

RadarCat: Radar Categorization for Input & Interaction

Hui-Shyong Yeo, Gergely Flamich, Patrick Schrempf, David Harris-Birtill and Aaron Quigley

SACHI, School of Computer Science
University of St Andrews, Scotland, UK
{hsy, gf38, ps90, dcchb, aquigley}@st-andrews.ac.uk



Figure 1. (left) The 26 materials/objects (including air and water) that we evaluate in the first study. (middle) 16 Transparent materials used in the second study, arranged from top to bottom according to the list in Figure 6. (right) 10 different body parts that are being classified in the third study.

ABSTRACT

In RadarCat we present a small, versatile radar-based system for material and object classification which enables new forms of everyday proximate interaction with digital devices. We demonstrate that we can train and classify different types of materials and objects which we can then recognize in real time. Based on established research designs, we report on the results of three studies, first with 26 materials (including complex composite objects), next with 16 transparent materials (with different thickness and varying dyes) and finally 10 body parts from 6 participants. Both leave one-out and 10-fold cross-validation demonstrate that our approach of classification of radar signals using random forest classifier is robust and accurate. We further demonstrate four working examples including a physical object dictionary, painting and photo editing application, body shortcuts and automatic refill based on RadarCat. We conclude with a discussion of our results, limitations and outline future directions.

Author Keywords

Radar; Context-Aware Interaction; Machine Learning; Material Classification; Object Recognition; Ubiquitous Computing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

UIST 2016, October 16-19, 2016, Tokyo, Japan.
Copyright © 2016 ACM ISBN 978-1-4503-4189-9/16/10...\$15.00
DOI: <http://dx.doi.org/10.1145/2984511.2984515>

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation (e.g. HCI): User interfaces - Input devices and strategies;

INTRODUCTION

Today we know more about our computing devices than they know about us, their environments, and their use. Existing visions of computing [24] assume knowledge of the world to realize their aims. For example, Weiser's vision of ubiquitous computing [34] relies on sensing, distributed through the fabric of life to help enable context-aware interaction [8]. Tangible user interfaces [11], rely on physical objects which often need to understand their collective configurations while Instrumental Interaction [4] offers an interaction model for post-WIMP interfaces. Richer sensing and understanding of the real world allows new forms and styles of interaction, and hence entirely new classes of user interface to emerge.

In this paper, we explore the potential of enabling computing devices to recognize proximate materials or objects they are touching with RadarCat. Our novel sensing approach exploits the multi-channel radar signals, emitted from a Project Soli [19, 30] sensor, that are highly characteristic when reflected from everyday objects; as different materials, thickness and geometry of the object will scatter, refract and reflect the radar signals differently. We employ machine learning and classification techniques on these signals, demonstrate that we can reliably extract rich information about the target material or object, and leverage this to enable novel interaction capabilities. Beyond HCI, RadarCat also opens up new opportunities in areas such as navigation and world knowl-

edge (e.g., low vision users), consumer interaction (e.g., scales), or industrial automation (e.g., recycling).

Although radar technology [29] has been used for decades in aircraft tracking, security scanners and non-destructive testing and evaluation [15], we are not aware of any previous attempt to explore this technology for enabling novel proximate interactions in the field of human-computer interaction and ubiquitous computing. As such, our contributions are:

1. Exploration of radar sensing to capture details of a proximate target and introduction of a technique that shows the potential of re-using tiny radar to:
 - (a) Classify everyday objects and recognize their orientation, or if a liquid is added.
 - (b) Classify and differentiate transparent yet visually similar materials.
 - (c) Classify different human body parts.
2. Series of studies showing that our sensing approach is accurate, which demonstrates the potential of RadarCat in a variety of real-world applications.
3. Identification of practical use-cases for real-world applications, and implementation of four example context-aware applications enabled by RadarCat.

As the use of sensing for material and object classification is not new, we first describe the related work. We draw on the experimental design in previous work to design our experiments. By replicating existing experimental approaches our results can be compared against previous and future studies.

Radar Primer

Radar uses an emission of electromagnetic radio waves, generally with a frequency within 1GHz-300GHz [1], which is then reflected back from an object and received by a detector. The time of flight can be used to calculate the distance to an object, and using the Doppler shift the velocity of the object can also be measured. The properties which effect the received radar intensity are the absorption and scattering properties of the material at the wavelengths used, and hence the reflection and transmission properties of the material, the material's thickness and shape, the refractive index and hence the specular reflection from the material, as well as the distance to the object from the emitter/receiver. The received signal has contributions from the reflection from the bottom surface, the scattering from the internal structure, and the reflection from the rear surface of the material. There are several physical properties of the material, such as the density, which effect these absorption and scattering coefficients, a review of which can be found in [12]. Objects may be composed of single (e.g., copper sheet) or composite materials (e.g., a mobile phone composed of a combination of materials such as glass, aluminium, plastics etc.) and have different received radar signals due to their physical properties.

RELATED WORK

Our research, and hence user studies, draw on three different bodies of research, including object recognition, material classification and approaches in context aware computing.

Object Recognition

Object recognition can be achieved by sampling the object in both destructive and non-destructive ways, or disruptive and non-disruptive ways. Destructive methods can involve taking a physical sample of the object and subjecting it to chemical analysis with different types of chromatography or spectroscopy. This relies on a single material or sampling sufficient aspects of composite materials to recognize the overall object. Methods which involve localised destruction (e.g., etching) are also possible [9] to realize acoustic barcodes.

Non-disruptive and non-destructive approaches rely on sensing the object from its real world use. Computer vision techniques, which require the object to be visible, well lit and within range for a suitable resolution have been explored in a range of object tracking approaches [37] (e.g., tracking rectangular objects [32] such as tablets), while depth sensing (Kinect) with infrared can overcome the issue of lighting.

Radar systems, have been used to recognize particular types of aircraft or materials in luggage or body scanners [26]. Ground-penetrating radar (GPR) can be used to detect buried objects [2, 14] such as utility pipes or bones. Object tracking systems which rely on measuring WiFi [3] or Bluetooth signals can also be employed to recognize objects.

Disruptive yet largely non-destructive approaches employ the addition of elements to the object which can be sensed (e.g., RFID [5], visual markers [25], QR codes) or by allowing the objects to emit visual or audio signals (e.g., ultrasonic [10]).

In practice, both destructive and disruptive object recognition approaches can have a significant impact and hence real-world disruption on the physical infrastructure, environment, computational system or services offered [15].

Material Classification

The destructive methods of object identification can also be employed for material classification. Further non-destructive methods such as near-infrared (NIR) spectroscopy is often utilized for analyzing pharmaceutical products. Likewise, millimeter wave and terahertz technology are being used to detect materials from a distance [13] for scientific exploration (e.g., planet hunting) or security purposes. By re-purposing an off-the-shelf radio chipset Zhu et al. [38] have used radar to recognize materials from a distance based on a database of material/radar signal loss from different distance and incident angle [16]. InfraStructs [36] suggest future applications for interaction, using terahertz imaging. While research into the object detection and material recognition has been undertaken on buried objects with GPR images [6, 23]. Regardless of the approach, these sensing methods are complex and costly, let alone the size and power requirements.

As with recognition, material classification can employ non-destructive and less disruptive vision-based approaches [28], although this can be challenging. However, in a controlled setup with sufficient light or self-illumination and at a close proximity, the problem is more tractable. Harrison and Hudson [8], employ a single photoresistor with multispectral illumination to identify the surface material property. Similar image-based surface classification techniques exist, such

as using a laser optical mouse sensor [20] for classification. More recent work in SpecTrans [27], is able to classify transparent materials in addition to surface material of everyday objects but does not report on different object states (e.g., filled/non-filled cup) or use with different body parts.

Vision based material classification suffers from being limited to the material qualities which are present on the surface of the object. This can result in confusion where a layer of opaque material (packaging) blocks the primary object of interest for classification. RadarCat, by contrast, provides a degree of surface level penetration as long as the outer layer is not highly reflective to the radar signal. This allows us to explore materials and classify object without being limited to just what is visible on the surface.

Context Recognition

The placement of a device on the body or within the environment can be seen as an aspect of context recognition. As such, existing wearable approaches suggest that material recognition to enable placement detection is valuable to both location and activity recognition [8]. Phoneprioception [35] further strengthens this, by suggesting that simple sensors such as an accelerometer and light sensor can be combined to achieve high accuracy in determining a phone placement location, both on body and within ones personal space. Using EMG sensor, Potential [21] is able to detect different placements on the human body. On the other hand, Lien et al. [19] and Song et al. [31] introduce a new approach to sensing finger gestures, with an end-to-end radar system (Soli) and classify the gestures using machine learning techniques such as random forest and deep neural network. Similarly, it is also possible to infer the 3D finger and hand position on top of a transparent electric field sensor [18], based on random forest regression.

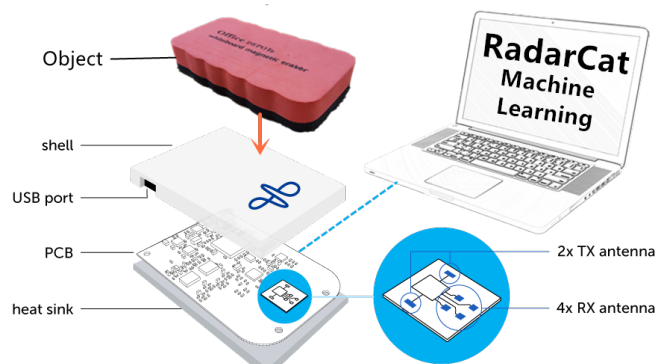


Figure 2. Soli alpha hardware exploded view (not to scale, image adapted from Soli alpha SDK). Object is placed on top of the sensor where raw radar signals are classified using machine learning (ML).

DESIGN OF RADARCAT: RADAR CATEGORIZATION

The development of RadarCat was based on iterative technical design decisions. We leveraged improved knowledge and analysis of radar signal signatures, machine learning and classification results in the development of our final approach. The signatures, unique to each object, are measured when the object is proximate to the sensor, and allows us to classify the object using a machine learning technique. Our goal is

to overcome the limitations of camera-based systems, with an embedded sensor that can detect surface material at high speed and accuracy. When the object is placed on or touched by the sensor, the near and fixed distance together with the fixed incident angle makes our classification task straightforward and allowing for accurate classification. The final design of RadarCat incorporates the following unique capabilities and aspects (a) non-destructive, non-tagging, no illumination (b) use with surface materials, composite objects and certain body parts and (c) identification of new sensing applications and practical use cases and interaction that are brought to bear by a portable radar technology.

RadarCat Hardware

Our system uses Soli [30] (Figure 1 & 2), a prototype radar device by Google ATAP, designed for capturing subtle finger motion for enabling interaction with computers. For detailed information, we refer the reader to the Soli paper [19]. Generally speaking, Soli is a monostatic radar device and contains multiple antennas (2 transmitters and 4 receivers), providing simultaneous operation of up to 8-channels, using frequency-modulated continuous wave (FMCW) operating in the 57-64 GHz range (center frequency of 60 GHz). The distance from the sensor top to plastic enclosure top is 6mm (Figure 2). In RadarCat, the object to be tested is placed on top of the enclosure, or touched by it as if using a stethoscope.

Currently, the Soli developer kit is only available to selected alpha developers. While Google's ambition is to have Soli embedded in mobile devices in the near future, we do not claim that our technique works with off-the-shelf hardware just yet, nor do we claim that the hardware is our contribution. Nonetheless, we suggest that our technique should work with other small radar systems such as Walabot [33], that will be available later in 2016. As noted previously, Zhu et al. [38] also explored the reuse of a off-the-shelf radio chipset working as synthetic aperture radar (SAR) for the recognition of four different materials from a distance.

Implementation

We implemented our system in two parts i) a graphical user interface (GUI) in C++ using Qt and ii) a classifier backend in Java using Weka [7] API. Communication between the GUI and classifier are with sockets, thus, the classifier can run on the same machine, or can be offloaded to a more powerful server via the network, suitable for scenarios involving wearable devices with limited computing resources, as we show.

Feature extraction

The received radar intensity is influenced by the reflection and transmission properties of the material. Reflected signals from many points both within and on the object surface are overlapping and hence contribute to the received signal.

As the radar signals are stable and highly discriminative (as shown in Figure 3 and Figure 4), we currently use all 8 channels as input features where each channel consists of 64 data points, yielding 512 features. We further extract the average (avg) and the absolute value (abs) along the signals from all 8 channels, yielding an additional 128 features. In addition, we extract common features such as absolute and root means

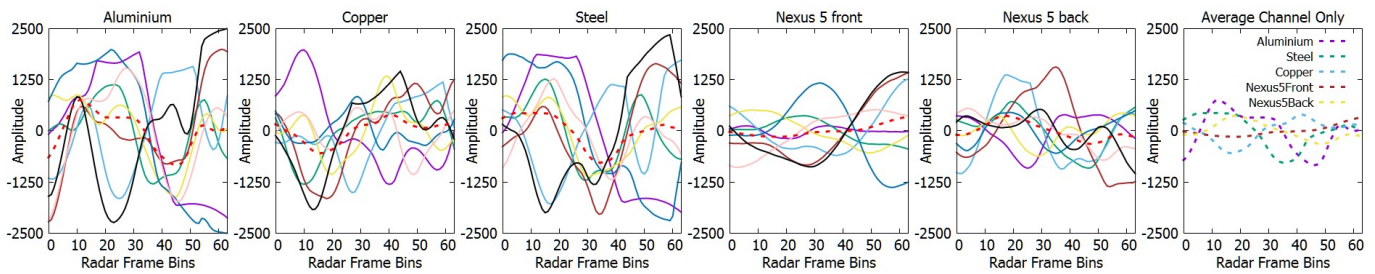


Figure 3. 8-channels raw radar signals for different materials, red dotted line is the average of 8 channels, from left to right i) aluminium ii) steel iii) copper iv) phone (Nexus5 front) v) phone (Nexus5 back) vi) combined view using the average channel. X-axis represents samples number, where each sample is 555 nanosecond long, as the sample rate is 1.8 mega samples per second.

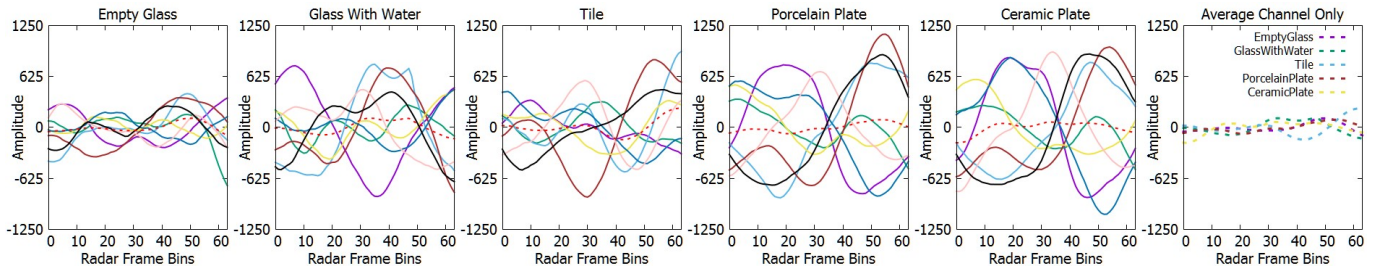


Figure 4. 8-channels raw radar signals for similar materials, red dotted line is the average of 8 channels, from left to right i) empty glass ii) glass filled with water iii) tile iv) porcelain plate v) ceramic plate vi) combined view using the average channel. As we can see, the radar signals are unique.

square (rms) for each channel (x8); global maxima, global minima, avg, abs and rms for all channels (x1), yielding extra 21 features, resulting in a total of 661 features. Through feature selection analysis, we found the derived features are highly ranked. Nonetheless, the remaining features are also important to fully capture the subtle signal behaviours and are important for training new objects. We experimented with different machine learning classifiers and ended up with two candidates: SVM and random forest. We finally selected random forest due to its established fast computation time, low memory footprint and in initial tests it outperformed SVM slightly. We trained our random forest classifier using the Weka API, with the default parameters. Once the classifier has been trained, classification can proceed in real-time.

EVALUATION

We conducted multiple studies, based on existing study designs, to evaluate several facets of RadarCat to support multiple purposes: i) everyday material and object classification [8] ii) transparent material classification [27] and iii) body parts classification [21]. Our results, both post-hoc and real-time analysis, show that it is accurate and robust. The studies were conducted in a quiet lab, with objects trained and tested in the same location. The Soli sensor was left powered-on for 10 minutes as a warm-up phase before the study.

Study One - Office and Kitchen Objects

This study aims to evaluate the classification accuracy and scalability of RadarCat on a broad range of everyday objects, such as those commonly found in the office and the kitchen. We selected 26 materials from our lab, as shown in Figure 1 and Figure 5. Following the procedure of lightweight material detection by Harrison and Hudson [8], we sampled the 26 materials twice a day for three days.

In each session, the material is being placed ten times on the sensor at different positions and orientation, by removing it and replacing it by hand. The material to be sampled is selected randomly from the pool of 26 materials, and no two materials were collected consecutively, to ensure that the sensor couple differently with the materials. Each time a material was sampled, five data points were recorded over a 0.17 second period (30Hz). After five iterations, the radar clutter map was rebuilt to reduce background noise. This produces 300 data points per material (6 sets of 50). Due to this large dataset, we performed offline analysis using the Weka toolkit.

We trained our random forest classifier using five of the six sessions of the collected data and then evaluate the classification accuracy using data from the remaining session. This leave-one-out process is repeated for all combinations of sessions (6 rotations), and the average accuracy is 96.0% (SD=1.3%). The confusion matrix is shown in Figure 5. Conventional 10-fold random holdout cross-validation using all samples which yields an optimistic accuracy of 99.97%.

Study Two - Transparent Materials

This study aims to evaluate the classification accuracy of RadarCat on transparent materials. We were able to source transparent materials from online plastic distributors and a local chemistry department, similar to that in SpecTrans [27], except cast acrylic. Figure 1 and Figure 6 lists all the transparent materials used in our study. All materials are in 3mm thickness and A4 size, except Borosilicate glass at 200x200mm, PVC at A3 size and microscope slide at 75x25mm. In addition, we add in extruded acrylic of different thickness (2,3,4,5,6,8,10mm) and extruded acrylic of same thickness (3mm) but with different dyes (red, green and blue), resulting in a total of 17 materials, including air.

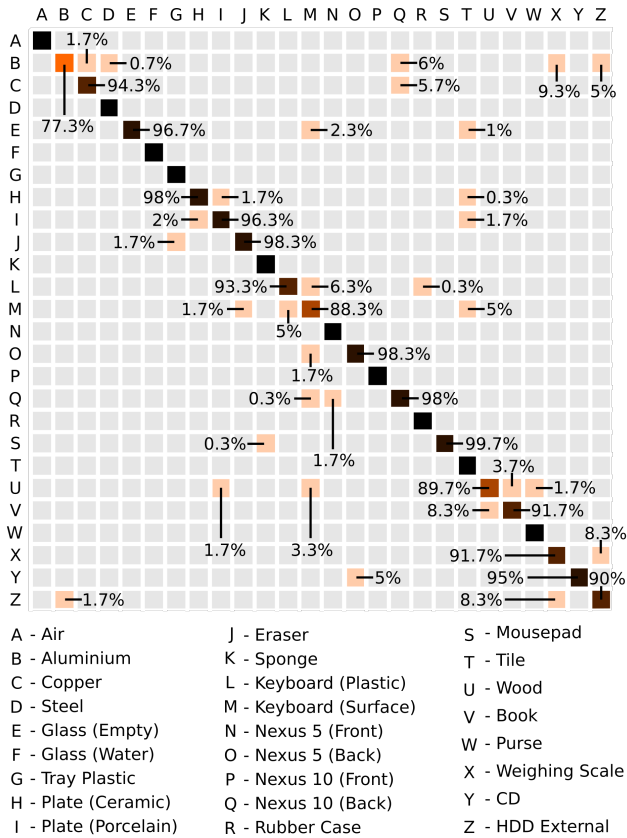


Figure 5. Confusion matrix for study one - 26 materials and objects commonly found, including air, as shown in Figure 1 (left).

We use the same procedure as the first study to collect the sample data (placing the objects ten times randomly, each time capturing five samples). Since the material is uniform when placed at different positions and orientations, we collected data for only three sessions, separated by one day each.

We trained our random forest classifier using two of the three sessions of the collected data and then evaluated the classification accuracy using data from the remaining session. This leave-one-out process is repeated for all combinations of sessions (3 rotations), and the average accuracy is 98.67% (SD=0.9%). The confusion matrix is shown in Figure 6. Conventional 10-fold random holdout cross-validation using all samples yields an optimistic accuracy of 100%.

Study Three - Body Parts

This study aims to evaluate the accuracy of RadarCat on classifying different body parts when they are touched. We initially selected a list of body parts to be tested following Bontential [21]. However, from a pilot test, we found that the upper arm and back of arm performed poorly in real time classification, even though post-hoc analysis shows a promising result [21]. Thus, we removed both from our experiment, leaving only palm, back of hand, finger, forearm, belly and calf. We further add in body parts covered by clothes: upper body wear, lower body wear, outerwear and glove, resulting in a total of 11 parts, including air, as shown in Figure 1 and Figure 7.

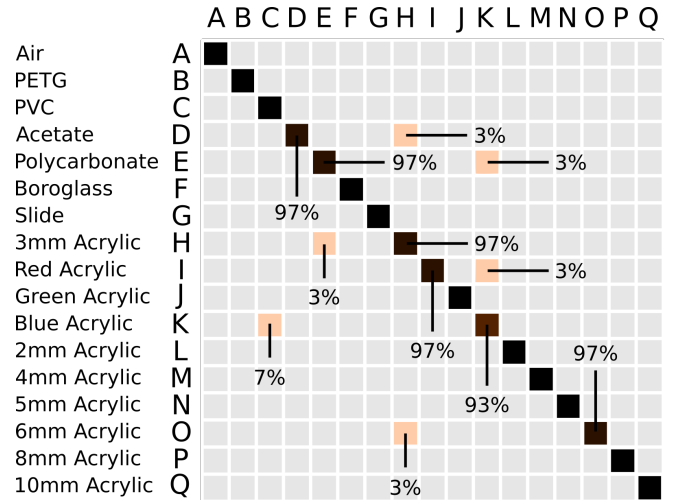


Figure 6. Confusion matrix for study two - 17 transparent materials.

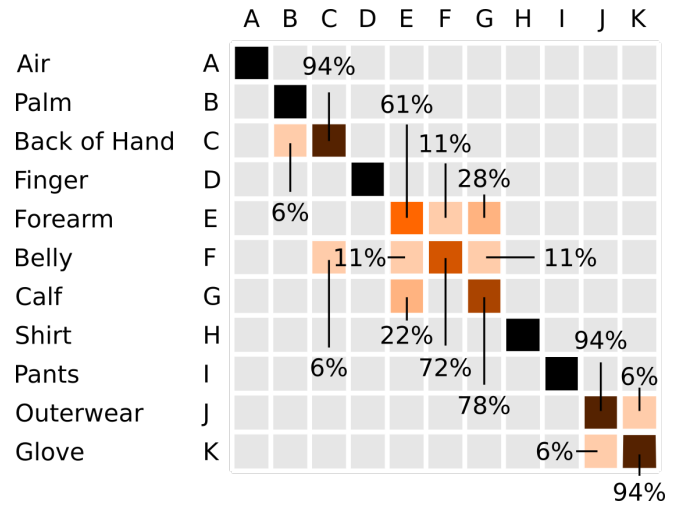


Figure 7. Confusion matrix for study three - 11 body parts.

We recruited 6 participants from local computer science department (2 females, mean age 20). Each study took about 30 minutes and participants were paid an Amazon voucher (5 GBP). During the data collection session, participants were instructed to put the sensor on different body parts and apply a small amount of pressure, as if they are using a stethoscope. We use the same procedure as the first and second study (placed ten times, each time capturing five samples). This procedure allows us to capture the variability performed by participant and we collected data for one session.

Because participants wore their own clothing, which are different among the participants, we employ per-user classification, where each participant had a custom classifier trained using his or her training data. This is ideal for personalized interaction with his or her own body parts but not for a generic classifier targeting all users. Following the training phase, we perform real-time classification evaluation, by using the collected data from the particular participant to initialize the system. Participants were requested to perform one of the gestures from the training set for three times. This was re-

peated for all the gestures. The experimenter then recorded the on-screen result, which was not visible to the participant.

Per-user Classifier

Real time evaluation using per-user classification shows an average accuracy of 90.4% (SD=13.6%) while post-hoc analysis using 10-fold cross validation yields 99.67% average accuracy. The confusion matrix is shown in Figure 7.

Generic Classifier

We also perform post-hoc analysis on the generic classifier (leave-one-out, using data from 5 users, test on remaining 1 user; note that they are all wearing different clothing) and the average accuracy is 62.15% (SD=11.69%, with clothing data) and 70.86% (SD=8.44%, clothing data removed), respectively. Although there are certain levels of cross user similarities, it is not reliable enough for general use, suggesting that per-user training is more appropriate for body parts classification. Nonetheless, conventional 10-fold cross validation with all six users data yields 99.82% accuracy (with clothing data) and 99.81% (clothing data removed).

DISCUSSION

While the levels of accuracy are very high overall, and the levels of confusion are very low as show in Figure 5, we can see where single material types and composites can be confused, for example, a macbook cover (aluminium) vs. weighing scale (contains aluminium). Likewise, wood can be occasionally be confused with a book. However, items (e) and (f), a filled and unfilled glass of water provide no confusion to each other, demonstrating the extent of the signal penetration and reflection required for RadarCat to disambiguate different materials and objects.

The results of this study further show that it is possible to recognize the front and the back of a mobile device, or recognize different models (e.g., Nexus5 and Nexus10). While the surface material of different models is the same (glass), the internal composition (different arrangement of the chipset) makes them differentiable by radar sensing. Some objects are composite (e.g., phone, tablet, eraser) which contains different materials in a thin form factor, while some materials are single and solid (e.g., glass, aluminium, plastic), while others have varying density throughout (e.g., wood). The states of electronic devices (switched on/off) did not affect the result.

Based on our testing, we can suggest that: i) For flat and solid materials (flush), very little training is needed to achieve the accuracy reported here, with the varieties of objects indicated. ii) For flat but low-density materials (sparse, hollow), more training from different positions and orientations of a single axis is needed but it is still possible to achieve high accuracy. iii) For non-flat materials, due to their geometry, more training from different positions and orientations from all three axes is required, which can limit the suitability of this training approach, but is an interesting direction for future research.

Experimentation with a smaller set of features, and a larger range of objects, in a wider set of scenarios is also required. Deep-learning methods are applicable to improve the scalability and generalization to everyday objects (e.g., Apples of

different sizes). In addition, the extent of material characteristics (e.g., types of liquid) is an area of rich future work.

Little can be added to the results presented in Figure 6 due to the high levels of accuracy reported. Given the frequency range of the radar we suggest it is the absorption/scattering properties and concentration of the dye we are classifying along with the different thicknesses of materials. This means, we can not only differentiate visually similar materials like different types of transparent plastics but also color. The extent of this, given different dye properties, thicknesses, material surface characteristics requires further exploration.

Finally, the results of the body parts study demonstrate that the forearm (e), belly (f) and calf(g) confused each other with, e - f (0.11), e to g(0.28), g to e (0.22), f - g (0.11), while the rest of body parts performed well. While this requires further study, we believe that this is due to these three body parts having somewhat similar structures (flat with mostly tissue and muscle) and hence appear similar to RadarCat, given that millimeter waves only penetrate shallowly into human tissue, typically less than 1mm [22].

In contrast, the palm, finger and back of hand each have very different structure near the 1mm range that the radar can “see”, due to the shape, nature of skin, bone and blood vessels near the surface, thus RadarCat differentiates this easily and has a higher accuracy. We also observe certain level of variability across users. For example, P1 has perfect accuracy on all trials on all body parts, while P4 has good accuracy on forearm and belly but 0 correct on calf and P5 has perfect accuracy on calf but 0 correct on forearm and belly (both are recognized as calf). Finally, body parts covered by clothes are often very accurate (Figure 7).

EXAMPLE USE SCENARIOS

There are many immediate applications that RadarCat can support (e.g., automatic waste sorting). Here we designed and implemented four example applications that demonstrate different interaction possibilities if the proximate target material or object is known (see Figure 8 and video figure). Following this, we speculate about potential applications and use cases.

Current Applications

Physical object dictionary - when an object is placed on the sensor (Figure 8a), the system can recognize the object and automatically search for relevant information or language translation, and then feedback this to the user. This can be useful because searching online often requires the user to know the name of an object in the first place. However, there are times when users are not aware of the name of an item - e.g., a specific phone model, which will make searching for it difficult. It also aids in learning environments because we can relate physical objects in-situ to improve learning efficiency.

Painting and photo editing application - users can use the RadarCat system as a physical probe instrument (as shown in Figure 8b), to quickly and intuitively change the operating mode (scale, rotate, pan) or the brush (size, color, style) depending on what the probe is sensing. For example, touching plastics of different materials switches the operating mode




Figure 8. Four example applications to demonstrate the interaction possibilities of RadarCat, from left to right a) physical object dictionary b) tangible painting app c) context-aware interaction and body shortcuts d) automatic refill.

while touching plastic of different color or thickness changes the brush's paint color and size.

Context-aware interaction and body shortcuts - with RadarCat attached to the back of a phone (as shown in Figure 8c), the system can tell whether the phone is held by bare palm or palm wearing a glove. This allows the phone to switch intelligently to easy mode - where the buttons are considerably larger to accommodate the fat finger problem when wearing a glove. In addition, touching different body parts activates different shortcut commands instantly. For example, touching the back of the hand, tummy (belly/trunk) and leg can be programmed to launch clock, food or map applications, respectively. It is also possible to know whether the phone is placed on the table, the sofa or inside the pocket (placement aware), facing up or facing down (situation aware), and allow the phone to switch into different modes automatically to adapt the environment, such as silent mode or loud speaker mode or turning the screen off to save battery.

Automatic refill - in a bar or restaurant scenario, where RadarCat sensors are embedded ubiquitously beneath the surface of a table, the system can tell whether a cup is full or empty, and if it is the latter, the system can alert the waiter for refill, all without intervention of the user (Figure 8d).

Future Applications

Recycle center - human intervention is often still needed to separate different types of waste, such as metal, glass and wood. With RadarCat, sorting waste can be automated.

Assisting the visually impaired - while one's sense of touch and smell can mitigate many of the challenges faced without sight, RadarCat once embedded in gloves or shoes can enhance one's understanding of the environment around you. Tactile paving on the sidewalk or limited braille interfaces might be replaced in the future.

Smart medical devices - current digital medical devices used outside the body, such as thermometers or stethoscopes still require the operator to manually note the different body parts being measured. Future devices with RadarCat embedded, can allow the automatic tagging of recorded temperature or sounds with the body part as it is detected.

LIMITATIONS

While RadarCat is a multi-purpose sensing system and achieves high accuracy in the three studies we conducted, it will not be suitable in all situations. For example, although our studies show that it can differentiate acrylic with different dyes with varying absorption/scattering properties (and hence

colors), this is mainly due to the high concentration of the dye component. In contrast, we were not able to differentiate single "stick-it note" of different colors. Therefore, in certain tasks, especially those involving identifying thin surface material based on color and texture we suggest imaging-based methods as described in our related work.

The radar hardware we used (Soli) is very sensitive, which introduces new problems. The high degree of sensitivity allows us to differentiate visually similar materials at high accuracy but this can be affected by background noise. In fact, the radar will gain clutter (reflections of unwanted objects) over time, especially before the hardware has reached a steady state temperature, or due to the movement or environmental changes. Radar clutter can impact the recognition rate, because the training data were collected with clutter removed. In practice, the clutter map can be easily rebuilt or by employing adaptive clutter removal [17]. Next version of Soli device removes this heat issue, and hence addressing the signal drift issue.

CONCLUSION AND FUTURE WORK

In this paper, we have presented RadarCat, a new sensing technique to provide versatile, multi-purpose material and object classification which enables novel forms of interaction. Our studies show that it is accurate and robust and we believe we have demonstrated its potential and implications in everyday interaction. Our technique can be used independently or combined with other sensing approaches. This can improve the sensing and computational edifice around which we realize new mobile, wearable and context-aware user interfaces.

Future work should explore a smaller set of channels, features and fewer sample points to explore the limits of object discrimination. For objects made of similar materials with strong radar reflections, investigation of its signature along different dimensions should be undertaken. In addition, the materials scattering and absorption properties at these wavelengths should be investigated further. We further wish to empirically validate the observed ability to recognize different fruits or credit cards, or counting the number of poker cards, or differentiating liquid content in a container. Finally, we would like to explore ways to encode information into an object, or stacking multiple layers of different materials which can be sensed, similar to [36] but in real time for interaction.

ACKNOWLEDGEMENT

The authors thank the Google Soli team for providing the hardware and software support. The research described here was supported by the University of St Andrews and the Scottish Informatics and Computer Science Alliance (SICSA).

REFERENCES

1. IEEE standard letter designations for radar-frequency bands. *IEEE Std 521-1984* (Nov 1984), 1–8.
2. Al-Nuaimy, W., Huang, Y., Nakhkash, M., Fang, M., Nguyen, V., and Eriksen, A. Automatic detection of buried utilities and solid objects with gpr using neural networks and pattern recognition. *Journal of Applied Geophysics* 43, 2 (2000), 157–165.
3. Bahl, P., and Padmanabhan, V. N. RADAR: an in-building rf-based user location and tracking system. In *Proceedings IEEE INFOCOM 2000, The Conference on Computer Communications, Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies, Reaching the Promised Land of Communications, Tel Aviv, Israel, March 26-30, 2000* (2000), 775–784.
4. Beaudouin-Lafon, M. Instrumental interaction: An interaction model for designing post-wimp user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '00, ACM* (New York, NY, USA, 2000), 446–453.
5. Buettner, M., Prasad, R., Philipose, M., and Wetherall, D. Recognizing daily activities with rfid-based sensors. In *Proceedings of the 11th International Conference on Ubiquitous Computing, UbiComp '09, ACM* (New York, NY, USA, 2009), 51–60.
6. El-Mahallawy, M. S., and Hashim, M. Material classification of underground utilities from gpr images using dct-based svm approach. *Geoscience and Remote Sensing Letters, IEEE* 10, 6 (2013), 1542–1546.
7. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., and Witten, I. H. The weka data mining software: An update. *SIGKDD Explor. Newsl.* 11, 1 (Nov. 2009), 10–18.
8. Harrison, C., and Hudson, S. E. Lightweight material detection for placement-aware mobile computing. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology, UIST '08, ACM* (New York, NY, USA, 2008), 279–282.
9. Harrison, C., Xiao, R., and Hudson, S. Acoustic barcodes: Passive, durable and inexpensive notched identification tags. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology, UIST '12, ACM* (New York, NY, USA, 2012), 563–568.
10. Harter, A., Hopper, A., Steggle, P., Ward, A., and Webster, P. The anatomy of a context-aware application. *Wirel. Netw.* 8, 2/3 (Mar. 2002), 187–197.
11. Ishii, H., and Ullmer, B. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '97, ACM* (New York, NY, USA, 1997), 234–241.
12. Jol, H. M. *Ground penetrating radar theory and applications*. Elsevier, 2008.
13. Kemp, M. C. Millimetre wave and terahertz technology for detection of concealed threats - a review. In *Infrared and Millimeter Waves, 2007 and the 2007 15th International Conference on Terahertz Electronics. IRMMW-THz. Joint 32nd International Conference on* (Sept 2007), 647–648.
14. Khan, U. S., Al-Nuaimy, W., and El-Samie, F. E. A. Detection of landmines and underground utilities from acoustic and gpr images with a cepstral approach. *Journal of Visual Communication and Image Representation* 21, 7 (2010), 731–740.
15. Kharkovsky, S., and Zoughi, R. Microwave and millimeter wave nondestructive testing and evaluation - overview and recent advances. *IEEE Instrumentation Measurement Magazine* 10, 2 (April 2007), 26–38.
16. Langen, B., Lober, G., and Herzig, W. Reflection and transmission behaviour of building materials at 60 ghz. In *Personal, Indoor and Mobile Radio Communications, 1994. Wireless Networks - Catching the Mobile Future., 5th IEEE International Symposium on* (Sep 1994), 505–509 vol.2.
17. Laput, G., Yang, C., Xiao, R., Sample, A., and Harrison, C. Em-sense: Touch recognition of uninstrumented, electrical and electromechanical objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology, UIST '15, ACM* (New York, NY, USA, 2015), 157–166.
18. Le Goc, M., Taylor, S., Izadi, S., and Keskin, C. A low-cost transparent electric field sensor for 3d interaction on mobile devices. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems, CHI '14, ACM* (New York, NY, USA, 2014), 3167–3170.
19. Lien, J., Gillian, N., Karagozler, M. E., Amihoud, P., Schwesig, C., Olson, E., Raja, H., and Poupyrev, I. Soli: Ubiquitous gesture sensing with millimeter wave radar. *ACM Trans. Graph.* 35, 4 (July 2016), 142:1–142:19.
20. Maldonado, A., Alvarez, H., and Beetz, M. Improving robot manipulation through fingertip perception. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on* (Oct 2012), 2947–2954.
21. Matthies, D. J. C., Perrault, S. T., Urban, B., and Zhao, S. Botential: Localizing on-body gestures by measuring electrical signatures on the human skin. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI '15, ACM* (New York, NY, USA, 2015), 207–216.
22. Nelson, D. A., Nelson, M. T., Walters, T. J., and Mason, P. A. Skin heating effects of millimeter-wave irradiation-thermal modeling results. *IEEE Transactions on Microwave Theory and Techniques* 48, 11 (Nov 2000), 2111–2120.

23. Pasolli, E., Melgani, F., and Donelli, M. Automatic analysis of gpr images: A pattern-recognition approach. *Geoscience and Remote Sensing, IEEE Transactions on* 47, 7 (2009), 2206–2217.
24. Quigley, A., Dix, A., Mackay, W. E., Ishii, H., and Steimle, J. Visions and visioning in chi: Chi 2013 special interest group meeting. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '13, ACM (New York, NY, USA, 2013), 2545–2548.
25. Rekimoto, J., and Ayatsuka, Y. Cybercode: Designing augmented reality environments with visual tags. In *Proceedings of DARE 2000 on Designing Augmented Reality Environments*, DARE '00, ACM (New York, NY, USA, 2000), 1–10.
26. Roodaki, P. M., Taghian, F., Bashirzadeh, S., and Jalaali, M. A survey of millimeter-wave technologies. In *Electrical and Control Engineering (ICECE), 2011 International Conference on* (Sept 2011), 5726–5728.
27. Sato, M., Yoshida, S., Olwal, A., Shi, B., Hiyama, A., Tanikawa, T., Hirose, M., and Raskar, R. Spectrans: Versatile material classification for interaction with textureless, specular and transparent surfaces. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, ACM (New York, NY, USA, 2015), 2191–2200.
28. Sharan, L., Liu, C., Rosenholtz, R., and Adelson, E. H. Recognizing materials using perceptually inspired features. *International Journal of Computer Vision* 103, 3 (2013), 348–371.
29. Skolnik, M. I. Introduction to radar. *Radar Handbook* 2 (1962).
30. Project Soli.
<https://www.google.com/atap/project-soli/>.
31. Song, J., Wang, S., Lien, J., Poupyrev, I., and Hilliges, O. Interacting with soli: Exploring fine-grained dynamic gesture recognition in the radio-frequency spectrum. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology*, UIST '16, ACM (New York, NY, USA, 2016).
32. Sugama, Y., Murase, T., and Fujii, Y. Projection based virtual tablets system involving robust tracking of rectangular objects and hands. In *Proceedings of the 7th Augmented Human International Conference 2016*, AH '16, ACM (New York, NY, USA, 2016), 33:1–33:2.
33. Walabot. <http://walabot.com/>.
34. Weiser, M. The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev.* 3, 3 (July 1999), 3–11.
35. Wiese, J., Saponas, T. S., and Brush, A. B. Phoneprioception: Enabling mobile phones to infer where they are kept. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, ACM (New York, NY, USA, 2013), 2157–2166.
36. Willis, K. D. D., and Wilson, A. D. Infrastructs: Fabricating information inside physical objects for imaging in the terahertz region. *ACM Trans. Graph.* 32, 4 (July 2013), 138:1–138:10.
37. Yilmaz, A., Javed, O., and Shah, M. Object tracking: A survey. *ACM Comput. Surv.* 38, 4 (Dec. 2006).
38. Zhu, Y., Zhu, Y., Zhao, B. Y., and Zheng, H. Reusing 60ghz radios for mobile radar imaging. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, MobiCom '15, ACM (New York, NY, USA, 2015), 103–116.