

ViBand: High-Fidelity Bio-Acoustic Sensing Using Commodity Smartwatch Accelerometers

Gierad Laput Robert Xiao Chris Harrison
Carnegie Mellon University, Human-Computer Interaction Institute
5000 Forbes Ave, Pittsburgh, PA 15213
{gierad.laput, brx, chris.harrison}@cs.cmu.edu




Figure 1. Comparison of 100 Hz vs. 4000 Hz accelerometer signals. At steady state, both signals look identical (A). However, high frequency micro-vibrations propagating through the arm are missed by the 100 Hz accelerometer (B). Characteristic vibrations can come from oscillating objects (B), hand gestures (C) and the operation of mechanical objects (D).

ABSTRACT

Smartwatches and wearables are unique in that they reside on the body, presenting great potential for always-available input and interaction. Their position on the wrist makes them ideal for capturing *bio-acoustic* signals. We developed a custom smartwatch kernel that boosts the sampling rate of a smartwatch's existing accelerometer to 4 kHz. Using this new source of high-fidelity data, we uncovered a wide range of applications. For example, we can use bio-acoustic data to classify hand gestures such as flicks, claps, scratches, and taps, which combine with on-device motion tracking to create a wide range of expressive input modalities. Bio-acoustic sensing can also detect the vibrations of grasped mechanical or motor-powered objects, enabling passive object recognition that can augment everyday experiences with context-aware functionality. Finally, we can generate structured vibrations using a transducer, and show that data can be transmitted through the human body. Overall, our contributions unlock user interface techniques that previously relied on special-purpose and/or cumbersome instrumentation, making such interactions considerably more feasible for inclusion in future consumer devices.

Author Keywords

Wearables; Gestures; Object Detection; Vibro-Tags

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
UIST 2016, October 16-19, 2016, Tokyo, Japan
© 2016 ACM. ISBN 978-1-4503-4189-9/16/10...\$15.00
DOI: <http://dx.doi.org/10.1145/2984511.2984582>

ACM Classification Keywords

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

INTRODUCTION

Watches are unique among computing devices in that they are worn, offering great potential to transform arms and hands into expressive input and sensing platforms. As people use their hands, tiny micro-vibrations propagate through the arm, carrying information about the objects they interact with and the activities they perform throughout the day. Smartwatches are ideally situated to capture these vibrations (Figures 1 and 2).

Although all modern smartwatches contain accelerometers, their APIs generally limit the sampling rate to around 100 Hz (Figure 1, top purple lines). This is sufficient for their main use: detecting the orientation of the watch (*e.g.*, to automatically activate the screen when raised). Some smartwatches also track step count (~ 2 Hz), which is also easily captured with 100 Hz sampling.

In this work, we use an off-the-shelf smartwatch with a modified OS kernel to capture accelerometer data at 4000 times per second (Figure 1, bottom purple lines). This fast sampling allows the smartwatch to not only capture coarse motions, but also rich bio-acoustic signals. For example, in Figure 1B, the sinusoidal oscillations of the toothbrush's motor are clearly visible. In Figure 1C (fingers rubbing) and D (pressing stapler), the 100 Hz signal captures the coarse impulse, but no useful spectral information is available.

Most smartwatches include microphones, which provide even higher sampling rates (typically 44.1 kHz). However, microphones are specifically designed to capture airborne vibrations, not contact vibrations, which means purposeful

signals must be segmented from background environmental noise. In contrast, our bio-acoustic approach only captures signals that are physically coupled to the body (Figures 1A and B, and Video Figure). This approach makes our technique naturally resistant to external environmental noise.

As we will discuss, our approach can be applied to a wide array of use domains; we selected three that we found to be particularly compelling. First, we use bio-acoustic data to classify hand gestures, which we combine with on-device motion tracking to enable a wide range of expressive input modalities. Second, we detect and classify vibrations of grasped mechanical or motor-powered objects, enabling uninstrumented object recognition. Finally, we explore structured vibrations and demonstrate reliable data transmission through the human body.

Our evaluations show that our sensing technique is accurate, robust to noise, relatively consistent across users, and independent of location or environment. Our system, which we call *ViBand*, makes the following contributions: 1) a system that performs bio-acoustic sensing using commodity accelerometers already present in modern smartwatches; 2) a set of example use domains enabled by our technique, including gesture detection, grasped object sensing, and data transmission; 3) a series of user studies evaluating the feasibility and accuracy of our sensing technique; and 4) a series of example applications for wrist-worn bio-acoustic sensing that illustrate the potential of our approach. Collectively, these bring to light novel and rich functionality for smartwatches, expanding their envelope of possible interactions.

RELATED WORK

ViBand intersects with a range of HCI and sensing topics, including worn hand input and gesture sensing, bio-acoustics, object recognition, and data transmission.

Worn Hand Input and Gesture Sensing

Hand gestures offer expressive input modalities that complement existing interfaces and devices. A popular approach for hand gesture recognition takes advantage of optical sensors such as cameras [25] and IR sensors [16, 34, 35, 52]. It is also possible to sense hand gestures by approximating skin contours and deformations. For instance, armbands instrumented with IR sensors [34, 35] or pressure sensors [10, 24] can measure skin contact variations whenever particular gestures are performed. Despite being low-cost, these approaches are highly dependent on contact conditions, which are inherently sensitive to periodic armband removal, and equally susceptible to unintentional arm movements.

Hand gestures can likewise be modeled by examining the internal anatomical configuration of the user's arm. Approaches can be passive, such as electromyography [45, 46], where gestures are classified by measuring the electrical signals caused by muscle activation, or active [9, 41], as in *Touché* [47] and *Tomo* [55], where a signal is injected into the body to detect hand gestures.




Figure 2. Smartwatches are capable computing devices that are always worn. *ViBand* allows us to transform the arm into an expressive input and sensing platform.

Finally, coarse and fine hand gestures indirectly induce arm motions which can be captured by inertial sensors *e.g.*, accelerometers and gyroscopes. Previous work introduced gloves equipped with accelerometers to model fine hand gestures [37]. Likewise, several techniques take advantage of the inertial sensors present in contemporary smartwatches. *Akl et al.* [1] and *Bernaerts et al.* [5] utilize wearable accelerometers to recognize gross-motor or whole hand motions. *Xu et al.* [54] used inertial sensors attached to the arm, wrist, and finger to detect three types of gestures (arm, wrist, and finger movements, respectively), although their system was trained and tested on one user sitting in an armchair. *Wen et al.* [53] introduced finger gesture recognition using commodity accelerometers on a smartwatch, supporting a maximum of five gestures. However, they caution that their technique is highly sensitive to arm orientation, and was never deployed in a real-time environment.

Bio-Acoustic Input and Sensing

Bio-acoustics has been studied in many fields [12, 15, 18, 32, 38, 48], including HCI. For instance, *Amento et al.* [2] placed contact microphones on the user's wrist to capture gross finger movement. Their work was first to demonstrate the use of on-body acoustic signals to passively recover finger gestures in one hand, although no formal evaluations were conducted. This became the direct inspiration to *Hambone* [11], which instrumented the user's limbs with piezo sensors to detect gestures (*e.g.*, finger flick, left foot rotate).

Likewise, *Skinput* [19] leveraged a similar technique, using an array of piezo sensors strapped onto the user's arm (above and below the elbow). Building on top of *Hambone*, *Skinput*'s sensor placement further expanded touch-interaction onto the arm, palm, and fingers. *The Sound of Touch* [33] employed a similar technique, using transdermal propagation of ultrasound across the user's arm to recover finger gestures and discrete touch points. A signal-emitting ring emitted ultrasound when the arm was touched or tapped, while an array of transducers monitored the transmitted signal. These bio-acoustic sensing approaches rely heavily on special-purpose sensors, increasing their invasiveness and ultimately limiting their practicality.

Object Recognition

Object recognition offers relevant information more closely matching a user's immediate context and environment [14, 29, 49]. Most approaches rely on markers [7, 20, 26, 40, 42]




Figure 3. Spectrogram of bio-acoustic signals from initial experiments. Note the distinctive signals and the lack of external noise (compared to e.g., microphones).

or special-purpose tags [28, 57, 58]. These offer robust recognition, but ultimately require every object to be instrumented. Further, these approaches approximate whether an object is nearby, and not when it is truly *grasped* or *handled*. Prior work has also leveraged acoustics to recognize objects [6, 8, 36]. For example, Ward *et al.* [51] built a worn necklace equipped with an accelerometer and a microphone to classify workshop tools, although the approach was susceptible to background noise.

Wearable devices are also increasingly being used for object sensing and recognition. Maekawa *et al.* [30, 31] utilized magnetic sensors and hand-worn coils to identify objects based on magnetic field changes. MagnifiSense [50] offered a similar approach, using three magneto-inductive sensors to identify objects during regular operation. Magnetic induction relies heavily on proximate contact between the sensor and the object, which is affected by posture, hand orientation, or even the inherent magnetic noise present in the human body. It is also possible to characteristically identify objects solely based on unintentionally emitted electromagnetic (EM) noise. EMISpy [56] detects different types of monitors and LCD screens by touching a screen surface while simultaneously holding a handheld EM sensor. Similarly, EM-Sense [27] augmented a smartwatch with a software-defined radio receiver, offering on-touch object detection in a wrist-worn form factor.

Through-Body Data Transmission

Data transmission through the body has been successfully demonstrated with radio frequency (RF) waves, in the form of “personal area networks.” Such networks can successfully transmit data at very high speeds amongst specially-equipped devices near the body [60]. More related to our technical approach are systems that use vibroacoustics to transmit data. Ripple [44], using an accelerometer and vibration motor mounted to a cantilevered metal arm (to amplify vibrations), transmitted data at about 200 bits/sec. Ripple II [43] utilized audible frequencies (2-10 KHz) to transmit data between a vibrating finger ring and a microphone at the finger tip. AT&T Labs publicly demonstrated a system that transmitted bio-acoustic data using a piezoelectric buzzer [3], although the technical details have not been published. Finally, and most similar to ViBand, is OsteoConduct [59], which transmits data through bone conduction. This system successfully demonstrated a data transmission rate of “almost 5 bits/sec” between the wrist, ear and lower back.

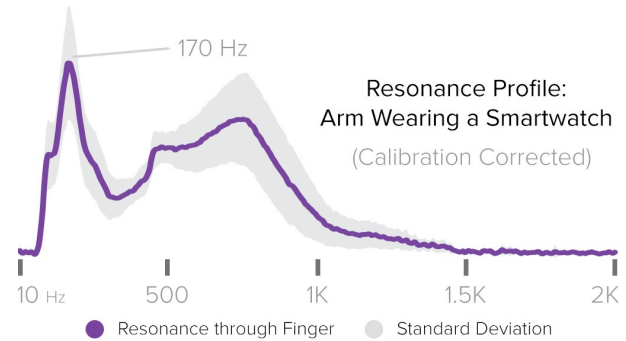


Figure 4. Resonance profile of an arm wearing a smartwatch (calibrated, watch+arm). Vibration frequencies between 20 Hz and 1 kHz transmit particularly well through the arm.

THEORY OF OPERATION

Although most mobile devices (including smartwatches) contain accelerometers and other inertial measurement sensors, existing APIs generally limit accelerometer data access to about 100 Hz. This rate is sufficient for detecting coarse movements such as changes in screen orientation or gross interactions such as walking, sitting, or standing. However, these accelerometers often support significantly higher sample rates – up to thousands of Hz. At these faster sampling speeds, the smartwatch can listen to nuanced and fine-grained movements that are initiated or experienced by the human arm. Like water, the human body is a non-compressible medium, making it an excellent vibration carrier. For example, when sampling at 4000 Hz, vibrations oscillating up to 2000 Hz (e.g., gestures, grasped objects) can be sensed and identified (per the Nyquist Theorem). This superior sensitivity transforms the smartwatch into a *bio-acoustic* sensor capable of detecting minute compressive waves propagating through the worn arm.

In our initial experiments, we sought to investigate whether our high-speed accelerometer signal was indeed bio-acoustic. To test this theory, we walked around our lab and performed a range of activities (e.g., tapping on table, scratching hand, grasping power tools) while simultaneously extracting accelerometer signals obtained from our prototype. Figure 3 is a spectrogram of some example signals. Each activity and object produces characteristic vibroacoustic signatures, and more critically, were only captured when in contact with the hand. These high-fidelity signals resemble those captured by a microphone, yet lack any audible external noise.

Like any medium, the human arm characteristically amplifies or attenuates vibrations at different frequencies. Therefore, we ran an experiment to identify the frequency transmission envelope of the human arm while wearing a smartwatch. First, we captured the resonance profile of an unworn smartwatch (LG G W100) placed directly on a transducer running a 0 to 2 kHz vibrational sweep. We then captured the resonance profile while the smartwatch was worn on the arm while pressing the transducer with the index finger. Figure 4 depicts the average resonance profile across

three users. These results suggest that oscillations between 20 Hz to 1 kHz transmit particularly well through the arm, with salient peaks at ~ 170 Hz and ~ 750 Hz.

IMPLEMENTATION

Our proof-of-concept system was developed on an LG G W100 smartwatch, which includes an InvenSense MPU6515 inertial measurement unit (IMU) capable of measuring acceleration at 4000 samples per second [22]. Of note, this is the same series of accelerometer used in many other popular smartwatches, including the Moto 360, LG Watch Urbane, Samsung Gear 2 and Gear Fit. However, the maximum rate obtainable through the Android Wear API [17] is 100 Hz. Therefore, we modified the Linux kernel on the device, replacing the existing accelerometer driver with our own custom driver.

Specifically, our kernel driver interfaces with the IMU via I²C, configuring the IMU registers to enable its documented high-speed operation [23]. Notably, this requires us to use the IMU’s onboard 4096-byte FIFO to avoid excessively waking up the system CPU. However, this FIFO only stores 160 ms of data—each data sample consists of a 16-bit sample for each of the three axes. Thus, we configured the driver to poll the accelerometer in a dedicated kernel thread, which reads the accelerometer FIFO into a larger buffer every 50 ms. Overall, this thread uses 9% of one of the watch’s four CPU cores.

We found that the accelerometer’s internal clock was not temperature-stabilized, resulting in higher sampling rates as the CPU temperature increased. We measured sampling rates varying between 3990 Hz (watch sleeping, off wrist) to 4080 Hz (on arm, high CPU activity). In response, we

augmented our kernel driver to compute the rate at which samples were written into the MPU’s FIFO buffer using a nanosecond-precision kernel timestamp. For applications requiring precise sampling rates, such as resonance profiling and data transmission, we normalized the input data to 4000 Hz using a sinc-based interpolator capable of supporting continuously variable input sample rates [13].

For prototyping purposes, we configured the watch to transmit all captured accelerometer data via Bluetooth to a paired Android phone, which then relayed the data to a laptop for analysis. This enabled rapid testing, iteration and development of our bio-acoustic applications. However, we also implemented data transmission and object classification on the watch itself for fully self-contained operation.

With this implementation, we unlock a wide variety of applications. In the next several sections, we describe how our technique enables novel interactions in three distinct application domains.

EXAMPLE DOMAIN 1: GESTURES

First, our technique can be used to classify unique hand gestures, such as flicks, claps, snaps and taps. These hand gestures create distinctive micro-vibrations that propagate through the arm. Depending on the location and type of gesture, different frequencies of vibrations are generated. Subsequently, various frequencies are attenuated during propagation (*e.g.*, anatomical features can act as passive vibroacoustic filters [19]). The resulting frequency profiles make many gestures uniquely identifiable. We explored three example gesture sets that offer distinctive bio-acoustic signals: one-handed gestures, two-handed gestures, and on-body touch input (see Figure 5).

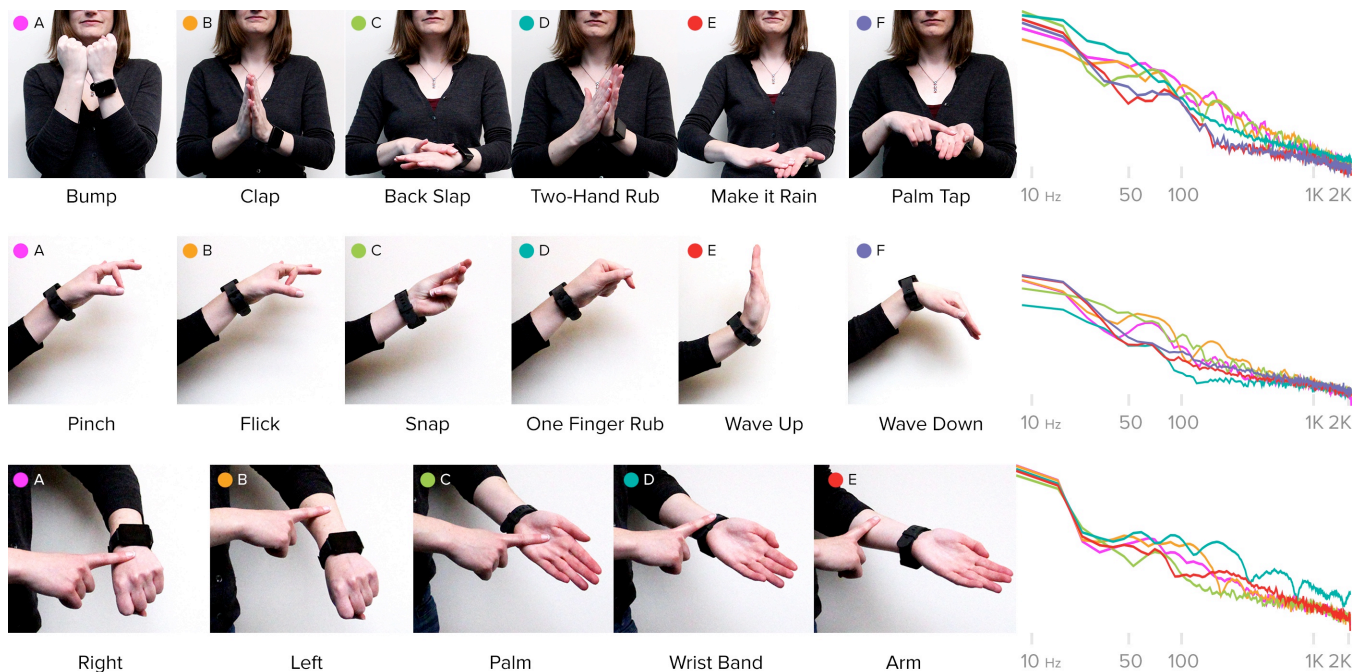


Figure 5. Example gesture sets: two-handed (top), one-handed (middle), and on-body touch input (bottom).




Figure 6. ViBand enables a wide range of interaction modalities when combined with coarse motion tracking information (from e.g., gyroscopes and accelerometers) natively available on existing smartwatches. Modalities include binary buttons (A), linear sliders (B), radial knobs (C), counters (D), hierarchical navigation (E), and even relative spatial tracking (F). In these examples, a pinching action serves as a gesture clutch. See also Video Figure.

Once the bio-acoustic signals are captured on the watch, we perform several signal processing operations to detect and classify hand gestures in real-time. For each incoming signal frame t , we first compute the power spectra of the fast Fourier transform (FFT) on data from each accelerometer axis, producing three spectra X_t, Y_t, Z_t . We use a Hamming window on the FFT to minimize spectral banding. To make sensing robust across hand orientations, we remove the DC component and combine the three FFTs into one by taking the max value across the axes ($F_{t,i} = \max(X_{t,i}, Y_{t,i}, Z_{t,i})$)

Next, we compute the average of the $w=20$ past FFT spectra ($S_i = \mu(F_{t,i}, F_{t-1,i}, \dots, F_{t-w+1,i})$) and extract statistical features from the averaged signal: mean, sum, min, max, 1st derivative, median, standard deviation, range, spectral band ratios, and the n highest peaks ($n=5$). These features form the input to a SMO-based SVM (poly kernel, $\epsilon=10^{-12}$, normalized) for real-time classification. From our experiments, we found that band ratios, peaks, mean, and standard deviation provide 90% of the bio-acoustic signal’s discriminative power. Table 1 describes these features and the motivations behind their use.

Feature Set	Operation	Justification
Power spectrum	S_i	Specific frequency data
Statistical	$\mu_S, \sigma_S, \Sigma_S, \max(S), \min(S), \text{centroid}, \text{peaks}$	Characterizes gross features of FFT signal
1 st Derivative	$\frac{d}{dt}(S_{t+1}) = S_{t+1} - S_t$	Encodes signal peaks and troughs
Band Ratios	$B_{j,k} = \frac{S_j}{S_k}$	Describes overall FFT shape, power distribution

Table 1. ViBand feature families and their respective merits.



Figure 7. A “vibro-tag” transmitting FSK bio-acoustic data through the user’s arm, and received by the smartwatch.



Figure 8. We implemented a range of encoding schemes to robustly transmit data-encoded vibrations.

When hand gestures are combined with relative motion tracking (e.g., native data from IMUs), our technique uncovers a range of interaction modalities (see Figure 6 and the Video Figure). These include: buttons, sliders, radial knobs, counters, hierarchical navigation, and positional tracking. On top of these, we can build applications that utilize these rich and expressive interaction modalities.

EXAMPLE DOMAIN 2: OBJECT DETECTION

Our sensing approach can also be used to identify grasped objects in order to e.g., launch context-relevant functionality or applications automatically. Specifically, when a user operates a mechanical or motor-powered device, the object produces characteristic vibrations, which transfer into the operator. Our bio-acoustic smartwatch captures and classifies these signals, allowing interactive applications to better understand their user’s context and further augment a wide range of everyday activities.

Worn microphones, which capture sounds produced by objects in operation, have been previously used for object recognition [51]. Because microphones are coupled through the air, they are particularly sensitive to ambient noise. Further, microphone-based techniques can only approximate when users are near to an active object, but not when they are truly interacting with an object. In contrast, our approach recognizes objects at the moment of touch, allowing us to generate *on-touch* and *on-release* events.

We utilize the same signal processing pipeline for both gestures and object detection, but with slightly tweaked parameters ($w=15, n=15$). We also apply a simple voting mechanism (size=10) to stabilize the recognition. Our setup recognizes a wide range of objects (see Figure 9), complement-

ing existing techniques (e.g., [27, 30, 50]), and further expanding capabilities for rich, context-sensitive applications.

EXAMPLE DOMAIN 3: STRUCTURED VIBRATIONS

In addition to being able to capture “passive” vibrations from objects and hand motions, we can also augment environments and objects with *structured vibrations*. We developed a “vibro-tag” consisting of a small (2.4 cm³) SparkFun COM-10917 Bone Conductor Transducer, powered by a standard audio amplifier. When a user touches the transducer, modulated vibrations are transmitted bio-acoustically to the smartwatch, which decodes the acoustic packet and extracts a data payload (Figures 7 and 8). Such tags could be used much like RFID or QR Codes while employing a totally orthogonal signaling means (vibroacoustic). A unique benefit of our approach is that it is only triggered upon physical touch (i.e., not just proximity) and is immune to variations in e.g., lighting condition.

Critically, we wanted to make our vibro tags *inaudible*, but still capable of transmitting data at high speed. Because the accelerometer can only sense frequencies up to 2 KHz, we cannot use ultrasound frequencies (e.g. frequencies above 16 kHz). We also ruled out frequencies above 300 Hz, as they would manifest as audible “buzzing” sounds. Using our transmission envelope data (Figure 4) and experiments with transmission frequencies, we ultimately selected 200 Hz as a suitable carrier frequency for data transmission.

Our data transmission system is a full stack signal pipeline, consisting of data packetization, error detection, error correction, and modulation layers. We first segment the input data stream into individually transmitted *data packets*. Applications are free to choose their own packet formats, but the recommended format consists of an 8-bit sequence number combined with a data payload. Packet size is constrained by the error detection and correction layers; in the current implementation, it can be up to 147 bits in length. In order to detect transmission errors and ensure that bad data is not accidentally accepted, we append an 8-bit *cyclic redundancy check* (CRC) to the message. In the present implementation, the CRC is computed by truncating the Adler-32 CRC of the message.

Next, error correction is applied. Although this stage also detects errors (like the CRC), its primary purpose is to mitigate the effects of minor transmission problems. We use a Reed-Solomon code [39] with 5 bits per symbol, allowing us to have 31 symbols per message (a total of 155 bits). These parameters were chosen to allow a single message to be transmitted in approximately one second using common modulation parameters. The number of ECC symbols can be tuned to compensate for noisier transmission schemes; see the evaluation for more details.

At this point, we transmit the full message+CRC+ECC, totaling 155 bits, as modulated vibrations. We experimented with four different classical modulation schemes [4], using binary Gray coding to encode bit strings as symbols:




Figure 9. Example objects and their bio-acoustic signatures. Note the x-axis is plotted in log-scale, from 10 Hz to 2 kHz. The light orange region is the spectrum below 50Hz, the maximum frequency detectable by most conventional smartwatch software. As seen in this Figure, a wide range of characteristic object oscillations fall well beyond the 50 Hz range.

- *Amplitude Shift Keying* (ASK): data is encoded by varying the amplitude of the carrier signal.
- *Frequency Shift Keying* (FSK): data is encoded by transmitting frequency multiples of the carrier signal (Figure 7, note spectrogram on laptop screen in background).
- *Phase Shift Keying* (PSK): adjusting the phase of the carrier signal, with respect to a fixed reference phase.
- *Quadrature Amplitude Modulation* (QAM): data encoded as variations in phase and amplitude, with symbols encoded according to a *constellation diagram* (Figure 12) mapping phase and amplitude combinations to bit sequences.

We prefix the message with a short header sequence consisting of three 20 ms chirps at 100 Hz, 300 Hz, and 200 Hz. This sequence is readily recognized and quite unlikely to occur by accident. Furthermore, the presence of a 300 Hz chirp in the header prevents accidental detection in the middle of a transmission. Finally, the 200 Hz chirp provides a phase and amplitude reference for the ASK, PSK and QAM transmission schemes, eliminating the need for clock synchronization between sender and receiver.

Decoding is performed on the watch itself (Figure 8), using an optimized decoding routine written in C. The decoder continuously reads samples from the accelerometer, converts the samples to 6400 Hz (to simplify FFT computations), and continuously searches for the header sequence. When found, the decoder demodulates the signal (using the amplitude and phase of the 200 Hz header chirp), performs decoding, verifies the CRC, and reports the resulting message to an application (if decoding was successful).

SYSTEM EVALUATION

Our user studies sought to address critical questions on feasibility, accuracy, and key operating parameters for bio-acoustic sensing across different application contexts. To push the limits of our system even further, we explored several questions relating to robustness and consistency: Are object vibration signatures consistent over time? How robust is the sensing accuracy when the watch is re-worn? Is sensing robust across different locations? Can sensing work on a model that was trained on a different device?

Participants

We recruited 18 participants (10 female, mean age 25.3, 17 right-handed) for a live user study in our lab. Participants were asked to perform a series of tasks while wearing our ViBand prototype. Users wore our ViBand prototype on whichever arm they preferred. Since variations in user anatomy could affect bio-acoustic signal propagation, we recorded user's body mass index (BMI, mean=22.3) to further explore the accuracy of our sensing technique.

The study had three distinct phases, which we discuss in detail subsequently, and lasted approximately 70 minutes in total; participants were paid \$20. Of note, one user had to be dropped from the study because the smartwatch strap did not have a notch that allowed the smartwatch to be ade-

quately secured to their thin arms. All subsequent discussion uses data from the remaining 17 participants.

Setup and Apparatus

The entire study took place in a mixed office-workshop building with two floors and rooms of varying function. Although parts of the study involved accessing tools spread across different rooms and floors, participants were initially welcomed in the lobby. From there, participants were briefed and eventually asked to wear the LG G Watch. Since we rely on physical coupling to the body, our system is susceptible to loose armband tightness. We instructed our participants to wear the watch in a "comfortable but firm" manner. At the end of the study, we had participants complete two Likert-scale questions: participants reported a mean tightness of 4.2 (1=loose, 5=tight), with a mean comfort rating of 3.5 (1=uncomfortable, 5=comfortable).

To verify the robustness of our classifiers across devices, we ran our study using two different smartwatches of the same model (Watch A and Watch B), randomized per user. All machine learning models were trained on Watch A, but deployed and tested on both watches. Data from our watches was streamed to a laptop via a Bluetooth bridge for data recording and live classification.

Study 1: Gesture Recognition

This part of the study aimed to validate if bio-acoustic signals on the body are indeed distinct, and whether they can be used to classify different gesture sets, seen in Figure 5.

Procedure. We trained different machine learning models for each gesture set (Figure 5). Each model was calibrated per-participant, *i.e.*, models were trained for each user. First, the presentation order of the three gesture sets was randomized, and participants were asked to perform a gesture within that set. We collected fifteen data instances per gesture and trained a model *in situ*. Once trained, the participant was asked to perform each gesture once (in random order). Once all gestures were performed for a given set, the participant was asked to remove the watch. After approximately five seconds, participants were asked to re-wear the watch. The participant then performed all of the gestures in that set again (random order), with the classifier output recorded. In total, participants performed two rounds, per gesture, per set.

Results. Across all 17 users and 17 gestures (in all three gesture sets), our system achieved a mean accuracy of 94.3% (SD=4.1%). Figure 11 offers a confusion matrix for each gesture set. We found no statistically significant differences between users and their BMI. There was a slight decrease in accuracy before and after the watch was removed (between rounds one and two), but this was not statistically significant, and so the results have been combined.

Study 2: Object Detection

This study aimed to evaluate whether, 1) bio-acoustic signals could be used to classify grasped objects, 2) object vibrations are consistent over time, and 3) how well the




Figure 10. Object confusion matrix across 29 objects and 17 participants. Results from both testing rounds (pre and post re-wearing of the smartwatch) are combined, yielding 34 trials per object. Chance is 3%.

approach works across different users. We conducted a month-long study to explore this. First, we collected data from one user on 29 objects using a single ViBand prototype (Watch A). The collected data was then used to train a machine learning model. Our example object set and their bio-acoustic signatures are shown in Figure 9.

Procedure. Four weeks later, the same model was used to perform real-time object classification for all 17 participants using the same 29 objects. Objects were spread across six locations to vary environmental conditions. These locations include: personal desk area, shared woodshop, office, kitchen and bathroom, public common area, and a parking space outside of the building. Further, all objects were tested in a location that was different from where it was trained (except the motorcycle). A single trial in our live object classification study involved a user interacting with one of our 29 objects. Participants were briefly shown how to operate the objects (for safety), but were free to grasp the object how-


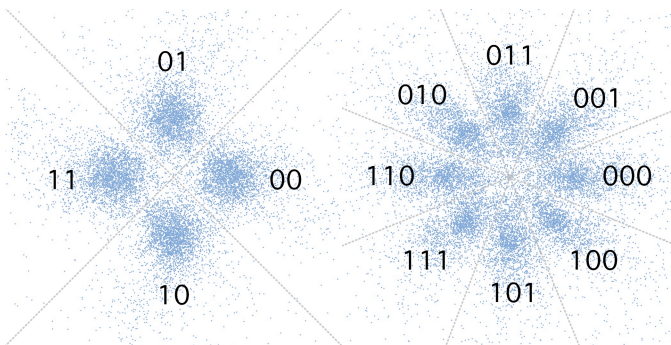


Figure 11. Confusion matrices for one-handed gestures (left, purple), two-handed gestures (mid), and on-body touch input locations (right). Across all users and all gestures, average accuracy was 94%.

ever they wished. Objects were randomized per location (rather than randomized globally). For each location, we performed two classification rounds per object (58 total trials for all 29 objects), with a quick break in between (*i.e.*, go to location one, test objects, break, test again). During each break, we asked the participant to remove the watch from their wrist, and wear it again after ten seconds. This routine is a more realistic measure of the system’s accuracy, as users wear and re-wear smartwatches in the real world.

Results. Across 29 objects, 17 users, and using data that was trained on a single person four weeks prior, we obtained an overall object detection accuracy of 91.5% (SD=4.3%). We found two outlier objects that were 3.5 standard deviations below the mean (using Iglewicz and Hoaglin’s outlier test [21]). When these two outlier objects are removed, we obtain an overall accuracy of 94.0% (27 objects). Figure 10 shows the confusion matrix for all 29 objects. Note that many objects achieve 100% accuracy, despite purposeful inclusion of experimental procedures that usually impact recognition accuracy, *e.g.*, no per-user calibration, significant time separation between train and test, and removal and replacement of the smartwatch *during* the experiment.

Additionally, we found no statistical differences between accuracies on the two watch prototypes. We also found no statistically significant differences between participant’s body-mass index, object location, and whether before/after the watch was removed. Overall, these results suggest that object detection is indeed accurate and robust across users and environment, and object bio-acoustic signatures are consistent over time.




Figure 12. QAM constellation diagrams plotting all data recorded in the study. From left to right: 4-PSK (showing phase boundaries), 8-PSK (showing phase boundaries), non-rectangular 8-QAM (showing symbol construction), and rectangular 16-QAM (showing grid construction).

Study 3: Structured Vibration Data Transfer

We sought to quantify the fidelity of our structured vibrations through a data transmission study. We first ran a pilot study in which we tested several variations of ASK, PSK, FSK and QAM modulation schemes over multiple symbol rate and bits-per-symbol configurations. We rejected configurations that resulted in higher than 10% bit error rates. We chose the five schemes with the highest raw transmission rates: 4-FSK (2 bits per symbol, transmitting frequencies of 50, 100, 150 and 200 Hz), 4-PSK (2 bits per symbol), 8-PSK (3 bits per symbol), 8-QAM (3 bits per symbol, non-rectangular constellation), 16-QAM (4 bits per symbol, non-rectangular constellation). FSK ran at 50 symbols per second, while the other four ran at 100 symbols per second – higher symbol rates were found to be too unreliable.

Procedure. We collected four rounds of data from each user. In each round, the user placed either their outstretched index finger (F) or their whole palm (P) against the transducer, with the watch on the same arm as the contacting hand. Users were randomly assigned one of four possible round orderings: FPPF, PFPF, PFFP, or PFPF. Between the second and third round, the watch was removed and put back on again. We tested both hand and finger contacts in order to determine the error rate difference between these two postures. The hand is a larger contact area close to the wrist, so we anticipated lower error rates there compared to the finger. Each round consisted of 5 data transmission trials for each condition, for a total of 25 trials. Trial order was fully randomized. In each trial, the experiment system (running on a laptop connected to the transducer) transmitted a single packet using the one of the five modulation schemes (Table 2) and waited for 0.5 seconds for the packet to be demodulated. In total, this yielded 1700 trials (17 participants x 4 rounds x 5 conditions x 5 trials per condition)

Results. Out of the 1700 trials collected, no header could be detected in 23 trials (1.4%). These trials were excluded from further analysis. For all remaining trials, we computed the bit error rate by comparing the received, demodulated message with the original transmitted message. The results are summarized in Table 2. Raw bit transmission rate indicates the modulation method’s data transmission speed, while bit error rate (BER) indicates the percentage of bits in the received message that were incorrect. The bit error distribution has a significant long tail across all conditions: most messages are received correctly, but a small number of messages are received with many errors. Therefore, we also computed the 80th percentile BER (BER₈₀), for parity with

	Bit Rate (bits/sec)	BER (hand)	BER (finger)	BER ₈₀ (hand)	BER ₈₀ (finger)
4-FSK	100	1.0%	2.9%	0.0%	0.3%
4-PSK	200	1.0%	3.1%	0.0%	0.6%
8-PSK	300	2.9%	5.8%	3.8%	7.1%
8-QAM	300	3.6%	7.9%	7.7%	15.3%
16-QAM	400	6.9%	8.6%	12.8%	16.0%

Table 2. Data Transmission Results

Ripple [44], to get a better sense of the distribution. This measurement has a practical impact on the choice of error correction parameter: if we choose an error correction scheme that can correct errors up to BER₈₀, then we can expect to successfully decode 80% of transmitted packets.

The results indicate that 4-PSK is the clear winner in terms of BER across all conditions, when considering the raw bit rate. With a BER₈₀ of 0.6% (0.93 message bits), we would need to add only 2 Reed-Solomon ECC symbols to our message in order to correct 80% of messages, leaving 137 bits for the payload. This payload takes 0.83 seconds to transmit (155 bits at 200 bits per second, plus header overhead), for an overall transmission rate of 165 bits per second (with a 20% packet loss rate), through the finger, hand and wrist. This significantly outperforms the most related prior work, OsteoConduct, which operated at “almost 5 bits/sec” [59]. In fact, this performance approaches that of Ripple [44], which obtained an effective bitrate of 196.6 bits per second (using correction up to the 80th BER percentile) transmitting through a cantilevered metal bar (which is obviously far superior to human tissue for transmitting mechanical vibrations).

Study 4: False Positive Rate

In a system that takes advantage of accelerometers, it is critically important to reduce the detection of false positives (*i.e.*, an action that is unintentionally triggered). To validate the resistance of our sensing approach to false positives, we trained our classifier with a large set of *background* data (*i.e.*, negative training examples) and tested the system *live* with our participants. Specifically, our 17 participants were asked to perform several mundane and physically rigorous activities in different locations. These activities included: walking for two minutes, jogging in place for 30 seconds, performing jumping jacks for 30 seconds, reading a magazine or book for one minute, and washing hands for 30 seconds. These five activities were randomly interspersed throughout the object detection study (*i.e.*, when users transitioned between each of the six building locations).




Figure 13. In a smartwatch launcher, we can place navigation controls on the skin. Users can traverse back up through the hierarchy with a flick gesture.



Figure 14. A snap turns on the nearest light. A pinch gesture followed by rotation of the wrist offers continuous brightness control. A flick confirms the manipulation.




Figure 15. Gestures can control remote devices. In the left, the user claps to turn on the TV. Wave gestures navigate and snaps select.

While participants performed these activities, we tallied the number of “false detections” triggered by the system (any prediction that was not “null” or “no object” was considered a false positive). Across 17 users, six random locations, and five activities, collectively spanning a total of 77 minutes, our system triggered a total of six false positive classifications. For 12 of 17 participants, the system triggered no false positives. These results suggest that false positives can be greatly reduced by exposing the machine-learning model to a large set of negative examples.

EXAMPLE APPLICATIONS

We created a series of example applications in the three use domains previously described: gestures, object detection, and data transmission. All example applications were functional, using real-time recognition of bio-acoustic signals (please also see Video Figure).

Expanding Smartwatch Input. Hand gestures can be used to appropriate the area around the watch for input and sensing. For example, in a smartwatch launcher, we can place navigation controls on the skin (*e.g.*, left, right, select), as well as enable users to traverse back up through the hierarchy with a flick gesture (Figure 13).

Controlling Remote Devices. Likewise, gestures can be used to control remote devices. For example, a user can clap to turn on a proximate appliance, such as a TV; wave gestures navigate and snaps offer input confirmation. Flick gestures can be used to navigate up the menu hierarchy (Figure 15).

Input for Infrastructure. Gestures can also be used to control nearby infrastructure. For example, a user can snap his fingers to turn on the nearest light. A pinching gesture can be used as a clutch for continuous brightness adjustment, and a flick confirms the manipulation (Figure 14).

Object-Aware Applications. Because our sensing approach can also be used to identify objects, we offer applications the ability to better understand context and augment every-



Figure 17. For an unpowered object *e.g.*, an acoustic guitar, we can detect the closest note to tune the instrument precisely.



Figure 16. Object sensing enables rich, context-sensitive applications. For example, we can sense equipment used in the preparation of a meal (left). A progress bar informs the user when to stop blending ingredients.

day activities. For example, we can augment the kitchen experience by sensing equipment used in the preparation of a meal and *e.g.*, offering a progress indicator for blending ingredients with an egg mixer (Figure 16).

Detecting Unpowered Objects. Our technique can also sense unpowered objects, such as an acoustic guitar. For example, we can detect the closest note whenever the guitar is grasped, and provide visual feedback to tune the instrument precisely (Figure 17). Detection happens *on touch*, which makes it robust to external noise in the environment.

Augmenting Analog Experiences. Through object sensing, we can also augment *analog* experiences with *digital* interactivity. For example, with a Nerf gun, we can detect the loading of a new ammo clip, and then keep count of the number of darts remaining (Figure 18).

Vibro-Tags. Many classes of objects do not emit characteristic vibrations, which means ViBand cannot detect them. However, we can instrument them with a *vibro-tag* that emits inaudible, structured vibrations containing data. For example, we can instrument a glue gun (non-mechanical but electrically powered) with a vibro tag. The tag broadcasts an object ID that enables the watch to know what object is being held. It also transmits metadata *e.g.*, its current temperature and ideal operating range (Figure 19, top).

Tagging Infrastructure. Structured vibrations are also valuable for augmenting fixed infrastructure with dynamic data or interactivity. For example, in an office setting, a user can retrieve more information about an occupant by touching the room nameplate augmented with a vibro tag, which transmits *e.g.*, the person’s contact details to the smartwatch (Figure 19, bottom).

DISCUSSION

We have demonstrated that high-speed accelerometer sampling in smartwatches offers new and compelling human interface possibilities, and we hope that this research en-



Figure 18. We can augment “analog” experiences with “digital” interactivity. With a Nerf gun, we can detect users loading ammo clips (left) and firing darts (middle).

courages manufacturers to expose this useful data source in future devices. Dedicated microcontrollers that sit between sensors and the application processor (“sensor hubs”) are already employed in many devices to help improve power efficiency. Sensor hubs could easily be used to buffer and process high-speed accelerometer data, enabling low-power, always-on, bio-acoustic applications.

Our technique should be readily portable to most smartwatches, as modern IMUs have comparable specs. Indeed, InvenSense is one of the largest IMU vendors, and as mentioned previously, the same series of accelerometers we use is also used in many other popular smartwatches. We also saw no performance difference when training or testing on two different watches of the same model. However, we suspect that entirely different models of smartwatch would alter the physical coupling slightly. We ran some basic tests, and this effect appears to be minor compared to the *active* signal of the gesture, object or data transmitting transducer.

We note that unintended and competing oscillations (e.g., bus/walking) inherently decrease the signal-to-noise ratio (SNR). For human actions, like locomotion, we observed these chiefly happen within lower frequency bands (roughly 0-20 Hz), which are easily filtered. Overall, similar to many deployed technologies, noise robustness can be improved through e.g., adaptive background subtraction or by incorporating diverse negative training examples.

Finally, applying structured vibrations to large objects or surfaces (e.g., tables) can result in audible noise (by essentially turning the surface into a amplifying diaphragm). For this reason, we used a vibration transducer with a small active area (as opposed to a voice coil with a large diaphragm), so airborne emissions were limited (as demonstrated in our Video Figure). Additionally, we note that malicious interception of vibration-borne data might be possible, especially with e.g., a high quality directional microphone or laser Doppler vibrometer. Whether all bits could be resolved is an open question.

CONCLUSION

To summarize, in this work, we explored bio-acoustic sensing on commodity smartwatches, introducing a wide range of novel interaction modalities and use cases. More importantly, our contributions unlock user interface techniques that previously relied on special-purpose hardware. The applications we describe and demonstrate could be deployed to existing smartwatches with an over-the-air update. Our evaluations show that our sensing technique can be accurate, robust to noise, and reliable across users. This has the potential to make smartwatches more useful, and complex interactions on them more practical.

ACKNOWLEDGMENTS

This research was generously supported by the David and Lucile Packard Foundation, a Google Faculty Research Award, and Qualcomm. We also thank Liz Carter and Yang Zhang for their generous help with user studies and demos.



Figure 19. Vibro-tags emit structured vibrations that carry data. For example, the user can touch an office nameplate instrumented with a vibro tag (top), which transmits the person’s contact details to the smartwatch. We can also instrument non-mechanical objects, e.g., a glue gun (bottom). The tag broadcasts an object’s ID and metadata — in this case, the current temperature and ideal operating range.

REFERENCES

1. Akl, A., Feng, C. and Valaee, S. A Novel Accelerometer-Based Gesture Recognition System. *IEEE Transactions on Signal Processing* '11, 12: 6197-6205.
2. Amento, B., Hill, W. and Terveen, L. The sound of one hand: a wrist-mounted bio-acoustic fingertip gesture interface. In *Proc. CHI EA '02*, 724-725.
3. AT&T Labs. AT&T research on the next secure data transfer device: your body. June 1, 2012. <http://innovationblogarchive.att.com/innovation/story/a/7782685>
4. Bellamy, J.C. Digital Telephony, 3rd Edition. Chapter 6.1 Digital Modulation, pp. 279-308.
5. Bernaerts, Y., Druwé, M., Steensels, S., Ermeulen, J. and Schöning, J. The office smartwatch: development and design of a smartwatch app to digitally augment interactions in an office environment. In *Proc. DIS '14*, 41-44.
6. Buchler, M.C. Algorithms for Sound Classification in Hearing Instruments. PhD thesis, ETH Zurich, 2002.
7. Buettner, M., Prasad, R., Philipose, M. and Wetherall, D. Recognizing daily activities with RFID-based sensors. In *Proc. UbiComp '09*, 51-60.
8. Clarkson, B., Sawhney, N. and Pentland, A. Auditory context awareness in wearable computing. In *Workshop on Perceptual User Interfaces*, November 1998.
9. Cornelius, C., Peterson, R., Skinner, J., Halter, R. and Kotz, D. A wearable system that knows who wears it. In *Proc. MobiSys '14*, 55-67.
10. Dementyev, A. and Paradiso, J.A. WristFlex: Low-power gesture input with wrist-worn pressure sensors. In *Proc. UIST '14*, 161-166.

11. Deyle, T., Palinko, S., Poole, E.S. and Starner, T. Hambone: A Bio-Acoustic Gesture Interface. In *Proc. ISWC '07*, 1-8.
12. Dimoulas, C., Kalliris, G., Papanikolaou, G. and Kalampakas, A. Long-term signal detection, segmentation and summarization using wavelets and fractal dimension: A bio-acoustics application in gastrointestinal-motility monitoring. *Comput. Biol. Med.* 37, 4 (April 2007), 438-462.
13. de Castro Lopo, E. libsamplerate (software). <http://www.mega-nerd.com/SRC/index.html>
14. Fogarty, J., Au, C. and Hudson, S.E.: Sensing from the Basement: A Feasibility Study of Unobtrusive and Low-Cost Home Activity Recognition. In *Proc. UIST '06*, 91-100.
15. Fry, W.J. (1958). Biological and medical acoustics. *The Journal of the Acoustical Society of America*, 30(5), 387-393.
16. Fukui, R., Watanabe, M., Gyota, T., Shimosaka, M. and Sato, T. Hand shape classification with a wrist contour sensor: development of a prototype device. In *Proc. Ubicomp '11*, 311-314.
17. Google, Inc. Compatibility Definition for Android 6.0. October 16, 2015. <http://source.android.com/compatibility/android-cdd.pdf>
18. Hansen, P. 2001. Recent Bio-acoustical Publications, 1999 and earlier. Part 1: Invertebrates - birds. *Bio-acoustics* 11(3), 223-262
19. Harrison, C., Tan, D. and Morris, D. Skinput: Appropriating the Body as an Input Surface. In *Proc. CHI '10*, 453-462.
20. Hodges, S., Thorne, A., Mallinson, H. and Floerkemeier, C. Assessing and optimizing the range of UHF RFID to enable real-world pervasive computing applications. In *Pervasive Computing '07*. 280-297.
21. Iglewicz, B. and Hoaglin, D. (1993), Volume 16: How to Detect and Handle Outliers., *The ASQC Basic References in Quality Control: Statistical Techniques*, Edward F. Mykytka, Ph.D., Editor.
22. InvenSense, Inc. MPU-6500 Product Specification Revision 1.0. September 18, 2013. https://store.invensense.com/datasheets/invensense/MPU_6500_Rev1.0.pdf
23. InvenSense, Inc. MPU-6500 Register Map And Descriptions Revision 2.1. September 16, 2013. <https://www.invensense.com/wp-content/uploads/2015/02/MPU-6500-Register-Map2.pdf>
24. Jung, P.G., Lim, G., Kim, S. and Kong, K. A Wearable Gesture Recognition Device for Detecting Muscular Activities Based on Air-Pressure Sensors. *IEEE Transactions on Industrial Informatics*, 11(2), 485-494. Feb. 2015.
25. Kim, D., Hilliges, O., Izadi, S., Butler, A., Chen, J., Oikonomidis, I. and Olivier, P. Digits: freehand 3D interactions anywhere using a wrist-worn gloveless sensor. In *Proc. UIST '12*, 167-176.
26. Lange, B.M., Jones, M.A. and Meyers, J.L. Insight lab: an immersive team environment linking paper, displays, and data. In *Proc. CHI '98*, 550-557.
27. Laput, G., Yang, C., Xiao, R., Sample, A. and Harrison, C. EM-Sense: Touch Recognition of Uninstrumented, Electrical and Electromechanical Objects. In *Proc. UIST '15*, 157-166.
28. Li, H., Ye, C. and Sample, A. IDSense: A Human Object Interaction Detection System Based on Passive UHF RFID. In *Proc. CHI '15*, 2555-2564
29. Lowe, D.G. Object recognition from local scale-invariant features. In *Proc. ICCV '99*, 1150-1157.
30. Maekawa, T., Kishino, Y., Sakurai, Y. and Suyama, T. Recognizing the use of portable electrical devices with hand-worn magnetic sensors. In *Proc. Pervasive '11*. 276-293
31. Maekawa, T., Kishino, Y., Yanagisawa, Y. and Sakurai, Y. Recognizing handheld electrical device usage with hand-worn coil of wire. In *Proc. Pervasive'12*. 234-252.
32. Medwin, H., Clay C.S. 1998. *Fundamentals of Acoustical Oceanography*, Academic Press.
33. Mujibiya, A., Cao, X., Tan, D.S., Morris, D., Patel, S.N. and Rekimoto, J. The sound of touch: on-body touch and gesture sensing based on transdermal ultrasound propagation. In *Proc. ITS '13*, 189-198.
34. Ogata, M. and Imai, M. SkinWatch: skin gesture interaction for smart watch. In *Proc. AH '15*, 21-24.
35. Ogata, M., Sugiura, Y., Makino, Y., Inami, M. and Imai, M. SenSkin: adapting skin as a soft interface. In *Proc. UIST '13*, 539-544.
36. Peltonen, V., Tuomi, J., Klapuri, A., Huopaniemi, J. and Sorsa, T. Computational auditory scene recognition. In *IEEE Acoust, Speech, and Signal Proc.* 1520-6149.
37. Perng, J.K., Fisher, B., Hollar, S. and Pister, K.S. Acceleration sensing glove. In *Proc. ISWC '09*, 178-180.
38. Qian, Z., Sagers, R.D. and Pitt, W.G. (1999), Investigation of the mechanism of the bio-acoustic effect. In *J. Biomed. Mater. Res.*, 44: 198-205.

39. Reed, I.S. and Solomon, G. Polynomial Codes over Certain Finite Fields. *J. Soc. Ind. Appl. Math.* 8(2), 300-304 (1960).
40. Rekimoto J. and Ayatsuka, Y. CyberCode: designing augmented reality environments with visual tags. In *Proc. DARE '00*, 1-10
41. Rekimoto, J. Gesturewrist and gesturepad: Unobtrusive wearable interaction devices. In *Proc. ISWC '01*, 21-27.
42. Ren, X. Egocentric recognition of handled objects: Benchmark and analysis. In *Proc. CVPR Workshops '09*, 2160-7508.
43. Roy, N. and Choudhury, R. Ripple II: Faster Communication through Physical Vibration. In *Proc. NSDI '16*, USENIX Association, 671-684.
44. Roy, N. Gowda, M. and Choudhury, R. Ripple: communicating through physical vibration. In *Proc. NSDI'15*, USENIX Association,, 265-278.
45. Saponas, T.S., Tan, D.S., Morris, D. and Balakrishnan, R. Demonstrating the feasibility of using forearm electromyography for muscle-computer interfaces. In *Proc. CHI '08*, 515-524.
46. Saponas, T.S., Tan, D.S., Morris, D., Balakrishnan, R., Turner, J. and Landay, J. A. Enabling always-available input with muscle-computer interfaces. In *Proc. UIST '09*, 167-176.
47. Sato, M., Poupyrev, I. and Harrison, C. Touché: enhancing touch interaction on humans, screens, liquids, and everyday objects. In *Proc. CHI '12*, 483-492.
48. Simmonds, J. and MacLennan, D. 2005. *Fisheries Acoustics: Theory and Practice, second edition*. Blackwell Press.
49. Torralba, A., Murphy, K.P., Freeman, W.T. and Rubin, M.A. Context-based vision system for place and object recognition. In *Proc. ICCV '03*, 273-280.
50. Wang, E.J., Lee, T.J., Mariakakis, A., Goel, M., Gupta, S. and Patel, S.N. MagnifiSense: inferring device interaction using wrist-worn passive magneto-inductive sensors. In *Proc. UbiComp '15*, 15-26.
51. Ward, J.A., Lukowicz, P., Tröster, G. and Starner, T.E. Activity recognition of assembly tasks using body-worn microphones and accelerometers. *IEEE Trans. on Pattern Analysis & Mach. Intelligence '06*. 1553-1567.
52. Way, D. and Paradiso, J. A usability user study concerning free-hand microgesture and wrist-worn sensors. In *Proc. BSN '14*, 138-142.
53. Wen, H., Rojas, J.R. and Dey, A. Serendipity: Finger Gesture Recognition using an Off-the-Shelf Smartwatch. In *Proc. CHI '16*, 3847-3851.
54. Xu, C., Pathak, P.H. and Mohapatra, P. Finger-writing with Smartwatch: A Case for Finger and Hand Gesture Recognition using Smartwatch. In *Proc. HotMobile '15*, 9-14.
55. Zhang, Y. and Harrison, C. Tomo: Wearable, Low-Cost, Electrical Impedance Tomography for Hand Gesture Recognition. In *Proc. UIST '15*, 167-173.
56. Zhao, N., Dublon, G., Gillian, N., Dementyev, A. and Paradiso, J.A. EMI Spy: Harnessing electromagnetic interference for low-cost, rapid prototyping of proxemic interaction. In *Proc. Wearable and Implantable Body Sensor Networks 2015*, 1-6.
57. Zhao, W., Chellappa, R., Phillips, P. and Rosenfeld, R. Face recognition: A literature survey. *ACM Comput. Surv. '03 Vol. 35-4*. 399-458.
58. Zhao, Y., LaMarca, A. and Smith, J.R. A battery-free object localization and motion sensing platform. In *Proc. UbiComp '14*. 255-259.
59. Zhong, L., El-Daye, D., Kaufman, B., Tobaoda, N., Mohamed, T. and Libschner, M. OsteoConduct: wireless body-area communication based on bone conduction. In *Proc. BodyNets '07*. ICST, Article 9, 8 pages.
60. Zimmerman, T.G. 1996. Personal area networks: near-field intrabody communication. *IBM Syst. J.* 35, 3-4 (September 1996), 609-617.