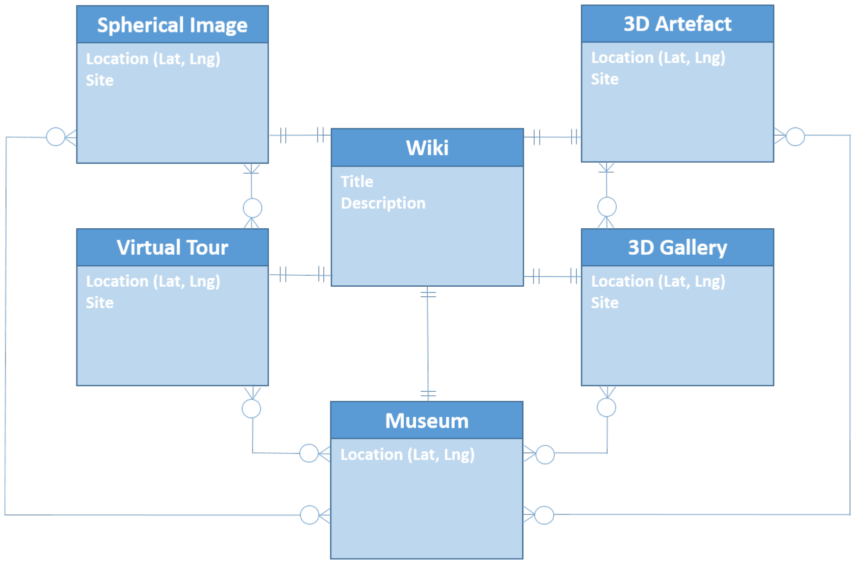

Figure 8.3: VMI entity relationship diagram

the use of a drone, they can facilitate virtual tours and live exploration of remote landscapes. The VMI facilitates content curation and management through an array of interfaces. A live uploader enables users to transfer digitised content to the data store. The content can be described using metadata, while narratives can be added and modified using a Wiki editor. Collections can be created and managed by thematically grouping uploaded items, while an exhibit builder facilitates the creation of exhibits that can be publicly available. The content that is created using the curation and management interfaces can be stored using both local and online storage solutions. Local storage facilitates access to resources in spite of Internet access while online storage facilitates interconnectivity and reach. The ability to ensure responsive delivery of resources is a desirable feature of the VMI design, which will enable the system to serve the appropriate fidelity of content to users depending on their operating platforms. The content stored in the system can be presented to users in multiple ways. Exhibits which feature images, videos, objects (and other media) can be presented on the web. A map-based interface can be used to geographically visualise museums, collections, exhibits and objects. VR-compatible mobile apps can be created to present already-existing content in the VMI storage as well as live content such as real-time videos streamed from drones; these mobile apps can then be disseminated on online app stores. Immersive exhibitions featuring VR technologies, spherical scenes, 3D objects and traditional media can be installed on-site in museums, and content can be disseminated traditionally through social media sites. The ability to upload and manage content using the live uploader and management interface respectively, and (re)use these content and resources in different ways emphasises the active nature of the VMI. Heritage practitioners and community members can continuously create and digitise content and once uploaded to

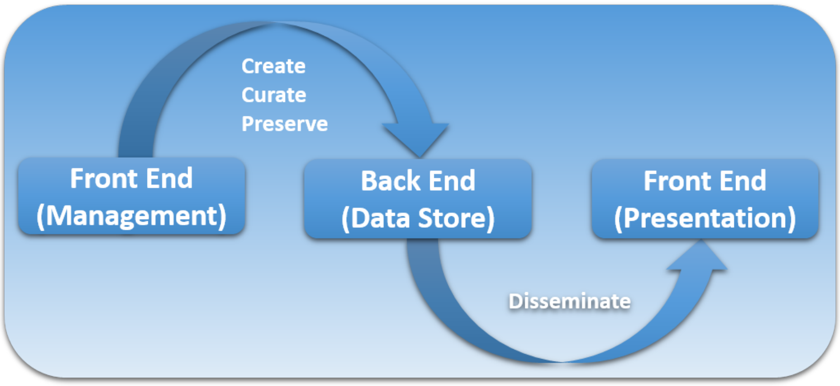

Figure 8.4: VMI components

the VMI, the same content can be used to build web exhibits and mobile apps, and can be accessible on a map interface and on social media. Metadata can be added to the content while it is being created/uploaded, or afterwards, and changes can be continuously made to the resource throughout its lifespan.

An entity relationship diagram which depicts how the content of the VMI fit together is shown in Fig. 8.3. Spherical images are used to depict locations to enable users undertake remote tours of places that they plan to visit (in anticipation), places that they have visited (in reminiscence or recollection) or places that they are unable to visit (in substitution). One or more spherical images, grouped together by a common theme (such as location or likeness) make a virtual tour, and each spherical image can belong to more than one virtual tour or none at all, hence this relationship is represented as a many to many association (one to many [spherical images] and zero to many [virtual tours]). Zero or more virtual tours may be associated with a museum (or other heritage organisation) which represents the top level entity of the VMI. This implies that each virtual tour may be associated with one or more museums or none at all hence this relationship is represented as a many to many association (zero to many [virtual tours] and zero to many [museums]). In addition, spherical images (whether they are part of virtual tours or otherwise) can also be associated with museums in a many to many manner (zero to many [spherical images] and zero to many [museums]).

3D artefacts are used to depict digital representations of physical artefacts or objects of interest. One or more of these 3D artefacts, grouped together (usually by location, provenance or composition) make a 3D gallery and each 3D artefact can belong to more than one 3D gallery or none at all hence this relationship is represented as a many to many association (one to many [3D artefacts] and zero to many [3D galleries]). Zero or more 3D galleries may be associated with a museum. This implies that each 3D gallery may be associated with one or more museums or none at all hence this relationship is represented as a many to many association (zero to many [3D galleries] and zero to many [museums]). In addition, 3D artefacts (whether they are part of 3D galleries or otherwise) can also be associated with museums in a many to many manner (zero to many [3D artefacts] and zero to many [museums]).

Each entity in the VMI has an associated wiki entry which comprehensively describes the entity with the text and metadata supplied using the archive form. Irrespective of type, each entity has one and only one wiki entry, however, the wiki entry of an entity may refer (or link) to the wiki entry of another entity (for example, a museum's wiki entry may refer to the wiki entry of a 3D gallery associated with the museum, or a 3D gallery's wiki entry may refer to the wiki entry of a 3D artefact associated with the 3D gallery). This implies a one to one relationship between an entity instance and a wiki entry (one and only one [entity instance] and one and only one [wiki entry]).

8.3 Implementation

The VMI has been implemented using a variety of technologies – web platforms, mobile platforms, VR headsets, and a Digital Asset Management System (DAMS). The VMI[1][2] features a back end which stores and organises data, a management front end which enables updating the data, a presentation front end which enables interacting with the data, and an integrated workflow by which the system components communicate. Fig. 8.4 shows the components of the VMI and the communication processes between them.

8.3.1 Back End

The VMI back end is implemented as a free and open-source Digital Asset Management System (DAMS) using Omeka; the Omeka installation is managed by the Open Virtual Worlds server administrator. The DAMS provides interoperability with popular web-based system and frameworks, and is supported by a Relational Database Management System (RDBMS) which stores the data, and a Representational State Transfer (ReST) Application Programmable Interface (API) which facilitates retrieving, modifying and adding to the data. Within the VMI, calls to the Omeka API are made using Javascript (from the online presentation front end) and Python (from the management front end). The code snippet in Listing 8.1 shows a Javascript function that can be used to retrieve the items in the Omeka repository using an AJAX call. Upon successful completion of the call, a *"json_response"* object is received, which

[1] Instantiation for the EU-LAC project: https://eu-lac.org
[2] Video summary: http://openvirtualworlds.org/cadam/eulac_virtual_museum.mp4

```
function load_data(api_url, api_key, element_type)
{
  //api_url e.g. "http://mydomain.org/omeka/api"
  //api_key e.g. "al0NgstR1ngwItHnumB3rsanDlett3rs"
  //element_type e.g. "items"
  $.ajax({
    url: api_url+'/'+element_type+'?key='+api_key,
    crossDomain: true,
    type: 'get',
    timeout: 6000,
    success: function(json_response)
    {
      parse_data(json_response, element_type);
    },
    error: function (xhr, ajaxOptions, thrownError)
    {
      console.log(thrownError);
    }
  });
}
```

Listing 8.1: API call to retrieve all items from the Omeka database

contains the result of the call in JSON format, and is then parsed as required by the *"parse_data"* function. If the AJAX call fails or time outs, an error is thrown instead. Data management and metadata handling are important features of cataloguing and archiving systems. The Dublin Core schema [316] was proposed as a solution to the problem of little interoperability between the vocabularies of cataloguing and archiving systems. The Europeana Data Model (EDM) [321] builds on the Dublin Core schema and is widely adopted across (Europe) hence it boasts a high level of familiarity with heritage practitioners. For this reason, information provided using the management front end is mapped to Europeana types (which are in turn described based on Dublin Core terms) so as to utilise the existing literacies and improve interoperability with existing systems.

8.3.2 Online Front End

The online front end is implemented in two facets: a management front end which is implemented as a web-based form (shown in Fig. 8.5) which contributors use to create and manage data, and a presentation front end which is implemented as a map-based interface (shown in Fig. 8.6). The map interface features icon-coded pins which represent entities in the VMI, such that users can click on a pin to reveal a popup (shown in Fig. 8.7) for viewing or interacting with the item. The

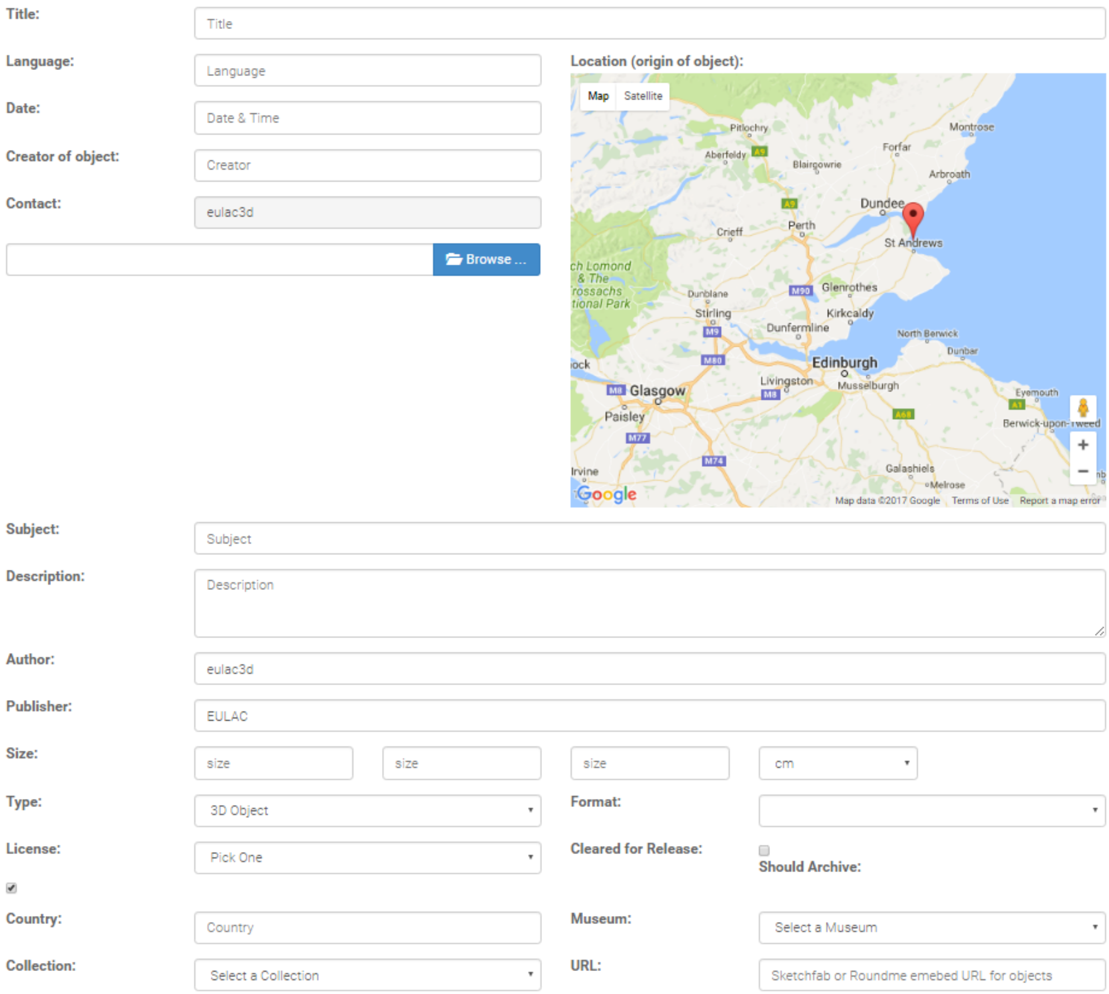

Figure 8.5: VMI management interface for creating and updating content

logic is implemented in Javascript and uses "GET" Asynchronous Javascript and XML (AJAX) requests to the Omeka API to retrieve data in Javascript Object Notation (JSON) format. The JSON data is then parsed to make sense of the data, which involves identifying distinct entities, categorising them by type and applying type-specific rendering. It is pertinent to note that the rendering of items depends on the information available. For instance, items without any location specific information may not be rendered on the map because of the requirement of latitude and longitude coordinates.

The map-based interface facilitates the *geographical visualisation* of data, as it represents entities on an interactive map of the world which enables the visualisation of the spatial and geographical relationships between entities. It also enables users to view entities by region. For example, users can view entities associated with the Caribbean only, or entities only found in Europe, or Latin America only, or a combination of these regions (recall that these regions are associated

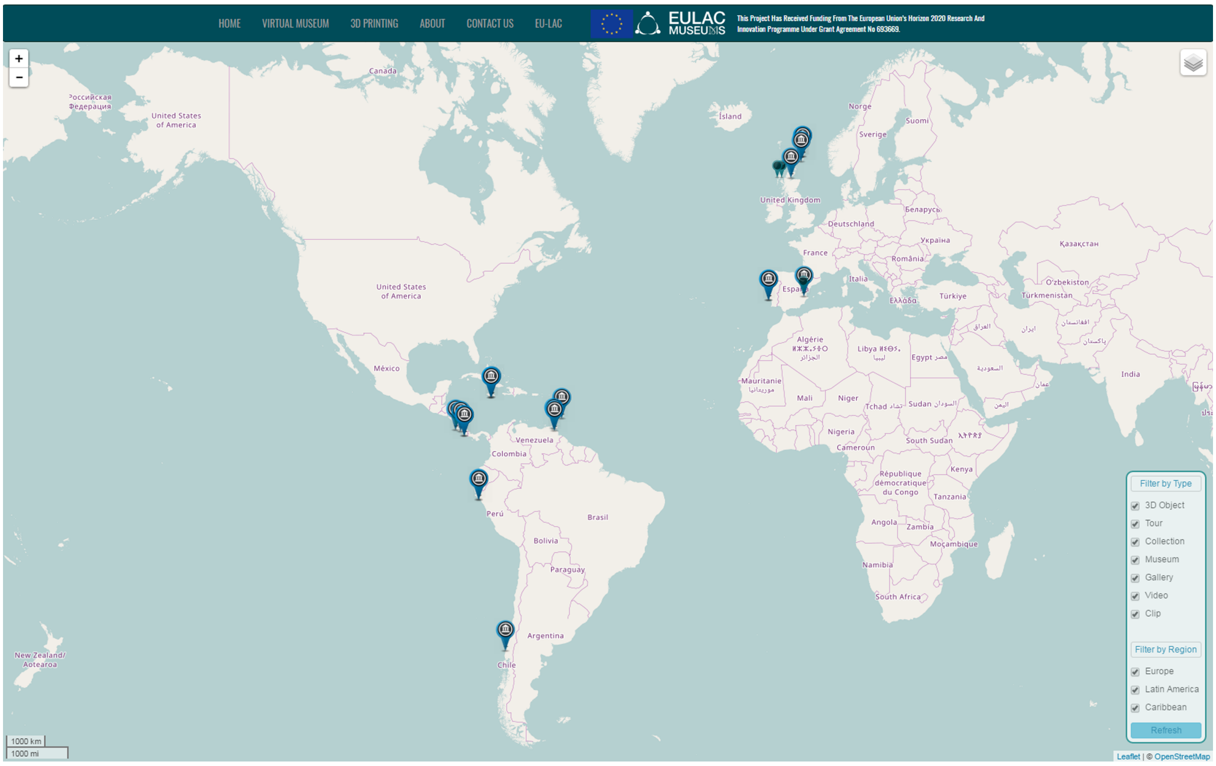

Figure 8.6: VMI map interface with options for viewing entities by location and type

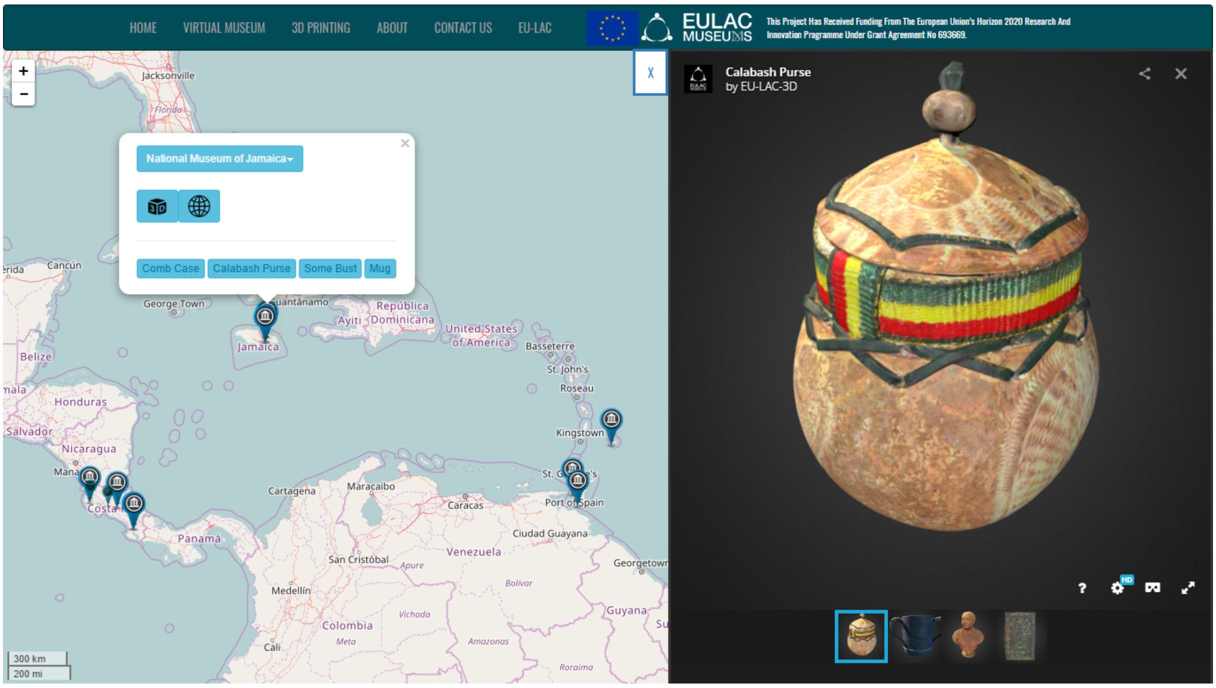

Figure 8.7: Viewing a 3D artefact in the VMI map interface

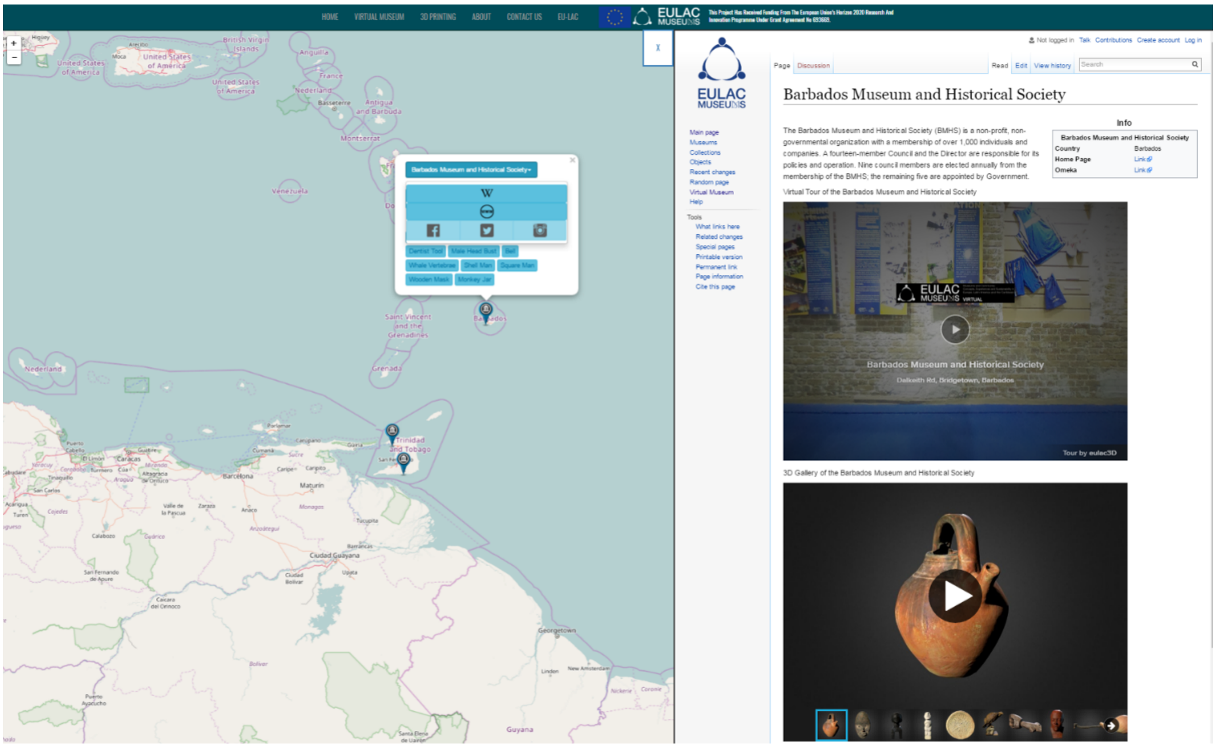

Figure 8.8: Viewing wiki content in the VMI map interface

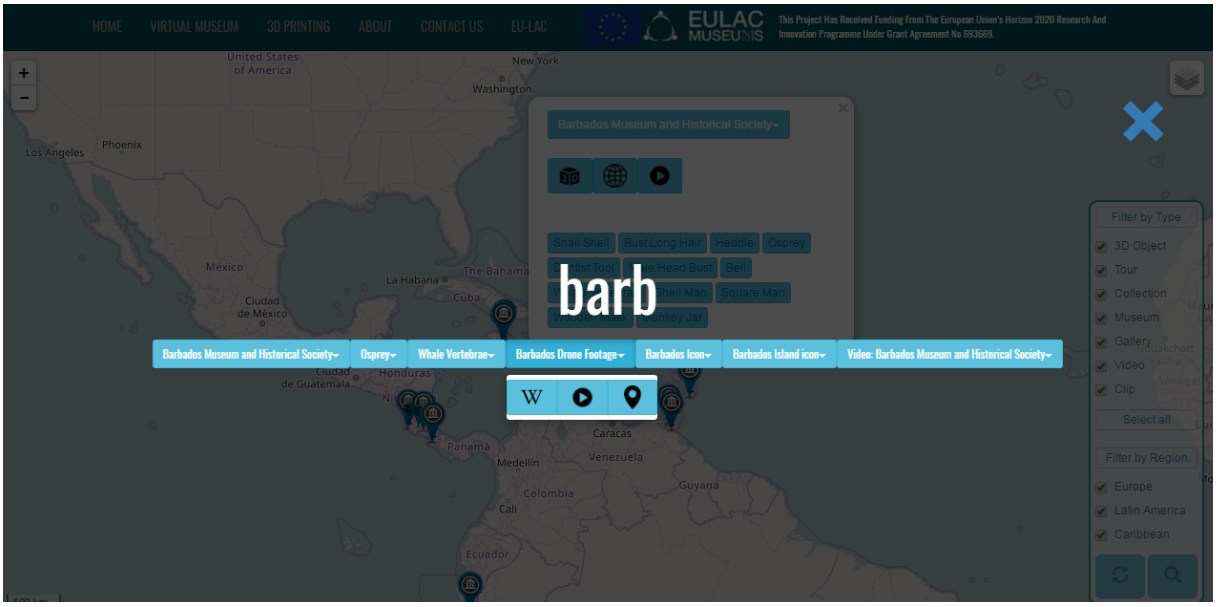

Figure 8.9: Searching for entities in the VMI map interface

with the EU-LAC project). This selective geographical visualisation is implemented using a bounding box model. For each identified region, four pairs of latitude and longitude coordinates are used to represent the four corners of an imaginary box. Users can select or deselect regions

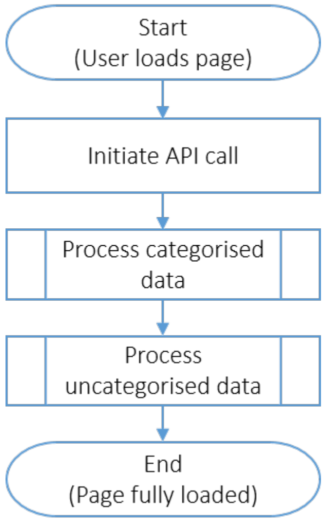

Figure 8.10: Condensed flow chart showing the operation of the map interface

using checkboxes on the map interface. When data is requested from the back end and before the entities are rendered on the map interface, an algorithm checks if each entity's coordinates fit into the bounding box defined by any selected regions. An entity will only be rendered on the map if it is found to be within a selected region. In addition to a region selector, the map interface features a panel that enables users to view content by type by selecting or deselecting checkboxes. This facilitates the *categorical visualisation* of data, such that users can choose to view only museums, 3D artefacts, tours, images, or any of the other entity types available, and any combination of these types (see the bottom right corner of Fig. 8.6). In addition to the geographical and categorical visualisation of content, the VMI supports *descriptive visualisation* of content using a wiki which collates the metadata that has been provided for each entity using the management interface. The wiki is displayed in a panel alongside the map interface (shown in Fig. 8.8) and it features any descriptive text and data that has been provided, as well as any associated multimedia in the form of 3D artefacts, virtual tours, images or video.

A fourth type of visualisation supported by the VMI is *exploratory visualisation*, which is facilitated using an instant search feature that enables users to query the data store for entities that match a given search string. The search results are updated after every key press that modifies the query string, and is performed on the title and description fields of entities. The instant search (see Fig. 8.9) is facilitated by locally storing representations of the entities contained in the data store so as to preclude the need to repeatedly access the server while searching for entities. The *exploratory* search feature also integrates with the other visualisation paradigms in the system so that users can visualise the results *geographically* by viewing an entity's location on the map, *descriptively* by viewing the wiki content for the entity and *categorically* by viewing type-specific

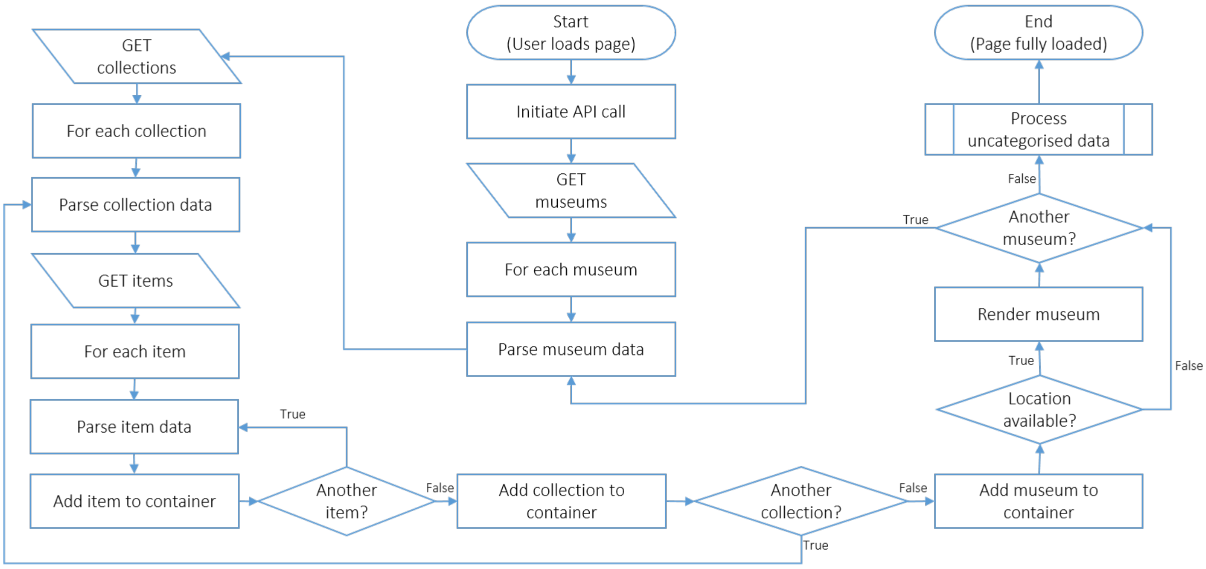

Figure 8.11: Flow chart showing the processing of categorised data on the map interface

information for the entity. These features work together to facilitate the visualisation of heritage content in four ways: *geographical*, *categorical*, *descriptive* and *exploratory*, thus providing breadth and depth in the dissemination of, and engagement with heritage content, as well as satisfying the information-seeking needs of virtual museum users [249, 250].

A condensed flowchart showing the operation of the web-based map interface at a glance is shown in Fig. 8.10. When the user attempts to load the web page by either clicking on a link to the URL or typing the URL into a web browser, the web page makes an Asynchronous Javascript and XML (AJAX) request to the server. Upon completion of this request, a Javascript Object Notation (JSON) representation of the entirety of the data store is returned. This JSON data is then processed in two stages; the categorised data is processed (see Fig. 8.11), followed by the processing of the uncategorised data (see Fig. 8.12). “Categorised data” in this context, refers to entities that are associated with other entities to form a hierarchical structure. For instance, an artefact may be part of a collection and that collection may be one of many that are associated with a museum in the system. Conversely, “uncategorised data” are made up of entities that are not explicitly associated with other entities in the system. For instance, little may be known about the provenance or origin of an artefact, hence it may not be part of a museum collection. There may also be collections of objects that are not associated with a museum or heritage organisation. “Processing” in this context refers to the parsing and rendering of entities on the map interface so that users can interact with them. The web page is fully loaded once the categorised and uncategorised data are processed, and all the data remain available for searching, exploring and interaction until the next AJAX request is made.

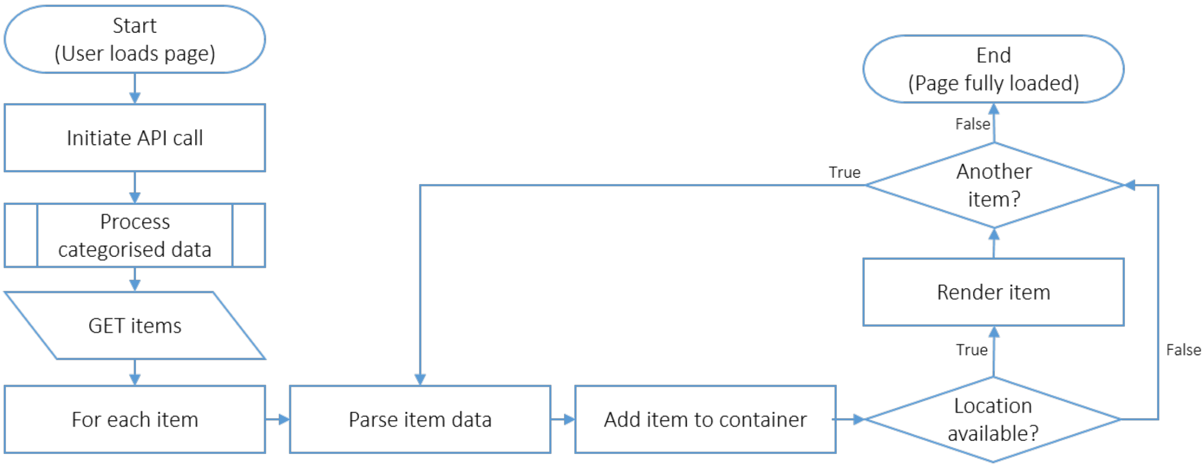

Figure 8.12: Flow chart showing the processing of uncategorised data on the map interface

The steps taken to process the categorised data are shown in detail in Fig. 8.11. Once the JSON representation of the data store is obtained, the top level entities which represent museums and heritage organisations are retrieved. The museums are iterated over so that for each one, each collection that is associated with it is obtained. The collections are then parsed and iterated over, so that for each collection, each item that is associated with it is retrieved. For each retrieved item, the data fields are extracted to obtain information such as its title, description, type, provenance, location of origin and so on. Each item, along with its associated data, is added to a container which represents the collection it belongs to, and similarly, each collection is added to a container which represents the museum it is associated with. Each museum that has location data (in the form of latitude/longitude coordinates), is then represented by a marker on the map interface, which when clicked, reveals a drop-down menu that is comprised of collections and items (see Figs. 8.6 – 8.9. At this point, the parsed data is held in the device's local storage so that it can be accessible even after connectivity is lost. The data is refreshed only when a new AJAX request is made to the server. After processing the categorised data, the uncategorised data is processed using a similar (but shorter) procedure, as shown in Fig. 8.12. Unlike the categorised data which have a hierarchical structure, the uncategorised data have a flat structure which are made up of items that are neither associated with museums nor collections in the system. Each item that is retrieved is parsed so that its data fields can be extracted, and then rendered as markers on the map if there is associated location data (in the form of latitude/longitude coordinates).

The operation of the management interface is depicted as a flow chart in Fig. 8.13. The process begins with a user creating content. This could be in the form of 2D images of a location or 3D images of an artefact. Content is created by providing files and describing the files by filling a form which contain fields such as a title, description, date, location, data type, and so on. The

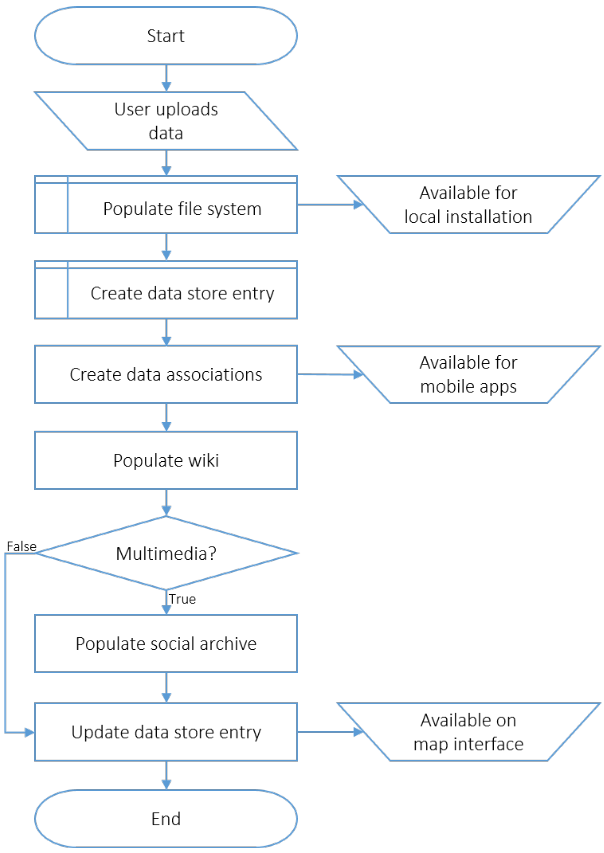

Figure 8.13: Flow chart showing the operation of the management interface

uploaded data is transferred to the file system and an entry with the appropriate data associations is created for the item in the data store. Once the data store entry exists, the contents of the filled form are used to create a wiki page and populate it with information. This is important so that the data (and in some cases the links to data) associated with a given item can be found in one place. If additional media files are provided with the content, the media files can be pushed to social archive sites such as Sketchfab [133] for 3D models, and Roundme [322] for spherical images, and the data store entries will be updated to contain pointers to the pushed content via URLs. The flow chart in Fig. 8.13 also shows three manual operations. Once data is uploaded using the management interface and transferred to the server, the content can be retrieved directly from the file system and used to manually set up a local exhibit or installation. Furthermore, once the data store entry exists and the associations have been made, an exhibit builder interface can be used to organise the content to make mobile app packages. Lastly, once the data store entry has been updated to contain pointers to any external sources, the data becomes available to users by loading the web-based map interface (see Figs. 8.10 and 8.11).

A collaboration diagram which depicts communication between system components is shown

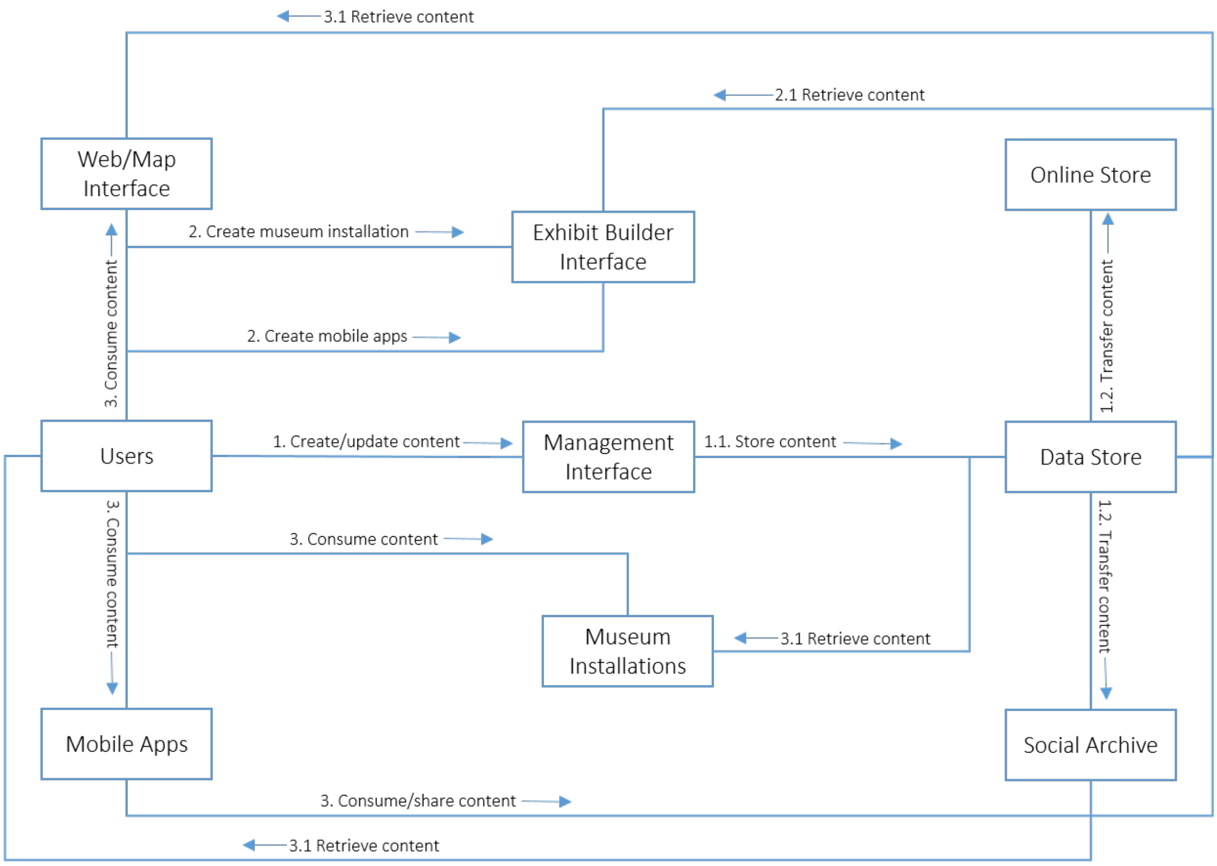

Figure 8.14: Collaboration diagram showing communication between system components

in Fig. 8.14. For simplicity, only the main system components are shown: *Users*, *Management Interface*, *Mobile Apps*, *Web Interface*, *Museum Installations*, *Exhibit Builder Interface*, *Social Archive*, *Data Store* and *Online Store*. Interaction typically begins with *Users* creating content using the *Management Interface*. The *Management Interface* is implemented as a web-based form that enables *Users* to upload and modify content, which is stored in the *Data Store*. Once in the *Data Store*, the data can be pushed to supported *Social Archive* platforms and backed up on an *Online Store*. The data that exists in the *Data Store* can be used to make exhibits either as VR and trail *Mobile Apps* or *Museum Installations*. The content created and curated using the *Management Interface* and the *Exhibit Builder Interface* can be consumed using the map-based *Web Interface*, *Mobile Apps*, *Museum Installations* and *Social Archive* sites, each of which draw content from the *Data Store*. The collaborative diagram in Fig. 8.14 shows the active nature of the VMI which supports not only the collaborative creation and management of content, but the presentation and reuse of this content in different ways such as web-based exhibits, mobile apps, museum installations and social media.

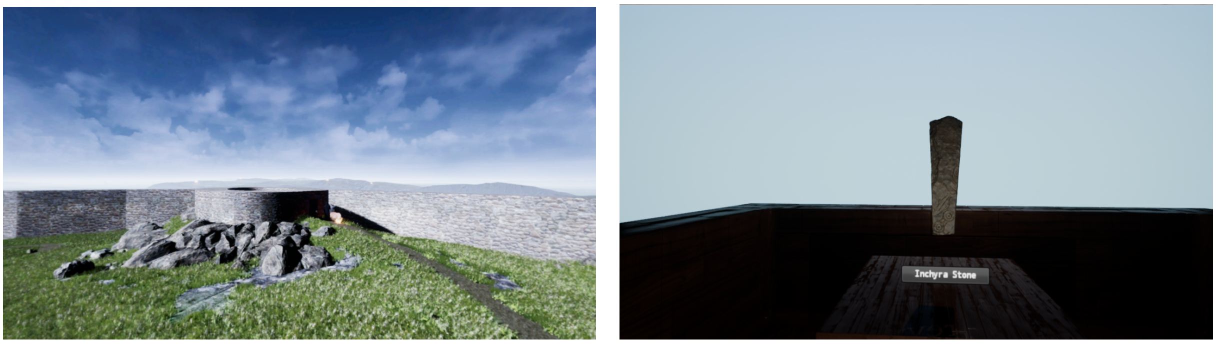

Figure 8.15: A reconstruction of an Iron-age fort (left) and a simplistic 3D object gallery (right)

8.3.3 On-site Exhibits

On-site exhibits, in this context, refer to the use of interactive, digital media for participatory heritage engagement in physical museum spaces. Chapter 5 introduced an on-site virtual museum installation which uses spherical media – extracted from a free-form 3D environment – to facilitate a comparison of the past and present from equivalent vantage points. The success of the exhibit as reported by both heritage practitioners and museum visitors (see section 5.5 makes a case for the adoption of digital means of heritage interpretation on-site, where resources are available. While free-form 3D environments can be built in virtual world simulators such as OpenSim[3] and SecondLife[4], there is a trend towards using free game engines such as Unity3D[5] and Unreal Engine 4[6]. Unreal Engine 4 is freely-available hence there are no upfront procurement or subscription costs; it has in-built support for integrating VR devices into 3D environments, as well as a proven track record in the creation of high-fidelity 3D environments as exemplified by recent versions of classic games such as Street Fighter V and Tekken 7 [323]. The free-form 3D environment (see Fig. 8.15) of an on-site installation is implemented in Unreal Engine 4 using the procedure discussed in section 5.4. In addition, spherical media can be extracted from these 3D environments using a similar procedure, and these can be used to make up the (spherical) scene components of an on-site exhibit, web-based virtual tour, or mobile VR app. Users can switch between an array of virtual environments either using game controllers (e.g. the Xbox controller) or by triggering hotspots using head movements, and these provide the illusion of moving from one place to another, similar to a virtual tour. Digital representations of objects can be imported as 3D meshes. These meshes are 3D artefacts that have been digitised using photogrammetry or scanning (as discussed in section 7.1).

[3]Online: http://opensimulator.org/wiki/Main_Page (Accessed: 2018-03-09).
[4]Online: http://secondlife.com/ (Accessed: 2018-03-09).
[5]Online: https://unity3d.com/ (Accessed: 2018-03-09).
[6]Online: https://www.unrealengine.com (Accessed: 2018-03-09).

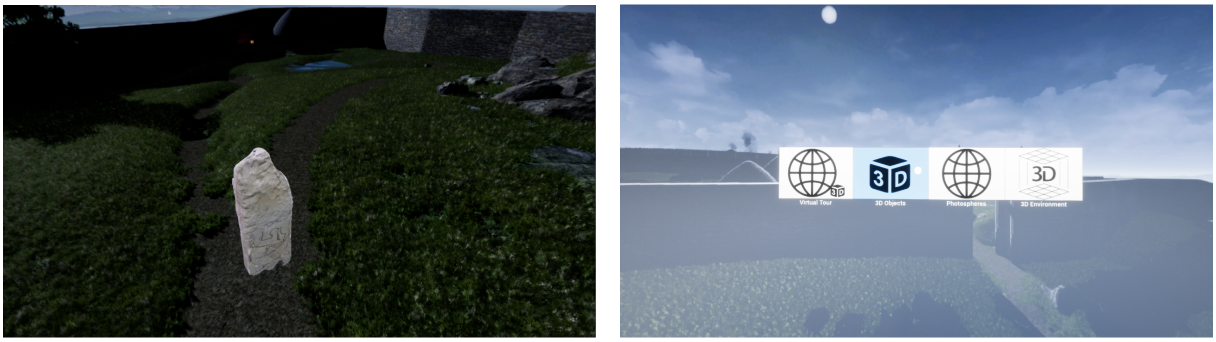

Figure 8.16: Contextualising objects in scenes (left) and a sample virtual exhibit menu (right)

A gallery-like 3D environment is used to house the digital artefacts. The 3D gallery is styled to place less emphasis on the level of detail in the virtual environment, and more emphasis on the 3D artefacts (see Fig. 8.15). Controller-based interaction is built into the 3D gallery so that users can rotate the artefacts around multiple axes, zoom in/out and move about the centrepiece. Users can also switch between an array of 3D artefacts in a manner that is similar to navigating through images in a photo gallery. In addition to representing scenes using spherical images and physical objects using 3D artefacts, the on-site exhibit allows for placing 3D artefacts in environments represented by spherical scenes (Fig. 8.16). This is useful for associating scenes and objects thereby providing context and interpretation. For instance, an object may be placed in a scene that represents its location of origin or find.

These components – spherical scenes, 3D object galleries, combined scenes and objects, and free-form 3D environments – and their underlying have been created as templates which are freely available online to facilitate reuse and accessibility. Thus, community members and heritage practitioners with little or no experience with 3D environments can download the template packs, "drag and drop" assets and functionality templates into game engine levels, populate these levels with content (such as spherical images, videos and 3D artefacts) and package them as immersive, museum installations with top-level menus that users can interact with (see Fig. 8.16). Spherical scenes provide the ability to represent landscapes, cityscapes and geological structures. 3D artefacts provide the ability to represent artefacts, specimens and sculptures. Spherical scenes and 3D artefacts can be combined to provide context and association between them. Audio and text can also be combined with these media to provide additional information and alternative means of content delivery (see section 1.1.3). Hence on-site exhibits add an addition dimension in the digital representation and interpretation of heritage.

S/N	Interaction	Description	Pros	Cons
1	Multi-button hand-held	Hold analog stick to move, release to stop	✓ Tactile feedback ✓ Multiple alternatives (buttons) for interaction	× Requires multi-button controller, and game controller proficiency × Occupies two hands
2	Single-button hand-held	Hold button to move, release to stop	✓ Tactile feedback ✓ Does not require game controller proficiency	× Requires single-button controller × Occupies one hand
3	Controller-free	Nod (or lean forward) to move, shake (or lean back) to stop	✓ Intuitive ✓ Hands-free	× No tactile feedback × May introduce fatigue

Table 8.2: Free-form locomotion in virtual environments

8.3.4 Interaction

The use of both controller-based and controller-free interaction techniques for navigating virtual environments have been investigated in chapters 4 and 5 respectively and a case can be made for the implementation of both paradigms in on-site instantiations of the VMI. Controller-based interaction taps into digital literacies, specifically games proficiencies and familiarities. This is useful because a significant portion of the population play video games and these games offer benefits for learning and social interaction [67, 324, 325, 326], hence by "gamifying" exhibits, museums can deliver interactive experiences to visitors using familiar technologies and devices [327, 328, 329]. That said, the merits of controller-free interaction lie in its simplicity and attractiveness to technology-averse museum visitors as well as those who are not proficient with gaming devices. The efficacy of this type of interaction has been demonstrated in chapter 5, in line with Kostoska et al. [199] who suggest that an interaction-free design may translate to an improved user experience for older adults. The use of controllers for interacting with games has been studied extensively in literature [330, 331, 332, 333, 334, 335]. While the choice of hand-held controllers can significantly impact player experience in games [330], the findings of Freeman et. al. [332] suggest that experiences can be designed to use hands-free interaction techniques (such as motion-based controllers) with little or no impact on games outcome. For these reasons, both controller-based and controller-free interaction paradigms have been designed, implemented and analysed as follows.

Multi-button hand-held This is a controller-based technique that uses game controllers for interaction. The buttons are mapped to multiple actions that users can perform while exploring a virtual environment.

Single-button hand-held This is a controller-based technique that uses point-and-click devices (such as the Oculus remote) for interaction. The single-button can be mapped to multiple actions by varying the number and/or duration of clicks.

Controller-free This is a hands-free technique that works by polling the sensors of VR systems to detect movements (such as nodding, shaking or leaning) which can be designated for interaction tasks.

Table 8.2 shows the characteristics of these interaction techniques, which makes them suited for different purposes. The availability of buttons on multi-button game controllers provides opportunities for performing multiple tasks and they provide tactile (touch-based) feedback [332]; however, their use is limited where game proficiency is lacking. Single-button controllers offer simplicity in interaction at the expense of multiple buttons for performing different tasks. Controller-free interaction techniques offer hands-free possibilities which may prove convenient but may sacrifice the tactile feedback provided by hand-held devices. The characteristics of these techniques can also be discussed in terms of comfort – the ease and intuitiveness of interaction – and performance – the potential for the interaction gesture to result in the desired outcome, as a body of work focus on evaluating these aspects for locomotion in virtual environments [336, 337, 338]. Repeatedly leaning one's body or shaking one's head may introduce fatigue which reduces comfort of controller-free techniques. Similarly, repeatedly changing grip when holding multi-button controllers may reduce comfort, especially if the controllers are too large to comfortably sit in users' hands. Llorach, Evans and Blat [338] report that a controller-free technique (using position sensing) produced less simulation sickness than a controller-based technique and may thus provide more comfort when navigating virtual environments, while, Cardoso [337], in slight contradiction, suggests that controllers (gamepads) are more comfortable. Both studies, however, conclude that controller-based techniques enable users to perform tasks faster than controller-free techniques [337, 338]. High performance enables museum visitors to engage with, and consume content in a timely manner and with little or no frustration. The ability of controller-based techniques to produce high performance may be due to their characteristics which translate to high accuracy in producing the required gesture. However, this high performance is typically observed when users possess the game proficiencies needed to operate controllers without looking at them (as is required when using VR headsets). Conversely, controller-free techniques may exhibit relatively lower performance due to the sensitivity of sensors and variations in body movements [337, 338]. However, they preclude the need for controller proficiency on the part of users. Therefore, both controller-free and controller-based interaction have been implemented as components of the VMI on-site exhibits, thus enabling museums to select and deploy interaction mechanisms that are appropriate for their target audience and project objectives, as controller-based techniques can be used to

S/N	Museum	Country	Attendees	3D Artefacts	Virtual Tours
1	Moruga Museum	Trinidad	52	3	Yes
2	CharlesTown Maroons Museum	Jamaica	16	3	Yes
3	Curre Museum	Costa Rica	22	5	Yes
4	Community Museum of Boruca Indians	Costa Rica	23	6	Yes
5	San Vicente Eco Museum	Costa Rica	6	9	Yes
6	Barbados National Museum and Historical Society	Barbados	14	13	Yes
7	National Museum and Art Gallery of Trinidad and Tobago	Trinidad	16	11	Yes
8	National Museum of Jamaica	Jamaica	3	5	Yes
9	University of West Indies Museum	Jamaica	8	6	Yes
10	Seixal Eco Museum	Portugal	16	9	Yes
11	National Archaeology Museum Lisbon	Portugal	42	7	Yes
12	Shetland Museum and Archive	Scotland	5	10	Yes
13	Unst Heritage Centre	Scotland	8	6	Yes
14	Unst Boat Haven	Scotland	18	1	Yes
15	Museo Comarcal de l' Horta Sud	Spain	8	8	Yes
16	Community Museum of Malalhue	Chile	15	11	Yes
17	Sican Museum	Peru	14	15	Yes
18	Tucume Museum	Peru	5	17	Yes
19	Museum of the Universidad Austral de Chile	Chile	14	14	Yes
20	Timespan Museum and Art Gallery	Scotland	8	14	Yes
21	Cortes de Pallas	Spain	15	N/A	Yes
22	Museo del Palmito d'Aldaio	Spain	N/A	N/A	Yes
23	Taigh Cheasabhagh Museum and Art Gallery	Scotland	12	4	Yes
	Total		**340**	**177**	

Table 8.3: Summary of museum workshops output

deliver interactive experiences to game-literate visitors while controller-free techniques can be used to deliver snackable content and foster passive interaction for technology-averse visitors.

8.4 Deployment and Evaluation – Phase 1

The VMI has been deployed in over 20 museums worldwide as part of the EU-LAC project and evaluated in terms of its value and impact with the workshop participants. Table 8.3 provides an overview of the workshops, highlighting the museums, their locations and the number of attendees for each workshop. Table 8.3 also shows that over 300 3D artefacts were created during the workshops and for each museum, virtual tours were created using spherical images captured during the workshops. These 3D artefacts and virtual tours were also uploaded to the data store using the content management interface (see Fig. 8.5) and subsequently presented using the front-end interface (see Figs. 8.6 – 8.9). Content that was uploaded using the management interface is also pushed to social archive sites depending on their type. For instance, 3D artefacts are available on Sketchfab[7] and spherical virtual tours are available on Roundme[8]. The following are excerpts[9] of testimonials gathered from museums in Latin America.

Community Museum of Malalhue *"I think it was a workshop of great significance, which looked to the future, to growth, and acted as window to the world. The workshop made possible the acquisition of new techniques for cataloguing and exhibiting our objects. The workshop ... engaged effectively with the attendees – creating new links, cementing trust, and delivering valuable experiences associated with the historic context of the artefacts, and their value to the wider community. The participation of the community in the workshop was successful. People attended from both local institutions and from Lanco...All attendees really valued the training, regarding it as a great step forwards in 3D technology."* – Nerys Mora, Museum Director.

Municipality of Lanco *"The workshop provided a globalised view of democratising new virtual tools, to foster learning and the exchange of knowledge about other cultures. These virtual museums potentially enhance understanding of the value of our artefacts. This is especially important for a museum such as Malalhue, which was set up by the local community. The workshop provided me with an understanding of different ways of depicting objects, and the techniques for producing a digital 3D artefact. It should be noted that the acquisition of the relevant equipment by the various museums is of great importance, to enable curators to maintain their skills, contribute to the new virtual museums, and place us in the vanguard of the latest curatorial practice."* – Diego Lira, Head of Culture.

Universidad Austral de Chile *"This tool is relevant for the management and dissemination of museum collections of different size and administrative structure. The technological applications proposed will contribute to improve the conditions of collections and museum spaces for the activities of dissemination and heritage education with children and young people who use the web."* – Marcelo Godoy, Anthropologist.

[7] https://sketchfab.com/eu-lac-3D/collections
[8] https://roundme.com/@eulac3d
[9] These excerpts were originally provided in Spanish and have been translated to English.

Universidad Austral de Chile *"The activity contributed to the knowledge of new technologies related to the documentation of objects and registration of museographic spaces to generate virtual tours. It was significant to see and practice working with these technologies, to familiarise ourselves with the technical language, as well as to see the examples of other community or small-scale museums that have implemented these technologies. It is surprising to know that the implementation of small changes in the way we use devices that we already know could imply great changes in the way we disseminate work, collections and museums to society and the communities in which they are inserted."* – Simon Urbina, Head of the Archaeology Laboratory.

These testimonials highlight the impact of the VMI on the museums and their respective communities, as well as the perceived value of the system for heritage management and dissemination The testimonials specifically demonstrate how: (1) the workshops have empowered museum staff and community members by making digitisation technologies accessible thus providing an infrastructure to facilitate the creation, curation and dissemination of community heritage, (2) the workshops have impacted a variety of heritage organisations, ranging from the community level (Community Museum of Malalhue) to the provincial level (Municipality of Lanco) to a larger university museum (Universidad Austral de Chile), (3) the workshops encourage active and continuous digitisation of heritage to further hone the participants' newly-acquired skills, and (4) by virtue of the workshops, ties within and between communities have been strengthened, as the workshops have fostered a strong sense of communities' heritage through digitisation and curation, while also enabling communities to disseminate their heritage and explore the heritage content disseminated by other communities.

8.5 Deployment and Evaluation – Phase 2

An instantiation of the on-site exhibit was implemented for the Picts & Pixels exhibition at the Perth Museum and Art Gallery, Perth, Scotland. The exhibition explored the combination of traditional and digital heritage interpretation in museums by combining VR headsets, game controllers, screens, projectors and physical objects strategically laid out in the exhibition space (see Fig. 8.17). The VR headset, screens and projectors facilitated the delivery of content (as output devices), the game controllers facilitated interaction with content (for instance navigating between scenes and rotating 3D objects) and the traditional museum artefacts remained fixed in physical space, with text panels for interpretation. This was done to combine scenes (spherical images) and objects (digital representation of stones), and it leverages digital literacies (game

Figure 8.17: Passive interaction (left) and active interaction (right)

controllers and computing devices) to deploy virtual museum content in a physical museum space alongside traditional objects.

Over a period of three (3) months, exhibit visitors were invited to provide feedback regarding their level of satisfaction with, and perceived value for money obtained from the exhibit. The feedback was gathered using questionnaires on iPads placed in the exhibition space, optimally positioned so that they were filled on the way out, after visitors had interacted with the exhibits and the data was collated and analysed in spreadsheets to observe the public perception of the exhibit at a glance. Two hundred and nine (209) respondents provided data during this period. Furthermore, an evaluation exercise of the on-site exhibit was conducted during an open-day event, on the closing weekend of the Picts & Pixels exhibition at the Perth Museum and Art Gallery. This took the form of an in-the-wild evaluation, where data was gathered as museum visitors interacted with the exhibits to investigate usability, value and impact. Because the closing weekend was expected to attract hundreds of visitors, steps were taken to select a subset of these visitors from which data would be gathered. The first point of elimination took place by observing museum visitors as they walked through the exhibit door. The decision as to whether to approach a visitor to take part in the study was made based on their apparent willingness to spend time in the exhibit; in other words, the visitors that did not appear to be in a hurry or to be merely passing through were considered as potential participants. These visitors were then observed more closely while they made their way through the exhibit, and the ones that appeared to have engaged with the system for a minimum duration of five (5) minutes were approached when it was apparent that their interaction had come to an end. They were then asked if they had more

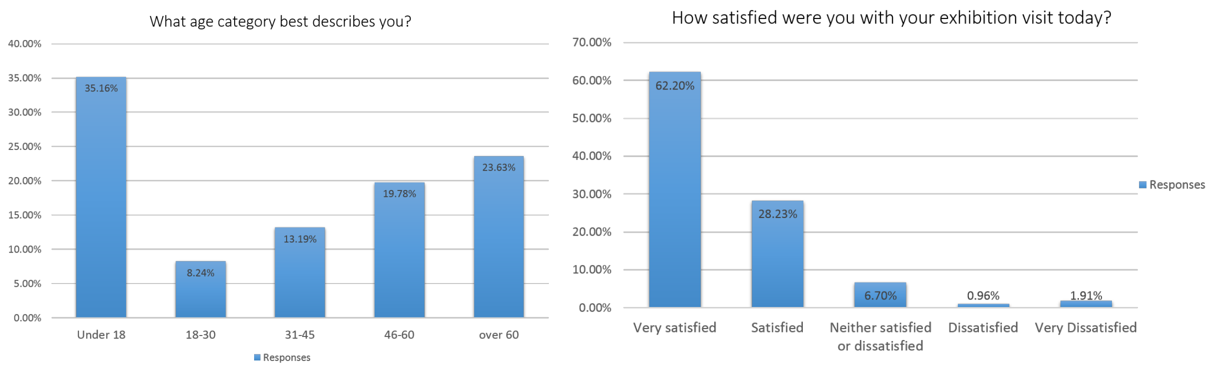

Figure 8.18: Age distribution (left) and visitor satisfaction (right) for the Picts & Pixels exhibit

time to spare to provide feedback. Consenting visitors were then engaged in a semi-structured interview while they filled out a Likert-scale questionnaire (see Appendix Table B.6). During the open-day exhibition, five hundred and four (504) museum visitors were recorded between 10AM and 4PM, and of these, quantitative data was gathered from twenty-two (22) of them. In addition to gathering quantitative data using Likert-scale questionnaires and qualitative data from the semi-structured interviews, the museum visitors were observed in the exhibition space and detailed notes were taken to supplement and make sense of the quantitative data. No video footage was recorded so as to mitigate the observer effect (a phenomenon which may cause participants to talk, act or behave differently due to the presence of recording equipment and/or obtrusive observers), the unsavoury perception that may be associated with perceived observation or surveillance in public spaces, and the ethical aspects associated with video recording such as participant identification through images and names [237, 238, 239, 240]. The findings of the evaluation exercise at the Perth Museum and Art Gallery are discussed as follows, with emphasis on the heritage value (and impact), usability and interaction.

8.5.1 Value and Impact

Fig. 8.18 shows the distribution of respondents age groups. Of the 209 visitors that responded to the survey, 182 provided information about their age group. The under 18 visitors represented the largest age group with 35.16% of the total respondents. Visitors aged over 60 represented the second largest group with 23.63% while visitors aged 46-60 represented the third largest group with 19.76% (see Fig. 8.18). 62.20% of respondents were very satisfied with the exhibition, 28.23% reported being satisfied, while less than 3% of the respondents were dissatisfied with the exhibition as shown in Fig. 8.18. Furthermore, of the 209 respondents, 82.30% felt that the exhibition provided value for money.

Fig. 8.19 shows the visitors' satisfaction levels across age groups. The chart shows that the

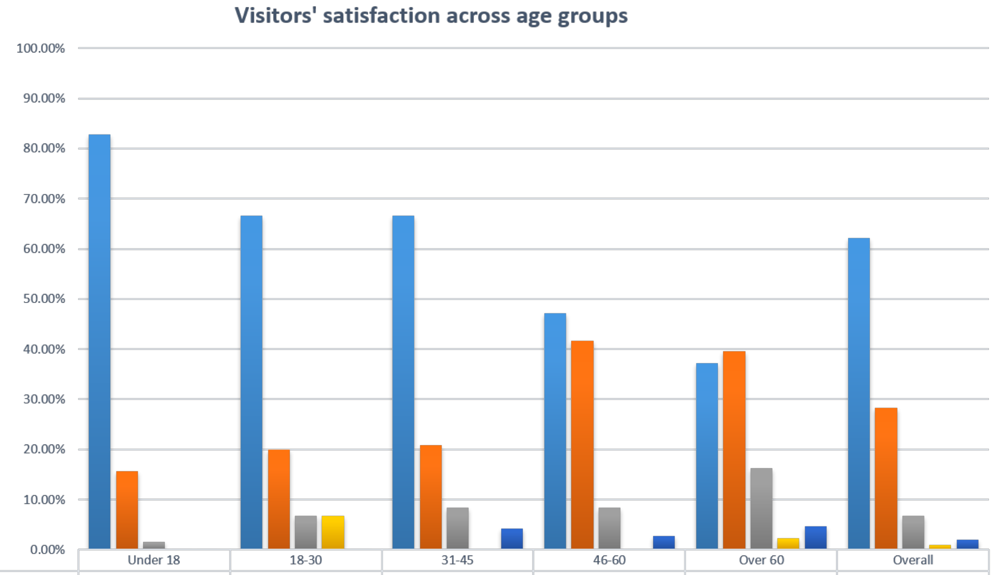

	Under 18	18-30	31-45	46-60	Over 60	Overall
Very satisfied	82.81%	66.67%	66.67%	47.22%	37.21%	62.20%
Satisfied	15.63%	20.00%	20.83%	41.67%	39.53%	28.23%
Neither satisfied or dissatisfied	1.56%	6.67%	8.33%	8.33%	16.28%	6.70%
Dissatisfied	0.00%	6.67%	0.00%	0.00%	2.33%	0.96%
Very Dissatisfied	0.00%	0.00%	4.17%	2.78%	4.65%	1.91%

Figure 8.19: Visitors' satisfaction across age groups

younger visitors were most satisfied with the exhibit, as the percentage of the respondents in each age group who reported being "very satisfied" was highest in the youngest age group (under 18) and decreased as the age group increased. The oldest visitors (over 60s) were the only age group who reported being "satisfied" more frequently than being "very satisfied" as a percentage of the total respondents in that age group. They also reported being "neither satisfied nor dissatisfied" more frequently than any other age group when the data is analysed based on the percentage of the total respondents in each age group. There were no reports of being "dissatisfied" or "very dissatisfied" in the under 18 age group, while less than 2% of the respondents in this age group were "neither satisfied nor dissatisfied", hence the under 18 visitors were the most satisfied with the exhibit, with over 98% reporting being "satisfied" or "very satisfied". Similarly, approximately 87% of the visitors aged 18-30 and 31-45 reported being "satisfied" or "very satisfied" with the exhibit. While similar figures hold for the 45-60 age group, the responses are more evenly split between the "satisfied" group and the "very satisfied" group, unlike the three younger age groups who reported being "very satisfied" far more than being "satisfied".

One hundred and sixty-four (164) respondents also provided feedback about what they liked about the exhibition. The response was overwhelmingly-positive, with emphases on the VR aspects, 3D content, interactive displays and the combination of the physical and digital content in the exhibition space. One respondent stated that it gave visitors the *"chance to see what it*

Theme		Observations
Technology	Liked the Virtual Reality and interactive components	51
	Liked the 3D artefacts and spherical imagery	17
	Felt educated and entertained by the exhibit	25
Heritage	Appreciated the depiction of local heritage and history	23
	Appreciated the combination of physical and digital content	9
	Remarked on the changed perception of heritage content	8

Table 8.4: Summary of thematic areas from qualitative evaluation

would be like looking at the village itself, using the virtual reality headset", another stated that it *"absolutely fantastic way to immerse yourself with an artefact!"*, and another stated that it is *"a good way to learn about the medieval houses on Moncrieff Hill"*. The qualitative data provided by the 164 respondents were analysed thematically using the process discussed by Nowell et. al. [274]. The process resulted in six main themes, and once these themes were identified, a secondary observation followed, that each theme was discussed in the context of either (1) technology as an entertainment/educational medium, or (2) the use of technology to interpret and visualise heritage. This results in the grouping of these themes under technology and heritage, as shown in Table 8.4. An appreciation of the VR and interactive displays was mentioned by 51 respondents, 17 respondents liked the use of 3D and spherical imagery, and 25 respondents liked the informative elements (such as the heritage interpretation) and entertainment elements (such as the game controllers) of the exhibition. Respondents also appreciated how local heritage was depicted (citing specific areas and local archaeological findings), appreciated on the combination of physical and digital content in the museum space (for instance by providing context through a comparison of the Pictish stones in the virtual environment and the physical space), and remarked on how the exhibition changed their perception of local history (for instance by reviving the past and by portraying old concepts like the Picts in a new light).

During the open-day exhibition, twenty-two (22) visitors filled out an experience questionnaire (shown in Appendix Table. B.6) to evaluate whether the system is easy to use, whether they would recommend it for learning history, whether it has changed their perception of the subject matter, whether it has stimulated their interest in learning and whether they felt immersed in the virtual environment. The responses were gathered on a 5-point Likert-scale questionnaire and Fig. 8.20 shows that the mean responses to all the questionnaire items were positive (i.e. above the point of neutrality) and that the visitors strongly agree that they would recommend the system for heritage learning. The questionnaire also elicited feedback concerning aspects that visitors liked or thought could be improved. Six (6) visitors commented on the use of the technology for learning history, for instance by highlighting the potential for use in schools and museums, and citing it as a potential attraction for the younger generation for heritage teaching and learning. Eight (8) visitors made comments regarding how the use of the system enables them to better appreciate local (Perth) history by providing concrete, pictorial illustrations of the Pictish stones, the hill forts, and the surrounding landscapes. Visitors' comments also bordered on the appreciation of the combination of history and technology in a museum setting. In terms of negative and advisory feedback, they cited the inclusion of audio – one participant cited "auditory input" as an addition they would like, and another stated "sound on the VR to explain what you are watching" – more interaction and immersion in places - for example, one participant stated "Add NPCs for more immersion, also interactive objects would make it more interesting", and another stated that "the lighting could be improved in some areas" – and VR headset compatibility with eye glasses – for example one participant stated "allow glasses users to continue to wear glasses" and another stated "if you could adjust the views for eyesight, can't use it with my glasses". There was also a solitary mention of dizziness, as one participant stated that the experience was "slightly disorientating".

8.5.2 Interaction

Overall, participants reported that the exhibit was easy to use, that it stimulated their interest in the subject matter and that they felt immersed in the virtual environment (see Fig. 8.20). A notable observation made during the open-day exhibition was that the visitors could not see the controller while they had the headset on. This may seem trivial and obvious because the view of the virtual environment in the headset occludes the participants' view of the physical world; however, the implications for interaction are significant. For instance, participants who are not familiar with controllers (enough to know where the buttons are without looking) may struggle to operate the exhibit, as was sometimes observed. This lends credence to the controller-free interaction design which was explored and adopted with the Curing Yard exhibit in chapter 5. The benefits of a controller-free design notwithstanding, there are limitations, one of which is the lack of added functionality required for interacting with digital objects and navigating between (and within) scenes. Given that both controller-free and controller-based systems have their benefits and limitations (see section 8.3.4), the decision regarding which one to adopt should ultimately depend on the target audience and objectives of the project or exhibit. If the target audience are known or believed to be familiar with game controllers, then a controller-based system that leverages this digital literacy may be beneficial. Otherwise, a controller-free system paired with snackable content may be more practical (see chapter 5). As a compromise between

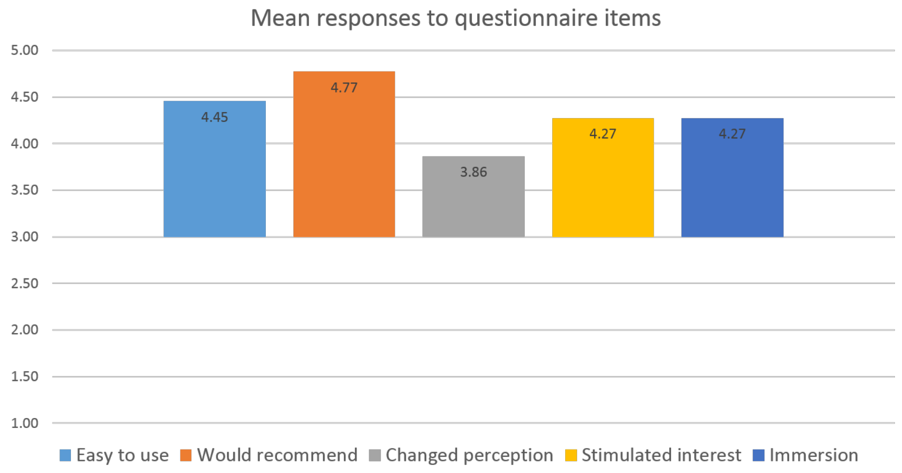

Figure 8.20: Participants' mean responses to the Likert-scale user experience questionnaire

a multi-button controller-based design and a controller-free design, a single-button controller can be adopted to provide more functionality and efficiency than a controller-free design, while exhibiting less complexity than a multi-button controller-based design.

8.5.3 Active vs Passive Engagement

Active and passive engagement are two of five basic needs fulfilled by public spaces [339]. The findings of Memarovic et. al. [340] demonstrate how large displays in public spaces cater to these needs, and suggests that passive engagement allows people to relax through observation of the environment while active engagement fosters intellectual stimulation. Similarly, some visitors during the open-day exhibit exhibited a preference for active engagement by interacting with the VR headset and controller, while others passively engaged by watching the idle content (videos and images) on the screens and observing other active visitors. When encouraged to actively interact with the systems, majority of the passive visitors declined, stating reasons such as *"I've never used one of these (VR devices) before"* and *"I'm fine just watching"*, or no reason at all. The inclusion of screen-based devices, either as a primary output device (in the case of the projectors with the Xbox controller) or accompanying output device (in the case of the screen with the VR headset) therefore extended the reach of the exhibit by delivering content to visitors who were unwilling or unable to actively engage with the technology. This is important because there is evidence to suggest that users may be unwilling to engage with unfamiliar technology [341, 342]. The inclusion of the screen as an accompanying output device to the VR headset also facilitated group-based exploration, as there were visitors who came in groups (i.e. with family and friends). Only one of the visitors could use the VR headset at a time, so the screen enabled the other members of the group to share a view with the headset user thus facilitating conversations between them. This is particularly important because museum visitors can benefit from shared views [343]. However, there is also evidence to suggest that a significant portion of museum visitors visit alone (and not with family or friends) [344], hence the design of digital exhibits should cater to these solo visitors just as much as the group visitors. The distinction between active and passive traditional museum visits is documented in literature [345, 346], but it is worth investigating the roles of immersive technologies in facilitating active and passive interaction in museum spaces, as this will enable exhibit designers to tailor these immersive technologies to project objectives.

8.5.4 Audio

The issue of whether (or not) to include audio in virtual environments was raised in previous case studies, as discussed in sections 4.6 and 5.5. This exhibit was no different, as there were participant comments that bordered on the inclusion of sound in the virtual environment for improved realism or information delivery. One participant stated that *"silence is not normal"*. According to them, some sort of ambient sound in the form of birds chipping, water flowing, or wind blowing would to make the VR experience more natural. It is pertinent to note that this participant suggested the inclusion of ambient sound not for improving the immersion in the virtual environment but to make the (already-felt) immersion more natural. Another participant stated that they would like *"...sound on the VR to explain what you are watching"*. This comment is along the lines of an omnipresent narrator that provides auditory content to complement the visual content already in the exhibit. The inclusion of audio content in a virtual environment has a significant effect on the perception of both the virtual environment and the physical environment of the user [275, 347, 348, 349], hence it may be worth investigating the effects of audio on user experience in immersive digital exhibitions; however this investigation is beyond the scope of this discussion. Rather, this discussion recognises the benefits that may accrue from the combination of audio with other media types as discussed in section 1.1.3, and adopts the position that the decision to include audio content in immersive museum exhibits is similar to the decision regarding a controller-free vs controller-based interaction design; the decision should be driven by the project objectives and target audience. For instance, if a primary goal of the museum exhibit is to immerse visitors in a virtual environment, then the inclusion of ambient sound may be in order. Conversely, immersion may be undesirable or impractical in a virtual museum exhibit, in which case ambient sound may not be beneficial. On another note, a virtual, omnipresent narrator may serve as an additional medium for content delivery in an exhibit, or may be used as an alternative to text interpretation for visually-impaired visitors.

8.5.5 Expert Evaluation

An interview was conducted with Mark Simmons, the Senior Collections Officer at the Perth Museum and Art Gallery and curator for the Picts & Pixels exhibition. The interview was conducted at the end of the three-month exhibition period in a debrief fashion, with questions asked about the team's initial expectations compared to eventual satisfaction with the project, as well as the challenges faced, while still allowing for flexibility for the discussion to evolve to allow for the discovery of additional information. The findings are discussed in terms of challenges (as well as how they were addressed) and expectations of the project.

Challenges: When asked about the challenges to the adoption of digital exhibition spaces in museums, Simmons highlighted conviction, cost and expertise.

- Conviction: A challenge early on when considering the adoption of digital media was the lack of evidence to demonstrate the feasibility of digital exhibits. Staff at the museum were sceptical about the use of digital exhibits owing to their perception of VR technology as clunky, rough around the edges, and unfinished as a product. When asked about the decision to proceed with the digital exhibition in spite of this conviction, Simmons stated that the museum decided to install a digital exhibit because technology can serve as a "strong and interpretive tool". He stated that traditionally, and without digital media, exhibitions are typically comprised of objects with text labels and interpretation. However, with digital media, content can be recreated, exhibits can be designed as interactive games, and visualised in engaging ways. These are examples of its potential, which make it worth exploring from the museum's perspective.

- Cost: The cost of purchasing equipment, maintaining equipment, and creating digital content for exhibitions was perceived as a major challenge to the development of digital exhibits. However, after sourcing for funds for the project, he reflected on how the reusability of digital content proved beneficial for the exhibition in retrospect. For instance, the digital representation of the stones that were created using photogrammetry were used in the on-site exhibition as well as in online archives, and the 3D environment that was generated using modelling software was used in the on-site exhibition, online virtual tours and on mobile apps. This demonstrates the reusability of digital content on multiple platforms and for multiple purposes, which significantly reduces the effective-cost of content creation.

- Technical Expertise: In addition to the cost, the technical expertise required to set up and maintain equipment may not be readily available in house, which may necessitate the outsourcing of tasks, as was the case with the Perth Museum and Art Gallery. Simmons

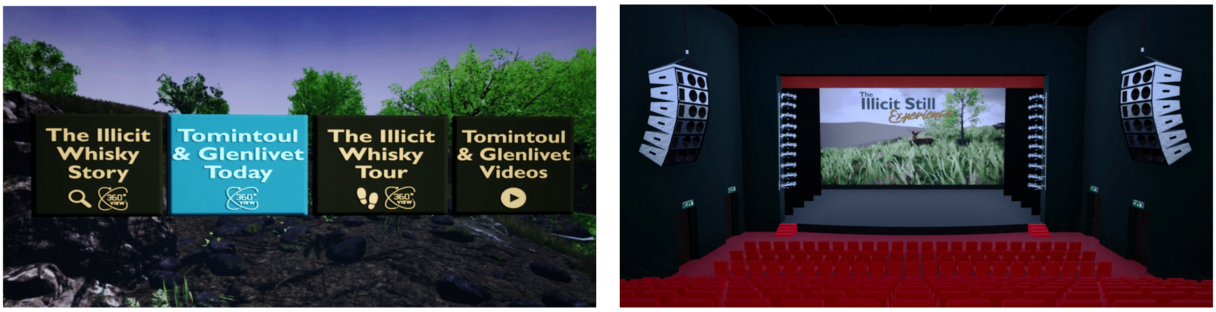

Figure 8.21: Tomintoul exhibit start menu (left) and view of the virtual theatre (right)

remarked that outsourcing is not always feasible or palatable in cases where it is expensive or where the technical experts are perceived as unapproachable and unfriendly, as they are sometimes perceived by heritage practitioners. This challenge can be addressed by: (1) effective interdisciplinary collaboration, which encourages technical experts and heritage experts to work together to bring the past to life, and (2) museum empowerment, by training heritage experts on the use of commodity technology to digitise museum collections. The workshops held during the EU-LAC project demonstrate the feasibility and value of this interdisciplinary collaboration and museum empowerment as discussed in section 7.1.

– Heritage Expertise: While the availability of technical expertise is important for the creation and management of digital exhibitions, the availability of heritage experts to liaise with the technical experts is just as important, as Simmons remarked. As obvious as this may seem, museums may be stretched thin and may not have the man power dedicated to working with technical experts, or they may just not have personnel that possess the knowledge of the content required. An awareness of this challenge is the first step in resolving it, because conscious efforts can be made in allocating resources during the planning phase of a project.

Expectations: When asked whether the end result of the exhibition exceeded, met, or fell short of expectations, Simmons stated that the exhibition exceeded expectations from the visitors' perspective, as demonstrated by the survey data, which shows that 82.30% thought it was good value for money and 90.43% expressed that they were either satisfied or very satisfied with the exhibition (see Fig. 8.18). Simmons remarked that one way in which visitor expectations were met was by combining physical and digital content in the same exhibition space as shown in Fig. 8.17. The essence was to provide a meaningful experience to different audiences. The combination of the digital and traditional content allows the museum to deliver a meaningful experience to visitors who are interested in novel technologies and digital content in museum exhibition spaces, as well as those who are interested in traditional heritage content.

Figure 8.22: Reconstructed Tomintoul landscape (left) and illustration of a historical still [350] (right)

8.6 Deployment and Evaluation – Phase 3

Section 8.3.3 discussed the design of templates that facilitate easy development of immersive digital exhibits in game engines. These templates and concomitant framework have been reused – in partnership with the Tomintoul & Glenlivet Development Trust – to create an immersive, on-site, whisky-themed exhibit in the Tomintoul Museum and Tourist Information Centre in Tomintoul, Scotland. The exhibit, which provides an informative resource on the 18th-century illicit stills in the Scottish Highlands, features similar components to that which was implemented in the Perth Museum and Art Gallery (such as a VR headset, a screen, spherical images and 3D environments), and to accommodate the stakeholders' requests, the exhibit replaces the 3D gallery level with a virtual theatre level in which videos of the surrounding landscape are shown in 3D space (see Fig. 8.21). Furthermore, the design of the interactive components has evolved based on the findings of the open-day evaluation at the Perth Museum, and thus features a single-button controller for interaction as a compromise between a multi-button controller-based design (as used in the Perth Museum [section 8.5]) and a controller-free design (as used in the Timespan Museum [chapter 5]).

The exhibit features a reconstruction of the landscape as it was in the 18th century, as well as cottages, caves, kilns and barley fields that would have existed in the area at the time as depicted by historical evidence and Ordinance Survey (OS) data (see Fig. 8.22). Users can walk-about the virtual environment by pressing and holding the single button on the controller which triggers movement in the direction that the user's head is facing, thus enabling them to move in any direction by simply turning their head (and body if required for comfort) towards that direction, or to take in a view from a stationary position by looking around when the button is not depressed. Spherical images have been extracted from the 18th-century reconstruction, as well as captured from present-day Tomintoul, and these have been used in conjunction with text hotspots to make informative virtual tours that tell stories about a whisky distillation process. As of this writing, the exhibit is yet to be publicly-launched, hence a large scale user evaluation has not been possible. Nonetheless, the feedback obtained during the stakeholders during the design stages and upon deployment and inspection of the final exhibit have been positive, and exude confidence in the success of the exhibit.

8.7 Discussion

A Virtual Museum Infrastructure (VMI) has been designed with input from heritage practitioners and community members in workshops held in over 20 museums, and implemented with affordable mobile and web technologies to enable dissemination of heritage content that are created by museums. Museums and their communities can also curate content using an online management interface using metadata. The design, implementation and evaluation of the VMI have focused on: (1) how to make the most of the infrastructure available to community museums, and (2) how to maximise experience for the users while minimising resource utilisation. To address these issues, this work has demonstrated how community members can be equipped with the skills required to digitise artefacts (via photogrammetry), create virtual tours (via spherical images) and document intangible heritage (via audio and video recordings) using commodity devices and already existing digital literacies. This is exemplified by the creation of over 100 3D models and virtual tours in over 20 community museums worldwide (see Table 8.3), thus museums can now digitally curate, preserve and disseminate their heritage with minimal costs (**Q1**). In addition, the created digital content can be reused in immersive on-site exhibits where resources and infrastructure permit, and this has been demonstrated to yield high satisfaction and positive impact for museum visitors (see section 8.5, **Q2**). Furthermore, these uses of technology in museums have the ability to transform museums practice. For instance, by virtue of the increasing accessibility of affordable technology, entry barriers have been lowered, such that heritage practice can now be collectively performed by the greater community through increasing participation, as opposed to being performed by a few experts (**Q3**). This is exemplified through the VMI features that espouse reproducibility such as a management interface for curating and managing metadata, an interface for making Virtual Reality mobile apps, and the design of "drag and drop" templates for replicating functions, characters and interaction paradigms in game engines and virtual environments.

The availability of resources imposed a major challenge on the use of the VMI in places. For instance, while photos can be captured using smartphones and commodity cameras, the photogrammetry processing requires graphics- and memory-capable computers which some community museums do not possess. By virtue of technology trends (see section 1.2), more capable computers are becoming more affordable, thus rendering this challenge less significant.

In the meantime, an automated photogrammetry processing system has been implemented as part of the management interface in the VMI, so that museums can upload batches of photos and receive processed 3D artefacts.

Moving forward, engaging with museums is essential to ensure the continuity of, and provide support for the processes of digitising, curating and disseminating heritage. The practice of using technology in museums will continue to develop as current technologies evolve and new ones emerge. It behooves both technology practitioners and heritage practitioners to work together to develop and refine paradigms for the curation, preservation and dissemination of heritage, as well as solve whatever socio-technical challenges that will undoubtedly arise.

8.8 Chapter Summary

This chapter has discussed the design, implementation and evaluation of the Virtual Museum Infrastructure (VMI) introduced in chapter 3. The VMI caters to the requirements of museums with varying amounts of resources by leveraging existing resources and using affordable technology. The design has been implemented as a multifunctional system which features a data store, a management interface and platforms for disseminating content via the web, mobile apps and through immersive on-site installations. The evaluation of the VMI reveals with heritage experts and museum visitors suggests that it makes positive impact and adds value in the curation, preservation and dissemination of heritage.

Part V

Final Thoughts

Chapter Nine

Discussion

A Virtual Museum Infrastructure (VMI) has been designed, implemented and evaluated. The findings – that affordable technologies can be adopted for participatory heritage engagement, that digital literacies can be leveraged where infrastructure and other resources are in short supply, and that mobile and web technologies can be used to foster global dissemination of resources – are all central to the design of the VMI. The implementation, deployment and evaluation of the VMI demonstrate how technology can facilitate the curation, preservation and dissemination of heritage in museums, how museums can maximise the resources at their disposal to facilitate these, and how technology has the potential to transform museums and heritage practice.

9.1 Introduction

This work has hitherto discussed the design, deployment and evaluation of a Virtual Museum Infrastructure (VMI) as a means of facilitating the curation, preservation and dissemination of heritage – functions which are performed by traditional museums [9, 11, 12]. This discussion therefore begins by exploring how these functions are facilitated by the VMI, and further explores how the three facets of this research – affordability, engagement and global access – identified in section 1.2 (see Fig. 1.3), are actualised. For affordability, the VMI has been designed to use the resources at museums' disposal including infrastructure (if present), manpower and digital literacies, hence a discussion of resource maximisation in the context of heritage management is provided. For engagement, this discussion highlights the ease of use, value and satisfaction from visitors' perspectives and the ease of management from heritage experts' perspectives. For global access, the VMI has been designed to put community museums on the map and encourage engagement between geographically-dispersed communities, and a recap of how emergent technologies can be used to facilitate the exploration of inaccessible sites irrespective of power and Internet access, is provided. This discussion proposes with some implications for museums in view of these findings, highlights some research challenges and proffers solutions.

9.2 Curation, Preservation and Dissemination

Museums are responsible for curating, preserving and disseminating heritage. This work has investigated the use of smartphones, VR devices and drones in chapters 4, 5 and 6 for disseminating heritage both within and outwith of museums. Spherical media and drone footage are also explored for heritage preservation, through an immersive on-site exhibit that facilitates a comparison of the past and present from equivalent vantage points (chapter 5), and an immersive aerial exploration system which facilitates remote access to endangered heritage sites thus ensuring their longevity (chapter 6). The workshops discussed in chapter 7, where heritage practitioners and community members are introduced to the digitisation of heritage content using commodity 3D technologies, provide insight into how emergent technologies can support the curation, preservation and dissemination of heritage in museums. The workshops demonstrate a variation in the level of resources available to museums, such that some museums are more equipped than others in terms of the infrastructure and equipment required to digitise and disseminate heritage. This variation is visualised as a resource spectrum, which is illustrated in Fig. 1.2. By virtue of this variation, the VMI has been designed to cater to museums on the spectrum, ranging from the far left where resources are scarce to the far right where there is a reasonable availability of resources. The design incorporates a web-based management interface which enables the creation of digital content, as well as a detailed description of these content using widely-known metadata schemas. The interface also provides opportunities to modify these data. Hence, the summation of the processes that span the creation and description of content with these metadata schemas address the curation function of museums using affordable technologies. The created data is housed in a data store which ensures that heritage content and its associated data will live on in digital form if the original artefact is lost or damaged, thus addressing the preservation function of museums. The stored data can be remotely disseminated on the web via a map-based interface, on mobile devices via app stores, and in museums as on-site exhibits. In this manner, the VMI allows communities to collectively create, curate and store content, which can be reused in multiple contexts and disseminated on multiple platforms.

The design, implementation and evaluation of the VMI constitute one of the contributions of this thesis, **C1**, which details a VMI that leverages emergent technologies for the curation, preservation and dissemination of heritage in exciting and interactive ways (**Q1**). The VMI supports active use, such that stakeholders can continuously create, maintain and manage content using an intuitive management interface that allows content description using metadata. The storage and management features of the VMI encourage content dissemination in immersive and engaging ways on mobile devices and the web, thus tapping into digital literacies and resources that are available to communities (**Q2**). This way, the VMI is akin to traditional museum archives where artefacts can be described, stored and reused in multiple exhibitions. Furthermore, the accessibility of management interfaces on web and mobile platforms enables community members and the general public to collectively create, curate and disseminate heritage, thus shifting the responsibility of heritage management from a few heritage experts to the greater heritage community. This represents a shift in heritage practice, as the proliferation of technology among heritage aficionados has resulted in new forms of heritage content, new relationships with exiting heritage content, and a redefinition of roles among stakeholders (**Q3**).

9.3 Resource Maximisation and Heritage Engagement

An investigation of the efficacy of an affordable (mobile) VR system as compared to more costly (desktop-based) alternatives is discussed in chapter 4. The design features a smartphone which renders immersive scenes using spherical imagery and combines this with intuitive head tracking. A comparative analysis of VR systems in the context of heritage education in secondary schools revealed that affordable VR systems can provide interaction and engagement levels that are comparable to more costly alternatives, hence they are cost- and task-effective for heritage learning. When coupled with the findings of Papachristos, Vrellis and Mikropoulos, that there is no significant difference in the levels of spatial presence and learning outcomes provided by mobile VR systems as compared to the Oculus Rift HMD [273], a case can be made for the use of these affordable, low-end mobile VR systems where resources are in short supply.

The design, implementation and evaluation of the mobile VR system for heritage learning (discussed in chapter 4) constitute one of the contributions of this thesis, **C2**, which explores how affordable technology and already-existing resources can facilitate a comparison of the past and present states of a heritage site, coupled with an evaluation of the systems to understand and push the boundaries of these systems thus maximising experience while minimising resource expenditure. It also goes towards addressing research questions **Q1** and **Q2**. **Q1** asks how emergent technologies can support and replicate heritage curation, preservation and dissemination, functions of traditional museums as identified in literature. Chapter 4 investigates the dissemination function of museums by exploring interactive ways to communicate heritage content. For instance, by using mobile devices to communicate heritage content, the approach leverages digital literacies and commodity hardware already possessed by the target audience, which has the advantage that there is little familiarisation required. Furthermore, mobile phones are combined with VR headsets for improved engagement with heritage content, and the availability of mobile app stores facilitates the widespread dissemination of content. **Q2** explores how heritage can be disseminated where resources may be limited. This provides the motivation for the comparative evaluation of the low-end system against more expensive alternatives. The findings of chapter 4 suggest that equivalent levels of engagement can be obtained from both systems, and this provides the impetus for further exploration of mobile VR systems and their incorporation into the VMI. Furthermore, the instantiation of the VMI as detailed in chapter 8 demonstrates how museums can deliver heritage content to visitors via mobile apps that they can download and run on their smartphones and tablets. This enables heritage organisations to foster engagement in the absence (or shortage) of dedicated infrastructure. In addition, chapter 8 also demonstrates how game proficiencies of museum visitors can be leveraged to foster engagement with virtual exhibits on-site. This takes advantage of already-existing skills and thus minimises the time required for visitors to familiarise themselves with the delivery platform, in line with the practice that digital exhibits should be designed so that visitors do not have to learn how to use them [351, 352, 343].

9.4 User Experience and Content Management

The use of spherical media, as introduced in chapter 4 is further explored in chapter 5, albeit with differences in context and platform. Spherical media are used to depict a scene in an on-site museum installation in form of snackable content which caters to less tech-savvy visitors, encourages relatively-short interaction times and facilitates easy absorption of content when visitors engage with museum installations. An evaluation of the on-site virtual museum installation is conducted from the perspective of museum visitors and staff. The findings suggest that spherical media can deliver immersive experiences in museums, and when used to implement controller-free exhibits, it can encourage quick and easy engagement with digital content, even for visitors with inclinations for technology aversion. This re-affirms the case for using spherical media and provides further impetus for the incorporation of spherical scenes in the VMI. Furthermore, chapter 5 serves as a starting point in the investigation of on-site museum exhibits where infrastructure permits.

From the perspective of museum staff, the cost of development and maintenance, and the ease of management of digital content in museums emerged as the determining factors in the adoption of technology-based content in museum spaces. The implication of this is twofold. Firstly, the emergence of cost as a determining factor in the adoption of digital content in museums, lends further credence to the concept of a resource spectrum (as introduced in chapter 1 and discussed further in chapter 8) which necessitates the provision of digital content in a cost-effective manner and in accordance with the resources available to museums. Secondly, the ease of management of digital content and infrastructure in museums is of the utmost importance. This can be actualised by providing support at each stage of the digital exhibition process, from creation to deployment to day-to-day maintenance. This is further explored in the design of a content management interface (discussed in chapter 8) that facilitates the creation, curation and preservation of heritage content.

The design, implementation and evaluation of this immersive, on-site virtual museum system (discussed in chapter 5) constitute one of the contributions of this thesis, **C3**, which explores how affordable technology can be used to foster an on-site comparison of the past and present states of a heritage site from equivalent vantage points, coupled with an evaluation which explores practices for deploying on-site exhibits and facilitating ease of management in museums. It also goes towards addressing research questions **Q1** and **Q2**. The dissemination aspect of **Q1** is further addressed, as the use of spherical imagery which was introduced in chapter 4 is further explored in chapter 5. Heritage content management by way of content creation and curation in museums, is also introduced given the emergence of ease of management as an important factor in the adoption of digital content in museums. **Q2** is addressed here through an approach that caters to museums with the resources required to deploy on-site virtual exhibitions. Chapter 5 explores the use of a controller-free interaction design with snackable content, with the aim of attracting less tech-savvy museum visitors and delivering a quick, yet impactful educational experience. The design decisions in chapter 5 also demonstrate the trade-offs between cost and user experience by adopting a desktop-based system with spherical environments in lieu of a mobile-based system with spherical environments or a desktop-based system with free-form 3D environments. The deployment of on-site virtual exhibits is further explored in the Picts & Pixels exhibition discussed in chapter 8 as a means of catering to museums with the resources required to leverage mid- to high-end devices to deliver compelling experiences to museums.

9.5 Social Exploration and Remote Access

While chapter 6 explores the dynamics of group interaction with HMD-based systems, it also explores a system for remote heritage exploration, so that the case studies discussed in chapters 4 – 6 collectively investigate the three types of virtual museum systems, classified based on location and mode of exploration (see section 2.3). Recall that virtual museum systems can be categorised into three types based on location of exploration: systems that facilitate both remote and on-site exploration, systems that facilitate on-site exploration, and systems that facilitate remote exploration [30]. The discussion in chapter 4 revolves around the exploration of heritage both at the remains (on-site) and in the classroom (remote), while the discussion in chapter 5 facilitates an on-site exploration of heritage (in a museum setting). This leaves the exploration of heritage in remote scenarios, hence the essence of chapter 6.

The design, implementation and evaluation of this VR system for remote, real-time exploration constitute one of the contributions of this thesis, **C4**, which explores how a drone-based VR system can be designed to facilitate remote exploration of heritage sites with an affordable, power- and Internet-independent infrastructure. The system features a drone which can be flown over a landscape, a Raspberry Pi which can draw electricity from a mobile power bank and one or more client phones in cardboard VR headsets. The work also serves as an initial exploration into how VR systems fare in group exploration of heritage sites. Given the shortage of literature in this niche area, the investigation begins with a cursory exploration of group dynamics by enabling pairs of users to share a view of a remote site in an immersive VR system. The system also enables one of the users to control this view intuitively with their head movements while the others act as passive receivers of the live footage being transmitted. The dynamics of group-based exploration with immersive VR headsets are further explored in the Picts & Pixels exhibition discussed in chapter 8, albeit with a combination of VR headsets and screens, and the findings demonstrate that while immersive VR can be an isolated experience [282], inclusive experiences can be designed for groups by mirroring content on large displays so that the headset user and other group members alike can engage in discussions.

9.6 Interpretation and Implications

An interpretation of the aforediscussed results is provided in this section with an emphasis on the digital literacies of museum visitors and the variations in the level of resources available to museums. Given these insights, the implications for museums practice are also discussed in the context of exhibit design in museums – specifically how museums should tailor the features of digital exhibits to the skills and preferences of their target audience – and resource availability – specifically how museums can deploy computing resources to leverage the proficiencies of visitors, and take advantage of widespread smartphone ownership to facilitate heritage dissemination with minimal resources.

9.6.1 Digital Literacies and Exhibit Design

Digital media can make heritage content more appealing to potential museum visitors and the public, as demonstrated in this work. Chapter 4 demonstrates the value of technology for classroom-based and on-site heritage learning, chapter 5 demonstrates how an on-site museum installation can add value to museums, chapter 6 highlights the appeal in the use of novel, albeit experimental systems for outdoor heritage exploration, and chapter 8 shows that on-site digital exhibits can be deployed in ways that result in high levels of satisfaction to the audience. The investigation of the case studies discussed in chapters 4 – 6 and the design and instantiation of the Virtual Museum Infrastructure in chapter 8 have demonstrated that the adoption of technology provides opportunities for situated, personalised and collaborative heritage learning [251, 252], as well as improved engagement with content [253]. These opportunities notwithstanding, it behoves museums to not adopt technology for its own sake because they risk falling into the "technology trap" [351, 352]. Rather, museums stand to gain from a thorough assessment of the needs and backgrounds of their potential visitors and take these into considerations when making decisions about digital exhibits in order to reap the aforementioned benefits [282, 351].

Exhibit design should be informed by the target audience of the museum. Museums may not be adequately aware of their visitors' needs and backgrounds, and may possess even less knowledge about the non-visiting public [353]. They should therefore devote effort towards understanding their visitors' requirements and these, coupled with the objectives of the project, should inform the nature of the digital components used in exhibits [282, 354]. For instance, spherical images and videos can be used to deliver snackable content where the museum wishes to encourage short dwell times or where the target audience is perceived to have little or no experience with VR or games. 3D environments with interactive objects can be used where the museum wishes to encourage more dwell time or where the target audience are perceived to possess the digital literacies required to interact with the exhibits. Visitors' smartphones can be used with downloadable mobile apps and VR headsets where resources are scarce, or entire physical spaces can be devoted to the installation of immersive digital exhibits where resources are available. Head Mounted Displays (HMDs) with audio headphones can be used to foster isolated immersive experiences for solo visitors as there is evidence to suggest that a significant portion of museum visitors visit alone [344], while Computer Automated Virtual Environments (CAVEs), domes and wide screens can be used to foster shared immersive experiences for group visitors [343]. The dichotomies depicted in the aforementioned examples highlight the different paths that museums can take depending on their project objectives and budget. Succinctly put, the design of digital exhibits should leverage the skills and resources of museum visitors (and staff) so as to preclude the need to expend considerable time and effort in learning how to use exhibits, and thus minimise the risk of visitors being put off by technology in museum spaces [351, 343].

9.6.2 Resource Availability

Chapter 5 identifies cost and ease of management as factors that influence the adoption of digital exhibits in museums, and provides a case study of how the design of on-site exhibitions can take these factors into consideration while providing positive experiences to visitors. This is revisited in the on-site installation discussed in chapter 8, with 90.43% of museum visitors reporting satisfaction with the exhibit and 82.30% stating that it is good value for money (Fig. 8.18). This demonstrates a value-maximising approach to the design and implementation of interactive, on-site exhibits in museums, which is important because of the resource constraints (such as a shortage or lack of power, networking and computing facilities) that some museums face. A discussion of these resource constraints is provided in chapters 1 and 8 with an emphasis on the variation that exists in the level of resources available to museums. The museum workshops discussed in chapter 7 evidences this variation, and this lends credence to the need for a VMI that caters to the significant variation in resources.

As demonstrated in chapter 8, compelling, virtual experiences can be facilitated on-site where resources are available. These resources may be obtained through fund-raising, grant revenues and admission fees [34]. While there are arguments for and against charging admission fees in museums, the extent to which these fees affect visitor numbers is unclear [355, 356]. What is clear is that museums, in alleviating their financial and economic pressures [34], aim to recoup the costs of digital exhibitions, and where the prospects of replenishing costs are not palatable, they may consider the more affordable options. This work has demonstrated that digital media as applied to heritage exploration can take various forms, with a significant range in cost. Chapters 4 and 5 demonstrate that compelling experiences can be delivered to users with low- to mid-end systems, and chapter 8 demonstrates how a museum's resources can be put towards delivering an immersive, on-site experience in ways that do not draw attention away from, but rather complement the physical artefacts and traditional museum offerings [357]. Museums with little or no resources can leverage visitors' digital literacies using the VMI to build VR apps that visitors can install on their smartphones and use with affordable VR headsets, while museums with substantial resources can deploy an array of systems such as VR headsets and large displays to render 3D scenes and interactive objects. These approaches collectively demonstrate how the VMI can cater to museums with varying levels of resources.

9.7 Challenges and Solutions

The methodology adopted for the research necessitated the evaluation of systems in physical-world settings, and this made the organisation of evaluation exercises more challenging in terms of the time and effort required to plan and find suitable participants. Also, while some noise was introduced into the data due to the physical-world nature of the studies, the natural setting makes the data representative of the context in which the evaluation took place. In addition, no video recordings were obtained as part of the data gathering process because of the unsavoury perception and ethical implications associated with recording participants in physical-world settings [237, 238, 239, 240]. As a replacement for video analysis, in-person observations with detailed notes were used in conjunction with participant comments during focus group sessions.

Like most relatively-new ideas, concepts and technologies, Virtual Reality (VR) still has challenges that need to be resolved before it can become truly mainstream. Even with well-designed and implemented VR headsets, some users report motion sickness. However, this may not be a shortcoming of the technology itself, as approximately one-third of the worlds population experience motion sickness from various activities including car rides, park rides and plane travel, yet its underlying causes are not well-known [358, 359]. In addition to motion sickness-prone users, visually-impaired users, particularly those with glasses may struggle with VR headsets, as they may either have to endure an uncomfortable experience with the headset strapped on top of visual aids, or have to endure blurry/fuzzy images without visual aids. Another group of users who may struggle to get a good experience from VR headsets are those with relatively-small heads, notably pre-teen users. While VR headset are usually designed with adjustable straps to account for different head sizes, the straps may not fasten the headset properly on users with smaller heads, which introduces two problems. Firstly, it introduces discomfort because it weighs down and moves around on the user's face. Secondly, it prevents proximity sensors (if available) on the headset from detecting faces, which may result in an intermittent loss of images if the headset is designed to only show images when a face is detected (as is the case with the Samsung Gear VR). The implication of this is that a user either has to endure sporadic images or be forced to hold up the headset on their face in order to get uninterrupted images. The former may result in an inferior experience where content is unintelligible, while the latter may result in an ineffective experience where the user's hands are unavailable for interaction or gesturing. Motion sickness while using VR systems can be mitigated by designing the experience to limit unnecessary movement, and by optimising the system to minimise head-tracking latency. In addition, VR headsets can be redesigned to account for a range of head sizes and visual impairments, for instance, by providing multiple form factors with different shapes and sizes, introducing a wide-range focus wheel that controls lens spacing, and/or providing allowance in the headset for glasses.

In addition to technical challenges, there are museological challenges to consider. One such challenge is the resistance of heritage practitioners in adopting 3D imagery due to the view that 3D models are inferior replicas that may undermine physical artefacts [130]. The adoption of an anti-materialist approach to the functions of museums, where museums are viewed as establishments that disseminate heritage instead of mere object repositories, can help to reduce the hostility towards 3D models [130]. It is hoped that this change in mindset will result in the treatment of physical artefacts and digital artefacts alike, as entities in their own rights with attribution, authorship and provenance, and the proliferation of 3D technologies and systems such as the Virtual Museum Infrastructure proposed in this thesis will catalyse this change. In addition, although this research has focused on resource maximisation through a Virtual Museum

Infrastructure (VMI) that can cater to both museums with substantial resources and those where resources are in short supply, it is important to recognise that there will always be competing needs in the day to day operations of museums, and by virtue of this, it can be argued that there may never be enough (financial) resources devoted to digital exhibits. Museums therefore need to prioritise when allocating resources, and it is hoped that the benefits to be had from the adoption of technology will be taken into consideration when making these decisions.

9.8 Chapter Summary

A discussion of the Virtual Museum Infrastructure (VMI) has been provided with a focus on how it facilitates the curation, preservation and dissemination of heritage both within and outwith physical museum spaces. Heritage curation is facilitated by empowering museums and communities with the skills and technical expertise required to digitise content, and a content management interface enables communities to upload data and metadata to a data store. By holding content in the data store, heritage is guaranteed to live on in some form even if the physical counterparts are lost or damaged, thus ensuring heritage preservation. The created content can also be disseminated via the web, mobile apps and on-site museum installations, and where infrastructure is lacking, digital literacies already possessed by the public can be leveraged. Instantiations of the VMI have been evaluated with heritage experts and museum visitors, and the findings demonstrate that the use of digital media in museum contexts can increase visitor satisfaction, add value and make positive impact in the heritage community.

CHAPTER TEN

CONCLUSION

This chapter provides concluding thoughts of this thesis. The methodology and work done hitherto are summarised, the thesis statement is reiterated, the activities that address the research questions are outlined, the resulting contributions are discussed, and a forward-looking discussion concludes the chapter.

10.1 The Story So Far

This thesis has discussed the investigation of emergent technologies for heritage curation, preservation and dissemination. A practice-orientated methodology has been adopted for the research, such that as part of the work done, the use of affordable, immersive technologies for heritage exploration has been investigated by means of underpinning theory, case studies, museum workshops and on-site installations. These collectively demonstrate the value of digital technologies for heritage management, and make a case for the use of technology for supporting the functions of traditional museums, hence a Virtual Museum Infrastructure (VMI) has been proposed to this end. The functions of museums which have been identified are content *curation*, *preservation* and *dissemination* [9, 11, 12]. The feasibility of a VMI has been investigated during workshops conducted with museums across Europe, Latin America and the Caribbean. These workshops demonstrate a variation in the resources available to museums, a phenomenon which has been visualised using the resource spectrum (Figs. 1.2 and 8.1). To this end, VMI design caters to the requirements of museums along the spectrum, through an inexpensive local infrastructure that leverages existing resources and digital literacies (towards the left end) and an immersive on-site installation that facilitates heritage exploration in engaging and exciting ways (towards the right end). The VMI design has been implemented as a multifunctional system. A management interface – which enables heritage experts and community members to create and update content (such as images, virtual tours and 3D artefacts), as well as supply metadata to describe these content – fulfils the functions of content creation and curation. The created content is stored in a remotely-accessible data store, which ensures the preservation of content in the event that the original artefacts or objects are lost or damaged. Furthermore, a map-based interface – which enables the general public to visualise content by geographical location and type (for example virtual tour, 3D artefact, image, audio and video), and presents descriptions for these content using a Wiki that features metadata supplied using the management interface – fulfils the function of content dissemination. In addition, the dissemination of these content can be actualised through the use of interactive and immersive media both on-site and online, as users can immerse themselves in virtual tours of remote locations, or interact with 3D artefacts using their smartphones and VR headsets. The VMI has been deployed in over 20 museums worldwide and evaluated in terms of its impact and value for heritage management. The findings demonstrate the feasibility of digital media for participatory heritage engagement.

10.2 Thesis Statement

The thesis of this dissertation is:

> *technological advances and concomitant upsurges in digital literacies facilitate the development of a novel Virtual Museum Infrastructure that enables museums to actively and continuously curate, preserve and disseminate local heritage on a global scale.*

Technology has advanced rapidly for decades, as projected by Moore's law [40] and emblematised by increases in processing capabilities, and reductions in cost and size. These have resulted in the ubiquity of computing devices to the point of permeating aspects of daily life, and thus provide opportunities that manifest in the form of widespread digital literacies, mass communication and collaboration platforms. These advances create the potential for the disruption of industries, as evidenced by the disruption of publishing by Facebook[1], hospitality by Airbnb[2] and transportation by Uber[3]. This work has not attempted to disrupt the heritage industry, but rather explores how to build on (and transform) existing practices and infrastructure through the use of technology to support participatory heritage curation, preservation and dissemination – the three main functions of museums. This work has therefore demonstrated how the curation, preservation and dissemination of heritage can be actualised by heritage practitioners and community members by facilitating collaborative creation and archiving through workshops and content management platforms, distributed storage and remote dissemination through web and mobile platforms, and interactive museum visits through immersive, on-site installations, thus transforming the nexus between museums, their respective communities, and the general public.

[1] Facebook©: https://www.facebook.com/ (Accessed: 2018-03-09).
[2] Airbnb©: https://www.airbnb.co.uk (Accessed: 2018-03-09).
[3] Uber©: https://www.uber.com (Accessed: 2018-03-09).

10.3 Research Questions

This research has investigated how emergent technologies can support the functions of traditional museums (**Q1**, Process), investigated how to utilise existing resources for heritage dissemination and management (**Q2**, Infrastructure), and investigated how these systems bring about change in heritage practice (**Q3**, Transformation). The questions that have been addressed in the course of this research are highlighted as follows.

Q1 How can emergent technologies support (and replicate) the functions – heritage curation, preservation and dissemination – of traditional museums?

The importance of this question stems from the need for digital heritage systems to replicate the functions of traditional museums as perceived by heritage practitioners. This question is addressed through the implementation of the features of the Virtual Museum Infrastructure (proposed in chapter 3 and implemented in chapter 8). The *curation* function is fulfilled through the use of a management interface which enables heritage experts and community members alike to add and update content, as well as metadata to describe the content. The created content and metadata are held in a distributed data store, which fulfils the *preservation* function. The *dissemination* function is fulfilled through an online map-based interface which displays heritage content geographically and features functions to view content and metadata (supplied using the management interface), as well as mobile apps and immersive on-site exhibits which facilitate interactive exploration of heritage content.

Q2 How can heritage dissemination systems be designed to maximise value in environments constrained by limited (computing and funding) resources?

This question is in two parts. The first part focuses on maximising value where computing resources are limited, and this is important because an emphasis on participatory heritage engagement where resources are limited implies the use of commodity hardware, either by purchasing affordable devices or by repurposing existing infrastructure which may be constrained by limited computing resources. This part of the question is addressed in chapter 4, through the use of low-end, spherical imagery running on smartphones coupled with affordable VR headsets to depict virtual environments, followed by a comparative evaluation of these low-end systems with high-fidelity systems depicting the same virtual environment using detailed 3D models. The findings of this evaluation show that the low-end systems can be just as effective in delivering an engaging and informative experience. Furthermore, the use of the low-end, mobile-based system opens up new use cases for on-site exploration of heritage owing to the portability and location-aware capabilities of mobile devices.

The second part focuses on maximising value (in terms of experience for users and satisfaction for heritage practitioners) while minimising cost. The importance of this question stems from the limited operating funds usually available to community heritage organisations, which necessitates the use of affordable technology for participatory heritage engagement. This is addressed through the design of the museum installation in chapter 5, by adopting a controller-free interaction design which features snackable content that confines users to a distinct virtual view point which corresponds to a physical world vantage point. Instead of deploying a virtual environment in the form of a 3D environment that requires high graphics and processing capabilities, a virtual environment can be depicted using a spherical image which corresponds to a physical world view that enables users to compare the past and present from equivalent vantage points.

The experience is further enriched by adding interpretation in the form of text snippets and historical images, which can be achieved using smartphones or commodity computer hardware with affordable VR headsets. This demonstrates that – and how – museums can adopt a value-maximising, cost-minimising system design (and implementation) for heritage engagement. In a similar vein, this research has explored how digital literacies – specifically smartphone ownership and gaming proficiency – can be used for participatory heritage engagement, as demonstrated in chapter 8 which shows how visitors' devices can be used in lieu of dedicated infrastructure by delivering content via VR apps that can downloaded on their smartphones and tablets, and explores the use of game controllers, menu systems and navigation paradigms that visitors may be familiar with to foster engagement with heritage content. This represents a viable alternative that may preclude the need for museums to purchase and maintain dedicated infrastructure on-site.

Furthermore, the design features introduced in chapters 4 and 5 are then adopted in the Virtual Museum Infrastructure (VMI) which is implemented in chapter 8, with a focus on how the VMI can cater to the varying needs of museums on the resource spectrum (see Fig. 1.2), thus maximising value and resource utilisation in constrained environments. Thus using the VMI, museums that are closer to the left end of the spectrum (where resources are scarce) can benefit from an affordable, low-powered infrastructure that leverages already existing resources and visitors' digital literacies for heritage engagement, while museums that are closer to the right end of the spectrum can maximise value by promoting interactive engagement with heritage content through immersive on-site installations.

Q3 How do technological advances transform museums practice?

The importance of this question stems from the disruptive nature of technology, which has the potential to bring about new opportunities and challenges for heritage practice. This question is addressed through an investigation of how technology impacts heritage dissemination and management as perceived by practitioners and the general public. The impact on the general public and heritage practitioners is investigated through a combination of questionnaires and semi-structured interviews to elicit feedback on a variety of heritage dissemination systems, and the findings of evaluation exercises discussed in chapters 4 and 5 suggest that technology has the ability to stimulate interest in heritage learning both in the classroom and in museums.

Furthermore, the findings of chapter 8 demonstrates an appreciation for the combination of digital and physical content in a museum exhibit, which may encourage the public to view digital content as entities in their own rights instead of viewing them as replicas of physical artefacts. Also, the proliferation of web and mobile technologies facilitate the development of decentralised Content Management Interfaces (CMIs) that allow groups of people to collectively create, curate and contribute to digital heritage content, thus shifting the responsibility of heritage management from a few heritage experts to the greater heritage community. This is demonstrated through the provision of a management interface discussed in chapter 8, which allows community members to upload content, add metadata, contribute narratives through a Wiki, and disseminate content through social archive platforms.

In summary, advances in, and consequent adoption of technology by museums and heritage practitioners brings about new forms of heritage content, new relationships with exiting heritage content, and a redefinition of roles in the heritage community.

10.4 Contributions

The contributions of this thesis are summarised below:

C1 A Virtual Museum Infrastructure (VMI), designed as a suite of tools for heritage curation, preservation and dissemination, coupled with system instantiations which demonstrate the feasibility of the VMI to support museums with both limited and ample resources [chapter 8]. The VMI has been demonstrated to facilitate active, collaborative use by museums and their communities, content (re)deployment on multiple platforms such as the web, mobile and on-site installations, and interactive experiences through immersive systems.

C2 An affordable VR system which exemplifies the instantiation of virtual museums for both remote (classroom) and on-site (relics) heritage exploration [chapter 4]. The feasibility of this affordable system for facilitating engaging experiences as compared to more costly alternatives is also demonstrated.

C3 An immersive, VR installation which exemplifies the instantiation of virtual museums for on-site access in a traditional museum setting, and facilitates comparison of the past and present from equivalent vantage points [chapter 5]. The feasibility of a resource-maximising, cost-minimising approach, coupled with a controller-free design for delivering snackable content in digital museum exhibits is also demonstrated.

C4 A VR system which exemplifies the instantiation of virtual museums for remote, real-time exploration of inaccessible heritage sites, designed to function irrespective of power and Internet access [chapter 6].

C5 An examination of the fields of virtual museums and digital heritage, and the relationship between them, coupled with taxonomies that demonstrate the varying technologies, use cases and characteristics that are evocative of work in these fields [chapter 2].

10.5 Future Work

If past technological and museological trends are any indication, one can envisage future museums to be forums of public creativity characterised by personalised experiences, collaboration and meaningful conversations on heritage. To remain relevant in the digital age, museums will likely evolve to embody digital and on-site participation. That said, the promise of the interactive museum where visitors can learn, play and connect is not without challenges. For instance, in order for museums to deliver personalised experiences, they need to know their audience. This requires large scale data gathering and analyses that may be time- and cost-intensive. In addition, public creativity and collaboration raise Intellectual Property (IP) issues that revolve around the ownership, distribution and changeability of heritage. Advances in pervasive computing, data science and machine learning can provide opportunities for museums to learn about visitor preferences and deliver personalised content, and advances in computer graphics can facilitate the consumption of these content using ultra-immersive VR systems, some of which may be wearable devices that are owned and regularly used by visitors. Similarly, suitable data management initiatives will emerge in response to IP issues, just as the Dublin Core Schema evolved to provide universal metadata terminology in response to the interoperability problems between the vocabularies of cataloguing and archiving systems. It is hoped that the incorporation of the Dublin Core Schema – which has been around for two decades and is still gaining popularity in the heritage community[4] – in the VMI will mitigate these interoperability issues that currently exist due to a shortage of widely-accepted standards in the heritage community. While scalability is arguably a crucial component of any system in the information age, it has not been the focus of this research. Rather, this work has focused on actualising the functionality of the local VMI, as opposed to implementing a fully-scalable infrastructure. That said, the incorporation of social archive sites (such as Sketchfab and Roundme) which have been demonstrated to be scalable, serves to bring the VMI closer to scalability. To further improve the reach of the VMI, the VMI architecture can be refactored to provide integral scalability, so that geographically-dispersed museums can share resources with each other and access heritage content irrespective of location. These allude to the continuous evolution of the fields of digital heritage and virtual museums as current technologies advance and new ones emerge. The onus is on both technology experts and heritage practitioners to work together to develop and refine paradigms for the curation, preservation and dissemination of heritage, and proffer solutions to the museological and socio-technical challenges that will undoubtedly arise.

[4]Dublin Core RFC: https://tools.ietf.org/html/rfc2413

Part VI

Appendix

Appendix A

Ethics

Ethics approval was obtained for the work done in Chapters 4– 8 due to their nature which involved trialling systems with users and eliciting feedback. The ethics approval documents are shown in Figs. A.1, A.2, A.3 and A.4.

University of St Andrews
Scotland's first university – 1413

University Teaching and Research Ethics Committee Sub-committee

12th June 2015
Adeola Fabola
School of Computer Science

Ethics Reference No: *Please quote this ref on all correspondence*	CS11556
Project Title:	Investigating quality of service and quality of experience HMD-Based virtual reality systems
Researchers Name(s):	Adeola Fabola
Supervisor(s):	Dr Alan Miller

Thank you for submitting your application which was considered at the Computer Science School Ethics Committee meeting on the 11th June 2015. The following documents were reviewed:

1. Ethical Application Form 5th June 2015
2. Participant Information Sheet 5th June 2015
3. Consent Form 5th June 2015
4. Debriefing Form 5th June 2015

The University Teaching and Research Ethics Committee (UTREC) approve this study from an ethical point of view. Please note that where approval is given by a School Ethics Committee that committee is part of UTREC and is delegated to act for UTREC.

Approval is given for three years. Projects, which have not commenced within two years of original approval, must be re-submitted to your School Ethics Committee.

You must inform your School Ethics Committee when the research has been completed. If you are unable to complete your research within the 3 three year validation period, you will be required to write to your School Ethics Committee and to UTREC (where approval was given by UTREC) to request an extension or you will need to re-apply.

Any serious adverse events or significant change which occurs in connection with this study and/or which may alter its ethical consideration must be reported immediately to the School Ethics Committee, and an Ethical Amendment Form submitted where appropriate.

Approval is given on the understanding that the 'Guidelines for Ethical Research Practice' https://www.st-andrews.ac.uk/utrec/guidelines/ are adhered to.

Yours sincerely

P.P.

Convenor of the School Ethics Committee

Ccs Supervisor
School Ethics Committee

ethics-cs@st-andrews.ac.uk

The University of St Andrews is a charity registered in Scotland: No SC013532

Table A.1: Ethics Approval for Pilot Study in Chapter 4

University of St Andrews | FOUNDED 1413

University Teaching and Research Ethics Committee

29 September 2015

Dear Adeola

Your ethical application has now been reviewed by the University Teaching and Research Ethics Committee (UTREC), alongside the following supporting documentation:

1. Ethical Application Form
2. Participant Information Sheets
3. Consent Forms
4. Debriefing Forms
5. Letter to Parents/children/headteacher
6. Questionnaires

I am pleased to confirm that UTREC has granted this application ethical approval and the particulars of the approved ethical application are as follows -

Approval Code:	CS11765	**Approved on:**	23 Sept 2015	**Approval expiry:**	23 Sept 2018
Project Title:	Investigating quality of service and quality of experience in Head Mounted Display (HMD-) based virtual reality systems.				
Researcher(s):	Adeola Fabola				
Supervisor(s):	Prof Alan Miller				

Approval is awarded for three years. Projects which have not commenced within two years of approval must be re-submitted for review by your School Ethics Committee, who may escalate your application to UTREC for review. If you are unable to complete your research within the 3 three year approval period, you are required to write to your School Ethics Committee Convener to request a discretionary extension of no greater than 6 months or to re-apply if directed to do so, and you should inform your School Ethics Committee when your project reaches completion.

If you make any changes to the project outlined in your approved ethical application form, you should inform your supervisor and seek advice on the ethical implications of those changes from the School Ethics Convener who may advise you to complete and submit an ethical amendment form for review.

Any adverse incident which occurs during the course of conducting your research must be reported immediately to the School Ethics Committee who will advise you on the appropriate action to be taken.

Approval is given on the understanding that you conduct your research as outlined in your application and in compliance with UTREC Guidelines and Policies (http://www.st-andrews.ac.uk/utrec/guidelinespolicies/). You are also advised to ensure that you procure and handle your research data within the provisions of the Data Provision Act 1998 and in accordance with any conditions of funding incumbent upon you.

If you have any questions in relation to this ethical approval then please do not hesitate to contact me.

Yours sincerely

Dr Emily Hearn, UTREC Secretary, on behalf of UTREC

cc Supervisor – Prof Alan Miller

University Teaching and Research Ethics Committee
Dr Emily Hearn (UTREC Secretary), Rm 106, College Gate, North Street, St Andrews, Fife, KY16 9AJ
T: 01334 462368 E: utrec@st-andrews.ac.uk
The University of St Andrews is a charity registered in Scotland: No SC013532

Table A.2: Ethics Approval for Chapter 4

University of St Andrews
Scotland's first university – 1413

University Teaching and Research Ethics Committee Sub-committee

12/08/2016
Adeola Fabola
School of Computer Science

Ethics Reference No: *Please quote this ref on all correspondence*	CS12329
Project Title:	Immersive Technologies for Heritage Exploration
Researchers Name(s):	Adeola Fabola
Supervisor(s):	Dr Alan Miller

Thank you for submitting your application which was considered at the Computer Science School Ethics Committee meeting on the 12/08/2016. The following documents were reviewed:

1. Ethical Application Form 12/08/2016
2. Participant Information Sheet 12/08/2016
3. Consent Form 12/08/2016
4. Debriefing Form 12/08/2016

The University Teaching and Research Ethics Committee (UTREC) approve this study from an ethical point of view. Please note that where approval is given by a School Ethics Committee that committee is part of UTREC and is delegated to act for UTREC.

Approval is given for three years. Projects, which have not commenced within two years of original approval, must be re-submitted to your School Ethics Committee.

You must inform your School Ethics Committee when the research has been completed. If you are unable to complete your research within the 3 three year validation period, you will be required to write to your School Ethics Committee and to UTREC (where approval was given by UTREC) to request an extension or you will need to re-apply.

Any serious adverse events or significant change which occurs in connection with this study and/or which may alter its ethical consideration must be reported immediately to the School Ethics Committee, and an Ethical Amendment Form submitted where appropriate.

Approval is given on the understanding that the 'Guidelines for Ethical Research Practice' https://www.st-andrews.ac.uk/utrec/guidelines/ are adhered to.

Yours sincerely

UNIVERSITY OF ST. ANDREWS SCHOOL OF COMPUTER SCIENCE

Convenor of the School Ethics Committee

Ccs Supervisor
School Ethics Committee

ethics-cs@st-andrews.ac.uk

Table A.3: Ethics Approval for Chapters 5 and 6

University Teaching and Research Ethics Committee

28 August 2017

Dear Adeola,

Thank you for submitting your ethical application which was considered at the School of Computer Science Ethics Committee meeting on 28th August when the following documents were reviewed:

1. Ethical Application Form
2. Participant Information Sheet
3. Consent Form
4. Debriefing Form

The School of Computer Science Ethics Committee has been delegated to act on behalf of the University Teaching and Research Ethics Committee (UTREC) and has granted this application ethical approval. The particulars relating to the approved project are as follows -

Approval Code:	CS13090	**Approved on:**	28.08.17	**Approval Expiry:**	28.08.22
Project Title:	VIRTUAL MUSEUMS FOR COMMUNITY ENGAGEMENT				
Researcher(s):	ADEOLA FABOLA				
Supervisor(s):	DR ALAN MILLER				

Approval is awarded for fiveyears. Projects which have not commenced within two years of approval must be re-submitted for review by your School Ethics Committee. If you are unable to complete your research within the five year approval period, you are required to write to your School Ethics Committee Convener to request a discretionary extension of no greater than 6 months or to re-apply if directed to do so, and you should inform your School Ethics Committee when your project reaches completion.

If you make any changes to the project outlined in your approved ethical application form, you should inform your supervisor and seek advice on the ethical implications of those changes from the School Ethics Convener who may advise you to complete and submit an ethical amendment form for review.

Any adverse incident which occurs during the course of conducting your research must be reported immediately to the School Ethics Committee who will advise you on the appropriate action to be taken.

Approval is given on the understanding that you conduct your research as outlined in your application and in compliance with UTREC Guidelines and Policies (http://www.st-andrews.ac.uk/utrec/guidelinespolicies/). You are also advised to ensure that you procure and handle your research data within the provisions of the Data Provision Act 1998 and in accordance with any conditions of funding incumbent upon you.

Yours sincerely

Alex Bain

Convener of the School Ethics Committee

ethics-cs@st-andrews.ac.uk

The University of St Andrews is a charity registered in Scotland: No SC013532

Table A.4: Ethics Approval for Chapter 8

APPENDIX B

QUESTIONNAIRES

The questionnaires used in the evaluation exercises undertaken as part of this research are shown in Tables. B.1, B.2, B.3, B.4, B.5 and B.6

USER EXPERIENCE QUESTIONNAIRE FOR ST ANDREWS CATHEDRAL VIRTUAL TOUR APP

Background Information (Please tick the appropriate column)		1	2	3	4	5
1	How good are your English skills?	Very Poor	Poor	Fair	Good	Very Good
2	How good are your IT skills?	Very Poor	Poor	Fair	Good	Very Good
3	How interested are you in History?	Very Low	Low	Some	High	Very High
4	Do you have previous experience with Virtual Reality (VR)?	None	Novice	Moderate	A lot	Expert
5	Do you have previous experience with Virtual Reality (VR) Headsets?	None	Novice	Moderate	A lot	Expert

Device Type (Please select the device you used to take part in the virtual tour)		
Google Cardboard	Samsung Gear VR	Screen-based Device

Experience Questionnaire (Please tick the appropriate column)		Strongly disagree	Probably disagree	Neither agree nor disagree	Probably agree	Strongly agree
	Statement	1	2	3	4	5
1	I think that this system is easy to use.					
2	I would recommend this virtual reality system for learning history.					
3	This system has changed how I think about the St Andrews Cathedral.					
4	I am now more interested in learning about local history.					
5	I felt like I was there in the virtual environment.					

Free Form Questions – Experience	
1	Please describe your experience with this system in three words: 1. 2. 3.
2	What do you like and what can be improved in the system?

Table B.1: User experience questionnaire used in Chapter 4

SCHOOL TEACHERS PARTICIPATION QUESTIONNAIRE

S/N	School Teachers Questionnaire
1	What is your impression of the virtual reality system for learning history?
2	What impact has this exercise had on the teaching and learning activities in your school?
3	Do you advocate for the continued use of these systems? If yes, how would you like to use these systems to support your teaching activities?

Table B.2: School teachers questionnaire used in Chapter 4

USER EXPERIENCE QUESTIONNAIRE FOR VIRTUAL CURING YARD APP

Background Information (Please tick the appropriate column)						
		1	2	3	4	5
1	How good are your IT skills?	Very Poor	Poor	Fair	Good	Very Good
2	How interested are you in History?	Very Low	Low	Some	High	Very High
3	Do you have previous experience with Virtual Reality (VR)?	None	Novice	Moderate	A lot	Expert
4	Do you have previous experience with Virtual Reality (VR) Headsets?	None	Novice	Moderate	A lot	Expert

Device Type (Please select the device you used to take part in the virtual tour)		
Google Cardboard	Samsung Gear VR	Screen-based Device

Experience Questionnaire (Please tick the appropriate column)						
		Strongly disagree	Probably disagree	Neither agree nor disagree	Probably agree	Strongly agree
	Statement	1	2	3	4	5
1	I think that this system is easy to use.					
2	I would recommend this virtual reality system for learning history.					
3	This system has changed how I think about Helmsdale.					
4	I am now more interested in learning about local history.					
5	I felt like I was there in the virtual environment.					

Free Form Questions – Experience	
1	Please describe your experience with this system in three words: 1. 2. 3.
2	What do you like and what can be improved in the system?

Table B.3: User experience questionnaire used in Chapter 5

EXPERT EVALUATION QUESTIONNAIRE

S/N	Free Form Questions – Expert Evaluation
1	**Why did you embark on this project? What did you hope to get out of it?**
2	**How do you see this being used in a museum? What role(s) do you think it will play?**
3	**Do you believe this system will add value to the museum visitors' experience? If yes, how? If no, why not?**
4	**What are the main challenges facing the use of such systems in museums?**

Table B.4: Expert evaluation questionnaire used in Chapter 5

USER EXPERIENCE QUESTIONNAIRE FOR REAL-TIME VIRTUAL TOUR APP

S/N	Free Form Questions – User Experience
1	**Is this system awesome or weird? Why?**
2	**Did you feel like you were in control of the system? Why?**
3	**Did you feel like this was a good social experience? Why?**
4	**Any other comments?**

Table B.5: User experience questionnaire used in Chapter 6

QUESTIONNAIRE FOR PICTS & PIXELS EXHIBIT

The aim of this study is to evaluate the user experience aspects of a digital exhibit on the Picts in the Perth area using a combination of virtual reality, 3D media and traditional media, as a project carried out by the School of Computer Science at the University of St Andrews. If you do not understand any of the terms used in the questionnaire, please ask the researcher for assistance.

Device Type (Please select the device(s) you used to take part in the virtual tour)			
VR Headset + Xbox controller	Screen + Xbox controller	Mobile VR Headset	Other

Experience Questionnaire (Please tick the appropriate column)						
		Strongly disagree	Probably disagree	Neither agree nor disagree	Probably agree	Strongly agree
	Statement	1	2	3	4	5
1	I think that this system is easy to use.					
2	I would recommend this system for learning history.					
3	This system has changed how I think about Perth					
4	I am now more interested in learning about local history.					
5	I felt like I was there in the virtual environment.					

Free Form Questions – Experience	
1	Please describe three things you learnt from this exhibition. 1. 2. 3.
2	What do you like, and what can be improved in the system?

Table B.6: User experience questionnaire used in Chapter 8

References

[1] Adeola Fabola, Alan Miller, and Richard Fawcett. Exploring the past with google cardboard. In *Digital Heritage, 2015*, volume 1, pages 277–284. IEEE, 2015.

[2] Adeola Fabola and Alan Miller. Virtual reality for early education: A study. In *International Conference on Immersive Learning*, pages 59–72. Springer, 2016.

[3] Adeola Fabola, Sarah Kennedy, Alan Miller, Iain Oliver, John McCaffery, Catherine Cassidy, Jo Clements, and Anna Vermehren. A virtual museum installation for time travel. In *International Conference on Immersive Learning*, pages 255–270. Springer, 2017.

[4] Adeola Fabola, Alan Miller, and Ishbel Duncan. Aerial virtual reality: Remote tourism with drones. In *Conference proceedings from the 7th European Immersive Education Summit*. iED, 2017.

[5] Catherine Cassidy, Adeola Fabola, and Alan Miller. A digital museum infrastructure for preserving community collections from climate change. In *Workshop, Long and Short Paper, and Poster Proceedings from the Third Immersive Learning Research Network Conference. Available: https://goo.gl/Lb81ik*, pages 170–177. iLRN, 2017.

[6] Catherine Anne Cassidy, Adeola Fabola, Elizabeth Rhodes, and Alan Miller. The making and evaluation of picts and pixels: Mixed exhibiting in the real and the unreal. In *International Conference on Immersive Learning*, pages 97–112. Springer, 2018.

[7] T. Pratchett. *I Shall Wear Midnight*. Discworld (Childrens). Doubleday Children's, 2010.

[8] Duncan F Cameron. The museum, a temple or the forum. *Curator: The Museum Journal*, 14(1):11–24, 1971.

[9] International Council of Museums (ICOM). International Council of Museums - Museum Definition. 2007. Accessed: 2018-03-09. Available: https://goo.gl/PtZ9he.

[10] Bloch, M. J. Reevaluating the ICOM Definition of the Museum: The Nonprofit Status of the Museum. 2004. Accessed: 2018-03-09. Available: https://goo.gl/VNmizo.

[11] Joseph Veach Noble. Museum manifesto. *Museum News*, 48(8):16–20, 1970.

[12] Peter VAN MENSCH. *Towards a methodology of museology. 1992*. PhD thesis, Tese (PHD)–Universidade de Zágreb, Zágreb, 1992.

[13] UNESCO. Tangible Cultural Heritage. 2017. Accessed: 2018-03-09. Available: https://goo.gl/UT96KD.

[14] UNESCO. Convention Concerning the Protection of the World Cultural and Natural Heritage. 1972. Accessed: 2018-03-09. Available: https://goo.gl/VFDzHz.

[15] UNESCO. What is Intangible Cultural Heritage?. 2017. Accessed: 2018-03-09. Available: https://goo.gl/SRg4Dh.

[16] Einar Bowitz and Karin Ibenholt. Economic impacts of cultural heritage - Research and perspectives. *Journal of Cultural Heritage*, 10(1):1–8, 2009.

[17] Dallen J Timothy, Gyan P Nyaupane, et al. *Cultural heritage and tourism in the developing world: A regional perspective*. Routledge, 2009.

[18] Emmanuel N Arinze. The role of the museum in society. *Lecture delivered at the National Museum, Georgetown, Guyana*, 17, 1999.

[19] Edward P Alexander, Mary Alexander, and Juilee Decker. *Museums in motion: An introduction to the history and functions of museums*. Rowman & Littlefield, 2017.

[20] Ze-Nian Li, Mark S Drew, and Jiangchuan Liu. *Fundamentals of multimedia*. Springer, 2004.

[21] Ashok Banerji and Ananda Mohan Ghosh. *Multimedia Technologies*. Tata McGraw Hill, 2010.

[22] Yehuda Kalay, Thomas Kvan, and Janice Affleck. *New heritage: New media and cultural heritage*. Routledge, 2007.

[23] Ralf Steinmetz and Klara Nahrstedt. Multimedia applications. In *Multimedia applications*, pages 197–214. Springer, 2004.

[24] Ranjan Parekh et al. *Principles of multimedia*. Tata McGraw-Hill Education, 2006.

[25] Mahesh Raisinghani. *Revolutionizing Education through Web-Based Instruction*. IGI Global, 2016.

[26] A. E Foni, G Papagiannakis, and N Magnenat-Thalmann. A taxonomy of visualization strategies for cultural heritage applications. *Journal on Computing and Cultural Heritage (JOCCH)*, 3(1):1, 2010.

[27] Oculus. Development Kits are Shipping. 2013. Accessed: 2018-03-09. Available: https://goo.gl/Baig3B.

[28] HTC. HTC Vive Update. 2015. Accessed: 2018-03-09. Available: https://goo.gl/Jvavsg.

[29] WebVR. Bringing Virtual Reality to the Web. 2017. Accessed: 2018-03-09. Available: https://goo.gl/whUVGz.

[30] Fabio Bruno, Stefano Bruno, Giovanna De Sensi, Maria Laura Luchi, Stefania Mancuso, and Maurizio Muzzupappa. From 3D reconstruction to virtual reality: A complete methodology for digital archaeological exhibition. *Journal of Cultural Heritage*, 11(1):42–49, 2010.

[31] Nina Simon. *The participatory museum*. Museum 2.0, 2010.

[32] Fredrick Nsibambi Ssenyonga. The emerging role of community museums in uganda: The need for capacity building among managers. *Museum International*, 68(1-2):125–129, 2016.

[33] David Chittenden. Commentary: Roles, opportunities, and challenges–science museums engaging the public in emerging science and technology. *Journal of Nanoparticle Research*, 13(4):1549–1556, 2011.

[34] Katja Lindqvist. Museum finances: Challenges beyond economic crises. *Museum Management and Curatorship*, 27(1):1–15, 2012.

[35] Osipova E. Lafrenz Samuels K. Markham, A. and A. Caldas. World heritage and tourism in a changing climate. *United Nations Environment Programme, Nairobi, Kenya and United Nations Educational, Scientific and Cultural Organization, Paris, France*, 2016.

[36] Helen Phillips. The capacity to adapt to climate change at heritage sitesâĂŤthe development of a conceptual framework. *Environmental Science & Policy*, 47:118–125, 2015.

[37] Zara Huijbregts, RP Kramer, MHJ Martens, AWM Van Schijndel, and HL Schellen. A proposed method to assess the damage risk of future climate change to museum objects in historic buildings. *Building and Environment*, 55:43–56, 2012.

[38] Paul Lankester and Peter Brimblecombe. The impact of future climate on historic interiors. *Science of the Total Environment*, 417:248–254, 2012.

[39] Michele Knobel. *Digital literacies: Concepts, policies and practices*, volume 30. Peter Lang, 2008.

[40] Gordon E. Moore. Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86(1):82–85, 1998.

[41] Marc Goodman and Parag Khanna. The power of moore's law in a world of geotechnology. *The National Interest*, 123:64–73, 2013.

[42] Chris A. MacK. Fifty years of Moore's law. *IEEE Transactions on Semiconductor Manufacturing*, 24(2):202–207, 2011.

[43] Ivan E. Sutherland. The ultimate display. *Proceedings of the Congress of the Internation Federation of Information Processing (IFIP)*, 21(3):506–508, 1996.

[44] Ivan E. Sutherland. A head-mounted three dimensional display. *Proceedings of the AFIPS '68 (Fall, part I)*, page 757, 1968.

[45] Frank Biocca and Mark R Levy. *Communication in the age of virtual reality*. Routledge, 2013.

[46] David Brewster. *The Stereoscope; Its History, Theory and Construction, with Its Application to the Fine and Useful Arts and to Education, Etc*. John Murray, 1856.

[47] Jon Peddie. The history of visual magic in computers. *Springer, ISBN*, 978:1–4471, 2013.

[48] William R Sherman and Alan B Craig. *Understanding virtual reality: Interface, application, and design*. Elsevier, 2002.

[49] C Burdea Grigore and P Coiffet. *Virtual reality technology*. London: Wiley-Interscience, 1994.

[50] Alan B Craig, William R Sherman, and Jeffrey D Will. *Developing virtual reality applications: Foundations of effective design*. Morgan Kaufmann, 2009.

[51] Shuhei Yoshida. Introducing: Project Morpheus. 2014. Accessed: 2018-03-09. Available: https://goo.gl/K22CYD.

[52] Fred Dutton. PlayStation VR is the official name for the PS4 virtual reality system. 2015. Accessed: 2018-03-09. Available: https://goo.gl/okawb8.

[53] Alex Kipman. Announcing Microsoft HoloLens Development Edition open for pre-order, shipping March 30. 2016. Accessed: 2018-03-09. Available: https://goo.gl/Vva3Az.

[54] Oculus. Oculus Rift Pre-Orders Now Open, First Shipments March 28. 2016. Accessed: 2018-03-09. Available: https://goo.gl/4ceJhP.

[55] Oculus. First Look at the Rift, Shipping Q1 2016. 2015. Accessed: 2018-03-09. Available: https://goo.gl/MfaM6i.

[56] Vive Team. Updates: Vive Availability and More. 2016. Accessed: 2018-03-09. Available: https://goo.gl/GDnVQk.

[57] Oculus. Samsung Gear VR now available for pre-orders at $99. 2015. Accessed: 2018-03-09. Available: https://goo.gl/TUAC6Q.

[58] Clay Bavor. (Un)folding a virtual journey with Google Cardboard. 2016. Accessed: 2018-03-09. Available: https://goo.gl/qKpuuT.

[59] Dan MacIsaac et al. Google cardboard: A virtual reality headset for $10? *The Physics Teacher*, 53(2):125–125, 2015.

[60] Paul Lamkin. Best VR headsets 2017: HTC Vive, Oculus, PlayStation VR compared. 2017 Accessed: 2018-03-09. Available: https://goo.gl/bc5ESr.

[61] The Guardian. Drone scores a first by successfully delivering package in Nevada town. 2016. Accessed: 2018-03-09. Available: https://goo.gl/B8kAAH.

[62] Dane Bamburry. Drones: Designed for product delivery. *Design Management Review*, 26(1):40–48, 2015.

[63] Walsh, Ryan and Falesch, Alex J. Landing Pad For Unmanned Aerial Vehicle Delivery. 2014. Accessed: 2018-03-09. US Patent App. 14/587,828. Available: https://goo.gl/YHCuXr.

[64] Raffaello D'Andrea. Guest editorial can drones deliver? *IEEE Transactions on Automation Science and Engineering*, 11(3):647–648, 2014.

[65] Raspberry Pi Foundation. Raspberry Pi 3 Model B. 2017. Accessed: 2018-03-09. Available: https://goo.gl/YYhsnA.

[66] Jacob Poushter. Smartphone Ownership and Internet Usage Continues to Climb in Emerging Economies. 2016. Accessed: 2018-03-09. Available: https://goo.gl/R1XhCI.

[67] Krista Lofgren. Video Game Trends and Statistics – Who's Playing What and Why?. 2017. Accessed: 2018-03-09. Available: https://goo.gl/AKnuVo.

[68] Linda Candy. Practice Based Research:A Guide. *CCS report*, V1(0), 2006.

[69] John W Creswell. *Qualitative inquiry and research design: Choosing among five traditions*. Sage publications, 1998.

[70] P Milgram, H Takemura, a Utsumi, and F Kishino. Mixed Reality (MR) Reality-Virtuality (RV) Continuum. *Systems Research*, 2351(Telemanipulator and Telepresence Technologies):282–292, 1994.

[71] Meredith Bricken. Virtual reality learning environments: potentials and challenges. *ACM SIGGRAPH Computer Graphics*, 25(3):178–184, 1991.

[72] Tassos A Mikropoulos and Antonis Natsis. Educational virtual environments: A ten-year review of empirical research (1999–2009). *Computers & Education*, 56(3):769–780, 2011.

[73] Zhigeng Pan, Adrian David Cheok, Hongwei Yang, Jiejie Zhu, and Jiaoying Shi. Virtual reality and mixed reality for virtual learning environments. *Computers & Graphics*, 30(1):20–28, 2006.

[74] Jeffrey Jacobson. *Ancient Architecture in Virtual Reality - Does Immersion Really Aid Learning?* PhD thesis, University of Pittsburgh, 2008.

[75] John McCaffery, Alan Miller, and Colin Allison. Extending the use of virtual worlds as an educational platform-network island: An advanced learning environment for teaching internet routing algorithms. In *CSEDU (1)*, pages 279–284, 2011.

[76] Indika Perera, Colin Allison, J Ross Nicoll, and Thomas Sturgeon. Towards successful 3d virtual learning-a case study on teaching human computer interaction. In *Internet Technology and Secured Transactions, 2009. ICITST 2009. International Conference for*, pages 1–6. IEEE, 2009.

[77] S Kennedy, R Fawcett, a Miller, L Dow, R Sweetman, a Field, a Campbell, I Oliver, J. McCaffery, and C Allison. Exploring canons & cathedrals with Open Virtual Worlds: The recreation of St Andrews Cathedral, St Andrews day, 1318. *2013 Digital Heritage International Congress (DigitalHeritage)*, pages 273–280, oct 2013.

[78] Teresa Monahan, Gavin McArdle, and Michela Bertolotto. Virtual reality for collaborative e-learning. *Computers & Education*, 50(4):1339–1353, 2008.

[79] Chris Davies, Alan Miller, and Colin Allison. Mobile cross reality for cultural heritage. In *Digital Heritage International Congress (DigitalHeritage), 2013*, volume 1, pages 331–338. IEEE, 2013.

[80] Anne-Cecilie Haugstvedt and John Krogstie. Mobile augmented reality for cultural heritage: A technology acceptance study. In *Mixed and Augmented Reality (ISMAR), 2012 IEEE International Symposium on*, pages 247–255. IEEE, 2012.

[81] Tanguy Coenen, Lien Mostmans, and Kris Naessens. Museus: Case study of a pervasive cultural heritage serious game. *Journal on Computing and Cultural Heritage (JOCCH)*, 6(2):8, 2013.

[82] Kristoffer Getchell, Alan Miller, J Ross Nicoll, Rebecca Sweetman, and Colin Allison. Games methodologies and immersive environments for virtual fieldwork. *IEEE Transactions on Learning Technologies*, 3(4):281–293, 2010.

[83] Kristoffer Getchell, Iain Oliver, Alan Miller, and Colin Allison. Metaverses as a platform for game based learning. In *Advanced Information Networking and Applications (AINA), 2010 24th IEEE International Conference on*, pages 1195–1202. IEEE, 2010.

[84] Hsiu-Mei Huang, Ulrich Rauch, and Shu-Sheng Liaw. Investigating learners' attitudes toward virtual reality learning environments: Based on a constructivist approach. *Computers & Education*, 55(3):1171–1182, 2010.

[85] Chairi Kiourt, Anestis Koutsoudis, and George Pavlidis. Dynamus: a fully dynamic 3d virtual museum framework. *Journal of Cultural Heritage*, 22:984–991, 2016.

[86] Marcos Balsa Rodriguez, Marco Agus, Fabio Bettio, Fabio Marton, and Enrico Gobbetti. Digital mont'e prama: 3d cultural heritage presentations in museums and anywhere. In *Digital Heritage, 2015*, volume 2, pages 557–564. IEEE, 2015.

[87] Stanislav Dimchev Kostadinov and Tzvetomir Ivanov Vassilev. Mobile virtual 3d model of a medieval town. *International Journal on Information Technologies & Security*, 5(4):13–18, 2013.

[88] Elizabeth Tait, Richard Laing, Andy Grinnall, Simon Burnett, and John Isaacs. (re) presenting heritage: Laser scanning and 3d visualisations for cultural resilience and community engagement. *Journal of information science*, 42(3):420–433, 2016.

[89] S Sylaiou, F Liarokapis, K Kotsakis, and P Patias. Virtual museums, a survey and some issues for consideration. *Journal of Cultural Heritage*, 10:520–528, 2009.

[90] Antonio Frisoli, Gunnar Jansson, Massimo Bergamasco, and Celine Loscos. Evaluation of the pure-form haptic displays used for exploration of works of art at museums. *Proc. World Haptics*, 5, 2005.

[91] Marcello Carrozzino, Chiara Evangelista, A Scucces, Franco Tecchia, G Tennirelli, and Massimo Bergamasco. The virtual museum of sculpture. In *Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts*, pages 100–106. ACM, 2008.

[92] J McCaffery and A Miller. The Virtual Museums of Caen: A case study on modes of representation of digital historical content. *2015 Digital Heritage*, 2015.

[93] William R Sherman and Alan B Craig. Literacy in virtual reality: a new medium. *ACM SIGGRAPH Computer Graphics*, 29(4):37–42, 1995.

[94] Iis P. Tussyadiah and Daniel R. Fesenmaier. Mediating Tourist Experiences. Access to Places via Shared Videos. *Annals of Tourism Research*, 36(1):24–40, 2009.

[95] Alberto Guarnieri, Francesco Pirotti, and Antonio Vettore. Cultural heritage interactive 3D models on the web: An approach using open source and free software. *Journal of Cultural Heritage*, 11(3):350–353, 2010.

[96] Massimiliano Albanese, Antonio d'Acierno, Vincenzo Moscato, Fabio Persia, and Antonio Picariello. A multimedia semantic recommender system for cultural heritage applications. In *Semantic Computing (ICSC), 2011 Fifth IEEE International Conference on*, pages 403–410. IEEE, 2011.

[97] Adeola Fabola, Chris Davies, Sarah Kennedy, Alan Miller, and Colin Allison. Mobile exploration of medieval st andrews. In *European Summit on Immersive Education*, pages 114–124. Springer, 2014.

[98] C Davies, A Miller, and R Fawcett. Mobile onsite exploration of parallel realities with Oculus Rift. *2015 Digital Heritage*, 1:301–304, 2015.

[99] J Tredinnick and P Richens. A fulldome interactive visitor experience a novel approach to delivering interactive virtual heritage experiences to group audiences in fulldome projection spaces, evaluated through spatial awareness and emotional response. In *Digital Heritage, 2015*, volume 1, pages 413–414. IEEE, 2015.

[100] Wan-yen Lo, Yu-pao Tsai, Chien-wei Chen, and Yi-ping Hung. Stereoscopic Kiosk for Virtual Museum. *Ics*, pages 1084–1090, 2004.

[101] International Telecommunications Union (ITU). Definition of Quality of Experience (QoE). 2007. Accessed: 2018-03-09. Available: https://goo.gl/EXU4KZ.

[102] Lutz Goldmann, Francesca De Simone, and Touradj Ebrahimi. A comprehensive database and subjective evaluation methodology for quality of experience in stereoscopic video. In *IS&T/SPIE Electronic Imaging*, pages 75260S–75260S. International Society for Optics and Photonics, 2010.

[103] Liyuan Xing, Junyong You, Touradj Ebrahimi, and Andrew Perkis. Estimating quality of experience on stereoscopic images. In *Intelligent Signal Processing and Communication Systems (ISPACS), 2010 International Symposium on*, pages 1–4. IEEE, 2010.

[104] Kandaraj Piamrat, Cesar Viho, Jean-Marie Bonnin, and Adlen Ksentini. Quality of experience measurements for video streaming over wireless networks. In *Information Technology: New Generations, 2009. ITNG'09. Sixth International Conference on*, pages 1184–1189. IEEE, 2009.

[105] Liyuan Xing, Junyong You, Touradj Ebrahimi, and Andrew Perkis. Factors impacting quality of experience in stereoscopic images. In *IS&T/SPIE Electronic Imaging*, pages 786304–786304. International Society for Optics and Photonics, 2011.

[106] Jukka Häkkinen, Takashi Kawai, Jari Takatalo, Tuomas Leisti, Jenni Radun, Anni Hirsaho, and Göte Nyman. Measuring stereoscopic image quality experience with interpretation based quality methodology. In *Electronic Imaging 2008*, pages 68081B–68081B. International Society for Optics and Photonics, 2008.

[107] WA IJsselsteijn, DG Bouwhuis, J Freeman, and H Ridder. P-15: Presence as an experiential metric for 3-d display evaluation. In *SID Symposium Digest of Technical Papers*, volume 33, pages 252–255. Wiley Online Library, 2002.

[108] Benjamin Petry and Jochen Huber. Towards effective interaction with omnidirectional videos using immersive virtual reality headsets. In *Proceedings of the 6th Augmented Human International Conference*, pages 217–218. ACM, 2015.

[109] Robert Anderson, David Gallup, Jonathan T Barron, Janne Kontkanen, Noah Snavely, Carlos Hernández, Sameer Agarwal, and Steven M Seitz. Jump: virtual reality video. *ACM Transactions on Graphics (TOG)*, 35(6):198, 2016.

[110] Robert E Wendrich, Kris-Howard Chambers, Wadee Al-Halabi, Eric J Seibel, Olaf Grevenstuk, David Ullman, and Hunter G Hoffman. Hybrid design tools in a social virtual reality using networked oculus rift: a feasibility study in remote real-time interaction. In

ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, pages 21–24. American Society of Mechanical Engineers, 2016.

[111] Facebook. Filming, Stitching and Editing 360 Video. 2017. Accessed: 2018-03-09. Available: https://goo.gl/ceax4Q.

[112] Sanjeev Verma. A new way to see and share your world with 360-degree video. 2015. Accessed: 2018-03-09. Available: https://goo.gl/j8wPJa.

[113] K Boudjit and Chevif Larbes. Detection and implementation autonomous target tracking with a quadrotor ar. drone. In *Informatics in Control, Automation and Robotics (ICINCO), 2015 12th International Conference on*, volume 2, pages 223–230. IEEE, 2015.

[114] Hyunwoong Choi, Mitchell Geeves, Bilal Alsalam, and Felipe Gonzalez. Open source computer-vision based guidance system for uavs on-board decision making. In *Aerospace Conference, 2016 IEEE*, pages 1–5. IEEE, 2016.

[115] Francisca Vasconcelos and Nuno Vasconcelos. Person-following uavs. In *Applications of Computer Vision (WACV), 2016 IEEE Winter Conference on*, pages 1–9. IEEE, 2016.

[116] MC Harvey, JV Rowland, and KM Luketina. Drone with thermal infrared camera provides high resolution georeferenced imagery of the waikite geothermal area, new zealand. *Journal of Volcanology and Geothermal Research*, 325:61–69, 2016.

[117] Valerio Baiocchi, Donatella Dominici, Maria Vittoria Milone, and Martina Mormile. Development of a software to plan uavs stereoscopic flight: An application on post earthquake scenario in l'aquila city. In *International Conference on Computational Science and Its Applications*, pages 150–165. Springer, 2013.

[118] Valentina Gatteschi, Fabrizio Lamberti, Gianluca Paravati, Andrea Sanna, Claudio Demartini, Alberto Lisanti, and Giorgio Venezia. New frontiers of delivery services using drones: A prototype system exploiting a quadcopter for autonomous drug shipments. In *Computer Software and Applications Conference (COMPSAC), 2015 IEEE 39th Annual*, volume 2, pages 920–927. IEEE, 2015.

[119] Javier Fernández-Lozano and Gabriel Gutiérrez-Alonso. Improving archaeological prospection using localized uavs assisted photogrammetry: An example from the roman gold district of the eria river valley (nw spain). *Journal of Archaeological Science: Reports*, 5:509–520, 2016.

[120] AT Mozas-Calvache, JL Pérez-García, FJ Cardenal-Escarcena, E Mata-Castro, and J Delgado-García. Method for photogrammetric surveying of archaeological sites with light aerial platforms. *Journal of Archaeological Science*, 39(2):521–530, 2012.

[121] Marijn Hendrickx, Wouter Gheyle, Johan Bonne, Jean Bourgeois, Alain De Wulf, and Rudi Goossens. The use of stereoscopic images taken from a microdrone for the documentation of heritage–an example from the tuekta burial mounds in the russian altay. *Journal of Archaeological Science*, 38(11):2968–2978, 2011.

[122] Filiberto Chiabrando, Francesco Nex, Dario Piatti, and Fulvio Rinaudo. Uav and rpv systems for photogrammetric surveys in archaelogical areas: two tests in the piedmont region (italy). *Journal of Archaeological Science*, 38(3):697–710, 2011.

[123] David Mirk and Helmut Hlavacs. Virtual tourism with drones: Experiments and lag compensation. In *Proceedings of the First Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use*, pages 45–50. ACM, 2015.

[124] Chenglei Wu, Zhi Wang, and Shiqiang Yang. Drone streaming with wi-fi grid aggregation for virtual tour. *arXiv preprint arXiv:1605.09486*, 2016.

[125] Sarah Kenderdine. Speaking in rama: Panoramic vision in cultural heritage visualization. In Fiona Cameron and Sarah Kenderdine, editors, *Theorizing Digital Cultural Heritage: A Critical Discourse*, chapter 15, pages 301–331. MIT Press, Cambridge, MA, England, 2007.

[126] Alan B Craig. *Understanding augmented reality: Concepts and applications*. Newnes, 2013.

[127] Oliver Grau. *Virtual Art: from illusion to immersion*. MIT press, 2003.

[128] Jean Sargent. You don't have to be Dr Who to travel back in time. 2014. Accessed: 2018-03-09. Available: https://goo.gl/fNjH8d.

[129] John Woodwark. Reconstructing history with computer graphics. *IEEE Computer Graphics and Applications*, 11(1):18–20, 1991.

[130] Fiona Cameron. Beyond the cult of the replicant: Museums and historical digital objects – traditional concerns, new discourses. In Fiona Cameron and Sarah Kenderdine, editors, *Theorizing Digital Cultural Heritage: A Critical Discourse*, chapter 3, pages 49–75. MIT Press, Cambridge, MA, England, 2007.

[131] George Basalla. Museums and technological utopianism. *Curator: The Museum Journal*, 17(2):105–118, 1974.

[132] British Museum. British Museum - About. 2017. Accessed: 2018-03-09. Available: https://goo.gl/XA5QFX.

[133] Sketchfab. Sketchfab - About. 2017. Accessed: 2018-03-09. Available: https://goo.gl/YvxJ6D.

[134] Thomas Flynn. Teaching Museums How to Produce 3D Content. 2017. Accessed: 2018-03-09. Available: https://goo.gl/Z8WdrQ.

[135] Megan Hancock. Museums and 3d printing: More than a workshop novelty, connecting to collections and the classroom. *Bulletin of the Association for Information Science and Technology*, 42(1):32–35, 2015.

[136] meSch project. 3D PRINTING IN MUSEUMS (WORKSHOP). 2017. Accessed: 2018-03-09. Available: https://goo.gl/LbWQNG.

[137] Natalia Krasnodebska. Sketchfab Supports Museums and Cultural Institutions. 2015. Accessed: 2018-03-09. Available: https://goo.gl/gwo5gk.

[138] Alban Denoyel. Facebook now supports Sketchfab 3D embeds. 2015. Accessed: 2018-03-09. Available: https://goo.gl/R9cQv4.

[139] Corentin. Twitter now supports Sketchfab 3D embeds. 2016. Accessed: 2018-03-09. Available: https://goo.gl/CwgLrs.

[140] Andrea Witcomb. The materiality of virtual technologies: A new approach to thinking about the impact of multimedia in museums. In Fiona Cameron and Sarah Kenderdine, editors, *Theorizing Digital Cultural Heritage: A Critical Discourse*, chapter 2, pages 35–48. MIT Press, Cambridge, MA, England, 2007.

[141] Geoffrey Lewis. The response of museums to the web. *Archives of MUSEUML@ HOME. EASE. LSOFT. COM Museum discussion list.*, 10:17, 1996.

[142] Werner Schweibenz. The "virtual museum": New perspectives for museums to present objects and information using the internet as a knowledge base and communication system. *ISI*, 34:185–200, 1998.

[143] David Bearman. Interactive multimedia in museums. *Studies in Multimedia: State-of-the-Art Solutions in Multimedia and Hypermedia*, pages 121–137, 1991.

[144] Werner Schweibenz. The Development of Virtual Museums. 2004. Accessed: 2018-03-09. Available: https://goo.gl/jemHpU.

[145] Filippo Costalli, Luisa Marucci, Giulio Mori, and Fabio Paterno. Design criteria for usable web-accessible virtual environments. In *ICHIM (1)*, pages 413–426, 2001.

[146] Ben Booth. Understanding the information needs of visitors to museums. *Museum Management and Curatorship*, 17(2):139–157, 1998.

[147] Dimitrios Charitos, George Lepouras, Costas Vassilakis, Vivi Katifori, Anna Charissi, and Leda Halatsi. Designing a virtual museum within a museum. In *Virtual reality, archeology, and cultural heritage: Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, volume 28, pages 284–284, 2001.

[148] Ross J Loomis, Steven M Elias, and Marcella Wells. Website availability and visitor motivation: An evaluation study for the colorado digitization project. *Unpublished Report. Fort Collins, CO: Colorado.*, 2003.

[149] S Sylaiou, F Liarokapis, L Sechidis, P Patias, and O Georgoula. Virtual museums: First results of a survey on methods and tools. In *Proceedings of the CIPA and the ISPRS Conference, Torino, Italy*, pages 1138–1143, 2005.

[150] Beng-Kiang Tan and Hafizur Rahaman. Virtual heritage: Reality and criticism. In *CAAD futures*, pages 143–156, 2009.

[151] Maria Roussou. Virtual heritage: from the research lab to the broad public. *Bar International Series*, 1075:93–100, 2002.

[152] Maria Roussou. Immersive interactive virtual reality in the museum. *Proc. of TiLE (Trends in Leisure Entertainment)*, 2001.

[153] Alonzo C Addison. Virtual heritage: technology in the service of culture. In *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 343–354. ACM, 2001.

[154] P Milgram and H Colquhoun. A Taxonomy of Real and Virtual World Display Integration. *Mixed Reality-Merging Real and Virtual Worlds*, pages 5–30, 1999.

[155] Steven Warburton. Second life in higher education: Assessing the potential for and the barriers to deploying virtual worlds in learning and teaching. *British journal of educational technology*, 40(3):414–426, 2009.

[156] Neal E Seymour, Anthony G Gallagher, Sanziana A Roman, Michael K O'brien, Vipin K Bansal, Dana K Andersen, and Richard M Satava. Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Annals of surgery*, 236(4):458, 2002.

[157] Anthony G Gallagher, E Matt Ritter, Howard Champion, Gerald Higgins, Marvin P Fried, Gerald Moses, C Daniel Smith, and Richard M Satava. Virtual reality simulation for the operating room: proficiency-based training as a paradigm shift in surgical skills training. *Annals of surgery*, 241(2):364, 2005.

[158] Daniel A Guttentag. Virtual reality: Applications and implications for tourism. *Tourism Management*, 31(5):637–651, 2010.

[159] Juan A Barceló, Maurizio Forte, and Donald H Sanders. *Virtual reality in archaeology*. ArchaeoPress Oxford, 2000.

[160] Tony Hall, Luigina Ciolfi, Liam Bannon, Mike Fraser, Steve Benford, John Bowers, Chris Greenhalgh, Sten-Olof Hellström, Shahram Izadi, Holger Schnädelbach, et al. The visitor as virtual archaeologist: explorations in mixed reality technology to enhance educational and social interaction in the museum. In *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 91–96. ACM, 2001.

[161] Zakiah Noh, Mohd Shahrizal Sunar, and Zhigeng Pan. A review on augmented reality for virtual heritage system. In *International Conference on Technologies for E-Learning and Digital Entertainment*, pages 50–61. Springer, 2009.

[162] Peter Dawson, Richard Levy, and Natasha Lyons. "breaking the fourth wall": 3d virtual worlds as tools for knowledge repatriation in archaeology. *Journal of Social Archaeology*, 11(3):387–402, 2011.

[163] Fausta Fiorillo, Belén Jiménez Fernández-Palacios, Fabio Remondino, and Salvatore Barba. 3d surveying and modelling of the archaeological area of paestum, italy. *Virtual Archaeology Review*, 4(8):55–60, 2013.

[164] John Cosmas, Take Itegaki, Damian Green, Edward Grabczewski, Fred Weimer, Luc Van Gool, Alexy Zalesny, Desi Vanrintel, Franz Leberl, Markus Grabner, et al. 3d murale: a multimedia system for archaeology. In *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 297–306. ACM, 2001.

[165] B. Thomas, B. Close, J. Donoghue, J. Squires, P. De Bondi, M. Morris, and W. Piekarski. ARQuake: an outdoor/indoor augmented reality first person application. *Digest of Papers. Fourth International Symposium on Wearable Computers*, pages 139–146, 2000.

[166] M Zoellner, A Pagani, Y Pastarmov, H Wuest, and D Stricker. Reality Filtering : A Visual Time Machine in Augmented Reality. *International Symposium on Virtual Reality Archaeology and Cultural Heritage (VAST)*, 24:137–143, 2008.

[167] BBC. Pokemon Go unleashed in the UK. 2016. Accessed: 2018-03-09. Available: https://goo.gl/sR1ceh.

[168] J Lifton and J. A. Paradiso. Dual reality: Merging the Real and Virtual. 2009.

[169] CJ Davies. Parallel Reality : Tandem Exploration of Real and Virtual Environments CJ Davies. (June), 2015.

[170] Hsin-Kai Wu, Silvia Wen-Yu Lee, Hsin-Yi Chang, and Jyh-Chong Liang. Current status, opportunities and challenges of augmented reality in education. *Computers & Education*, 62:41–49, 2013.

[171] M Lombard, T. B Ditton, D Crane, B Davis, G Gil-Egui, K Horvath, and J Rossman. Measuring presence: A literature-based approach to the development of a standardized paper-and-pencil instrument. *Presence 2000: The Third International Workshop on Presence*, page 13, 2000.

[172] W. A IJsselsteijn, H de Ridder, J Freeman, and S. E Avons. Presence: concept, determinants, and measurement. In *Electronic Imaging*, pages 520–529. International Society for Optics and Photonics, 2000.

[173] T Schubert, F Friedmann, and H Regenbrecht. The Experience of Presence: Factor Analytic Insights. *Presence: Teleoperators and Virtual Environments*, 10(3):266–281, 2001.

[174] M Slater and S Wilbur. A Framework for Immersive Virtual Environments (FIVE): Speculations on the Role of Presence in Virtual Environments. pages 1–20, 1997.

[175] M Slater and M Usoh. Representation Systems, Perceptual Position and Presence in Virtual Environments. *Presence-Teleoperators and Virtual Environments*, 2(3):221–234, 1994.

[176] M Slater. Measuring presence: A response to the Witmer and Singer presence questionnaire. pages 1–13, 1999.

[177] B. G Witmer and M. J Singer. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and virtual environments*, 7(3):225–240, 1998.

[178] E Lee and K Wong. A Review of Using Virtual Reality for Learning. *Transactions on Edutainment I. Lecture Notes in Computer Science*, 5080/2008:231–241, 2008.

[179] Fernando Loizides, Alexios El Kater, Constantinos Terlikas, Andreas Lanitis, and Despina Michael. Presenting cypriot cultural heritage in virtual reality: A user evaluation. In *Euro-Mediterranean Conference*, pages 572–579. Springer, 2014.

[180] F Rasheed, P Onkar, and M Narula. Immersive virtual reality to enhance the spatial awareness of students. In *Proceedings of the 7th International Conference on HCI, IndiaHCI 2015*, pages 154–160. ACM, 2015.

[181] Adri Gabriel Sooai, Surya Sumpeno, and Mauridhi Hery Purnomo. User perception on 3d stereoscopic cultural heritage ancient collection. In *Proceedings of the The 2th International Conference in HCI and UX on Indonesia 2016*, pages 112–119. ACM, 2016.

[182] S Smith, S Cutchin, A Rockwood, A Saad, N. G Smith, and T. E Levy. Demo paper: virtual and immersive experience presentation of cultural heritage sites. In *Virtual Systems and Multimedia (VSMM), 2012 18th International Conference on*, pages 645–648. IEEE, 2012.

[183] Daniel Acevedo, Eileen Vote, David H Laidlaw, and Martha S Joukowsky. Archaeological data visualization in vr: Analysis of lamp finds at the great temple of petra, a case study. In *Proceedings of the conference on Visualization'01*, pages 493–496. IEEE Computer Society, 2001.

[184] Athanasios Gaitatzes, Dimitrios Christopoulos, and Maria Roussou. Reviving the past: cultural heritage meets virtual reality. In *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 103–110. ACM, 2001.

[185] Sarah Kenderdine. "pure land": Inhabiting the mogao caves at dunhuang. *Curator: The Museum Journal*, 56(2):199–218, 2013.

[186] Dawei Li, Mooi-Choo Chuah, and Li Tian. Demo: lehigh explorer augmented campus tour (lact). In *Proceedings of the 2014 workshop on Mobile augmented reality and robotic technology-based systems*, pages 15–16. ACM, 2014.

[187] Mila Koeva, Mila Luleva, and Plamen Maldjanski. Integrating spherical panoramas and maps for visualization of cultural heritage objects using virtual reality technology. *Sensors*, 17(4):829, 2017.

[188] Miri Barak, Orit Herscoviz, Zvia Kaberman, and Yehudit J Dori. Mosaica: A web-2.0 based system for the preservation and presentation of cultural heritage. *Computers & Education*, 53(3):841–852, 2009.

[189] Oculus. Rift + Touch. 2017. Accessed: 2018-03-09. Available: https://goo.gl/XjcL23.

[190] Google. Google Cardboard. 2017. Accessed: 2018-03-09. Available: https://goo.gl/Q2oEuu.

[191] Helena Rua and Pedro Alvito. Living the past: 3d models, virtual reality and game engines as tools for supporting archaeology and the reconstruction of cultural heritage–the case-study of the roman villa of casal de freiria. *Journal of Archaeological Science*, 38(12):3296–3308, 2011.

[192] Andrea Gaucci, Simone Garagnani, and Anna Maria Manferdini. Reconstructing the lost reality archaeological analysis and transmedial technologies for a perspective of virtual reality in the etruscan city of kainua. In *Digital Heritage, 2015*, volume 2, pages 227–234. IEEE, 2015.

[193] Manjula Patel, Martin White, Krzysztof Walczak, and Partick Sayd. Digitisation to presentation: building virtual museum exhibitions. *Vision, Video and Graphics 2003*, 2003.

[194] Eva Pietroni, Christie Ray, Claudio Rufa, Daniel Pletinckx, and Iefke Van Kampen. Natural interaction in vr environments for cultural heritage and its impact inside museums: The etruscanning project. In *Virtual Systems and Multimedia (VSMM), 2012 18th International Conference on*, pages 339–346. IEEE, 2012.

[195] Dennis Tsichritzis, Simon J Gibbs, et al. Virtual museums and virtual realities. In *ICHIM*, pages 17–25. Citeseer, 1991.

[196] Paolo Paolini, Thimoty Barbieri, Paolo Loiudice, Francesca Alonzo, Marco Zanti, and Giuliano Gaia. Visiting a museum together: how to share a visit to a virtual world. *Journal of the American society for information science*, 51(1):33–38, 2000.

[197] Maria José Aurindo and Conceição Machado. Muvitur®(virtual museum of tourism): a new approach to tourism history. *Journal of Tourism History*, pages 1–10, 2017.

[198] Chairi Kiourt, Anestis Koutsoudis, Fotis Arnaoutoglou, Georgia Petsa, Stella Markantonatou, and George Pavlidis. A dynamic web-based 3d virtual museum framework based on open data. In *Digital Heritage, 2015*, volume 2, pages 647–650. IEEE, 2015.

[199] Galena Kostoska, Marcos Baez, Florian Daniel, and Fabio Casati. Virtual, remote participation in museum visits by older adults: a feasibility study. In *Proceedings of the 8th International Conference on Personalized Access to Cultural Heritage-Volume 1352*, pages 10–13. CEUR-WS. org, 2015.

[200] Athanasios Fevgas, Nikolaos Fraggogiannis, Panagiota Tsompanopoulou, and Panayiotis Bozanis. The imuse virtual museum: towards a cultural education platform. In *Information, Intelligence, Systems and Applications, IISA 2014, The 5th International Conference on*, pages 171–175. IEEE, 2014.

[201] Masaki Hayashi, Steven Bachelder, Masayuki Nakajima, and Akihiko Iguchi. A new virtual museum equipped with automatic video content generator. In *Cyberworlds (CW), 2014 International Conference on*, pages 377–383. IEEE, 2014.

[202] Steven Neale, Winyu Chinthammit, Christopher Lueg, and Paddy Nixon. Facilitating learning through hands-on engagement with virtual museum artefacts. In *Proceedings of the 28th International BCS Human Computer Interaction Conference on HCI 2014-Sand, Sea and Sky-Holiday HCI*, pages 222–227. BCS, 2014.

[203] Jakub Flotyński, Jacek Dalkowski, and Krzysztof Walczak. Building multi-platform 3d virtual museum exhibitions with flex-vr. In *Virtual Systems and Multimedia (VSMM), 2012 18th International Conference on*, pages 391–398. IEEE, 2012.

[204] Felix Mata, Christophe Claramunt, and Alberto Juarez. An experimental virtual museum based on augmented reality and navigation. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 497–500. ACM, 2011.

[205] José Ignacio Rojas-Sola, Miguel Castro-García, and María del Pilar Carranza-Cañadas. Content management system incorporated in a virtual museum hosting. *Journal of Cultural Heritage*, 12(1):74–81, 2011.

[206] William Swartout, David Traum, Ron Artstein, Dan Noren, Paul Debevec, Kerry Bronnenkant, Josh Williams, Anton Leuski, Shrikanth Narayanan, Diane Piepol, et al. Ada and grace: Toward realistic and engaging virtual museum guides. In *International Conference on Intelligent Virtual Agents*, pages 286–300. Springer, 2010.

[207] JianPing Zhang and Yuhui Yang. Design and implementation of virtual museum based on web3d. In *Transactions on Edutainment III*, pages 154–165. Springer, 2009.

[208] Victor Mateevitsi, Michael Sfakianos, George Lepouras, and Costas Vassilakis. A game-engine based virtual museum authoring and presentation system. In *Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts*, pages 451–457. ACM, 2008.

[209] Volkan Isler, Bradford Wilson, and Ruzena Bajcsy. Building a 3d virtual museum of native american baskets. In *3D Data Processing, Visualization, and Transmission, Third International Symposium on*, pages 954–961. IEEE, 2006.

[210] Chun-Rong Huang, Chu-Song Chen, and Pau-Choo Chung. Tangible photorealistic virtual museum. *IEEE computer graphics and applications*, 25(1):15–17, 2005.

[211] George Lepouras and Costas Vassilakis. Virtual museums for all: employing game technology for edutainment. *Virtual reality*, 8(2):96–106, 2004.

[212] Céline Loscos, Franco Tecchia, Antonio Frisoli, Marcello Carrozzino, Hila Ritter Widenfeld, David Swapp, and Massimo Bergamasco. The museum of pure form: touching real statues in an immersive virtual museum. In *VAST*, pages 271–279, 2004.

[213] Rafal Wojciechowski, Krzysztof Walczak, Martin White, and Wojciech Cellary. Building virtual and augmented reality museum exhibitions. In *Proceedings of the ninth international conference on 3D Web technology*, pages 135–144. ACM, 2004.

[214] Yong-Moo Kwon, Jie-Eun Hwang, Tae-Sung Lee, Min-Jeong Lee, Jai-Kyung Suhl, and Sae-Woon Ryu. Toward the synchronized experiences between real and virtual museum. In *APAN 2003 Conference in Fukuoka*, 2003.

[215] Ana Carolina Bertoletti, Marcia Cristina Moraes, and Antonio Carlos da Rocha Costa. Providing personal assistance in the sagres virtual museum. 2001.

[216] E Ciabatti, Paolo Cignoni, Claudio Montani, and Roberto Scopigno. Towards a distributed 3d virtual museum. In *Proceedings of the working conference on Advanced visual interfaces*, pages 264–266. ACM, 1998.

[217] Epic Games. Unreal Engine 4. 2017. Accessed: 2018-03-09. Available: `https://goo.gl/rbDmoK`.

[218] P Petridis, M White, N Mourkousis, F Liarokapis, M Sifiniotis, A Basu, and C Gatzidis. Exploring and interacting with virtual museums. *Proc. Of Computer Applications and Quantitative Methods in Archaeology (CAA)*, 2005.

[219] Martin White, Panagiotis Petridis, Fotis Liarokapis, and Daniel Plecinckx. Multimodal mixed reality interfaces for visualizing digital heritage. *International Journal of Architectural Computing*, 5(2):321–337, 2007.

[220] Robin Milner. Is computing an experimental science? In *Journal of Information Technology*, volume 2, pages 58–66. Springer, 1987.

[221] Computer Science and CORPORATE Telecommunications Board. Academic careers for experimental computer scientists and engineers. *Communications of the ACM*, 37(4):87–90, 1994.

[222] Hossein Hassani. Research methods in computer science: The challenges and issues. *arXiv preprint arXiv:1703.04080*, 2017.

[223] D Johnson, C Papadimitriou, A Wigderson, and M Yannakakis. Challenges for theoretical computer science. *report, June*, 2000.

[224] Allen Newell and Herbert A Simon. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3):113–126, 1976.

[225] Ian M. Leslie, Derek McAuley, Richard Black, Timothy Roscoe, Paul Barham, David Evers, Robin Fairbairns, and Eoin Hyden. The design and implementation of an operating system to support distributed multimedia applications. *IEEE journal on selected areas in communications*, 14(7):1280–1297, 1996.

[226] Derek Robert McAuley. *Protocol design for high speed networks*. PhD thesis, University of Cambridge, 1990.

[227] David J Greaves. *Multi-access metropolitan area networks*. PhD thesis, University of Cambridge, 1989.

[228] Van Jacobson. Congestion avoidance and control. In *ACM SIGCOMM computer communication review*, volume 18, pages 314–329. ACM, 1988.

[229] Jerome H Saltzer, David P Reed, and David D Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems (TOCS)*, 2(4):277–288, 1984.

[230] Ian Malcom Leslie. Extending the local area network. Technical Report UCAM-CL-TR-43, University of Cambridge, Computer Laboratory, February 1983.

[231] Marja Sarvimaki. *Case Study Strategies for Architects and Designers: Integrative Data Research Methods*. Taylor & Francis, 2017.

[232] Roger J Gagnon. Empirical research: The burdens and the benefits. *Interfaces*, 12(4):98–102, 1982.

[233] David Crowther and Geoff Lancaster. *Research methods*. Routledge, 2012.

[234] Dennis Howitt and Duncan Cramer. *Introduction to research methods in psychology*. Pearson Education, 2007.

[235] Patrick McNeill and Steve Chapman. *Research methods*. Psychology Press, 2005.

[236] John Brooke et al. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.

[237] Marcus Bowman. *Using video in research*. Scottish Council for Research in Education, 1994.

[238] Andrea Garcez, Rosalia Duarte, and Zena Eisenberg. Production and analysis of video recordings in qualitative research. *Educação e Pesquisa*, 37(2):249–261, 2011.

[239] Carey Jewitt. An introduction to using video for research. 2012.

[240] Rachel Lofthouse and Peter Birmingham. The camera in the classroom: video-recording as a tool for professional development of student teachers. *Teacher Education Advancement Network Journal*, 1(2), 2010.

[241] Athanasis Karoulis, Stella Sylaiou, and Martin White. Usability evaluation of a virtual museum interface. *Informatica*, 17(3):363–380, 2006.

[242] Elizabeth J Halcomb and Louise Hickman. Mixed methods research. 2015.

[243] Courtney A McKim. The value of mixed methods research: A mixed methods study. *Journal of Mixed Methods Research*, 11(2):202–222, 2017.

[244] Viswanath Venkatesh, Susan A Brown, and Hillol Bala. Bridging the qualitative-quantitative divide: Guidelines for conducting mixed methods research in information systems. *MIS quarterly*, 37(1), 2013.

[245] John W Creswell and Vicki L Plano Clark. *Designing and conducting mixed methods research*. Sage publications, 2017.

[246] Anthony J Onwuegbuzie, Hannah R Gerber, and Sandra Schamroth Abrams. Mixed methods research. *The International Encyclopedia of Communication Research Methods*, pages 1–33, 2017.

[247] Andreas M Riege. Validity and reliability tests in case study research: a literature review with ”hands-on” applications for each research phase. *Qualitative market research: An international journal*, 6(2):75–86, 2003.

[248] Helen Noble and Joanna Smith. Issues of validity and reliability in qualitative research. *Evidence-Based Nursing*, pages ebnurs–2015, 2015.

[249] Mette Skov and Peter Ingwersen. Exploring information seeking behaviour in a digital museum context. In *Proceedings of the second international symposium on Information interaction in context*, pages 110–115. ACM, 2008.

[250] Mette Skov and Peter Ingwersen. Museum web search behavior of special interest visitors. *Library & Information Science Research*, 36(2):91–98, 2014.

[251] Michela Ott and Francesca Pozzi. Towards a new era for cultural heritage education: Discussing the role of ict. *Computers in Human Behavior*, 27(4):1365–1371, 2011.

[252] Margee Hume and Michael Mills. Building the sustainable imuseum: is the virtual museum leaving our museums virtually empty? *International Journal of Nonprofit and Voluntary Sector Marketing*, 16(3):275–289, 2011.

[253] Dirk Vom Lehn and Christian Heath. Accounting for new technology in museum exhibitions. *International Journal of Arts Management*, pages 11–21, 2005.

[254] Soojeong Yoo and Callum Parker. Controller-less interaction methods for google cardboard. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*, pages 127–127. ACM, 2015.

[255] Shenchang Eric Chen. Quicktime vr: An image-based approach to virtual environment navigation. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 29–38. ACM, 1995.

[256] Mark Claypool and Kajal Claypool. Perspectives, frame rates and resolutions: it’s all in the game. In *Proceedings of the 4th International Conference on Foundations of Digital Games*, pages 42–49. ACM, 2009.

[257] Mark Claypool, Kajal Claypool, and Feissal Damaa. The effects of frame rate and resolution on users playing first person shooter games. In *Proceedings of SPIE*, volume 6071, pages 1–11, 2006.

[258] John D McCarthy, M Angela Sasse, and Dimitrios Miras. Sharp or smooth?: comparing the effects of quantization vs. frame rate for streamed video. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 535–542. ACM, 2004.

[259] Ronnie T Apteker, James A Fisher, Valentin S Kisimov, and Hanoch Neishlos. Video acceptability and frame rate. *IEEE multimedia*, 2(3):32–40, 1995.

[260] Anne Oeldorf-Hirsch, Jonathan Donner, and Edward Cutrell. How bad is good enough?: exploring mobile video quality trade-offs for bandwidth-constrained consumers. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*, pages 49–58. ACM, 2012.

[261] Hendrik Knoche, John D McCarthy, and M Angela Sasse. Can small be beautiful?: assessing image resolution requirements for mobile tv. In *Proceedings of the 13th annual ACM international conference on Multimedia*, pages 829–838. ACM, 2005.

[262] Ernst Kruijff, Bernhard Riecke, Christina Trekowski, and Alexandra Kitson. Upper body leaning can affect forward self-motion perception in virtual environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*, pages 103–112. ACM, 2015.

[263] Abraham M Hashemian and Bernhard E Riecke. Leaning-based 360Âř interfaces: Investigating virtual reality navigation interfaces with leaning-based-translation and full-rotation. In *International Conference on Virtual, Augmented and Mixed Reality*, pages 15–32. Springer, 2017.

[264] Bernhard E Riecke. Simple user-generated motion cueing can enhance self-motion perception (vection) in virtual reality. In *Proceedings of the ACM symposium on Virtual reality software and technology*, pages 104–107. ACM, 2006.

[265] Steffi Beckhaus, Kristopher J Blom, and Matthias Haringer. Intuitive, hands-free travel interfaces for virtual environments. In *New Directions in 3D User Interfaces Workshop of IEEE VR*, pages 57–60, 2005.

[266] EL Groen and W Bles. How to use body tilt for the simulation of linear self motion. *Journal of Vestibular Research*, 14(5):375–385, 2004.

[267] Dimitri. Unity3D: Using textures larger than 4096x4096. 2011. Accessed: 2018-03-09. Available: https://goo.gl/yGrvc9.

[268] Alan Zucconi. A gentle introduction to shaders in Unity3D. 2015. Accessed: 2018-03-09. Available: https://goo.gl/vXwXbo.

[269] Jacqueline Murray. Likert data: what to use, parametric or non-parametric? *International Journal of Business and Social Science*, 4(11), 2013.

[270] Constantin Mircioiu and Jeffrey Atkinson. A comparison of parametric and non-parametric methods applied to a likert scale. *Pharmacy*, 5(2):26, 2017.

[271] Gail M Sullivan and Anthony R Artino Jr. Analyzing and interpreting data from likert-type scales. *Journal of graduate medical education*, 5(4):541–542, 2013.

[272] Ashfaq Amin, Diane Gromala, Xin Tong, and Chris Shaw. Immersion in cardboard vr compared to a traditional head-mounted display. In *International Conference on Virtual, Augmented and Mixed Reality*, pages 269–276. Springer, 2016.

[273] Nikiforos M Papachristos, Ioannis Vrellis, and Tassos A Mikropoulos. A comparison between oculus rift and a low-cost smartphone vr headset: Immersive user experience and learning. In *Advanced Learning Technologies (ICALT), 2017 IEEE 17th International Conference on*, pages 477–481. IEEE, 2017.

[274] Lorelli S Nowell, Jill M Norris, Deborah E White, and Nancy J Moules. Thematic analysis: Striving to meet the trustworthiness criteria. *International Journal of Qualitative Methods*, 16(1):1609406917733847, 2017.

[275] Liz Falconer. Experiencing sense of place in virtual and physical avebury. *Personal and Ubiquitous Computing*, 21(6):977–988, 2017.

[276] Liz Falconer and Julian Green. Exploring the learning potential of acoustic design in 3d virtual environments. *Literacy Inf Comput Educ J*, 6(1):1273–1276, 2015.

[277] Huong Q Dinh, Neff Walker, Larry F Hodges, Chang Song, and Akira Kobayashi. Evaluating the importance of multi-sensory input on memory and the sense of presence in virtual environments. In *Virtual Reality, 1999. Proceedings., IEEE*, pages 222–228. IEEE, 1999.

[278] Youngjung Suh, Choonsung Shin, and Woontack Woo. A mobile phone guide: spatial, personal, and social experience for cultural heritage. *IEEE Transactions on Consumer Electronics*, 55(4), 2009.

[279] Laia Pujol Tost and Maria Economou. Worth a thousand words? the usefulness of immersive virtual reality for learning in cultural heritage settings. *International Journal of Architectural Computing*, 7(1):157–176, 2009.

[280] Maria Economou and L Pujol. Educational tool or expensive toy? evaluating vr evaluation and its relevance for virtual heritage. *New Heritage: New Media and Cultural Heritage, London, Routledge*, pages 242–260, 2008.

[281] Tony Hall and Liam Bannon. Designing ubiquitous computing to enhance children's learning in museums. *Journal of Computer Assisted Learning*, 22(4):231–243, 2006.

[282] Marcello Carrozzino and Massimo Bergamasco. Beyond virtual museums: Experiencing immersive virtual reality in real museums. *Journal of Cultural Heritage*, 11(4):452–458, 2010.

[283] Oliver Grau. Intermedia Stages of Virtual Reality in the Twentieth Century: Art as Inspiration of Evolving Media. *Virtual art : from illusion to immersion*, pages 140–190, 2003.

[284] Neil Selwyn. The information aged: A qualitative study of older adults' use of information and communications technology. *Journal of Aging studies*, 18(4):369–384, 2004.

[285] Anabela Berenguer, Jorge Goncalves, Simo Hosio, Denzil Ferreira, Theodoros Anagnostopoulos, and Vassilis Kostakos. Are smartphones ubiquitous?: An in-depth survey of smartphone adoption by seniors. *IEEE Consumer Electronics Magazine*, 6(1):104–110, 2017.

[286] Tracy L Mitzner, Julie B Boron, Cara Bailey Fausset, Anne E Adams, Neil Charness, Sara J Czaja, Katinka Dijkstra, Arthur D Fisk, Wendy A Rogers, and Joseph Sharit. Older adults talk technology: Technology usage and attitudes. *Computers in human behavior*, 26(6):1710–1721, 2010.

[287] James Loch. *An Account of the Improvements on the Estates of the Marquess of Stafford: In the Counties of Stafford an Salop, and on the Estate of Sutherland. With Remarks*. Longman, Hurst, Rees, Orme & Brown, 1820.

[288] FSBI. FSBI Publications. 2017. Accessed: 2018-03-09. Available: https://goo.gl/1JW7Jg.

[289] SketchUp. The easiest way to draw in 3D. 2017. Accessed: 2018-03-09. Available: https://goo.gl/eFTjUC.

[290] HTC. Vive. 2017. Accessed: 2018-03-09. Available: https://goo.gl/UW641J.

[291] Oculus. Recommended System Specifications. 2018. Accessed: 2018-03-09. Available: https://goo.gl/xxRGYy.

[292] Krpano. Krpano Panorama Viewer. 2017. Accessed: 2018-03-09. Available: https://goo.gl/VwcfWw.

[293] Aristotelis Kasomoulis, Maria Vayanou, Akrivi Katifori, and Yannis Ioannidis. Magicholo: A collaborative 3d experience in the museum. In *Proceedings of the 20th Pan-Hellenic Conference on Informatics*, page 74. ACM, 2016.

[294] Peter Tolmie, Steve Benford, Chris Greenhalgh, Tom Rodden, and Stuart Reeves. Supporting group interactions in museum visiting. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 1049–1059. ACM, 2014.

[295] Rebecca E Grinter, Paul M Aoki, Margaret H Szymanski, James D Thornton, Allison Woodruff, and Amy Hurst. Revisiting the visit:: understanding how technology can shape the museum visit. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 146–155. ACM, 2002.

[296] Dirk Vom Lehn, Christian Heath, and Jon Hindmarsh. Exhibiting interaction: Conduct and collaboration in museums and galleries. *Symbolic interaction*, 24(2):189–216, 2001.

[297] Eva Hornecker and Matthias Stifter. Learning from interactive museum installations about interaction design for public settings. In *Proceedings of the 18th Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments*, pages 135–142. ACM, 2006.

[298] Yann Laurillau and Fabio Paternò. Supporting museum co-visits using mobile devices. In *International Conference on Mobile Human-Computer Interaction*, pages 451–455. Springer, 2004.

[299] DJI. Phantom 3 Advanced Specs. 2017. Accessed: 2018-03-09. Available: https://goo.gl/G7dBeF.

[300] DJI. Inspire 1 Specs. 2017. Accessed: 2018-03-09. Available: https://goo.gl/Vy42pA.

[301] Hostapd. Hostapd. 2017. Accessed: 2018-03-09. Available: https://goo.gl/wBda1i.

[302] FFmpeg. FFmpeg. 2017. Accessed: 2018-03-09. Available: https://goo.gl/7dZAWi.

[303] Bigflake. Android MediaCodec stuff. 2016. Accessed: 2018-03-09. Available: https://goo.gl/dvY7qa.

[304] Android Documentation. MediaCodec. 2017. Accessed: 2018-03-09. Available: https://goo.gl/pnnSFL.

[305] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7):560–576, 2003.

[306] Jon Postel. RFC 768 – User Datagram Protocol. 1980. Accessed: 2018-03-09. Available: https://goo.gl/WDKoKs.

[307] DJI. Drones. 2017. Accessed: 2018-03-09. Available: https://goo.gl/YokJYr.

[308] Xingyin Wang, Stefan Poikonen, and Bruce Golden. The vehicle routing problem with drones: Several worst-case results. *Optimization Letters*, 11(4):679–697, 2017.

[309] MJ Westoby, J Brasington, NF Glasser, MJ Hambrey, and JM Reynolds. 'structure-from-motion' photogrammetry: A low-cost, effective tool for geoscience applications. *Geomorphology*, 179:300–314, 2012.

[310] Susie Green, Andrew Bevan, and Michael Shapland. A comparative assessment of structure from motion methods for archaeological research. *Journal of Archaeological Science*, 46:173–181, 2014.

[311] Giovanna Sansoni, Marco Trebeschi, and Franco Docchio. State-of-the-art and applications of 3d imaging sensors in industry, cultural heritage, medicine, and criminal investigation. *Sensors*, 9(1):568–601, 2009.

[312] Jeroen De Reu, Gertjan Plets, Geert Verhoeven, Philippe De Smedt, Machteld Bats, Bart Cherretté, Wouter De Maeyer, Jasper Deconynck, Davy Herremans, Pieter Laloo, et al. Towards a three-dimensional cost-effective registration of the archaeological heritage. *Journal of Archaeological Science*, 40(2):1108–1121, 2013.

[313] Fabio Remondino. Heritage recording and 3d modeling with photogrammetry and 3d scanning. *Remote Sensing*, 3(6):1104–1138, 2011.

[314] Markos Katsianis, Spyros Tsipidis, Kostas Kotsakis, and Alexandra Kousoulakou. A 3d digital workflow for archaeological intra-site research using gis. *Journal of Archaeological Science*, 35(3):655–667, 2008.

[315] David Koller, Bernard Frischer, and Greg Humphreys. Research challenges for digital archives of 3d cultural heritage models. *Journal on Computing and Cultural Heritage (JOCCH)*, 2(3):7, 2009.

[316] Dublin Core. Dublin Core. 2017. Accessed: 2018-03-09. Available: https://goo.gl/QBdzrL.

[317] WebGL. Get WebGL. 2017. Accessed: 2018-03-09. Available: https://goo.gl/YaUJGX.

[318] Krzysztof Walczak and Wojciech R Wiza. Designing behaviour-rich interactive virtual museum exhibitions. In *Proceedings of the 8th International conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage*, pages 101–108. Eurographics Association, 2007.

[319] Anestis Koutsoudis, Fotis Arnaoutoglou, and Christodoulos Chamzas. On 3d reconstruction of the old city of xanthi. a minimum budget approach to virtual touring based on photogrammetry. *Journal of Cultural Heritage*, 8(1):26–31, 2007.

[320] Alex Galuzin. 23 Recommended and Available 3D Game Engines (Updated). 2016. Accessed: 2018-03-09. Available: https://goo.gl/DXHe2Y.

[321] Europeana. Europeana. 2017. Accessed: 2018-03-09. Available: https://goo.gl/ki5tC1.

[322] Roundme. Virtual Tours Made Simple. 2017. Accessed: 2018-03-09. Available: https://goo.gl/5AVSZx.

[323] Epic Games. Unreal Engine 4 Games. 2017. Accessed: 2018-03-09. Available: https://goo.gl/xeNFk1.

[324] James Paul Gee. *What video games have to teach us about learning and literacy*. Macmillan, 2014.

[325] Isabela Granic, Adam Lobel, and Rutger CME Engels. The benefits of playing video games. *American psychologist*, 69(1):66, 2014.

[326] Peter Vorderer, Tilo Hartmann, and Christoph Klimmt. Explaining the enjoyment of playing video games: the role of competition. In *Proceedings of the second international conference on Entertainment computing*, pages 1–9. Carnegie Mellon University, 2003.

[327] Andreas Döpker, Tobias Brockmann, Stefan Stieglitz, and Matthias Horbach. Use cases for gamification in virtual museums. In *GI-Jahrestagung*, pages 2308–2320, 2013.

[328] Julita Vassileva. Motivating participation in social computing applications: a user modeling perspective. *User Modeling and User-Adapted Interaction*, 22(1-2):177–201, 2012.

[329] Sebastian Sauer and Stefan Göbel. Dinohunter: Game based learn experience in museums. In *ICHIM'03*. Citeseer, 2003.

[330] Max Birk and Regan L Mandryk. Control your game-self: effects of controller type on enjoyment, motivation, and personality in game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 685–694. ACM, 2013.

[331] Michelle A Brown and I Scott MacKenzie. Evaluating video game controller usability as related to user hand size. In *Proceedings of the International Conference on Multimedia and Human Computer Interaction–MHCI 2013*, pages 114–1, 2013.

[332] Dustin Freeman, Otmar Hilliges, Abigail Sellen, Kenton O'Hara, Shahram Izadi, and Kenneth Wood. The role of physical controllers in motion video gaming. In *Proceedings of the Designing Interactive Systems Conference*, pages 701–710. ACM, 2012.

[333] Daniel Natapov, Steven J Castellucci, and I Scott MacKenzie. Iso 9241-9 evaluation of video game controllers. In *Proceedings of Graphics Interface 2009*, pages 223–230. Canadian Information Processing Society, 2009.

[334] Astrid Twenebowa Larssen, Toni Robertson, and Jenny Edwards. The feel dimension of technology interaction: exploring tangibles through movement and touch. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 271–278. ACM, 2007.

[335] Pippin Barr, James Noble, and Robert Biddle. Video game values: Human–computer interaction and games. *Interacting with Computers*, 19(2):180–195, 2006.

[336] Preston Tunnell Wilson, William Kalescky, Ansel MacLaughlin, and Betsy Williams. Vr locomotion: walking> walking in place> arm swinging. In *Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry-Volume 1*, pages 243–249. ACM, 2016.

[337] Jorge Cardoso. Comparison of gesture, gamepad, and gaze-based locomotion for vr worlds. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*, pages 319–320. ACM, 2016.

[338] Gerard Llorach, Alun Evans, and Josep Blat. Simulator sickness and presence using hmds: comparing use of a game controller and a position estimation system. In *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology*, pages 137–140. ACM, 2014.

[339] Stephen Carr. *Public space*. Cambridge University Press, 1992.

[340] Nemanja Memarovic, Marc Langheinrich, Florian Alt, Ivan Elhart, Simo Hosio, and Elisa Rubegni. Using public displays to stimulate passive engagement, active engagement, and discovery in public spaces. In *Proceedings of the 4th Media Architecture Biennale Conference: Participation*, pages 55–64. ACM, 2012.

[341] Pablo Figueroa, Mauricio Coral, Pierre Boulanger, Juan Borda, Eduardo Londoño, Felipe Vega, Flavio Prieto, and Diego Restrepo. Multi-modal exploration of small artifacts: an exhibition at the gold museum in bogota. In *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology*, pages 67–74. ACM, 2009.

[342] SA Brewster. The impact of haptic 'touching' technology on cultural applications in proceedings of eva 2001. *Glasgow, UK), Vasari UK, s28 pp1-14, also available on line: http://www. dcs. gla. ac. uk/˜ stephen/papers/EVA2001. pdf*, 2001.

[343] Areti Galani and Matthew Chalmers. Far away is close at hand: shared mixed reality museum experiences for local and remote museum companions. *Archives & Museum Informatics*, page 2, 2003.

[344] Areti Damala, Pierre Cubaud, Anne Bationo, Pascal Houlier, and Isabelle Marchal. Bridging the gap between the digital and the physical: design and evaluation of a mobile augmented reality guide for the museum visit. In *Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts*, pages 120–127. ACM, 2008.

[345] Richard Voase. Rediscovering the imagination: investigating active and passive visitor experience in the 21st century. *International Journal of Tourism Research*, 4(5):391–399, 2002.

[346] Chieh-Wen Sheng and Ming-Chia Chen. A study of experience expectations of museum visitors. *Tourism management*, 33(1):53–60, 2012.

[347] Sarah Johnson and Matthew Coxon. Sound can enhance the analgesic effect of virtual reality. *Royal Society open science*, 3(3):150567, 2016.

[348] Evelyn K Orman. Effect of virtual reality exposure and aural stimuli on eye contact, directional focus, and focus of attention of novice wind band conductors. *International Journal of Music Education*, 34(3):263–270, 2016.

[349] Wendy Powell, Brett Stevens, Steve Hand, and Maureen Simmonds. Sounding better: Fast audio cues increase walk speed in treadmill-mediated virtual rehabilitation environments. *Studies in health technology and informatics*, 154:202–207, 2010.

[350] McIan. The Whiskey Still. Illustration from R R McIan's Picturesque gatherings of the Scottish Highlanders at home, on the heath, the river, and the loch. 1848. London.

[351] Ross Parry. Digital heritage and the rise of theory in museum computing. *Museum management and Curatorship*, 20(4):333–348, 2005.

[352] Tomislav Šola. Essays on museums and their theory-towards the cybernetic museum. 2003.

[353] Fiona McLean. Services marketing: the case of museums. *Service Industries Journal*, 14(2):190–203, 1994.

[354] Maria Roussou. Learning by doing and learning through play: an exploration of interactivity in virtual environments for children. *Computers in Entertainment (CIE)*, 2(1):10–10, 2004.

[355] Stephen Bailey, Peter Falconer, Malcolm Foley, Gayle McPherson, and Margaret Graham. Charging for admission to museums and galleries: arguments and evidence. *Museum Management and Curatorship*, 16(4):355–369, 1997.

[356] Stephen J Bailey and Peter Falconer. Charging for admission to museums and galleries: A framework for analysing the impact on access. *Journal of Cultural Economics*, 22(2-3):167–177, 1998.

[357] Mariza Dima, Linda Hurcombe, and Mark Wright. Touching the past: haptic augmented reality for museum artefacts. In *International Conference on Virtual, Augmented and Mixed Reality*, pages 3–14. Springer, 2014.

[358] Bethann S Hromatka, Joyce Y Tung, Amy K Kiefer, Chuong B Do, David A Hinds, and Nicholas Eriksson. Genetic variants associated with motion sickness point to roles for inner ear development, neurological processes and glucose homeostasis. *Human molecular genetics*, 24(9):2700–2708, 2015.

[359] Joseph J LaViola Jr. A discussion of cybersickness in virtual environments. *ACM SIGCHI Bulletin*, 32(1):47–56, 2000.