# Inferring TV Content from Electrical Noise

Miro Enev[1], Sidhant Gupta[1], Tadayoshi Kohno[1], Shwetak N. Patel[1, 2]
[1]Computer Science & Engineering, [2]Electrical Engineering
University of Washington, Seattle, WA, 98195
{mir0, sidhant, kohno, shwetak} @ uw.edu

## ABSTRACT
In this paper, we critically examine how modern switching power supplies found in many new consumer electronic devices are a source of information leakage through the powerline infrastructure. Unlike current consumption-based security vulnerabilities we show how a single easy to install plug-in device can infer the content of what is being watched on television by simply monitoring the electrical noise generated by the TVs power supply. We show that given a 5 minute recording of the electrical noise of a particular DVD movie, we can infer the movie from a database of noise signatures. In addition, we demonstrate this phenomenon in real homes despite the presence of noise from other electronic devices. We also discuss potential defenses one could employ to prevent others from eavesdropping over the power line.

## Categories and Subject Descriptors
K.6.5 Security and Protection: Unauthorized access; K.4.1 Public Policy Issues: Privacy

## General Terms
Algorithms, Experimentation, Security

## Keywords
Information Leakage, Electromagnetic Interference, Activity Recognition, Pattern Classification

## 1. INTRODUCTION
There is growing concern that a home's power consumption data could reveal private information about the occupants' personal activities. Indeed, just last month the Electronic Frontier Foundation (EFF) submitted a request to the California Public Utilities Commission, to petition the adoption of stronger laws to protect sensitive energy consumption data [21,22]. EFF's motivation for the policy change is based on discoveries revealing that power consumption information can be used to recognize "the use of most major home appliances" and more alarmingly to track "sleep, work, and travel habits" [22].

The EFF's claims are backed by recent findings that power consumption data can be used to infer appliances use. The rapidly evolving research strand of electrical sensing has reached a high level of specificity showing that it is possible to tell the difference between multiple devices used in a home simultaneously (by analyzing their unique noise signatures over the powerline) [6,16]. The principal goal of prior research in electrical sensing has been to aid users in adopting efficient energy habits and for developing activity recognition applications. As an example of a possible use case, consider a home monitoring device which determines that every evening between 8-10pm a single-occupant homeowner leaves the lights on in her bedroom and bathroom while watching TV in the living room. Having reached this conclusion the device informs the homeowner of the monetary benefits of turning off those unutilized lights and reducing her energy footprint. Such an advanced level of energy tracking and inference also enables numerous other applications which correlate consumption to activity. For example, the activation of a series of lights and electrical devices can help determine one's path or location within the home to aid in elder care by allowing a remote caretaker to assess the amount and nature of activity [17].

While such monitoring devices have the potential to increase efficiency and lead to quality of life improvements, the underlying methods are clearly unsettling when viewed through a privacy lens. Unfortunately, a privacy centric security analysis has been lacking in the energy sensing community which has thus far been exclusively focused on developing novel technologies while helping people become more conscientious consumers. Our present work flips this situation around and asks: how much information could one learn from monitoring a home's power line infrastructure? Is the electrical signal used for tracking consumption also capable of revealing private activity data? Said another way: are currently unknown forms of sensitive information leaking out over our power lines, waiting to be discovered?

To deeply explore into this question, we have chosen to extensively study the power-line information leakage due to the incidental electromagnetic noise generated from a single class of home appliances: televisions (TVs). We chose to focus on TVs because they are a nearly ubiquitous, high-end technology. Past research has shown that it is possible to detect when a TV is operating in a home [6,16]; but could the electromagnetic noise from a TV's switching power supply leak information beyond its on/off power state? We find that the answer is a definitive yes. Moreover, we show that given a 5 minute recording of the electrical noise unintentionally produced by the TV it is possible to infer exactly what someone is watching (with an average accuracy of 96% when considered over a possible set of 20 movies) by matching it to a database of content signatures. Even more surprising about our findings is that our sensor can be installed anywhere along the power line and does not require installing a power sensor in-line with the electrical power source of the home or the device of interest.

Given the potential exposure of private information over the powerline, a natural question is determining the actual risk to homeowners might be. We go into this question more deeply in the body of this paper, but stress two important points here. First, there are already natural entities that might be able to mount the attacks we consider here. For example, a power company with a modern smart power meter can remotely collect sufficient information to mount an attack. Moreover, anyone capable of attaching a device to a home's power line would be able to mount this attack (e.g., a parent wishing to track a child's TV viewing habits when the parent is not home, a neighbor plugging a device into an external power outlet, or the manufacturer of a Trojan appliance like a picture frame with wireless capabilities to exfiltrate data).

But, more importantly, we conjecture that (1) future appliances may leak even more information over the power line (a conjecture we support in the discussion sections of this paper given recent power efficiency mandates like Energy Star), and (2) as future homes become increasingly networked, new measurement vectors may appear over time. There are still challenges with developing defense mechanisms, because a tension begins to develop between the need for more energy efficient devices and preserving one's privacy.

For the remainder of the paper, we begin with a discussion of relevant prior work and go on to present the key concepts needed to understand the information leakage phenomenon over the power line. Next we shift our focus to detailing the experimental data collection and analysis workflow necessary to infer TV content from electrical noise. We then briefly sketch several motivating examples of threat models. In the last two sections we describe a theoretical model that can learn to mimic the electrical noise produced by a TV and conclude with a discussion of the universality of our approach, possible obfuscation mechanisms, and interesting security challenges.

# 2. RELATED WORK

The computer security literature has long been fascinated with information leakage through non-obvious channels. Although not brought to the public's attention until 1985 [23], evidence suggests that the government have long known that ancillary electromagnetic emissions from CRT devices can leak private information about what those devices might be displaying [7, 11]. This early work on studying electromagnetic information leakage from CRTs has since been extended to flat-panel displays [12] and wired and wireless keyboards [24]. The principal differences between this prior work and our own is that all the prior work uses electromagnetic interference that is emitted, that is, it travels through air and can be picked up wirelessly over a short range. Our work uses conducted electromagnetic interference which propagates from the device over to the power lines of a home.

Related to power consumption, but slightly further afield, is the broader area of power analysis and differential analysis for cryptographic processors [9]. Other examples of information leakage vectors include the time to perform various tasks (e.g., [10]), optical emanations (e.g., [14] for network appliances and [13] for CRTs), acoustic emanations (e.g., for printers [2], CPUs [19], and keyboards [4]), and reflections (e.g., [1]). In the modern television space, past work has also shown that it is possible to infer what someone might be watching over a wireless video stream from the size of the transmitted packets [18]; that approach exploits information leakage through variable bitrate encoding schemes, which was concurrently pioneered in [25].

Detecting electrical device activity and power consumption in the home has generally been done in the 'distributed sensing model' wherein each device being monitored is equipped with a separate sensor. This one sensor per device model is limiting because as the name suggests each monitored device requires separate instrumentation. Researchers in the ubiquitous computing field have been trying to use a single sensor approach in the home to infer human activity from the incidental noise produced by devices in the home as their signal. Gupta et al. accomplished this by using a single sensor that can be plugged into any available electrical outlet and analyzing the conducted electromagnetic interference (EMI) present on the power line in the frequency domain [6]. In this transformed space different devices occupy different frequency ranges centered around the switching frequencies of their power supplies. The presence or absence of such EMI is a direct consequence of the on or off state of a device respectively. In this work, we leverage the same fundamental phenomenon but move beyond detecting the power state of a device to infer the content being shown on the screen.

# 3. THEORY OF OPERATION

In this section we describe the fundamental theory behind the power line information leakage phenomenon which is made possible by unintentional electrical noise production in modern appliances. We also highlight the pros and cons of alternative methods that can provide access to the same information source.

## 3.1 EMI

The drive to produce smaller, cheaper and more efficient consumer electronics has made the use of Switched Mode Power Supplies (SMPS) increasingly prevalent. The adoption of SMPS is also spurred by policy guidelines as manufacturers strive to provide products which meet the efficiency requirements set by the Department of Energy's Energy Star program. In contrast to linear power regulation based supplies, SMPS do not dissipate excess energy as heat but rather store it in the magnetic field of an inductor. The load output of the inductor can be modulated by using a switch that allows current to flow when the circuit is closed (switch is on); thus by modulating the opening and closing of the switch the circuit is able to regulate the amount of power output. In modern SMPS this modulation, also known as the 'switching frequency,' happens at a very high rate (typically tens to hundreds of KHz). 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 and around the switching frequency. Due to the physical contact between the power line and the device 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. Because such EMI is undesired, in the US, the Federal Communications Commission (FCC) sets rules for any device that connects to the power line and limits the amount of EMI it can conduct (47CFR part 15/18 Consumer Emission Limits). This limit is set to -40dBm for a frequency range between 150 KHz to 500 KHz (which is much higher than the lowest levels of EMI that our prototype system can sense and capture effectively -100dBm). Figure 1 shows the EMI as captured by our system for various devices in a home. These include a compact fluorescent lamp (CFL), a modern LCD television, and other SMPS based devices.

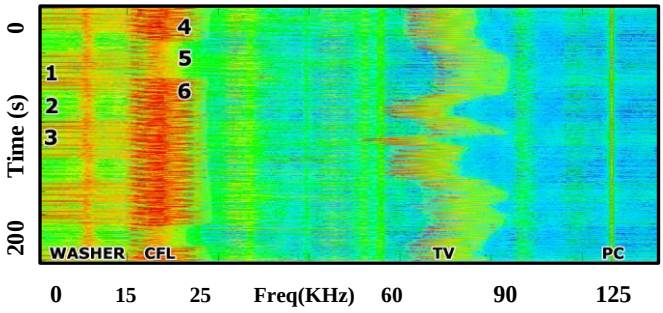


**Figure 1: Frequency domain Waterfall plot of the EMI spectrum in a home (red = high amplitude signal, blue = low amplitude). Note the dynamic nature of the TV EMI from 60 KHz (dark scene/low power consumption) to 90 KHz (bright scene/high power consumption). Interesting events for other devices are labeled numerically**

**{1: Washer Spin Cycle Left, 2: Off, 3: Spin Right}**

**{4: CFL Light ON, 5: OFF, 6: ON}**

Modern high definition liquid crystal display (LCD) televisions, which are of particular interest to our study, dominate today's consumer market and are almost always based on switched mode power supplies. As a result, the majority of modern TVs have power supplies that produce unintentional EMI. We implemented a system that records this signal [6] and in our experiments, we found that the TV produces a static band of EMI centered at the switching frequency of the SMPS. Furthermore we observed that the switching frequency (and EMI band) can be translated by altering the brightness setting of the television. Even more interesting is that the dynamic video content on the TV screen produces fluctuations in the EMI which leads to a time varying signal that fluctuates in a +/- 20 KHz window centered at the switching frequency (Figure 1). To better understand this phenomenon, we used an inline power sensor to determine the consumption of the TV in real time. We observed two things. First that the power consumption changes as a function of the screen brightness (menu setting) and second that it also fluctuates as a function of change in screen content.

To summarize, the brightness menu setting of the TV determines the baseline power consumption of the device, and changing this setting requires the SMPS to alter its switching frequency to match the load. In addition the dynamic visual content on the screen causes systematic fluctuations around this baseline since darker images require less energy while lighter screen content requires more. These content driven consumption changes manifest themselves as fluctuations in the EMI (which to reiterate, is an artifact of the SMPS's adjustments to match the power draw). In the case of our Sharp 42" LCD TV changes in brightness setting cause the center frequency of the EMI to be translated between 65 KHz and 75 KHz while modulating screen content cause the EMI to sway around this center (between 60 KHz and 90 KHz).

In the analysis that follows, we use the time varying EMI as a source of information about on screen-content and track this feature to determine what is being watched.

## 3.2 Current Consumption as a Feature

As described above, the power consumption of the TV is modulated by the nature of the dynamic screen content. Because power is the product of voltage and current, screen content changes should be manifested as changes in the amount of current that the TV draws.

To validate this hypothesis, we collected current consumption data alongside the EMI trace and found the signals to be identical; suggesting the validity of either approach for inferring the contents on a TV screen.

Though the current consumption data carries information about screen content, it comes with its own disadvantages in that current sensors have to be installed 'in line' with the TV. Ideally this means the sensor is attached to the power cord of the TV itself, or alternatively is instrumented inside the breaker panel. If the latter of these options is chosen, the sensor would also be reporting the current draw from all other devices in the home. Such an additive mixture of current consumption greatly complicates the isolation of the TV's signal. In contrast, the voltage EMI approach offers greater flexibility as it relies on a voltage sensor which could be plugged into any electrical outlet in the home. Moreover our EMI collection technique utilizes frequency domain analysis which allows us to simultaneously track multiple devices with low probability of signal clutter.

# 4. PROOF OF PLAUSIBILITY

To use EMI as a tool for inferring what is watched on a TV we needed to ensure that multiple recordings of the same visual inputs led to repeatable EMI signals while differing video content produced dissimilar EMI traces. To test whether these conditions were met we recorded data from four movies (60 minutes of data per movie) and repeated the recording three times (for a total of 3 recording sessions).

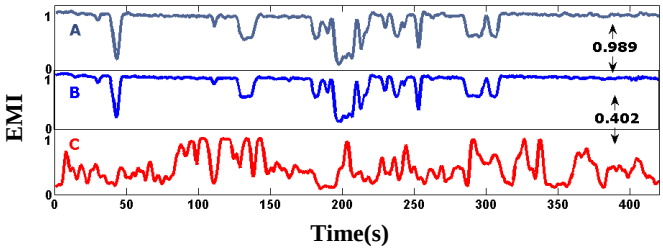


**Figure 2. Repeated recordings of identical screen content lead to nearly equivalent EMI traces (Panels A vs. B, Lion King, correlation = .989), while content from different movies produce distinct electrical noise patterns (B vs. C; Lion King vs. Bourne Ultimatum, correlation = .402).**

Next we analyzed the cross-correlation of the same movie between sessions and found that the similarity was consistently over 98% in all possible session pairings. This finding validated our requirement for signal consistency and the result is visually apparent in the top two panels (A, and B) of Figure 2, which represent 7 minutes of sample data from two recordings of the same movie (The Lion King).

When different movies are compared, the amount of cross correlation between their EMI signal traces is a function of the similarity of their content. Panel C of Figure 2 depicts a 7 minute trace from The Bourne Ultimatum which is apparently different from the data recorded from The Lion King (Panels A and B).

## 5. SYSTEM OVERVIEW

Our early experiments convinced us that there exists a strong relationship between EMI and screen content and that we had a sufficient platform to derive an algorithm capable of matching EMI traces from sections of movies to a film database in order to infer what is being watched. The next step was to build a recording setup that captures EMI from multiple movies, processes the signal, and populates a database with the EMI trace. We expected to process a large number of movies (multiple times) so we opted to create an automated data collection environment to guarantee consistency across recording sessions.

### 5.1 Hardware and Signal Processing

Our prototype consists of three main components (Figure 3). First, we connect a power line interface module (PLI) to any electrical outlet of the recording environment to gather the conducted EMI signal. Second, a high speed data acquisition module is used to digitize the incoming analog signals from the PLI. Lastly, a data collection and analysis PC running our custom software conditions and processes the incoming signals from the digitizer. We also connect a spectrum analyzer for debugging purposes and for visualizing the real-time EMI signal as a waterfall plot.

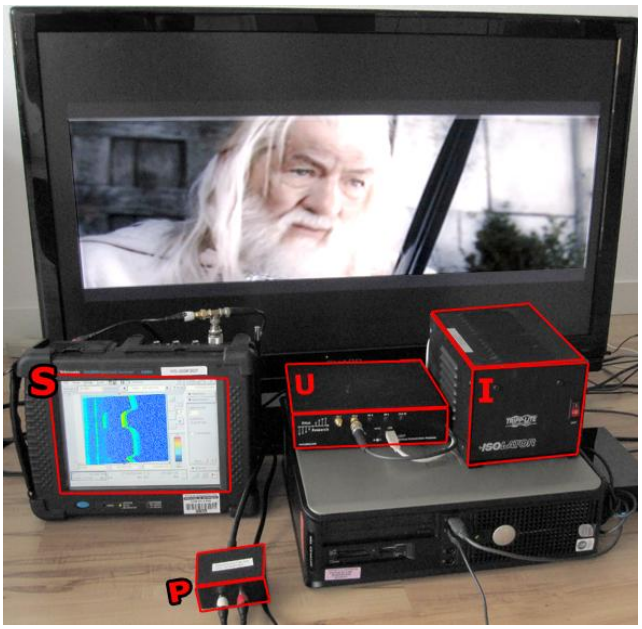


**Figure 3. Recording Hardware Setup. S = Spectrum Analyzer, P = Power Line Interface (PLI), U = Universal Software Radio (USRP), I = Isolating Transformer. The Sharp 42" LCD TV and the data logging PC are also visible.**

Of the components we use for data collection the only custom hardware is found within the PLI. The analog frontend PLI module is essentially a voltage sensor with a high pass filter that removes the AC line frequency (60 Hz in the US). This is necessary so that the dynamic range of the digitizer and the spectrum analyzer are not overwhelmed by the strong amplitude of the 60 Hz carrier wave and its harmonics (including the hazardous 120V output). The PLI's high-pass filter has a flat frequency response from 50 KHz to 30 MHz, allowing us to capture the entire range of conducted EMI. The analog signal from the PLI is then fed into a USRP (Universal Software Radio Peripheral) which acts as a high speed digitizer. We set the sampling rate of the USRP to 500 KHz, which (under the Nyquist Theorem) allows us to effectively analyze the spectrum from 0 to 250 KHz. The digitized data from the USRP is then streamed in real time over a USB connection to a PC.

We developed software on the PC which extends upon the GNU Radio Companion platform. Our system processes the incoming data and performs a real time Fast Fourier Transform (FFT) on the time domain signal arriving from the USRP. The output of the FFT is a frequency domain signal (or an FFT vector) of 2048 points which are spread uniformly over the entire spectral range from 0 to 250 KHz. The FFT vector is computed 122 times per second and its contents corresponds to the magnitude of the frequency strength along the range. The stream of FFT vectors is stored on the data recording PC for post-processing by our feature extraction algorithm. Figure 1 depicts a waterfall plot of a sequence of FFT vectors captured over a 200 second window.

### 5.2 Feature Extraction

Since we are only interested in tracking the TV, we post process the raw FFT bins and only retain the region around the TV's central frequency. This means that we reduce the 2048 element vector to 122 points in the 60-90KHz range wherein the EMI signal fluctuates (Figure 4).

In order to reduce the dimensionality of the data we perform a decimation to reduce the rate at which FFT vectors are processed. We found that to capture the variability in the EMI signal from the TV, using every 40$^{th}$ FFT vector (a decimation factor of 40) was sufficient.

Next we iterate through each time sample of the [abridged] 122 element FFT and extract the maximal element. We do this because we seek to compress the signal to a single point per time sample. The extracted maxes are then filtered using a 2nd order low-pass digital Butterworth filter with normalized cutoff frequency of .05 (frequency where the magnitude response of the filter is (1/2)^(1/2)). This removes the oscillation artifacts in the EMI and yields a smooth timeseries (EMI trace) whose shape tracks the fluctuations in the raw data (Figure 4 - blue overlay).

Prior to storing the EMI trace we perform mean removal (centering) and normalization to capture relative differences around the center frequency.

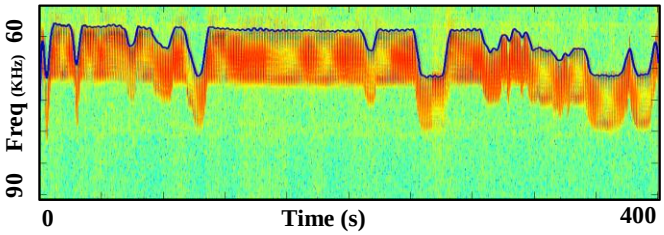


**Figure 4. Raw FFT signal around the EMI band of the TV. The result of post-processing this signal to extract the EMI trace features is overlaid as the blue time series. Note that unlike Figure 1 which also shows raw FFT data, the time axis is now along the horizontal.**

## 6. DATA COLLECTION

To validate our content inference approach, we recorded EMI data during the playback of 20 different movies on a Sharp LC-SB45U LCD 42" TV. At the end of the data collection our database contained three sessions of recordings (each session contained 20

movies, and only the first 60 minutes of each movie were considered to ensure data length consistency).

We hypothesized that there may be differences in the EMI features between genres so we tailored our choices to include 5 genres with 4 representative films per category. Our selection was informed by genre labels gathered from the internet movie database (IMDB, imdb.org) and in general we opted to choose titles which spanned a range of years and were among the most popular in their respective categories (see table below).

**Table 1: Movie Database Contents**

| | |
|---|---|
| **Action** | Lord of the Rings: Return of the King, Star Wars V: Empire Strikes Back, The Bourne Ultimatum, The Matrix |
| **Animation** | Wall-E, Shrek 2, The Lion King, Aladdin |
| **Comedy** | Office Space, Meet the Parents, The Hangover, Wedding Crashers |
| **Documentary** | Planet Earth: Fresh Waters, Food Inc., An Inconvenient Truth, Top Gear (s.14;ep.7) |
| **Drama** | The Shawshank Redemption, American Beauty, Titanic, Requiem for a Dream, |

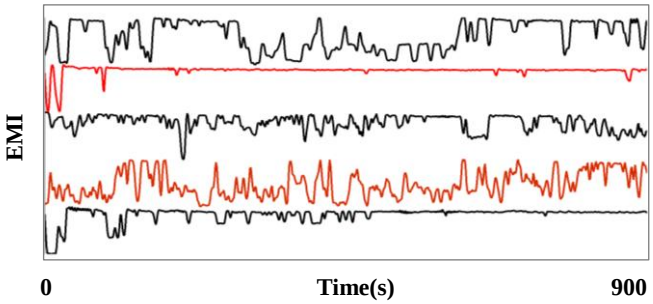


**Figure 5: Selected Movie Traces.**

**From Top To Bottom these are: An Inconvenient Truth, Meet the Parents, Wedding Crashers, The Bourne Ultimatum, and Shrek 2.**

For illustrative purposes the first 15 minutes of 4 movie EMI traces are shown in Figure 5. Note the elevated level of EMI fluctuation in the Bourne Ultimatum; this is typical of action movies which have a consistently high rate of scenes changes. Other than this observation we did not find any statistically significant differences between movie genres in our database.

## 6.1 Lab and Home Data

The three data recording sessions described above were performed in a lab environment. To demonstrate the applicability of our approach in naturalistic settings, we also recorded the same 20 movies in three homes.

The key difference between the data collection in the lab and home deployments was the use of a line isolation transformer in the lab setting (Tripp Lite 250W isolation transformer). A line isolator is essentially a broadband filter that removes any EMI present on the power line and presents an EMI free power output (are often used in audio/video recording studios and other high end applications). In the lab environment, we plugged the TV and the PLI into the line isolator's output to ensure that the PLI would have exclusive access to the EMI from the TV without interference from other electrical devices. In the naturalistic case, we collected data from the three homes *without* the line isolator, and the PLI captured EMI generated from the TV as well as myriad other devices (power adapters, CFL and dimmer based lighting, appliances etc.). In some instances, we found devices which generated electrical noise in the same range as the signal we were tracking (TV's EMI). As we show in section 8.2, despite such overlaps, the ability to infer screen content was relatively unhindered.

## 6.2 Automation

We opted to create an automated data collection environment to guarantee consistency across recording sessions. To this end we created a system which synchronized movie playback and data logging. The software running on the PC sent video content to the TV via a composite connection and simultaneously recorded data samples (computed FFT vectors) streaming in from the USRP to a binary file for post-processing and analysis.

# 7. ANALYSIS AND RESULTS

Once we constructed our database of reference EMI traces we could focus on designing a search method to find matches given a query trace. We crafted an algorithm that would take as input a query (snippet from an EMI movie trace), traverse the database, and return the movie with the highest similarity to the input.

## 7.1 Query Method

The query search problem we are faced with is an instance of subsequence matching [5]. The existing methods for problems of this type include spectral and statistical techniques as well as more recent approaches such as Dynamic Time Warping (DTW) and Semblance matching [3,8]. Due to the repeatability of the EMI signal we observed across recording sessions (see Section 4) we decided to forgo using dynamic programming measures of matching costs (i.e., DTW) since the signals we were comparing were not stretched in time. Furthermore we decided not to use spectral techniques (which would shift our analysis into the frequency domain) and instead found that the most natural way to express similarity between EMI traces was to use the cross-correlation coefficient (CCF).

The cross correlation coefficient (CCF) offers a statistical measure of the similarity between two timeseries and produces a numerical value ranging between -1 and 1 (higher values representing higher similarity; a CCF of 1 indicates identity) [20]. The inputs to the CCF computation are two time series of equal length; hence we used a sliding window approach to extract sequential snippets (of query length size) from the reference trace and for each sub-segment computed the CCF to the query. This results in a similarity vector whose maximum value represents the highest similarity between the query and movie pair; the index of the maximum represents the point within the movie EMI trace at which the best matching to the query occurs.

To obtain a query's best match, we compute the maximum CCF value across all movies in the database and declare the winner to be the movie with the highest CCF.

## 7.2 Experimental Evaluation

Successful matches were defined to be search instances whose winning match was the same movie that the query itself was extracted from. Consequently, accuracy was defined as the number of successful matches divided by the total number of

searches. As long as the query data was generated from an EMI trace of a movie included in our database we expected to have high classification accuracy.

To test this hypothesis we designed experiments to evaluate the effectiveness of our inference algorithm as we varied relevant parameters. We conducted a set of experiments in which we manipulated the following variables: query length, starting query location, and combinations of data sources for the query and database.

In order to investigate the effect of query length on accuracy we chose 9 monotonically increasing query lengths ranging from 15 seconds to 20 minutes (15s, 30s, 60s, 120s, 240s, 300s, 600s, 900s, 1200s**).** For each query length we generated 10 randomly chosen indexes (ranging between 0 and 3600 seconds) as query starting locations. Lastly to ensure that our metric is consistent across recordings we enumerated all possible pairings of sessions for query and database sources (Query from Session1: DB from Session2 , Q S2: DB S1, Q S1: DB S3, Q S3: DB S1, Q S2: DB S3, Q S3: DB S2). We then invoked the matching algorithm once for each possible parameter combination (9 * 10 * 6 = 540 runs)

## 7.3 Lab Results

A plot of the average accuracy (across session combinations and query start indexes) as a function of query length is shown in Figure 6. From this curve we can deduce that even short length queries lead to high accuracy classification. In particular, once the query length exceeds 4 minutes the accuracy reaches a rate of 95.7% (regardless from which part of the movie the query segment is chosen). Performance improvements due to extended query lengths (4 minutes and beyond) do not significantly change the average accuracy but they do reduce the variability in the results. This can be seen in Figure 7 which depicts averaged confusion matrices for selected query lengths (averaging is done across session combinations and query start indexes). The diagonal entries represent successful matches. Note the decrease in the perceived similarity of off-diagonal entries as the query length increases.

Movies 11 (Office Space) and 12 (Meet the Parents) were the worst performers and we believe that this is due to their consistently high brightness which produced very little fluctuation in their EMI traces.

## 7.4 Home Results

Having found convincing results in the lab setting, we were interested in validating our approach in naturalistic deployments. We setup our system in three different home environments and in each context recorded a smaller version of our database. All three homes were in the Seattle area; Home 1 was a typical suburban house in Lake City, Home 2 was a townhouse in the University District, and Home 3 was an apartment building in the Green Lake area.

The home data collection consisted of 10 minute segments collected from each of the 20 movies. Using this database we repeated the experiments described in Section 7.2 with the caveat that we fixed the query length to a 10 minute EMI trace (to exploit the entire recording from the home). The need for this longer query length was intended to offset the the increased noise conditions in the homes. The majority of appliances in a home do not disturb the signal quality of the TV EMI which we track however there are certain devices which produce obscuring noise (i.e. dimmer switches, washers, and vaccums). We did not limit the use of these and appliances and asked the residents of the home to ignore the recording system.

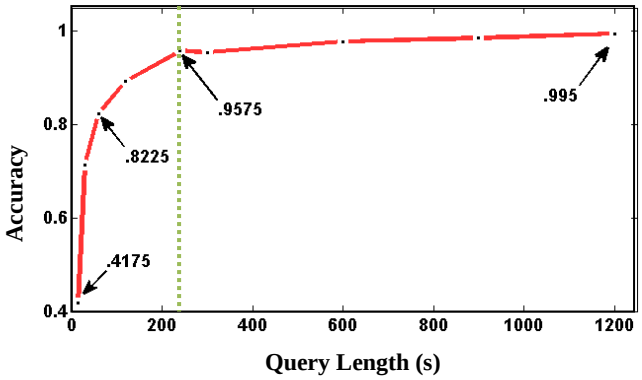


**Figure 6. Accuracy as a function of Query Length. Note that the accuracy improvements are minimal once the query length reaches 4 minutes (240 s) - indicated by the dashed line.**

In these home deployments the average accuracy was notably degraded (98% accuracy in lab for a 10 minute query vs 85.8% accuracy in homes). A thorough investigation of our approach in naturalistic settings is beyond the scope of the current work yet we feel that our preliminary study suggests the feasibility of EMI-based content inference in residential deployments.

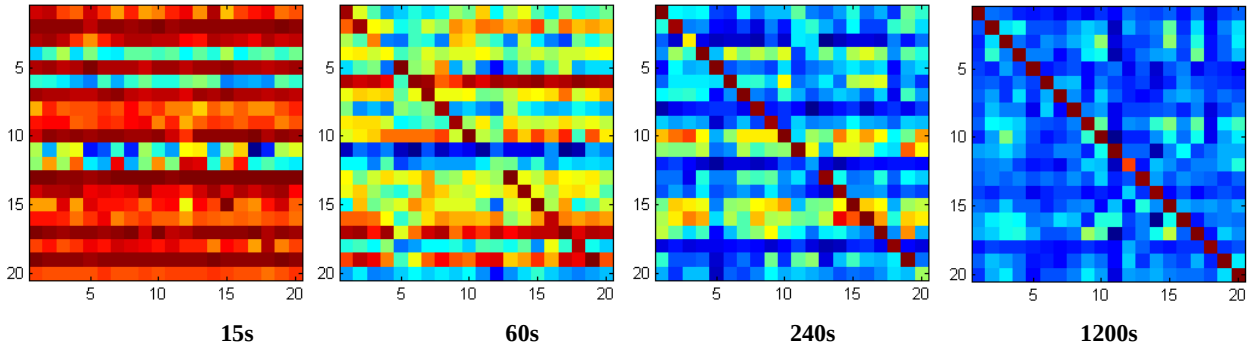


**Figure 7: Average confusion matrices for selected query lengths for the entire database. Red represent a high level of similarity and blue a low level.**

**Table 2: Content Classification Accuracy in Homes**

| Home # | Avg. Accuracy |
|---|---|
| 1 | 93.2% |
| 2 | 76.4% |
| 3 | 87.8% |

# 8. THREAT MODELS EXAMPLES

Although it is beyond the scope of our work to provide a thorough threat model analysis, we feel that it is informative to present several scenarios which highlight the diversity and scale of privacy breaking methods which leverage EMI content inference.

## 8.1 Surveying Burglary Targets

One potential adversary we consider is a thief interested in picking lucrative targets for burglary. Rob the robber, could build several databases of EMI noise from popular commercials recorded from the emissions of expensive televisions (one database of commercials per TV model of interest). Rob could then deploy a small form factor sensor in the external outlets of homes he is considering as targets for his next heist. Since the sensor equipment is relatively cheap, Rob could easily have a wide deployment that allows for the monitoring of multiple locations.

After deploying his web of sensors for a week, Rob could collect the 'field equipment' and analyze the data from each home using an algorithm similar to the one we've presented. More specifically, Rob could look for signs of strong correlations between the home's EMI data and any of the commercials for which he has stored traces in his database. Whenever he detects a match, he is able to conclude that with high probability the home in question houses an expensive television.

As an added bonus to the robber, the raw sensor data could reveal the number of EMI producing appliances in the home (a home with many appliances will be a more attractive target). If Rob was savvy he might even be able to detect activity patterns in the energy usage so that he can plan his attack when no one is likely to be home.

## 8.2 Large Scale Monitoring

A wealthy and trusted organization could track the EMI streams collected from smart sensors to infer what large numbers of people are watching. This data might then be sold for targeted advertisement. The monitoring could be accomplished by matching the EMI from a home against a database constructed 'on the fly' that mines the EMI streams of popular channels.

An adversary seeking to implement this large scale threat model would need physical access to X copies of Y television models for which the content tracking is desired (where X denotes the number of content streams/channels being monitored and Y denotes the number of television brands to track). Due to the large diversity of television manufacturers and the number of channels in the spectrum, the operational costs of the attack at first appear impractical for all but the wealthiest agencies.

Interestingly, the limitation of physicality can be surmounted if the adversary is armed with a quantitative model capable of generating an EMI trace given only an RGB video frame sequence for a given movie. In Section 9 we describe our success in developing such a generative model for the Sharp LCD TV using a recurrent artificial neural network.

# 9. LEARNING MODELS OF EMI

Motivated by the robust relationship we found between screen content and EMI we sought to reverse engineer the method by which electromagnetic noise is produced as a function of changing video input. Access to this transfer function would allow us to predict the EMI without actually laying out content on the screen and hence bypass the need for physical access to the target device. In the following section we investigate the plausibility of finding this function by framing the problem as an instance of supervised learning using a recurrent neural network with compressed input features.

## 9.1 Input Features

The transfer function we seek to approximate takes in as input a sequence of 3 dimensional RGB matrices (one per frame) and produces as output a time series of EMI (normalized between 0:1). The full input matrix is extremely high dimensional (~10^6 elements - **color** {R, G, B}* **screen width** {pixels} * **screen height** {pixels}) and prohibitively large for use in its full state. Thus we opted to compress each video frame into a 10 element vector which extracts selected features from the visual content and greatly reduces the complexity of the learning problem. Since we did not know which aspects of the screen content contribute most to the EMI signal we chose varied features in hopes that they would be sufficient to drive the learning. The features we derived from each video frame are as follows:

- **Brightness**: Cumulative sum of averaged RGB intensities
- **Flux**: average change in brightness b/w consecutive frames
- **Edge**: intensity - cumulative sum of Canny Edge filter output
- **FFT**: slope of the best fit line to an FFT of the image (the FFT shape becomes nearly linear after the frequency and amplitude axes are converted using a log-log scale)
- **Color**: mean and standard deviation of fitted gaussians for the R, G, and B color histograms (6 parameters total)

## 9.2 Neural Network

Since we are dealing with a function fitting problem of unknown complexity, we chose to use a recurrent neural network (RNN) model in order to accommodate for possible dynamic and non-linear effects. RNNs are a class of neural networks in which intermediate layers (i.e. those separating input and output) have connections to neighboring layers as well as (re)connections to themselves; these properties lead to self feedback (i.e. memory) which enable dynamic temporal behavior [15]

At time ***t*** the network input layer consisted of a video frame represented as a 10 element feature vector (section 9.1). The input layer was connected to the first of 3 hidden layers (connected in succession, each composed of 10 neurons to match the dimensionality of the input)[3] and the final hidden layer was connected to a scalar output layer representing the EMI at time ***t***.

The training phase began with randomly initialized network parameters which were tuned using backpropagation through time (BPTT) via the Levenberg-Marquardt gradient method. The

[3] Hidden layers used the hyperbolic tangent sigmoid as their transfer function.

criterion for performance was the how well the network output matched desired EMI for a given video input (measured as mean squared normalized error). Each training session concluded when the optimization converged or after 50 epochs (whichever came first). We ran several hundred training experiments and chose the network which performed best on sets of test inputs.

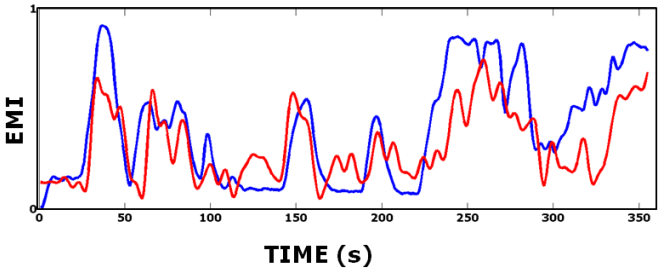


**Figure 8. Neural Net output (red) vs Ground Truth EMI (blue).**

Although there is much more that can be done in this line of analysis, our preliminary results are promising. (Figure 8) shows RNN predictions (driven via visual features) vs actual EMI of a 6 minute trace recorded from the opening segment of Lord of the Rings: The Two Towers. Though not perfect, the fit above clearly suggests that supervised methods can be used to train generative models of EMI.

# 10. DISCUSSION

## 10.1 Other TVs and Devices

Although we are only focused on a single TV, these results extend to other TVs and consumer electronic devices that employ similar power supplies (DVRs, PCs, power adaptors, CFLs, etc). The trend towards more efficient Energy Star compliant power supplies and even states mandating the use of switching power supplies implies an increase in these vulnerabilities in the near future.

As we demonstrated earlier in the paper, different devices exhibit EMI at varying center frequencies depending on the switching characteristics of the SMPS. The tolerance in the internal electronics that make of the SMPS can provide enough difference in the frequency domain to allow multiple similar devices to be observed simultaneously. Depending on the load characteristics of the electronic device, the switching frequency can range from anywhere between a few KHz to 1 MHZ.

LCD TVs tend to exhibit similar behavior between models and brands because of its general functionally. Newer LED TVs are also similar to LCDs, but the resonant frequency may be slightly different and the noise model may need modifications. The challenge with tracking new devices, however, is that they need to be tested to ensure that a strong relationship between EMI and screen content changes is present.

Similar to TVs, this is typically plausible in most consumer electronic devices. Often the power draw of a device can be strong indicator of its activity, which has been confirmed in prior work from the security community. Another way to think about SMPS EMI is that it is likely directly related to information leakage from the power draw, but can be inferred from line voltage which is much easier to tap into.

Beyond TVs, another popular device that we have observed information leakage is from home theater audio systems, where the output volume typically modulates the SMPS switching frequency. Some hi-end audio receivers also employ multiple power supplies, which would allow us to further infer the state of the receiver. Similar observations have been made with DVD players and power adaptors to most consumer electronic devices.

## 10.2 Potential Defenses

There are a number of potential defense mechanisms that could be used to minimize information leakage through EMI. The simplest is the use of a powerline isolator similar to the one used in our laboratory experiments. The internal transformer provides enough isolation that the high frequency noise does not pass back over the powerline. This does assume that the isolator itself has not been comprised. We have observed this isolation phenomenon in some, but not all, uninterruptable power supplies (UPSs). Most power strips only offer transient noise suppression and rarely offer any high frequency noise rejection. Newer home theatre line conditioners, which have a build in power bar, do offer some isolation capabilities.

A potential whole home solution, which does not require installing a device behind every electronic appliance, would be to inject random broadband noise over the powerline. The challenge with this approach is that it must conform to FCC regulations. In addition, this would cause problems with legitimate powerline-based communication systems like broadband over powerline and X10 home automatic systems. A more practical could identify potential devices that may be leaking information by observing the power line and only blocking certain frequency bands.

The other defense may be to employ new regulation on how SMPS power supplies are built. One critical observation, however, is that it may be impossible to fully defend against such information disclosure while still being in compliance with Energy STAR. Said another way, new government regulations may make it impossible or infeasible to protect privacy. The reason for this is most changes that would need to be incorporated into these device would likely cause an increase in power consumption and a reduction in efficiency, which is in conflict with recent legislation.

# 11. CONCLUSION

We have demonstrated that significant information leakage is present in modern switching power supplies found in many new consumer electronic devices. We have found that a single easy to install plug-in device can infer the content of what is being watched on television by simply monitoring the electrical noise generated by the TVs power supply. Only a 5 minute recording of the electrical noise of a particular movie, is needed to infer the movie from a database of noise signatures with up to 93% accuracy in actual homes. Although we have only demonstrated this with TVs, we believe our approach extends to other devices that employ SMPS. DVD players, power adaptors, and home theater systems all modulate their power draw during their operation, which can be used to infer its activity.

# 12. REFERENCES

[1] Backes, M., Dürmuth, M., and Unruh, D. 2008. Compromising Reflections-or-How to Read LCD Monitors around the Corner. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy* (May 18 - 21, 2008).

[2] Briol, R., "Emanations: How to Keep Your Data Confidential". In *Proceedings of Symposium on Electromagnetic Security for Information Protection*. 1991.

[3] Cooper, G. R. and Cowan, D. R. 2008. Comparing time series using wavelet-based semblance analysis. *Comput. Geosci.* 34, 2 (Feb. 2008), 95-102.

[4] Dmitri Asonov, Rakesh Agrawal, Keyboard Acoustic Emanations, *Security and Privacy*, IEEE Symposium on, p. 3, 2004 IEEE Symposium on Security and Privacy, 2004

[5] Faloutsos, C., Ranganathan, M., and Manolopoulos, Y. 1994. Fast subsequence matching in time-series databases. *SIGMOD Rec.* 23, 2 (Jun. 1994), 419-429.

[6] Gupta, S., Patel, S.N., Reynolds, M.S., ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home. *Under review in UbiComp 2010*.

[7] History of Tempest. http://cryptome.org/tempest-old.htm. Last accessed on 15th April 2010.

[8] Keogh, E. 2002. Exact indexing of dynamic time warping. In *Proceedings of the 28th international Conference on Very Large Data Bases* (Hong Kong, China, August 20 - 23, 2002).

[9] Kocher, P. C., Jaffe, J., and Jun, B. 1999. Differential Power Analysis. In Proceedings of the *19th Annual international Cryptology Conference on Advances in Cryptology* (August 15 - 19, 1999.

[10] Kocher, P. C. 1996. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Proceedings of the 16th Annual international Cryptology Conference on Advances in Cryptology* (August 18 - 22, 1996).

[11] Kuhn, M.G., Compromising emanations: eavesdropping risks of computer displays, 2003.

[12] Kuhn, M.G., Electromagnetic Eavesdropping Risks of Flat-Panel Displays. *4th Workshop on Privacy Enhancing Technologies*, 23–25 May 2004

[13] Kuhn, M. G. 2002. Optical Time-Domain Eavesdropping Risks of CRT Displays. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy* (May 12 - 15, 2002).

[14] Loughry, J. and Umphress, D. A. 2002. Information leakage from optical emanations. *ACM Trans. Inf. Syst. Secur.* 5, 3 (Aug. 2002), 262-289.

[15] Mandic, D. P. and Chambers, J. 2001 *Recurrent Neural Networks for Prediction: Learning Algorithms,Architectures and Stability*. John Wiley & Sons, Inc.

[16] 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: UbiComp 2007, pp. 271--288 (2007)

[17] Rowan, J., Mynatt, E.D.: Digital Family Portrait Field Trial: Support for Aging in Place. In: Proc of the ACM Conference on Human Factors in Computing Systems (CHI 2005), pp. 521–530. ACM Press, New York (2005).

[18] Saponas, T. S., Lester, J., Hartung, C., Agarwal, S., and Kohno, T. 2007. Devices that tell on you: privacy trends in consumer ubiquitous computing. In *Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium* (Boston, MA, August 06 - 10, 2007).

[19] Shamir, A., Tromer, E., Acoustic cryptanalysis on nosy people and noisy machines. http://people.csail.mit.edu/tromer/acoustic/. Last accessed: 15th April 2010.

[20] Spiegel, M. R. "Correlation Theory." Ch. 14 in *Theory and Problems of Probability and Statistics*, 2nd ed. New York: McGraw-Hill, pp. 294-323, 1992.

[21] Tien, L., EEF Comments filing to California Public Utilities Commission. http://www.eff.org/files/CDTEFFJointComment030910.pdf. Last accessed: 15th April 2010.

[22] Tien, L., New "Smart Meters" for Energy Use Put Privacy at Risk. http://www.eff.org/deeplinks/2010/03/new-smart-meters-energy-use-put-privacy-risk, EEF Blog. Last accessed: 15th April 2010.

[23] van Eck, W. 1985. Electromagnetic radiation from video display units: an eavesdropping risk?. *Computer &. Security*. 4, 4 (Dec. 1985), 269-28

[24] Vuoagnoux, M., Pasini, S., Compromising Electromagnetic Emanations of Wired and Wireless Keyboards. *In Proceedings of USENIX 2009*

[25] Wright, C. V., Ballard, L., Monrose, F., and Masson, G. M. 2007. Language identification of encrypted VoIP traffic: Alejandra y Roberto or Alice and Bob?. In *Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium*(Boston, MA, August 06 - 10, 2007)

## APPENDIX

**DEMO VIDEO OF EMI FROM TV**:

http://abstract.cs.washington.edu/~miro/sec_power/emi.avi

**ANALYSIS & VISUALIZATION TOOL**:

Note that the full movie traces are shown in the left panel

The source code (MATLAB) is available upon request from miro@cs.washington.edu