EchoPrint: Two-factor Authentication using Acoustics and Vision on Smartphones

Bing Zhou

Stony Brook University
Department of Electrical and Computer Engineering
bing.zhou@stonybrook.edu

Ruipeng Gao

Beijing Jiaotong University
School of Software Engineering
rpgao@bjtu.edu.cn

Jay Lohokare

Stony Brook University
Department of Computer Science
jay.lohokare@stonybrook.edu

Fan Ye

Stony Brook University
Department of Electrical and Computer Engineering
fan.ye@stonybrook.edu

ABSTRACT

User authentication on smartphones must satisfy both security and convenience, an inherently difficult balancing act. Apple's FaceID is arguably the latest of such efforts, at the cost of additional hardware (e.g., dot projector, flood illuminator and infrared camera). We propose a novel user authentication system *EchoPrint*, which leverages acoustics and vision for secure and convenient user authentication, without requiring any special hardware. *EchoPrint* actively emits almost inaudible acoustic signals from the earpiece speaker to "illuminate" the user's face and authenticates the user by the unique features extracted from the echoes bouncing off the 3D facial contour. To combat changes in phone-holding poses thus echoes, a Convolutional Neural Network (CNN) is trained to extract reliable acoustic features, which are further combined with visual facial landmark locations to feed a binary Support Vector Machine (SVM) classifier for final authentication. Because the echo features depend on 3D facial geometries, *EchoPrint* is not easily spoofed by images or videos like 2D visual face recognition systems. It needs only commodity hardware, thus avoiding the extra costs of special sensors in solutions like FaceID. Experiments with 62 volunteers and non-human objects such as images, photos, and sculptures show that *EchoPrint* achieves 93.75% balanced accuracy and 93.50% F-score, while the average precision

is 98.05%, and no image/video based attack is observed to succeed in spoofing.

CCS CONCEPTS

• **Security and privacy** → **Multi-factor authentication**;
• **Human-centered computing** → **Mobile computing**;
Smartphones; • **Computing methodologies** → *Supervised learning by classification*;

KEYWORDS

mobile sensing; acoustics; authentication

ACM Reference Format:

Bing Zhou, Jay Lohokare, Ruipeng Gao, and Fan Ye. 2018. EchoPrint: Two-factor Authentication using Acoustics and Vision on Smartphones. In *The 24th Annual International Conference on Mobile Computing and Networking (MobiCom '18), October 29–November 2, 2018, New Delhi, India*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3241539.3241575>

1 INTRODUCTION

User authentication on smartphones is pivotal to many important daily apps, such as social networks, shopping and banking [11, 21]. Central to the user authentication is the balancing act between security and convenience. A solution must be secure while easy to use. A series of efforts have been undertaken to address this problem.

The most basic and traditional method, PIN number, has both usability (e.g., the pass code forgotten by the user) and security (e.g., shoulder-surfing [46] attacks) issues. Face recognition based authentication can be easily spoofed by images or videos of the user [20]. Simple twists such as requiring eye blinks are vulnerable to video attacks [4]. Iris scan [17] is probably the most secure way, however it requires special sensors unavailable on most mobile devices.

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom '18, October 29–November 2, 2018, New Delhi, India

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5903-0/18/10...\$15.00

<https://doi.org/10.1145/3241539.3241575>

Figure 1: *EchoPrint* emits nearly inaudible sound signals from the earpiece speaker to “illuminate” the user’s face. The extracted acoustic features from the echoes are combined with visual facial landmarks from the frontal camera to authenticate the user.

Fingerprint sensors [3, 44], while convenient for authentication, are facing the challenge posed by the trend of ever-increasing screen size, which leaves little space for fingerprint sensors.

The latest effort, Apple’s FaceID [7], packs a dot projector, a flood illuminator and an infrared depth sensor in a small area to sense the 3D shape of the face, thus achieving high security while saving space. However, the special sensors still take precious frontal space and cost extra ($\sim 5\%$ of its bill of materials) [2]. Thus, we ask this question: is an alternative using existing sensors possible?

In this paper, we propose a novel user authentication system *EchoPrint*, which leverages existing earpiece speaker and frontal camera thus can be readily deployed on most phones. It does not require costly special sensors (e.g., depth or iris) that take more spaces. *EchoPrint* combines acoustic features from a customized CNN feature extractor and facial landmark features from vision algorithms as the joint feature description of the user. It does not require the user to remember any passcode, thus avoiding the usability issues as PIN numbers. The acoustic features depend on 3D facial geometries, thus it is resilient to image/video attacks easily spoofing 2D visual based approaches. Similar to FaceID, it does not require direct touch from the user, thus avoiding issues like wet fingers that pose difficulties to fingerprint sensors.

To achieve resilient, secure and easy-to-use authentication using acoustic and vision, we must address several challenges: i) echo signals are highly sensitive to the relative position between the user’s face and the device (i.e., pose), which makes it extremely hard to extract reliable pose-insensitive features for robust authentication; ii) smartphones come with multiple speakers and microphones - which ones are most suitable, and what are the proper sound signals, are critical to authentication performance; iii) sophisticated signal

processing, feature extraction and machine learning techniques are needed for fast user registration and real-time authentication.

We make the following contributions in this work:

- We design acoustic emitting signal suitable for considerations including hardware limitation, sensing resolution, and audibility to humans. We also create acoustic signal processing techniques for reliable segmentation of echoes from the face.
- We propose an end-to-end hybrid machine learning framework, which extracts representative acoustic features using a convolutional neural network, and fuses vision and acoustic features to SVM for final authentication.
- We design a data augmentation scheme for generating “synthesized” training samples, which reduces false negatives significantly with limited training sample size, thus saving the user efforts in new profile registration.
- We propose three authentication modes, ultra low-power presence detection, low-power continuous authentication and two-factor one-pass authentication for different application scenarios, balancing the trade-offs between security, user convenience and power consumption.
- We build a prototype, conduct extensive experiments and find that *EchoPrint* achieves 93.75% balanced accuracy and 93.50% F-score, while the precision is up to 98.05%. No image/video based attack is observed to succeed in spoofing our system.

To the best of our knowledge, *EchoPrint* is the first to leverage active acoustic sensing combined with vision features for smartphone user authentication, demonstrating robust performance without requiring any additional special sensor.

2 BACKGROUND

2.1 Attack Scenarios

We summarize typical attack scenarios for main existing authentication methods.

Information Leak. Critical PIN numbers or passwords can be leaked or stolen easily. E.g., in shoulder-surfing attacks [46], it is not uncommon for someone standing close by to peek the whole PIN typing.

Replay Attacks. 2D image based face recognition systems suffer from replay attacks by images or videos of the user face. The face recognition system on Samsung’s flagship Galaxy S8 is reported to be spoofed by a simple picture [6].

Biometric Duplication. Fingerprint is the mainstream biometric used for authentication solutions. However, fingerprints are widely left on objects (e.g., glasses) touched by

the user, and can be duplicated with reasonable efforts and skill [3] to fool the sensor.

2.2 Design Considerations

We make the following considerations when designing *EchoPrint*.

Universal. We prefer to use existing hardware widely available on most smartphones, so that it can be deployed at large scale rapidly with minimum hardware costs. Besides, we want to use a biometric that is pervasive to every human being.

Unique. The human face has been widely used as a biometric because it is distinctive. However, most existing 2D visual based systems can be spoofed by images or videos. Thus we leverage the 3D information of the facial contour for much higher security.

Persistent. The biometric must not change much over time. Biometrics such as heart beat, breathing, gait are highly affected by the user's physical conditions (e.g., running vs. walking), thus not optimal choices for robust authentication. In contrast, the human face geometries are not likely to change too much over short time periods. However, daily changes like wearing hats or glasses must be easily accommodated.

Difficult to Circumvent. This is essential for any authentication system to ensure a high security level. Existing authentication approaches, such as PIN numbers, 2D based face recognition, fingerprint sensors still have quite some risks to be circumvented. Because our two-factor authentication examines both acoustic and visual features simultaneously, circumventing would require duplicating both 3D facial geometries and acoustic reflections properties close enough to the human face, which will be much more difficult than needed in circumventing other methods.

2.3 Design Goal

Based on the above discussions, our goal is to build a highly secure, resilient two-factor authentication system that is available to most existing smartphones without requiring any special sensors. To suit different application scenarios, we design different authentication modes that have proper trade-offs between security, convenience and power consumption. While light weight vision algorithms should be used to minimize computation complexity, we keep the option open for future integration with state-of-the-art vision solutions. We believe such "free" acoustic-aided authentication will play an important role in mobile authentication developments.

3 OVERVIEW

EchoPrint uses speakers/microphones for acoustic sensing and the frontal camera for facial landmarks detection. It

Figure 2: *EchoPrint* uses a pre-trained CNN to extract acoustic features, which are combined with visual facial landmarks as joint features. An SVM classifier is trained to authenticate a registered user.

extracts acoustic features from echo signals using a deep learning approach and fuses such features with facial landmarks as a joint representation for authentication. Figure 2 shows the overview of the system design, which consists of two major phases: user *registration* and user *authentication*.

In the registration phase, *EchoPrint* detects facial landmarks (e.g., eyes, mouth) using the frontal camera. Meanwhile, the earpiece speaker emits designed acoustic signals to "illuminate" the user's face. Echoes bouncing back are received by the microphone. A pre-trained CNN model is used to extract acoustic features resilient to phone pose changes, which are combined with facial landmarks as joint feature representation, and fed into an SVM classifier for model training.

In the authentication phase, the user just needs to hold the smartphone in front of the face for facial landmarks detection and acoustic sensing. The joint features are extracted and fed into the trained SVM classifier for final authentication. Optionally, an existing image-based face recognition system can be integrated in our system for pre-screening.

4 ACOUSTIC SENSING

Acoustic echoes from the human face are highly distinctive: i) the echoes are very sensitive to the relative position between the user face and device. ii) each 3D facial contour is a unique set of multiple reflecting surfaces [9], which create a unique sum of individual echoes. iii) different materials absorb, attenuate sound waves differently [28], allowing us to distinguish objects of similar geometry but different materials (e.g., a stone sculpture).

4.1 Speaker/Microphone Selection

Figure 3 shows the speakers, microphones and camera layout on a typical smartphone. There are two speakers, a main one at the bottom and an earpiece speaker at the top for making phone calls. There are also one microphone at the bottom, and another at the top for noise cancellation.

Figure 3: Typical layout of speakers, microphones and camera on smartphones.

We select the earpiece speaker, top microphone, and frontal camera combination for robust acoustic/visual sensing. Earpiece speaker is chosen for sound emitting for two reasons: i) it's a highly standard design on almost all existing smartphones. Its location is suitable for "illuminating" the user's face; whereas the main speaker has a more diverse design, either located at the bottom or on the back; ii) earpiece speaker is close to frontal camera, which minimizes alignment errors when the frontal camera is used for adjusting the phone pose.

Top microphone is chosen as the receiver because it is close to earpiece speaker, and it's less affected by the user's hand holding the device. As shown in Figure 3, the hand is close to the bottom microphone. Even slight hand movements can create noises and variations in sound signals received by the bottom microphone.

4.2 Acoustic Signal Design

There are several considerations in the emitting signal design. First, it should facilitate isolation of the segment of interest (i.e., echoes from the face) from the other reflections, such as interferences from clutters and self-interference from the speaker. This requires the signal be short enough so that echoes from objects at different distances have little overlap in time domain. Second, the acoustic signal should be as inaudible as possible to human ears to minimize annoyance. An ideal frequency range should be over 20KHz. Lastly, the designed signal frequency range should be apart from ambient noises (usually under 8KHz), to enable noise removal (e.g., using band-pass filters) and improve robustness.

According to our survey, a comfortable distance from human eyes to the phone is 25 - 50cm, corresponding to a time delay of $\sim 1.4 - 2.8ms$ at the speed of sound. From our experiments, when the frequency goes above 20KHz, serious power attenuation and worse signal to noise ratio occur, thus echoes from faces are buried under noises. Considering all these facts, we choose a pulse signal with a length of 1ms with linear increasing frequencies from 16 - 22KHz. A Hanning window [24] is applied to reshape the pulse envelop to increase its peak to side lobe ratio, thus producing higher

Figure 4: Sample recording segment of a received signal after noise removal.

SNR for echoes. For authentication modes that require continuous sound emitting, we leave a delay of 50ms for each pulse such that echoes from two consecutive pulses do not overlap.

4.3 Acoustic Signal Pre-processing

4.3.1 Background Noise Removal. The received raw signal goes through a 16 - 22KHz Butterworth band-pass filter to remove background noises, such that weak echoes from human faces will not be buried in the noise.

A sample recording segment of a received signal after noise removal is shown in Figure 4. The *direct path* segment is the emitting signal traveling from speaker to the microphone directly, which ideally should be a copy of the emitting signal and has the highest amplitude. The *major echo* corresponds to the mix of echoes from the major surfaces (e.g., cheek, forehead) of the face. Other surfaces of the face (e.g., nose, chin) at different distances to the phone also produce echoes, arriving earlier/later than the major echo. The *face region echoes* include all these echoes, capturing the full information of the face. Accurate segmenting the face region echoes is critical to minimize the disturbances from dynamic clutters around the phone, and reduce the data dimension for model training and performance.

4.3.2 Signal Segmentation. There are two steps extracting the face region segment: locating the direct path segment in raw recordings, then locating the major echo thus face region segment after the direct path segment.

Locating the Direct Path. An easy but naive assumption is that a constant gap exists between the emitting and recording, thus the direct path can be located after that constant gap. However, both emitting and recording must go through multiple layers of hardware and software processing in the OS, many of which have unpredictable, varying delays. Thus locating the direct path using a constant delay is extremely unreliable.

Figure 5: Distance measurements from acoustics, vision, and calibrated acoustics.

Instead, since the direct path signal usually has the highest amplitude, using cross-correlation to locate it is more reliable [56]. From our experiments, occasional offsets of direct path signal still happen after cross-correlation, due to ambiguities from comparable peak values in the cross-correlation result. We propose two techniques to enhance the stability:

i) *Template Signal Calibration.* Due to the hardware (speaker/microphone) imperfection, the received sound signal is usually slightly different from the designed emitting signal. To get an accurate “template” signal for cross-correlation, we perform emitting and recording in a quiet environment, so that the direct path signal can be reliably detected and saved as a calibrated template for future cross-correlation.

ii) *Signal Fine-tuning.* In addition to the Hanning window, we manually tune the signal slightly to make the key peaks/valleys more prominent, which reduces cross-correlation ambiguity significantly. Only the central portion (15 samples) of the template signal is used in cross-correlation, further enhancing resilience to residual noises.

Locating the Major Echo. A straightforward way for locating the major echo is to find cross-correlation peak location corresponding to typical phone holding distance (e.g., 25 - 50cm) after the direct path location. However, human face echoes can be so weak that echoes from larger obstacles faraway can have comparable amplitudes. This makes the estimation unstable and leads to occasional location “jumping”, thus outliers in distance measurements. The dotted line in Figure 5 shows the distance measurements from acoustic while the device is moving back and forth from the face. We can observe quite some outliers due to such “jumping” outliers. To solve this problem, we propose a vision-aided major echo locating technique of two steps:

i) *Vision Measurement Calibration.* From camera image projection principle, the closer the device to the face, the larger the image and larger distances between facial landmarks, and vice versa. Thus the distance from face to device d_v can be formulated as $d_v = \tau \cdot \frac{1}{d_p}$, where d_p is the distance

Figure 6: Illustration of FMCW.

between two facial landmarks and τ is an unknown scale factor specific to the user. We choose d_p as the pixel distance between two eye landmarks as they are widely separated and can be detected reliably. To estimate the scale factor τ , we calculate τ_i for each pair-wise $d'_{v,i}$ from acoustic distance measurement and $d_{p,i}$ in pixels. To eliminate errors caused by acoustic distance measurement outliers, we first find the major cluster of $\{\tau_i\}$ using density-based spatial clustering algorithm DBSCAN [13], then leverage linear regression to find the best τ that minimizes the offset between d'_v and $\tau \cdot \frac{1}{d_p}$. Figure 5 shows that outliers are removed in vision calibrated acoustic distance measurements.

ii) *Vision-aided Major Echo Locating.* Although vision based distance measurement is more stable than acoustics, it can not capture the error caused by rotations of smartphone or user’s face. Thus the vision calibrated distance measurement is used to narrow down the major echo searching range and reduce outliers. We still use cross-correlation to find the exact major peak location within this range. Note that the user face cannot rotate to extreme angles, otherwise facial landmark detection may fail.

Face Region Echoes. Since the depth of human face is limited, we extend 10 sample points before and after the major echo segment to cover the whole face region (allowing a depth range of $\sim 7cm$), which are later used as inputs for machine models for authentication.

4.4 Segmented Signal Analysis

The face region echoes are a combination of individual echoes with different amplitudes and phases, thus isolating individual echoes in time domain can be very hard due to noises. Instead, we measure the arrival time of each echo by a technique Frequency-Modulated Continuous Wave (FMCW) [18] used in radars. In traditional FMCW, the speaker transmits continuous chirp signals with linear increasing frequency, from f_{min} to f_{max} . To estimate the distance from an object, FMCW compares the frequency of the echo signal to that of a reference signal using a technique called signal mixing,

to find the frequency shift Δf (shown in Figure 6), which is proportional to the distance. Thus finding Δf gives the distance (i.e., Δf multiplying a constant coefficient).

To capture minute surface geometries on the face, the FMCW distance measurement resolution is critical. The resolution in Δf is equal to the size of one bin in the FFT, which depends on the bandwidth used. This is why we use a wide frequency of 16 - 22KHz, though it may be slightly audible to some users. In Figure 6, the FFT is taken over a duration of the face region with length T and hence the size of one FFT bin is $1/T$. Given a minimum measurable frequency shift $\Delta f_{min} = 1/T$, the minimum measurable distance resolution can be computed using the slope of signals (see Figure 6), which is the total swept bandwidth B divided by the sweep time T . Thus the distance resolution:

$$d_r = C \frac{TOF_{min}}{2} = C \frac{\Delta f_{min}}{2 \times slope} = \frac{C}{2B} \quad (1)$$

where C is the speed of sound. Assuming $C = 343m/s$ at 20° Celsius, thus d_r is $\frac{343m/s}{2 \times 6000s^{-1}} = 2.88cm$. Note that this is the resolution that FMCW can separate mixed echoes. The resolution of major echo location corresponds to one single acoustic sample, which is $\frac{C}{2F_s} = 3.57mm$, where $F_s = 48KHz$ is the recording sampling frequency. The spectrogram of the segmented face region echoes after FMCW signal mixing is then used as input for CNN training in Section 5.1.

5 AUTHENTICATION MODEL

We design an end-to-end hybrid machine learning framework for authentication, which consists of two major components (shown in Figure 7): a CNN based acoustic representation learning and an SVM based two-factor authentication.

5.1 Acoustic Representation Learning

Traditional acoustic features such as mel-frequency cepstral coefficients [32], chromagram [37] and spectral contrast [25] have been proven to be effective in human speech recognition and voice-based authentication, but not in active acoustic sensing as in our case. Recently, deep learning approaches (especially CNNs) have shown great successes in a variety of challenging tasks such as image classification due to their powerful automatic feature extraction [23, 43]. We design a CNN based neural network which takes the spectrogram of the segmented signal as input, and train it on a large data set collected from users. We find that such extracted features outperform all traditional features (in Section 8.2).

A customized CNN architecture designed for acoustic feature learning is shown in Table 1. We use rectified linear unit (ReLU) as activation function for convolutional layers, a popular choice especially for deep networks to speed up training. Two max pooling layers with a size of 2×2 are used to down-sample the input representations from their

Figure 7: The authentication framework consists of two major components: acoustic representation learning and two-factor authentication.

previous activation layers. This saves computational costs by reducing the number of parameters for both training and inference, which is critical when the model needs to be deployed on mobile devices. Dropout layers are added after each max pooling layer to prevent over-fitting. Batch normalization normalizes the output of a previous layer by subtracting the batch mean and dividing by the batch standard deviation, which increase the stability of the neural network and speed up training ($\sim 6\times$ speedup in our case). Categorical cross-entropy is used as the loss function. The dense layer with softmax activation function outputs the probability of each class. The CNN is trained on a data set that contains acoustic samples from 50 classes (45 users and 5 non-human classes). Note that although the CNN is trained for 50 classes, the objective of the trained model is to extract features that can be used to distinguish far more classes beyond those 50. To use the trained model as a general acoustic feature extractor, the last layer, which is used for final classification, is removed. Thus the remaining network outputs a 128 dimensional feature vector. The trained model has 710539 parameters, and a size of 5.47MB, which is portable enough for mobile devices for real-time inference.

5.2 Facial Landmark Detection

We also extract lightweight visual features of the face to complement acoustic ones. The vision techniques serve two purposes: i) we detect facial landmarks which are later used as basic visual features. ii) we track the user's face on the smartphone screen so that the user can hold the device within

Table 1: CNN layers and parameter amounts.

Layer	Layer Type	Output Shape	# Param
1	Conv2D + ReLU	(33,61,32)	320
2	Conv2D + ReLU	(31,59,32)	9248
3	Max Pooling	(15,29,32)	
4	Dropout	(15,29,32)	
5	Batch Normalization	(15,29,32)	128
6	Conv2D + ReLU	(15,29,64)	18496
7	Conv2D + ReLU	(13,27,64)	36928
8	Max Pooling	(6,13,64)	
9	Dropout	(6,13,64)	
10	Batch Normalization	(6,13,64)	256
11	Flatten	(4992)	
12	Dense + ReLU	(128)	639104
13	Batch Normalization	(128)	512
14	Dense + Softmax	(50)	5547

Figure 8: Facial landmarks and face tracking, and landmark transformation between two camera positions.

some “valid” zone (thus distance and orientation) for data collection.

We detect the 2D coordinates of facial landmarks (e.g., corners/tips of eyes, nose and mouth) on the image as features, using the mobile vision API from Google [22] on Android platform (shown in Figure 8(a)). The face is tracked by a bounding rectangle. We observe that these landmarks describe critical geometry features on the face, and their locations are associated with the relative position from the device to the face. Detailed description of the implementation is presented in Section 7.

5.3 Two-factor Authentication

The joint acoustic and visual features are used for two-factor authentication.

Features Summary. The 2D coordinates of facial landmarks on the image are concatenated with the corresponding 128-dimensional CNN features as the joint features for final authentication. Both acoustic and vision data are collected simultaneously so that they are well synchronized, which

ensures the correspondence between facial landmarks distribution on the screen and relative device position, thus echo signals.

Classifier. One-class SVM is an unsupervised algorithm that learns a decision function for novelty detection: classifying new data as similar or different to the training set. It detects the soft boundary of the training set so as to classify new samples as belonging to that set or not. We use one-class SVM with radial basis function (RBF) kernel function for final classification. It allows us to train an SVM classifying model for a new user (or the same user wearing new hats or glasses) on mobile devices easily, without requiring large amounts of training data as in CNN.

Ideally, a user should move the device at various relative positions to the face so as to collect sufficient training data during user registration. In practice, this imposes more efforts on the user, and it is hard to tell when sufficient data has been collected. Insufficient training data will cause higher false negatives (i.e., denial of the legitimate user). Thus we propose a data augmentation technique, which populates the training data by generating “synthesized” training samples based on facial landmark transformation and acoustic signal prediction. During this process, we transform measured facial landmarks and acoustic signals into synthesized ones, by assuming different poses of the phone.

5.4 Data Augmentation

In projective geometry, the projection matrix P of a 3D point (x_w, y_w, z_w) in the world coordinate system onto the image plane in camera is modeled as Equation 2:

$$\lambda \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_1 \\ r_{21} & r_{22} & r_{23} & t_2 \\ r_{31} & r_{32} & r_{33} & t_3 \end{bmatrix} \begin{bmatrix} x_w \\ y_w \\ z_w \\ 1 \end{bmatrix} \quad (2)$$

$$= K \cdot [R|T] \cdot \begin{bmatrix} x_w \\ y_w \\ z_w \\ 1 \end{bmatrix} = P \cdot \begin{bmatrix} x_w \\ y_w \\ z_w \\ 1 \end{bmatrix}$$

where λ is the scale factor for homogeneous coordinates, (u, v) denotes its pixel coordinate on image, $K = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$

is the intrinsic matrix of the camera, e.g., the focal length f_x and f_y , skew s , and image center (c_x, c_y) in pixels. $[R|T]$ represents the extrinsic matrix of the camera, i.e., camera’s pose in the world coordinate system, where R is a 3×3 matrix for its 3D orientation, and T is a 3×1 matrix for its 3D translation.

As shown in Figure 8(b), assume two cameras take images of the same object at different distances/angles, and $x =$

$[u, v, 1]^T$ and $\mathbf{x}' = [u', v', 1]^T$ represent the object's pixel coordinates on two images. Without loss of generality, we define the first camera as the world origin, thus the projection matrix of two cameras are:

$$P = K[I|0], P' = K'[R|T] \quad (3)$$

where I is a 3×3 identity matrix.

Based on the above background of projective geometry, we can transform the landmark pixel coordinates in one camera to those of any new camera pose, thus augmenting our training set.

Step 1: Compute the landmark's world coordinates. Given the projection matrix P and landmark pixel coordinates $\mathbf{x}$ of the first camera, we can compute the landmark's world coordinates as $(x_w, y_w, z_w)^T = z_w K^{-1} \mathbf{x}$, where z_w is the distance of landmark from camera center, which can be measured via our acoustic sensing.

Step 2: Transform the landmark onto new images. From the projection matrix of the new camera pose, we can compute the corresponding pixel coordinates of the landmark as:

$$\mathbf{x}' = K' R K^{-1} \mathbf{x} + K' T / z_w \quad (4)$$

This transform equation consists of two parts: the first term depends on the image position alone, i.e., $\mathbf{x}$, but not the landmark's depth z_w ; the second term depends on the depth and takes account of the camera translation. In the case of pure translation ($R = I$, $K' = K$), Equation 4 reduces to $\mathbf{x}' = \mathbf{x} + K T / z_w$.

Step 3: Data augmentation. We augment our training set based on Equation 4. i) Before data collection, we calibrate the camera with a benchmark paper printing of a chessboard with known size, thus obtain its intrinsic matrix K . ii) Given a new camera pose of $\theta = (T, \phi)$, where T for its 3D coordinates and $\phi = (\alpha, \beta, \gamma)$ for its rotation angles along three axes of the smartphone, we transform ϕ to the 3×3 rotation matrix R based on Rodrigues's Formula [50], then compute $\mathbf{x}'$ via Equation 4.

Accordingly, following the sound propagation inverse-square law, the face region signal segment is shifted by the same distances, with the amplitude adjusted by the scale equal to the inverse of the square of distance. Due to the omni-directional property of smartphone speaker and microphone, a slight device rotation at a fixed position causes negligible changes in the signal, thus only device position change accounts for acoustic signal transform.

6 AUTHENTICATION MODES

We propose three authentication modes: two-factor one-pass authentication, low-power continuous authentication and ultra low-power presence detection, suitable for scenarios

requiring progressively less security level but more user convenience and power efficiency.

Two-factor One-pass Authentication. In this mode, the user must hold the phone properly to align his face within the valid area rectangle as shown on the screen (see Figure 9). Both visual facial landmarks from camera images and acoustic features extracted by the trained CNN are fed to the SVM for recognition. This incurs the most computation, energy costs, providing the highest security level suitable for scenarios such as phone unlock, account log in.

Low-power Continuous Authentication (LP mode). In this mode, acoustic features extracted from the CNN are used in one-class SVM classification. It avoids power hungry cameras and real-time video processing, providing reduced security level suitable for scenarios such as continuous access/browse of private data in banking transactions after login is completed. The user needs to hold the phone in position ranges similar to training data collection.

Ultra Low-power Presence Detection (ULP mode). This mode uses acoustic signal only and an SVM model to detect the presence of the user face. To minimize computation and energy costs, the spectrum of a set of samples (e.g., the first 80 after the direct path signal) instead of CNN extracted features is fed to SVM. Data collected to train the SVM include positive samples while holding the device before the user's face, negative samples when putting the device on tables, in pockets, or holding it away from the user. This mode consumes the least power and is suitable for scenarios like auto screen lockup when the user face is not present.

7 IMPLEMENTATION

We implement *EchoPrint* on multiple Android smartphones, including Samsung S7 Edge, Samsung S8, and Huawei P9. Figure 9 shows the UI. The prototype consists of three major modules: facial landmark detection, acoustic sensing, and machine learning pipeline for authentication. The "predict/train" switch is used for authentication/registration selection ¹.

Facial Landmark Detection. We use Google mobile vision API [22] for real-time facial landmark detection and face tracking. The frame rate is set at 30 fps with a resolution of 1024×768 . The middle red rectangle (Figure 9) denotes the detected face area, and two green rectangles are the inner and outer bounds of face valid areas, which are fixed. The user face must be aligned within the two green rectangles during data collection; otherwise the acoustic data are discarded. Yellow dots are detected facial landmarks, saved in pixel coordinates.

Acoustic Sensing. The acoustic signal is pre-processed and displayed on the screen in real-time, and the segmented

¹"P9/S8" switch is used for selecting microphones (top vs. bottom). P9 and S8 are just opposite, thus requiring manual selection.

Figure 9: EchoPrint prototype application UI.

signal from the face is highlighted in blue. For better visualization, we amplify the signal by $3\times$ after the direct path signal.

Machine Learning Pipeline. The machine learning pipeline requires one CNN acoustic feature extractor and one SVM classifier. We train the CNN model off-line on a PC with Intel i7-8700K CPU, 64GB memory and GTX 1080 Ti GPU. Keras [16] with Tensorflow [8] backend is used for CNN construction and training. The trained model is frozen and deployed on mobile devices. Using acoustic features extracted from CNN, an SVM classifier using LibSVM [14] is trained on mobile devices. Both CNN and SVM inferences are performed on mobile devices in real-time.

8 EVALUATION

8.1 Data Collection

We obtain the required human subjects training certificate from our institution before data collection. 45 participants of different ages, genders, and skin colors are recruited in experiments. The diversity in physical appearances of participant faces help us capture sufficient data to create a strong feature extraction model. We also include 5 non-human classes: printed/displayed human faces on different materials such as paper, desktop monitor, photo on paper box, wall and a marble sculpture. During data collection, each participant is asked to hold the smartphone in front of his/her face to ensure face alignment. To accommodate slight phone movements, participants are encouraged to move the phone slowly to cover different poses. Data captured while the face is out of the valid area are discarded automatically. About 120 seconds' data is captured from each user, at around 7 - 8MB and containing ~ 2000 samples. To ensure diversity, the data is collected in multiple uncontrolled environments (e.g., quiet

Figure 10: Different classifiers performance on extracted features from CNN, and SVM performance using different features.

laboratories, noisy classrooms, and outdoor environments) under different background noises and lighting conditions. A portion of the participants who are more accessible to us collected data in multiple sessions at different times and locations. Facial landmarks are also detected and recorded simultaneously, but no facial images are recorded to protect the participants' privacy. In total, the data set contains 91708 valid samples from 50 classes. We divide it in three parts, 70% for model training, 15% each for model validation and testing. Additionally, 12 more volunteers join as new users for model evaluation.

8.2 CNN Feature Extractor Performance

We compare the performance of different classifiers and feature extraction methods using the test data set.

Different Classifiers. The last fully connected layer of our trained CNN is removed so that the remaining network is used as a general feature extractor. Such extracted features are then fed to different classifiers for final classification. We compare Linear Regression (LR), Linear Discriminant Analysis (LDA), K-Nearest Neighbor (KNN), Decision Tree (DT), Naive Bayesian (NB), Support Vector Machine (SVM) and a standalone Neural Network (NN). The box plot in Figure 10(a) shows the lower and upper quartiles, and the median. The whiskers extend from the box show the range of accuracy, and outliers beyond the whiskers are marked as circles. We find SVM outperforms all other classifiers, and it takes short time (15.06s compared to 65.38s of NN which has second best performance) for training. Thus we use SVM as the final classifier for authentication.

Different Features. We compare different commonly used acoustic features: spectrogram (SPEC), mel-frequency cepstral coefficients (MFCC) [32], chromagram (CHRO) [37], spectral contrast (CONT) [25] and our CNN features. Figure 10(b) shows their accuracies using SVM classifier. Our CNN extractor outperforms all other features and achieves the highest accuracy of $\sim 95\%$, which show the effectiveness and necessity of the CNN feature extractor. Spectrogram

has less accuracy at $\sim 85\%$, and chromagram 67% . MFCC and CONT have much lower accuracy $\sim 30\%$, which is what we expected because they are mostly used for human voice recognition, not active acoustic sensing used in *EchoPrint*. Besides, the 15.06s using CNN features to train the SVM model is a fraction of the 134s needed when training with spectrogram. This is a significant improvement when training a model on resource-constraint mobile devices, which is critical for the speed of user registration.

8.3 Authentication Accuracy

In a binary classification problem, there are four results: true positive (TP), positive samples correctly classified as positive class; true negative (TN), negative samples correctly classified as negative class; false positive (FP), negative samples wrongly classified as positive class; false negative (FN), positive sample wrongly classified as negative class. Specifically, in authentication scenarios, a high TP means the authorized user can get access easily, a high TN means the system can block most attacks. The worst case is high FP, which means unauthorized users gain access. A high FN means the authorized user may be denied access, which is annoying and not user-friendly. In this evaluation, we train a one-class SVM for each subject and attack the model using the data from the rest users. Note that the model is trained on positive samples only, it does not have negative samples from attackers during training.

8.3.1 Precision, Recall, F-score and BAC. We introduce precision, recall, F-score and balanced accuracy (BAC) as metrics. Precision is the fraction of true positives among all samples classified as positive, defined as $P = \frac{TP}{TP+FP}$; recall is the fraction of true positives among all positive samples, defined as $R = \frac{TP}{TP+FN}$. A high precision means the authorized user can pass easily, a high recall means the authorized user is seldom denied. When the class distribution is imbalanced, precision and recall alone can be misleading. We also introduce F-score and balanced accuracy (BAC), both insensitive to class distribution. F-score is the harmonic mean of precision and recall with a best value of 1 and worst value of 0, defined as $F\text{-score} = 2 \frac{P \cdot R}{P+R}$. BAC is the average of true positive rate ($TPR = \frac{TP}{TP+FN}$) and true negative rate ($TNR = \frac{TN}{TN+FP}$), defined as $BAC = \frac{1}{2} \cdot (TPR + TNR)$. A BAC of 1 means no false positive (i.e., successful attack) or false negative (i.e., denied access of legitimate users).

Table 2 shows the mean and median accuracies using vision, acoustic, and joint features. Vision (2D coordinates of a few facial landmarks like the corners/tips of eyes, nose and mouth) is the worst with a low average precision of $\sim 72\%$. Acoustic achieves 86% , and joint features further increase it to 88% while also decreasing recall by $\sim 6\%$. That is because simple 2D coordinates of facial features do not capture

Table 2: Mean/median accuracy with vision, acoustic and joint features.

	Vision	Acoustic	Joint
Precision (%)	72.53 / 80.32	86.06 / 99.41	88.19 / 99.75
Recall (%)	64.05 / 64.04	89.82 / 89.84	84.08 / 90.10
F-score (%)	65.17 / 69.19	85.39 / 94.31	83.74 / 93.23
BAC (%)	81.78 / 81.83	94.79 / 94.88	91.92 / 95.04

Figure 11: The precision, recall, F-score and BAC of one-class SVM model using acoustic and joint features.

the full characteristics of the face, thus alone they do not perform well when many test subjects exist. They can help “block” unauthorized users which happen to have similar acoustic features, thus increasing precision. But they also make it harder for the authorized user to pass, thus decreasing recall. Both acoustic and joint features have an average F-score $\sim 85\%$ and BAC above 90% . The vision features used are not sophisticated and detailed visual features (e.g., the contour of face) of facial appearances as used in state-of-the-art vision-based face recognition systems. These basic face landmarks are mainly used for face alignment, which is critical for robust acoustic sensing. While such facial landmarks are not intended to greatly improve recognition accuracy, *EchoPrint* as an acoustic based approach is free to incorporate more sophisticated facial features, e.g., features from a deep neural network trained on a huge face image dataset [10]. Those would have a much higher impact on performance improvements.

Note that the median precision ($\sim 99\%$) and F-score ($\sim 94\%$) for both acoustic and joint features are much higher than the respective average ($83 \sim 88\%$). This is caused by outliers. Figure 11 shows the box plot of all four metrics of acoustic and joint features. A few outlier classes with very low precision cause low average but do not affect the median. Such outliers are mainly non-human noise classes or human classes with very limited valid samples. When such outliers are excluded, the averages will increase significantly to above $\sim 95\%$.

8.3.2 Performance on New Users. To evaluate how well the pre-trained CNN can extract features for new users, we invite 12 additional volunteers whose data are not used in CNN training. Each volunteer follows the same data collection process for ~ 2 minutes' data, half of which are used for SVM training and the other half for testing. We train a one-class SVM model for each volunteer, and test the model with positive samples from the user and negative samples from all other users, including the data from 50 classes used in CNN training. Table 3 shows that the average precision is over 98%, about 10% increase compared to results in Table 2 due to the absence of outlier classes. Similarly the average recall, F-score and BAC are all improved compared to those in Table 2.

Table 3: Authentication accuracy of new users.

	Mean	Median	Standard Deviation
Precision (%)	98.05	99.21	2.78
Recall (%)	89.36	89.91	1.62
F-score (%)	93.50	94.33	1.68
BAC (%)	93.75	94.52	0.85

8.3.3 Data Augmentation Evaluation. We evaluate how effective data augmentation can improve the performance by generating “synthesized” training samples when training data is limited. We split 20% samples from the *2min* data as testing set, and vary the size of training set from 20% to 80%. The data set is shuffled before the splitting to make it more balanced. Figure 12 shows the precision, recall, F-score and BAC under different amounts of training samples from 50 to 600, which are tested against another 1054 positive testing samples and all the 91708 negative samples from existing 50 classes. It is obvious that data augmentation improves recall significantly, thus F-score and BAC, especially when the training samples are very limited (e.g., <100). As the size grows, the recall with data augmentation is always higher. However the precision decreases to $\sim 95\%$, which is because “synthesized” training samples have more noises, making it easier to have false positives. The performance becomes stable with more than 400 training samples, which can be collected within one minute when registering a new user.

8.3.4 Continuous Modes Evaluation. We evaluate the two continuous modes of presence detection and continuous authentication that uses only acoustics.

Presence Detection. We put the smartphone at different locations as well as holding it in front of the user's face. The detection result is shown on the screen in real-time so that we know the correctness. From our experiments, it can differentiate putting on a table and holding in front of the user with nearly 100% accuracy with unnoticeable

Figure 12: Classification performance comparison of data augmentation (DA) under different training data amounts.

Figure 13: Continuous authentication performance with different number of trials.

delay. Holding in the air sometimes may be detected as user presence when the device is close to some major objects, which may affect timely screen lockup.

Continuous Authentication. To ensure friendly user experience during continuous authentication, a low false negative rate is very important. One volunteer participates in this experiment with a trained model using data when the face is aligned. In the authentication phase, the volunteer keeps using the device as normal and tries to keep it within positions where the face is aligned, with the camera disabled. We evaluate the precision, recall, F-score and BAC when multiple authentication trials are conducted for each cycle. Authentication trial happens every 100ms thus one verdict from multiple trials is fast enough, causing no noticeable delay to the user. At least one trial must pass in a cycle to declare authentication success. Figure 13 shows that more trials increase the recall rapidly while decreasing the precision. This is because more trials give the user more chances to pass, thus reducing denials while increasing false positives. We choose 3 trials for each authentication circle to balance all the metrics.

8.3.5 Miscellaneous. We evaluate the following factors that have direct impacts on practical use.

Figure 14: Performance under difference noises.

Figure 15: Average recall of 5 users before/after model updating with new training data.

Robustness Against Background Noise. We evaluate the robustness again background noise under different conditions: quiet room, with ambient noise (playing pop music nearby), and with ambient plus self noise (playing music through earpiece speaker on the *same device* during data collection, an extreme condition). Figure 14 shows the results. Except for a slightly lower recall, there is no major difference between quiet and ambient noise conditions, which demonstrates *EchoPrint* is very robust to ambient noise. The ambient plus self noise brings down to recall to $\sim 70\%$, but still the precision remains above 95%.

Image Spoofing Attacks. We print color photos of 5 volunteers in 10 different sizes on paper, and also display the photos on desktop monitors while zooming in/out gradually, both at various distances between 20 - 50cm to the smartphone. The printed and displayed photos can easily pass the system if only vision features are used, but none of them can pass the acoustic or two-factor authentication.

User Appearance Changes. Appearance changes such as wearing glasses/hats cause changes in the reflected acoustic signals, thus more false negatives and low recall. To combat such problems, we retrain the SVM model with data samples of new appearances in addition to the existing training data. Figure 15 shows the average recall of 5 users with different appearance changes before/after model update using additional ~ 1 minute's data. Without retraining, the recall drops to single digits. After the retraining, it increases back to normal levels, so correct users can pass easily. This shows retraining is effective combating such changes.

Table 4: Mean/max resource consumption.

Device	Memory (MB)	CPU (ms)	Delay (ms)
S7	22.0 / 50.0	6.42 / 31.59	44.87 / 91
S8	20.0 / 45.0	5.14 / 29.04	15.33 / 35
P9	24.0 / 53.0	7.18 / 23.87	32.68 / 86

User Experience Study. We conduct a survey with 20 users (mostly graduate and undergraduate students) to collect their feedback, mainly on two aspects that directly impact user experience, the sensitivity to the emitted sound signal and the effort for new user registration. Out of 20 users, only 4 reported able to hear the high frequency sound from the earpiece when holding the smartphone at a normal distance. Out of 20 users, 9 rated *EchoPrint* equally easy to register as other authentication systems such as image-based face recognition and fingerprint sensor, 6 rated it harder and 5 rated it easier.

8.4 Resource Consumption

We evaluate memory, CPU usage using the Android Studio IDE Profiler tool, and power consumption using Qualcomm's Trepro Profiler tool [5] on Samsung S7 edge, Samsung S8, and Huawei P9.

Memory & CPU Consumption. Table 4 shows the resource consumption on three smartphones. The memory consumption has an average $\sim 22MB$ and max $\sim 50MB$, which appears when CNN feature extraction using tensorflow inference is running. The average amount of time for the CPU to complete all the machine learning inferences is low on all phones (5 $\sim 7ms$). The max CPU time is around $\sim 30ms$, still very low. Such low memory and CPU usage makes it possible to deploy *EchoPrint* on most existing devices.

Response Delay. Response delay is the time needed for the system to produce an authentication result after the raw input signal is ready (Table 4). Samsung S8 has the least delay with an average of $\sim 15ms$, and the other two 32 - 45ms. The delay approaches maximum when the user keeps moving the phone trying to align the face in the valid area, which incurs a lot of camera preview refreshing and rendering. It is also affected by other computation heavy background apps. For real-time continuous authentication, the delay between consecutive sound signal emitting is 50ms. We choose to do authentication every other emitting, leaving sufficient time for processing.

Power Consumption. We test the three modes and pure vision based authentication using 2D coordinates of facial landmarks, each for 30 minutes to measure power consumption on Samsung S7 Edge, S8 and Huawei P9. We use Qualcomm's Trepro Profiler tool [5], which provides power consumption in *mW* for a chosen application. We subtract the

Table 5: Power consumption of different modes.

Device	ULP (mW)	LP (mW)	Two-factor (mW)	Vision (mW)
S7	305	1560	2485	1815
S8	215	1500	2255	1655
P9	265	1510	2375	1725

background power consumption while the screen is on, the increased power consumption caused by different modes are shown in Table 5. The results show that presence detection consumes minimum power, while low power continuous authentication takes less than that of pure light weight vision based authentication. Two-factor authentication has the highest battery consumption; but it is also designed for occasional one-pass authentication finishing in just a few seconds, not long time continuous operation. The slight power increase of vision based mode over LP is due to the simple form of facial landmarks used, which are much lighter weight compared to more sophisticated ones such as those in OpenFace [10].

9 RELATED WORK

Smartphone Authentication. Personal Identification Number (PIN) or a text/graphical password are the earliest and still most widely used smartphone user authentication methods. Despite the simplicity, the PIN or password can be easily peeked by someone close by [46]. Speech recognition is easy to spoof when the voice is recorded, or closely imitated by advanced learning algorithms [41]. BreathPrint [15] senses the user's breath sound, which may change significantly when the user has intense exercises. Vision based face recognition is vulnerable to camouflaged images. Although eye blinks can enhance its security [4], a recorded video can still spoof the system. Fingerprint sensors have achieved great security and convenience. However the sensor takes a lot of precious space, and forging one from fingerprints left by the user is proven practical [3]. More advanced fingerprint sensors use ultrasonics [42] to penetrate the skin and construct 3D imaging, but such sensors are unavailable on most smartphones. Apple's FaceID [7] uses special TrueDepth sensors, bringing extra hardware costs and requiring significant design changes. Intel's RealSense [27] is a similar technology, but it is costly and power-computation heavy, unsuitable for mobile devices. Unlike all the above solutions, *EchoPrint* is the first to leverage active acoustic sensing combined with visual features for user authentication. It achieves high balanced accuracy ($\sim 95\%$) using existing hardware.

Acoustic-based Face Recognition. Acoustics has been used for face recognition in some prior work [19, 26, 35, 36]. I. E. Dror *et al.* [19] recognize a limited number of five human faces with an accuracy over 96% and the gender of 16 faces

with an accuracy of 88% using bat-like sonar input from special ultrasonic sensors. K. Kalgaonkar *et al.* [26] propose a sensing mechanism based on the Doppler effect to capture the patterns of motion of talking faces using ultrasound. K.K. Yoong *et al.* [35, 36] classify up to 10 still faces with an accuracy of 99.73% using hand-crafted features from ultrasound echo signals. Compared to all the above work using special ultrasonic sensors which are not available in consumer electronics, *EchoPrint* uses commodity smartphone speakers and microphones not intended for ultrasonic frequencies. This puts a lot of challenges on the signal design and processing, and much more experiments and tests to find out the best acoustic signal design providing required sensing resolution within hardware limitations, while minimizing the audibility to users. Besides, such prior work uses pure ultrasonic sensing without the aid from vision, thus creating major limitations (e.g., requiring the user to move the head at a fixed location and angle). While *EchoPrint* leverages the vision to align faces using face tracking algorithms for practical two-factor vision-acoustic authentication. *EchoPrint* is the first mobile device based approach. New users can register using a pre-trained CNN model to extract features and train a standalone SVM model on-device. While prior work uses hand-crafted features or needs re-training the whole neural network, inefficient and infeasible on mobile devices.

Acoustic Sensing on Smartphones. Acoustic sensing is widely used for distance measurement, thus applications in localization, tracking, stress and encounter detection. Beep-Beep [40] and SwordFight [55] measure the distance between two smartphones directly; Liu *et al.* [30] leverage cross-correlation to compute the arrival time difference for keystroke snooping; Yang *et al.* [51] detect driver phone usage in a vehicle; EchoTag [52] recognizes different locations and BatMapper [56, 58] builds indoor floor plans using echo signals. Besides, acoustic ranging can significantly improve smartphone localization accuracy, e.g., adding constraints among peer phones [29], deploying an anchor network that transmits spatial beacon signals [31], or enabling high-precision infrastructure-free mobile device tracking [57]. UbiK [47], AAMouse [53], FingerIO [39], and LLAP [48] leverage phase shift in received signals for near field finger gesture tracking, achieving $\sim 1\text{cm}$ or higher accuracy. StressSense [33] detects personal stress using smartphones in unconstrained acoustic environments. ApenaApp [38] monitors the minute chest and abdomen breathing movements using FMCW [45], and SonarBeat [49] monitors breathing beat using signal phase shifts. CAT [34] leverages FMCW with external speakers for smartphone movement tracking and achieves mm-level accuracy. DopEnc [54] extracts acoustic features to identify encountered persons. Compared to them, *EchoPrint* leverages acoustic features from deep neural networks for a different purpose of user authentication.

10 DISCUSSION

Leveraging Sophisticated Vision Features. *EchoPrint* is open for integration with the state-of-the-art image-based face recognition algorithms for more practical use, such as OpenFace [12], a state-of-the-art open source face recognition system based on neural networks. We implement a preliminary prototype leveraging OpenFace, where a 128-dimensional feature vector is generated as the image representation. We test it with 5 volunteers and attack the system with user images printed on paper or images/videos displayed on desktop monitor. Despite of 100% accuracy recognizing each face, we find OpenFace has almost no capability of identifying images against real human faces. In contrast, the two-factor approach blocks all images attacks due to the significant differences in acoustic features, while retaining high recognition accuracy. We realize that the results are preliminary due to limited time and resources. Nevertheless, they are promising, showing great potential of combining advanced visual features. We leave more mature design and systematic evaluation as future work.

Comparison with FaceID. FaceID is a mature commercial product, the outcome of a large team of highly skilled engineers from a technology giant. In terms of engineering maturity, which matters greatly in final performance, there is no way *EchoPrint* can approach that of FaceID. Thus, our position is never to beat or replace FaceID. Rather, our goal is to present an alternative low-cost acoustic-based approach that is promising to have similar performance while much lower costs, and possibly some advantages over FaceID (e.g., FaceID may fail in direct sunlight [1]; while *EchoPrint* uses acoustics and is not affected by strong sunlight.). There is great space of further improvements in *EchoPrint*, such as obtaining data from large populations to train a more robust model, incorporating acoustic data from many different angles, all of which will further improve its performance.

Limitations. *EchoPrint* is only a research prototype and far from a well engineered product. It has several main limitations: *i) requirement of face alignment.* To ensure both high true positive and low false negative, face alignment is required for authentication. It may be inconvenient to users to hold the phone in such positions. In contrast, FaceID has much more flexibility in device holding positions. *ii) limitations from vision.* *EchoPrint* leverages vision algorithms for facial landmark detection, thus inheriting their limitations. We notice that face tracking is not stable under poor lighting, which makes it hard for face alignment, thus more false negatives. *iii) user appearance changes.* The current CNN feature extractor is trained on limited data, far from exhaustive to be robust against various appearance changes such as hats, glasses or hair styles. Retraining the SVM model with new data is promising for combating such changes. Similar to

FaceID, an online model updating mechanism is needed to address such changes dynamically. *iv) continuous authentication usability.* Although we demonstrate *EchoPrint* has the potential for pure acoustic-based continuous authentication, it requires the smartphone aligned and in front of the user's face, which impacts the usability. To mitigate such problem, we can either decrease the frequency of acoustic sampling and authentication at a cost of decreased security, or register users with data from different angles thus enlarging the range of effective authentication. We believe that continuous authentication, although an additional feature with some robustness issues in its current form, is indeed valuable for many application scenarios, e.g., doing banking transactions in public places.

Future Work. *i) enhancing CNN acoustic feature extractor.* We will collect more training data from more users (e.g., by crowdsourcing) with larger variety. This will further improve performance together with more sophisticated neural network design. *ii) integration with existing solutions.* Due to the rapid progress of deep learning, image based face recognition has achieved unprecedented accuracy. *EchoPrint* can be integrated with existing pure image based solutions to enhance their anti-spoofing capability. *iii) feasibility of replay attacks.* In theory the echoes can be recorded and replayed. However a successful attack is far from trivial: weak, high frequency echoes are difficult to record at sufficient fidelity, and they must be replayed at proper timing after signal emitting. The amount of hardware resources and human efforts needed for successful attacks could be huge, and we plan to find it out. *iv) large scale experiment.* We only have ~ 50 users in the current experiments, while large scale experiment (e.g., thousands or more) is needed for a mature solution. We will seek ways to do experiments at such scale.

11 CONCLUSION

In this paper, we propose *EchoPrint*, which leverages acoustics and vision on commodity smartphones for two-factor authentication. To combat phone pose changes and extract reliable acoustic features that best distinguish different users, a convolutional neural network is trained on a large acoustic data set. The CNN is then used as general acoustic feature extractor to feed an SVM based classifier for authentication. Experiments show that *EchoPrint* achieves 93.75% balanced accuracy and 93.50% F-score, while the average precision is 98.05%.

ACKNOWLEDGMENTS

We thank our anonymous shepherd and reviewers for their invaluable feedback and comments. This work is supported in part by US NSF 1730291 and China NSF 61702035.

REFERENCES

- [1] 10/31/2017. Does Apple iPhone X Face ID Have a Sunlight Problem? <http://fortune.com/2017/10/31/apple-iphone-x-face-idsunlight/>.
- [2] 11/9/2017. iPhone X: This is how much it costs to make one, in components., <http://www.zdnet.com/article/iphone-x-this-is-how-much-it-costs-to-make-in-components/>.
- [3] 12/3/2007. How to fool a fingerprint security system as easy as ABC. <http://www.instructables.com/id/How-To-Fool-a-Fingerprint-Security-System-As-Easy-/>.
- [4] 2011. Fast face recognition: Eye blink as a reliable behavioral response. *Neuroscience Letters* 504, 1 (2011), 49 – 52.
- [5] 2018. Qualcomm Trepp power profiler. <https://developer.qualcomm.com/software/trepp-power-profiler>.
- [6] 3/31/2017. Galaxy S8 face recognition already defeated with a simple picture. <https://arstechnica.com/gadgets/2017/03/video-shows-galaxy-s8-face-recognition-can-be-defeated-with-a-picture/>.
- [7] 6/8/2018. About Face ID advanced technology. <https://support.apple.com/en-us/HT208108>.
- [8] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*, Vol. 16. 265–283.
- [9] JM Martín Abreu, T Freire Bastos, and L Calderón. 1992. Ultrasonic echoes from complex surfaces: an application to object recognition. *Sensors and Actuators A: Physical* 31, 1-3 (1992), 182–187.
- [10] Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. 2016. Openface: A general-purpose face recognition library with mobile applications. *CMU School of Computer Science* (2016).
- [11] Apple. 2018. Apple pay. <https://www.apple.com/apple-pay/>.
- [12] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. 2016. Openface: an open source facial behavior analysis toolkit. In *Applications of Computer Vision (WACV), 2016 IEEE Winter Conference on*. IEEE, 1–10.
- [13] Derya Birant and Alp Kut. 2007. ST-DBSCAN: An algorithm for clustering spatial-temporal data. *Data & Knowledge Engineering* 60, 1 (2007), 208–221.
- [14] Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* 2 (2011), 27:1–27:27. Issue 3.
- [15] Jagmohan Chauhan, Yining Hu, Suranga Seneviratne, Archan Misra, Aruna Seneviratne, and Youngki Lee. 2017. BreathPrint: Breathing Acoustics-based User Authentication. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. 278–291.
- [16] François Chollet et al. 2015. Keras. [urlhttps://github.com/keras-team/keras](https://github.com/keras-team/keras).
- [17] John G Daugman. 1994. Biometric personal identification system based on iris analysis. US Patent 5,291,560.
- [18] Konstantinos G Derpanis. 2010. Overview of the RANSAC Algorithm. *Image Rochester NY* 4, 1 (2010), 2–3.
- [19] Itiel E Dror, Faith L Florer, Damien Rios, and Mark Zagaeski. 1996. Using artificial bat sonar neural networks for complex pattern recognition: Recognizing faces and the speed of a moving target. *Biological Cybernetics* 74, 4 (1996), 331–338.
- [20] Nguyen Minh Duc and Bui Quang Minh. 2009. Your face is not your password face authentication bypassing lenovo-asus-toshiba. *Black Hat Briefings* 4 (2009), 158.
- [21] Google. 2014. Android pay. <https://www.android.com/pay/>.
- [22] Google. 2017. Introduction to Mobile Vision. <https://developers.google.com/vision/introduction>.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [24] Emmanuel C Ifeachor and Barrie W Jervis. 2002. *Digital signal processing: a practical approach*. Pearson Education.
- [25] Dan-Ning Jiang, Lie Lu, Hong-Jiang Zhang, Jian-Hua Tao, and Lian-Hong Cai. 2002. Music type classification by spectral contrast feature. In *Multimedia and Expo, 2002. ICME'02. Proceedings. 2002 IEEE International Conference on*, Vol. 1. IEEE, 113–116.
- [26] Kaustubh Kalgaonkar and Bhiksha Raj. 2008. Recognizing talking faces from acoustic doppler reflections. In *Automatic Face & Gesture Recognition, 2008. FG'08. 8th IEEE International Conference on*. IEEE, 1–6.
- [27] Leonid Keselman, John Iselin Woodfill, Anders Grunnet-Jepsen, and Achintya Bhowmik. 2017. Intel RealSense Stereoscopic Depth Cameras. *arXiv preprint arXiv:1705.05548* (2017).
- [28] Kai Kunze and Paul Lukowicz. 2007. Symbolic object localization through active sampling of acceleration and sound signatures. In *International Conference on Ubiquitous Computing*. Springer, 163–180.
- [29] Hongbo Liu, Yu Gan, Jie Yang, Simon Sidhom, Yan Wang, Yingying Chen, and Fan Ye. 2012. Push the limit of WiFi based localization for smartphones. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 305–316.
- [30] Jian Liu, Yan Wang, Gorkem Kar, Yingying Chen, Jie Yang, and Marco Gruteser. 2015. Snooping keystrokes with mm-level audio ranging on a single phone. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 142–154.
- [31] Kaikai Liu, Xinxin Liu, and Xiaolin Li. 2013. Guoguo: Enabling fine-grained indoor localization via smartphone. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. ACM, 235–248.
- [32] Beth Logan et al. 2000. Mel Frequency Cepstral Coefficients for Music Modeling.. In *ISMIR*, Vol. 270. 1–11.
- [33] Hong Lu, Denise Frauendorfer, Mashfiqui Rabbi, Marianne Schmid Mast, Gokul T Chittaranjan, Andrew T Campbell, Daniel Gatica-Perez, and Tanzeem Choudhury. 2012. Stresssense: Detecting stress in unconstrained acoustic environments using smartphones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 351–360.
- [34] Wenguang Mao, Jian He, and Lili Qiu. 2016. CAT: high-precision acoustic motion tracking. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. ACM, 69–81.
- [35] Phillip McKerrow and Kok Kai Yoong. 2007. Classifying still faces with ultrasonic sensing. *Robotics and Autonomous Systems* 55, 9 (2007), 702–710.
- [36] Phillip J McKerrow and Kok Kai Yoong. 2006. Face classification with ultrasonic sensing. (2006).
- [37] Meinard Müller, Frank Kurth, and Michael Clausen. 2005. Audio Matching via Chroma-Based Statistical Features.. In *ISMIR*, Vol. 2005. 6th.
- [38] Rajalakshmi Nandakumar, Shyamnath Gollakota, and Nathaniel Watson. 2015. Contactless sleep apnea detection on smartphones. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 45–57.
- [39] Rajalakshmi Nandakumar, Vikram Iyer, Desney Tan, and Shyamnath Gollakota. 2016. FingerIO: Using Active Sonar for Fine-Grained Finger Tracking. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1515–1525.
- [40] Chunyi Peng, Guobin Shen, Yongguang Zhang, Yanlin Li, and Kun Tan. 2007. BeepBeep: A High Accuracy Acoustic Ranging System using COTS Mobile Devices. In *ACM SenSys*.
- [41] Wei Ping, Kainan Peng, Andrew Gibiansky, Sercan O. Arik, Ajay Kannan, Sharan Narang, Jonathan Raiman, and John Miller. 2018. Deep

- Voice 3: Scaling Text-to-Speech with Convolutional Sequence Learning. In *Proceedings of the Sixth International Conference on Learning Representations (ICLR)*.
- [42] Qualcomm. 2018. Fingerprint Sensors. <https://www.qualcomm.com/solutions/mobile-computing/features/security/fingerprint-sensors>.
- [43] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [44] C. Stein, C. Nickel, and C. Busch. 2012. Fingerphoto recognition with smartphone cameras. In *2012 BIOSIG - Proceedings of the International Conference of Biometrics Special Interest Group (BIOSIG)*. 1–12.
- [45] Andrew G Stove. 1992. Linear FMCW radar techniques. In *IEE Proceedings F-Radar and Signal Processing*, Vol. 139. IET, 343–350.
- [46] Furkan Tari, Ant Ozok, and Stephen H Holden. 2006. A comparison of perceived and real shoulder-surfing risks between alphanumeric and graphical passwords. In *Proceedings of the second symposium on Usable privacy and security*. ACM, 56–66.
- [47] Junjue Wang, Kaichen Zhao, Xinyu Zhang, and Chunyi Peng. 2014. Ubiquitous Keyboard for Small Mobile Devices: Harnessing Multipath Fading for Fine-grained Keystroke Localization. In *Proceedings of ACM MobiSys*. 14–27.
- [48] 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*. ACM, 82–94.
- [49] Xuyu Wang, Runze Huang, and Shiwen Mao. 2017. SonarBeat: Sonar phase for breathing beat monitoring with smartphones. In *Computer Communication and Networks (ICCCN), 2017 26th International Conference on*. IEEE, 1–8.
- [50] Maria Weber and Arthur Erdélyi. 1952. On the finite difference analogue of Rodrigues' formula. *The American Mathematical Monthly* 59, 3 (1952), 163–168.
- [51] Jie Yang, Simon Sidhom, Gayathri Chandrasekaran, Tam Vu, Hongbo Liu, Nicolae Cekan, Yingying Chen, Marco Gruteser, and Richard P. Martin. 2011. Detecting Driver Phone Use Leveraging Car Speakers. In *ACM MobiCom*.
- [52] Kang G. Shin Yu-Chih Tung. 2015. EchoTag: Accurate Infrastructure-Free Indoor Location Tagging with Smartphones. *MobiCom '15 Proceedings of the 21st Annual International Conference on Mobile Computing and Networking* (2015), 525–536.
- [53] Sangki Yun, Yi-Chao Chen, and Lili Qiu. 2015. Turning a Mobile Device into a Mouse in the Air. In *Proceedings of ACM MobiSys*. 15–29.
- [54] Huanle Zhang, Wan Du, Pengfei Zhou, Mo Li, and Prasant Mohapatra. 2016. DopEnc: acoustic-based encounter profiling using smartphones. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. ACM, 294–307.
- [55] Zengbin Zhang, David Chu, Xiaomeng Chen, and Thomas Moscibroda. 2012. SwordFight: enabling a new class of phone-to-phone action games on commodity phones. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 1–14.
- [56] Bing Zhou, Mohammed Elbadry, Ruipeng Gao, and Fan Ye. 2017. BatMapper: Acoustic Sensing Based Indoor Floor Plan Construction Using Smartphones. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 42–55.
- [57] Bing Zhou, Mohammed Elbadry, Ruipeng Gao, and Fan Ye. 2017. BatTracker: High Precision Infrastructure-free Mobile Device Tracking in Indoor Environments. In *Proceedings of the 15th ACM Conference on Embedded Networked Sensor Systems (SenSys 2017)*. ACM.
- [58] Bing Zhou, Mohammed Elbadry, Ruipeng Gao, and Fan Ye. 2017. Demo: Acoustic Sensing Based Indoor Floor Plan Construction Using Smartphones. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. ACM, 519–521.