EM-Sense: Touch Recognition of Uninstrumented, Electrical and Electromechanical Objects

Gierad Laput[1] **Chouchang Yang**[2] **Robert Xiao**[1] **Alanson Sample**[2] **Chris Harrison**[1]

[1] Carnegie Mellon University, HCI Institute
5000 Forbes Avenue, Pittsburgh, PA 15213
{gierad.laput, brx, chris.harrison}@cs.cmu.edu

[2] Disney Research Pittsburgh
4720 Forbes Ave, Suite 110, Pittsburgh, PA 15213
{jack.yang, alanson.sample}@disneyresearch.com

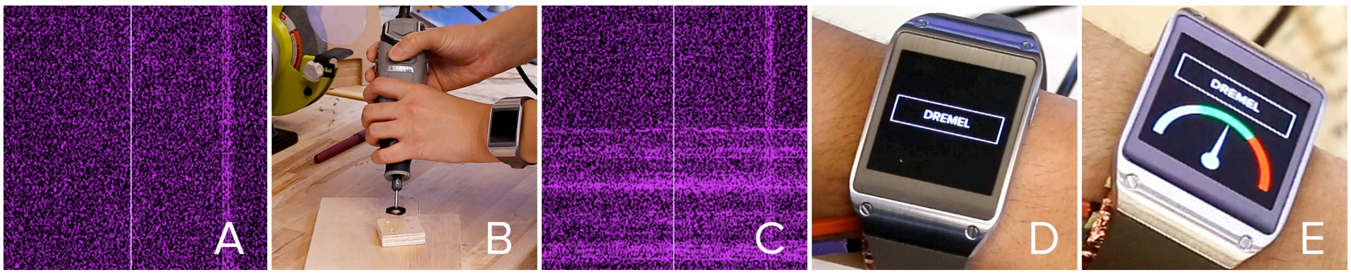

Figure 1. Spectrogram of ambient electromagnetic noise (A). When a user operates an electrical or electromechanical object, such as a Dremel (B), it emits EM noise (C), which we classify (D) and use to enable rich contextual applications (E).

ABSTRACT

Most everyday electrical and electromechanical objects emit small amounts of electromagnetic (EM) noise during regular operation. When a user makes physical contact with such an object, this EM signal propagates through the user, owing to the conductivity of the human body. By modifying a small, low-cost, software-defined radio, we can detect and classify these signals in real-time, enabling robust on-touch object detection. Unlike prior work, our approach requires no instrumentation of objects or the environment; our sensor is self-contained and can be worn unobtrusively on the body. We call our technique *EM-Sense* and built a proof-of-concept smartwatch implementation. Our studies show that discrimination between dozens of objects is feasible, independent of wearer, time and local environment.

Author Keywords

Context Sensitive; Object Detection; Smartwatch; EMI; Wearable Computing; Smart Clothes; Tags.

ACM Classification Keywords

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

INTRODUCTION

For years, intelligent systems have promised to improve people's lives by inferring context and activities in diverse environments. In particular, people's interactions with objects offer rich, contextual information closely reflecting one's immediate activity. Yet practical detection and recognition of object interactions remains an elusive research goal. For example, although RFIDs can provide object recognition capabilities, the technology requires all desired objects to be physically tagged and it is unknown if users are simply nearby or truly touching an object.

We propose a novel sensing approach for object detection, triggered only when objects are *physically touched*. Our approach exploits unintentional EM noise emitted by many everyday electrical and electromechanical objects, such as kitchen appliances, computing devices, power tools and automobiles. These signals tend to be highly characteristic, owing to unique internal operations (e.g., brushless motors, capacitive touchscreens) and different enclosure designs, material composition and shielding. When a user makes physical contact with these objects, electrical signals propagate through the user's body, as it is conductive. By modifying a commodity software-defined radio receiver, we can detect and classify these signals in real time, enabling robust, on-touch object detection.

Our approach, which we call *EM-Sense*, utilizes low-cost, commodity hardware and is small enough to be worn on the wrist or, in the near future, integrated into smartwatches. We draw inspiration from the sensing principles introduced in Humantenna [8], and move beyond environment localization and gesture recognition, to focus instead on context and activity sensing made possible through object interaction detection. Unlike existing approaches requiring object instrumentation (RFIDs, barcodes, BLE beacons, etc.), EM-Sense can identify objects solely on their EM signatures.

UIST '15, November 8-11, Charlotte, NC, USA

ACM 978-1-4503-3779-3/15/11 $15.00. http://dx.doi.org/10.1145/2807442.2807481

Our work makes the following contributions:

- A sensing approach utilizing low-cost, wearable hardware for recognizing handled electrical and electromechanical objects based on body-coupled EM signatures;
- A novel hardware and software implementation, transforming a low-cost software defined radio receiver into a fast, wideband, general purpose EM sensor;
- A series of studies showing that our sensing approach is accurate, robust, and consistent across users, environments and time; and
- A series of example applications illustrating powerful assistive, context-sensing, and communication applications across a wide range of activities and environments.

BACKGROUND

Electronic devices, especially those driven by motors (e.g., power drills) or switching power supplies (e.g., LCD screens), produce significant levels of electromagnetic noise. These unwanted signals propagate as radio frequency (RF) waves and can disrupt nearby devices operating with similar frequency bands. Beginning in the late 1970s, the US Federal Communications Commission (FCC) chartered mandates to regulate the susceptibility of consumer devices to EM noise [33, 43]. These were also established to prevent EM noise from interfering with other electronics, utilities, and purposeful broadcasts, such as TV and radio.

Infrastructures and environments also produce EM noise. For example, AC electricity and devices connected to the power line contribute to majority of electrical noise at home. In general, EM noise propagates through conduction over circuits and power lines (1kHz - 30MHz) or through radiation in free space (30MHz to 10GHz).

Additionally, a few classes of non-electromechanical objects can have unique EM signatures. Most notable among these are large, metallic objects, such as structural members in buildings, doors, ladders, furniture, and window framework. These are sufficiently large that they act as antennas, capturing energy radiated by proximate noisy devices and wiring. The amalgamation of these signals, in our experiences, is fairly unique and particular to a location.

RELATED WORK

Our work intersects with several bodies of research, and we now summarize key areas.

Acitivity and Object Recognition

Traditional activity recognition systems infer user state based on temporal data of physical movement (e.g., accelerometers). These require individuals to wear sensors or have a smartphone continuously monitoring data. Extensive prior work [2, 18, 38] has demonstrated promising results for determining e.g., running, walking and sitting. However, motion-driven approaches by themselves lack context to infer higher-level activities.

For this reason, we pursue a complimentary approach that recognizes *handled objects*. This provides relevant information more closely reflecting a user's immediate environment and activity [13]. Many approaches have been considered for object recognition, though most methods require objects to be instrumented with some form of marker or sensor [4, 23, 39, 42, 46, 47]. These can provide robust recognition, but as there are many objects in the world, installation and maintenance is troublesome and costly.

Recent work from Maekawa and colleagues cleverly utilized magnetic sensors [30] and hand-worn coils [31] to detect objects based on temporal changes in the magnetic field during an object's operation. Although related, magnetic induction relies on proximate contact between objects and the sensing apparatus, which means object detection is strongly affected by hand posture and inherent magnetic noise in the body, or even diamagnetic properties of hands and fingers. Conversely, as we will show, our approach is robust across users, time and hand/body posture.

Visual, Acoustic and RF-Based Approaches

Early research into visual markers used 1D barcodes, and more recently, fiducial markers [26, 40] as unique identifiers. Further, there is considerable work in the computer vision domain for object recognition in natural scenes without artificial markers [9, 41, 49], as well as efforts that leverage crowd workers [27, 28]. These schemes require cameras, line of sight, and suitable lighting conditions.

Acoustic-based object recognition has also been explored extensively [3, 6, 36, 45]. For example, Acoustic Barcodes [20] described tags with sound-producing physical notches that resolve to a binary ID. More related to our approach are acoustic methods that attempt to recognize objects from vibro-acoustic information generated by operation of a device. For example, Ward et al. [45] used worn accelerometers and microphones to classify workshop tools.

Also popular are RFID-based approaches. Example systems include a wrist-worn, near-field RFID reading system that could identify objects affixed with tiny RFID tags [29,37]. Similarly, Buettner et al. [4] used the Wireless Identification and Sensing Platform (WISP), which is a battery-free, long range RFID tag enhanced with an accelerometer to detect movement of a tagged object. Other object recognition efforts exist that use wifi sensing [44], NFCs [15], Bluetooth Low Energy [12], and body-area networks [32].

EM-Based Sensing

There are two main classes of EM-based sensing techniques: 1) infrastructure-mediated sensing and 2) using the human body as an antenna. The former instruments the infrastructure, while the second instruments the user.

Infrastructure-Mediated Sensing. Early work by Abott [1] and Hart [21, 22] in the 1980s used metering devices attached to a building's electrical lines to detect "events" caused by home appliances. Because the electrical lines in a house are shared infrastructure, a single sensor can observe activity across the entire home. These pioneering efforts inspired infrastructure-mediated sensing (IMS), i.e., attach-

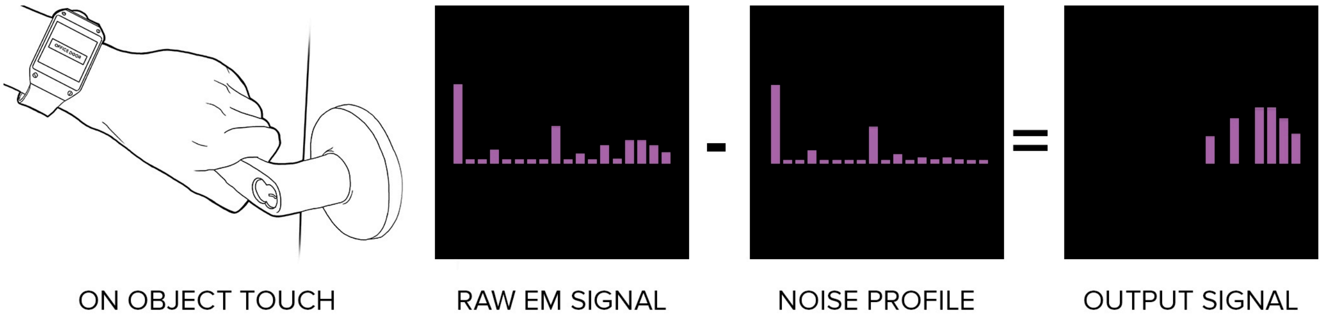

Figure 2. EM-Sense captures EM signals emitted by electrical and electromechanical objects. Raw EM signals are subtracted against an adaptive noise profile. Next, frequencies whose magnitudes are above a z-score threshold (e.g., 3.5 above the standard deviation) are filtered and amplified. Features are extracted from the resulting output signal for machine learning classification.

ing probes to a variety of utility infrastructures, including HVACs [34], plumbing [14], natural gas lines [7], lighting [16] and electrical wiring [10, 17, 25, 35].

Using the Human Body as an Antenna. Because the human body is conductive, it has electrical properties that allow it to behave much like an antenna. Pioneering work in HCI has exploited this "body antenna effect." For example, in DiamondTouch [11], the human body is used as an electrical conductor, which allows the system to differentiate touches between users. More recently, in "Your Noise is My Command," Cohn et al. [9] utilize the human body as an antenna for detecting EMI signals in the home. A small electrode is attached behind the neck of the user and connected to a backpack-bounded A/D converter. As the user moves around the home, the system captures all recorded EM noise received by the human antenna. With this setup, they inferred user location within a home, as well as detect different gestures and continuous touch tracking along a wall. A later extension enabled free-space, whole body gestures by utilizing EM Doppler shifts [8]. Unlike infrastructure-mediated sensing, body-antenna EM sensing requires no instrumentation of the environment.

EM-SENSE IMPLEMENTATION

As mentioned, EM-Sense exploits the unintentional electromagnetic noise emitted by everyday electrical and electromechanical objects during regular operation. In this section, we describe our sensing technique and processing framework for real-time on-touch object detection.

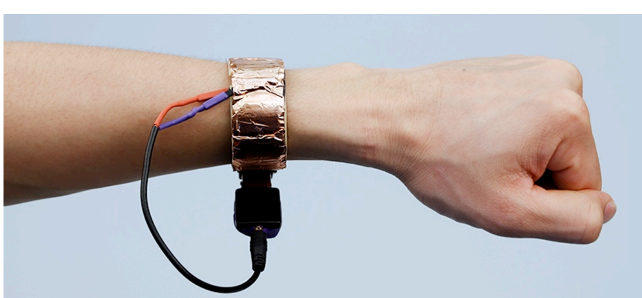

Figure 3. EM-Sense wrist-worn prototype. Body-coupled EM signals are captured through a modified SDR receiver.

Proof-of-Concept Setup

For our proof-of-concept hardware implementation (Figure 3), we modified a software-defined radio receiver (RTL-SDR) to function as an inexpensive, but extremely fast A/D converter (Figure 4). Originally, RTL-SDRs use a tuning chipset to listen for frequencies covering FM bands and beyond (25 – 1766 MHz). However, useful EM emissions for most objects fall well below this operating range. To address this limitation, we modified the RTL-SDR's circuitry by bypassing its tuner and routing raw antenna signals directly into the chipset's main A/D converter. It has an 8-bit resolution with ~2Vpp.

As a result, this modification re-adjusts the device's sensing range to 1Hz – 28.8MHz, making it possible to detect low-band EM signals present in many electrical and electromechanical objects. Figure 5 details three simple steps describing our modifications, which researchers and hobbyists can replicate. Our sensing setup costs under $10, two orders of magnitude cheaper than previous EM sensing approaches.

To make the prototype wearable, we retrofitted the interior of an armband with copper tape, and connected it to the RTL-SDR's antenna terminal. Data received from the RTL-SDR is further processed through a software pipeline. First, we read from the RTL-SDR's input channel through a physical USB connection. At the time of research, no smartwatch on the market was capable of hosting a USB-OTG interface. In response, we offload USB-reading to a smartphone (Nexus 5), which is clipped to waist of the wearer, and uses an open source RTL-SDR driver (osmocom.org) we ported. In addition to reading the SDR, our smartwatch software also streams incoming data to a laptop computer over wifi, which in turn performs signal processing and live classification. With this setup, we can wirelessly stream data sampled at 1MHz with minimal packet loss. All of the physical components of EM-Sense fit into a wrist-worn apparatus, which could be easily integrated into future smartwatches.

Sensing Raw EM Signals

Whenever a user makes physical contact with an electrical or electromechanical object, its EM signal propagates

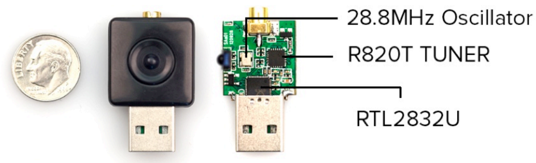

Figure 4. The software-defined radio (SDR) receiver we used for our EM-Sense prototype. 8-bit resolution, 2Vpp.

through the body and is sensed by a conducting electrode worn on the user's wrist (Figures 2 and 3). Connected to the electrode is our modified software-defined radio, which converts this analog signal into digital data. We sample incoming signals at 1MHz; thus our theoretical Nyquist limit is 500kHz.

We note that neighboring objects and signals (e.g., objects proximate to the user, but not in direct physical contact) can introduce interference through capacitive coupling and the body antenna effect. However, these signals are comparatively weak compared to those transmitted by actual physical contact, and do not appear to affect detection.

Baseband Shift

To extend our effective bandwidth, we shift the SDR receiver's baseband frequency to 500kHz. Without shifting, the frequency spectrum is symmetric because it is a real signal. In this mode, the effective bandwidth for a signal sampled at 1Ms/s is -0.5MHz to 0.5MHz (i.e., see Fig 6 raw FFT, where left half is redundant). Shifting to 500kHz moves the bandpass sampling window from 0 to 1MHz (0.5MHz and above will be undersampled, but still useful). As a result, the left-shifted spectrum contains no redundant information. We then apply a fast Fourier Transform (with an FFT size of 16384 bins i.e., 61 Hz per band), and the resulting frequency domain values become the primary input for our sensing and classification pipeline.

Enviornmental Noise Rejection

To enable robust object detection, our sensing approach must differentiate between environmental EM *noise* and EM *signals* from objects. In addition to differences in amplitude (touched objects generally transmit more signal), we also take advantage of the fact that environmental EM noise tends to change at a slower rate, while EM signals change rapidly at the moment an object is touched or released (or the object is turned on/off). These events appear as as high delta, "spiking" events in the signal.

We build a model of environmental EM noise using an adaptive background subtraction approach: an average frequency spectrum derived from a six-second rolling window, updated every 100ms (Figures 2 and 6). This provides a baseline "noise profile" from which we can subtract the live signal, amplifying transitions in touch state. In this particular implementation, if an object is held for a few seconds, its EM signal is integrated into the noise profile. The release of an object thus generates a large negative change, which is interpreted as a "touch up" event, signifying the object is no longer held.

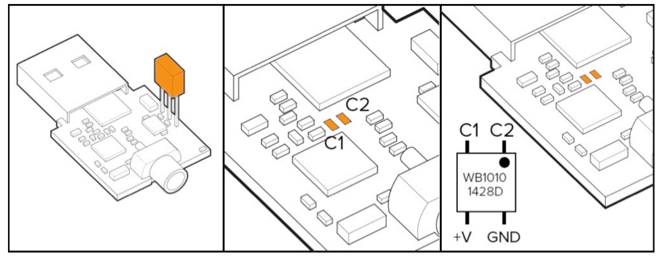

Figure 5. We modified an RTL2832U SDR receiver into a low-cost EMI sensor. First, we removed the IR receiver (left). Next, we removed C1 and C2 capacitors (mid), and attached a wideband transformer (right).

Object Signal Extraction

To extract EM signals generated from object on-touch events, we perform real-time statistical analysis between the modeled noise profile and all incoming EM readings. We compute a baseline *threshold signal* based on the statistical *Z-score* of the individual frequency bands in the noise profile. Essentially, frequency bands whose values are above a specified *Z-score* (e.g., 3.5 standard deviations above the noise profile) are amplified, while frequencies below the threshold are set to zero.

Thus, a frequency band at index **n** of the extracted EM signal, ***S***, can be characterized as:

$$S_n = A \times \max(0, F_n - (G_n + z\sigma_n))$$

where ***F*** is the incoming EM reading, ***G*** is the noise profile, **σ** holds the standard deviations for each frequency band at index **n**, ***A*** denotes the amplification factor, and **z** is a constant that denotes a statistical z-score parameter. In our implementation, we use an amplification factor of 18 and a *z-score* of +3.5 (upper 0.1% of a normal distribution curve).

Live Object Classification

Once an object's EM signature is decoupled from environmental noise (Figure 2), we use it as input for live object classification. First, we downsample the EM signature's FFT into 512 frequency bands. From this, we generate ~2K additional features based on: 1^{st} and 2^{nd} Derivatives (1021), min index, max index, RMS, center of mass, standard deviation, area under the curve, pair-wise band ratios (496), spectral kurtosis, crest factor, and 2^{nd} order FFT (512).

These features are fed into a SMO-trained Support Vector Machine (c=1.0, $\epsilon=1^{-12}$, poly kernel) provided by the Weka Toolkit [19]. Feature selection analysis revealed that derivatives, band ratios, 2^{nd} order FFTs, and max index serve as the important distinguishing features (providing 80% merit), but the remaining features nonetheless are important to fully capturing nuanced signal behaviors. Other machine learning techniques could potentially allow EM-Sense to scale to larger collections of objects. Object classification can be treated as an "information retrieval" problem, which means that techniques such as clustering, similarity metrics, and deep-learning methods are applicable.

EXAMPLE USE SCENARIOS

To demonstrate how EM-Sense can augment activities across a wide range of contexts and environments, we provide a usage narrative contextualized in a hypothetical user's day. Although meant to be illustrative, we built fully functional versions of every application described here (see Video Figure, filmed live, and Figures 7 through 9). This narrative includes five categories of objects: *home, office, workshop*, *fixed structures*, and *transportation* – a taxonomy we employ in our subsequent evaluation.

Home – At home, Julia wakes up and gets ready for another productive day at work. Her EM-Sense-capable smartwatch informs and augments her activities throughout the day. For instance, when Julia grabs her electric toothbrush, EM-Sense automatically starts a timer (Figure 7A). When she steps on a scale, a scrollable history of her weight is displayed on her smartwatch automatically (Figure 7B). Down in the kitchen, EM-Sense detects patterns of appliance touches, such as the refrigerator and the stove. From this and the time of day, EM-Sense infers that Julia is cooking breakfast and fetches the morning news, which can be played from her smartwatch (Figure 7D).

Fixed Structures – When Julia arrives at the office, EM-Sense detects when she grasps the handle of her office door. She is then notified about imminent calendar events and waiting messages: *"You have 12 messages and a meeting in 8 minutes"* (Figure 8A). Julia then leaves a reminder – tagged to the door handle – to be played at the end of the day: *"Don't forget to pick up milk on the way home."*

Workshop – In the workshop, EM-Sense assists Julia in her fabrication project. First, Julia checks the remaining time of a 3D print by touching anywhere on the print bed – "five minutes left" (Figure 9A) – perfect timing to finish a complementary wood base. Next, Julia uses a Dremel to cut a piece of wood. EM-Sense detects the tool and displays its rotatory speed on the smartwatch screen (Figure 9B). If it knows the task, it can even recommend the ideal speed. Similarly, as Julia uses other tools in the workshop, a tutorial displayed on the smartwatch automatically advances (Figures 9C and 9D). Finally, the 3D print is done and the finished pieces are fitted together.

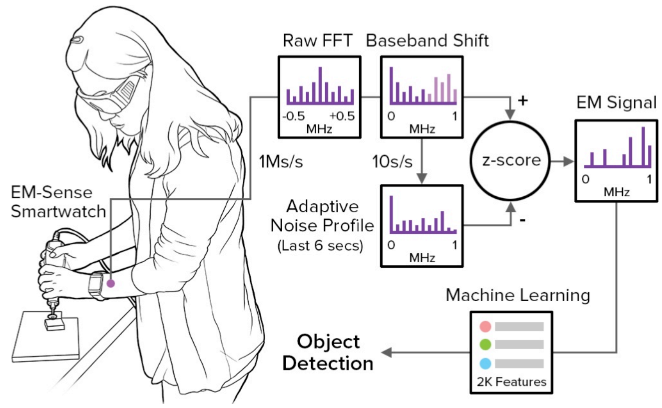

Figure 6. EM signals coupled through the body are captured at 1M samples per second, and baseband shifted to 500kHz. Next, it is compared against an adaptive noise profile using Z-score thresholding. The resulting signal is fed to a machine learning pipeline for object classification.

Office – Back at her desk, Julia continues work on her laptop. By simply touching the trackpad, EM-Sense automatically authenticates Julia without needing a password (Figure 8B). Later in the day, Julia meets with a colleague to work on a collaborative task. They use a large multitouch screen to brainstorm ideas. Their EM-Sense-capable smartwatches make it possible to know when each user makes contact with the screen. This information is then transmitted to the large touchscreen, allowing it to differentiate their touch inputs. With this, both Julia and her colleague can use distinct tools (e.g., pens with different colors); their smartwatches provide personal color selection, tools, and settings (Figure 8C).

Transportation – At the end of the day, Julia closes her office door and the reminder she left earlier is played back: *"Don't forget to pick up milk on the way home."* In the parking lot, Julia starts her motorcycle. EM-Sense detects her mode of transportation automatically (e.g., bus, car, bicycle) and provides her with a route overview: *"You are 10 minutes from home, with light traffic"* (Figure 8D).

EM-SENSE INTERACTIONS

We built our example use scenario, described above, around six interaction categories, which we describe briefly.

Figure 7. EM-Sense can augment activities in the home. For example, EM-Sense can launch a timer (inset) when the user is brushing his teeth (A), or display the user's data when stepping on a scale (B). Next, EM-Sense knows that the user is making breakfast and fetches the morning news (C and D).

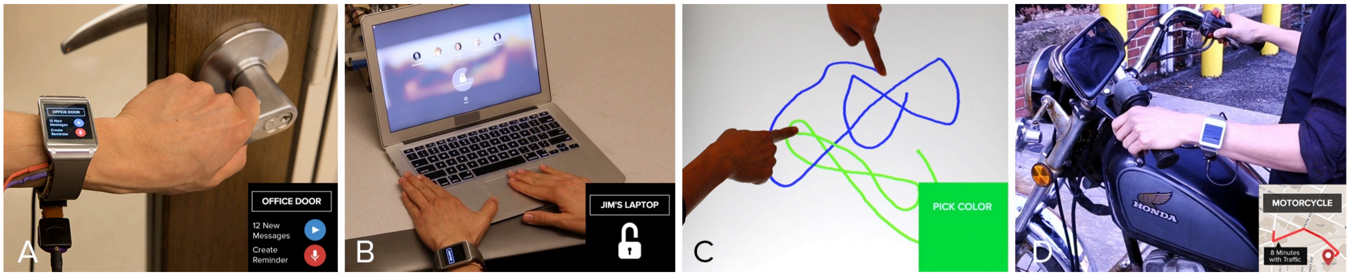

Figure 8. In the office, the use cases for EM-Sense are diverse. EM-Sense can be used for context-based communication (A), and for password-less authentication (B). In collaborative tasks, users with EM-Sense-capable devices enable user differentiation on touchscreens. When EM-Sense detects that the user is riding a motorcycle on the way home, a map is displayed (D).

Object-Specific Applications – When a user handles objects with known EM signatures, EM-Sense can launch object-specific applications. For example, our electric toothbrush example launched a timer application.

Object Sequence Applications – It is also possible to launch applications based on sequences and patterns of object events. Combined with other readily accessible features, such as time of day and rough geospatial location, activity and context recognition is possible. For example, a pattern of activation in the kitchen suggesting dinner preparation can launch music, recipe, and other applications.

Object State Recognition – We can further extend object-specific applications by utilizing changes in an object's EM signature in different operational modes. We demonstrated this in our Dremel application depicting a "speedometer".

Authentication **–** A smartwatch with EM-Sense could allow users to authenticate across devices and applications, potentially without passwords. For example, to log in into a laptop, a user can simply touch the trackpad. Because the smartwatch knows that a trackpad is being touched, and the trackpad knows that it is being touched, a handshake mediated by the cloud could proceed (using e.g., temporal co-occurrence of events). For added security, a confirmation button can be displayed on the owner's smartwatch.

User Differentiation – Similar to the authentication interaction above, knowledge of touchscreen events provided by EM-Sense could be used to differentiate users in groupware applications, which have many uses (see e.g., [11,24]). Specifically, a wearer's smartwatch knows the time of touch contact, which can be paired (e.g., in the cloud) to a touch event registered on a shared screen. Because the smartwatch knows its owner, touches can be attributed and parameterized to a specific user – in our example day, we used the watch display for a personalized color selector.

Object-Tagged Messaging – Knowledge of which objects are being handled also enables tagging of items with media, such as text and voice messages. In our example day, Julia leaves a message for herself by tagging her office's door handle. By using names, it would also be possible to leave messages for particular people.

EVALUATION

We ran multiple studies evaluating several facets of EM-Sense. These studies serve several purposes: 1) to evaluate the accuracy and robustness of our sensing approach across different users, 2) to observe the longitudinal consistency of object EM signatures over time, and 3) to form a baseline understanding of the uniqueness of EM signatures across a wide range of objects. We also conducted several smaller supporting studies, which explore other important aspects of EM-Sense, including signal similarity across similar devices and object state recognition. Overall, EM-Sense enables an expansive range of applications (see Applications section), and in our evaluation, we endeavored to select objects reflecting diverse contexts and environments.

Accuracy and Longitudinal Consistency

This study aims to evaluate the sensing accuracy of EM-Sense across different users, and determine whether object EM signatures are consistent over time. Because our sensing technique relies on the conductivity of the human body,

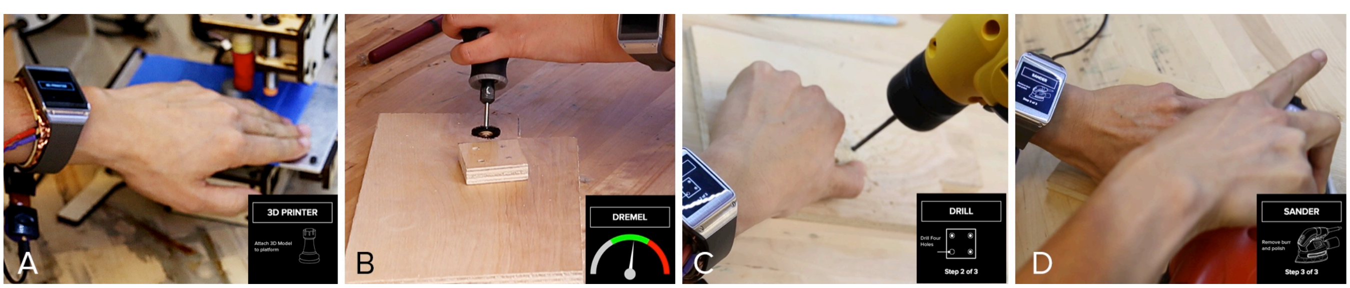

Figure 9. In the workshop, EM-Sense can assist in fabrication activities. Here, a tutorial is displayed on the watch. When the user grasps or operates a tool, the tutorial is automatically advanced (A, C, D). For some tools, EM-Sense can detect the operational state. For example, the speed of the Dremel is displayed on the watch (B).

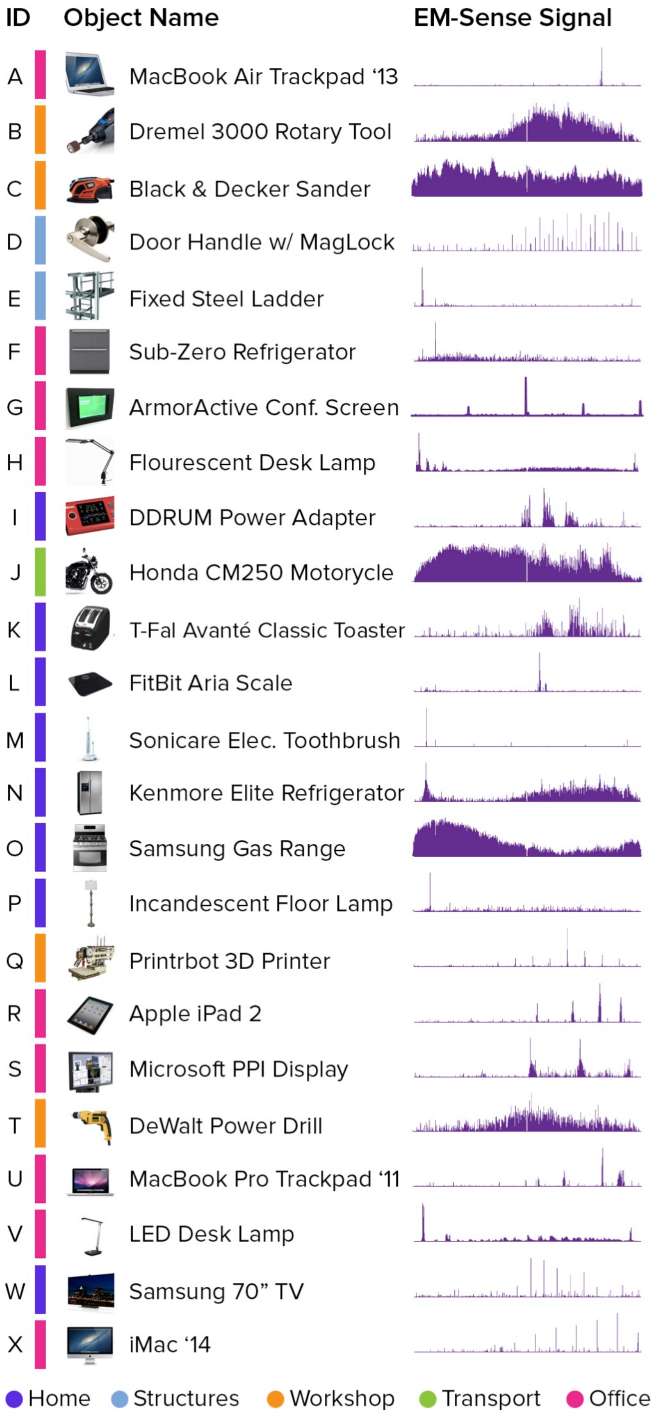

Figure 10. EM spectrums (0Hz to 1MHz) for the objects in our study as captured by our sensor.

our approach can be sensitive to differences in anatomy. Thus, we recruited 12 adult participants (5 female, age range 22 – 40, 1 left-handed), encompassing different statures and body types (mean height = 67 in., mean weight = 148 lbs., BMI range 20 – 28).

To further test sensing robustness, each study session was split across two different buildings, and we used data collected from *a single* user *six weeks* prior to the user study (i.e., no per user training or calibration). Nine objects were evaluated, dispersed across our two locations: MacBook Air trackpad, mouse sander, door handle with an electromagnetic lock, fixed steel ladder, refrigerator, ArmorActive conference room touchscreen affixed to a wall, fluorescent desk lamp, power adapter, and a Dremel rotary tool.

For logistical convenience, all experiments started in our first location. Participants were asked to wear our prototype on their preferred arm (anecdotally, we noticed participants preferred to wear the device on their non-dominant arm, as is the norm for watches). For each trial, an experimenter announced an object name (e.g., "Dremel"), and participants were asked to touch, grasp, or operate the object. The experimenter recorded the *real-time* prediction made by EM-Sense. Objects were requested in a random order, appearing five times each in total. Participants were free to interact with objects with either or both hands. Each session took approximately 45 minutes to complete, and participants were paid $15 for their time.

Across nine objects, 12 users, two locations, and using data trained on one user collected six weeks prior, EM-Sense achieved an average overall accuracy of 96.1% (see Figure 11, STDEV=4.9%, chance 11%). This result is promising given the strict constraints imposed on our training data. Some objects achieved an accuracy of 100% (lowest is 85%, Sub-Zero Refrigerator). While not the focus of the study, we can report anecdotally that signal magnitudes appear stronger when the prototype is worn on the same arm as the hand touching the object (consistent with prior findings [9]). Overall, our results indicate that sensing is accurate and robust across different users and that object EM signatures are consistent over time.

Signal Uniqueness Across Objects

To more fully explore the uniqueness of EM Signatures across many objects, we ran a second study that collected data from 23 objects across four locations. This set was composed of our initial nine objects, plus fourteen new objects that spanned a wider range of contexts and environments, including the home, office, workshop, large structural features, and transportation (see Figure 10). We also included *similar* objects (same category, but different models) to see if this caused classification confusion. Specifically, we include two refrigerators, two Apple laptops, three lamps, and four devices where the LCD display is touched (ArmorActive conference room touchscreen, iPad, Samsung TV, and Microsoft PPI display).

Due to the large number of objects and locations, we performed an offline analysis. Three rounds of data were collected for each object, with at least 15 minutes between rounds. We utilized the first two rounds of data for training and the third and final round for testing. This procedure prevents the model from over-fitting on time-adjacent instances (e.g., inherent similarities in touches when performed back-to-back). For each object, a round consisted of

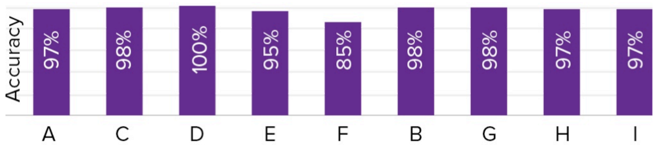

Figure 11. Across 12 users and 9 objects, real-time accuracy was 96.1%. Data were trained on 1 user, 6 weeks prior. Objects were trained in one location and tested in another.

collecting 250 instances with various hand poses to aid classifier generality. In total, we collected 17,250 data points (23 objects x 3 rounds x 250 instances). We also added a null, "no object touched" class, increasing the set to 24. We then trained a single SVM model using the aforementioned features and parameters.

Across these 24 classes, our system achieved an overall accuracy of 97.9% (STDEV=4.3%, chance 4%), which suggests object EM signatures are reasonably unique and discriminable (see Figure 12). Note that the majority of objects (18 of the 24) reach an accuracy of 100%, while the lowest object accuracy is at 85.6% (Samsung Gas Range). These results are promising given the large number of classes and the low volume of training instances per object.

EM Signatures of Similar Objects

As noted previously, we purposely included multiple objects of the same category, but different models, to see if *similar* devices would produce similar EM signals, and thus result in classification confusion. For our two refrigerators (Figure 12, F and N), two Apple laptops (A and U), and four LCD devices (G, R, S and W), there was 0% confusion. We found 1% confusion between the incandescent lamp (P) and the fluorescent lamp (H). These results strongly suggest that objects within a common category still have their own unique EM signatures.

EM Signatures of Identical Objects

We ran a supplemental study to determine if EM signals are consistent across *identical* objects. For this, we used the ArmorActive touchscreens installed at four conference rooms in an office setting. We used the 24-object classifier from our second study, which was trained on *one* of these devices six weeks prior. We then evaluated real-time classification accuracy on all *four* units. We ran 10 randomized touch trials per device, for a total of 40 trials. Our EM-Sense classifier correctly identified the object as the ArmorActive touchscreen 100% of the time (chance is 4%).

We used the same procedure for five iMac 2014 computers. We gathered training data on one machine, and ran 10 randomized classification trials on all five machines, for a total of 50 trials. Similarly, our classifier correctly classified these as iMacs with 98% accuracy (chance 4%).

Overall, these results suggest that the EM signatures of identical devices are very similar, allowing for object recognition even when that *particular* instance of the object has never been touched before. This outcome is beneficial, as it means EM-Sense capable devices could be *preloaded*

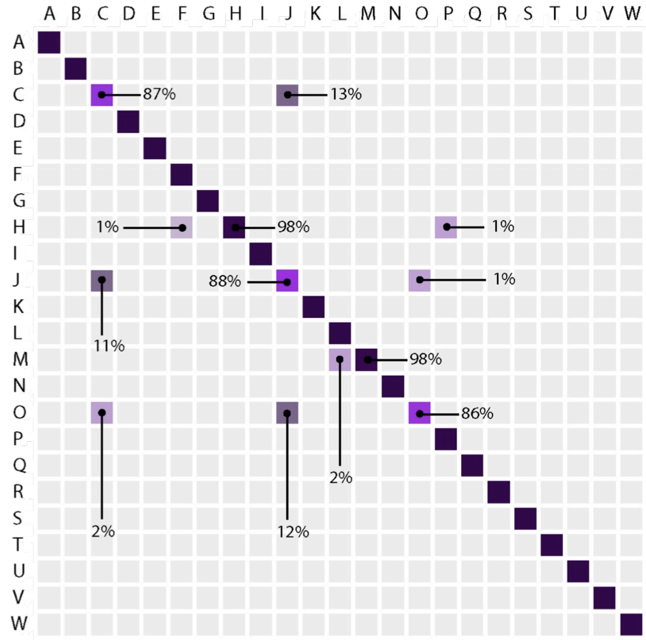

Figure 12. Object confusion matrix. Accuracy is 100%, unless indicated otherwise. Across 24 classes (including null class), average prediction accuracy was 97.9%. Figure 10 provides a key to the letters used on the axes.

with an EM signature database of known objects (or e.g., use a database in the cloud, which could grow overtime as users and companies add newly encountered objects).

Infering Object State

For some objects, it is also possible to infer the operational state based on EM signature. For example, the magnitude of a power drill's EM signal is generally proportional to the rotational speed of its motor. In response, we ran another supplemental study to determine whether EM-Sense can exploit this phenomenon.

We trained an EM-Sense classifier to detect four operational speeds of a Dremel 3000 rotary tool: OFF, LOW, MID, and HIGH. A total of 200 instances were collected per state. Of note, we tweaked our EM-Sense noise-detection parameters (e.g., from 6s to 60s) to delay the system from integrating EM signals into its background noise profile. Across 40 trials (10 trials per state), our system achieved a real-time classification accuracy of 92.5% across the four speeds, suggesting that variations in EM signal can also reveal object state.

DISCUSSION AND LIMITATIONS

Because we perform adaptive background subtraction, our technique is location independent. In fact, most portable objects in our study (Dremel, laptop, iPad, etc.) were trained in one location (again, 6 weeks prior), and tested in another location without issue. Throughout piloting, we never observed a location effect. However, large passive objects, like our ladder, which are an amalgamation of EM signals from their respective local environments, would change if relocated.

Our approach is passive, capturing noise, but not generating any signals itself. As we have discussed, this limits us to certain classes of objects. Indeed, most objects do not generate EM signals (e.g., chairs, cups, books). Thus, our sensing scope is generally limited to electrical and electromechanical objects (and some large static objects, as discussed previously). Even still, not all of these objects are detectable, as the strength of EM signals is subject to the physical design of objects (e.g., variations in electrical shielding and grounding). Moreover, some frequencies of noise may not be (faithfully) conducted through the human body and thus not reach our smartwatch-bound sensor.

Additionally, high fidelity analog sensing requires a stable and strong electrical ground as a reference. In our prototype, we tried to faithfully replicate the grounding conditions of a smartwatch, which contains a small battery. Additionally, our SDR receiver only provided 8-bit ADC resolution. With a superior ground reference and increased resolution (i.e., a commercial-level implementation), EM-Sense may support even larger sets of objects.

Finally, we noticed in some cases that very strong environment noise (e.g., transmitters that broadcast in overlapping bands of interest, or a microwave oven in operation) raised the noise floor and overpowered local EM signals. Even when noise subtraction is applied, high intensity noise can blanket subtle but discriminative EM signals. Additionally, because our sensor is worn on the arm, it is subject to frequent movements, which can cause unintended electrical effects (e.g., Doppler shifts). Movement information from e.g., accelerometers and gyroscopes could compensate for e.g., sudden arm movements, or simply pause classification. Anecdotally, however, these effects appear to be minor.

CONCLUSION

We have proposed a novel sensing approach for on-touch object detection that exploits the unintentional electromagnetic noise generated by commonplace objects. By modifying a small, low-cost, embedded software-defined radio receiver, we can detect and classify EM signals in real time, enabling quick and robust detection of *when* an object is touched and *what* that object is. Our experiments show that sensing can be accurate and robust. We also highlighted the wide variety of objects that EM-Sense can detect in several example contexts and environments, which point towards more powerful assistive, context sensing, and communication applications.

ACKNOWLEDGEMENTS

This research was generously supported by the David and Lucile Packard Foundation, Google, Yahoo! InMind, and Qualcomm. We would like to thank Eric Brockmeyer and Liz Carter for their invaluable help and feedback.

REFERENCES

1. Abbott, R.E. and S.C. Hadden, Product Specification for a Nonintrusive Appliance Load Monitoring System. EPRI Report #NI-101, 1990.
2. Bao, L. and Intille, S.S. Activity Recognition from User-Annotated Acceleration Data. In *Proc. of Pervasive Computing '04*. 1-17.
3. Buchler, M.C. Algorithms for Sound Classification in Hearing Instruments. PhD thesis, ETH Zurich, 2002.
4. Buettner, M., Prasad, R., Philipose, M. and Wetherall, D. Recognizing daily activities with RFID-based sensors. In *Proc. UbiComp '09.* 51-60.
5. Chen, K.Y., Cohn, G., Gupta, S. and Patel. S.N. uTouch: sensing touch gestures on unmodified LCDs. In *Proc. of CHI '13*. 2581-2584.
6. Clarkson, B., Sawhney, N. and Pentland, A. Auditory context awareness in wearable computing. In *Workshop on Perceptual User Interfaces*, *November 1998*.
7. Cohen, G., Gupta, S., Froehlich, J., Larson, E. and Patel, S. GasSense: Appliance-Level, Single-Point Sensing of Gas Activity in the Home. In *Proc. of Pervasive '10*. 265-282.
8. Cohn, C., Morris, D., Patel, S. and Tan, D. Humantenna: using the body as an antenna for real-time whole-body interaction. In *Proc. of CHI'12.* 1901-1910.
9. Cohn, G., Morris, D., Patel, S.N. and Tan, D.S. Your noise is my command: sensing gestures using the body as an antenna. In *Proc. of CHI '11.* 791-800.
10. Cohn, G., Stuntebeck, E., Pandey, J., Otis, B., Abowd, G.D. and Patel, S.N. SNUPI: sensor nodes utilizing powerline infrastructure. In *Proc. UbiComp '10.* 159-168.
11. Dietz, P. and Leigh, D. DiamondTouch: a multi-user touch technology. In *Proc. of UIST '01.* 219-226.
12. Do, T.M.T., Kalimeri, K., Lepri, B., Pianesi, F. and Gatica-Perez, D. Inferring social activities with mobile sensor networks. In *Proc of ICMI '13.* 405-412.
13. 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. of UIST '06.* 91-100.
14. Froehlich, J., Larson, E., Campbell, T., Haggerty, C., Fogarty, J. and Patel, S.N.: HydroSense: infrastructure-mediated singlepoint sensing of whole-home water activity. In *Proc. of UbiComp '09.* 235-244.
15. Grosse-Puppendahl, T., Herber, S., Wimmer, R., Englert, F., Beck, S., von Wilmsdorff, J., Wichert, R., and Kuijper, A. Capacitive near-field communication for ubiquitous interaction and perception. In *Proc. of UbiComp '14.* 231-242.
16. Gupta, S., Chen, K.Y., Reynolds, M.S., and Patel, S.N. LightWave: using compact fluorescent lights as sensors. In *Proc. UbiComp '11.* 65-74.
17. Gupta, S., Reynolds, M.S. and Patel, S.N. ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. In *Proc. UbiComp '10.* 139-148.

18. Haché, G., Lemaire, E.D and Baddour, N. Wearable mobility monitoring using a multimedia smartphone platform. In *IEEE Trans. on Instrumentation and Measurement '11*. 3153-3161.

19. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P. and Witten, I. The WEKA Data Mining Software: An Update. *SIGKDD Explor.,* 11(1), 2009.

20. Harrison, C., Xiao, R. and Hudson, S.E. Acoustic barcodes: passive, durable and inexpensive notched identification tags. In *Proc. UIST '12.* 563-568.

21. Hart, G. Advances in Nonintrusive Appliance Load Monitoring. In *Proc. EPRI Information and Automation Conference* '91.

22. Hart, G., Nonintrusive appliance load monitoring. In *Proc. of the IEEE*, 1992. 1870-1891.

23. Hodges, S., Thorne, A., Mallinson, H. and Floerkemeier, C. Assessing and optimizing the range of UHF RFID to enable real-world pervasive computing applications. *Pervasive Computing '07.* 280-297.

24. Holz, C., and Baudisch, P. Fiberio: A Touchscreen That Senses Fingerprints. In *Proc. of UIST'13.* 41-50.

25. Kim, Y., Schmid, T., Charbiwala, Z. and Srivastava, M.B. ViridiScope: design and implementation of a fine-grained power monitoring system for homes. In *Proc. of UbiComp '09.* 245-254.

26. Lange, B.M., Jones, M.A., Meyers, J.L. Insight lab: an immersive team environment linking paper, displays, and data. In *Proc. of CHI '98.* 550-557.

27. Laput, G., Lasecki, W.S., Wiese, J., Xiao, R., Wiese, J., Bigham, J.P. and Harrison, C. Zensors: Adaptive, Rapidly Deployable, Human-Intelligent Sensor Feeds. In *Proc. of CHI '15.* 1935-1944.

28. Lasecki, W.S., Song, Y., Kautz, H., and Bigham, J.P. Real-time crowd labeling for deployable activity recognition. In *Proc. of CSCW '13.* 1203-1212.

29. Li, H., Ye, C. and Sample, A.P. IDSense: A Human Object Interaction Detection System Based on Passive UHF RFID. In *Proc. of CHI '15.* 2555-2564.

30. Maekawa, T., Kishino, Y., Sakurai, Y. and Suyama, T. Recognizing the use of portable electrical devices with hand-worn magnetic sensors. In *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. of Pervasive'12.* 234-252.

32. Movassaghi, S., Abolhasan, M., Lipman, J., Smith, D., Jamalipour, A. 2014. Wireless Body Area Networks: A Survey. In *Comm. Surveys & Tutorials*. 1658 - 1686.

33. Paglin, M.D., Hobson, J.R. and Rosenbloom, J. (1999). The Communications Act: A Legislative History of the Major Amendments, 1934-1996. Pike & Fische.

34. Patel, S.N., Reynolds, M.S. and Abowd, G.D. Detecting Human Movement by Differential Air Pressure Sensing in HVAC System Ductwork: An Exploration in Infrastructure Mediated Sensing. In *Pervasive '08.* 1-18.

35. Patel, S.N., Robertson, T., Kientz, J.A., Reynolds, M.S. and Abowd, G.D. At the flick of a switch: detecting and classifying unique electrical events on the residential power line. In *Proc. of UbiComp '07.* 271-288.

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. Philipose, M. Large-scale human activity recognition using ultra-dense sensing. *The Bridge*, National Academy of Engineering. 35, 4 (2005).

38. Randell, C. and Muller, H. Context awareness by analysing accelerometer data. In *Proc. ISWC* '00. 175-176

39. Ranjan, J., Yao, Y., Griffiths, E. and Whitehouse, K. Using mid-range RFID for location based activity recognition. In *Proc. of UbiComp '12.* 647-648.

40. Rekimoto J. and Ayatsuka, Y. CyberCode: designing augmented reality environments with visual tags. In *Proc. DARE '00.* 1-10.

41. Ren, X. Egocentric recognition of handled objects: Benchmark and analysis. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. 2009. 2160-7508.

42. Smith, J.R., Fishkin, K.P., Jiang, B., Mamishev, A., Philipose, M., Rea, A.D., Roy, S. and Sundara-Rajan. K. RFID-based techniques for human-activity detection. In *Commun. ACM*. Vol. 48-9, 2005. 39-44.

43. The Federal Communications Commission. Code of Federal Regulations, Title 47, Part 15.

44. Wang, Y., Liu, J., Chen, Y., Gruteser, M., Yang, J. and Liu, H. E-eyes: device-free location-oriented activity identification using fine-grained WiFi signatures. In *Proc. of MobiCom '14. 617-628.*

45. 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.

46. Zhao, W., Chellappa, R., Phillips, P. and Rosenfeld, R. Face recognition: A literature survey. *ACM Comput. Surv. '03* Vol. 35-4. 399-458.

47. Zhao, Y., LaMarca, A. and Smith, J.R. A battery-free object localization and motion sensing platform. In *Proc of UbiComp '14.* 255-259.