SelfSync: Exploring Self-Synchronous Body-Based Hotword Gestures for Initiating Interaction

Juyoung Lee
KAIST
Daejeon, Republic of Korea
ejuyoung@kaist.ac.kr

Shaurye Aggarwal
Georgia Institute of Technology
Atlanta, GA, United States
saggarwal46@gatech.edu

Jason Wu
Carnegie Mellon University
Pittsburgh, PA, United States
jsonwu@cmu.edu

Thad Starner
Georgia Institute of Technology
Atlanta, GA, United States
thad@gatech.edu

Woontack Woo
KAIST
Daejeon, Republic of Korea
wwoo@kaist.ac.kr

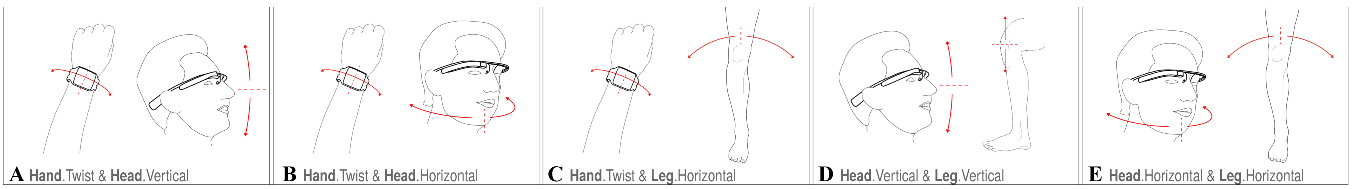

Figure 1: SelfSync gestures can unlock a gesture interface or immediately trigger an action through coordinated body movements such as hand and head gestures (A, B), hand and leg gestures (C), and head and leg gestures (D, E).

ABSTRACT

SelfSync enables rapid, robust initiation of a gesture interface using synchronized movement of different body parts. SelfSync is the gestural equivalent of a hotword such as OK-Google in a speech interface and is enabled by the increasing trend where a user wears two or more wearables, such as a smartwatch, wireless earbuds, or a smartphone. In a user study comparing five potential SelfSync gestures in isolation, our system averages 96%, 98% and 88% for user dependent, user adapted, and user independent accuracy, respectively. For when the user has a phone in a pocket and a smartwatch, we suggest twisting the hand about the wrist while moving the leg with the phone in synchrony left and right. When the user has a head worn device and a smartwatch, we suggest twisting the hand while twisting the head left and right.

CCS CONCEPTS

• **Human-centered computing → Gestural input**.

KEYWORDS

wearable computing; interaction; synchronous gestures; smartwatches; head-worn displays

ISWC '19, September 9–13, 2019, London, United Kingdom

ACM ISBN 978-1-4503-6870-4/19/09...$15.00
https://doi.org/10.1145/3341163.3347745

ACM Reference Format:
Juyoung Lee, Shaurye Aggarwal, Jason Wu, Thad Starner, and Woontack Woo. 2019. SelfSync: Exploring Self-Synchronous Body-Based Hotword Gestures for Initiating Interaction. In *Proceedings of the 2019 International Symposium on Wearable Computers (ISWC '19), September 9–13, 2019, London, United Kingdom.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3341163.3347745

1 INTRODUCTION

Many aspects of mobile and wearable interaction rely on accurate and reliable input initiation. For example, most voice assistants use certain trigger phrases (*e.g.,* "OK Google") to activate functionality. Other wearables that dim or disable the screen require users to "unlock" the device before use. However, these activation mechanisms may not be socially acceptable or usable in many cases where the user is situationally impaired. For example, activating voice assistants in public is difficult and can be socially awkward. Also, smartwatches can occupy both hands and require conspicuous movement to unlock. Alternatively, many "gesture delimiter" approaches, where a special unlock gesture indicates the interface should look for a subsequent gestural command, can require distinctive gestures that draw significant attention, as the gestures must be differentiable from everyday motion with very low false positive rates [27]. Finally, speech-based triggers are not socially acceptable in many cases and can cause the undesirable result of simultaneously activating nearby devices.

Our approach to interface initiation can take advantage of combinations of wearable devices commonly worn today (wireless earbuds, smartphone, smartwatch, etc.) and coordinates their input for gesture recognition. Specifically, we explore self-synchronous body-based gestures for the purpose of accurate and robust input initialization using the SelfSync gesture interface. We first present

SelfSync in the context of related work on subtle activation gestures and social acceptability. We describe the design and implementation of the SelfSync gesture interaction and the underlying recognition method. Our system is then evaluated in an offline and online experiment, quantifying accuracy, speed, task workload, and social acceptability.

2 RELATED WORK

2.1 Activation Gesture Input

Gesture recognition systems often use gestural delimiters to segment different portions of user interaction [16]. These "activation" gestures must have very low false-positive rates while retaining high reliability and detection accuracy [27]. Systems such as Whack and DoubleFlip introduce gestures that are easily distinguishable from other motion when characterized by accelerometry data [12, 27]. Gesture On explores touchscreen-based gesture recognition on locked mobile devices in "standby" mode without needing to toggle between modes of interaction [20]. A similar concept of initializing a locked device to a targeted application state has also been explored on tablets [29]. Activation gestures for wearable devices such as smartwatches and smartglasses employ similar techniques. WristRotate is a personalized motion delimiter for wrist-worn devices [15], and Google Glass uses a head-tilting gesture for device waking. Activation gestures using non-voice acoustics [25] and gaze-tracking have also been explored [23].

SelfSync creates an activation gesture input method which is sufficiently robust to false triggering such that it is adequate for daily usage even when providing multiple commands.

2.2 Synchronous Gestures and Motion Correlation

Detecting correlated movement between multiple devices has been explored as a natural mechanism for pairing and initiating information transfer. Hinckley investigated correlating the sensor values of multiple tablet computers as a distributed sensing method to detect gestures for initiating information transfer and changing display settings [11]. Other work enabling device pairing [40] and authentication [18, 37] have found the explicit, localized interaction as beneficial for such applications.

Velloso et al. used correlated movement as an input method for wearable and small form-factor devices, where users express intent for selecting an item by mimicking the movement of a cursor that orbits the item [34]. Some synchronous motion eye interfaces take advantage of the natural tendency for users' eyes to track moving targets of interest [7, 8, 31, 35, 36]. Selection of both virtual UI elements and real-world objects can be performed by tracking body-based synchronous gestures as well [3, 4]. Finally, synchronous gestures have recently been explored as hand gestures for subtle control and one-handed smartwatch input [9, 24, 38, 39].

In contrast to these previous systems, SelfSync enables motion-correlation interfaces without external stimuli such as periodic sounds or visual elements. Such a system may reduce or eliminate the visual distraction of moving cursors, improve recognition accuracy, decrease training, increase speed and reduce the cognitive load caused by matching an external rhythm.

Body Part	Movements	Gesture
Head	Horizontal (**Head.**H) Vertical (**Head.**V)	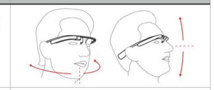
Hand	Twisting (**Hand.**T)	
Leg	Horizontal (**Leg.**H) Vertical (**Leg.**V)	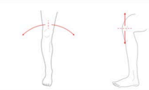

Figure 2: SelfSync gesture set abbreviations

2.3 Subtle Input Interfaces

The usefulness and adoption of gesture interfaces are highly dependent on the social acceptability of the gestures. Previous research evaluated the social acceptability of various device-based and body-based gestures in common settings and found that participants strongly preferred subtle gestures that required small and unobtrusive motion [26]. To this end, input interfaces have been developed that seek to minimize the amount of discernible movement required to trigger the system. Costanza et al. explored the use of an input device based on the electromyographic (EMG) signal for motionless interaction [5]. The bone-conduction microphone used in Bitey and the Outer-Ear Interface (OEI) recognize jaw and tongue gestures for subtle interaction [1, 2, 19]. Research in developing silent speech interfaces are similarly motivated by subtle interaction [6] by minimizing visible movement [14, 19, 21] and sound [2, 10, 32]. Another strategy for facilitating socially acceptable gestures is to disguise the gesture as an everyday action or interaction with another "conventional" technology [22, 26, 33]. Itchy Nose presents a subtle interface for wearable computers by detecting nose flicks, taps, and rubs using EOG glasses [17]. Blinking and winking gestures have also been used as a mode of unobtrusive input interaction that can be disguised as normal activity [13, 30].

In our research, we use multiple everyday actions such as hand twists or head nodding for input and achieve sufficient accuracy using commodity wearable devices.

3 SELFSYNC

SelfSync is designed as a synchronous body-based gesture performed across multiple body parts. When creating our system's gesture set, we considered common locations where a wearable or sensor-enabled device might be present on the user's body.

- **Head** - Many wearable accessories such as headphones, earbuds, and smart eyewear have integrated sensors such accelerometers, gyroscopes, and even electrooculography (EOG) electrodes.
- **Hand** - Devices such as smartwatches, smart rings, and sensor-enabled armbands can capture both gross arm movements and subtle finger movements.
- **Leg** - Smartphones are often placed in the user's pants pocket and the area is well positioned for capturing information about the orientation and movement of the leg [28].

We chose five common gestures that can be performed easily with periodic movement using the aforementioned body parts (Figure 1): leg Left-Right (toe rotation) and Up-Down (dorsiflexion); head Left-Right and Up-Down; and Hand Twist. For SelfSync, we initially considered seven combinations of these gestures: Hand.T & Head.V, Hand.T & Head.H, Hand.T & Leg.H, Head.V & Leg.V, Head.H & Leg.H, Hand.T & Leg.V, and Head.H & Leg.V. In our pilot studies, the last two gestures had the lowest true-positive accuracy, and we found them to be much harder than the rest of the gestures. Hence, we omitted them from our gesture set and decided to explore the first five gestures in our evaluations (Figure 1).

3.1 System Overview

Most commodity smart devices have in-built Inertial Measurement Unit (IMU) sensors to track the movement of devices. We used gyroscopes which measure the rate of rotation to track the movement of each body part. Hence, SelfSync's gestures are rotational motion. For both of our studies, we decided to use common off-the-shelf devices: Sony Smartwatch 3 SWR50 for the hand, Google Glass Explorer Edition for the head, and an Android phone for the leg. We sampled gyroscope data at 33Hz for the watch and 100Hz for Glass and the phone. The data streams through a UDP socket via WiFi to a central server for processing. We used a Macbook Pro as the central server. We implemented the system in Python with Pygame for visualization in data gathering and Scikit-learn for offline training and testing of machine learning classifiers.

In the data collection, to help participants to know whether they are doing well, we used a simple threshold-based classifier to show users visual feedback of their trial. The interface only checks whether a Pearson correlation coefficient value has passed a given threshold. We used 1.5-seconds (1-second for the offline classifier) of data for calculation. We ran a Principal Component Analysis that converted the three axes of each gyroscope to one dominant axis for segmented windows of data. Then, we ran cross-correlation to align the time series from two different devices and ultimately calculated the Pearson correlation coefficient. We did this process for each pair: head & hand, head & leg, and hand & leg. With this system, we were able to distinguish which body parts were moving synchronously.

Per-Device Features (x3 Devices)			Coordinated Features
Window Statistics	**First Difference Statistics**	**Custom Features**	**Signal Correlation**
i) max *ii)* min *iii)* range *iv)* mean *v)* std. dev. *vi)* root mean squared	*vii)* max *viii)* min *ix)* mean *x)* std. dev. *xi)* root mean squared	*xii)* # positive peaks *xiii)* # negative peaks *xiv) - xvi)* abs. value of 3-D principal component vector *xvii)* index of largest axis of principal component vector	*i) - iii)* pairwise correlation values *iv) - vi)* pairwise correlation values after time-delay correction *vii)* index of pair w/ maximum correlation after time-delay correction

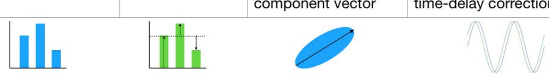

Figure 3: Feature set for the SelfSync gesture interface.

After data collection, to classify different SelfSync gestures and make it robust to daily-life actions, we trained a Random Decision Forest classifier with extracted features, including the correlation coefficient. We extracted a set of per-device features and coordinated (cross-device) features, which are shown in Figure 3. In total, we used 58 features (3 devices x 17 per-device features + 7 coordinated features).

3.2 Data Collection

To collect data for the classifier training and testing, we conducted a within-subject offline study in our laboratory. We collected both false and true positive data. The false data was collected while participants were completing paperwork for the study. True trial data was collected with displayed prompts over two conditions: standing and sitting.

3.2.1 Participants. We conducted the study with 10 students (all male, ages 20-25) from our institution in the United States, recruited via word of mouth. Only two participants regularly used wearables (smartwatches) for tasks such as notification updates, screening calls, controlling music players, and monitoring health. Participants received $10 compensation for their time.

3.2.2 Procedure. Upon arrival, the participants wore Google Glass. The smartwatch was worn on the left wrist, and the smartphone was kept in the right-leg pocket. The participants walked around the lab with the three devices while reading the instructions and consent form. The experimenter walked alongside the participants and explained the details of the study. Approximately five minutes of false-positive data was collected from each participant.

For true data collection, the experimenter demonstrated each gesture in no specific order. The participants were asked to practice until they could confidently perform the gesture. Visual feedback of the gesture detection by the threshold-based classifier was shown during this phase. The participants then got acquainted with the experimental setup. For the main study, each gesture was presented in random order and participants were asked to perform the target gesture as naturally as possible. Users began performing the gesture during a five seconds warmup period that commenced after acknowledging the target gesture by pressing a key on the keyboard. Then a two seconds recording period started automatically after the warmup period ended. The participants were notified of the start and end by beeps. Participants performed five repetitions for five gestures across two conditions giving a total of 50 gestures per participant.

4 SELFSYNC CLASSIFIER

We removed eight outliers from our data set which sensor value was too short or small, or was performed with different body part. We took two one-second windows from each recorded trial for true positive data, and cut the collected false positive data into one second windows avoiding overlap. Next, we trained a Random Decision Forest (RDF) classifier on this dataset and evaluated the gestures based on accuracies for user-independent, user-dependent, and user-adaptive models. The false positive rate was calculated using the user-adaptive model and applying a threshold.

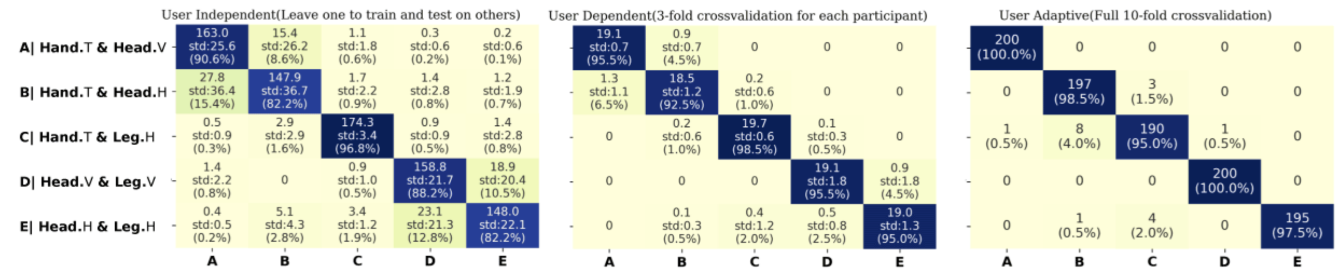

Figure 4: Confusion matrices for user independent (left), user dependent (center), and user adaptive (right) tests.

4.1 Result

User-dependent models were trained using a subset of each user's data and then tested on an independent test set. We performed 3-fold cross validation on each user's data and achieved 96.0% accuracy across all five gestures. Details are shown in Figure 4. We built user-independent models by training the classifier on data from all users holding out one user at a time and then testing the classifier on the holdout's data. The results were averaged across all combinations of leave-one-user-out cross validation. In our results, Hand.T & Leg.H had the highest average accuracy of 96.8%, followed by Hand.T & Head.V with 90.6% accuracy. Training a user-independent model with additional training instances from a specific user results in a user-adaptive model. By performing 10-fold cross validation on the whole dataset, we achieved 98.2% average accuracy.

We tested the user-adaptive models for false positives at four confidence level thresholds: 0.6, 0.7, 0.8, and 0.9. At 0.6 confidence threshold, when we ran all the false-positive data we collected at the beginning of our user study, the false-positive rate was 5.29 errors per hour. The classifier triggered four false detections as either Hand.T & Head.V or Head.V & Leg.V. For all other thresholds, the SelfSync classifiers did not have any false positive errors (0.00 errors per hour). There was not a considerable change in the other true positive accuracies measured above.

A gesture for input initialization should achieve high accuracy with a low number of false positives to be practical for everyday use. Our user-adaptive models recognized a few false-positive instances as gestures involving the leg. Classification of non-gesture data as leg gestures could be due the fact that users were mainly involved with walking and reading during the false-positive data collection phase of the user study. However, for thresholds ≥ 0.7, our classifier was capable of differentiating every day action from actual gestures and resulted in zero false-positives without affecting gesture recognition accuracies considerably. We believe that everyday actions rarely include synchronous motion of two body parts and hence SelfSync should be robust against false-positives.

Among hand and head gestures, Hand.T & Head.V has higher accuracy in user-independent and user-dependent models. However, since Hand.T & Head.H achieved similar and higher accuracies in user-dependent and user-adaptive cases respectively and was more suitable in terms of workload, acceptability, and user preference, we think it is the better gesture among the two.

For combined head and leg gestures, leg up-down had worse synchronization issues compared to leg left-right. However, it achieved better accuracy in all cases. Moreover, Hand.T & Leg.H achieved the best accuracy among the leg gestures in all cases. Hence, when results from the previous section are taken into account, Hand.T & Leg.H comes out as the most optimal gesture among SelfSync gestures involving leg.

4.2 Discussion

When accuracies are higher for user adapted training than user dependent, it often indicates that each user did not provide enough data to span the space of probable input. In this case, that hypothesis is highly probable. However, the user independent rates are significantly less than the user dependent rates. That result indicates that there is variability in how the gestures are performed across users, perhaps due to improper training or physiological differences. However, it could also mean that there was not enough training to cover all the situations in which the interfaces was used. Adding some user specific training improves recognition significantly. In addition, Hand.T & Leg.H is surprisingly robust. Perhaps there is less variability in how users move their legs back and forth horizontally. We only need one gesture with high accuracy to be the equivalent of a hotword, and one can imagine a subtle version of this gesture, with slight twisting of the wrist and moving of the leg, to initiate a silent interaction while in a meeting.

5 CONCLUSION

Most gesture systems focus on gestures for responding to notifications or for giving commands once an interface is triggered and listening. Few gesture recognizers have low enough false positive rates such that the gesture can be used to initiate an interaction. SelfSync introduces the concept of having the user move two body parts in synchrony to indicate a desire to initiate communication with their computer. However, how subtle can these gestures be made? Would they be unnoticeable to a bystander sharing an elevator or to a conversational partner? One of the best uses of wearable computers is to provide aid during face to face conversation. Creating subtle gestural interfaces that can be used without distracting from a face-to-face conversation would help further that vision significantly.

ACKNOWLEDGMENTS

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2019-0-01270, WISE AR UI/UX Platform Development for Smartglasses)

REFERENCES

[1] Daniel Ashbrook, Carlos Tejada, Dhwanit Mehta, Anthony Jiminez, Goudam Muralitharam, Sangeeta Gajendra, and Ross Tallents. 2016. Bitey: An Exploration of Tooth Click Gestures for Hands-free User Interface Control. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '16)*. ACM, New York, NY, USA, 158–169. https://doi.org/10.1145/2935334.2935389

[2] Abdelkareem Bedri, David Byrd, Peter Presti, Himanshu Sahni, Zehua Gue, and Thad Starner. 2015. Stick It in Your Ear: Building an In-ear Jaw Movement Sensor. In *ACM International Symposium on Wearable Computers (ISWC) (UbiComp/ISWC'15 Adjunct)*. ACM, New York, NY, USA, 1333–1338. https://doi.org/10.1145/2800835.2807933

[3] Marcus Carter, Eduardo Velloso, John Downs, Abigail Sellen, Kenton O'Hara, and Frank Vetere. 2016. PathSync: Multi-User Gestural Interaction with Touchless Rhythmic Path Mimicry. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3415–3427. https://doi.org/10.1145/2858036.2858284

[4] Christopher Clarke, Alessio Bellino, Augusto Esteves, Eduardo Velloso, and Hans Gellersen. 2016. TraceMatch: A Computer Vision Technique for User Input by Tracing of Animated Controls. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 298–303. https://doi.org/10.1145/2971648.2971714

[5] Enrico Costanza, Alberto Perdomo, Samuel A. Inverso, and Rebecca Allen. 2004. EMG as a subtle input interface for mobile computing. In *Mobile Human-Computer Interaction - MobileHCI 2004*, Stephen Brewster and Mark Dunlop (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 426–430.

[6] B. Denby, T. Schultz, K. Honda, T. Hueber, J. M. Gilbert, and J. S. Brumberg. 2010. Silent Speech Interfaces. *Speech Commun.* 52, 4 (April 2010), 270–287. https://doi.org/10.1016/j.specom.2009.08.002

[7] Murtaza Dhuliawala, Juyoung Lee, Junichi Shimizu, Andreas Bulling, Kai Kunze, Thad Starner, and Woontack Woo. 2016. Smooth Eye Movement Interaction Using EOG Glasses. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction (ICMI '16)*. ACM, New York, NY, USA, 307–311. https://doi.org/10.1145/2993148.2993181

[8] Augusto Esteves, Eduardo Velloso, Andreas Bulling, and Hans Gellersen. 2015. Orbits: Gaze Interaction for Smart Watches Using Smooth Pursuit Eye Movements. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 457–466. https://doi.org/10.1145/2807442.2807499

[9] Euan Freeman, Gareth Griffiths, and Stephen A. Brewster. 2017. Rhythmic Micro-gestures: Discreet Interaction On-the-go. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction (ICMI '17)*. ACM, New York, NY, USA, 115–119. https://doi.org/10.1145/3136755.3136815

[10] Masaaki Fukumoto. 2018. SilentVoice: Unnoticeable Voice Input by Ingressive Speech. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18)*. ACM, New York, NY, USA, 237–246. https://doi.org/10.1145/3242587.3242603

[11] Ken Hinckley. 2003. Synchronous Gestures for Multiple Persons and Computers. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology (UIST '03)*. ACM, New York, NY, USA, 149–158. https://doi.org/10.1145/964696.964713

[12] Scott E. Hudson, Chris Harrison, Beverly L. Harrison, and Anthony LaMarca. 2010. Whack Gestures: Inexact and Inattentive Interaction with Mobile Devices. In *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '10)*. ACM, New York, NY, USA, 109–112. https://doi.org/10.1145/1709886.1709906

[13] Ricardo Jota and Daniel Wigdor. 2015. Palpebrae Superioris: Exploring the Design Space of Eyelid Gestures. In *Proceedings of the 41st Graphics Interface Conference (GI '15)*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 273–280. http://dl.acm.org/citation.cfm?id=2788890.2788938

[14] Arnav Kapur, Shreyas Kapur, and Pattie Maes. 2018. Alterego: A personalized wearable silent speech interface. In *23rd International Conference on Intelligent User Interfaces (IUI '18)*. ACM, New York, NY, USA, 43–53. https://doi.org/10.1145/3172944.3172977

[15] Frederic Kerber, Philipp Schardt, and Markus Löchtefeld. 2015. WristRotate: A Personalized Motion Gesture Delimiter for Wrist-worn Devices. In *Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia (MUM '15)*. ACM, New York, NY, USA, 218–222. https://doi.org/10.1145/2836041.2836063

[16] Sven Kratz and Jason Wiese. 2016. GestureSeg: Developing a Gesture Segmentation System Using Gesture Execution Phase Labeling by Crowd Workers. In *Proceedings of the 8th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS '16)*. ACM, New York, NY, USA, 61–72. https://doi.org/10.1145/2933242.2933261

[17] Juyoung Lee, Hui-Shyong Yeo, Murtaza Dhuliawala, Jedidiah Akano, Junichi Shimizu, Thad Starner, Aaron Quigley, Woontack Woo, and Kai Kunze. 2017. Itchy Nose: Discreet Gesture Interaction Using EOG Sensors in Smart Eyewear. In *Proceedings of the 2017 ACM International Symposium on Wearable Computers (ISWC '17)*. ACM, New York, NY, USA, 94–97. https://doi.org/10.1145/3123021.3123060

[18] Ho-Man Colman Leung, Chi-Wing Fu, and Pheng-Ann Heng. 2018. TwistIn: Tangible Authentication of Smart Devices via Motion Co-analysis with a Smartwatch. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 2, Article 72 (July 2018), 24 pages. https://doi.org/10.1145/3214275

[19] Richard Li, Jason Wu, and Thad Starner. 2019. TongueBoard: An Oral Interface for Subtle Input. In *Proceedings of the 10th Augmented Human International Conference 2019 (AH2019)*. ACM, New York, NY, USA, 1–9. https://doi.org/10.1145/3311823.3311831

[20] Hao Lu and Yang Li. 2015. Gesture On: Enabling Always-On Touch Gestures for Fast Mobile Access from the Device Standby Mode. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3355–3364. https://doi.org/10.1145/2702123.2702610

[21] Geoffrey S Meltzner, James T Heaton, Yunbin Deng, Gianluca De Luca, Serge H Roy, and Joshua C Kline. 2018. Development of sEMG sensors and algorithms for silent speech recognition. *Journal of neural engineering* 15, 4 (2018), 046031.

[22] Halley P. Profita, James Clawson, Scott Gilliland, Clint Zeagler, Thad Starner, Jim Budd, and Ellen Yi-Luen Do. 2013. Don'T Mind Me Touching My Wrist: A Case Study of Interacting with On-body Technology in Public. In *Proceedings of the 2013 International Symposium on Wearable Computers (ISWC '13)*. ACM, New York, NY, USA, 89–96. https://doi.org/10.1145/2493988.2494331

[23] Hayes Solos Raffle, Adrian Wong, and Ryan Geiss. 2012. Unlocking a screen using eye tracking information. US Patent 8,235,529.

[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. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 4, Article 158 (Jan. 2018), 26 pages. https://doi.org/10.1145/3161162

[25] Gabriel Reyes, Dingtian Zhang, Sarthak Ghosh, Pratik Shah, Jason Wu, Aman Parnami, Bailey Bercik, Thad Starner, Gregory D. Abowd, and W. Keith Edwards. 2016. Whoosh: Non-voice Acoustics for Low-cost, Hands-free, and Rapid Input on Smartwatches. In *Proceedings of the 2016 ACM International Symposium on Wearable Computers (ISWC '16)*. ACM, New York, NY, USA, 120–127. https://doi.org/10.1145/2971763.2971765

[26] Julie Rico and Stephen Brewster. 2010. Usable Gestures for Mobile Interfaces: Evaluating Social Acceptability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 887–896. https://doi.org/10.1145/1753326.1753458

[27] Jaime Ruiz and Yang Li. 2010. DoubleFlip: A Motion Gesture Delimiter for Interaction. In *Adjunct Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 449–450. https://doi.org/10.1145/1866218.1866265

[28] Jeremy Scott, David Dearman, Koji Yatani, and Khai N. Truong. 2010. Sensing Foot Gestures from the Pocket. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 199–208. https://doi.org/10.1145/1866029.1866063

[29] 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 (CHI '13)*. ACM, New York, NY, USA, 3027–3036. https://doi.org/10.1145/2470654.2481421

[30] R. Shaw, E. Crisman, A. Loomis, and Z. Laszewski. 1990. The eye wink control interface: using the computer to provide the severely disabled with increased flexibility and comfort. In *[1990] Proceedings. Third Annual IEEE Symposium on Computer-Based Medical Systems*. IEEE, Chapel Hill, NC, USA, 105–111. https://doi.org/10.1109/CBMSYS.1990.109386

[31] Junichi Shimizu, Juyoung Lee, Murtaza Dhuliawala, Andreas Bulling, Thad Starner, Woontack Woo, and Kai Kunze. 2016. Solar System: Smooth Pursuit Interactions Using EOG Glasses. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, USA, 369–372. https://doi.org/10.1145/2968219.2971376

[32] Ke Sun, Chun Yu, Weinan Shi, Lan Liu, and Yuanchun Shi. 2018. Lip-Interact: Improving Mobile Device Interaction with Silent Speech Commands. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18)*. ACM, New York, NY, USA, 581–593. https://doi.org/10.1145/3242587.3242599

[33] Aaron Toney, Barrie Mulley, Bruce H. Thomas, and Wayne Piekarski. 2003. Social Weight: Designing to Minimise the Social Consequences Arising from Technology Use by the Mobile Professional. *Personal Ubiquitous Comput.* 7, 5 (Oct.

2003), 309–320. https://doi.org/10.1007/s00779-003-0245-8

[34] Eduardo Velloso, Marcus Carter, Joshua Newn, Augusto Esteves, Christopher Clarke, and Hans Gellersen. 2017. Motion Correlation: Selecting Objects by Matching Their Movement. *ACM Trans. Comput.-Hum. Interact.* 24, 3, Article 22 (April 2017), 35 pages. https://doi.org/10.1145/3064937

[35] Eduardo Velloso, Markus Wirth, Christian Weichel, Augusto Esteves, and Hans Gellersen. 2016. AmbiGaze: Direct Control of Ambient Devices by Gaze. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 812–817. https://doi.org/10.1145/2901790.2901867

[36] Mélodie Vidal, Andreas Bulling, and Hans Gellersen. 2013. Pursuits: Spontaneous Interaction with Displays Based on Smooth Pursuit Eye Movement and Moving Targets. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, New York, NY, USA, 439–448. https://doi.org/10.1145/2493432.2493477

[37] Jacob Otto Wobbrock. 2009. TapSongs: Tapping Rhythm-based Passwords on a Single Binary Sensor. In *Proceedings of the 22Nd Annual ACM Symposium on User Interface Software and Technology (UIST '09)*. ACM, New York, NY, USA, 93–96. https://doi.org/10.1145/1622176.1622194

[38] Jason Wu, Cooper Colglazier, Adhithya Ravishankar, Yuyan Duan, Yuanbo Wang, Thomas Ploetz, and Thad Starner. 2018. Seesaw: Rapid One-Handed Synchronous Gesture Interface for Smartwatches. In *Proceedings of the 2018 ACM International Symposium on Wearable Computers (ISWC '18)*. ACM, New York, NY, USA, 17–20. https://doi.org/10.1145/3267242.3267251

[39] Cheng Zhang, Xiaoxuan Wang, Anandghan Waghmare, Sumeet Jain, Thomas Ploetz, Omer T. Inan, Thad E. Starner, and Gregory D. Abowd. 2017. FingOrbits: Interaction with Wearables Using Synchronized Thumb Movements. In *Proceedings of the 2017 ACM International Symposium on Wearable Computers (ISWC '17)*. ACM, New York, NY, USA, 62–65. https://doi.org/10.1145/3123021.3123041

[40] Tengxiang Zhang, Xin Yi, Ruolin Wang, Yuntao Wang, Chun Yu, Yiqin Lu, and Yuanchun Shi. 2018. Tap-to-Pair: Associating Wireless Devices with Synchronous Tapping. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 4, Article 201 (Dec. 2018), 21 pages. https://doi.org/10.1145/3287079