

Ancient Roman Coin Retrieval: A Systematic Examination of the Effects of Coin Grade

Callum Fare and Ognjen Arandjelović

University of St Andrews, UK

Abstract. Ancient coins are historical artefacts of great significance which attract the interest of scholars, and a large and growing number of amateur collectors. Computer vision based analysis and retrieval of ancient coins holds much promise in this realm, and has been the subject of an increasing amount of research. The present work is in great part motivated by the lack of systematic evaluation of the existing methods in the context of coin grade which is one of the key challenges both to humans and automatic methods. We describe a series of methods – some being adopted from previous work and others as extensions thereof – and perform the first thorough analysis to date.

1 Introduction

The present is an exciting time for computer vision: the field itself has matured, the hardware needed to support developed algorithms is affordable and pervasive, and the potential user base is greater than ever owing to the increasing recognition of the benefits that machine intelligence can offer. This technological and social climate has opened a vast field of potential new applications for computer vision, with many attractive and exciting problems emerging from its applications in arts and humanities. In this work we are interested in the application of computer vision to ancient numismatics.

1.1 Terminology

Considering the interdisciplinary nature of the present paper, it is important to explain the relevant numismatic terminology so that the specific task at hand and its challenges can be clearly understood. A succinct summary is presented next.

Firstly, when referring to a *coin*, the reference is made to a specific physical object i.e. a specimen. This is to be contrasted with a coin *type*. A coin type is a more abstract concept which is characterized by the semantic features shown on both sides of the coin (the obverse i.e. the “front”, and the reverse i.e. the “back”). Multiple coins of the same type have the same visual elements e.g. the head or bust of a particular emperor with specific clothing (e.g. drapery or cuirass, crowned or laureate) and *legends* (textual inscriptions), a particular reverse motif etc. Notice that although the visual elements on coins of the same type are semantically the same, their depictions may differ somewhat. The reason lies in the fact that the same coin type was minted using dies created by different engravers. For example, observe in Fig. 1 which shows three specimens of the same type, that the spatial arrangements of the legend (by

definition the same in all cases) is different between the very fine example in Fig. 1(b) and the extra fine example in Fig. 1(c). In the former case the break (space) in the legend is AEQVITA-SAVG, and in the latter AEQVI-TASAVG. Nevertheless the type is the same.

Condition grades As noted previously, at the focal point of the present work is the *condition* of a coin. Succinctly put the condition describes the degree of preservation of a coin, or equivalently the amount of damage it suffered since it was minted. The usual grading scale adopted in ancient numismatics includes the following main grades: (i) poor, (ii) fair, (iii) good, (iv) very good, (v) fine, (vi) very fine, and (vii) extremely fine. Virtually universally (i.e. save for extremely rare coin types) only the last three are considered of interest to collectors, that is fine (F), very fine (VF), and extremely fine (EF or XF). Note that less frequently used transitional grades can be derived from the main seven by qualifiers e.g. near or almost fine (nF, aF), better than fine (F+) etc.

An ancient coin in a fine condition displays all the main visual elements of the type, as illustrated with an example in Fig. 1(a). A very fine coin also has more subtle elements preserved such as clothing creases as exemplified in Fig. 1(b). An extremely fine condition coin is in approximately the same condition in which it was when it was minted, showing the entirety of the original detail, as can be seen in Fig. 1(c).

Miscellaneous In order to appreciate the challenge of the task at hand, it is important to recognize a number of factors other than the condition which affect the appearance of a coin. These include die *centring*, surface metal changes (due to oxidation or other chemical reactions), and die wear.

Die centring refers to the degree to which the centre of the die coincides with the centre of the actual piece of metal against which it is struck to create the coin. A coin with poor centring may have salient design elements missing e.g. a part of the legend. An example of a somewhat poorly centred obverse can be seen in Fig. 1(c) and of a reverse in Fig. 2(a).

Depending on the presence of different substances in a coin's environment (soil, air etc), the surface metal can change its colour and tone as it reacts with chemicals it is exposed to. Observe the difference in the tone of the coins in Fig. 1 as well as of those in Fig. 2.

Finally, it is worth noting that the appearance of a coin can be affected by die wear. Just as coins experience physical damage when handled and used, repeated use of a die in the minting process effects damage on the die. To a non-trained eye a coin minted with a worn die can seem identical to a worn coin minted with an intact die. However, a reasonably skilled (but not necessarily expert) numismatist can readily make a distinction, as subtler patterns of damage in the two cases are quite unlike one another. In addition, close inspection and the presence of oxidation or particles in ridges can be used for conclusive verification.

1.2 Previous work

Most early and some more recent attempts at the use of computer vision for coin analysis have concentrated on modern coins [1–3]. This is understandable considering that modern coins are machine produced and as such pose less of a challenge than

ancient coins. Modern coins do not exhibit variation due to centring issues, shape, different depictions of semantically identical elements, etc. From the point of view of computer vision, two modern coins at the time of production are identical. This far more restricted problem setting allows for visual analysis to be conducted using holistic representations such as raw appearance [4] or edges [5], and off-the-shelf learning methods such as principal component analysis [4] or conventional neural networks [6]. However such approaches offer little promise in the context of ancient numismatics.

The existing work on computer vision based ancient coin analysis can be categorized by the specific problem addressed as well as by the technical methodology. As regards the former categorization, some prior work focuses on coin instance recognition i.e. the recognition of a specific coin rather than a coin type. This problem is of limited practical interest, its use being limited to such tasks as the identification of stolen coins or the detection of repeated entries in digital collections. Other works focus on coin type recognition, which is a far more difficult problem [7–9]. Most of these methods are local feature based, employing local feature descriptors such as SIFT [10] or SURF [11]. The reported performance of these methods has been rather disappointing and a major factor appears to be the loss of spatial, geometric relationship in the aforementioned representations [12, 13]. In an effort to overcome this limitation, a number of approaches which divide a coin into segments have been described [14]. These methods implicitly assume that coins have perfect centring, are registered accurately, and are nearly circular in shape. None of these assumptions are realistic. The sole method which does not make this set of assumptions builds meta-features which combine local appearance descriptors with their geometric relationships [9]. Though much more successful than the alternatives, the performance of this method is still insufficiently good for most practical applications.

All of the aforementioned work shares the same limitation of little use of domain knowledge. In particular, the general layout of the key elements of Roman imperial coins is generally fixed, save for few rare exceptions. Hence it makes sense to try to use this knowledge in analysis. The few attempts in the existing literature generally focus on the coin legend [15]. In broad terms this appears sensible as the legend carries a lot of information, much of which is shared with the coin’s pictorial elements. For example, the obverse legend in almost all cases contains the name of the emperor depicted, and the reverse the name of the deity shown. The denarius of Antoninus Pius with Aequitas (goddess of justice and equality) in Fig. 1 illustrates this well, the obverse legend being **ANTONINVS AVG PIVS P P TR P COS III**, and the reverse **AEQVITAS AVG**. However, in spite of this, methods such as that described in [15] offer little promise for practical use. The key reason for this lies in the fact that the legend, with its fine detail, is one of the first elements of the coin to experience damage and wear. Coins with clearly legible legends are generally expensive and rare, and thus of little interest to most collectors. They are also the easiest to identify, by the very nature of their good preservation, and hence do not represent the target data well. Consequently, this class of algorithms is not of interest in the present paper.

The main purpose of this work is to provide a clear picture of the performance of existing methods on data representative of images likely to be used in practice. Moreover our aim is to provide the first systematic evaluation which looks specifically at

the effects that coin grade has on coin type recognition accuracy. In particular, all work to date has been highly unstructured and *ad hoc* in its evaluation methodology. Some authors use data sets with coins in different conditions and unstated distributions thereof [9], and others very small data sets with coins in extremely rare, museum grade [8]. Hence the current understanding of different methods’ behaviour is not very well understood at all.



Fig. 1. Specimens of Antoninus Pius’s denarius (RIC 61) from our data set.

2 Data

As noted in the previous section, one of the key motivating factors for the present work can be found in the lack of systematic evaluation of different algorithms described in the literature with respect to the condition of the coins present in the specific data sets used. Given that the condition of a coin by definition affects the visibility and even the very presence of elements depicted on the coin, it is unsurprising that it is a major factor which governs the ease (or lack thereof) that a human experiences when attempting to identify a coin. Understanding the behaviour of different methods when presented with this challenge, and in particular the effects of both the condition of the query coin as well as of the distribution of coin conditions in the so-called gallery corpus, should be a crucial consideration in directing future research efforts.

At this point in time there does not exist a data set structured in a manner which allows for the analysis outlined above to be conducted: none of the corpora used in previous work can be readily adopted for use to this end, nor are there any other readily available sources, to the best of our knowledge. Hence we collected a novel data set which we introduce for the first time in this paper – it will be made freely available after anonymity is no longer needed.

We collected our data by searching for images of coins sold by well known auction houses. In this manner we achieved two goals. Firstly, we could ensure that the images are in the public domain and can thereafter be shared without restriction. Secondly, having been put up for sale by well known auction houses, the coins have been graded by professionals allowing us to associate reliable meta data with all images.

We collected 600 images in total. These represent 100 types of Roman imperial denarii, with six exemplars for each type: two in fine condition, two in very fine, and two in extremely fine. The period covered by the coins included in the data set starts with the beginning of the Empire and the rule of Octavian in 27 BC and ends with the end of the rule of Philip II (Philip the Arab) in 249 AD when the denarius ceases to be used due to economic and political crises. A few representative examples of different coin types in different grades from our corpus are shown in Figs 1–3.



Fig. 2. Specimens of Caligula's denarius (RIC 2) from our data set.



Fig. 3. Specimens of Octavian's denarius (RIC 102) from our data set.

3 Methods

This section describes the methods used in our evaluation.

3.1 Histogram distance measures

Given that all methods described in the sections that follow (and indeed most of the methods in the existing literature on computer vision based ancient coin analysis) employ histogram based representations, we start by detailing the histogram distance measures used in our experiments.

For the sake of continuity with previous work, the first distance measure we adopt is the standard Euclidean distance [16]. Given two L_2 normalized histograms \mathbf{h}_1 and \mathbf{h}_2 , defined over some vocabulary of size n , the Euclidean distance $d_E(\mathbf{h}_1, \mathbf{h}_2)$ between them is given by the following expression:

$$d_E(\mathbf{h}_1, \mathbf{h}_2)^2 = \sum_{i=1}^n [h_1(i) - h_2(i)]^2, \quad (1)$$

where $h_j(i)$ denotes the i -th entry in histogram \mathbf{h}_j .

Notwithstanding the widespread use of the Euclidean distance metric, a recently proposed alternative which is based on Hellinger distance has universally been shown to yield superior performance [17]. The metric is just as simple and efficient to evaluate as the Euclidean one and is given by:

$$d_H(\mathbf{h}_1, \mathbf{h}_2)^2 = \sum_{i=1}^n [\sqrt{h_1(i)} - \sqrt{h_2(i)}]^2, \quad (2)$$

with an important difference that histograms should be L_1 normalized.

3.2 Baseline SIFT

The first algorithm we implemented and evaluated in this work is what we term 'Baseline SIFT' on the account of its widespread use in the existing literature [10,

18]. As different elements of this algorithm are employed by the other approaches we also evaluated, we explain the key steps in some detail. In summary, Baseline SIFT first creates a visual dictionary by clustering SIFT descriptors from the coin gallery, uses the constructed dictionary to represent a single coin as a histogram of visual words, and performs matching using one of the distance metrics described previously.

Visual dictionary construction Baseline SIFT starts the construction of a visual dictionary by detecting keypoints and extracting the corresponding SIFT descriptors from all coin images in the gallery of ‘known’ coins [10]. The extracted descriptors are then clustered using k -means clustering, with the parameter k set *a priori* (we will discuss the choice of k shortly). Given the stochastic nature of k -means, in order to obtain the best (most descriptive, for a set value of k) clustering, in this work we perform several clustering attempts and of those choose the one with the least average L_2 error measured between individual descriptors and their assigned cluster centres. The final k cluster centres are deemed the visual vocabulary which allows a single image of a coin to be represented using a fixed length representation. In particular, given the set of SIFT descriptors from a single coin, each descriptor is taken to be a representative of the visual word given by the closest cluster centre, and the entirety of the coin image represented by a histogram over the visual vocabulary.

On the choice of vocabulary size The choice of the visual vocabulary size k is an interesting and practically important one. Two different views on the approach taken can be put forward. A large value of k , commonly used in instance retrieval applications [19], can be considered as a way of hashing and matching noisy descriptors. Alternatively, smaller values of k , more often used for object class recognition, can be seen as a means of generalization. Although this choice has not been explicitly discussed in the existing literature of automatic ancient coin analysis, implicitly the latter view seems to be dominant [9]. Considering the lack of systematic analysis to date, we take no *a priori* stance and instead conduct experiments for a range of values of k .

3.3 Wedge SIFT

As discussed previously, Baseline SIFT and similar approaches suffer from a major limitation caused by the loss of spatial information. A specific SIFT descriptor affects the overall representation of the coin in the same manner regardless of its absolute location or indeed location relative to other descriptor loci. A number of methods in the literature attempt to address this problem by dividing a coin into segments. Given the approximately circular shape of coins, a natural way of segmenting a coin is into radial segments or wedges, constructing a histogram for each segment, and concatenating these into a single, higher dimensional vector used to represent the entirety of a coin [20]. We refer to this method as Wedge SIFT. An example is shown in Fig. 4(a) which displays keypoint loci colour coded for the corresponding segments, using four wedges.

3.4 Soft Wedge SIFT

All existing methods in the literature which attempt to combine sparse local feature based appearance information with geometric information by segmenting a coin do

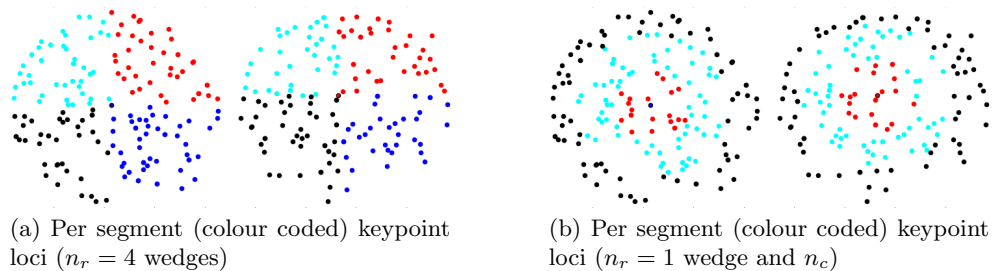


Fig. 4. Different coin segmentation approaches used to incorporate geometric information by grouping spatially related keypoints and the corresponding local appearance descriptors.

so using manually predefined segments and ‘hard’ segment membership i.e. a specific keypoint and the corresponding descriptor are strictly considered either to fall within a segment or not. Several problems emerge from this approach. Firstly, these methods implicitly assume that coins are perfectly registered both in terms of their translation and rotation. This is difficult to achieve by automatic means and indeed none of the existing work discusses this challenge. Yet, to perform this manually defeats the very premise of automatic coin analysis. What is more, the problem of exact rotational alignment is not even well posed as it is not objectively clear what the precise ‘up’ direction is in the first place. The hard membership of features compounds this problem – even a slight misalignment, translational or rotational, can greatly affect feature distributions in different segments. Hence in this work we also evaluate an extension of the original method by allowing soft feature membership within a wedge. In particular, we apply weighting using a triangular fuzzy number function which reaches its maximum value of 1 for a feature at the centre of a wedge (i.e. in the bisecting direction from the centre of the coin) and its minimum of 0 at the centres of the two neighbouring wedges [21, 22].

3.5 Wedge-Sector SIFT

The Wedge SIFT approach creates segments by dividing the image of a coin (that is, more precisely, the images of its obverse and reverse) using radial boundaries. A complementary dividing methodology uses concentric circular boundaries using different radii, creating sectors. Therefore this algorithm has two free parameters, namely the number of radial boundaries n_r and the number of circular boundaries n_c , thereby dividing a coin into $n_r \times n_c$ segments. An example of keypoint assignments is shown in Fig. 4(b), for clarity using $n_r = 1$ wedges and $n_c = 4$ sectors.

3.6 Soft Wedge-Sector SIFT

The hard membership based Wedge-Sector SIFT segmentation suffers from the same limitations as those highlighted in the case Wedge SIFT. Hence we apply the same idea of soft membership of local features in Wedge-Sector SIFT by using weighting both in radial and angular directions.

3.7 Local binary patterns (LBPs)

Most methods in the existing literature on computer vision based ancient coin analysis rely on the use of sparse, local features. This is a reasonable choice given that the

precise geometric layout between different elements of the same coin type can vary considerably across specimens minted with different dies. However, state of the art performance in problem domains where similar geometric flexibility is present, such as face recognition, has been achieved with the use of dense local features in the form of local binary patterns (LPBs) [23]. The LPB representation has proven to be very effective across a range of applications, including texture and face recognition [24, 25], and numerous others.

The elementary local LBP descriptor considers an image patch of size 3×3 pixels. By comparing the values of the 8 neighbouring pixels with the value of the central pixel, the neighbourhood is mapped to a series of binary digits (0 or 1) depending on whether a specific pixel has a smaller value than the central pixel or not, as illustrated. The 8 bit sequence corresponds to an integer in the range $[0, 127]$ and describes the local appearance. The description of an entire image (or a region of interest within it) is then obtained by creating histograms over local LBP descriptors within blocks into which the image is divided (the number of blocks is a free parameter, examined empirically in the next section).

3.8 Local ternary patterns (LTPs)

The thresholding of pixels at the heart of LPBs, similarly to the hard thresholding in terms of the spatial layout of local features discussed previously, is vulnerable to small perturbations when neighbouring pixels have values close to the central pixel. A generalization of the LPB descriptor in the form of a local ternary pattern (LTP) has demonstrated effectiveness in addressing this problem. In particular, instead of mapping neighbourhood pixels to binary digits, to produce a LTP the mapping is done to a ternary digit i.e. without loss of generality, to 0, 1, or 2. A pixel is mapped onto 0 or 2 respectively if its value is smaller or greater than that of the central pixel by at least a certain amount (this threshold is a free parameter), and to 1 otherwise. The latter, additional value can be seen as representing neighbourhood pixels sufficiently similar to the central one. The remainder of the method, that is the aggregation of local descriptors into histograms over blocks, and the concatenation of these to form a holistic representation, is performed just as in the original LBP based method.

4 Experiments

In this section we present our experiments. Specifically, we begin with a summary of automatic data preprocessing needed to prepare images for use in the algorithms described in the previous section, go on to explain our experimental methodology, and finally present and discuss our findings.

4.1 Automatic data preprocessing

The images in our database, being originally acquired for use at auctions, require additional processing before they can be used by the methods described in the previous section: their size is non-uniform, the photographs of a coin's obverse and reverse are shown within the same image and their locations are not in *a priori* known locations within the image. Here we describe a series of automatic pre-processing steps which normalize for these confounding sources of variation. In summary, we (i) detect and

segment out image regions which correspond to the obverse and the reverse, (ii) estimate the size of the coin in the segmented images, and (iii) perform image rescaling to the canonical scale.

Obverse/reverse segmentation The first step of our pre-processing pipeline concerns the separation of a coin’s obverse and its reverse. This is achieved by keypoint localization using a Gaussian scale-space as described by Lowe [10], and then by clustering the loci using k -means for $k = 2$. This process readily leads to the identification of image areas which correspond to the two sides of a coin, as illustrated in Fig. 5.




Fig. 5. Keypoint detection and clustering using k -means with $k = 2$ readily allows the areas of the image which correspond to the two coin sides to be separated. In this image keypoint loci are colour coded by their final cluster assignments.

Scale normalization The second pre-processing step we conduct concerns image scale canonization i.e. rescaling to the uniform scale. This is needed because higher resolution images tend to produce higher numbers of keypoints which can clearly affect the representations described in the previous section. Following k -means clustering of loci of keypoints detected in raw images, coin scale can be determined in a simple manner. In particular we consider the median of the cluster to be the centre of the coin and estimate the average diameter of the coin by computing the mean distance of convex hull defining keypoints from this centre.

4.2 Experimental methodology

In order to facilitate as thorough understanding of the effects of coin grade on the performance of different methods as possible, we conducted a series of experiments which vary the grade of coins used as query and as gallery. In particular, we evaluated all algorithms first using three experiments in which different coin grades (F, VF, or XF) were used to query a gallery with all conditions of coins present in it, and then three further experiments in which a query coin of a specific grade was matched against a gallery which includes coins of that grade only.

4.3 Results and discussion

We began our analysis by looking at the simplest, Baseline SIFT method. The key results using different parameter values (in particular, the visual vocabulary size) for the two adopted histogram metrics are summarized in Tables 1 and 2. The first observation that emerges from the tables concerns the poor performance of the method which correctly recognized no more than 5% of query coins in the best case. This is consistent with previous reports in the literature [9], with the method showing any promise only in the context of the far simpler problem of coin instance recognition.

Table 1. Summary of results of the Baseline SIFT method (Euclidean histogram distance).

k	All	F	VF	XF
100	4.0%	3.1%	3.1%	2.7%
1000	3.1%	1.8%	2.1%	1.0%
10000	1.4%	1.1%	1.4%	0.8%

Table 2. Summary of results of the Baseline SIFT method (Hellinger histogram distance).

k	All	F	VF	XF
100	5.5%	3.4%	3.6%	3.6%
1000	3.4%	1.9%	2.6%	1.3%
10000	1.4%	1.0%	1.4%	0.6%

The second clear observation concerns the superiority of the Hellinger distance based histogram metric as compared with the more conventional Euclidean distance. This finding too is consistent with the reports in the literature on other recognition problems [17]. Therefore henceforth we adopt the use of this metric exclusively.

Next, notice that in both Tables 1 and 2, the best performances were achieved using the smallest value of k . This supports the idea of using coarse feature discretization as a means of providing generalization robustness, as discussed previously. Considering that the same trend was found in all experiments we conducted, henceforth all reported results are for $k = 100$.

As expected, results superior to those obtained with Baseline SIFT were attained through the use of geometric information and the two variants of Wedge SIFT, as illustrated in Tables 3 and 4. Nevertheless, on the absolute scale in the context of practical applicability, the recognition rates remain poor, not exceeding 10%.

Unsurprisingly, recognition was worst when poorest condition coins (F) were used. Interestingly though, in the case of the original Wedge SIFT, the use of the best condition coins (XF) did not effect an improvement over medium grade coins (VF) – rather, the performance worsened. A possible explanation for this may lie in the greater number of keypoints and the corresponding features detected in extremely fine coins. These features often correspond to idiosyncratic details specific to individual die engravers, rather than discriminative features in the context of coin type recognition. This is something that future research should bear in mind and which may be an interesting avenue to explore in the analysis of engraving style patterns.

Table 3. Summary of results of the Wedge SIFT method, using the Hellinger histogram distance metric and the visual vocabulary size $k = 100$ (see Table 2).

n_r	All	F	VF	XF
2	7.4%	6.6%	9.3%	6.2%
3	6.4%	7.7%	6.9%	4.8%
4	5.3%	5.6%	6.5%	3.8%

Finally, the results obtained using the two Sector SIFT methods are summarized in Tables 5 and 6. While both methods performed better than Baseline SIFT, their recognition rates were lower than of Wedge SIFT and Soft Wedge SIFT. We expect that the key reason lies in the weaker geometric constraint imposed by this repre-

Table 4. Summary of results of the Soft Wedge SIFT method, using the Hellinger histogram distance metric and the visual vocabulary size $k = 100$ (see Table 2).

n_r	All	F	VF	XF
2	7.9%	6.6%	8.3%	8.7%
3	8.2%	6.4%	8.8%	9.1%
4	8.7%	6.1%	9.7%	10.1%

sensation – in the context of the problem at hand, angular displacement is more informative than radial, as well as less sensitive to the precise localization of a coin’s centre. Indeed, we highlighted the importance of the latter in our coverage of previous work, and the lack of consideration thereof in the previous work which proposed methods predicated on this information.

Table 5. Summary of results of the Sector SIFT method, using the Hellinger histogram distance metric and the visual vocabulary size $k = 100$ (see Table 2).

n_c	All	F	VF	XF
2	6.8%	5.6%	8.3%	6.2%
3	7.2%	6.1%	8.8%	6.7%
4	5.3%	7.1%	5.1%	3.8%

Table 6. Summary of results of the LBP based method.

n_b	All	F	VF	XF
3	3.6%	3.6%	4.6%	2.4%
4	4.8%	4.6%	5.1%	4.8%
5	4.4%	4.6%	4.6%	3.8%

5 Conclusions and future work

In this paper we focused on the problem of recognizing Roman imperial denarii – a difficult computer vision problem which is of much interest to communities interested in ancient numismatics. In particular our work was motivated by the lack of systematic evaluation of the effects that coin grade has on the performance of different algorithms. We described a series of different methods, some adopted from previous work and others proposed as extensions thereof, and performed the first thorough analysis in the existing literature. Our findings demonstrate the difficulty of the problem and suggest that the existing methods still perform very poorly on real world data. We analysed and discussed the behaviour of different algorithms and their parameters, and highlighted a series of observations which should guide future work. In particular our results suggest a focus on the use of prior knowledge of the coin layout and line or edge based features [26, 27] rather than appearance.

References

1. P. Davidsson. Coin classification using a novel technique for learning characteristic decision trees by controlling the degree of generalization. *International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, 1996.

Table 7. Summary of results of the LTP based method.

t	All	F	VF	XF
5	5.6%	4.1%	6.9%	5.8%
6	5.8%	4.6%	6.0%	6.7%
7	4.5%	4.1%	4.2%	5.3%

2. M. Nölle *et al.* Dagobert - a new coin recognition and sorting system. *International Conference on Digital Image Computing*, 2003.
3. X. Pan and L. Tougne. Topology-based character recognition method for coin date detection. *CAIP*, 2016.
4. R. Huber *et al.* Classification of coins using an eigenspace approach. *PRL*, 2005.
5. L. van der Maaten and P. Boon. COIN-O-MATIC: A fast system for reliable coin classification. *MUSCLE CIS Coin Recognition Competition Workshop*, 2006.
6. Y. Mitsukura *et al.* Design and evaluation of neural networks for coin recognition by using GA and SA. *IJCNN*, 2000.
7. M. Zaharieva *et al.* Image based recognition of ancient coins. *CAIP*, 2007.
8. M. Kampel and M. Zaharieva. Recognizing ancient coins based on local features. *ISVC*, 2008.
9. O. Arandjelović. Automatic attribution of ancient Roman imperial coins. *CVPR*, 2010.
10. D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2003.
11. H. Bay *et al.* SURF: Speeded up robust features. *CVIU*, 2008.
12. W. Rieutort-Louis and O. Arandjelović. Bo(V)W models for object recognition from video. *IWSSIP*, 2015.
13. O. Arandjelović. Object matching using boundary descriptors. *BMVC*, 2012.
14. H. Anwar *et al.* Coarse-grained ancient coin classification using image-based reverse side motif recognition. *MVA*, 2015.
15. O. Arandjelović. Reading ancient coins: automatically identifying denarii using obverse legend seeded retrieval. *ECCV*, 2012.
16. O. Arandjelović. Matching objects across the textured-smooth continuum. *ACRA*, 2012.
17. R. Arandjelović and A. Zisserman. Three things everyone should know to improve object retrieval. *CVPR*, 2012.
18. W. Rieutort-Louis and O. Arandjelović. Description transition tables for object retrieval using unconstrained cluttered video acquired using a consumer level handheld mobile device. *IJCNN*, 2016.
19. J. Sivic and A. Zisserman. Video Google: A text retrieval approach to object matching in videos. *ICCV*, 2003.
20. H. Anwar *et al.* Supporting ancient coin classification by image-based reverse side symbol recognition. *CAIP*, 2013.
21. W. Pedrycz. Why triangular membership functions? *Fuzzy Sets and Systems*, 1994.
22. O. R. Bidder *et al.* A risky business or a safe BET? A fuzzy set event tree for estimating hazard in biotelemetry studies. *Animal Behaviour*, 2014.
23. H. Tang *et al.* 3D face recognition using local binary patterns. *SP*, 2013.
24. M. Heikkilä *et al.* Description of interest regions with local binary patterns. *PR*, 2009.
25. R. S. Ghiass *et al.* Infrared face recognition: a literature review. *IJCNN*, 2013.
26. O. Arandjelović. Gradient edge map features for frontal face recognition under extreme illumination changes. *BMVC*, 2012.
27. O. Arandjelović. Making the most of the self-quotient image in face recognition. *FG*, 2013.