TapNet: The Design, Training, Implementation, and Applications of a Multi-Task Learning CNN for Off-Screen Mobile Input

Michael Xuelin Huang
Google, Inc.
Mountain View, CA, USA
mxhuang@google.com

Yang Li
Google, Inc.
Mountain View, CA, USA
liyang@google.com

Nazneen Nazneen
Google, Inc.
Mountain View, CA, USA
fnazneen@google.com

Alexander Chao
Google, Inc.
Mountain View, CA, USA
alexanderchao@google.com

Shumin Zhai
Google, Inc.
Mountain View, CA, USA
zhai@google.com

ABSTRACT

To make off-screen interaction without specialized hardware practical, we investigate using deep learning methods to process the common built-in IMU sensor (accelerometers and gyroscopes) on mobile phones into a useful set of one-handed interaction events. We present the design, training, implementation and applications of TapNet, a multi-task network that detects tapping on the smart-phone. With phone form factor as auxiliary information, TapNet can jointly learn from data across devices and simultaneously recognize multiple tap properties, including tap direction and tap location. We developed two datasets consisting of over 135K training samples, 38K testing samples, and 32 participants in total. Experimental evaluation demonstrated the effectiveness of the TapNet design and its significant improvement over the state of the art. Along with the datasets, codebase¹, and extensive experiments, TapNet establishes a new technical foundation for off-screen mobile input.

CCS CONCEPTS

• Human-centered computing → Gestural input.

KEYWORDS

Back-of-device, gesture recognition, IMU

ACM Reference Format:

Michael Xuelin Huang, Yang Li, Nazneen Nazneen, Alexander Chao, and Shumin Zhai. 2021. TapNet: The Design, Training, Implementation, and Applications of a Multi-Task Learning CNN for Off-Screen Mobile Input. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3411764.3445626>

¹TapNet Dataset

This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI '21, May 8–13, 2021, Yokohama, Japan
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8096-6/21/05.
<https://doi.org/10.1145/3411764.3445626>

Figure 1: TapNet is a one-model-for-all solution that jointly learns from cross-device data and performs multiple tap recognition tasks at a time.

1 INTRODUCTION

Touchscreen as the sole method of mobile device input is increasingly challenged by its inherent limitations. The difficulty of one-handed use and the visual occlusion by the operating finger are two of them. Beyond research explorations, we began to see off-screen occlusion-free alternatives in mainstream products, such as double pressing power button to activate the camera, squeezing the Active Edge [23] of Google Pixel phones and long-press power button [30] of the iPhones for voice command invocation.

In addition to a *tap event* (presence of a tap) alone, potentially useful *tap properties* include *tap direction* [20, 34] (i.e. whether the tap is on the front, back, or side of the phone), *tap location* [17, 21] (i.e. which region the tap falls on the devices), and *tap finger part* [9] (i.e. tapping with finger pad vs fingernail). Prior studies has focused on the recognition of a single tap property, based on a limited amount of data as a proof-of-concept. This paper aims to address the needs in practice, by predicting comprehensive tap properties for diverse application purposes and advancing the state of the art of off-screen interaction towards a practical level of performance. Note that the state of the art for this topic can only be advanced

if there are well-established benchmarks with accessible datasets and codebases. Our work is set to help create a benchmark with extensive dataset development, neural network model design, model training, and experimentation.

After experimenting with multiple methods and architectures for tap detection, we found that a multi-input, multi-output (MIMO) convolutional neural network (CNN) gave the best results. The resulting method, TapNet, enables joint learning on cross-device data and joint prediction of multiple tap properties. Compared with the solution of one model for each task, a joint model (see Figure 1) can exploit the interrelation among multiple tap properties during training and also saves run-time computation and memory. TapNet contains shared convolutional layers for task-agnostic knowledge extraction, meanwhile each of its output branch retains the task-specific information. TapNet uses the inertial measurement unit (IMU) signals as primary input and the phone form factor as “auxiliary information” [28] for *cross-device training*, i.e. joint training on cross-device data. Together, TapNet offers a one-model-for-all solution [13] across devices.

Our technical investigation and evaluation shed lights on two important aspects for machine learning (ML) in HCI development: methods and data. First, thanks to the MIMO architecture, TapNet increased tap property recognition accuracy over the prior art methods, particularly for the more difficult tasks such as tap location recognition. Specifically, TapNet yields a mean distance error around 10% of the screen diagonal, similar to the between-icon distance on mobile phone home page. Although this resolution is much less precise compared to the touchscreen, such resolution can already enlarge the interaction design space, by for example enabling users to perform selection by tapping on the back of the phone, even when they wear gloves.

Second, we discuss an efficient data strategy for developing a ML-based HCI application. Different from most computer vision tasks that heavily rely on multi-person data to achieve user generalizability [12, 44], a well-performing ML model for HCI (interactive) tasks may not necessarily demand multi-person training data. The key is to ensure the training data diversity. To test this hypothesis, we collected a one-person dataset for training and a separate multi-person dataset for testing (and optionally model adaptation). The results show that the one-person model could achieve a comparable level of user generalizability as a model tuned on multi-person data.

Our contribution is four-fold. We 1) developed a multi-task network that can be jointly trained across devices. The network could simultaneously recognize a set of tap properties, including tap direction and location; 2) developed two datasets and conducted extensive evaluations that advanced the state of the art; 3) offered new perspectives on alleviating the data hurdle in ML-based HCI research; 4) established a benchmark with open-source code and datasets for off-screen tapping recognition, which will be reproducible and accessible by others.

2 RELATED WORK

This study is related to off-screen interaction, in particular, tap-based interaction, as well as multi-task neural networks [3].

2.1 Off-Screen Interaction

Back-of-device (BoD) [5, 6, 14, 16] and edge-of-device interactions [10, 16, 37, 41] have attracted much attention, however, most of them relied on specialized sensors that are not readily available on phones. For instance, BackXPRESS studied finger pressure for BoD interaction using a sandwiched smartphone [5]. Infini-Touch recognized finger location from capacitive image outside the touchscreen [16]. Some exploited small widgets to enable off-screen gesture sensing. Wearing a magnet can allow magnetometer to track the 3D finger movement [4, 24]. Adding a mirror can make camera to detect finger location on the back [33] or around the device [38]. Acoustic sensing exploits sound propagation properties to recognize BoD gesture [29], grip force [32], and contact finger part [9], and electric field sensing detects around-device gestures in a non-intrusive manner [45, 47]. Unlike using the IMU, these detection methods require additional hardware installed on the already very compact mobile devices, increasing manufacturing and material cost and potentially reducing available space to the largest possible battery.

2.2 Tap Detection from IMU Signals

Despite the line of research on detecting tapping from motion signals captured by the built-in IMU sensors on smartphones [8, 26, 39, 40], there is room for improvement. SecTap allowed users to move a cursor by tilting and select by back tapping [19]. Bezel-Tap detected tap on the edge of the devices [27]. Granell and Leiva conducted feature engineering for tap detection [8]. In contrast to the detection of tap event alone, BackPat recognized tapping in three locations based on gyroscope and microphone signals [26]. BackTap [40] and BeyondTouch [39] classified four-corner tap based on accelerometer, gyroscope, and microphone signals. These studies relied on simple features (e.g. mean, kurtosis, and skewness) and statistical methods typically based on shallow neural networks [8, 19, 26, 39, 40], thus only achieved limited precision of tap location detection.

More recently, researchers started to explore neural networks for tap location classification for PIN code inference in the field of privacy and security [17, 21]. PINlogger classified tapping on ten buttons using a single layer neural network [21]. Liang et al. applied a two-layer CNN model to estimate tap location from z-axis signal of accelerometer [17]. Although these studies yielded promising result, their network capacity was relatively small and the reported accuracy was still far from practical level.

2.3 Multi-Task Learning and Training with Auxiliary Information

To improve tap recognition, we develop a multi-input and multi-output a neural network that can estimate multiple tap properties. This section reviews two related concepts: multi-task learning [3] and learning with auxiliary information. Multi-task learning with neural network leads to a multi-branch architecture. Each branch addresses one of the recognition tasks, such as different targets in tracking [22], languages in translation [13], and head poses in gaze estimation [43]. The shared layers of all branches extract the common knowledge across tasks, whose generalizability is ensured

by the regularization effect across tasks. Multi-task models employ multiple loss functions thus more supervisions during training that can potentially helps learning. To our knowledge, no attempt has been made to apply multi-task learning for off-screen tap recognition and the multi-task architecture can conduce to the learning of each tap recognition task from their interrelationship.

We are also the first to exploit the form factor of the phone as auxiliary information, which has been shown to be beneficial for training. Zhang et al. investigated the benefit of using auxiliary information in training and found that it can improve performance in testing even without using the auxiliary information as input [42]. Stephenson et al. pointed out that conditioning on auxiliary information can achieve higher robustness than that of appending auxiliary information directly to the main features [28]. Liao et al. showed that integrating simple but essential auxiliary information can increase prediction accuracy [18].

3 RECOGNIZING TAP PROPERTIES

Moving beyond the prior work in this space, this project aims at developing IMU-based input methods that meet the requirement of practical applications, by means of deep neural network design and training. The key objective is to achieve the five recognition tasks (i.e. the five network outputs) with one network (TapNet) as shown in Table 1. Each *task* aims to recognize one *tap property*, such as direction and location. We first describe the pipeline overview, followed by the core of the method (a MIMO network). We then discuss the gating component and signal filtering for recognition.

Table 1: The five tasks of TapNet: four classification tasks to detect different tap properties, such as tap direction and location, and one regression task to estimate tap location.

Task	Paradigm	TapNet Output
Event	2-class	tap, non-tap
Direction	6-class	front, back, left, right, top, bottom
Finger part	2-class	finger pad, finger nail
Location	35-class	the region ID in a 5x7 grid
Location	regression	x- and y- ratio compared with screen width and height

3.1 Pipeline Overview

As shown in Figure 2, the system listens to the six-channel IMU (three-axis accelerometer and three-axis gyroscope) data and maintains a 150-ms data window. To avoid unnecessary down-stream computation for recognizing tap properties, a heuristic-based gating mechanism using the z-axis signal of the accelerometer is performed to reject obvious non-tap motions. If the signal passes the gating, the six-channel IMU signals within a 120-ms feature window will be concatenated into a one-dimensional feature vector. The feature window is identified and aligned across tap samples by the tap-induced peak in the z-axis signal. TapNet, a multi-input, multi-output convolutional neural network, takes this feature vector and a device vector as input. The *device vector* depicts the phone size, IMU pointing direction, and installation location relative to the

upper left corner of the device housing. TapNet then outputs the predictions of five tap properties at a time as shown in Table 1.

3.2 Design of TapNet Architecture

TapNet uses a multi-input and multi-output architecture. We use IMU signals as primary input and device vector that describes the phone form factor as auxiliary input. Using this auxiliary information helps to accommodate the difference of device form factor and achieves the cross-device training.

The output of TapNet contains multiple branches. Since different tap properties have a confound impact on IMU responses, we jointly learn the mappings from IMU signals onto these interrelated tap properties using a multi-layer CNN. In this architecture, different tap-related tasks share four convolutional layers, which extract the common shape patterns that are indicative across tap properties. Following the shared layers, there are branches of fully connected layers, each of which extracts property-specific patterns for individual task. In practice, shared layers for tap properties also means shared computation. Therefore, TapNet requires less computation and memory than multiple networks each trained for a single task.

Training over an intersection of multiple tap tasks confines the learning in a restrained feature embedding space and thus allows it to converge to a solution for all related tasks [13]. Multi-task learning allows for good alignment between feature embeddings of different tasks. Such additional guidance or supervision can potentially prevent over-fitting especially given few training samples.

3.3 One-Channel Convolutional Layers

As illustrated in Figure 2, a one-channel CNN is used for tap recognition. We exploit convolutional layers to extract shape features from IMU signals, because the convolutional filter offers temporal locality to capture signal dynamics and shared weights to reduce trainable parameters. In general, a convolutional filter in early layers describes the basic shape features, such as a peak, a valley, or a certain degree of slope. A convolutional filter in later layers, with a large receptive field, is more likely to contain shape semantics, such as a large magnitude peak or an impulse with double peaks.

Regarding the options between one-channel vs multi-channel network, we see that the one-channel network can be more efficient as it allows for filter reuse across IMU channels. Conventional methods applied multi-channel CNN to describe signal alignment across channels, and it has been widely use to handle EEG [25] and IMU [36] data. However, a large number of filters are required to depict the fine-grained signal alignment combinations of the six-channel signals, and thus increases the model training difficulty and the demand for data. In contrast, we propose to concatenate the six-channel data into a one-dimensional vector and apply one-channel convolutional layers in TapNet. By doing so, each one-channel convolutional filter can be reused to detect shape features across channels, and then rely on the fully connected layers to draw decision by associating filter activations in different parts of the signals. Dividing the shape extraction and alignment analysis conceivably mitigates the model training difficulty and achieves an efficient use of data. We also evaluated this in an experiment.

Figure 2: Accelerometer and gyroscope signals are processed in the Gating component and passed to TapNet if met set criteria. Using phone form factor (embedded in the device vector) as auxiliary information, TapNet jointly recognizes multiple tap properties, including tap event, direction, finger part, and location.

3.4 Signal Gating for Neural Network

Although TapNet is rather lightweight compared with vision-based convolutional networks, performing recognition per frame generates unnecessary overheads. Our observation, as well as previous findings [17], suggests that a peak or valley in the z-axis output is a necessary (high recall) but not sufficient (low precision) signal to a tapping event on the device. As such, we use the z-axis signal of accelerometer as a *gating signal* to the CNN in order to minimize the amount of computation (hence power consumption) on the mobile device. Gating signal alone may not be enough for accurate recognition of tap properties, it is sufficient for tap-like (high recall low precision) event detection, and thus qualified as gating for the CNN recognition. Therefore, we only perform CNN recognition when we observe a tap-like signal from the gating signal.

Specifically, we first perform a simple, linear complexity peak detection on the gating signal at the per-frame level (Figure 2 the pink region). This peak detection yields a set of peaks and valleys and their corresponding timestamps. We define a *tap-like signal* as the one has an impulse that contains at least one peak, and an *impulse* as a group of peaks between each pair the interval is less than the threshold, T_v , which is empirically set to 80 ms. We also apply a magnitude threshold on the peak and valley to control the sensitivity of the gating component.

3.5 Sensor Signal and Temporal Alignment

Sensor system on smartphones provides motion data in a number of formats, including the raw signal, its gravity excluded counterpart, and rotation vector. As we aim to detect tap event and identify tap properties in an orientation-invariant manner, we leverage the raw accelerometer and gyroscope data and compute their first-order derivative. Therefore, these signals represent the change of force on the housing of the phone and the resulting change of rotation velocity, and they stay at zero when the phone is stationary. In addition to orientation invariance, we also see that the first-order derivatives accelerometer and gyroscope signals have high shape

similarity. As such, using this representation conduces to the reuse of convolutional filters and alleviates network training difficulty.

All tap feature vectors are temporally aligned. Specifically, The first major peak in the accelerometer z-axis is aligned to a designated frame (=106th) in the feature vector, and this determines the feature window location. Then 50 sensor samples (~ 120 ms) from each IMU channel are concatenated to form a 300-element feature vector for each tap-like event. Such temporal alignment allows the network to focus on the shape differences induced by a tap.

4 DATASET DEVELOPMENT

This section describes two datasets: the one-person dataset (135,260 samples) was used for TapNet training only (not testing) and the multi-person dataset (38,545 samples) was for performance testing and user generalizability tuning. Our hypothesis is that *a well-designed data collection protocol can ensure training data diversity even from one person and be sufficient for some tap recognition tasks*. When we need to improve the model, it is relatively easy to enlarge the one-person training set, evaluate on the test set, and repeat this process in an iterative fashion. This is more efficient than collecting additional multi-person data.

4.1 One-Person Training Dataset

We used a single researcher training strategy. One of us produced the entire training data following a comprehensive data collection protocol, which aims to cover the diversity in real use. Tap data were collected through multiple sessions on Phone A and Phone B². Each session aims to collect tap in a specific condition described by a set of characteristics of the phone, grip gestures, and the tapping itself (see Table 2). Visual indicator was presented on the screen to guide the experimenter. The experimenter can examine the finger alignment with the target location by turning the phone back and forth to ensure the annotation correctness. Please refer to the supplemental file for a detailed description.

²Phone A = Pixel 3 and Phone B = Pixel 3 XL

Table 2: The one-person dataset contains samples in different conditions of phone, grip gestures, and tapping actions.

Condition	#Opt.	Option
Phone size	2	Phone A (5.5"), Phone B (6.3")
Phone case	2	with case, without case
Orientation	4	portrait, portrait downward, landscape +90°, landscape -90°
Grip and tap manner	7	tapping with one hand (the four combinations of holding in left/right hand and tapping in left/right hand), both hands with thumbs, tapping on the phone which is lying on a solid surface, and on a soft one.
Grip force	3	phone rests on the palm, grasp the phone with normal force, grasp with strong force.
Grip location	5	bottom, mid-bottom, middle, mid-top, top
Grip gesture	4	grasping the phone against to the palm with one, two, three, and four fingers
Thumb gesture	2	resting on the screen edge, hovering above the screen
Tap event	2	tap, non-tap
Tap finger part	2	finger pad, finger nail
Tap direction	6	front, back, left, right, top, bottom
Tap location	35	5x7 grid on the housing of the phone
Tap force	5	extremely gentle, gentle, normal, strong, extremely strong

The advantages of this strategy are three-fold: 1) it is low-cost and convenient for rapid development iteration, 2) it produces systematic and diverse data, and 3) the data quality is manageable and assessable. This strategy allows to use the researcher’s intuition of the problem space to push the recognition envelope to the fullest degree, for example, by including the training data with different grip gestures and forces. The risks of this strategy include the potential lack of diverse tapping patterns that a single person can not fully represent as well as the impact of the biomechanics. We mitigate the risks by evaluating on a separate multi-person dataset.

We repeated data collection under the same condition to mitigate bias. In the end, the one-person data collection took over 30 sessions (around half an hour for each) over 69 days. In total, the one-person training dataset contains 109,200 tapping actions, including 45,500 front taps, 48,450 back taps, 4,020 left taps, 4,020 right taps, 3,610 top taps, 3,600 bottom taps. Among these tapping actions, 94,764 were collected from a Phone A and 14,440 from a Phone B, and 85.7% of them are tapping with finger pad and the rest are with finger nail. We also collected non-tap motions (26,060 in total) in a number of scenarios: grasping (4,700) the phone, rubbing the phone case (6,000), releasing a tap (6,160), knocking on the phone placing surface (4,400), shaking and moving with the phone in pocket and in bag (4,800). The usage of this dataset (135,260 samples including tap and non-tap) was solely in training (or “teaching”). None was used in any testing to ensure the validity and generalizability of the efficacy measures reported later.

Table 3: The multi-person dataset contains samples in different conditions of phone, grip gestures, and tapping actions. The number in the parenthesis shows the data percentage.

Condition	Option
Phone size	Phone A (61%), Phone B (39%)
Phone condition	with case (58%), without case (42%)
Grip and tap manner	one-handed/portrait (thumb & index), two-handed/portrait (thumbs & indexes), two-handed/landscape (thumbs & indexes), two-handed/portrait (index)
Handedness	right (90%), left (10%)
Hand size	small (32%), medium (48%), large (20%)
Tap direction	front (54%), back (28%), left (4%), right (4%), top (5%), bottom (5%)
Tap location	comfort range in the 5x7 grid of the screen/back

Figure 3: Four common phone grip gestures investigated in the multi-person dataset. Only natural tapping actions with these grip gestures were collected for testing.

4.2 Multi-Person Testing Dataset

To evaluate TapNet performance, we collected a multi-person ($n=31$) dataset. We focused on collecting the natural tapping actions to simulate tapping data distribution in real use, in particular, with the common finger (thumb and index), in the finger comfort range, and for four most common one-handed and two-handed gestures [11, 15] as shown in Figure 3.

4.2.1 Data collection design. The four grip gestures we studied are in a combination conditions of phone, grip gestures, and tapping actions as shown in Table 3. To keep data collection natural, we avoid continuous and extensive tapping. For every condition, participants were asked to tap for five to ten times and take rest at various intervals to avoid fatigue. In addition, data collection was informed by research conducted on the most common hand grips and finger placements [11, 15]. Unrealistic reach for targets on the back surface of the phone were excluded.

4.2.2 Apparatus. We used a 5.5" Phone A and a 6.3" Phone B throughout the collection. To provide the most natural feeling, we assigned phone size in the collection based on participants’ personal phone size. To achieve a balance ratio, we selected participants according to their personal phone size in recruitment.

4.2.3 Participants. We recruited 31 participants (13 female, age: 18-54). We marked down a number of factors that may affect tapping signal responses, including finger length, finger nail length,

hand size, and handedness. During data collection, the experiment interface would adapt to the handedness of individual participant to ensure the comfort range was valid.

Overall, the multi-person dataset contains 38,545 taps, including 20,615 taps from front, 10,505 back, 1,705 left, 1,705 right, 1,975 top, and 2,040 bottom, among which 58.3% was collected on phones with rubber cases, and the rest without.

5 MACHINE LEARNING EXPERIMENTS

This section evaluates the overall performance of TapNet, followed by the comparison between one- and multi- channel CNN as well as ablation studies on the MIMO network design.

We measured classification performance by F1 score and regression by mean absolute error (MAE) and r^2 score. F1 score is an weighted average metric of precision and recall, ranging from [0, 1].

The MAE of tap location, i.e. $\sqrt{[(g_x - t_x)/w]^2 + [(g_y - t_y)/h]^2}$, is computed by the Euclidean distance between ground truth location (g_x, g_y) and TapNet output (t_x, t_y) normalized by screen width w and height h . Its range is [0, 1.414], and a baseline (always predicting the center) is 0.707. Similar to standard deviation, r^2 measures how well the regression line fits the data with a range of [0, 1].

In the implementation, training to recognize the presence of tap event requires tap and non-tap data (e.g. phone shaking and rubbing motions). As non-tap data does not have annotations about tap finger part, direction, and location, it cannot be used to train the rest of the network. As such, the two parts of the network (tap event v.s. the rest of the tap properties) were trained in turn. The tap event branch was trained once after every ten epochs of training the rest of the tap property branches.

We applied ReLU and batch normalization after each convolutional layer and used Adam optimizer with a learning rate of $1e-4$ and a momentum decay ($1e-6$). Each model was trained with sufficient number of epochs until the validation (5% data) loss converges, and training was done using Tensorflow [1] on a 4G memory GPU.

The datasets were collected on Phone A and Phone B that ran on Android 9, and the IMU sampling rate was approximately 416Hz. In run time, it took 0.56 ms on the Phone B main CPU to finish one TapNet inference. A simplified TapNet can also be run in real time (9ms) on the embedding DSP using Tensorflow Lite for microcontroller [31], but this is beyond the scope of this paper. The major latency (105 ms) of the whole pipeline lies in waiting to observe the complete tap signal in the feature window.

Due to space limitation, we have put some of the implementation evaluations into the supplemental file. Interesting findings include: 1) high sensor sampling rate is crucial to tap recognition accuracy; 2) the use of gyroscope in our task configuration is important; and 3) data augmentation by temporally shifting the samples conduces to the performance gain but scaling the samples does not.

5.1 Performance on Tap Recognition Tasks

This section evaluates TapNet’s improvement over prior art and user generalizability when training on one-person data.

5.1.1 Improving the state of the art. We implemented and trained four ML algorithms, including TapNet, to compare their relative performance. The overall performance measured by the weighted

average F1 score and MAE across participants and devices are shown in Table 4. Support Vector Machine (SVM) represents a line of studies [19, 39] that used traditional machine learning methods. TinyCNN is a replication of Liang et al.’s two-layer CNN [17]. SISO refers to a truncated single-input and single-output (SISO) TapNet. It uses the same training configuration (Adam optimizer and learning rate) as the MIMO TapNet, and it gives performance reference if TapNet is configured for a single task. All the methods were trained on the one-person dataset and tested on the multi-person dataset.

Table 4: Weighted average F1 score of four classification tasks (no color shading) as well as MAE and r^2 score of the tap location regression task (purple shading). SVM and TinyCNN are our re-implementation of related works [17, 19, 39]. SISO TapNet is the truncated single-input and single-output variant. TapNet is the proposed MIMO method. Note that we have built one model per task for the single-output models. In contrast, TapNet is a multi-task network that gives predictions for all five tasks at a time. Finger part, direction, and location are not conditional on the event detection. The numbers are the higher the better for F1 and r^2 scores, while lower the better for MAE.

Method	F1 score				MAE (r^2)
	2-class Event	2-class Finger part	6-class Direction	35-class Location	Location regression
SVM [19, 39]	.73	.93	.55	.13	.20 (.28)
TinyCNN [17]	.89	.93	.47	.03	.40 (-2.4)
SISO TapNet	.88	.89	.82	.30	.19 (.12)
TapNet	.87	.94	.85	.34	.14 (.49)

Overall, MIMO TapNet significantly outperforms the prior art [17, 19, 39]. Compared with the best performance among SVM and TinyCNN, TapNet achieved considerable improvements on tap direction (by 51.0%) and location (classification: 161.5% and regression: 30%). The SISO TapNet variant also outperforms related works by a marked margin. TapNet generally outperforms its SISO variant, implying that the MIMO architecture also contributes to performance improvement in addition to the computation and memory benefits.

A close examination of the 35-class tap location classification reveals that TapNet can be usable despite its seemingly limited F1 score (0.34). Figure 4(a) shows the normalized confusion matrix of tap location classification. Neighboring region above or below the target in a 5x7 grid has an index offset of five. The three parallel lines along the diagonal with a five-cell offset in the confusion matrix indicates TapNet either predicts correctly or to nearby regions (see Figure 4(c)). This is in good agreement with the regression results (MAE: .14; ~10% of the screen diagonal). Such location error is similar to the distance between two icons on the phone home screen (see the supplemental video figure). It thus indicates that IMU-based tap location recognition can be viable and useful in situations which does not require very high resolution and when capacitive sensing is inadequate (e.g. wearing gloves or under water).

Figure 4: (a) Normalized confusion matrix of the 35-class location classification while training on one person and testing on multiple. (b) An enlarged part of the confusion matrix. (c) The region IDs in a five-by-seven grid. Each region refers to a cell in the grid. Neighboring areas above or below the target has an index offset of five. Predicting to the nearby areas therefore leads to three parallel lines along the diagonal in the confusion matrix.

However, the performance difference on simple tasks (event and finger part classification) among different methods is marginal. This suggests that the minimal model capacity can be task-dependent; simple tasks may not be benefited from a model capacity increase, but the more complicated tasks can be.

5.1.2 Achieving person generalizability with one-person data. We hypothesized that training on diverse but one-person data can still achieve generalizability for unseen users in some tap recognition tasks. To verify this, we performed evaluation in two paradigms:

- *One-to-n-participant* evaluation: TapNet was trained on the one-person dataset and tested on the multi-person dataset. This is the same as the previous evaluation.
- *Leave-one-participant-out* cross-validation: TapNet was pre-trained on the one-person dataset, fine tuned on $n-1$ participants from the multi-person dataset, and tested on the leftout participant in an iterative fashion.

By comparing the one-to-n-participant with the leave-one-out evaluation (see Table 5), we see that training on a one-person data can still be effective. As expected, fine tuning on the $n-1$ participants further improved the performance, but some of the improvements are relatively moderate, for example, by 1% for tap event and 3.2% for finger part classification. This corroborates the hypothesis that training on the diverse data even from a single person can achieve viable user generalizability for tasks, such as event (F1:.92), finger part (F1:.93), and direction classification (F1:.85). That said, we also learned that tap location recognition relies on very subtle IMU responses, which may relate to personal biomechanics and are hard to simulate by a single person.

5.2 Ablation Studies: Multi- Input and Output

This section presents the ablation studies on the TapNet architecture: the use of multi-task learning (i.e. multi-output) and cross-device training with auxiliary information (i.e. multi-input). We evaluated on tap direction classification as a representative task.

Table 5: Performance comparison between the one-to-n-participant and leave-one-participant-out evaluation. The columns with no color shading are classification tasks, and the one with purple shading is a regression task.

Evaluation Paradigm	F1 score				MAE (r^2)
	2-class Event	2-class Finger part	6-class Direction	35-class Location	Location regression
one-to-n	.92	.93	.85	.42	.15 (.52)
leave-one-out	.93	.96	.92	.54	.11 (.73)

5.2.1 Multi-task learning helps training with limited data. We first present the evaluation on multi-task learning. We compared against the single-output counterpart (SISO; a truncated TapNet variant) and recent related works, including a SVM [19, 39] and a tiny CNN model [17]. Limited by the tap samples on Phone B (~15K) in the training dataset, we evaluated model performance averaged over two devices with incremental training samples from 1K to 15K. Figure 5 shows the comparison results.

Multi-task learning contributes to tap direction classification, especially with a small amount (no more than 3K) of training samples. Nevertheless, with sufficient training samples (over 3K) performance of SIMO and SISO TapNets flats out with a comparable F1 score (SIMO:.81, SISO:.80). Although the performance improvement of SIMO over SISO is modest in this specific task, its advantages of reusing computation (for the shared layers) across recognition tasks and memory saving to avoid running multiple models on small processing units are essential. One the other hand, both SIMO and SISO TapNets can considerably outperform TinyCNN (purple dashed) starting from 3K training samples, and outperform SVM (blue dashed) starting from 6K. Taken together, a multi-output TapNet is favorable over the state of art [17, 19, 39].

Figure 5: Weighted average F1 score of tap direction classification using multi-task learning (SIMO TapNet) and single-task learning (SISO TapNet). TapNet (either SIMO or its truncated variant) can considerably outperform a SVM [19, 39] and a tiny CNN model [17]. TapNet with multi-task learning shows advantage given small training data (e.g. 3K samples).

Figure 6: Performance comparison of tap direction classification with auxiliary information (A+B) against fine tuning (B->A, A->B), and training on device-specific data (A, B) alone. X-axis shows the number of training samples per device and y-axis indicates the weighted average F1 score. The jointly trained TapNet yields comparable performance to that of the fine tuned models on specific devices and outperforms the models trained on device-specific data alone.

5.2.2 Cross-device training utilizes data efficiently. To investigate cross-device model training, we evaluated model performance with incremental training samples from 1K to 15K, as previous experiments suggest that TapNet converges with around 15K training samples. Note that in the case of training with 15K samples, TapNet was jointly trained on 7.5K samples from Phone A and another 7.5K samples from Phone B. Similarly, its counterparts were pre-trained on 7.5K samples from one device and fine tuned on the other. As baselines, we also evaluated training on the 7.5K device-specific samples alone for each device.

Figure 6 gives the performance comparison. A+B represents TapNet jointly trained on cross-device data; B->A and A->B denote tuning the pre-trained model from one device on the other; and A and B indicate models trained on device-specific data. The overall performance of different models increases with training samples and they flats out after 4.5K samples per device. More interestingly, cross-device training with sensor location as auxiliary information (A+B) can approach the performance upper bound earlier (with 1.5K samples per device) than the rest of the models.

5.3 Signal Alignment in One-Channel CNN

To verify the efficacy of one-channel CNN for tap recognition, we performed a comparison on tap direction classification, which is a representative task and demonstrated to require subtle signal alignment across channels. As input data dimension affects the number of trainable parameters (i.e. model capacity) of a specific network architecture, we evaluated similar architectures with commensurate number of trainable parameters, so as to compare models with comparable capacity. This experiment was conducted on the Phone A data using the one-to-n evaluation paradigm.

Figure 7 shows the performance comparison of one-channel CNN against its multi-channel counterparts. The x-axis shows the number of training samples, and the y-axis the weighted average of

Figure 7: Performance comparison of one-channel CNN (TapNet) against its six-channel counterpart with two levels of model capacity (large models: 163K, 144K trainable parameters; small models: 11K and 9K). These are the average results over three runs to cancel out training randomness. The one-channel CNN with 11K parameters can achieves comparable performance as that of the six-channel CNN with a larger number of parameter (144K). We tuned the number of filters and maintained its ratios across layers to roughly match the total number of trainable parameters between the one-channel and six-channel models.

F1 score. The solid lines denote the performance of large-capacity models with 163K (one-channel) and 144K (six-channel) trainable parameters, while the dashed lines those of small-capacity models with 11K (one-channel) and 9K (six-channel) trainable parameters. Most importantly, small-capacity one-channel CNN (purple dashed) and large-capacity six channel CNN (blue) have almost equivalent performance with 6K-12K training samples. Further, comparing the solid and dashed lines, large-capacity models generally outperforms small-capacity models with similar architecture and input format. Taken together, we conclude that the one-channel CNN can address across-channel signal alignment more efficiently than its multi-channel counterpart for tap property recognition.

6 DISCUSSION

We designed, developed, and evaluated a set of deep learning methods for on-device IMU signal based off-screen input, in particular TapNet, a multi-task network that allows for cross-device training with phone form factor as auxiliary information and joint prediction of tap-related tasks, including tap event, direction, finger part, location classification, and the regression of tap location. This architecture not only shares knowledge across tap properties and devices during training, but also shares computation and memory at run time. Some of the TapNet building blocks are not novel in fields such as computer vision, but to bring them into building an interaction gesture to a practical performance level required original research. TapNet achieved a marked improvement over the prior art [17, 19, 39] on tap direction (by 51.0%) and location (161.5%) classifications as well as tap location regression (30%).

We also discovered encouraging generalizability across users when the model was trained (or "taught") by a one person. We have

Figure 8: Conceptual relation between performance and training strategy with one-person and multi-person data.

tested TapNet on those who had not been in any set of the data collection and encouraged them to explore the effects of TapNet freely on their own device in daily use (for such functions as taking screenshot). Their experience matched with the leave-one-participant-out test results reported here. This high generalizability is probably because the inter-class differences (e.g. finger pad vs. nail in the finger part classification task) are generally much greater than the inter-person differences, such that the personal biomechanics only imposes a neglectable effect.

The other advantage of our approach is that an expert design could intentionally push the variations of an intentional tap gesture in terms of speed, strength, angle, and hand posture. However, it is possible certain recognition tasks, such as higher resolution tap location classification, could demand more fine-grained information only available in person-specific data. Training on incremental one-person data can increase performance up to a certain level and then plateaus (see Figure 8). Further adapting the model to multi-person data may teach the model to understand the artifacts of personal biomechanics, and thus further improves the performance until its next plateau. This performance gain is from knowing how much people can vary. To reach the ideal performance, it becomes a must to know the person-specific information, i.e. how exactly the target user acts. In practice, especially for learning-based gesture studies, it can be beneficial to identify the curve in Figure 8, and then decide the needed number of training participants for specific tasks.

TapNet opens up new interaction opportunities such as one-handed interaction that uses off-screen tapping or designated tap gestures. TapNet is robust to phone case and fabric, and thus can be used for wearable interaction without specialized smart fabrics [7]. This study also offers new potential for other research domains, including biometrics. For instance, it is possible to improve continuous and passive authentication [2, 35, 46] by using tap properties that contain biometric information, i.e. with a clear gap between the first two plateaus (orange area) in Figure 8.

Although evaluation results demonstrated that TapNet benefited from cross-device training, we do not expect the current TapNet trained on these two devices would directly apply to unseen devices without further training, i.e. a cross-device model. However, the joint training architecture has been shown effective and this is a

Figure 9: Four example use cases enabled by TapNet. AssistiveTap uses back tap and tilting for one-handed interaction. ExplorativeTap introduces a new two-sided interaction paradigm for users with visual impairments. Interactive wallpapers enables interaction with UI background objects. Inertial Touch senses tapping by force and angle changes and provides auxiliary information for capacitive sensing.

promising step toward a device adaptive model. We have set up the infrastructure and we plan to further investigate along this line by adding new devices and also opensource our implementation.

7 HCI APPLICATIONS OF TAPNET

Without adding new sensor hardware, TapNet enables many new input possibilities on smartphones. In addition to providing shortcuts to quick activations of apps or functions such as camera and screenshot, this section sketches out four applications enabled by TapNet’s capability of measuring multiple proprieties in addition to the presence of tap event, such as direction and location. Please refer to our supplemental video for these applications in action.

7.1 AssistiveTap

AssistiveTap allows users to complete a number of system interactions using back tap and tilting. Users can perform a back double tap to invoke the AssistiveTap interface (see Figure 9a), then tilt the phone (based on the same IMU signal input to TapNet) to select gesture, and back tap to perform the selected gesture (e.g. ‘back’ gesture, scrolling, app switch). In other words, AssistiveTap provides the back-of-device alternatives for the most commonly used on-screen gestures, by using the gesture combination of back tap and tilting. It can be beneficial in situations, where one-handed interaction is preferable or even required.

7.2 ExplorativeTap

ExplorativeTap is a two-handed interaction method that exploits signals from both the front (touch) and back of the phone tap (see Figure 9b). Users can use the exploration finger on the screen to glide over the on-screen objects, hear them, and perform a back tap to confirm selection. This is an improvement for the visual impairment accessibility modes, such as Voice Over on iOS and Talk Back on Android. These accessibility modes occupy the common on-screen gestures (e.g. swipe and touch), and their users need to learn a more complicated system navigation gestures. In contrast, ExplorativeTap starts the accessibility mode by a double back tap, selects object by a single back tap, and exits by lifting the exploration finger. It, therefore, solves the conflicts with system navigation gestures, and thus can be helpful for users with low

vision and print disability, who need quick and temporary access to the accessibility mode.

7.3 Interactive Wallpaper

Off-screen tap recognition makes it possible to interact with objects living in different interface layers, such as background targets that do not react to on-screen touch. For instance, TapNet can enable users to interact with flashcards or news feeds shown in the wallpaper (see Figure 9c). They can back tap to change the flashcard or news feeds, or edge tap to switch a different set of them, and these can be done even on the lock screen. It enables the user to utilize the bits and pieces of time for their favorite spare time activity.

7.4 Inertial Touch

Inertial Touch estimates tap force from the IMU signals and tap location from the TapNet output. We can define Inertial Touch as a *fast touch event with a strong tapping momentum*. As TapNet estimates tap location from force and angle changes, it can function even when capacitive sensing fails (Figure 9d). Common use situations include when users are wearing gloves, having long fingernails, or when there are waterdrops on the touchscreen.

7.5 Summary of Use Cases

The aforementioned use cases can be particularly useful in challenging situations when one-handed and non-contact interactions are preferable or required. For example, AssistiveTap exploits back tap and titling to partially address the issue of limited thumb reaching area during one-handed interaction. ExplorativeTap allows for the coordination between on-screen and off-screen interactions and thus saves the need for learning additional set of on-screen gestures. Inertial Touch offers auxiliary impact information beyond capacitive sensing. These are just a few application examples of tapping into TapNet's ability to detect multiple tap properties.

8 CONCLUSION

This paper presents the design, training, implementation and application of TapNet for off-screen mobile input. TapNet employs a multi-task convolutional neural network with motion signals as primary input and phone form factor as auxiliary information. It allows for joint learning on data across devices and simultaneous estimation of multiple tap properties. In comparison to many alternatives, this neural network architecture worked the best towards our goal of increasing the UI design space of off-screen interactions. To train and optimize this and other alternative ML models, we developed a one-person dataset for training and a multi-person dataset for testing. The evaluation results show that 1) TapNet significantly outperformed the state of the art especially in difficult recognition tasks such as tap location estimation; 2) multi-task learning is more data efficient, showing greater advantage particularly with a limited amount of training data; 3) cross-device training with phone form factor increases the efficiency of data utilization; and 4) one-channel CNN can achieve cross-channel signal alignment more efficiently than its multi-channel counterpart. We verified the hypothesis that training on the one-person data can still generalize well across users if the diversity of the training data can be ensured. This sheds light on the conceptual relation between model performance and

ML training strategy in IMU-based input systems. We demonstrated that many new interaction use cases could be enabled by TapNet. Taken together, the TapNet project made significant progress towards practically enabling and enlarging the off-screen interaction design space by deep learning from on-device IMU signals, establishing new benchmarks with reproducible results, datasets and codebase.

REFERENCES

- [1] 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 *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*. USENIX Association, USA, 265–283.
- [2] Sara Amini, Vahid Noroozi, Sara Bahaadini, S Yu Philip, and Chris Kanich. 2018. DeepFP: A Deep Learning Framework For User Fingerprinting via Mobile Motion Sensors. In *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, USA, 84–91.
- [3] Rich Caruana. 1997. Multitask learning. *Machine learning* 28, 1 (1997), 41–75.
- [4] Ke-Yu Chen, Shwetak N Patel, and Sean Keller. 2016. Finexus: Tracking precise motions of multiple fingertips using magnetic sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, USA, 1504–1514.
- [5] Christian Corsten, Bjoern Daehlmann, Simon Voelker, and Jan Borchers. 2017. BackXPRESS: Using back-of-device finger pressure to augment touchscreen input on smartphones. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, USA, 4654–4666.
- [6] Alexander De Luca, Emanuel Von Zeischwitz, Ngo Dieu Huong Nguyen, Max-Emanuel Maurer, Elisa Rubegni, Marcello Paolo Scipioni, and Marc Langheinrich. 2013. Back-of-device authentication on smartphones. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, USA, 2389–2398.
- [7] David Dobbeltstein, Christian Winkler, Gabriel Haas, and Enrico Rukzio. 2017. PocketThumb: a Wearable Dual-Sided Touch Interface for Cursor-based Control of Smart-Eyewear. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 2 (2017), 9.
- [8] Emilio Granell and Luis A Leiva. 2016. Less is more: Efficient back-of-device tap input detection using built-in smartphone sensors. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces*. ACM, USA, 5–11.
- [9] Chris Harrison, Julia Schwarz, and Scott E Hudson. 2011. TapSense: enhancing finger interaction on touch surfaces. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, USA, 627–636.
- [10] David Holman, Andreas Hollatz, Amartya Banerjee, and Roel Vertegaal. 2013. Unifone: designing for auxiliary finger input in one-handed mobile interactions. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction*. ACM, USA, 177–184.
- [11] Steven Hooper. 2013. How do users really hold mobile devices. *Uxmatters* (<http://www.uxmatter.com>). Published: February 18 (2013), 2327–4662.
- [12] Michael Xuelin Huang, Jiajia Li, Grace Ngai, and Hong Va Leong. 2018. Quick Bootstrapping of a Personalized Gaze Model from Real-Use Interactions. *ACM Transactions on Intelligent Systems and Technology (TIST)* 9, 4 (2018), 43.
- [13] Lukasz Kaiser, Aidan N Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. 2017. One Model To Learn Them All. *arXiv preprint arXiv:1706.05137* 0, 0 (2017), 1–10.
- [14] Huy Viet Le, Patrick Bader, Thomas Kosch, and Niels Henze. 2016. Investigating Screen Shifting Techniques to Improve One-Handed Smartphone Usage. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. ACM, USA, 27.
- [15] Huy Viet Le, Sven Mayer, Patrick Bader, and Niels Henze. 2018. Fingers' Range and Comfortable Area for One-Handed Smartphone Interaction Beyond the Touchscreen. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, USA, 31.
- [16] Huy Viet Le, Sven Mayer, and Niels Henze. 2018. InfiniTouch: Finger-aware interaction on fully touch sensitive smartphones. In *The 31st Annual ACM Symposium on User Interface Software and Technology*. ACM, USA, 779–792.
- [17] Yi Liang, Zhipeng Cai, Jiguo Yu, Qilong Han, and Yingshu Li. 2018. Deep learning based inference of private information using embedded sensors in smart devices. *IEEE Network* 32, 4 (2018), 8–14.
- [18] Binbing Liao, Jingqing Zhang, Chao Wu, Douglas McIlwraith, Tong Chen, Shengwen Yang, Yike Guo, and Fei Wu. 2018. Deep sequence learning with auxiliary information for traffic prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, USA, 537–546.
- [19] Zhen Ling, Junzhou Luo, Yaowen Liu, Ming Yang, Kui Wu, and Xinwen Fu. 2018. SecTap: Secure Back of Device Input System for Mobile Devices. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, USA, 1520–1528.

- [20] William McGrath and Yang Li. 2014. Detecting tapping motion on the side of mobile devices by probabilistically combining hand postures. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, USA, 215–219.
- [21] Maryam Mehrnezhad, Ehsan Toreini, Siamak F Shahandashti, and Feng Hao. 2018. Stealing PINs via mobile sensors: actual risk versus user perception. *International Journal of Information Security* 17, 3 (2018), 291–313.
- [22] Hyeonseob Nam and Bohyung Han. 2016. Learning multi-domain convolutional neural networks for visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, USA, 4293–4302.
- [23] Philip Quinn, Seungyon Claire Lee, Melissa Barnhart, and Shumin Zhai. 2019. Active Edge: Designing Squeeze Gestures for the Google Pixel 2. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, USA, 274.
- [24] Gabriel Reyes, Jason Wu, Nikita Juneja, Maxim Goldshtein, W Keith Edwards, Gregory D Abowd, and Thad Starner. 2018. SynchroWatch: One-handed synchronous smartwatch gestures using correlation and magnetic sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 4 (2018), 158.
- [25] Robin Tibor Schirmer, Jost Tobias Springenberg, Lukas Dominique Josef Fiederer, Martin Glasstetter, Katharina Eggensperger, Michael Tangermann, Frank Hutter, Wolfram Burgard, and Tonio Ball. 2017. Deep learning with convolutional neural networks for EEG decoding and visualization. *Human brain mapping* 38, 11 (2017), 5391–5420.
- [26] Karsten Seipp and Kate Devlin. 2014. BackPat: one-handed off-screen patting gestures. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. ACM, USA, 77–80.
- [27] Marcos Serrano, Eric Lecolinet, and Yves Guiard. 2013. Bezel-Tap gestures: quick activation of commands from sleep mode on tablets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, USA, 3027–3036.
- [28] Todd A Stephenson, Mathew Magimai Doss, and Hervé Bourlard. 2004. Speech recognition with auxiliary information. *IEEE transactions on speech and audio processing* 12, 3 (2004), 189–203.
- [29] Ke Sun, Ting Zhao, Wei Wang, and Lei Xie. 2018. Vskin: Sensing touch gestures on surfaces of mobile devices using acoustic signals. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. ACM, USA, 591–605.
- [30] Apple Support. 2019. Use Siri on all your Apple devices. <https://support.apple.com/en-us/HT204389>.
- [31] TensorFlow. 2020. TensorFlow Lite for Microcontrollers. <https://www.tensorflow.org/lite/microcontrollers>.
- [32] Yu-Chih Tung and Kang G Shin. 2016. Expansion of human-phone interface by sensing structure-borne sound propagation. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, USA, 277–289.
- [33] Pui Chung Wong, Hongbo Fu, and Kening Zhu. 2016. Back-Mirror: back-of-device one-handed interaction on smartphones. In *SIGGRAPH ASIA 2016 Mobile Graphics and Interactive Applications*. ACM, USA, 10.
- [34] xda-developers. 2020. Tap, Tap brings iOS 14/Android 11's Back Tap gesture to any Android device. <https://www.xda-developers.com/tap-tap-brings-ios-14-android-11-back-tap-gesture-any-android-device/>.
- [35] Hui Xu, Yangfan Zhou, and Michael R Lyu. 2014. Towards continuous and passive authentication via touch biometrics: An experimental study on smartphones. In *10th Symposium On Usable Privacy and Security (SOUPS) 2014*. ACM, USA, 187–198.
- [36] Jianbo Yang, Minh Nhut Nguyen, Phyo Phyo San, Xiao Li Li, and Shonali Krishnaswamy. 2015. Deep convolutional neural networks on multichannel time series for human activity recognition. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*. AAAI press, USA, 3995–4001.
- [37] Hui-Shyong Yeo, Juyoung Lee, Andrea Bianchi, and Aaron Quigley. 2016. Sidetap & slingshot gestures on unmodified smartwatches. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, USA, 189–190.
- [38] Chun Yu, Xiaoying Wei, Shubh Vachher, Yue Qin, Chen Liang, Yueting Weng, Yizheng Gu, and Yuanchun Shi. 2019. HandSee: Enabling Full Hand Interaction on Smartphone with Front Camera-based Stereo Vision. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, USA, 705.
- [39] Cheng Zhang, Anhong Guo, Dingtian Zhang, Yang Li, Caleb Southern, Rosa I Arriaga, and Gregory D Abowd. 2016. Beyond the touchscreen: an exploration of extending interactions on commodity smartphones. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 6, 2 (2016), 16.
- [40] Cheng Zhang, Aman Parnami, Caleb Southern, Edison Thomaz, Gabriel Reyes, Rosa Arriaga, and Gregory D Abowd. 2013. BackTap: robust four-point tapping on the back of an off-the-shelf smartphone. In *Proceedings of the adjunct publication of the 26th annual ACM symposium on User interface software and technology*. ACM, USA, 111–112.
- [41] Cheng Zhang, Junrui Yang, Caleb Southern, Thad E Starner, and Gregory D Abowd. 2016. WatchOut: extending interactions on a smartwatch with inertial sensing. In *Proceedings of the 2016 ACM International Symposium on Wearable Computers*. ACM, USA, 136–143.
- [42] Qilin Zhang, Gang Hua, Wei Liu, Zicheng Liu, and Zhengyou Zhang. 2014. Can visual recognition benefit from auxiliary information in training?. In *Asian Conference on Computer Vision*. Springer, USA, 65–80.
- [43] Xucong Zhang, Michael Xuelin Huang, Yusuke Sugano, and Andreas Bulling. 2018. Training person-specific gaze estimators from user interactions with multiple devices. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, USA, 624.
- [44] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. 2017. Mpiigaze: Real-world dataset and deep appearance-based gaze estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 1 (2017), 162–175.
- [45] Yang Zhang, Gierad Laput, and Chris Harrison. 2017. Electrick: Low-cost touch sensing using electric field tomography. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, USA, 1–14.
- [46] Nan Zheng, Kun Bai, Hai Huang, and Haining Wang. 2014. You are how you touch: User verification on smartphones via tapping behaviors. In *2014 IEEE 22nd International Conference on Network Protocols*. IEEE, USA, 221–232.
- [47] Junhan Zhou, Yang Zhang, Gierad Laput, and Chris Harrison. 2016. AuraSense: enabling expressive around-smartwatch interactions with electric field sensing. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, USA, 81–86.