

Exploring Tangible Interactions with Radar Sensing

HUI-SHYONG YEO, University of St Andrews, United Kingdom

RYOSUKE MINAMI, University of St Andrews, United Kingdom

KIRILL RODRIGUEZ, University of St Andrews, United Kingdom

GEORGE SHAKER, University of Waterloo, Canada

AARON QUIGLEY, University of St Andrews, United Kingdom

Research has explored miniature radar as a promising sensing technique for the recognition of gestures, objects, users' presence and activity. However, within Human-Computer Interaction (HCI), its use remains underexplored, in particular in Tangible User Interface (TUI). In this paper, we explore two research questions with radar as a platform for sensing tangible interaction with the counting, ordering, identification of objects and tracking the orientation, movement and distance of these objects. We detail the design space and practical use-cases for such interaction which allows us to identify a series of design patterns, beyond static interaction, which are continuous and dynamic. With a focus on planar objects, we report on a series of studies which demonstrate the suitability of this approach. This exploration is grounded in both a characterization of the radar sensing and our rigorous experiments which show that such sensing is accurate with minimal training. With these techniques, we envision both realistic and future applications and scenarios. The motivation for what we refer to as Solinteraction, is to demonstrate the potential for radar-based interaction with objects in HCI and TUI.

CCS Concepts: • **Human-centered computing** → *Interaction techniques; Interface design prototyping; Ubiquitous and mobile computing design and evaluation methods;*

Additional Key Words and Phrases: Tangible Interaction, Tangible User Interface, Radar Sensing, Soli, Ubiquitous Computing, Context-Aware Interaction, Token+Constraint, Machine Learning

ACM Reference Format:

Hui-Shyong Yeo, Ryosuke Minami, Kirill Rodriguez, George Shaker, and Aaron Quigley. 2018. Exploring Tangible Interactions with Radar Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 4, Article 200 (December 2018), 25 pages.

1 INTRODUCTION

The challenges in sensing and tracking day-to-day objects to facilitate Tangible User Interfaces (TUIs) has been explored in prior work. In a TUI, people interact with physical objects which are coupled to the control of computational artifacts and digital information. Such tangible interactions have been explored in education [6], musical expressiveness [24, 43], input [12, 14] and data manipulation [47, 55]. This paper proposes the use of miniature radar sensing [69] to enhance such interactions.

Research on sensing and tracking objects to facilitate tangible interactions, is a well established field with techniques relying on the use of cameras [6, 9, 15], capacitive sensing [14, 50], fiducial markers [24], sound [38], magnets [11, 12, 30] and tags (e.g., RFID) [28, 29, 31, 53]. However, these techniques can often involve complex tracking technologies, bespoke infrastructures or object modification. In addition, object occlusion, multiple

Authors' addresses: Hui-Shyong Yeo, University of St Andrews, School of Computer Science, Fife, Scotland, KY16 9SX, United Kingdom, hsy@st-andrews.ac.uk; Ryosuke Minami, University of St Andrews, School of Computer Science, Fife, Scotland, KY16 9SX, United Kingdom, rm264@st-andrews.ac.uk; Kirill Rodriguez, University of St Andrews, School of Computer Science, Fife, Scotland, KY16 9SX, United Kingdom, krb5@st-andrews.ac.uk; George Shaker, University of Waterloo, Department of Electrical and Computer Engineering, Waterloo, Ontario, ON N2L 3G1, Canada, gshaker@uwaterloo.ca; Aaron Quigley, University of St Andrews, School of Computer Science, Fife, Scotland, KY16 9SX, United Kingdom, aquigley@st-andrews.ac.uk.

2018. 2474-9567/2018/12-ART200

<https://doi.org/>

Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., Vol. 2, No. 4, Article 200. Publication date: December 2018.




Fig. 1. Common everyday objects tested in Solinteraction.

objects and objects at varying distances can present challenges. This range of issues has motivated the exploration of alternative forms of sensing, such as the use of miniature radar in this work, to help address such limitations.

The identification of such problems gives rise to the first of the research questions investigated here, namely, does miniature radar sensing provide an accurate measure of the precise stacking, ordering and tracking the movement of planar objects and fabricated objects. To answer this research question, a study has been carried out for evaluating the six proposed sensing mechanisms, namely - counting, ordering, identification, flipping, movement and distance. These experiments are grounded in both a characterization of the radar sensing and our rigorous experiments which show that such sensing is accurate with minimal training.

The second research question aims to demonstrate that such sensing can provide useful information for new forms of tangible interaction. A series of design attributes, application scenarios and future applications have been defined to answer this question.

In this paper, we employ radar sensing and machine learning for the exploration of tangible interactions with a family of techniques we term, Solinteraction. The use of radar is not new, indeed automotive, ground and weather radar are commonly employed at varying scales and with variable environmental and lighting conditions. The potential of radar sensing in HCI also has been demonstrated in detecting micro hand gestures [32, 61] and in our prior work in recognizing objects and materials [68]. However, other facets of radar have remained underexplored in the context of HCI, especially in tangible interaction. We propose using radar as a sensing platform for tracking the identity of an object or number of proximate objects to enable TUI. In the section on Radar, we first study the permittivity of objects with a Microwave Network Analyzer. Once we can recognize, count and order objects we can determine movement (i.e., sliding or rotation), identity and activity with common, largely planar objects, as shown in Figures 1, 6, 7 and video figure and as described in the prototype section.

To realize our prototype, we employ the Google ATAP Project Soli radar sensor [20, 32], extend the RadarCat [68] system and by building on the Token+Constraint [56] paradigm. We show that by adding a “guide case” that

act as constraints, we can greatly expand the interaction capabilities and improve the object sensing accuracy. We then utilize these capabilities to explore the design space of tangible input and interaction and suggest scenarios that can be enabled. We evaluate the robustness and accuracy of the family of techniques and discuss the potential design space and applications. This paper makes the following contributions:

- An exploration of radar-based sensing for tangible interaction and extending it to support six new sensing modes. We demonstrate fabricating constraints and tokens of various designs, which further extends the sensing capabilities and accuracy.
- A lab measurement accompanied by an EM full-wave simulation, followed up by real-world evaluation of the robustness and accuracy of the family of techniques, with a focus on counting and ordering a stack. We also provide the dataset¹ of raw signals for others to draw upon.
- An exploration of the design space and demonstration of the potential of creating applications with rich, tangible interactive functionality. We also suggest scenarios that can be enabled with our techniques.

2 RELATED WORK

This form of tangible user interface, with Solinteraction, builds on a wide-range of related work from sensing to radar. Here we summarize key areas, focusing on (i) digital-physical approaches to recognize and track objects (ii) the use of constraints to augment sensing (iii) fabricating and altering internal structure and (iv) the use of radar sensing in HCI.

2.1 Tracking Tangible Tokens

Robust tracking of objects is essential to enable seamless TUIs. While a variety of object or token tracking and sensing techniques have been proposed, this remains an active area of research.

Vision-based tracking: Vision-based methods such as PlayAnywhere [66] and Portico [6] use an overhead camera to visually track arbitrary objects. Similarly, Augmented Surfaces [46] and iCon [15] also use cameras but require augmenting everyday objects with visual markers. Alternative vision-based approaches use a camera beneath the surface for tracking objects or markers attached to objects placed on a tabletop (Microsoft Surface Table [18] and reacTable [24]). In particular, Lumino [9] and Stacks on the Surface [8] allow multiple objects and markers to be stacked and tracked to achieve more complex interactions.

Capacitive sensing: The widespread adoption of tablet devices has fostered research and commercial efforts (e.g., Disney appMATes) for tracking objects near or around such devices. Popular techniques include using conductive material to create a unique “footprint” that can be sensed by the capacitive screen. Identifying and tracking both passive [45, 59] and active (powered) tokens [58, 70] are well explored methods. In particular, CapStones [14] demonstrate stackable blocks with a system of wires and connectors inside that, once placed on a capacitive screen, allows the device underneath to discern the block’s identity.

Magnetic sensing: The naturally occurring or induced magnetic fields of some objects can also be sensed to recognize and track objects. For example, Sensetable [42] uses individually tuned resonant coils to track and differentiate objects. GaussBits [30] consist of an array of hall-effect sensors placed underneath a tablet device and is used for simultaneous tracking of multiple magnetic objects on-screen, while Magnetic Appcessories [11] use only the built-in magnetometer of a mobile device. MagGetz [23] and MagnID [12] extend interaction by allowing tracking of multiple objects around the device in fixed or arbitrary locations.

RFID or NFC: Radio Frequency Identification (RFID) [62] or Near Field Communication (NFC) based approaches are relatively common and are employed in consumer products such as the toys-to-life series [63] (Skylanders, Amiibo). Such tags allow for the recognition of different tagged objects and motions undertaken with them [28, 29, 47], thereby turning objects into interactive elements [22, 53].

¹<https://github.com/tcboy88/solinteractiondata>

Hybrid methods: However, each of the above techniques suffers from different limitations such as large form factors (vision, RFID), lighting and occlusion problems (vision), requirement for tagging (capacitive, RFID), interference (magnetic, RFID). Therefore, hybrid methods combine multiple sensing techniques with the goal to achieve better tracking with fewer limitations. For example, GaussRFID [31] combines magnetic sensing with RFID whereas Project Zanzibar [57] combines capacitive sensing and NFC to track tokens placed or hovering on a flexible mat. Yet, all these approaches cannot support more than a few stacks and ordering of objects. Solinteraction offers a solution to many of these limitations such as no tagging required and supports stacking multiple layers, albeit with its own limitations such as a limited sensing area.

2.2 Constrained Sensing

Some approaches rely on adding constraints to augment the sensing capabilities of their underlying systems [25, 39, 49]. For example, Osmo [39] is a simple mirror module that can be attached to the front camera of a mobile device. With the mirror to guide the front camera, it is able to track various pieces that come with the package. The Token+Constraint [56] paradigm describes constraints that can mechanically channel how tokens can be manipulated, reducing ambiguity and challenges in tracking. Acoustruments [25] consist of an add-on phone case with tubes that guide the acoustic pathway between the speaker and the microphone of a mobile device. Sauron [49] augments a camera with a 3D printed case that encloses it, for sensing human input on physical controls, such as buttons, sliders and a joystick. FlexiBend [16] is a shape-sensing strip that can be enclosed in different 3D printed widgets to sense various user inputs such as buttons, slider and knobs. These approaches can simplify the extent of the sensing or the range of inputs resulting in improvements to recognition or interaction. Solinteraction adopts this approach to improve its sensing capabilities, but can also work without it.

2.3 Fabrication and Altering Structure

To improve the tracking of tokens, some approaches rely on altering the tokens such as adding electronic components [26, 27] or altering the external [34] and internal structure [14, 25, 50, 65]. For example, TouchTokens [34] are fabricated with notches that constrain a users' grasp so that the relative spatial configuration of fingers on the surface can be easily tracked. Capricate [50] proposed 3D printing pieces with embedded capacitive parts so they can be tracked by capacitive touch screens. Other approaches such as designing inner structures and pathways to guide optical [64], acoustic [25] or even electromagnetic wave [65] properties as these sources are passed through the structures. By decoding the resulted signal it is possible to extract information embedded in the way inner layers are constructed. As a result, the tangibility, manipulability and multiplicity of tokens can be impacted by such approaches. Solinteraction also involves fabrication of the constraint case and tokens of various designs.

2.4 Radar in HCI

While Paradiso et al. [40, 41] employed radar to track motion for musical interaction and performance over 20 years ago, radar sensing in the context of HCI has received comparatively little attention, until the recent availability of small low-cost hardware elements [13, 20, 60] which has enabled preliminary explorations [4, 69]. Due to its higher sensitivity and wider sensing range, low-cost radar is increasingly replacing PIR sensors in smart lighting products that sense motion to turn on the light automatically, for example. In computing, recent research has focused on using radar for small finger motion recognition [32, 61], augmented reality interactions [19] and musical applications [10, 48], activity recognition [5] and human pose tracking through walls [1, 72]. In the health domain, radar related sensing has been explored for monitoring glucose levels in blood samples non-invasively [37, 52], breathing and heart rate monitoring [2], sleep sensing [44] and emotion tracking [71] in a smart home environment. In terms of object and material recognition, prior work in RadarCat [67, 68] and

Deskwave [33] have demonstrated the ability to disambiguate passive and unmodified objects and their materials by analyzing the characteristics of the radio frequencies reflected.

3 RADAR

With radar, a target is illuminated with an electromagnetic signal which is then scattered, absorbed and reflected back by the target to a radar's receiving antenna for measurement. There are various type of radars (e.g., impulse or Doppler radar) but our focus here is on frequency-modulated continuous-wave (FMCW) radar, as it has been shown that this can be used to recognize objects and materials with high accuracy at a close distance [68]. In FMCW mode, transmission is a continuously increasing frequency signal until it reaches the maximum frequency within a predefined time period (known as a chirp), whose duration is much greater than the radar return time. The received signal at the output port of the receiver antenna is amplified and correlated with the transmit signal, which results in a signal called a beat signal. The beat signal contains information about the objects in the scene. Particularly, the delay in the reflected signal is translated to an instantaneous frequency difference between the transmitted and the received chirps. Due to the property of a chirp, the range information is contained in the intermediate frequency component and can be obtained by the spectral analysis of the beat signal (i.e., interference between the two signals). FMCW has a range resolution ($res = \frac{c}{2BW}$, where c is the speed of light and BW is the total bandwidth of the transmitted waveform) that varies with the range of frequencies used.

For example, the radar employed here in the Soli [32], is a high resolution radar so when it illuminates a thick target object (more than the range resolution) radially the target is resolved into more than a single range bin at a frequency peak. Even a minute change in thickness of the detected object is capable of changing the phase of spectral content of the detected beat signal. Thus, the time domain intermediate frequency (IF) signal, represents the properties of the detected target (thickness/height, radar cross section, material, etc.). This presents a unique signature which was exploited previously in [68] to classify everyday objects and materials, whereas here in Soliinteraction we aim to support new sensing modalities and hence to enable novel tangible interaction.

3.1 Radar and Minute Planar Changes

Prior to exploring the use of radar for stacking, ordering and movement of planar objects, fabricated objects and identification we first aimed to quantify the extent to which such a radar sensing approach might be able to differentiate between even minute changes. This measurement campaign starts with a simple question: "Can we use a multi-channel radar, such as the Soli, to differentiate among different numbers of playing cards?" To help answer this question, we completed an investigative study where we used a Keysight N5227A PNA-67GHz Microwave Network Analyzer fitted with two V-band adapters connected to WR-15 waveguides. A single playing card was sandwiched between the waveguides following the process in [7] to extract the permittivity of the cards, namely, a figure of merit quantifying how the cards react to electromagnetic energy.

Once the permittivity of a card was extracted, we ran an exploratory electromagnetic study, where we placed two half-wavelength separated printed patch antennas at a 6mm from a numerical computer aided design (CAD) model of one playing card. The 6mm separation was chosen following the distance between the Soli antennas and the plastic enclosure of the Soli developer kit. Next, we started changing the number of cards in simulation. This allowed us to study different scattering parameters when energy is emitted from one of the antennas towards the playing card deck, and we observe the energy scattered by the deck to be received by the other antenna. The electromagnetic simulation was conducted using a finite element numerical solver available from Ansys [3]. Figure 2 (a) shows the CAD setup for a single card case and Figure 2 (b) shows a multiple-card stack. The 3D tool allows us to position the cards exactly as in real-life use.

The physical scattering behavior between the two antennas in presence of different cards can be assessed by comparing the amount of power received by the receiver antenna with the amount of transmitted power by




Fig. 2. (left) Simulation setup in Ansys (right) Scattering parameters between the two antennas.

the transmitter antenna. This power ratio S_{RT} can be represented on a decibel scale (Figure 2 (c)). The phase difference between the transmitted and received signals can be represented as $\angle S_{RT}$ (Figure 2 (d)). For these two figures, one can note that across a wide band of frequencies, there is a clear difference between different thickness of cards (i.e., 1 card, 2 cards, 3 cards, etc...). The noted scattering parameter differences between this one antenna pair, implies that a multi-channel low-cost system like Soli (featuring eight channels/antenna pairs) should be capable of identifying a number of cards once the system was trained to recognize their scattering signature. This forms the basis for our investigation of planar (and 3D) objects for tangible interaction in HCI.

4 SOLINTERACTION PROTOTYPE

For the hardware in Solinteraction system, we use the Google ATAP Project Soli alpha developer kit [20, 32]. Soli [32] is a mono-static, multi-channel (8-channels, 2 transmitters and 4 receivers) radar device, operating in the 57-64 GHz range (center frequency of 60 GHz), using frequency modulated continuous wave (FMCW) where the radar transmits and receives continuously. Soli [32] is designed for capturing micro finger motion for enabling interaction with computers but RadarCat [68] exploited it for object and material recognition. To improve reproducibility, the approach in RadarCat [68] has also been shown to work with commercially available systems such as the Infineon Technologies BGT60 millimeter-wave RF transceiver [35]. In this work, Solinteraction greatly extends the capabilities previously demonstrated in RadarCat [68] with six sensing modalities, enabling new forms of tangible interaction.

4.1 Radar Signal and Machine Learning

When object(s) are placed on the top or nearby the Soli sensor, the energy transmitted from the sensor is absorbed and scattered by the object, which varies depending on its distance, thickness, shape, density, internal composition




Fig. 3. From left: 8-channels raw radar signals for 1, 2, 3, 4, 5, 10, 25, 50 playing cards placed on top of the Soli sensor.

and surface properties, commonly described as the radar cross section (RCS). As a result, the signals that are reflected back represent rich information and the contributions from a range of surface and internal properties. As these raw signals (Figure 3) are unique to different objects and the different configuration of the objects (stack, order, orientation), we can exploit these to enable new sensing mechanisms. The basic sensing layer in Solinteraction relies on recording these multi-channel raw radar signals, training and classifying these signals using a machine learning technique. Following a supervised learning approach, our software extracts a set of features, and classifies objects which have different physical characteristics without requiring any instrumentation. Our basic implementation follows the approach described in [68]. We use the raw signal from 8 channels (512 features) and extract extra statistical features (absolute and root mean square for each channel and all channels, global maxima, global minima, mean) resulting in a total of 661 features. In addition, we also employ the range feature provided in the Soli SDK. We employ a moving average of 10 frames within a 1 second window, where the sensor is running at 300 fps. The features are fed into the Weka toolkit [21], both for classification and regression in real-time. Finally, the classification result is stabilized by taking the most common classification amongst the last 10 outputs. Our final system and classification run in real-time on a low-power Intel NUC with an i3 processor.

4.2 Token+Constraint

For the hardware constraints in Solinteraction, we used a 3D printed case and swappable modules. Note that the six sensing modes described in the next section work *without* such constraints (Figure 4) but the sensing accuracy is enhanced with them. The case and modules act as channels for constraints within the Token+Constraint paradigm [56]. As articulated by Ullmer et al. [56],

“Constraints are confining regions that are mapped to digital operations. These are frequently embodied as structures that mechanically channel how tokens can be manipulated, often limiting their movement to a single degree of freedom.”

In Solinteraction, such constraints serve as well-defined sensing zones, help to affix the object in position and afford numerous benefits including degrees of freedom, compatibility, embodiment, reduced ambiguity, discrete and continuous interactions. Here we further benefit from a reduced amount of training data required for all dimensions.

4.3 Fabricating Tokens and Constraints

Informed by previous work on adding constraints [49, 54, 56] and modifying the internal structure to alter acoustic [25] or optical properties [64], our fabrication approach is motivated by two goals. First, we want to explore different kinds of designs and how the internals and externals of objects affect the radar response, so that they can be uniquely differentiated, even though they are fabricated using the same material. Secondly, we want to quickly and cheaply create snap-on modules that can be easily swapped out for augmenting different functionality (akin to a lens mount) without replacing the entire case. In particular, a case holds different modules,




Fig. 4. 3D printed cases and modules for guiding objects and signals, e.g., numpad, rotation, sliding, Lego holder, card holder and slotting disc holder.

and tokens are constrained within the opening of a module. Thus, we explore how these modules can afford various types of interaction with primitive building blocks (tokens) such as Lego or even the users' finger.

We 3D print the tokens (varying shapes and holes) and constraints (varying openings and support structures) using the Replicator 2X printer. For adding tags, we primarily apply copper tape but other materials are possible, such as conductive ink, or by etching a PCB. While in this work we rapidly prototype with 3D printing, other means of low-cost manufacturing techniques, such as machining or injection molding, are possible.

Literature in instrument design [16, 25, 49, 64], informs our approach to the design of the add-on modules. Through a structured exploration, we identified a range of properties and geometries that can be altered during the fabrication process to allow us to control and isolate the radar effects, such as a hole, shape, thickness and height. Our mechanical mechanisms allow us to explore the effects on radar signal propagation pathways. From this, we are able to determine a series of design recommendations which affect the scattering and reflection of the transmitted signals (Figure 5). To further aid this exploration, we measure the physical dimensions of the device and objects and model these with CAD software (e.g., Tinkercad). From our experience, if objects are larger than the sensor area and are solid, then they can be placed on top of the sensor directly without constraint. However, for smaller objects, these can be affixed within the sensing area to improve recognition rates.

5 SENSING CAPABILITIES OF SOLIINTERACTION

The ability to classify materials and objects in real time enables the use of many physical items as a way to interact with computation. This is achieved without altering objects and as such allows existing objects to form components of new types of interaction. Once the objects are in proximity of the requisite sensing capabilities we can enable triggers, cues, actions and entirely new forms of tangible interaction. Specifically, a non-tagging (i.e., non-invasive) method such as RadarCat [68] opens up new potential that we wish to further explore.

By extending this technique, not only to static individual objects, but to multiple objects simultaneously and in dynamic situations (e.g., movements), we can identify a design space, not previously explored. This paper describes six sensing modes: the (i) count (ii) order of stack (iii) identity (iv) orientation (v) movement and (vi) distance of object placed on and around the Soli sensor. These perspectives collectively constitute the design space of radar-sensing for tangible interaction - a "*hidden dimension of Soli*" that we seek to explore and exploit.

5.1 Counting

As the radar signal is able to penetrate through non-metallic material, and the returned signal is highly characteristic depending on the layers of thin objects, we can leverage this to enable the counting of number of visually and physically similar thin objects (e.g., paper sheets, playing cards, poker chips or Lego blocks) stacked on top




Fig. 5. Our fabricated designs and corresponding radar signals which are unique to each design. The y-axis’s range is ± 800 for each graph.

of the sensor (Figure 6 (a)). Prior work [8, 9, 14] has recognized stacked objects using optical sensing but cannot support more than a few items. Our system relies on the fact that the radar signal can penetrate multiple layers, where the returned signal through multiple layers can be seen in Figure 3. A limitation of this counting technique is that objects should not fully obstruct the radio signal propagation, for example, stacks of metal sheets cannot be counted. Counting with radar opens up new possibilities, for example the amount of sand in an hourglass can be “measured” too, as demonstrated in Figure 6 (b) and the accompanying video figure.

5.2 Ordering

The order in which items are stacked can also be recognized, because our technique is based on the analysis of the signal corresponding to the entire group of objects rather than of each individual item, much like structural transparency [8, 9, 14]. We experimented with different credit cards, cup coasters and Lego blocks stacked in different orders. The x number of ways n cards can be ordered, including sub-stacks, can be calculated with the Equation 1 [8].

$$x = \sum_{r=1}^n \frac{n!}{(n-r)!} \quad (1)$$

As detailed in our evaluation section, using 4 credit cards we were able to distinguish all the 64 different combinations (including sub-stacks) with very high accuracy when using a card holder, but with lower accuracy without. We can also recognize different combinations of Lego blocks stacked on top of the sensor (with constraint). As shown in Figure 7 (b) and video figure, four “double” bricks combined with two “single” bricks, that are placed at different levels can be uniquely identified.

5.3 Identification

Prior work with radar has demonstrated that proximate objects can be distinguished [68] with over 96% accuracy. However, the study was limited to objects with different materials or thicknesses. To push the boundary of such sensing, we attempt to recognize objects of the same material, but with varying geometry (e.g., shape), internal structure (e.g., holes) or surface properties (e.g., texture, bumps). Our initial exploration shows that unmodified




Fig. 6. From left: (a) counting playing cards (b) counting sand in hour glass matched by timer on the smartphone (c) identify Domino pieces (d) identify touch location on a numeric keypad sheet as input for calculator (e) tracking rotation of tagged game piece to control the clock hand and (f) tracking rotation of untagged smartphone using regression.




Fig. 7. From left: (a) inside-out movement tracking on a desk made from compressed wood (b) ordering Lego blocks (c) distance tracking of three mugs made from different material (d) music player by slotting chips at different height (e) casino application with multiple players and (f) mathematics learning applications.

daily objects with varying geometry (e.g., cutlery) can be easily recognized, even though they are made from the same material. Interestingly, we found unmodified, yet similar daily objects (e.g., credit cards from different banks or Domino pieces) can also be identified with very high accuracy. This is largely due to two factors - the surface coating and the embossed marks.

For a single material such as the plastic filament used in 3D printing, we explored various designs that can be altered during the manufacturing process. For example, we 3D printed pieces with i) holes ii) surface texture or bump and iii) altering the shape, as shown in Figure 5 with its corresponding radar response. In addition, the exact same object can be minimally modified and made differentiable, specifically by adding reflective material (e.g., copper tape) or by etching, as shown in Figure 1 (bottom). Each design acts as a waveguide, altering the returned signal. This further extends the range of tokens or physical icon (“phicons”) [55] which we can explore.

Furthermore, with the help of a thin tactile guidance layer (Figure 4 (b)), the location of a finger touching the surface can be identified. It also works with non-flat surfaces such as a hemisphere (Figure 5), which effectively turns Soli into a low-fidelity touch sensor. With touch locations identified, further information regarding the touch (e.g., direction, orientation and velocity) can be extracted.

5.4 Flipping

The orientation (facing up or down) of objects can be detected. Many planar objects can be flipped and sensed, if the opposite side has different surface properties (material, shape, texture, etc). Examples include, a smartphone, credit card, dice, mirror, book, Lego block, Domino piece, cutlery and crockery. Flipping can enable a binary control mechanism such as turning on or off a light. While knowing the orientation of an object could also be useful for privacy and context awareness, such as setting the phone to silent mode when facing down or

folding a card in Poker games. Flipping further affords the exploration of symmetric and asymmetric interaction techniques.

5.5 Movement

It is also possible to distinguish different parts of the same object due to the varying internal structure of the object, depending on the area being considered, which in turn produces different radar responses. This knowledge allows us to train and infer the location of the object relative to the sensor (Figure 7 (a)) in real-time, hence also track the movement of the object, which can be used to enable fluid and dynamic interactions. We have explored the possibility of using objects as continuous controls, such as sliding or rotation input (Figure 6 (e) and Figure 6 (f)). This technique also allows us to use objects as multi-purpose dynamic input controllers. A wide range of movements (step-wise, ordered, patterned and continuous), direction, velocity, acceleration and force applied can be considered in the design of new forms of interaction.

In the current prototype, there are two ways to achieve movement input, for which we borrow terminology from computer vision and classify them as outside-in and inside-out tracking. With outside-in, we can use a composite object sliding on top of the sensor, and infer the location of the object relative to the sensor. Whereas with inside-out, we use the sensor itself as a probe to scan a surface (e.g., a desk made from compressed wood or a wall with hidden pipes). By extending movement to two dimensions (x and y), it is possible to achieve 2D positioning, akin to using a low-resolution computer mouse.

5.6 Distance

Solinteraction also supports recognition of objects placed from a distance, where we dynamically detect the distance of the object from the sensor. For example, the distance of the recognized object made of different materials (e.g., a mug made of plastic, glass or metal) to the sensor can be measured and used as a continuous one-dimensional input (Figure 7 (c)). A small object (e.g., a disk) can be placed (e.g., slotted) at a different height from the sensor with the help of a constraint such as a 3D printed case, which triggers different actions, as shown in our Jukebox application (Figure 7 (d)). Each object can correspond to different controls such as adjusting volume or playback speed. Multiple objects can also together correspond to a single control such as each mug controlling an aspect of the color (red, green or blue) of the environment's lighting.

6 EVALUATION

Prior to the development of such interactions in practice, an evaluation of each mechanism is required to validate these new interaction building blocks. To evaluate the accuracy and feasibility of our approach, we conducted various experiments, covering a broad range of use cases and objects which exemplify different design dimensions. However, the combination of our mechanisms, objects and number of objects results in a combinatorial explosion of possible experiments. Instead, we report on the close and careful study of Solinteraction with a series of representative cases. In this paper, our contribution is to specifically and selectively focus on (i) counting (ii) ordering and (iii) identification by collecting rigorous experimental data. Whereas for the rest of the experiments - (iv) orientation (v) movement and (vi) distance, we collect a less exhaustive dataset but remain sufficient for showing the accuracy and feasibility. All the experiments and results are summarized in Table 1. These experiments show our system requires little training data to support the sensing and interaction. And for each experiment, we also visually report on its performance, shown in real-time in our video figure. Our dataset is publicly available at <https://github.com/tcboy88/solinteractiondata>. To avoid over-fitting, we average over 10 frames as single data input. We then apply offline analysis using Weka [21] toolkit, by default using the Random Forest (RF) classifier with ten-fold cross-validation unless otherwise specified, e.g., leave-one-out and support vector machine (SVM) classifier in some cases for more thorough comparison. We chose objects such as playing

cards, Poker chips and Lego blocks, as we believe these are very common and representative objects used in tangible user interface.

6.1 Procedures for Data Collection

6.1.1 Counting. We start with counting (i) **52 playing cards**. Two authors take turns to collect an extensive 60 rounds of data (6 sessions x 10 rounds, each session was collected at different times of day, which took 3 days). In each round of data collection, 52 cards are placed on top of the Soli and one card is removed at a time until no more cards remain. During each card removal, 10 frames of the raw radar signal are averaged and saved. For counting (ii) **20 sheets of paper** (A4, 90gsm), a similar procedure is used, except sheets are added to the stack instead of being removed. 10 rounds of data are collected. Due to the large surface area of the paper, support structures are placed around the sensor to reduce paper curling. For counting (iii) **20 poker chips**, we use 2 conditions: with and without a constraint (3D printed case with circle opening, Figure 4 (i)) to hold the chips in place. We decided that more than 20 chips (roughly 5cm in height, made of solid plastic) are considered very challenging for the low-power radar signal to penetrate. 30 rounds of data for each condition are collected, where in each round, each chip is added on top of the sensor. For counting (iv) **5 Lego blocks**, initial testing of counting them without using a constraint to hold them resulted in low accuracy, mainly because a Lego block has complex internal structure (studs on top and a round hollow tube on the bottom) that reflects the radar signal quite differently at different positions with respect to the radar antenna at such close distance. Thus we only evaluated it by using a constraint to hold them in place and collected 10 rounds of data. For counting (v) **sand in an hourglass**, we do not evaluate its accuracy but instead demonstrate this in the video figure.

6.1.2 Ordering. For ordering (i) **credit cards**, we chose 4 cards that have different radar signatures by visually inspecting the radar response using our GUI visualizer. As there are various designs of credit cards, e.g., 3 out of the 4 cards that we chose have raised embossed characters and 1 card is without. We chose cards that do not have metallic coating, to allow the radar signal to propagate multiple layers. We use two conditions - with and without a constraint (3D printed case with square opening, Figure 4 (g)) to hold the cards in place. With 4 cards there are 64 possible combinations (Equation 1). For the sake of completeness in this evaluation we tested all of them. The procedure when using a constraint is as follows: place each card combination once in the case, and save 1 dataset, until all combinations are collected. This process is repeated for 6 rounds and evaluated using leave-1-round-out validation process that is repeated for all 6 rotations. The procedure without using constraint is as follows: for each combination, repetitively remove and place it on the sensor at slightly different positions each time, for 10 times, which is evaluated using ten-fold cross-validation. We use a similar procedure for ordering (ii) **coasters**. In addition, we combine both counting and ordering with the (iii) **Lego blocks**.

6.1.3 Identification. For identifying (i) **10 credit cards** with and without a constraint, the procedure is similar: repetitively remove and place card on the sensor at the same position (with constraint) and slightly different position (without constraint) each time, for all 10 cards, for 10 rounds. The procedure is similar for identifying (ii) **7 Domino blocks** (iii) **12 3D printed designs** and (iv) **6 tagged cards**. For identifying (v) different **color chips**, we gather 20 poker chips of 5 colors each, and collect data once for each chip with a constraint, for both sides. For identifying (vi) **users** and (vii) **finger touch locations**, we recruited 6 participants from our department. The process took roughly 30 minutes and they were compensated with a voucher worth 5 pounds. We ask each participant to repetitively place their upper palm (finger area) on top of the sensor with a case but without module (Figure 4 (a)), and data is collected for 30 times. For identifying touch locations, we use a 3D printed sheet to provide tactile guidance (see Figure 4 (b)). We collect data for 12 locations (the middle 3 columns x 4 rows, which represents the layout of a numeric keypad), once for each location, and repeated the process for 10 rounds. We ask each participant to place their left index finger in each position, without any limitation in their hand

posture. We then repeat the touch experiment with a different module - a hemisphere with 9 locations (Figure 4 (b)), which has no tactile landmarks but each location is marked with red ink.

6.1.4 Orientation. For evaluating orientation (flipping of the object), we collect data from the rear of (i) **10 credit cards**, and combine this with data from the previous experiment (identification), therefore we have data for both sides of the credit cards. (ii) We also collect both sides of a deck of **52 playing cards**. We do not further systematically evaluate flipping with other objects but we tested with success on flipping phones, cutlery or opening a book.

6.1.5 Movement. For (i) **rotation**, we place a poker chip on the constraint (3D printed case with an open circle in it, Figure 4 (d) and Figure 6 (e)). Since the poker chip is uniform, it has to be tagged with a small piece of copper tape, which alters the radar signal when the chip is rotated. We then collect data for 12 intervals at 30 degrees each, twice. First we start at 0 degree and second starting at 15 degrees. This approach allows us to explore regression to infer angular movements, for which we do not have training data (further explained in the results section). In addition, while we did not evaluate rotating a phone we demonstrate its feasibility in the video figure. For (ii) **outside-in sliding**, we use a similar procedure except with a 3D printed and copper tagged rectangular bar, for 10 intervals at 1cm each. For (iii) **inside-out sliding**, we slide the sensor along a desk surface with the same 10 intervals at 1cm each.

6.1.6 Distance. Holding an object at a stable distance from the sensor proved difficult, due to hand tremor and the high sensitivity of the radar with respect to small movements. Therefore, for experimentation we rely on two approaches: (i) use of a constraint (Figure 4 (h)) that permits small objects to be placed (in slots) at different heights (ii) placing the sensor facing sideways on a table, so that objects can be stably placed at different distances from the sensor. For (i), we slot a **poker chip at 7 different heights** and we collected 10 rounds of data. For (ii), we chose 3 objects of different materials, specifically **a steel mug, a glass and a plastic bottle**. We then place each object at different distances (1 inch to 10 inches, measured by a ruler) and collected 3 rounds of data. We then attempt to classify the distance. Later, instead of distance classification, we also group all distances from the same object into the same class and simply classify them by object type. In addition, we used the range feature from the Soli SDK, however we did not evaluate its accuracy due to its black-box implementation.

6.2 Results and Discussion

The results shown in Table 1, are both with and without a constraint (the 3D printed case and modules). For those with very high accuracy we omit the confusion matrix for brevity, while those of note can be seen in Figure 9 and Figure 10. Of particular note are the ordering results with 4 credit cards experiment which show a 43.14% difference with and without the constraint (56.40% vs. 99.54%), when using Random Forest (RF) classifier. This shows that a constraint indeed helps improve the sensing accuracy. In counting though, the improvement when using the constraint is less pronounced at approximately 7% to 17% (64.39% vs. 71.27% using RF and 63.12% vs. 80.16% using SVM). Some experiments already work very well without a constraint (e.g., counting and flipping playing cards) and yield very high accuracy, thus we decide not to repeat the study with a constraint, where we can expect higher accuracy. Some experiments require a constraint in order to work properly, because the small but irregular shapes affect the signal considerably even with small movements (e.g., Lego blocks, printed designs, tagged designs and rotation). In addition, a constraint is needed to hold the tactile guidance sheet (modules). However, some experiments are less practical with a constraint due to the large surface area (e.g., counting paper sheets and ordering coasters), where the confusion matrix can be seen in Figure 9.

We typically employ a Random Forest (RF) classifier in our real-time demo and off-line evaluation due to its fast speed and low memory footprint. However, a more complex machine learning classifier might yield better accuracy. For example, a Support Vector Machine (SVM) constantly outperforms Random Forest (RF) in most

Table 1. Classification accuracy of different experiments on different conditions, by default using Random Forest classifier with 10-fold cross-validation unless stated otherwise in parentheses, RF - Random Forest, SVM - Support Vector Machine.

<i>Experiment</i>	<i>Test Conditions</i>	<i>Without case</i>	<i>With case</i>
Count	52 playing cards (10-fold CV)	99.50%	n/a
Count	52 playing cards (Leave-1session-out)	99.53%	n/a
Count	20 paper sheets (RF)	77.62%	n/a
Count	20 paper sheets (SVM)	88.57%	n/a
Count	20 poker chips (RF)	64.39%	71.27%
Count	20 poker chips (SVM)	63.12%	80.16%
Order	3 credit cards (RF)	85.37%	100%
Order	4 credit cards (RF)	56.40%	99.54%
Order	4 credit cards (SVM)	62.05%	99.54%
Order	3 cup coasters (RF)	79.17%	n/a
Order	3 cup coasters (SVM)	81.67%	n/a
Count+Order	5 double + 2 single Lego blocks	n/a	96.23%
Identify	10 credit cards (different banks)	87.27%	100%
Identify	7 Domino blocks	n/a	99.29%
Identify	12 3D printed designs	n/a	100%
Identify	6 tagged cards	n/a	100%
Identify	5 color chips (RF)	n/a	75.00%
Identify	5 color chips (SVM)	n/a	87.08%
Identify	12 touches on numpad (10-fold CV)	n/a	98.59%
Identify	12 touches on numpad (leave-1user-out)	n/a	94.87%
Identify	9 touches on half-sphere (10-fold CV)	n/a	98.50%
Identify	9 touches on half-sphere (leave-1user-out)	n/a	95.50%
Identify	6 users palm (RF)	n/a	85.71%
Identify	6 users palm (SVM)	n/a	91.90%
Identify	6 users touches (grouped by user)	n/a	82.82%
Orientation	10 different credit cards	83.18%	100%
Orientation	52 playing cards	97.14%	n/a
Movement	Rotation (classification)	n/a	100%
Movement	Rotation (regression)	n/a	99.78%
Movement	Sliding (outside-in, ruler)	n/a	100%
Movement	Sliding (inside-out, desk)	100%	n/a
Distance	Slot chip at 7 different heights	n/a	100%
Distance	3 mugs (10 distances)	44.09%	n/a
Distance	3 mugs (grouped by material type)	83.87%	n/a

experiments where the improvement in accuracy can be up to 11% (77.62% vs. 88.57%) in counting papers and 9% (71.27% vs. 80.16%) in counting poker chips. In ordering and identification, the difference is between 2% to 6% and




Fig. 8. (left) For counting 20 chips, the accuracy decreases with increasing number of chips, with and without using a case. (right) For ordering 4 credit cards **without** a case, the accuracy decreases with increasing number of permutations. Data points only at air 5, 11, 16, 17, 41, 65 permutations (including air), by removing certain cards from the dataset.

6% to 12%, respectively. Further exploration of different machine learning classifiers (e.g., deep neural network) is left for future work.

Our exploration demonstrates that items with a uniform shape afford the most reliable counting. As our results show, we can count playing cards very accurately, even **without** using a constraint. This is due to the nature of playing cards which are flat, rigid and have similar material across the entire surface. The cards also cover the field of view of the radar well, which can be observed when the card is moved slightly and the signal remains unchanged. Of note is the very thin layer of color printing does **not** impact the signal sufficiently to confuse the recognition - hence we also cannot identify an individual card. In contrast, sheets of paper and banknotes are too flexible, they curl when placed on top of the sensor which allows non-uniform air-gaps to be formed between the stacked layers, making counting of multiple layers less accurate, but possible nonetheless under controlled setup. For thicker objects (e.g., poker chips), the accuracy decreases with increasing numbers, as expected. For example, in counting chips, under 5 chips there is not much difference in using a constraint or not, as they are both above 90% accuracy. However, we can see from Figure 8 (left) that with increasing number of chips, the effect of constraint comes in. With a constraint, we maintain above 80% accuracy (SVM) up to 20 chips. Without a constraint, the accuracy drops to near 60%.

The accuracy measured when ordering 4 credit cards using a constraint/case, was very high (99.54%) with only two errors. Further analysis of the data identified one of these errors as human error, due to mislabeling the data during collection of the numerous permutations explored. We also identified that certain credit cards work well for this ordering experiment but this is not true for all cards because banks use different coating. Some cards with metallic coating will reflect most of the radar signal, making detection of multiple layers difficult. Ordering without using a constraint proved to be accurate when there are fewer combinations (over 85% accuracy with 3 cards, 15 combinations including sub-stacks). However, as the number of combinations grows, the accuracy drops significantly, reaching 56% to 62% when all 4 credit cards (64 combinations including sub-stacks) are tested. A plot of accuracy vs. number of combinations can be seen in Figure 8 (right).

Identification with both existing objects (credit cards, Domino blocks) and fabricated objects (printed designs, tagged cards) were very accurate (near 100%), with one exception in identifying 5 color chips where the accuracy is between 75% (RF) to 87% (SVM) with a confusion matrix shown in Figure 10 (left). The relative lack of accuracy

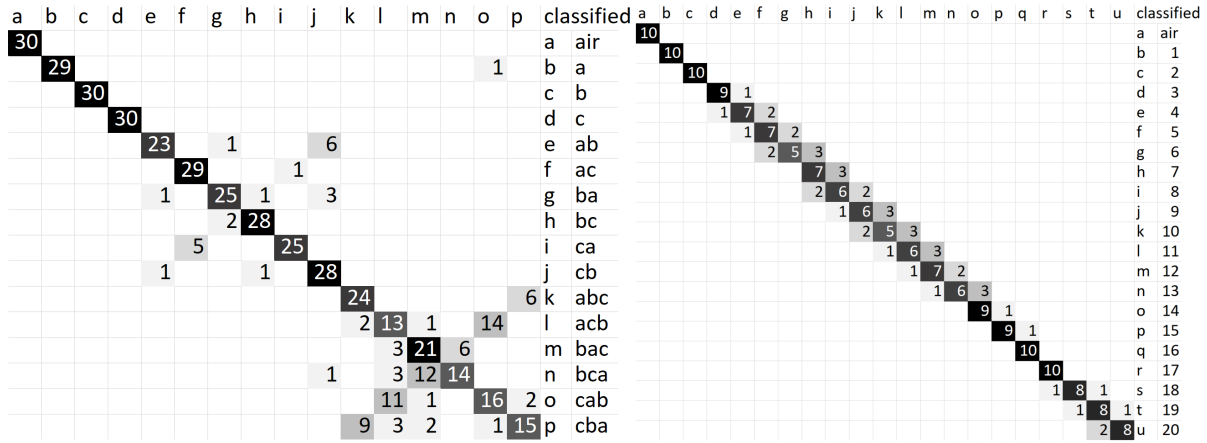


Fig. 9. Confusion matrices for different experiments: (left) ordering 3 coasters (RF) and (right) counting 20 papers (RF).

is largely due to the uneven surface and imperfections in manufacturing, as the chips we used are low-cost toy product rather than casino quality.

We can identify finger touch locations on flat (98.59%) and non-flat, hemispherical (98.50%) surfaces, even though we observed that participants use different postures (touch with a pointed, inclined finger from the top or a flat, horizontal finger). Using a leave-one-user-out evaluation (trained with $n - 1$ participants and tested with the remaining one participant), the average results are 94.87% (flat) and 95.50% (non-flat, hemisphere). This suggests that with enough training data, we can generalize to other users without needing to collect and train with a new user every time. Future work aims to extract multi-touch, finger pose and orientation.

We can also identify different participants by their upper palm with rather promising accuracy, from 85.71% (RF) to 91.90% (SVM), where the confusion matrix is shown in Figure 10 (middle right, RF only). However, the palm data (30 rounds) is collected within the same session, although we do ask participants to pause and stand up after 10 rounds. As such, it is worth highlighting that the recognition might not generalize over time or session.

Most surprisingly, we also found that by grouping all the touches by participants (ignoring location information), it is possible to differentiate the user with rather good accuracy (82.82%) where the confusion matrix is shown in Figure 10 (rightmost, RF only). We suggest this is because users have different touch behavior when using a Numeric keypad (numpad), and these can be extracted for additional input during an authentication process. P4 has the lowest precision potentially because she wore a metal ring that might have affected the radar reflection.

In flipping credit cards, the accuracy increases from 83.18% to 100% when the constraint is being used. It is because the radar signal reflected by a credit card varies depends on the spot being observed, due to uneven surface, embossed letters, integrated circuit and surface coating material. In the 10 rounds of data collection without constraint, the card was being randomly placed at a slightly different position each time, therefore the lower accuracy. When using a constraint, this issue is virtually gone. In contrast, flipping playing cards is very accurate (97.14%) even without using a constraint, because of the similar signal response on any location, as explained previously.

In classifying movement (rotation and sliding), we achieve 100% accuracy with our 10 rounds of data collected using 10-fold cross-validation. We also evaluate rotation by using regression, achieving a high correlation coefficient of 0.9978 with a mean absolute error of 0.15 and RMS error of 0.27. We also collected a new set of data at half intervals and use it as a test set (not trained) for regression to estimate these values. In this case, the




Fig. 10. Confusion matrices for different experiments: (left) counting 5x20 color chips (most left: SVM, middle left: RF) and (right) identify users (middle right: by palm, most right: by touch behavior, both using RF).

correlation coefficient is lower at 0.9327 with a mean absolute error of 1.02 and RMS error of 1.35. However, we found that only a single position constantly gives a very high error ($M=3.7$). By removing this outlier position from the result, the mean absolute error decreases to 0.77 and RMS error to 0.85. Therefore, it is possible for Solinteraction to infer values at positions that have not been previously trained using regression.

For sensing distance with our constraint slotting scheme, we are able to recognize the height of the object being slotted very accurately, albeit this is limited to fixed intervals and to a single type of object. We extend it to multiple types of objects (3 mugs with varying materials - metal, plastic, glass) and larger distances (from 1 to 10 inches away) without using a constraint, but then the recognition rate was only 44.09%. We then discard the distance information and group the observations into material only, and the accuracy increases to 83.87%. Therefore, we decided to employ the range feature provided in the Soli SDK. Since radar is not a vision sensor, and we are only using the raw signal without high levels of abstraction and processing, we cannot reconstruct the exact shape of the object from a distance. However, we can tell the material of the object, based on the overall reflected signal response.

7 DESIGN ATTRIBUTES OF SOLINTERACTION

The design of interactive systems can harness the six Solinteraction sensing mechanisms in a variety of ways. Individually they can be mapped to simple forms of interaction with objects in terms of movement, flipping or ordering, etc. Such simple mechanisms can afford common and familiar interactions with everyday objects.

Solinteraction affords basic interactions but it is the combination of mechanisms which affords new opportunities in the synthesis of new forms of interaction. By imagining how interaction with everyday things might be coupled to computing, based on the six building blocks of Solinteraction for creating radar-driven interactive elements, here we detail five second order dimensions, **Mixed-mode**, **Expressive**, **Discreet**, **Interleaving** and **Context (MEDIC)**, which describe this current design space of radar-based interaction.

DA1: Mixed-mode tangible & gesture interaction. When there is no object or only a thin layer is placed on top of the sensor, the mid-air gesture sensing (the original purpose of Google Project Soli) can continue to function, thus allowing mixed-mode interaction. Combining both types of input modalities greatly expands the possible interaction with a single, generic sensor supporting multiple purposes.

DA2: Expressive tangible interaction. Tangible interaction is not limited to a single token, but multiple tokens can be added or stacked in different orders to enable more expressive input. For example, 3 tokens will allow 15 combinations (a, b, c, ab, ac, ba, bc, ca, cb, abc, acb, bac, bca, cab, cba), whereas increasing to 4 tokens increases the combinations to 64 (Equation 1).

DA3: Discreet and subtle interaction. As radar signals can penetrate through material such as clothing, interaction with devices can be very subtle or even hidden from plain sight, e.g., when the device is covered by

cloth or hidden inside a box. Users can then perform subtle gestures or employ nearby physical objects that act as tangible tokens for interaction.

DA4: Interleaving static and dynamic interaction. An object can be recognized when placed statically on the sensor or touched by the sensor as if using a probe. In addition, an object can be moved and continuously tracked. This means a single object can serve as a site for static and dynamic interaction for multi-purpose input control.

DA5: Contextual commands. Instead of simply recognizing tokens placed on top of it, the sensor can be placed downward to sense the current surface, which could enable a mobile device to switch its context automatically, e.g., on a kitchen table it will show the most relevant cooking app or recipe. Thus, any surface now becomes a canvas for exploration.

8 APPLICATION SCENARIOS

Solinteraction offers designers the creative freedom to design many new forms of tangible interaction without tagging and with stacks of planar objects. The ability to sense everyday objects, from a small device, combined with interaction allows for numerous “interaction zones” around sensing surfaces. Examining these new capabilities within the design space allows us to consider which techniques might be suitable for particular scenarios in an interaction zone. Here, with reference to the **MEDIC** design attributes, labeled DA1-DA5, we describe a range of application scenarios which can help designers understand the relative merits of the techniques proposed. To achieve this, we created two example applications which couple and combine our basic mechanisms into a single blended interaction. We then further populate the identified design space with three speculative scenarios, where we might expect this kind of sensing to be commonplace in the future.

8.1 Board and Card Games

We can consider our system in use in traditional board games (e.g., Settlers of Catan) or professional card game environments such as cards (e.g., Poker or Blackjack), playing pieces (e.g., Domino or Mahjong), where the sensing is embedded beneath the surface, for example in a playing board or casino. This can enable a system with Solinteraction to automatically count cards or chips (DA2), plotting the identity of Mahjong or Domino pieces (DA4), determining which side of dice is facing up (DA2), thus avoiding cheating or human mistakes (DA3). In a Poker game, it could also recognize the player’s turn (DA1), or automatically update the game without player or dealer intervention (DA3). Since card flipping is also recognized, this could enable a seamless play experience (DA3). Players can take turns by placing their hand on the sensor and being authenticated (Figure 7 (e) and video figure) (DA4). Players can also touch, knock or swipe on the sensor which are common casino gestures for hit, stand or split, respectively (DA1).

8.2 Learning and Education Tools

Such interaction may also enable education and learning applications to be more engaging and fun for children. For example, in a simple mathematics application, children could be required to place the correct number of objects or cards to solve a question (DA2). Different gestures (DA1), or as shown, objects can be used to represent mathematical symbols such as addition, subtraction, multiplication or division (Figure 7 (f) and video figure). The ordering of pieces in puzzle development can be determined to ensure the correct mathematical operation sequence is followed (e.g., multiplication before addition) (DA2). Understanding the importance of the sequence of moves is also important in games such as the Tower of Hanoi or the Soma cube. Such information can also be used in games which require dexterity (DA4) and careful order of play piece movements such as with Jenga.

8.3 Assembling Models

The ability to sense Lego blocks of different sizes, thickness, shapes and how they are stacked opens up considerable potential for interaction (DA2), in part, due to how common Lego is and the many combinations of how Lego blocks can be combined. For example, large model kits (e.g., Lego spaceship or Gundam) usually consist of many pieces, sometimes over a thousand of them. Assembling these model kits is a complex task, where commonly each part is being assembled separately and joined together later. Our system can examine if each subpart is correctly assembled, before joining to the base model for completion (DA2). While moving or manipulating aspects of a physical model (DA4) can afford new forms of interaction. Outside a toy scenario, this might be adopted for assembly line in factories (e.g., electronics or automotive).

8.4 Smart Home and Office

Smart work and home environments offer a wide range of opportunities for Solinteraction. For example, context-aware computing where devices are placed on different surfaces (working desk or a kitchen table) and different interaction zones, tailored personally to this context can be used (DA5). Imagine a smart desk that can detect whether a book is flipped open, followed by a simple gesture (DA1). We can then count the number of pages being flipped (DA2) and save the bookmark automatically or display companion content of different pages on another screen or through augmented reality (projector or head-mounted display). The dynamic placement and combination of ornaments within a home can control lighting, heating and the audio setup (DA4). Physical items can be linked to audio, video and multimedia items or presentations, e.g., the jukebox example we created based on slotting discs (DA2). Thus, the system brings life to existing pieces - a cup naturally becomes a rotating controller for adjusting the volume (DA4). We might also consider a social setting where using digital devices in front of other people appears to be rude (DA3). People sometimes politely turn their phone upside down on the table to show respect, yet they still want to interact with it in subtle and discreet ways (DA3). Users can then use nearby physical objects (e.g., office stationery, cutlery) that are disguised as tangible tokens for interaction, to control simple tasks (e.g., start a voice recorder or silence the phone) (DA2).

8.5 Retail or Dining

This setting can be tailored to specific retail or dining scenarios. For example, a restaurant presents a range of opportunities for new forms of digital-physical interaction. In many restaurants, typically fast food and sushi restaurants, the plates are color-coded to represent the price. Similarly, different plates can be tagged and if the user wants to order more food of this type, they just place this token on the sensor, and the order will be processed automatically (DA2). Placing a menu face down might call for a waiter while a specific drink coaster might be tied to a cocktail, wine or beer selection (DA3). Moving a bottle to the edge of the table might indicate more wine (DA5). While stacking a set of sushi plates might indicate the customer is leaving and trigger a bill (DA2). A gesture (DA1) might call for the bill, while physically ordering credit cards may allow a group of friends to split the bill unevenly, where the bottom card pays for a bigger portion (DA2).

8.6 Other Potential Applications

Soli is a low-cost sensor that can be embedded in everyday devices. Further development of such sensors with higher bandwidth or at different frequencies (e.g., Terahertz), means more powerful sensing modes can be achieved. We suggest that with deep learning approaches, the system can scale to a even wider range of daily objects. Here we envision future applications that are intended to show the potential of Solinteraction when such sensing becomes ubiquitous.

Biometrics - user authentication: To differentiate and authenticate users by their hands [17], much like blood vein recognition, might be feasible. We have shown that it is possible to identify different users just by

touching the upper palm on the sensor, but the accuracy is limited (85.71% to 91.9%). In parallel, Diederichs et al. [17] show more than 98% accuracy using the same sensor. At 60GHz, the radar signal can only penetrate the skin at about 1mm [36]. We also show that it is possible to differentiate users based on how they touch a numeric keypad on top of the sensor. This enables on-the-fly authentication when users are requested to press a few digits, by judging how they touch. Such biometric measures can be useful additional input when combined with other means of biometric authentication.

Food and drinks: We have conducted a preliminary study in our lab to differentiate i) liquids with different dye and contamination and ii) powder like foods under highly controlled environments and the results are promising [35]. Others [52] have also used the Soli sensor for differentiating types of drinks (coke, diet coke, and zero coke). In the future, we envision that the nutrition level inside the foods or drinks can be also measured [51]. Recognizing ingredients along with the nutrition and freshness can enable a truly smart kitchen.

Wearables: We can also imagine this kind of sensing being woven into material and fabric that people are wearing. When embedded in a shoe, for example, it is possible to know ones current location in a building (location awareness) based on the tiles or carpet being stepped on. It can perhaps replace the tactile paving on the road that is used for assisting the blind, with material that is embedded beneath the ground.

9 LIMITATIONS AND FUTURE WORK

There remain several practical limitations in our work. Credit to the high sensitivity of the radar sensor, we can detect subtle differences between objects and layers. However, this high sensitivity also presents challenges. For an object to be recognized in a stable manner, it has to be placed directly on top or at a stable distance from the sensor, which limits the field of view of the sensor. Other dimensions such as the air gap, bending, position, height, all contribute to the overall recognition accuracy. As shown in the confusion matrix in the counting paper experiment, there is some confusion with increasing layers, but the confusion is often between neighbours. This is because the sheets of paper are very thin and flexible. In contrast, counting the full set of playing cards (up to 52) was very accurate. In the flipping playing card experiment, we realized that when we intentionally bend the card, the classified class is temporarily inverted. However, it returns to the correct class after a short while, due to the flexibility of the card (as long the the card is not bent to an extent where it cannot return to its original shape).

While we have demonstrated that we can recognize and interact with objects from a distance, it is currently most robust with objects placed directly on top of the sensor. Similarly, the signal propagation level has limits, which decreases with increasing layers, as seen in the chips counting experiment (Figure 8), where the accuracy drops with increasing number of chips. As a result, we recommend limited stacking layers until future version of the radar hardware with stronger signal penetration strength becomes available.

Following the Token+Constraint paradigm, we used the 3D printed case which allows us to extend the functionality and sensing accuracy. As the sensor we used (Soli) employs a broad beam of 150 degrees, a small object (e.g., a Lego piece) cannot cover the entire field. At such a short distance (6mm from antenna to object), even a tiny movement alters the raw signal. Therefore, the case serves as a simple constraint mechanism [56] to fix the object in place, reducing the degrees of freedom of the data required to compensate for all tiny changes in location and orientation. While using a case might reduce the practicality of the system, it is not uncommon for such sensing systems to have a case or holder [25, 49]. For example, interacting with toys-to-life products that use NFC-tagged toy figures is limited to placing and removing them on a NFC reader “portal”. Other real-world examples such as a credit card reader often requires the card to be fully slotted in, or swiped through a magnetic reader; vending machine requires coins to be push through a small peephole while a money counting machine requires bank notes to be placed exactly within the mechanical guide, etc.

In future work, we can consider an array of such sensors which may overcome the small sensing area each has. Recent work, such as DeskWave [33] demonstrates multiple microwave sensors beneath a table for gestures and

tangible object recognition. We can also consider a sensor fusion approach combining radar with other low-cost sensors such as RFID/NFC or capacitive touch array as demonstrated in Project Zanzibar [57].

We have explored both pre-existing “everyday” objects and purpose-fabricated tokens with several designs. Future work aims to investigate more designs and characterize how such structures (hole, shape, thickness) affect the radar response, perhaps structures with heightened legibility both to humans and the radar sensor for creating tokens that are more visually and aesthetically pleasing.

As a proof-of-concept prototype, the current training process requires considerable human effort to collect the data. However, for a commercial product, it can be factory trained and calibrated before shipping to end-users because the sensing modes (counting, ordering, etc) do not vary with users. For touch location sensing (numpad), we also see that it can be generalized to other users when only trained with a group of users.

Nonetheless, it is also possible to reduce the training effort by using interpolation or regression techniques. As we have shown, regression works for a slider and rotating dial when data is collected at small to medium intervals. One exception where regression fails even when trained with data collected at small intervals in the desk inside-out tracking example. The desk is made of compressed wood and as such there is no linearity of the signal change between different locations.

Other means of reducing training effort can be seen as automation. For example, in the hourglass example, the data collection was synced with a computer timer, thus this is only a one click process. In the rotation example, by pairing the system with a device which contains an inertial measurement unit (IMU) such as the phone (Figure 6 (f)), it can provide ground truth orientation label for real-time training. Furthermore, the training can be automated with the help of a mechanical device such as using a robotic hand or paper dispenser. Our future work aims to extend the sensing capabilities such as environmental sensing and adopting a deep learning approach to scale it to more real-world objects.

10 CONCLUSION

The essence of Solinteraction, is to allow users to use everyday objects as tangible tokens. We have shown that many day-to-day objects can be recognized, with or without tagging. Further, as radar signals penetrate through many type of materials, we can achieve counting and ordering with high accuracy and for many layers, which has received little attention in prior work. In this paper we have explored several facets of radar-based tangible interaction, having carefully situated it with respect to related work. Reflecting on our design space and application scenarios, we show that radar sensing can be useful as a new means to interact with computation using everyday objects, not limited to static interaction only, but also dynamic interaction. We suggest the insights presented are useful for the community to build on and learn from, as we believe there remain considerable unexplored potential with radar-based sensing in Human-Computer Interaction (HCI).

ACKNOWLEDGMENTS

We thank the anonymous reviewers and associate editor for their insightful feedback which has helped us to improve this paper greatly. In addition, we thank the Google ATAP Soli team for providing us access to the alpha developer kit and SDK. We also gratefully acknowledge the support of the NVIDIA Corporation with the donation of a Titan X Pascal GPU used in this research. We also thank Avik Santra from Infineon for his help in revising the draft and Patrick Schrempf for his help in the development of the card game application.

REFERENCES

- [1] Fadel Adib, Chen-Yu Hsu, Hongzi Mao, Dina Katabi, and Frédo Durand. 2015. Capturing the Human Figure Through a Wall. *ACM Trans. Graph.* 34, 6, Article 219 (Oct. 2015), 13 pages. <https://doi.org/10.1145/2816795.2818072>
- [2] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C. Miller. 2015. Smart Homes That Monitor Breathing and Heart Rate. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA,

- 837–846. <https://doi.org/10.1145/2702123.2702200>
- [3] Ansys. 2017. Engineering Simulation & 3D Design Software | ANSYS. Website. Retrieved September 14, 2017 from <http://www.ansys.com/>.
- [4] Google ATAP. 2018. Project Soli: Alpha Developers Showcase. Website. Retrieved October 10, 2018 from https://www.youtube.com/watch?v=H41A_IWZwZl.
- [5] Daniel Avrahami, Mitesh Patel, Yusuke Yamaura, and Sven Kratz. 2018. Below the Surface: Unobtrusive Activity Recognition for Work Surfaces Using RF-radar Sensing. In *23rd International Conference on Intelligent User Interfaces (IUI '18)*. ACM, New York, NY, USA, 439–451. <https://doi.org/10.1145/3172944.3172962>
- [6] Daniel Avrahami, Jacob O. Wobbrock, and Shahram Izadi. 2011. Portico: Tangible Interaction on and Around a Tablet. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 347–356. <https://doi.org/10.1145/2047196.2047241>
- [7] M. E. Baginski, D. L. Faircloth, and M. D. Deshpande. 2005. Comparison of two optimization techniques for the estimation of complex permittivities of multilayered structures using waveguide measurements. *IEEE Transactions on Microwave Theory and Techniques* 53, 10 (Oct 2005), 3251–3259. <https://doi.org/10.1109/TMTT.2005.855133>
- [8] Tom Bartindale and Chris Harrison. 2009. Stacks on the Surface: Resolving Physical Order Using Fiducial Markers with Structured Transparency. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '09)*. ACM, New York, NY, USA, 57–60. <https://doi.org/10.1145/1731903.1731916>
- [9] Patrick Baudisch, Torsten Becker, and Frederik Rudeck. 2010. Lumino: Tangible Blocks for Tabletop Computers Based on Glass Fiber Bundles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1165–1174. <https://doi.org/10.1145/1753326.1753500>
- [10] Francisco Bernardo, Nicholas Arner, and Paul Batchelor. 2017. O soli mio: exploring millimeter wave radar for musical interaction.. In *NIME*, Vol. 17. 283–286.
- [11] Andrea Bianchi and Ian Oakley. 2013. Designing Tangible Magnetic Appcessories. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (TEI '13)*. ACM, New York, NY, USA, 255–258. <https://doi.org/10.1145/2460625.2460667>
- [12] Andrea Bianchi and Ian Oakley. 2015. MagnID: Tracking Multiple Magnetic Tokens. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '15)*. ACM, New York, NY, USA, 61–68. <https://doi.org/10.1145/2677199.2680582>
- [13] Xethru by Novelda. 2017. Single-chip radar sensors with sub-mm resolution. Website. Retrieved September 14, 2017 from <https://www.xethru.com/>.
- [14] Liwei Chan, Stefanie Müller, Anne Roudaut, and Patrick Baudisch. 2012. CapStones and ZebraWidgets: Sensing Stacks of Building Blocks, Dials and Sliders on Capacitive Touch Screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2189–2192. <https://doi.org/10.1145/2207676.2208371>
- [15] Kai-Yin Cheng, Rong-Hao Liang, Bing-Yu Chen, Rung-Huei Laing, and Sy-Yen Kuo. 2010. iCon: Utilizing Everyday Objects As Additional, Auxiliary and Instant Tabletop Controllers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1155–1164. <https://doi.org/10.1145/1753326.1753499>
- [16] Chin-yu Chien, Rong-Hao Liang, Long-Fei Lin, Liwei Chan, and Bing-Yu Chen. 2015. FlexiBend: Enabling Interactivity of Multi-Part, Deformable Fabrications Using Single Shape-Sensing Strip. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 659–663. <https://doi.org/10.1145/2807442.2807456>
- [17] K. Diederichs, A. Qiu, and G. Shaker. 2017. Wireless Biometric Individual Identification Utilizing Millimeter Waves. *IEEE Sensors Letters* 1, 1 (Feb 2017), 1–4. <https://doi.org/10.1109/LESENS.2017.2673551>
- [18] Paul H. Dietz and Benjamin D. Eidelson. 2009. SurfaceWare: Dynamic Tagging for Microsoft Surface. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction (TEI '09)*. ACM, New York, NY, USA, 249–254. <https://doi.org/10.1145/1517664.1517717>
- [19] Barrett Ens, Aaron Quigley, Hui-Shyong Yeo, Pourang Irani, Thammathip Piumsomboon, and Mark Billinghurst. 2017. Exploring Mixed-scale Gesture Interaction. In *SIGGRAPH Asia 2017 Posters (SA '17)*. ACM, New York, NY, USA, Article 27, 2 pages. <https://doi.org/10.1145/3145690.3145740>
- [20] Google. 2016. Project Soli. Website. Retrieved September 16, 2016 from <https://atap.google.com/soli/>.
- [21] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA Data Mining Software: An Update. *SIGKDD Explor. Newsl.* 11, 1 (Nov. 2009), 10–18. <https://doi.org/10.1145/1656274.1656278>
- [22] Meng-Ju Hsieh, Rong-Hao Liang, Da-Yuan Huang, Jheng-You Ke, and Bing-Yu Chen. 2018. RFIBricks: Interactive Building Blocks Based on RFID. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 189, 10 pages. <https://doi.org/10.1145/3173574.3173763>
- [23] Sungjae Hwang, Myungwook Ahn, and Kwang-yun Wohn. 2013. MagGetz: Customizable Passive Tangible Controllers on and Around Conventional Mobile Devices. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, NY, USA, 411–416. <https://doi.org/10.1145/2501988.2501991>
- [24] Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. 2007. The reacTable: Exploring the Synergy Between Live Music Performance and Tabletop Tangible Interfaces. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction (TEI '07)*. ACM, New York, NY, USA, 139–146. <https://doi.org/10.1145/1226969.1226998>

- [25] Gierad Laput, Eric Brockmeyer, Scott E. Hudson, and Chris Harrison. 2015. Acoustruments: Passive, Acoustically-Driven, Interactive Controls for Handheld Devices. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2161–2170. <https://doi.org/10.1145/2702123.2702414>
- [26] Mathieu Le Goc, Pierre Dragicevic, Samuel Huron, Jeremy Boy, and Jean-Daniel Fekete. 2015. SmartTokens: Embedding Motion and Grip Sensing in Small Tangible Objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology (UIST '15)*. ACM, New York, NY, USA, 357–362. <https://doi.org/10.1145/2807442.2807488>
- [27] Jinha Lee, Yasuaki Kakehi, and Takeshi Naemura. 2009. Bloxels: Glowing Blocks As Volumetric Pixels. In *ACM SIGGRAPH 2009 Emerging Technologies (SIGGRAPH '09)*. ACM, New York, NY, USA, Article 5, 1 pages. <https://doi.org/10.1145/1597956.1597961>
- [28] Hanchuan Li, Eric Brockmeyer, Elizabeth J. Carter, Josh Fromm, Scott E. Hudson, Shwetak N. Patel, and Alanson Sample. 2016. PaperID: A Technique for Drawing Functional Battery-Free Wireless Interfaces on Paper. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5885–5896. <https://doi.org/10.1145/2858036.2858249>
- [29] Hanchuan Li, Can Ye, and Alanson P. Sample. 2015. IDSense: A Human Object Interaction Detection System Based on Passive UHF RFID. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2555–2564. <https://doi.org/10.1145/2702123.2702178>
- [30] Rong-Hao Liang, Kai-Yin Cheng, Liwei Chan, Chuan-Xhyuan Peng, Mike Y. Chen, Rung-Huei Liang, De-Nian Yang, and Bing-Yu Chen. 2013. GaussBits: Magnetic Tangible Bits for Portable and Occlusion-free Near-surface Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1391–1400. <https://doi.org/10.1145/2470654.2466185>
- [31] Rong-Hao Liang, Han-Chih Kuo, and Bing-Yu Chen. 2016. GaussRFID: Reinventing Physical Toys Using Magnetic RFID Development Kits. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4233–4237. <https://doi.org/10.1145/2858036.2858527>
- [32] Jaime Lien, Nicholas Gillian, M. Emre Karagozler, Patrick Amihoud, Carsten Schwesig, Erik Olson, Hakim Raja, and Ivan Poupyrev. 2016. Soli: Ubiquitous Gesture Sensing with Millimeter Wave Radar. *ACM Trans. Graph.* 35, 4, Article 142 (July 2016), 19 pages. <https://doi.org/10.1145/2897824.2925953>
- [33] Jess McIntosh, Mike Fraser, Paul Worgan, and Asier Marzo. 2017. DeskWave: Desktop Interactions Using Low-cost Microwave Doppler Arrays. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, New York, NY, USA, 1885–1892. <https://doi.org/10.1145/3027063.3053152>
- [34] Rafael Morales González, Caroline Appert, Gilles Bailly, and Emmanuel Pietriga. 2016. TouchTokens: Guiding Touch Patterns with Passive Tokens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4189–4202. <https://doi.org/10.1145/2858036.2858041>
- [35] David Morrison, Alasdair Houston, N. Aaron Quigley, and David Harris-Birtill. 2017. AquaCat: Radar and Machine Learning for Fluid and Powder Identification. In *Proceedings of the MobileHCI'17 Workshop on Object Recognition for Input and Mobile Interaction (MobileHCI '17)*. 6. https://sachi.cs.st-andrews.ac.uk/wp-content/uploads/2017/09/David_Aquacat.pdf
- [36] D. A. Nelson, M. T. Nelson, T. J. Walters, and P. A. Mason. 2000. Skin heating effects of millimeter-wave irradiation-thermal modeling results. *IEEE Transactions on Microwave Theory and Techniques* 48, 11 (Nov 2000), 2111–2120. <https://doi.org/10.1109/22.884202>
- [37] A. E. Omer, G. Shaker, S. Safavi-Naeini, K. Murray, and R. Hughson. 2018. Glucose Levels Detection Using mm-Wave Radar. *IEEE Sensors Letters* 2, 3 (Sept 2018), 1–4. <https://doi.org/10.1109/LESENS.2018.2865165>
- [38] Makoto Ono, Buntarou Shizuki, and Jiro Tanaka. 2013. Touch & Activate: Adding Interactivity to Existing Objects Using Active Acoustic Sensing. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, NY, USA, 31–40. <https://doi.org/10.1145/2501988.2501989>
- [39] Osmo. 2017. Osmo - Award-Winning Educational Games System for iPad. Website. Retrieved March 30, 2017 from <https://www.playosmo.com>.
- [40] Joseph Paradiso, Craig Abler, Kai-yuh Hsiao, and Matthew Reynolds. 1997. The Magic Carpet: Physical Sensing for Immersive Environments. In *CHI '97 Extended Abstracts on Human Factors in Computing Systems (CHI EA '97)*. ACM, New York, NY, USA, 277–278. <https://doi.org/10.1145/1120212.1120391>
- [41] Joseph A Paradiso. 1999. The brain opera technology: New instruments and gestural sensors for musical interaction and performance. *Journal of New Music Research* 28, 2 (1999), 130–149.
- [42] James Patten, Hiroshi Ishii, Jim Hines, and Gian Pangaro. 2001. Sensetable: A Wireless Object Tracking Platform for Tangible User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. ACM, New York, NY, USA, 253–260. <https://doi.org/10.1145/365024.365112>
- [43] James Patten, Ben Recht, and Hiroshi Ishii. 2002. Audiopad: A Tag-based Interface for Musical Performance. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME '02)*. National University of Singapore, Singapore, Singapore, 1–6. <http://dl.acm.org/citation.cfm?id=1085171.1085175>
- [44] Tauhidur Rahman, Alexander T. Adams, Ruth Vinisha Ravichandran, Mi Zhang, Shwetak N. Patel, Julie A. Kientz, and Tanzeem Choudhury. 2015. Dopplesleep: A Contactless Unobtrusive Sleep Sensing System Using Short-range Doppler Radar. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 39–50.

- <https://doi.org/10.1145/2750858.2804280>
- [45] Jun Rekimoto. 2002. SmartSkin: An Infrastructure for Freehand Manipulation on Interactive Surfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. ACM, New York, NY, USA, 113–120. <https://doi.org/10.1145/503376.503397>
- [46] Jun Rekimoto and Masanori Saitoh. 1999. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 378–385. <https://doi.org/10.1145/302979.303113>
- [47] Jun Rekimoto, Brygg Ullmer, and Haruo Oba. 2001. DataTiles: A Modular Platform for Mixed Physical and Graphical Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. ACM, New York, NY, USA, 269–276. <https://doi.org/10.1145/365024.365115>
- [48] Christian Sandor and Hiraku Nakamura. 2018. SoliScratch: A Radar Interface for Scratch DJs. In *Adjunct Proceedings of the IEEE International Symposium for Mixed and Augmented Reality 2018 (To appear)*.
- [49] Valkyrie Savage, Colin Chang, and Björn Hartmann. 2013. Sauron: Embedded Single-camera Sensing of Printed Physical User Interfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, NY, USA, 447–456. <https://doi.org/10.1145/2501988.2501992>
- [50] Martin Schmitz, Mohammadreza Khalilbeigi, Matthias Balwierz, Roman Lissermann, Max Mühlhäuser, and Jürgen Steimle. 2015. Capricate: A Fabrication Pipeline to Design and 3D Print Capacitive Touch Sensors for Interactive Objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology (UIST '15)*. ACM, New York, NY, USA, 253–258. <https://doi.org/10.1145/2807442.2807503>
- [51] Scio. 2017. ScIO - The World's First Pocket Sized Molecular Sensor. Website. Retrieved March 30, 2017 from <https://www.consumerphysics.com/>.
- [52] George Shaker, Karly Smith, Ala Eldin Omer, Shuo Liu, Clement Csech, Udeshtaya Wadhwa, Safieddin Safavi-Naeini, and Richard Hughson. 2018. Non-Invasive Monitoring of Glucose Level Changes Utilizing a mm-Wave Radar System. *Int. J. Mob. Hum. Comput. Interact.* 10, 3 (July 2018), 10–29. <https://doi.org/10.4018/IJMHCI.2018070102>
- [53] Andrew Spielberg, Alanson Sample, Scott E. Hudson, Jennifer Mankoff, and James McCann. 2016. RapID: A Framework for Fabricating Low-Latency Interactive Objects with RFID Tags. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5897–5908. <https://doi.org/10.1145/2858036.2858243>
- [54] Brygg Ullmer and Hiroshi Ishii. 1997. The metaDESK: Models and Prototypes for Tangible User Interfaces. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology (UIST '97)*. ACM, New York, NY, USA, 223–232. <https://doi.org/10.1145/263407.263551>
- [55] Brygg Ullmer, Hiroshi Ishii, and Dylan Glas. 1998. mediaBlocks: Physical Containers, Transports, and Controls for Online Media. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. ACM, New York, NY, USA, 379–386. <https://doi.org/10.1145/280814.280940>
- [56] Brygg Ullmer, Hiroshi Ishii, and Robert J. K. Jacob. 2005. Token+Constraint Systems for Tangible Interaction with Digital Information. *ACM Trans. Comput.-Hum. Interact.* 12, 1 (March 2005), 81–118. <https://doi.org/10.1145/1057237.1057242>
- [57] Nicolas Villar, Daniel Cletheroe, Greg Saul, Christian Holz, Tim Regan, Oscar Salandin, Misha Sra, Hui-Shyong Yeo, William Field, and Haiyan Zhang. 2018. Project Zanzibar: A Portable and Flexible Tangible Interaction Platform. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 515, 13 pages. <https://doi.org/10.1145/3173574.3174089>
- [58] Simon Voelker, Christian Cherek, Jan Thar, Thorsten Karrer, Christian Thoresen, Kjell Ivar Øvergård, and Jan Borchers. 2015. PERCs: Persistently Trackable Tangibles on Capacitive Multi-Touch Displays. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 351–356. <https://doi.org/10.1145/2807442.2807466>
- [59] Simon Voelker, Kosuke Nakajima, Christian Thoresen, Yuichi Itoh, Kjell Ivar Øvergård, and Jan Borchers. 2013. PUCs: Detecting Transparent, Passive Untouched Capacitive Widgets on Unmodified Multi-touch Displays. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces (ITS '13)*. ACM, New York, NY, USA, 101–104. <https://doi.org/10.1145/2512349.2512791>
- [60] Walabot. 2016. Walabot Sensor | Radio Frequency technology for Advanced Detection. Website. Retrieved September 16, 2016 from <https://walabot.com/>.
- [61] Saiwen Wang, Jie Song, Jaime Lien, Ivan Poupyrev, and Otmar Hilliges. 2016. Interacting with Soli: Exploring Fine-Grained Dynamic Gesture Recognition in the Radio-Frequency Spectrum. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 851–860. <https://doi.org/10.1145/2984511.2984565>
- [62] Roy Want, Kenneth P. Fishkin, Anuj Gujar, and Beverly L. Harrison. 1999. Bridging Physical and Virtual Worlds with Electronic Tags. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 370–377. <https://doi.org/10.1145/302979.303111>
- [63] Wikipedia. 2018. Toys-to-life. Website. Retrieved August 1, 2018 from <https://en.wikipedia.org/wiki/Toys-to-life>.
- [64] Karl Willis, Eric Brockmeyer, Scott Hudson, and Ivan Poupyrev. 2012. Printed Optics: 3D Printing of Embedded Optical Elements for Interactive Devices. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*. ACM, New York, NY, USA, 589–598. <https://doi.org/10.1145/2380116.2380190>

- [65] Karl D. D. Willis and Andrew D. Wilson. 2013. InfraStructs: Fabricating Information Inside Physical Objects for Imaging in the Terahertz Region. *ACM Trans. Graph.* 32, 4, Article 138 (July 2013), 10 pages. <https://doi.org/10.1145/2461912.2461936>
- [66] Andrew D. Wilson. 2005. PlayAnywhere: A Compact Interactive Tabletop Projection-vision System. In *Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology (UIST '05)*. ACM, New York, NY, USA, 83–92. <https://doi.org/10.1145/1095034.1095047>
- [67] Hui-Shyong Yeo, Barrett Ens, and Aaron Quigley. 2017. Tangible UI by Object and Material Classification with Radar. In *SIGGRAPH Asia 2017 Emerging Technologies (SA '17)*. ACM, New York, NY, USA, Article 14, 2 pages. <https://doi.org/10.1145/3132818.3132824>
- [68] Hui-Shyong Yeo, Gergely Flamich, Patrick Schrempf, David Harris-Birtill, and Aaron Quigley. 2016. RadarCat: Radar Categorization for Input & Interaction. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 833–841. <https://doi.org/10.1145/2984511.2984515>
- [69] Hui-Shyong Yeo and Aaron Quigley. 2017. Radar Sensing in Human-computer Interaction. *Interactions* 25, 1 (Dec. 2017), 70–73. <https://doi.org/10.1145/3159651>
- [70] Neng-Hao Yu, Li-Wei Chan, Seng Yong Lau, Sung-Sheng Tsai, I-Chun Hsiao, Dian-Je Tsai, Fang-I Hsiao, Lung-Pan Cheng, Mike Chen, Polly Huang, and Yi-Ping Hung. 2011. TUIC: Enabling Tangible Interaction on Capacitive Multi-touch Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2995–3004. <https://doi.org/10.1145/1978942.1979386>
- [71] Mingmin Zhao, Fadel Adib, and Dina Katabi. 2016. Emotion Recognition Using Wireless Signals. In *Proceedings of the 22Nd Annual International Conference on Mobile Computing and Networking (MobiCom '16)*. ACM, New York, NY, USA, 95–108. <https://doi.org/10.1145/2973750.2973762>
- [72] Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, and Dina Katabi. 2018. Through-Wall Human Pose Estimation Using Radio Signals. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Received February 2018; revised August 2018; accepted October 2018