

Pyro: Thumb-Tip Gesture Recognition Using Pyroelectric Infrared Sensing

Jun Gong¹, Yang Zhang², Xia Zhou¹, Xing-Dong Yang¹

Dartmouth College¹, Carnegie Mellon University²

{jun.gong.gr; xia.zhou; xing-dong.yang}@dartmouth.edu, yang.zhang@cs.cmu.edu

ABSTRACT

We present Pyro, a micro thumb-tip gesture recognition technique based on thermal infrared signals radiating from the fingers. Pyro uses a compact, low-power passive sensor, making it suitable for wearable and mobile applications. To demonstrate the feasibility of Pyro, we developed a self-contained prototype consisting of the infrared pyroelectric sensor, a custom sensing circuit, and software for signal processing and machine learning. A ten-participant user study yielded a 93.9% cross-validation accuracy and 84.9% leave-one-session-out accuracy on six thumb-tip gestures. Subsequent lab studies demonstrated Pyro's robustness to varying light conditions, hand temperatures, and background motion. We conclude by discussing the insights we gained from this work and future research questions.

Author Keywords

Micro thumb-tip gesture; PIR; pyroelectric; wearable

ACM Classification Keywords

H.5.2. [User interfaces] – Input devices and strategies.

INTRODUCTION

Micro finger gestures [1, 34, 62] offer new opportunities for natural, subtle, fast, and unobtrusive interactions in wearable, mobile, and ubiquitous computing applications. For example, gesturing the thumb tip against the tip of the index finger [1] is a natural method of performing input, requiring little effort from users because the index finger serves as a supporting surface to naturally provide haptic feedback. This motion introduces less fatigue over time compared with traditional gestural input methods, which often require moving the finger, hand, or even the entire arm in mid-air [16, 33].

Despite the known benefits of this new input modality, tracking fine-grained thumb-tip gestures remains very challenging due to the small magnitude of finger motions and frequent occurrences of self-occlusion. Existing studies have exploited magnetic sensing [1, 19], which achieves a relatively high tracking precision but requires fingers to be

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 2017, October 22–25, 2017, Quebec City, QC, Canada

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4981-9/17/10...\$15.00

<https://doi.org/10.1145/3126594.3126615>

instrumented with magnets and sensors. The Soli project [34, 62] explored the use of millimeter-wave radar to sense subtle finger movement without instrumenting the user. The active sensor's energy consumption, however, is a concern, especially for small wearable devices (e.g., smart watches).

In this paper, we propose an alternative approach that senses thermal infrared signals radiating from fingers to recognize micro thumb-tip gestures (Figure 1). We sense these signals using a passive infrared (PIR) sensor made of pyroelectric materials. A PIR sensor is highly sensitive to subtle motion and thus enables recognition of fine gestures. This passive sensing approach provides two unique benefits. First, by eliminating the need to generate active signals, the sensing technique itself is energy-efficient. It is preferable for small wearable devices. Second, the PIR sensor generates very little heat and thus requires no cooling [26]. It is an important benefit for wearable devices since cooling is a known challenge in engineering small consumer devices [48].




Figure 1. Sensing micro thumb-tip gesture using a PIR sensor.

We demonstrate the technical feasibility through Pyro, a proof-of-concept prototype developed using a low-cost, off-the-shelf PIR sensor (Figure 1). We augment the PIR with customized sensor electronics and optimize Pyro for detecting micro thumb-tip gestures performed close to the sensor. We test the system using six thumb-tip gestures: a triangle, rectangle, circle, question mark, check mark, and finger rub (Figure 3). Results from ten participants show 93.9% cross-validation accuracy and 84.9% leave-one-session-out accuracy. Additionally, our study provides insights into the robustness of this approach under environmental noises such as ambient light interference, hand temperature variations, and background hand movement. Our work provides the first evidence to support pyroelectric infrared sensing as a promising alternative for detecting micro finger gestures.

Our primary contributions include: 1) an approach to detect micro thumb-tip gestures using pyroelectric infrared sensing; 2) development of a prototype using an off-the-shelf sensor and customized hardware and software; and 3) initial validation of this approach through a series of experiments.

RELATED WORK

In this section, we briefly summarize previous research based on various sensing techniques.

Camera-based Sensing

Camera-based approaches have shown good accuracy in tracking small finger motions. For example, 2D images of a hand in different angles can be used to query a database of existing hand models to find a best match [31, 57]. Recent work by Song et al. [46] shows a technique for 3D hand gesture recognition using a single 2D camera. Wrist-worn [33], finger-worn [18] and head-mounted [21] cameras have also been used to track small finger motions. In recent work, depth cameras have been widely used for improved accuracy. Much of the existing work uses machine learning classifiers to recognize hand postures [4, 36, 37, 45]. RetroDepth [9] took a different approach by sensing 3D hand postures using the silhouettes of the hands. Although tracking is precise, camera-based approaches have been criticized for being bulky and power consuming, making them hard to integrate into small wearable devices.

RF Sensing

RF signals (e.g., Wi-Fi, GSM, radar signal) have also been shown to be effective for detecting finger gestures. Coarse-grained hand gestures (e.g., flick, slide, hover) can be sensed using Wi-Fi [8, 50] and GSM [17]. Mudra [64] is a fine-grained finger gesture recognition system which senses finger motion using Wi-Fi signals. Soli [34] provides a promising alternative approach. The technology tracks very small finger movements using 60-GHz radar signals. Soli is capable of detecting 11 hand gestures [62], although only four of them are micro (pinch index, pinch pinky, finger slide, and finger rub). A common concern of 60-GHz radar, however, is power consumption, especially in the context of wearables. Compared to Soli, our pyroelectric infrared sensing is passive, which significantly reduces the power consumption.

Pyroelectric Infrared Sensing

Pyroelectric infrared sensors are sensitive to thermal radiation emitted by the human body (8 - 14 μm) [13]. Tiny deviations from the thermal equilibrium of the surrounding environment can be detected [13, 63]. Pyroelectric infrared sensing is commonly used in commercial applications to detect the presence of humans or trigger alarms. PIR sensors have also been explored for much more complex applications such as human localization [14, 27, 30, 32, 41, 47, 52, 53, 54, 56], motion direction detection [10, 23, 25, 42], thermal imaging [11], radiometry [38], thermometers [15], and biometry [10, 26, 49, 55, 56]. Most prior work in this space has focused on detecting large and coarse-grained body movement happening at a relatively long distance from the

PIR sensor ($>2\text{m}$). For shorter-distance sensing, a 4×4 PIR sensor array has been used to identify hand motion in four directions at a distance of tens of centimeters [61]. In our work, we explore PIR sensing for detecting nearby micro and fine-grained thumb-tip gestures for wearable applications.

Other Sensing Techniques

Thumb-tip movements can also be sensed using magnetic sensors [1, 19]. The limitation of this approach, however, is the need to instrument the fingertips with magnets and sensors. Acoustic sensing [39, 58] also shows potential, but no existing system has demonstrated feasibility in recognizing thumb-tip gestures with micro finger movement. A variety of sensing techniques have been developed to detect the commonly-used pinch gestures (e.g., thumb touching the other fingers) [2, 12, 22, 29, 35, 44, 65]. GestureWrist [43] is one of the earliest examples, which uses an array of capacitive sensors to detect the changes in the shape of the forearm to identify different finger pinches. Recent research has shown that the forearm shape can also be detected using infrared photo reflectors [24, 40]. Sensing resolution can be further improved using pressure sensors [22] or electrical impedance tomography sensors [65].

SENSING PRINCIPLE

PIR sensors are made of pyroelectric crystals, a material that generates a surface electric charge when exposed to heat in the form of infrared radiation. Commercial PIR sensors are typically tuned for human detection by adding a bandpass filter window which only passes the infrared wavelengths emitted by the human body (e.g., 8 - 14 μm) (Figure 2). In the presence of a thermal object (e.g., a finger), PIR sensors convert the thermal radiation into an electrical current proportional to the difference in temperature between the finger and the environment [13].

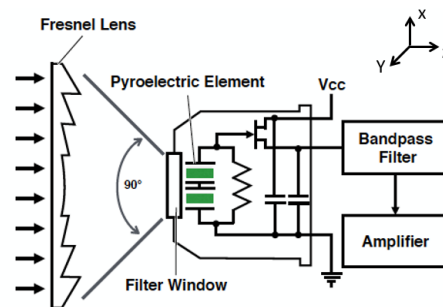


Figure 2. Pyroelectric infrared sensing principle.

A PIR sensor commonly arranges two sensing elements side by side, connected to a differential amplifier to cancel common-mode noise caused by environmental temperature change, vibration, and sunlight, since these simultaneously affect both elements. When a finger passes by, though, it is observed by one element first and then the other, which causes a positive differential change between the two crystals (e.g., generating a sinusoidal swing). When the object crosses from the opposite direction, it intercepts the elements in a reverse order, thus generating a negative differential change (e.g., a flip of the sinusoidal swing). When the change

in thermal infrared has stabilized between the two crystals, the signal returns to its baseline voltage. Thus, if the finger remains still, no output signal will be generated. PIR sensors are less responsive to motion towards or away from the sensor since the motion in z-axis causes a smaller difference in temperature between the two crystals (Figure 2).

To make the sensor responsive to tiny movements, a Fresnel lens can be added to concentrate incoming radiation on the sensing elements (Figure 2). To further improve sensitivity, the Fresnel lens can be split into multiple zones, each with its own sub-lens focused on all sensing elements. The downside of using a multi-zone Fresnel lens, however, is that the finger’s movement direction cannot be reliably detected due to the mixture of multiple signals coming from different zones. Thus, we used a single-zone Fresnel lens for the Pyro prototype.

GESTURE SET

Thumb-tip gestures are performed by moving the thumb tip against the tip of the index finger, which is natural, subtle, fast, and unobtrusive [1]. While the design space of thumb-tip gestures is large, we focus our exploration on free-form shape gestures carried out on the distal segment of the index finger as it is the most common and intuitive way to perform the gestures. Since drawing the thumb on the index finger resembles gesturing on a touchscreen, we choose five gestures from known unistroke gestures shown to be useful on touchscreen devices [59, 60] (Figure 3).

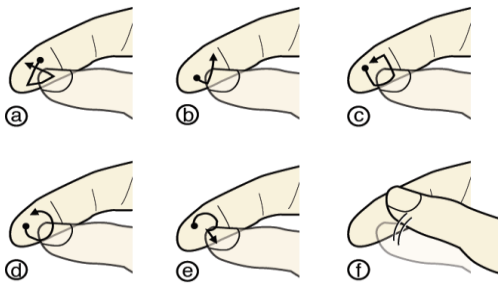


Figure 3. Gesture set: (a) triangle; (b) check mark; (c) rectangle; (d) circle; (e) question mark; (f) finger rub.

To ensure diversity, we picked unistroke gestures with straight lines and corners of different degrees (counter clockwise triangle, check mark, and counter clockwise rectangle), one with a curved path (counter clockwise circle) and one mixing a curve, straight line, and corner (question mark). We also added the finger rub gesture from [62]. Although this set of gestures is not exhaustive, it is so far the largest micro-gesture set that has been used to validate a sensing technique.

PYRO IMPLEMENTATION

We created a self-contained prototype using our customized hardware and software. This section describes our implementation details.

PIR Sensor and Fresnel Lens

We optimized our hardware for finger motion close to the sensor. To achieve this, we chose a single-zone Fresnel lens (IML-0637 from Murata Manufacturing Co.) and a PIR sensor (IRA-E710 from Murata Manufacturing Co.) without the built-in amplifier and bandpass filter. As mentioned previously, the single-zone Fresnel lens is chosen over the multiple-zone lens to preclude interference from multiple monitoring zones. Our system’s horizontal and vertical field of view are both 90 degrees. Figure 4 shows a smartwatch prototype augmented with Pyro. A pilot study with 3 participants suggested that the orientation of the crystal elements does not affect gesture recognition accuracy, so we aligned the elements parallel to the table.




Figure 4. A smartwatch prototype augmented with Pyro.

Sensing Board

We built our customized sensing board (Figure 5) around a Cortex M4 micro-controller (MK20DX256VLH7 [3]) running at 96MHz, powered by the Teensy 3.2 firmware [5]. The board has an LM324 [6] based ADC preamp, a power management circuit, and a Bluetooth module. To reduce the dominant noise (50 kHz - 300 kHz) caused by powerline and fluorescent light ballasts, we implemented a bandpass filter with cut-off frequencies of 1.59 Hz and 486.75 Hz. The relatively wide bandwidth gives us the flexibility to explore sampling rates. After the noise is removed, the input signal is amplified with a gain of 33 and biased by AREF/2 (1.5 V) to preserve the fidelity of the analog signal. The gain value is carefully tuned to have an optimal sensing range of approximately 0.5 cm to 30 cm away from the PIR sensor. This design mitigates the background thermal infrared signals from the human body minimizing the impact on the foreground finger gesture signal.




Figure 5. Pyro sensing board.

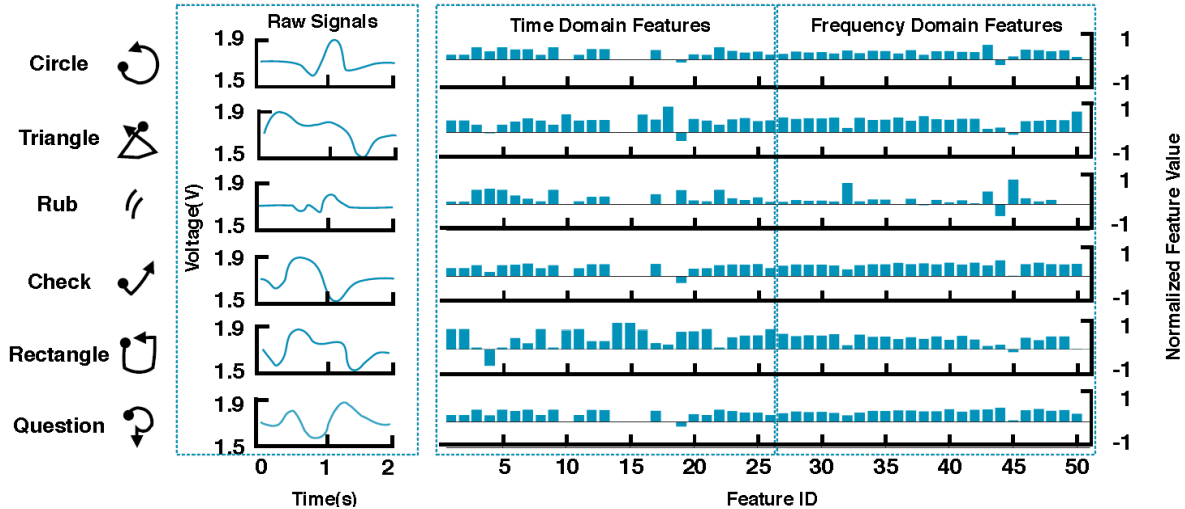


Figure 6. Top 50 features of six thumb-tip gestures.

Although existing literature suggests that the PIR signals should be better sampled at 10 Hz for detecting human body movement [51], we found 20 Hz works better for micro finger gestures. This is because the frequency of PIR signals generated by nearby-finger movement is between 2 Hz and 10 Hz. Finally, PIR signals are sent to a laptop through Bluetooth for further computation. In total, our prototype costs \$24. It can be made smaller and cheaper in high volume commercial applications.

Machine Learning

We use machine learning to classify thumb-tip gestures. While there are many options for classification algorithms (e.g., Hidden Markov Models and Convolutional Neural Networks), many of them are computationally expensive, and therefore potentially unsuitable for real-time applications on low-power platforms such as smartwatches [34]. We aim to strike a balance between recognition accuracy and computation efficiency. As such, we narrowed the candidate gesture recognition methods to Random Forest, Support Vector Machine, and Logistic Regression. After comparing their recognition accuracy (e.g., results shown in Figure 9), we decided to use Random Forest in our implementation. Random Forest has previously been found to be accurate, robust, scalable, and cost-efficient in computation when tracking micro gestures using radar [34] or computer vision [18] techniques.

Feature Extraction

Like any machine learning application, extracting relevant features is critical to the success of Pyro. The challenge, however, lies in the fact that selecting the right feature set is not obvious. Although features like FFT, peak amplitude or first-order derivative are commonly used in various applications, we found that using them directly to train a Random Forest model led to a rather low accuracy and none of the existing research provided insights into suitable features for characterizing micro thumb-tip gestures using pyroelectric infrared signals. We decided to use *tsfresh* [7], a

feature extraction toolbox, to extract hundreds of features from time and frequency domains. We sampled PIR signals, made them equal length with zero padding, and normalized them. We then extracted features and used these features to train and test the models. Results are reported in the later sections. Table 1 shows the top-50 most effective and relevant features ranked by Random Forest. Interestingly, half of them are from the time domain and the remaining half are from the frequency domain. This confirms that data from both domains are treated equally important by Random Forest. Figure 6 presents the normalized values of the top-50 features (same order as in Table 1) and raw signals for the six thumb-tip gestures.

Time Domain (26 features)	<ul style="list-style-type: none"> • Statistical Functions (21): Sum, Mean, Median, Standard Deviation, Skewness, Quantiles (4), Kurtosis, Longest strike above/below mean, Count above/below mean, mean autocorrelation, mean absolute change quantiles (3), autocorrelation of lag, ratio of unique values, Variance • Peak (1): Number of values between max and min • Entropy (3): Binned Entropy, Sample Entropy, Approximate Entropy • Energy (1): Absolute energy
Frequency Domain (24 features)	<ul style="list-style-type: none"> • Continuous Wavelet Transform (21) • Fast Fourier Transform (1) • Autoregressive (1) • Welch (1)

Table 1. Top-50 features ranked by Random Forest.

USER EVALUATION

The goal of this study is to validate Pyro’s gesture recognition accuracy, as well as its robustness against individual variance and among different users.

Participants

Ten right-handed participants (average age: 26.4, two female) were recruited to participate in this study.

Participants' finger temperatures measured between 24.1 °C and 34.4 °C (SD = 4.6). The room temperature was 24 °C.

Data Collection

Each participant was instructed to sit in front of the PIR sensor placed on a desk. Before a session started, participants were given several minutes to learn the six unistroke gestures (triangle, rectangle, circle, question mark, check mark, and finger rub). After the short training session, each participant performed the gestures roughly 0.5 cm to 7 cm in front of the PIR sensor using their right hand. Participants were not given any instruction on how to perform the gestures (e.g. magnitude or duration), except the direction in which the gestures should be drawn. The start and end of each gesture was indicated by clicking a computer mouse using their left hand. Each gesture was repeated 20 times in each session [19, 31, 50, 51], which took about 15 minutes to complete. A five-minute break was given between sessions, where participants were asked to leave the desk and walk around the lab. Data collection finished after three sessions. The study took about an hour to complete for each participant. In total, we collected 3600 samples (10 participants × 6 gestures × 20 repetitions × 3 sessions) for analysis.

Result

We present experiment results to demonstrate the accuracy and reliability of our system.

Within-User Accuracy

Within-user accuracy measures the prediction accuracy where the training and testing data are from the same user. For each participant, we conducted a twofold cross validation, where half of the data was used for training and the remaining half used for testing. The overall within-user accuracy was calculated by averaging the results from all the participants. The result yielded an accuracy of 93.9% (SD = 0.9%). Figure 7 left shows the confusion matrix.

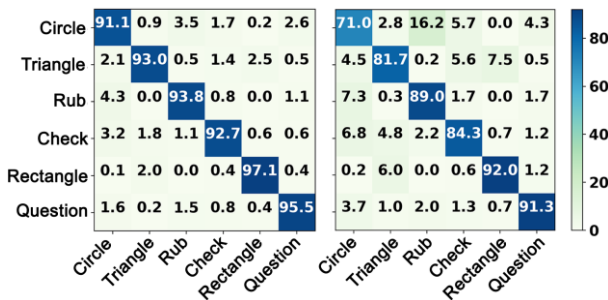


Figure 7. Confusion matrices. Left: cross validation accuracies; Right: leave-one-session-out accuracies.

Reproducibility

Reproducibility measures how stable and scalable the system is against the data collected from a different session. To measure the system reproducibility, we calculated the leave-one-session-out accuracy for each participant by training the model using the data from two sessions and testing it using the remaining session. The average accuracy for each participant was calculated by averaging all possible

combinations of training and test data. The overall accuracy was then calculated by averaging the accuracy from all participants. The result yields 84.9% accuracy (SD = 3.5%). Compared with cross-validation accuracy, this result reflects a more realistic situation. Figure 7 right shows the confusion matrix. Rectangle received the highest accuracy (i.e., 92%) among all six gestures. A potential reason is that the rectangular trajectory has many sharp turns that make the signal more distinguishable than others. The mix of curves and a sharp turn in the question mark may also contribute to the higher accuracy. Most gestures (except rectangle) are more likely to be confused with circle, and vice versa (Figure 7 left). This can be attributed to many factors (e.g. gesture geometry, how gestures were drawn, and recognition algorithm) and requires further investigation. The trend is similar between within-user accuracy and leave-one-session-out accuracy, where rectangle and question mark received higher scores than others, while circle remained the most confusing gesture. These results suggest that gestures with higher accuracy were also drawn more consistently across sessions.

Universality

Universality measures whether an existing model works across different users. To calculate the accuracy, we used the data from nine participants for training and the remaining one for testing. The overall accuracy was then calculated by averaging the results from all ten combinations of training and test data. The overall accuracy is 69% (SD = 11.2%), which indicates that different users performed gestures differently even though the internal consistency is quite high for each individual participant. Figure 8 left shows the confusion matrix of all six gestures, from which we found that check mark (48.2%) and circle (58.5%) contributed the most to the error. We then removed them and calculated the accuracies using the remaining data. The result yielded a higher accuracy of 76.3% (SD = 6.8%) without check mark and 87.6% (SD = 6.7%) without both (Figure 8 right).

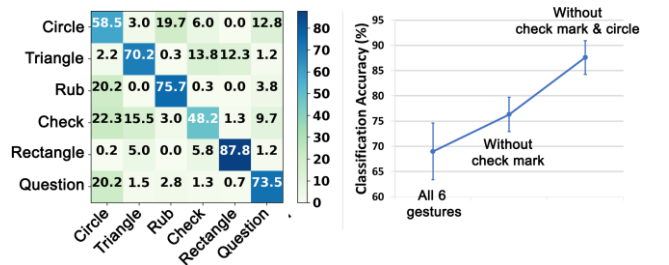


Figure 8. Left: confusion matrix of cross-user accuracies; Right: cross-user accuracy with gesture sets of different sizes.

Prediction Methods

With the number of different options available for prediction methods, we were also interested in measuring how well they perform on our data. We ran our data with four additional methods, including Poly Kernel Support Vector Machine (SVM), RBF Kernel Support Vector Machine, Logistic Regression, and Dynamic Time Warping (DTW), each with different strengths and weaknesses. Similar to [34], we did

not try Hidden Markov Models and Convolutional Neural Networks as they require significant computational power, making them less suitable for small computing devices. We report the prediction accuracy obtained from each method by showing the cross-validation accuracy, leave-one-session-out accuracy, and leave-one-subject-out accuracy (Figure 9). The result shows that Random Forest outperformed all other tested methods on all three metrics, followed by SVM with a Poly Kernel.

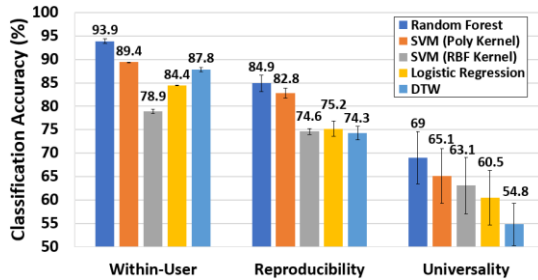


Figure 9. Recognition accuracy under various prediction methods.

SUPPLEMENTARY STUDY: ENVIRONMENTAL NOISE

Micro finger gestures will be performed in noisy and inconsistent environments. Thus, we conducted initial experiments in a controlled lab environment to evaluate how robust our system is against common environmental noises, such as ambient light and nearby hand movements. Additionally, we also measured the impact of rapid changes in hand temperature. This study was carried out with a single participant (male, right-handed, 25 years old).

Data Collection

The data collection procedure was similar to the user evaluation, except that we collected only two sessions of data. Both sessions were used for training. Since no ambient noise was presented, the prediction model was created under a clean and controlled environment, which we believe is the easiest way to model in real practice. Our goal was to test the performance of this model under varying noise conditions. In total, we collected 240 (6 gestures \times 20 repetitions \times 2 sessions) gestures to train our prediction model. Test data was collected in separate sessions under different noise conditions. For both training and testing, the participant performed the gestures roughly 0.5cm to 7cm in front of the PIR sensor using his right hand. Room and finger temperatures measured around 23°C and 35°C respectively prior to the experiment.

Ambient Light

A PIR sensor senses thermal infrared with wavelengths ranging from 8 μ m to 14 μ m, which is not emitted by most indoor light sources (e.g., LED, fluorescent lights) and yet is contained in sunlight. Thus, we focused on understanding how much sunlight affects the sensing performance. We collected test data (6 gestures \times 20 repetitions \times 2 sessions) under two lighting conditions: *dark* (0 lx – 20 lx, a dark room without any sunlight) and *bright* (200 lx – 300 lx, under sunlight leaked through a window). Data for both conditions

were collected indoors to ensure the consistency of the environmental temperature.

The result shows that the clean model achieves 82.5% and 84.2% accuracy in dark and bright condition respectively. This is similar to the leave-one-session-out accuracy in Study 2, indicating that interferences from ambient thermal infrared have little effect on the sensing performance in our set-up. This is expected because the differential amplifier of our PIR sensor cancels out any ambient interference that equally affects both sensing elements. More evaluation, however, should be done outdoors to fully understand the effect of ambient light (e.g. whether the sensor will be saturated when sun light is too strong).

Nearby Hand Movement

We also tested the robustness of our system against background hand movements. Another person waved their hand in random trajectories behind the participant’s fingers in a distance no further than 30 cm away from the sensor to create background noise. In total, 120 gesture instances (6 gestures \times 20 repetitions \times 1 session) were collected for testing. The result was 86.7% accuracy, which is again similar to those found in the other conditions, indicating that background hand movement does not have a negative impact on sensing micro thumb-tip movement in our settings. We believe it is because 1) the foreground hand blocks background objects from the sensor’s viewing angle and 2) the amplifier gain was adjusted to limit sensing long-range motion.

Hand Temperature

Hand temperature may change drastically after the hand holds a hot or cold object (e.g., a cup of a hot or cold drink). To understand whether the rapid, significant change in finger temperature affects sensing performance, we varied the temperature of the participants’ fingers by asking the participant to hold a cup of hot water or soak fingers in ice water before performing gestures. In the hot condition, the fingertips measured around 41°C after holding a cup of hot water for several minutes whereas in the cold condition, the fingertips measured around 24°C after soaking fingers in ice water for several minutes. The participant started gesturing immediately after the temperature was set. The finger temperature returned to around 36°C at the end of the hot finger session and 34°C at the end of the cold finger session.

We observed that hot fingers did not introduce a visible impact on the analog signal. The resulting 85.8% accuracy further confirmed that a rapid increase in finger temperature does not negatively affect recognition accuracy. In contrast, when the hand was cold, the analog signal became visually weaker. However, the signal quickly returned to the normal scale after the hand temperature reached to 27°C (within roughly 3 seconds in a room temperature of 23°C). Although we found that the overall prediction accuracy was not affected (i.e., 83.3%), the hand temperature increased too quickly to allow us to draw a conclusion. To extend our understanding on the effect of cold fingers, we collected

another set of gestural data, where we controlled the finger temperature within a range between 24°C and 26°C. The result yields 53% accuracy, which suggests that recognition accuracy was affected by the significant drop of hand temperature. It is because a smaller temperature difference between the finger and environment causes weaker signals when hand temperature drops significantly. Thus, the system performance will likely be affected if our model is used in cold temperature conditions, but the issue may go away quickly once the hand returns to a normal temperature.

Overall, the results of this study are encouraging. They provide insights into the pyroelectric infrared sensing in varying usage conditions, and the robustness of our system against tested noises.

DEMO APPLICATIONS

We implemented two demo applications to showcase Pyro’s potential on wearable devices. Our first application is a video player on a smartwatch. We created a smartwatch prototype using a 2” TFT display, a 3D printed case, and the Pyro system. First, the user can draw a circle on their index finger as a shortcut to launch the video player app. This way the user does not need to browse the app list to find the app. Unlike the existing video players on smartwatches, where the control panel can occlude the screen content, our application allows the user to draw thumb-tip gestures to control the video. For example, the user can rub their finger to play or pause the video (Figure 10 left). Drawing a question mark shows the information of the video, such as title and year.

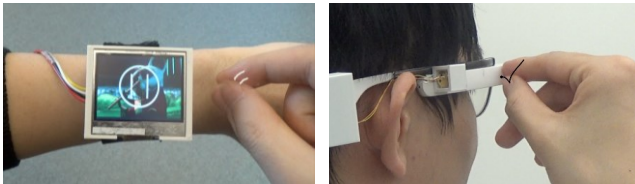


Figure 10. Left: A user rubs the fingers to play/pause a video; Right: Drawing a check mark with touch takes a photo and shares it on Facebook.

Our second application allows the user to interact with a head-worn display using the thumb-tip gestures. We augmented a Google Glass using Pyro. The sensor is placed beside the touchpad near the ear. This provides a new input channel on Google Glass. Additionally, it also allows the touchpad and thumb-tip input to be used jointly. With this new style of joint input, many novel interactions can be performed. For example, thumb-tip gestures performed with and without the index finger touching the touchpad can lead to different actions. Touching the touchpad in different locations may also lead to different actions. In our application, a check mark gesture is a shortcut for taking a photo while a check mark gesture with the index finger touching the touchpad will take the photo and share it on Facebook (Figure 10 right). Alternatively, performing a thumb-tip gesture before or after gesturing on the touchpad can trigger different actions. This style of input is similar to

Air+Touch [20], but without the need of an expensive camera-based sensing technique. In our application, rubbing the thumb and index finger before swiping the touchpad zooms the map in or out whereas swiping without rubbing pans the map.

DISCUSSION AND LIMITATIONS

In this section, we discuss the insights gained from this work, propose future research, and acknowledge the limitations.

Gesture delimiter. The focus of this work is the sensing technique. The gesture delimiter, however is an important topic to study in the future. A number of options exist. For example, distinguishable signals from the hand entering or leaving the sensor’s active region can be used as an explicit delimiter. To quickly validate this method, we conducted an informal study, where we recruited 3 male participants (average age: 26.7) and trained a two-class classifier (6 micro gestures vs hand-in/out) using 120 samples for each class. Overall, we collected 720 samples (2 class × 120 samples × 3 participants) for analysis. A two-fold cross validation yields a 98.6% (SD = 0.4%) mean accuracy. The result is very promising. Future implementations include developing a hierarchical classifier, where the first classification layer determines the start or end of a gesture, and the second layer predicts micro gestures that the user performs.

False positives. Coarse-grained movements, such as a person passing by the sensor, may generate signals similar to hand motions, and so future research should focus on reducing false positives. Our initial tests indicate that body movement more than 40 cm away from the sensor generates much weaker signals that can be distinguished from hand-in/out. We believe this can filter out many ambient motion noise in public settings. According to Edward Hall’s theory of interpersonal spatial relationships, 40 cm is still within the distance between people in a social environment [28], so body movements from a nearby colleague or friend may accidentally trigger the delimiter. A potential solution is to reduce the focal distance of the Fresnel lens to around 10 cm, which filters out motion noises in many social activities.

Additionally, smartwatches have a built-in mechanism to turn on the screen by detecting the user’s intention to use the smartwatch. Pyro can leverage this mechanism and only activate the sensor when the smartwatch screen turns on. Whirling the wrist of the hand wearing the smartwatch might introduce false positives. Activating the sensor only when the touchscreen is on can reduce the error. Interacting with the touchscreen might also cause false positives but the PIR sensor can be deactivated if the smartwatch detects a touch event. Future research will carefully validate the effectiveness and usability of different options and techniques to avoid false positives.

Evaluation. Although our supplemental studies show some promising system robustness against different lighting conditions, hand temperatures, and background motion noises, further evaluation should be done in more diverse and

realistic settings (e.g., outdoors). Since the amplitude of the analog output of a PIR sensor is proportional to the temperature difference between the finger and the surrounding environment, it is interesting to test our system in various environmental temperatures, such as in extremely hot or cold days. It is also interesting to validate whether a model trained in one temperature condition (e.g., hand and environment) works in a very different temperature condition. Additionally, Pyro's tracking accuracy can also be evaluated when the user is on-the-move (e.g., walking or running). A potential research direction is to reduce the impact of physical activities on sensor data. Future research can focus on studying how much the signal can be affected by a shaking wrist when walking or running. The result can help us design and validate solutions to filter out the motion noise from gesture signals.

Cross-user model. Our study shows that people may perform the same gesture in different ways. This means that a model needs to be trained for each user in order to make use of all six tested gestures. In our future work, we will seek to enhance our machine learning model to better deal with user diversity. We will also explore additional thumb-tip gestures and examine gesture parameters that vary across users. Future research could focus on exploring alternative micro gestures and understand the parameters, in which gestures from different users may vary. Signal variance may also appear between users with and without long fingernails. Future research will help us identify and extract ad-hoc features to improve the cross-user accuracy.

Customizing PIR sensor. In this work, we used an off-the-shelf PIR sensor with a pre-configured Fresnel lens. An interesting research direction is to customize the inner configurations of a PIR sensor for detecting micro thumb-tip gestures. Future work will include building a PIR sensor from scratch, so that we can test different crystal alignments and electronic designs. We also plan to test Fresnel lenses with different focal lengths to optimize sensing performance. Conversely, it will be also interesting to test more off-the-shelf infrared sensors (e.g., thermopile and quantum-type infrared sensors).

Power. We examined the power consumption of our current prototype. Overall, our sensing board consumes 148.1 mW, excluding the Bluetooth radio (99 mW) used to transfer PIR data to an external laptop for feature extraction and gesture classification. The sensing component (PIR sensor and its analog frontend) alone consumes 2.6 mW.

The current power number is dominated by the Teensy framework. In particular, the micro-controller [3] in the framework is the most power-consuming, as it contains two ADC components each operating at a 20-KHz sampling rate at a minimum. Given that Pyro requires only 20-Hz sampling, the system can consume significantly less power by using low-power ADC (for example, the ADS7042 from Texas Instruments supports 1 kHz sampling rate with less than 1 microwatt). Furthermore, our feature extraction and

gesture classification algorithm are lightweight. Thus, it holds the potential to be run on lower-power micro-controllers. Future research will explore porting these components to micro-controllers to make the system stand-alone and measure the system's total power consumption.

CONCLUSION

In this paper, we demonstrated the feasibility of recognizing micro thumb-tip gestures through sensing changes in thermal infrared signals emitted from our fingers. We developed a self-contained, proof-of-concept prototype in a wearable form factor using off-the-shelf PIR sensor and electronics. We used a Random Forest classifier to recognize six thumb-tip gestures, including triangle, rectangle, circle, question mark, check mark, and finger rub. We evaluated system performance with ten participants, yielding a 93.9% cross-validation accuracy and 84.9% leave-one-session-out accuracy on the six thumb-tip gestures. Additionally, we initially demonstrated our system's robustness against different lighting conditions, hand temperatures, and background motion noises. Our work presents a passive sensing methodology for detecting micro thumb-tip gestures. We believe it holds the potential to be applied in a wide range of wearable and mobile devices.

REFERENCES

1. AFB. Facts and Figures on American Adults with Vision Loss. 2017. Retrieved July 15, 2017 from <http://www.afb.org/info/blindness-statistics/adults/facts-and-figures/235>
2. Aria Wearable. 2017. Retrieved July 15, 2017 from <http://www.flicktek.com/>
3. Freescale Microcontroller Datasheet. 2017. Retrieved July 4, 2017 from http://cache.freescale.com/files/32bit/doc/data_sheet/K20P64M72SF1.pdf
4. JAWS. Screen reader from Freedom Scientific. 2017. Retrieved July 5, 2017 from <http://www.freedomscientific.com/Products/Blindness/JAWS>
5. PJRC Electronic Projects - Teensy 3.2 & 3.1 s. 2017. Retrieved July 5, 2017 from <https://www.pjrc.com/teensy/teensy31.html>
6. Texas Instruments. 2017. Retrieved July 4, 2017 from <http://www.ti.com/lit/ds/snosc16d/snosc16d.pdf>
7. tsfresh Toolkit. 2017. Retrieved July 4, 2017 from <https://tsfresh.readthedocs.io/en/latest/text/introduction.html>
8. VoiceOver, Screen reader from Apple. 2017. Retrieved July 4, 2017 from <https://www.apple.com/accessibility/mac/vision/>
9. Window-Eyes, Screen Reader GW Micro. Window-Eyes, Screen Reader GW Micro. 2017
10. Ali Abdolrahmani, Ravi Kuber and Amy Hurst. 2016. An empirical investigation of the situationally-induced impairments experienced by blind mobile device users. In *Proceedings of the 13th Web for All Conference (W4A '16)*, Article 21, 8 pages.

- DOI=<https://doi.org/10.1145/2899475.2899482>
11. Troy Allman, Rupinder K Dhillon, Molly AE Landau and Sri H Kurniawan. 2009. Rock Vibe: Rock Band® computer games for people with no or limited vision. In *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '09), 51-58.
 12. Brian Amento, Will Hill and Loren Terveen. 2002. The sound of one hand: a wrist-mounted bio-acoustic fingertip gesture interface. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '02), 724-725.
DOI=<http://dx.doi.org/10.1145/506443.506566>
 13. Chieko Asakawa, Hironobu Takagi, Shuichi Ino and Tohru Ifukube. 2002. Auditory and tactile interfaces for representing the visual effects on the web. In *Proceedings of the Fifth International ACM SIGACCESS Conference on Assistive Technologies* (ASSETS '02), 65-72.
DOI=<http://dx.doi.org/10.1145/638249.638263>
 14. Shiri Azenkot and Emily Fortuna. 2010. Improving public transit usability for blind and deaf-blind people by connecting a braille display to a smartphone. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '10), 317-318.
DOI=<http://dx.doi.org/10.1145/1878803.1878890>
 15. Shiri Azenkot, Richard E Ladner and Jacob O Wobbrock. 2011. Smartphone haptic feedback for nonvisual wayfinding. In *proceedings of the 13th international ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '11), 281-282.
DOI=<http://dx.doi.org/10.1145/2049536.2049607>
 16. Shiri Azenkot and Nicole B Lee. 2013. Exploring the use of speech input by blind people on mobile devices. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '13), Article 11, 8 pages.
DOI=<http://dx.doi.org/10.1145/2513383.2513440>
 17. Stephen Brewster and Lorna M Brown. 2004. Tactons: structured tactile messages for non-visual information display. In *Proceedings of the Fifth Conference on Australasian User Interface-Volume 28* (AUIC '04), 15-23.
 18. Liwei Chan, Yi-Ling Chen, Chi-Hao Hsieh, Rong-Hao Liang and Bing-Yu Chen. 2015. CyclopsRing: Enabling Whole-Hand and Context-Aware Interactions Through a Fisheye Ring. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '15), 549-556.
DOI=<https://doi.org/10.1145/2807442.2807450>
 19. Ke-Yu Chen, Kent Lyons, Sean White and Shwetak Patel. 2013. uTrack: 3D input using two magnetic sensors. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), 237-244.
DOI=<http://dx.doi.org/10.1145/2501988.2502035>
 20. Xiang'Anthony' Chen, Julia Schwarz, Chris Harrison, Jennifer Mankoff and Scott E Hudson. 2014. Air+ touch: interweaving touch & in-air gestures. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), 519-525.
DOI=<https://doi.org/10.1145/2642918.2647392>
 21. Andrea Colaço, Ahmed Kirmani, Hye Soo Yang, Nan-Wei Gong, Chris Schmandt and Vivek K Goyal. 2013. Mime: compact, low power 3D gesture sensing for interaction with head mounted displays. In *Proceedings of the 26th Annual ACM Symposium on User interface Software and Technology* (UIST '13), 227-236.
DOI=<http://dx.doi.org/10.1145/2501988.2502042>
 22. Artem Dementyev and Joseph A. Paradiso. 2014. WristFlex: low-power gesture input with wrist-worn pressure sensors. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), 161-166.
DOI=<https://doi.org/10.1145/2642918.2647396>
 23. Yasmine N El-Glaly, Francis Quek, Tonya Smith-Jackson and Gurjot Dhillon. 2013. Touch-screens are not tangible: Fusing tangible interaction with touch glass in readers for the blind. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction* (TEI '13), 245-253.
DOI=<http://dx.doi.org/10.1145/2460625.2460665>
 24. Rui Fukui, Masahiko Watanabe, Tomoaki Gyota, Masamichi Shimosaka and Tomomasa Sato. 2011. Hand shape classification with a wrist contour sensor: development of a prototype device. In *Proceedings of the 13th International Conference on Ubiquitous Computing* (UbiComp '11), 311-314.
DOI=<https://doi.org/10.1145/2030112.2030154>
 25. Jun Gong, Lan Li, Daniel Vogel and Xing-Dong Yang. 2017. Cito: An Actuated Smartwatch for Extended Interactions. In *Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems* (CHI '17), 5331-5345.
DOI=<https://doi.org/10.1145/3025453.3025568>
 26. Timo Götzelmann. 2016. LucentMaps: 3D Printed Audiovisual Tactile Maps for Blind and Visually Impaired People. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '16), 81-90.
DOI=<https://doi.org/10.1145/2982142.2982163>
 27. Sidhant Gupta, Tim Campbell, Jeffrey R Hightower and Shwetak N Patel. 2010. SqueezeBlock: using virtual springs in mobile devices for eyes-free interaction. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (UIST '10), 101-104.
DOI=<https://doi.org/10.1145/1866029.1866046>
 28. Edward Twitchell Hall. 1966. The hidden dimension.
 29. Chris Harrison, Desney S. Tan and Dan Morris. 2010. Skinput: appropriating the body as an input surface. In

- 28th Annual ACM Conference on Human Factors in Computing Systems (CHI '10), 453-462.
DOI=https://doi.org/10.1145/1753326.1753394
30. Victoria E Hribar, Laura G Deal and Dianne TV Pawluk. 2012. Displaying braille and graphics with a tactile mouse. In *Proceedings of the 14th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '11), 251-252.
DOI=http://dx.doi.org/10.1145/2049536.2049584
 31. Rabia Jafri, Asmaa Mohammed Aljuhani and Syed Abid Ali. 2015. A tangible interface-based application for teaching tactual shape perception and spatial awareness sub-concepts to visually impaired children. *Procedia Manufacturing*, 3, 5562-5569.
 32. Shaun K Kane, Meredith Ringel Morris and Jacob O Wobbrock. 2013. Touchplates: low-cost tactile overlays for visually impaired touch screen users. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '13), 22.
DOI=http://dx.doi.org/10.1145/2513383.2513442
 33. David Kim, Otmar Hilliges, Shahram Izadi, Alex D. Butler, Jiawen Chen, Iason Oikonomidis and Patrick Olivier. 2012. Digits: freehand 3D interactions anywhere using a wrist-worn gloveless sensor. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (UIST '12), 167-176.
DOI=https://doi.org/10.1145/2380116.2380139
 34. Jaime Lien, Nicholas Gillian, M. Emre Karagozler, Patrick Amihood, Carsten Schwesig, Erik Olson, Hakim Raja and Ivan Poupyrev. 2016. Soli: Ubiquitous Gesture Sensing with Millimeter Wave Radar. In *Proceedings of ACM SIGGRAPH (SIGGRAPH '16)*, 142.
DOI=https://doi.org/10.1145/2897824.2925953
 35. Christian Loclair, Sean Gustafson and Patrick Baudisch. 2010. PinchWatch: a wearable device for one-handed microinteractions. In *proceedings of the 12th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '10).
DOI=https://doi.org/10.1.1.453.6052
 36. Muhanad S Manshad, Enrico Pontelli and Shakir J Manshad. 2011. MICOO (multimodal interactive cubes for object orientation): a tangible user interface for the blind and visually impaired. In *proceedings of the 13th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '11), 261-262.
DOI=http://dx.doi.org/10.1145/2049536.2049597
 37. David McGookin, Euan Robertson and Stephen Brewster. 2010. Clutching at straws: using tangible interaction to provide non-visual access to graphs. In *28th Annual ACM Conference on Human Factors in Computing Systems* (CHI '10), 1715-1724.
DOI=https://doi.org/10.1145/1753326.1753583
 38. Tony Morelli, John Foley and Eelke Folmer. 2010. Vi-bowling: a tactile spatial exergame for individuals with visual impairments. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '10), 179-186.
DOI=http://dx.doi.org/10.1145/1878803.1878836
 39. Rajalakshmi Nandakumar, Vikram Iyer, Desney Tan and Shyamnath Gollakota. 2016. Fingerio: Using active sonar for fine-grained finger tracking. In *34th Annual ACM Conference on Human Factors in Computing Systems* (CHI '16), 1515-1525.
DOI=https://doi.org/10.1145/2858036.2858580
 40. Santiago Ortega-Avila, Bogdana Rakova, Sajid Sadi and Pranav Mistry. 2015. Non-invasive optical detection of hand gestures. In *Proceedings of the 6th Augmented Human International Conference* (AH '15), 179-180.
DOI=http://dx.doi.org/10.1145/2735711.2735801
 41. Denise Prescher, Oliver Nadig and Gerhard Weber. 2010. Reading braille and tactile ink-print on a planar tactile display. *Computers Helping People with Special Needs*. 482-489.
 42. Denise Prescher, Gerhard Weber and Martin Spindler. 2010. A tactile windowing system for blind users. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '10), 91-98.
DOI=http://dx.doi.org/10.1145/1878803.1878821
 43. Jun Rekimoto. 2001. GestureWrist and GesturePad: Unobtrusive Wearable Interaction Devices. In *Proceedings of the 5th IEEE International Symposium on Wearable Computers*, 21.
DOI=10.1109/ISWC.2001.962092
 44. T. Scott Saponas, Desney S. Tan, Dan Morris, Ravin Balakrishnan, Jim Turner and James A. Landay. 2009. Enabling always-available input with muscle-computer interfaces. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology* (UIST '09), 167-176.
DOI=https://doi.org/10.1145/1622176.1622208
 45. Toby Sharp, Cem Keskin, Duncan Robertson, Jonathan Taylor, Jamie Shotton, David Kim, Christoph Rhemann, Ido Leichter, Alon Vinnikov, Yichen Wei, Daniel Freedman, Pushmeet Kohli, Eyal Krupka, Andrew Fitzgibbon and Shahram Izadi. 2015. Accurate, Robust, and Flexible Real-time Hand Tracking. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 3633-3642.
DOI=https://doi.org/10.1145/2702123.2702179
 46. Jie Song, Gabor Soros, Fabrizio Pece, Sean Ryan Fanello, Shahram Izadi, Cem Keskin and Otmar Hilliges. 2014. In-air gestures around unmodified mobile devices. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), 319-329.
DOI=https://doi.org/10.1145/2642918.2647373

47. Andrii Soviak, Anatoliy Borodin, Vikas Ashok, Yevgen Borodin, Yury Puzis and IV Ramakrishnan. 2016. Tactile Accessibility: Does Anyone Need a Haptic Glove? In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '16)*, 101-109. DOI=<https://doi.org/10.1145/2982142.2982175>
48. Thad Starner. 2001. The challenges of wearable computing: Part 1. *Ieee Micro*, 21 (4). 44-52.
49. Burkay Sucu and Eelke Folmer. 2014. The blind driver challenge: steering using haptic cues. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '14)*, 3-10. DOI=<http://dx.doi.org/10.1145/2661334.2661357>
50. Li Sun, Souvik Sen, Dimitrios Koutsonikolas and Kyu-Han Kim. 2015. WiDraw: Enabling Hands-free Drawing in the Air on Commodity WiFi Devices. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom '15)*, 7-89. DOI=<http://dx.doi.org/10.1145/2789168.2790129>
51. Shuai Tao, Mineichi Kudo, Hidetoshi Nonaka and Jun Toyama. 2011. Person authentication and activities analysis in an office environment using a sensor network. In *International Joint Conference on Ambient Intelligence*, Springer, 119-127.
52. Christiane Taras and Thomas Ertl. 2009. Interaction with Colored Graphical Representations on Braille Devices. *Universal Access in Human-Computer Interaction. Addressing Diversity*. 164-173. DOI= https://doi.org/10.1007/978-3-642-31479-7_19
53. Marlon Twyman, Joe Mullenbach, Craig Shultz, J Edward Colgate and Anne Marie Piper. 2015. Designing wearable haptic information displays for people with vision impairments. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '15)*. DOI=<https://doi.org/10.1145/2677199.2680578>
54. Fernando Vidal-Verdú and Moustapha Hafez. 2007. Graphical tactile displays for visually-impaired people. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 15 (1). 119-130. DOI=10.1109/TNSRE.2007.891375
55. Anita Vogl, Patrick Parzer, Teo Babic, Joanne Leong, Alex Olwal and Michael Haller. 2017. StretchEBand: Enabling Fabric-based Interactions through Rapid Fabrication of Textile Stretch Sensors. In *Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems (CHI '17)*, 2617-2627. DOI=<https://doi.org/10.1145/3025453.3025938>
56. Thorsten Völkel, Gerhard Weber and Ulrich Baumann. 2008. Tactile graphics revised: the novel brailledis 9000 pin-matrix device with multitouch input. In *Proceedings of the 11th International Conference on Computers Helping People with Special Needs (ICCHP '08)*. 835-842. DOI=http://dx.doi.org/10.1007/978-3-540-70540-6_124
57. Robert Y. Wang, Sylvain Paris and Jovan Popovic. 2011. 6D hands: markerless hand-tracking for computer aided design. In *Proceedings of the 24th annual ACM Symposium on User Interface Software and Technology (UIST '11)*, 549-558. DOI=<https://doi.org/10.1145/2047196.2047269>
58. Wei Wang, Alex X. Liu and Ke Sun. 2016. Device-free gesture tracking using acoustic signals. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking (MobiCom '16)*. DOI=<https://doi.org/10.1145/2973750.2987385>
59. Jacob O. Wobbrock, Meredith Ringel Morris and Andrew D. Wilson. 2009. User-defined gestures for surface computing. In *Proceedings of the 27th Annual ACM Conference on Human Factors in Computing Systems (CHI '09)*, 1083-1092. DOI=<https://doi.org/10.1145/1518701.1518866>
60. Jacob Wobbrock, Andy Wilson and Yang Li. 2007. Gestures without libraries, toolkits or Training: a \$1.00 Recognizer for User Interface Prototypes. In *Proceedings of the 20th annual ACM Symposium on User Interface Software and Technology (UIST '07)*, 159-168.
61. Piotr Wojtczuk, Alistair Armitage, T David Binnie and Tim Chamberlain. 2012. Recognition of simple gestures using a PIR sensor array. *Sensors and Transducers*, 14 (1). 83.
62. Hanlu Ye, Meethu Malu, Uran Oh and Leah Findlater. 2014. Current and future mobile and wearable device use by people with visual impairments. In *Proceedings of the 32th Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*, 3123-3132. DOI=<https://doi.org/10.1145/2556288.2557085>
63. Bei Yuan and Eelke Folmer. 2008. Blind hero: enabling guitar hero for the visually impaired. In *Proceedings of the 10th international ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '08)*, 169-176. DOI=<http://dx.doi.org/10.1145/1414471.1414503>
64. Ouyang Zhang and Kannan Srinivasan. 2016. Mudra: User-friendly Fine-grained Gesture Recognition using WiFi Signals. In *Proceedings of the 12th International Conference on Emerging Networking Experiments and Technologies (CoNEXT '16)*, 83-96. DOI=<https://doi.org/10.1145/2999572.2999582>
65. Yang Zhang and Chris Harrison. 2015. Tomo: Wearable, Low-Cost Electrical Impedance Tomography for Hand Gesture Recognition. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology (UIST '15)*, 167-173. DOI=<https://doi.org/10.1145/2807442.2807480>