Your Noise is My Command:
Sensing Gestures Using the Body as an Antenna

Gabe Cohn[1,2], Dan Morris[1], Shwetak N. Patel[1,2,3], Desney S. Tan[1]

[1]Microsoft Research
Redmond, WA (USA)
{dan, desney}@microsoft.com

[2]Electrical Eng., [3]Computer Science & Eng.
UbiComp Lab, DUB Group, Univ. of Washington
Seattle, WA (USA)
{gabecohn, shwetak}@uw.edu

ABSTRACT

Touch sensing and computer vision have made human-computer interaction possible in environments where keyboards, mice, or other handheld implements are not available or desirable. However, the high cost of instrumenting environments limits the ubiquity of these technologies, particularly in home scenarios where cost constraints dominate installation decisions. Fortunately, home environments frequently offer a signal that is unique to locations and objects within the home: electromagnetic noise. In this work, we use the body as a receiving antenna and leverage this noise for gestural interaction. We demonstrate that it is possible to robustly recognize touched locations on an uninstrumented home wall using no specialized sensors. We conduct a series of experiments to explore the capabilities that this new sensing modality may offer. Specifically, we show robust classification of gestures such as the position of discrete touches around light switches, the particular light switch being touched, which appliances are touched, differentiation between hands, as well as continuous proximity of hand to the switch, among others. We close by discussing opportunities, limitations, and future work.

Author Keywords
Input, touch interaction, surface interaction, electrical noise

ACM Classification Keywords
H.5.2 Information interfaces and presentation: User Interfaces - Input devices and strategies.

General Terms
Design, Experimentation, Measurement

INTRODUCTION

As computers become more mobile and more ubiquitous, people increasingly expect always-available computing, either with devices that they carry on their bodies, or using devices embedded in the environment. We see an increasing need for interaction modalities that go beyond the key-

Figure 1: The human body behaves as an antenna in the presence of noise radiated by power lines and appliances. Our approach analyzes this noise, turning the whole home into an interaction surface.

board and mouse, and furthermore, that do not require mediated interaction with specialized devices such as styluses.

Researchers have addressed this need through a variety of input channels. Speech recognition enables hands-free interaction for a variety of desktop and mobile applications. Similarly, computer vision enables machines to recognize faces, track movement, recognize gestures, and reconstruct 3D scenes. Various techniques, most notably capacitive sensing, have been used to instrument surfaces such a tables, walls, and mobile devices in order to provide touch sensing. In addition, specialized depth cameras that allow users to interact with their computers using whole-body gestures are becoming commercially available to consumers (e.g., Microsoft Kinect).

Speech input comes at a relatively low cost of instrumentation, but is limited in input bandwidth and may not be appropriate in many scenarios. Vision- and touch-based technologies offer an array of subtle, natural interaction techniques, but are limited in the potential scale of deployment due to their associated installation burden and cost. Consequently, we will likely not see homes or workplaces that allow truly ubiquitous input in the foreseeable future using these modalities.

Other researchers, realizing these limitations, have explored sensors that exploit characteristics of the human body itself to turn it into an inherently portable interaction device. Harrison et al. utilize bio-acoustic sensors to determine the location of taps on the body, and thereby turn it into a touchscreen [12]. Saponas et al. use electrical recordings of forearm muscles to sense muscle activity and infer finger gestures [16]. However, these on-body input systems are to date limited to a small number of discrete inputs, and do not offer the large-scale interaction that is provided by touch-sensitive surfaces.

In this paper, we present a novel interaction modality that utilizes the human body as an antenna to receive electromagnetic (EM) noise that already exists in our environments (Figure 1). While this noise is bothersome to nearly every other EM sensing application, we treat it as the core of our signal. By observing the properties of the noise picked up by the body, we can infer gestures on and around existing surfaces and objects, specifically the walls and appliances in the home. We targeted our initial proof-of-concept work to the home both because there is growing demand for computing in the home, but also because power lines in the home have been shown to be a relatively good transmitting antenna that creates a particularly noisy (i.e., signal-rich) environment for us.

Specifically, the contributions of this paper are:

1) A novel technique that uses the human body to sense EM noise already present in the environment to infer gestures such as touch and hover around various uninstrumented objects.

2) A core experiment conducted by 10 people in 10 homes validating the operation and robustness of the technique. Results from this experiment show that we can accurately classify the location in the home where the interaction occurred and the contact positions around light switches.

3) A set of smaller experiments to explore additional capabilities and limitations of our approach. Results from these experiments suggest that we can infer proximity to walls, multi-hand gestures, touched appliances, and continuous position along a touched wall.

BACKGROUND AND RELATED WORK

The Human Body as an Antenna
A basic receiving antenna can be thought of as an apparatus that converts electromagnetic waves into electrical current. An antenna consists of a set of conductors that can be arranged in a variety different ways, where the size, geometry, and material dictate its effectiveness at receiving a particular frequency. One of the simplest antennas is just a loop of wire (commonly used for AM and FM radio), whose length determines its effective frequency response. In addition, any wire or conductor carrying current or a wire exposed to an electromagnetic field may exhibit unintentional antenna characteristics. For instance, it is not un-

common to hear AM or CB radio through a set of speakers that is not plugged into a radio. Home electrical wiring also makes an effective antenna [4, 13, 15, 20], a phenomenon which we leverage in this work.

It turns out that the human body is also a very effective antenna over a broad frequency range. The human body is an electrical conductor, and thus when exposed to electromagnetic fields, it behaves as an antenna with a frequency resonance determined by various factors including height, posture, etc. Research suggests that since the human body is a lossy conductor (dielectric) with a complex geometry, it does not have a single effective resonance frequency, but rather a broadly distributed response [7], capturing frequencies from 40 Hz all the way to 400 MHz [8].

Sometimes referred to as the "body antenna effect," this phenomenon poses significant problems for systems employing body area networks (i.e., using the body as a conductor to send data from one part of the body to another) and for systems analyzing electrical phenomena within the body (e.g., muscle activity). Consequently, researchers have gone through great lengths to mitigate the problems of ambient electromagnetic noise being coupled to the body [2]. In contrast, our focus in the present work is to *leverage* the ambient electromagnetic noise picked up by the human body as a signal for classifying human interaction with the environment.

Related research has focused on using the human body as a conductor for body area networks [6, 10, 22]. For example, in the HCI community, Fukumoto and Tonomura demonstrated the FingeRing [6], a finger-worn sensor that communicates to a wrist mounted receiver by using the body as an "electric wire." Although their approach did not use the body as an antenna, they noted that touching a surface greatly reduced communication reliability because of the body being grounded and acting as a human antenna. Other work has explored using the human body as a transmission/reception antenna for inter-body communication [1, 3, 10].

Also in the HCI community, near-field electric field sensing has been a popular approach for touch and motion gestures, where the human body has been used as a disturbance and a radiator of an electric field [18, 21]. The DiamondTouch [5] employs the human body as both a signal conductor and an antenna by passing electromagnetic waves from a conductive pad into the body; these waves are in turn picked up by antennas embedded in a tabletop interaction surface. In the present work, we are interested in picking up noise signals using the body *without* instrumenting the environment. We use the particular properties of the measured noise to infer the gestures performed by the human user. We are not aware of other work that has explicitly looked at using ambient electromagnetic noise picked up by the human body for user interaction.

Home Power Infrastructure
There are many sources of electromagnetic noise in the

environment, but the home power line infrastructure is a major source. A home typically consists of electrical wiring that supplies power to outlets, appliances, and lighting via wall switches. The electrical wiring branches from a central circuit breaker, but the ground and neutral wires in the home are all tied together. Thus, signals occurring on the power line in one part of the house can be measured in other parts as well. In addition, the walls of the home are dielectrics and will radiate electromagnetic fields even if there are no power lines in the wall.

Also, as mentioned earlier, the electrical wiring in the home can act as both a reception and transmission antenna. Past work has used this phenomenon for indoor location tracking, where a tracking signal is radiated off the power line [13, 20]. Similarly, [4] used the power line as a large antenna for receiving data wirelessly from ultra-low-power sensor nodes. All of these approaches use a known signal that is either injected through or received by the power lines. Other work has looked at passively monitoring the power line using a high-frequency sampling of the voltage at a single electrical outlet to infer the activation of appliances and electrical devices in the home based on the appearance of electrical noise from those devices [9, 14].

Although similar in spirit, our work monitors only the electromagnetic noise radiated off of the power lines and received by the human body to determine where in the home the person is and what type of gesture they are performing.

Electrical Noise

The AC signal itself is one of the largest sources of electromagnetic noise in the home: this signal typically oscillates at 60 Hz[1]. However, appliances and electronic devices attached to the power line also contribute some noise. There are roughly three general classes of electrical noise sources that may be found in a home: resistive loads, inductive loads such as motors, and loads with solid state switching (also known as switched-mode power supplies).

Purely resistive loads, such as incandescent lamps or electric stoves, do not create detectable amounts of electrical noise while in operation, although just like a resistor, they can be expected to produce trace amounts of thermal noise at an undetectable level. A motor, such as in a fan or a blender, is modeled as both a resistive and inductive load. The continuous breaking and connecting by the motor brushes creates a voltage noise synchronous to the AC power at 60 Hz (and at 120 Hz). Solid state switching devices, such as those found in computer power supplies, compact fluorescent light (CFL) bulbs, modern TVs, TRIAC dimmer switches and microwave ovens, emit noise that varies among devices and whose frequency is determined by an internal oscillator [14].

The drive towards smaller and more efficient consumer

electronics has made use of switched-mode power supplies (SMPS) increasingly prevalent. In a modern SMPS this modulation happens at a very high rate (10 kHz – 1 MHz). A side effect of an SMPS's operation is that the modulation of the inductor's magnetic field produces large amounts of unintentional electromagnetic interference (EMI) centered at or around the modulation frequency. Due to the physical contact between the power line and the power supply, this EMI gets coupled onto the power line, which then propagates the noise throughout the entire electrical infrastructure of a home. This is known as conducted EMI, which in turn is radiated by the power line as radiated EMI. The appliance or device itself can also exhibit radiated EMI. Because such EMI may cause problems in the operation of certain electronic devices, the US Federal Communications Commission (FCC) sets rules for any device that connects to the power line and limits the amount of EMI it can conduct and radiate. However, despite these limits, significant and detectable EMI is still coupled back over the power line.

There are also several significant sources of electrical noise on the power line which originate outside the home. Radio broadcasts, including commercial AM and FM radio, are picked up by the power line, which acts as a receiving antenna over a wide range of frequencies. In addition, noise from elsewhere in the neighborhood is often coupled through the earth ground connection as well. Pilot tests showed that even when we turned off the main power coming into a home, there was still significant baseline noise present in the home, and radiated from the power line.

Combining Power Line Noise and "Body as an Antenna"

Past work in power line noise analysis and using the human body as an antenna has largely explored disparate applications. Recognizing the potential of using electrical noise as a signal and the human body as a receiving antenna, our work seeks to enable new user interaction capabilities in the home that require no additional instrumentation to the environment, and only a simple analog-to-digital converter on the body itself. Based on the prior work, we hypothesized that the complex shape of the power line infrastructure provides enough spatial differentiability in the signal space to allow us to uniquely identify locations and contact points relative to electrical devices and wiring. In other words, by looking at a various characteristics of the frequencies (presence, amplitude, shape, etc.) observed on body, it is possible to detect gestures.

CORE EXPERIMENT: VALIDATING THE TECHNIQUE

Participants and Homes

We conducted the experiment in 10 homes selected to represent a variety of constructions, in the Pacific Northwest region of the United States. These homes were single-family and townhouses built between 1948 and 2006 (μ=1981). They ranged in size between 120 and 290 square meters (μ=215), and had between 1 and 3 floors, some of them basements. The owner of each of these homes participated in our experiment. These 10 participants (5 female) were between 28 and 61 years old (μ=38), weighed be-

---

[1] This work was conducted in North America and thus refers to 60 Hz AC power; other parts of the world use 50 Hz.

tween 52 and 82 kg (μ=64), and were between 150 and 188 cm tall (μ=169 cm).

Apparatus

Electromagnetic signals radiating from the power lines and walls and picked up by the human body antenna can be measured as voltages. Since the body is relatively conductive, we can measure these voltages by placing a conductive pad, connected by a wire to an analog-to-digital converter, nearly anywhere on the body. In this experiment, we chose to measure voltages on the back of the neck because it is a stable point on the body that does not move significantly while a person is gesturing with their hands. The neck was also a convenient place because it is near our data collection equipment, which was housed in a backpack worn by the participant (Figure 2). We will validate in additional experiments that wearing the contact pad on a different part of the body still permits robust classification of hand gestures.

We made electrical contact to the skin using a standard grounding strap, typically worn around the wrist when working with sensitive electronics. We ran a small wire from the contact pad to a National Instruments USB-6216 data acquisition unit, which sampled the voltages at 400 kS/s. We biased the voltage on the contact point to a local ground signal on the data acquisition unit through a 10 MΩ resistor in order to remove most of the DC offset of the single-ended voltage. The data acquisition unit's local ground is internally isolated from the laptop's ground so that the measurements are referenced only to the small local ground on the data acquisition unit. The signal was digitized at 16-bit resolution and streamed to disk on an attached laptop for subsequent processing.

Experimental Procedure

We selected 5 light switches and 1 spot above an electrical outlet on a blank wall for testing in each of the 10 homes. In order to test whether or not we could differentiate between locations in close proximity, we ensured that two of the chosen light switches were located in the same room. The

Figure 2: Experimental setup. A laptop and a USB data acquisition device are worn in a backpack. A wire connects the data acquisition device to a conductive pad in contact with the back of the participant's neck.

other locations were distributed around the home with at least one location on each floor.

To minimize the number of variables that changed during the experimental session, we turned off appliances to which we had reasonable access and that periodically change their state, including most computers, as well as heating and air conditioning units. We left all light switches used in the experiment on, and we did not change the state of any lights or appliances once the experiment started.

Participants stood at arm's length away from the wall and performed 6 specific gestures around each interaction point (i.e., light switch or wall). The first was a "rest" gesture in which participants placed both hands at their sides. The other five involved contacting the wall with the right palm, placed flat against the wall for 6 seconds at different positions around the switch. These positions included directly on the light switch plate and at points approximately 20 cm above, below, right of, and left of the light switch. In the case of the blank wall, the same positions were used, but in reference to an arbitrary point at about the height of a light switch above the outlet marked on the wall. Each participant performed these six gestures at all six locations (5 switches, 1 wall) around their home. Figure 2 shows a participant in contact with the wall in the "above light switch" position.

To help participants and to ensure consistency, we marked each of the contact points with tape. When obstacles prevented the touch from occurring at 20 cm in any direction from the center position, we placed the tape as close as possible to the target position and noted this. We also taped over the ground screws on each light switch to ensure that the participant would not be shorted to ground while touching the switch. This was done to ensure that each contact with the switch was conducted under the same known conditions. Subsequent experiments confirmed that the ground screw provides a unique signal unto itself that is easy to robustly discriminate from the other positions in our experiment.

Software running on the data collection laptop issued verbal commands in order to guide participants through the experiment. In addition, all experiments involved a second person as an observer to check for mistakes and inconsistencies. The observer stood at least 1 m away from the participant to ensure that his presence did not significantly alter the received signals. At each location, the software issued commands about which position around the switch the participant should touch, followed by a 2-second beep, allowing the participant time to move to that position. Data was then collected for 6 seconds before the next command was issued. We randomized the order of the gestures at each light switch to eliminate any potential temporal bias and to ensure that the participant remained cognitively engaged. Participants moved from location to location in a predetermined order, and repeated the entire procedure 4 times (144 total gestures performed).

Figure 3: Signal captured from a participant during two touches on a wall (red) and a light switch (blue). Green highlights (left and right) indicate periods of contact, the orange highlight (center) indicates the "rest" period of non-contact. Differences between these signals, and between the contact and no-contact "rest" states, form the basis of our approach.

Analysis

Our primary goal is to inform the development of interactive systems that use the sampled signal to recognize gestures. Hence we treated our analysis as a machine learning classification problem. Specifically, we used the Sequential Minimal Optimization (SMO) implementation of the support vector machine (SVM) found in the Weka machine learning toolkit [11]. An SVM uses labeled data to construct a set of hyperplanes that separate labels in a high-dimensional feature space, which can then be used for classification. Fully exploring possible machine learning techniques is outside the scope of this paper; our results can be treated as a baseline that may be further optimized.

In order to prepare data for the SVM, we first segmented the 6-second gestures, removing a half-second from the front and end of this period to account for potential reaction time and anticipatory effects. We then divided the raw voltage signal into consecutive 82-millisecond windows. This window size allows for very low latency in gesture detection; however the results of the classification can be improved by smoothing over longer windows. In our analyses, each of these windows was treated as being independent data points. We then generated the following 1002 features for each window, which we used to train our SVM.

Time-Domain Features (2)

The most basic feature was the mean of the voltage (DC value). We also calculated the root-mean-square, or RMS value. The RMS value represents the AC amplitude of the voltage, which changes significantly between different gestures, as shown in Figure 3.

Low-Frequency Features (582)

Since the power lines are used to carry low-frequency AC power (at 60 Hz), it is not surprising that most of the energy radiated off of the power line and received by the human body antenna is in the low-frequency range. Figure 4 shows that the power spectrum is dominated by 60 Hz and its harmonics. As a result, these frequencies are important for machine learning. We used all of the raw frequency bins between DC and 2 kHz (12 Hz resolution) produced from a


Figure 4: A frequency-domain representation of a signal captured during contact between a participant and a light switch. The 60 Hz peak and its harmonics are clearly visible throughout the low-frequency spectrum.

32768-point FFT as features. Since the SVM is a linear classifier, we included both the linear and log (dB) versions of these features (332 features total). In addition, the harmonics of 60 Hz seem to become negligible at frequencies higher than about 25 kHz, and hence we create a second set of low frequency features containing FFT bins between DC and 25 kHz at 200 Hz resolution, again using both the linear and log (dB) versions (250 features total).

High-Frequency Peak Features (18)

Through an initial exploration of the signals received on the human body antenna, it became obvious that several high-frequency peaks were indicative of certain types of gestures. As a result, we also include the maximum value of several of these high-frequency peaks as features. The peaks that we included are 20k, 30k, 50k, 60k, 80k, 90k, 110k, 140k, and 150 kHz, again using both the linear and log (dB) versions (18 features total).

Full Frequency Range Features (400)

In order to encode the general shape of the frequency spectrum, we use features containing frequency samples between DC and 200 kHz with a 1 kHz resolution, again using both the linear and log (dB) versions (400 features total).

Results

To calculate how accurately we could classify various conditions, we conducted multiple four-fold cross-validations. Each fold was made up of data points from a single "round" through all locations in the home. This ensured that training and testing data points were separated within a fold, and that training data and testing data were separated by several minutes in time (to avoid over-fitting to transient variations in the environment). These numbers are representative of what we would expect to see in an interactive system. We report average accuracies and standard deviations.

Wall Touch

Looking only at the time-domain signals (Figure 3), it is easy to see the difference between the time when the participant is touching the wall (green) and the time when they are not (orange). Therefore, not surprisingly, the classification results for this kind of analysis were quite high.

Two-class classification of wall-touch vs. no-wall-touch

performed at 98.5%, σ=4.1 (chance=50%) when averaged across the participants in all 10 homes. Since the strength of the signal received on the body is related to the proximity to the radiating source, in this case the power lines, we expected that our wall touch classification would perform better on light switches than on the blank walls above outlets. However, our results show that the classification worked just as well on the blank walls, indicating that gestures do not need to be confined to the area around light switches. In fact, touches on most walls are detectable because of the wiring elsewhere in the wall.

Location in Home
The 6-location classification of interaction location in the home performed at 99.1%, σ=1.3 (chance=16.7%) when using data from all gestures around each light switch in each home. This is a very impressive result, made even more impressive by noting that by experimental design, two of the walls in each classification were located in the *same* room. This suggests the possibility to classify which wall a user is interacting with, rather than just which room.

Perhaps even more interestingly, the same level of accuracy can be obtained without even touching the wall. Using only the data from when the participant was standing at rest at arm's length from the wall (hands at sides), the 6-location classification performed at 99.5% σ=1.2 (chance=16.7%). This is a promising result, because it hints at the possibility of determining the location of people throughout the home, even when they are not interacting directly with the walls. This could enable location-aware systems that use in-air gestures in addition to on-wall gestures.

Touch Position on Wall
The 5-position classification of gesture position around the light switches performed at 87.4%, σ=10.9% (chance=20%). The touch position on the blank walls can be classified at 74.3%, σ=16.1% (chance=20%). This is an interesting result, as it suggests that it may be possible to classify arbitrary touch positions on blank walls, not just touches that are near light switches. We explore this further in the exploratory experiments described in the next section.

Location in Home and Touch Position on Wall
By combining the classification of *both* the location in home *and* the touch position on the wall, we have a 30-class problem, which performed at 79.8%, σ=7.0 (chance=3.3%). While this number may not seem high, recall that these are unoptimized classifications on individual time-windows and that these numbers should increase for entire touches, even simply using naïve voting schemes across multiple windows. With no additional instrumentation to the home, these results are quite promising in terms of the ability to both classify touch locations in the home as well as the absolute position on the wall.

Summary
Based on the classification results presented in this section, we can determine the location of the user in the home, with near 100% accuracy, and can determine whether the user is touching a wall or not with 98% accuracy. With 87% accuracy, we are able to determine the position around a light switch on a wall. We can even simultaneously identify the location in the home and the position on the wall of a given touch with 80% accuracy.

EXPLORING ADDITIONAL CAPABILITIES
Our core experiments, described above, confirmed our hypothesis that electromagnetic noise in the home is unique to specific locations, allowing discrimination of locations within the home and touched wall positions. In order to guide future work and understand the boundaries of this approach, we conducted a series of additional experiments to determine other capabilities that our approach might offer. We also wanted to confirm that decisions we made for consistency in experimental design (e.g., choosing the neck as the body contact location) are not restrictions that are fundamental to our approach.

We used the data acquisition system and methodology presented in the previous sections, but the following experiments were performed by two participants (instead of ten), each in one (different) home. This reduced participant pool allowed us to explore a variety of techniques, while still ensuring that results were not unique to a single person or home. Unless otherwise indicated, classification results are based on the same SVM described above, classifying 82-millisecond windows within our 5-second gestures, with 4-fold cross-validation across 4 "rounds" through all gestures in an experiment.

Each subsection introduces an experimental question that we aimed to answer in this additional exploration, along with results that provide preliminary answers to those questions.

Experiments and Results

Body Contact Location
In our core experiment, we placed the contact pad on the participant's neck, which allowed us to eliminate movement of the pad for the sake of experimental consistency. However, for real-world scenarios, connecting a computing device to a user's skin would much more likely sit near a location on which we already wear computing devices (e.g., the wrist). We thus repeated our core 5-position classification experiment around a single light switch in each of 2 homes, with the contact pad placed on the participant's forearm, instead of the neck. The 5-position classification of position around the switch performed at 98% and 97% (chance=20%) for our two participants.

This indicates that our approach performs well with the contact pad placed on the arm (where it might be connected to a watch), which is more practical for consumer scenarios.

Appliance Classification
Home appliances are known to emit a significant amount of electromagnetic noise. In addition, they have large sections of metal which are well grounded to the home's power line

Figure 5: Frequency-domain differences of measured signal between a refrigerator, microwave, and a wall touch. Peaks show characteristic changes in amplitude that allow our classifiers to discriminate among touched objects. Here we show just one set of peaks from the high-frequency spectrum; similar characteristic peaks occur throughout the spectrum.


Figure 6: Regressing captured data onto a participant's distance from the wall (hand movement only). Error bars represent standard deviations. Light gray line indicates correct prediction values.

infrastructure. We hypothesized that this would allow us to robustly classify contact between a user and a home appliance, suggesting a variety of interaction techniques based on touching uninstrumented appliances. In order to address this hypothesis, participants touched each of six appliances in the same kitchen: refrigerator, freezer, stove, microwave, dishwasher, and faucet. All appliances were plugged in, but not actively running, during the experiment.

Consistent with our hypothesis, the measured electromagnetic noise while touching these appliances was quite dramatic, even compared to the noise observed during wall touches, and showed strong differences in our feature space among appliances (Figure 5). Consequently, classification among these six appliances was 100% for both participants (chance=16.7%).

This indicates that direct contact with appliances provides a robust signal for classification, suggesting the potential to turn uninstrumented appliances into real-world "buttons".

Number of Hands
In the experiments presented in the previous section, participants used only their right hand. We hypothesized that asymmetries in body conductivity, subtle asymmetries in contact pad placement, and differences in contact area would allow us to robustly discriminate left-, right-, and two-handed contact with a wall or light switch. As a preliminary investigation of this hypothesis, participants made left-, right-, or two-handed contact with a single light switch, and we attempt to classify among these contact types.

Consistent with our hypothesis, dramatic differences in our acquired signal were visible among these states, resulting in classification accuracies of 96% and 99% for our two participants (chance=33.3%).

This indicates that our approach allows robust discrimination among left-, right-, and two-handed contact.

Proximity to Wall
Based on initial observations that the amplitude of the sig-

nal received on the human body antenna changes continuously as participants' hands approached the wall prior to a touch, we hypothesized that the captured signal would provide a robust indication of the distance between a person's hand and the wall when no touch is occurring.

To investigate this hypothesis, participants rested their right hands above a measuring device for several seconds at 5, 10, 15, 20, 25, and 30 cm away from a light switch (along a line perpendicular to the wall). The participant's body did not move throughout this experiment, only the hand.

We used the same features as in our core experiment, and the same cross-validation procedure, but within each fold, rather than training a support vector machine classifier (which discriminates among discrete states), we trained a regression of our features onto the user's distance to the wall, using Weka's implementation of the SMO regression algorithm [17], which uses a support vector machine to map features into a high-dimensional space (as in an SVM classifier) and performs a linear regression in that space. Figure 6 shows the results across all regressions. The overall RMS error was 4.1 cm.

This indicates that our approach provides an indication of a user's distance from a wall containing electrical wires, with a resolution on the order of several centimeters.

Continuous Position Along Wall
The success of the "blank wall" results in the core experiment described earlier suggested that noise radiated from the power lines would vary continuously and predictably *along* a wall, offering the possibility of *continuous* touch localization. To assess this hypothesis, we again used a regression approach. In this case, participants rested their right-hand index finger against a wall at distances from 10 cm to 60 cm away from a light switch, *along* the wall, in one horizontal direction, at 10 cm increments (i.e., 10, 20, 30, 40, 50, and 60 cm). Figure 7 shows the results across all regressions. The overall RMS error was 8.1 cm.
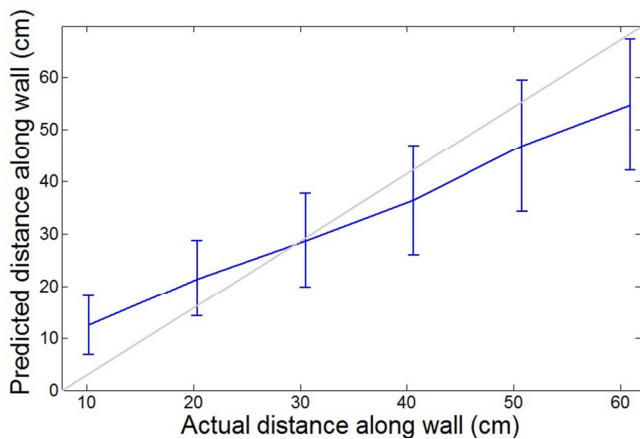
Figure 7: Regressing captured data onto a participant's touch position along a wall; distances indicate cm away from a light switch. Error bars represent standard deviations. Light gray line indicates correct prediction values.

This indicates that our approach provides an indication of a user's position along a noise-radiating wall, with a resolution on the order of several centimeters.

Interaction Switch Off
Our core experiments focused solely on light switches that were on (i.e., passing current). The noise environment at a light switch, however, is a function of the complete wiring pattern in the wall and the house, suggesting that passing current through the light switch at the center of interaction may not be necessary. We thus repeated our 5-position classification experiment around a single light switch in each of two homes, with the contact switch *off*. Classification performed at 99% and 97% (chance=20%) for our two participants.

This indicates that our approach performs well even for determining position relative to a light switch that is off.

Number of Fingers
Our ability to discriminate among one- and two-handed contact suggested that changes in contact area would be reflected in our signal, and that it might be possible to discriminate finer changes in contact area, such as the difference between single-finger and multi-finger contact.

To assess this hypothesis, our two participants conducted 5-second touches on a light switch using 1, 2, 3, 4, and 5 fingers on a single hand. Classification performed at 53% and 59% (chance=20%).

This is significantly lower than classification performance for the other capabilities explored in this section, but indicates that there is a relationship between contact area and our signal features. Additional work may be required to robustly discriminate fine changes in contact area.

DISCUSSION AND FUTURE WORK
Through our core experiment, we found that by measuring the electromagnetic noise received by the human body antenna, we can classify a person's location in the home with nearly 100% accuracy, and the position that a person is touching around a light switch with 87% accuracy on average. Through our additional explorations, we have demonstrated the ability to differentiate right and left hands, and determine which appliance is being touched. In addition, we have shown the ability to estimate the distance between a wall and a hand in the air, as well as the position of a hand along a wall.

These are promising results, which we believe can be integrated into an interactive real-time system for gesture sensing in uninstrumented homes. However, note that all of the classification accuracies reported in this work are on 82-millisecond windows. In an interactive system, it would be more natural to classify at the level of individual touches, which will likely *improve* classification performance, since smoothing can be performed over multiple windows.

Classification Features
We analyzed the classifiers built for our experiments to determine which features were most informative, by ranking the feature weights generated for each classifier. Over several hundred classifiers, the raw FFT amplitude spanning 60 Hz was the feature most frequently in the top ten highest-weighted features. Other features frequently ranked highly in our classifiers included the RMS amplitude and the FFT amplitude at 48 Hz (most likely a distortion of the 60 Hz signal). Perhaps more surprisingly, though, is the fact that several of our "high frequency peak" features in the kilohertz range were assigned large weights by most classifiers, motivating further exploration of these informative signals (Figure 5).

In addition to further exploring the "high frequency peaks", we would like to conduct an in-depth analysis of the feature space for this type of gesture sensing. We are currently using a naïve feature set, and are confident that our results can be improved – and several of our current limitations can be addressed – by using more appropriate features. For example, we have found examples of high-frequency peaks whose center frequency shifts as a function of the gesture. Exploration of these features remains future work.

Limitations and Future Improvements
Although our approach holds some promise, there are a number of limitations that we observed which will require further examination. The first is the generalizability of the observed noise signals. Since the electrical noise is a side effect of the power line infrastructure, there are no simple predictive models to infer what the signal will look like at different locations. Thus, the fingerprinting (classification) approach used here appears to be the most viable solution for now. This would require a user to calibrate and train the relevant gestures and locations for their home.

However, there are some simple techniques that could be used to help inform the user if a particular location would perform well. For example, the relative noise spread at a given location compared to the observed baseline noise in the house, and the observed signal strength, could be used to indicate a good location. Since it is possible to model

certain electronic devices in the home, such as the noise emitted from a switching power supply, this could allow us to actively reject or select certain features to improve the performance at a particular location.

Alternatively and perhaps more interestingly, if we allow ourselves to marginally instrument the environment, previous work shows that we may be able to infer activities such as which light switch has been flipped (and hence touched) [15]. Using this system, we might be able to passively train the system as the user goes about their daily activities, turning lights on and off.

With the electrical wiring being embedded in the walls, it may be difficult to determine good locations for interaction. Appliances, TVs, light switches, and outlets can all be used as a general guide and anchor points. For blank walls, since it is hard to see the actual electrical wiring, we could imagine using a "stud finder" approach with our system, in which the user runs their hand on the surface and audible feedback is played to indicate that the area under the hand contains enough signal for interaction.

Large inductive loads or very noisy dimmer switches may pose problems since they generate broadband noise that can potentially mask or overwhelm a previously observed noise signal. Some of these may only cause localized problems, while others may propagate throughout the entire home. Similarly, other loads operating in the home may temporally change the signature of a particular location. A potential solution may be to have the wearable device synchronized to a single device plugged into an outlet that is continuously monitoring the power line noise and computing difference vectors for classification. This may only work for a certain number of noisy devices, since the response of the power line is not necessarily linear and superposition may not hold. Another approach may be to have a plug-in device that generates a known broadband signal and injects it into the power line, similar to [13].

In all of our experiments, the electrical state of the home (i.e., which appliances and lights were turned on) remained constant throughout the testing session. We performed an informal exploration of the effects associated with changing the electrical state of the home. We found that the classification works well when the home is in the same state as it was during training; however, causing large changes in state (i.e., turning on the air conditioning, or all lights in the home) causes the classification accuracy to drop. Smaller changes to the electrical state of the home appear to reduce the classification accuracy to a lesser extent. We plan to develop a feature set that is more robust to these kinds of changes. Additionally, the user could train the system using a few significantly different electrical states (e.g., all lights off, all lights on, and air conditioning on). Using existing systems [9, 14, 15] is possible to passively determine when each of the electrical appliances in the home change states, and therefore our classifier can change its model based on the electrical state of the home.

During our core experiment, we taped over the ground screws of all light switches to ensure consistency between these touches. We hypothesized that touching the ground screw would short the human body antenna to ground, and would therefore significantly affect the received signal. We performed an informal experiment to test this theory. Based on our results, it appears that the received signal is significantly different when the body is grounded through the screws on the light switch. Not surprisingly, the signal looks similar to those seen when touching grounded appliances. Since the signal is significantly different when touching the ground screws, these can be considered an additional type of gestural input on each switch.

Our initial results suggest the contact pad may be located anywhere on the body. We would like to further explore non-intrusive form factors for connecting a computing device to the body, such as a wrist watch. Another possibility is to explore how well a short range air-coupled connection to the body may work. In this case, we could imagine having sensors integrated into a mobile device, which could either reside on belt clip or in a pocket. It would need to be able to sense the electromagnetic noise received by the human body without any physical contact to the body.

As we described earlier, the shape of an antenna dictates the received frequency response. We intend to explore free-space gestures conducted near the electrical power lines. Changes in posture and hand gestures could result in discernable frequency shifts, which may indicate a gesture.

We attempted to characterize the side-effects of our particular hardware, and minimize all effects of noise and interference; however, it is possible that the data collection hardware (i.e., the data acquisition unit and laptop) may have been injecting noise in the environment (though almost certainly not in a gesture-specific manner). We plan to more carefully characterize the side-effects of our equipment, and if it turns out that noise generated by the current hardware is beneficial to the operation or robustness of the system, future designs can intentionally generate these signals.

Applications
The ability to turn almost any wall surface or electrical device in the home into an interactive input system enables a breadth of applications. In the light switch scenario, we can imagine mapping a collection of gestures to digital lighting in a room without having to add additional physical switches to the space. Since we are able to identify the location of the gesture, the interaction can be mapped to specific parts of the home. This enables having arbitrary widgets being placed in the environment. Another application is a simple gesture that can be used to control the home's thermostat from anywhere in the house. For instance, tapping on the wall above and below any light switch could be mapped to increasing and decreasing the thermostat temperature. Similarly, these gestures could be mapped to controlling the music playing through the intercom or whole-home audio system.

In addition, the wearable computing unit has the side benefit of identifying the user performing the gesture. Thus, each person in a home could have their own device and can create custom gestures and map them to their own applications. The ability to scale the entire home into an input system also enables a breadth of new gaming and general computing applications that could easily be deployed in any home. Exploring and building these applications remains future work.

CONCLUSION

We have demonstrated the feasibility of a new interaction modality that utilizes the human body as a receiving antenna for ambient electromagnetic noise already in existence in our environments. While this noise poses problems for many sensing applications, we have used this phenomenon as our signal, thereby reducing the need to instrument the environment. By examining the noise picked up by the body, we have shown that we can infer the absolute touch position around a light switch or blank wall near electrical wiring within the home with nearly 87% accuracy. The location of which wall in the home a person is touching has nearly 100% classification accuracy. We also demonstrated the potential for hovering and continuously tracking a hand on a wall to enable a richer set of interaction techniques. Interacting with electrical devices and appliances also produces discernable changes in the received signal which could provide additional opportunities for further exploration. Although our initial experiments were conducted with rather bulky test equipment, this sensing modality only requires a wearable contact pad and an analog-to-digital converter, suggesting incorporation into an easy-to-deploy form factor such as a watch or mobile phone.

REFERENCES

1. Chen, W.T. and Chuang, H.R. Numerical Computation of the EM Coupling between a Circular Loop Antenna and a Full-Scale Human-Body Model. *IEEE Trans Microwave theory and Techniques,* 46.10 (1998).

2. Cho, N., Yan, L., Bae, J., Yoo, H.J. 60 kb/s-10 Mb/s Adaptive Frequency Hopping Transceiver for Interference-Resilient Body Channel Communication. *IEEE Jour Solid-State Circuits*, 44.3 (2009).

3. Chuang, H.R. and Chen, W.T. Computer Simulation of the Human-Body Effects on a Circular-Loop-Wire Antenna for Radio-Pager Communications at 152, 280, and 400 MHz. *IEEE Trans Vehicular Tech*, 36.3 (1997).

4. Cohn, G., Stuntebeck, E., Pandey, J., Otis, B., Abowd, G.D., Patel, S.N. SNUPI: Sensor Nodes Utilizing Powerline Infrastructure. *Proc ACM Ubicomp 2010*.

5. Dietz, P. and Leigh, D. DiamondTouch: A Multi-User Touch Technology. *Proc ACM UIST 2001*.

6. Fukumoto, M. and Tonomura, Y. Body Coupled FingeRing: Wireless Wearable Keyboard. *Proc ACM CHI 1997*.

7. Gabriel, S., Lau, R.W., Gabriel, C. The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz. *Phys Med Biol*, 41 (1996).

8. Gabriel, S., Lau, R.W., Gabriel, C. The dielectric properties of biological tissues: I. Literature Survey. *Phys Med Biol* 41 (1996).

9. Gupta, S., Reynolds, M.S., Patel, S.N. ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home. *Proc ACM Ubicomp 2010*.

10. Hall, P.S. and Hoa, Y. Antennas and Propagation for Body Centric Communications. *Proc Euro Conf on Antennas and Propagation 2006*.

11. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H. The WEKA Data Mining Software: An Update. *SIGKDD Explorations*, 11.1 (2009).

12. Harrison, C., Tan, D. Morris, D. Skinput: Appropriating the Body as an Input Surface. *Proc ACM CHI 2010*.

13. Patel, S.N., Truong, K.N., Abowd, G.D. PowerLine Positioning: A Practical Sub-Room-Level Indoor Location System for Domestic Use. *Proc ACM Ubicomp 2006*.

14. Patel, S.N., Robertson, T., Kientz, J.A., Reynolds, M.S., Abowd, G.D. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. *Proc ACM Ubicomp 2007*.

15. Patel, S.N., Stuntebeck, E.P., Robertson, T. PL-Tags: Detecting Batteryless Tags through the Power lines in a Building. *Proc Pervasive 2009*.

16. Saponas, S., Tan, D., Morris, D., Balakrishnan, R., Turner, J., Landay, J. Enabling Always-Available Input with Muscle-Computer Interfaces. *Proc ACM UIST 2009*.

17. Shevade, S.K., Keerthi, S.S., Bhattacharyya, C., Murthy, K.R.K. Improvements to the SMO Algorithm for SVM Regression. *IEEE Trans Neural Networks*, 11.5 (2000).

18. Smith, J.R. Field Mice: Extracting hand geometry from electric field measurements. *IBM Systems Journal*, 35.3 (1996).

19. Smola, A.J. and Schoelkopf, B. A tutorial on support vector regression. *Statistics and Computing*, 14.3 (2004).

20. Stuntebeck, E.P., Patel, S.N., Robertson, T., Reynolds, M.S., Abowd, G.D. Wideband power line positioning for indoor localization. *Proc ACM Ubicomp 2008*.

21. Zimmerman, T., Smith, J., Paradiso, J., Allport, D., Gershenfeld, N. Applying Electric Field Sensing to Human-Computer Interfaces. *Proc ACM CHI 1995*.

22. Zimmerman, T.G., Personal Area Networks: Nearfield intrabody communication. *IBM Systems Journal*, 35.3 (1996).