Bodyprint: Biometric User Identification on Mobile Devices Using the Capacitive Touchscreen to Scan Body Parts

Christian Holz, Senaka Buthpitiya, Marius Knaust
Yahoo Labs, Sunnyvale, CA
{christianh, senaka, mariusk} @ yahoo-inc.com

ABSTRACT
Recent mobile phones integrate fingerprint scanners to authenticate users biometrically and replace passwords, making authentication more convenient for users. However, due to their cost, capacitive fingerprint scanners have been limited to top-of-the-line phones, a result of the required resolution and quality of the sensor. We present *Bodyprint*, a biometric authentication system that detects users' biometric features using the same type of capacitive sensing, but uses the touchscreen as the image sensor instead. While the input resolution of a touchscreen is ~6 dpi, the surface area is larger, allowing the touch sensor to scan users' body parts, such as ears, fingers, fists, and palms by pressing them against the display. Bodyprint compensates for the low input resolution with an increased false rejection rate, but does not compromise on authentication precision: In our evaluation with 12 participants, Bodyprint classified body parts with 99.98% accuracy and identifies users with 99.52% accuracy with a false rejection rate of 26.82% to prevent false positives, thereby bringing reliable biometric user authentication to a vast number of commodity devices.

INTRODUCTION
Mobile phones now store a manifold of sensitive user data, such as photos, emails as well as login credentials to access personal data on the web, including finances and shopping portals. To protect such data from theft and unauthorized access, mobile devices implement lock screens to verify the user's identity and authenticate their use of the device.

Lock screens commonly ask the user to enter a PIN to unlock [4,8]. Unfortunately, only a small percentage of mobile phone users actually protect their device using a PIN [13]. A commonly cited reason is that PIN codes impede convenience and ease of access [3]. Graphical passwords ease the memorization and entry of PIN codes, but are subject to eavesdropping and find only limited acceptance [8].

Researchers have thus sought to replace PIN codes to protect mobile devices, such as to identify users from behavioral biometrics [5] or analyzing the characteristics of gestures users perform with the device [12]. Alternative approaches leverage the front-facing cameras in mobile devices to implement face [1] or iris recognition [7].

To balance reliable and strong authentication with the convenience of use, some recent phones have started to integrate capacitive fingerprint scanners as part of the enclosure (e.g., iPhone 5S and up, Samsung Galaxy S5). However, capturing fingerprints requires high-quality sensors, which incurs considerable cost in the manufacturing process. Fingerprint scanners have thus been reserved for top-of-the line mobile devices only. Several research prototypes have further explored touch-based authentication by identifying users during contact with the touchscreen itself [9,10,11].

In this note, we present *Bodyprint*, a biometric authentication system that uses the capacitive sensor of commodity touch*screens* to detect users' biometric features—the same type of sensing that fingerprint scanners use. Whereas fingerprint scanners exist only in a small number of devices, all current smartphones feature capacitive touchscreens, ready to run Bodyprint for authenticating user access.

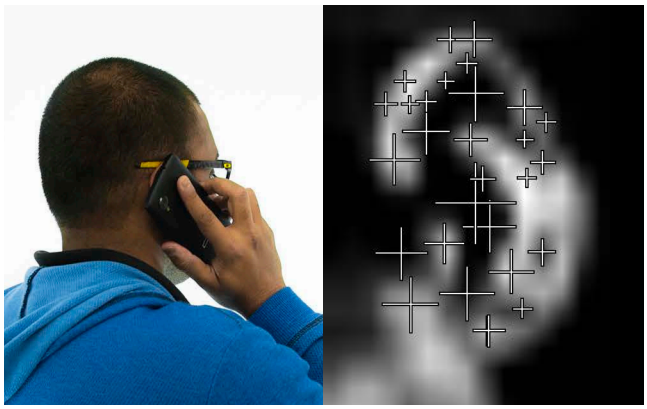

Figure 1: *Bodyprint* brings biometric authentication to commodity mobile devices using the capacitive touchscreen as a low-resolution, but large-area image sensor to reliably identify users based on their ears, fists, or grips when pressed against the touchscreen. (left) To accept an incoming call, the user places the touchscreen onto his ear. (right) Bodyprint extracts features from the raw capacitive image to identify the user.

BODYPRINT: SENSING BIOMETRICS WITH THE SCREEN
Figure 1 shows the use of Bodyprint to accept an incoming call. (a) The user touches the phone to his ear and (b) Bodyprint captures the raw capacitive image from the touch sensor, extracts features and matches them against a user database, and identifies the user to authenticate access.

Bodyprint replaces password entry on mobile devices with biometric authentication. Since the input resolution of

CHI 2015, April 18 – 23 2015, Seoul, Republic of Korea
Copyright 2015 ACM 978-1-4503-3145-6/15/04…$15.00
http://dx.doi.org/10.1145/2702123.2702518

commodity touchscreens at ~6 dpi is two orders of magnitude lower than that of fingerprint scanners, Bodyprint scans body parts with proportionally bigger structural features as shown in Figure 2: ears, fists, phalanges, fingers and palms. Bodyprint appropriates the capacitive touchscreen as an image sensor and thus has the potential to run on most if not all currently available touch-based phones.

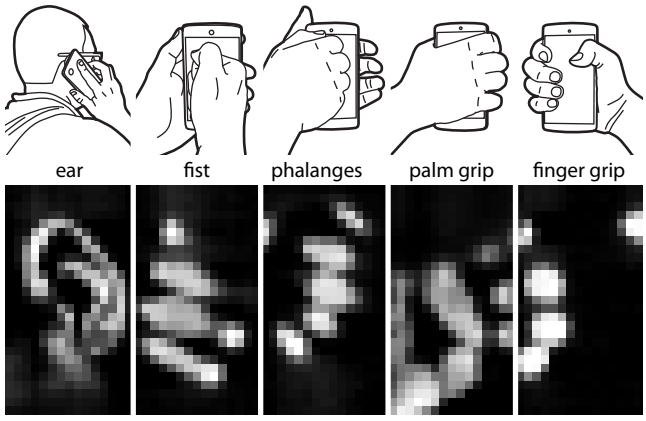

Figure 2: Bodyprint identifies users through these body parts and characteristic poses when touching the screen.

A key benefit of Bodyprint is the usage of the flat touchscreen *surface* to capture an image of the body part, which conveniently solves the challenge of projection for matching images. When pressed against the screen, a body part has a 1–1 mapping to the screen's sensor cells. This mapping remains *constant* over time, such that scanned parts will always exhibit the same dimensions on the device. Using the contact surface for scanning also frees Bodyprint from accommodating image captures from various angles, as is a requirement when using camera-based scanning [1].

CONTRIBUTION

Our main contribution is technical: We demonstrate the feasibility and reliability of the touchscreen as an image sensor with sufficient resolution to capture users' biometric features that are adequate for robust user identification. We require no additional hardware or dedicated sensors, just the capacitive touchscreen, which is the largest component in smartphones today. Our identification algorithm is designed to identify users with high precision and robustness using low-resolution images as input, running in real-time on commodity phones to authenticate access, which brings biometric user identification as a replacement for passwords from high-end devices to *all* current smartphones.

BODYPRINT'S ALGORITHM AND IMPLEMENTATION

We implemented Bodyprint on an LG Nexus 5 phone, which features a Synaptics ClearPad 3350 touch sensor. Through rooting the Android phone and modifying the touchscreen module in the kernel source, we activated Synaptics debug mode and now obtain the capacitive values from the sensor: a 27×15px 8-bit image across a 4.95″ surface (6.24 dpi) at 30fps. Figure 2 visualizes the raw data from the sensor as a "depth" image that captures the proximity of skin to the surface with limited depth resolution.

Bodyprint robustly identifies users despite the low resolution of the input sensor. We accomplish this by processing the sequence of consecutive raw images that result from a *trial*, gathering images between activating the debug mode and detecting touches. Our algorithm has three steps: preprocessing, body part classification, and user identification. Figure 3 shows an overview over our implementation.

Step 1: Preprocessing (raw values to 3×4 key frames)

For each raw image in a trial, Bodyprint first generates three representations using logarithmic, linear and exponential tone mapping. Each mapping accentuates contours and gradients at different intensity levels, enhancing the information in the low-resolution images from the touch sensor.

Depending on the average pixel intensity of an image, the image is sorted into one of five brightness buckets. The purpose of the five buckets is to collect sensor images depending on the body part's proximity to the touchscreen during a trial (i.e., low-intensity images correspond to hovering body parts, high intensities to body parts in contact). We then merge all frames within each bucket into a single image, stretch intensities to the full range, discard the lowest-intensity bucket and scale the remaining four images 16 times using cubic interpolation. This results in 3 tone mappings × 4 buckets = 12 images. Finally, we derive the SURF descriptors [2] of the features resulting from Hessian corner detection, arriving at a set of 12 *key frames* per trial.

During training, Bodyprint inserts each of the 12 key frames into two databases. The user database stores each key frame along with the user's name, body part, and trial ID. The body part database contains 12 groups of features for each body part, corresponding to each of the 12 key frames, thereby combining all features across all trials.

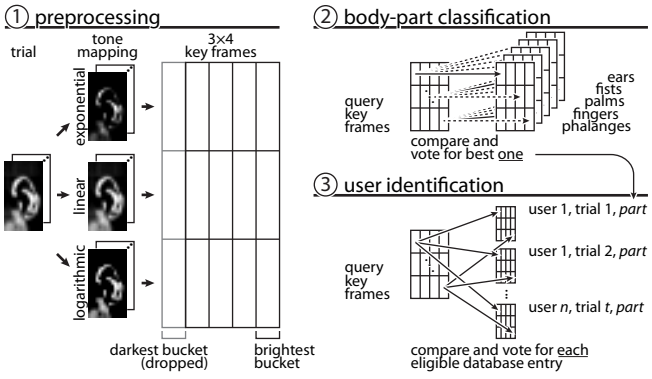

Figure 3: Overview over Bodyprint's algorithm.

Step 2: Body-part classification

The purpose of body part classification is to reduce the search space of the ensuing identification. Although trials may be misclassified as wrong body parts at this stage, we discard false positive matches during user identification.

To classify the body part of a query trial, Bodyprint assigns a body part from the database to each query key frame by comparing it to the respective key frame group in the database. For each body part, we search the database for a best-

matching feature f_j of a SURF feature f_i in the query key frame, then search the query key frame for a best match of f_j using L2 distance, and obtain $f_{i'}$. If the two searches produce the same result, i.e., if $f_i = f_{i'}$, we add (f_i, f_j) to the set of matched features for that key frame. If more than half of the features in a query key frame match, we derive the average L2 distance of the matches of that body part. Otherwise, we reject the key frame from voting for that body part. After repeating this procedure for all body parts and key frames, we assign the body part with the lowest L2 distance to each key frame, resulting in a 3×4 set of voted-on body parts.

Each of the 12 key frames now votes for one body part in the database. Bodyprint classifies the query trial as the body part with the highest number of votes if the ratio of votes and number to all voting key frames is above the threshold "body part vote percentage." Otherwise, Bodyprint discards the entire trial and stops the subsequent user identification.

Step 3: User identification using classified body parts

To identify a user, Bodyprint implements a multi-step voting procedure. Each key frame in the query trial can vote for multiple users in the database (similar to Bootstrapper [14]). We additionally check the transformation between the query trial and database trials to reject false positive matches. We repeat this process for each key frame k_l in the database where tuple l = (user, body part, trial ID), but limit body part to the body part determined in Step 2.

First, we compare each query key frame to each key frame k_l. For each feature f_j in a query key frame, we determine the two features f_{j1} and f_{j2} in k_l with the lowest and the second-lowest L2 distance. If the distance between (f_i, f_{j1}) is smaller than 70% of the distance between (f_i, f_{j2}), we add f_i to the set of good matches [2]. If this results in less than 10 good matches, we prevent this query key frame from voting for the currently-tested combination of (user, trial ID).

Next, we calculate the best rigid 2D transformation to rotate the good matches in the query key frame onto the respective features in the database key frame, but exclude other transformations, such as scaling and shearing. If the average error of the transformation is less than the "rotational error threshold," we count a vote for that user.

By iterating across all (user, trial ID) combinations in the database, each key frame can vote as many times for a user as there are trials for that user's body part. We thus normalize each user's votes by the number of those trials. Finally, we sum the votes per user across all key frames. If the vote winner receives less than a fraction of all votes ("user vote percentage threshold"), we reject the query trial. This is how Bodyprint robustly rejects false positive matches.

The use of the previous three thresholds allows us to balance between two extremes: high security of the system, resulting in a higher false rejection rate, and low false rejection rate, which always produces a match after each trial, but risks lower authentication precision. If a trial is rejected, Bodyprint prompts the user to retry. Upon repeated rejection, Bodyprint resorts to regular password entry.

APPLICATIONS

We implemented two applications that use Bodyprint, which we demonstrate in our video figure. Figure 1 shows our first application, which authenticates the user by their ear for an incoming call. Once the user touches their ear onto the screen, Bodyprint identifies the user, confirms or rejects them for the call, and plays auditory confirmation.

Our second application extends Bodyprint to two people and implements the four-eye principle for locking sensitive documents; accessing them then requires the presence of both people. To lock a document, a user pushes the 'lock' button for an open document. Bodyprint requests that all parties shake the phone, which provides the touchscreen with their palm prints. Bodyprint identifies each user and locks the document using their identities. To unlock a document later, a user needs to provide their palm print to the touchscreen again and Bodyprint will prompt for the other user to also be authenticated to reveal the document.

TECHNICAL EVALUATION

We evaluated Bodyprint's performance for each of the 5 body parts shown in Figure 2. We recruited 12 participants (ages 24–53, 4 female) and demonstrated each of the 5 poses to them once. During the evaluation, participants held the Nexus 5 phone as demonstrated and performed 12 trial repetitions. Between trials, participants put the phone down on a table. We did not verify the correctness of performed poses or ask for retries. Overall, we collected 864 trials.

Evaluation methodology

We tested the performance of our algorithm with 12-fold cross validation, which evaluated two aspects: identifying a known user and rejecting an unknown user who is trying to authenticate. In each fold, all of one participant's trials were withheld from the training data set, such that Bodyprint's body part and user database contained 11 trials for each of the remaining 11 participants. We then tested the trained system with the 11 withheld trials from participants in the databases as well as all 12 trials of the withheld participant. Bodyprint then attempted to correctly identify the user of the system or reject the user when it determined the trial to represent an unknown user. We repeated the full 12-fold cross validation varying the three tunable thresholds explained above in unison linearly through each one's range.

Results

Figure 4 plots Bodyprint's false rejection rate against authentication precision. The false rejection rate represents the likelihood of a user to be prompted to retry. Precision is the probability of a known user being correctly identified and unknown users being rejected. The ROC-like curves illustrate the tradeoff for all threshold settings: A precision of 99.5% leads to a false rejection rate of 26.8% using any of the five body parts. Or, for a false rejection rate of 0.2%, Bodyprint achieves 86.1% precision. Limiting the body parts to the ear only boosts the precision of our algorithm to 99.8% with a false rejection rate of only 7.8%.

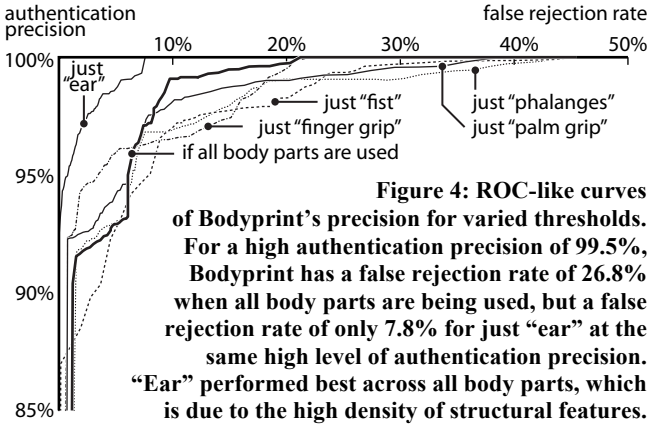

Figure 4: ROC-like curves of Bodyprint's precision for varied thresholds. For a high authentication precision of 99.5%, Bodyprint has a false rejection rate of 26.8% when all body parts are being used, but a false rejection rate of only 7.8% for just "ear" at the same high level of authentication precision. "Ear" performed best across all body parts, which is due to the high density of structural features.

Battery consumption of the touchscreen debug mode

We repeatedly measured the time to fully discharge the Nexus 5 running Bodyprint fulltime with and without Synaptics' debug mode enabled. In both cases, the phone was in airplane mode and a wake lock kept the screen on. In debug mode, the phone turned off after ~15 hours on average compared to ~18 hours on average in regular mode.

Discussion and limitations

The results from our evaluation demonstrate that Bodyprint can be tuned for very high identification precision (99.6%) for applications that require secure authentication, such as accessing sensitive data. In this case, Bodyprint achieves this high precision by strictly rejecting potential false positives, which naturally leads to increased false rejection rate for known users. Alternatively, Bodyprint's thresholds can be set to facilitate quick unlocking, such that each trial will be matched (low false rejection rate), albeit at reduced precision (86.1%). To achieve the results shown in Figure 4, applications that have ~10 known users require a minimum of 11 training samples per user and body part.

When run fulltime, Bodyprint accounts for a 17% overhead on the battery of the device. In reality, however, Bodyprint activates the touch sensor's debug mode only when a user is authenticating, which takes only a few seconds.

CONCLUSIONS

We presented Bodyprint, a password replacement for commodity mobile devices to identify users biometrically from the features of their body, thus possibly increasing the convenience of logging in and authenticating. Users thereby press a body part against the large touchscreen of the device, from which Bodyprint obtains the raw capacitive image, extracts features, and identifies body parts and users.

Bodyprint appropriates the capacitive touchscreen of mobile devices as an image sensor and accommodates their low input resolutions through an increased false rejection rate, but not reduced precision or increased false positives. Bodyprint identified users with 99.5% precision in an evaluation with 12 participants with a false rejection rate of 26.8% across all body parts, but as low as 7.8% for ear-only matching, which is explained through the increased presence of structural features in the human ear [6].

Since all current touchscreens use capacitive sensing, Bodyprint brings reliable biometric authentication to a vast number of commodity devices. In the case that future touchscreens support higher input resolutions, up to a point where they may detect the fine structure of fingerprints, Bodyprint will readily incorporate the higher level of detail of sensor data, which will not only extend our approach to further body parts, but likely reduce false rejection rates at the same high levels of authentication precision.

ACKNOWLEDGEMENTS

We thank Dan Odell, Scott Lin and Mitesh Patel for feedback.

REFERENCES

1. Abate, A.F., Nappi, M., Riccio, D., Sabatino, G. 2D and 3D face recognition. *Pattern Recognition Letters (28):14*, 2007.
2. Bay, H., Ess, A., Tuytelaars, T., Van Gool, L. Speeded-up Robust Features. *Proc. CVIU '07,* 346–359.
3. Ben-Asher, N., Kirschnick, N., Sieger, H., Meyer, J., Ben-Oved, A., Möller, S. On the need for different security methods on mobile phones. *Proc. MobileHCI '11*.
4. Biddle, R., Chiasson, S., Van Oorschot, P. Graphical passwords: Learning from the first twelve years. *ACM Computing Surveys*, (44)4, 2012.
5. Bo, C., Zhang, L., Li, X., Huang, Q., Wang, Y. SilentSense: silent user identification via touch and movement behavioral biometrics. *Proc. MobiCom '13*.
6. Burge, M. and Burger, W. Ear biometrics. *Biometrics* 273–285, Springer US, 1996.
7. Cho, D., Park, K.R., Rhee, D.W., Kim, Y., Yang, J. Pupil and Iris Localization for Iris Recognition in Mobile Phones. *Proc. SNPD '06*, 197–201.
8. Everitt, K., Bragin, T., Fogarty, J., Kohno, T. A comprehensive study of frequency, interference, and training of multiple graphical passwords. *Proc. CHI '09*.
9. Harrison, C., Sato, M., Poupyrev, I. Capacitive Fingerprinting: Exploring User Differentiation by Sensing Electrical Properties of the Human Body. *Proc. UIST '12*.
10. Holz, C. and Baudisch, P. Fiberio: A Touchscreen that Senses Fingerprints. *Proc. UIST '13*, 41–50.
11. Holz, C. and Baudisch, P. The Generalized Perceived Input Point Model and How to Double Touch Accuracy by Extracting Fingerprints. *Proc. CHI '10*, 581–590.
12. Kirschnick, N., Kratz, S. and Moller, S. An improved approach to gesture-based authentication for mobile devices. *Proc. SOUPS '10*.
13. Kurkovsky, S. and Syta, E. Digital natives and mobile phones: A survey of practices and attitudes about privacy and security. *Proc. ISTAS '10,* 441–449.
14. Richter, S., Holz, C., Baudisch, P. Bootstrapper: Recognizing Tabletop Users by their Shoes. *Proc CHI '12*.