

Inferring Personal Information from Demand-Response Systems

Demand-response systems provide detailed power-consumption data to utilities and those angling to assist consumers in understanding and managing energy use. Such data reveals information about in-home activities that can be mined and combined with other readily available information to discover more about occupants' activities.

MIKHAIL
A. LISOVICH
Cornell
University

DEIRDRE
K. MULLIGAN
University of
California,
Berkeley

STEPHEN
B. WICKER
Cornell
University

In the US, a radical transformation of power-distribution systems is well under way. *Next-generation supervisory control and data acquisition* (NG-SCADA) architectures, now in development, will precipitate an exponential increase in both data collection and the extent of control available to consumers and utilities. The latter are increasingly adopting automated metering, advanced demand-response architectures, microgrids, and other systems that will provide cost savings in power generation, increase grid reliability and flexibility, and create new modes of consumer-utility interaction.

Large utility companies have deployed several pilot micro-grid projects (<http://certs.aeptechlab.com>) in recent years; they have also increasingly deployed *advanced metering infrastructure* (AMI) systems across the US. According to a 2008 Federal Energy Regulatory Commission staff report,¹ 5 percent of meters installed in the US are “smart” meters, and 8 percent of US customers participate in demand-response programs. Furthermore, the smart grid has recently become a presidential priority, receiving \$4.5 billion that must be spent on its deployment within the next two years. This infusion of funding means that market penetration for both smart meters and demand response is likely to increase dramatically in the near term.

NG-SCADA projects could benefit utilities, consumers, and new market players. For power companies, automated metering will reduce data collection costs and improve large-scale load planning and long-term research through real-time energy consumption feeds. This research will let utilities improve planning and test the effects of various demand-side manage-

ment programs.

For consumers, the projects will result in potentially lower costs, more information about consumption patterns, more control over power use, and the ability to actively participate in power generation. In addition, increased knowledge and management tools could reduce overall consumption. However, the engine for these activities—per-household consumption data—poses both privacy and security risks.

The Team for Research in Trustworthy Systems (TRUST) Science and Technology Center aims to promote a robust, secure, and trustworthy smart grid. (TRUST is a multi-university National Science Foundation center focused on trustworthy systems; see www.truststc.org.) Its teams focus on the confluence of sensor networking, power distribution, and policy to address the privacy and security issues that emerge from a substantial increase in power system monitoring at the consumer level. Our main claim here is that in the present regulatory and judicial environment, it's both possible and probable that interested parties will repurpose the household consumption data gathered via advanced metering projects to reveal and exploit personally identifying information.

Here, we explore the technical aspects of this claim, focusing on data generated, what it reveals, and how and why it might be collected and repurposed. (The long-term legal and privacy implications of in-home monitoring are discussed elsewhere.²) We highlight the importance of certain algorithms for extrapolating activity information from power-consumption data, present a formal way to evaluate information dis-

closure, and provide an illustrative proof-of-concept technical study. We also develop and substantiate certain aspects of our claim.

Exploiting the Smart Grid

The smart grid is bringing new, nontraditional players into the energy-consumption market, and many of them aren't abiding by existing privacy and security regulations. The data flows' real-time nature and increasing granularity will generate new interest in access and reuse by these players, which include law enforcement, marketing, and nefarious individuals.

Without proper technical, procedural, and legal safeguards, access to detailed household consumption data raises ethical concerns. Criminals could use it to facilitate burglary, marketers to initiate targeted advertising based on activities occurring wholly within the home, and law enforcement to monitor home-based activities in real time.

Although the sanctity of the home holds a special place among constitutional privacy protections in the US, businesses' capture and storage of information about private activities have eroded the protections afforded to it. In *United States v. Miller*, the Supreme Court held that individuals have no reasonable expectation of privacy in data voluntarily given to and held by third parties.³ Since this ruling, several state constitutions have been interpreted to provide some privacy protection for information about consumers captured in business records. Additionally, a leading Supreme Court case interpreted the Fourth Amendment as requiring law enforcement to obtain a warrant prior to aiming a thermal imaging device at a residence because it revealed detailed information about occupants' activities.⁴ However, it's unclear what, if any, constitutional limitations will apply to the government's potential access to and use of detailed energy records. Regardless, robust privacy protections are best produced through a mix of technology, regulations aimed at the private sector, and regulations—hopefully codified as constitutional limitations—on government use.

This issue came to prominence in a 2009 report by the US National Institute for Standards and Technology (NIST)⁵ that highlights the need for industry and policymakers to attend to privacy while developing smart grid standards and implementing plans. In addition, a recent ruling from the California Public Utilities Commission (CPUC)⁶ grants consumers the right to control many uses and disclosures of household-level consumption data. The technical design of and policies for the smart grid that the US and other areas of the world are currently working out will have a lasting impact on the privacy of in-home behavior. Consequently, we need a sustained and thoughtful research-based approach to establishing appropriate technical and legal protections.

Technology Overview

To familiarize readers with the technical aspects of this issue, let's first examine demand-response technologies, including AMI systems and non-intrusive load-monitoring (NILM) systems, a fundamental tool for extrapolating in-home activity. (More complete overviews of AMI and NILM are available elsewhere.^{7,8})

Advanced Metering

In a typical advanced metering setup, the customer is equipped with solid-state electronic meters that collect time-based consumption data at daily, hourly, or sub-hourly intervals. These meters then transmit the collected data to the *meter-data management system* (MDMS), which manages data storage and analysis, shaping the information into a form useful for the utility.⁷

Non-Intrusive Load Monitoring

An NILM system collects data much like its AMI counterpart but goes a step further by processing the data to determine individual electrical loads' operating schedules. The system typically does this by disaggregating the collected data stream into individual load signatures and matching each signature with reference signatures stored in a database. For private residences, these loads are usually appliances such as refrigerators, air conditioners, or water heaters. These systems are used for a wide variety of purposes, including collecting load research and implementing incentive programs for particular appliance usage patterns.⁹

Current NILM systems require electrical data sampled at second or sub-second intervals. Consequently, processing usually occurs locally at the electricity meter. However, we run data extraction remotely and obtain useful results even with the sparse data an AMI system provides. So, when considering how various players can repurpose power-consumption data and what kinds of information they can extract from it, we should consider an NILM algorithm as an essential building block. We develop this thought further in a later section.

Players, Use Cases, and Motivations

Utilities typically have policies that protect utility records and personal information. For example, the California Energy Commission requires the consumer's written consent before utilities can release personal data related to billing, credit, and power usage.¹⁰ Companies may release utility records in certain circumstances if the customer isn't identified, and exceptions exist for law enforcement access.

Given these policies, agencies, organizations, and individuals have natural motives for using power-consumption data for purposes other than load research and demand response.

Law Enforcement Agencies

Federal and state law enforcement agencies currently access utility records for a range of purposes. Current jurisprudence allows easy access to public utility records and provides legal precedent for their use in prosecuting criminal cases. (See the *Stanford Law Review* article² for a more in-depth discussion of current law on this topic.)

Police routinely use public utility records to seek out drug producers. *The Austin Chronicle* recently reported that the Austin Police Department has an agreement that lets it access Austin Energy power-usage records without a search warrant.¹¹ Investigators have used their access to screen consumers for possible drug production, relying on the fact that the heat lamps and watering systems used to grow marijuana indoors can increase a consumer's energy consumption far beyond the norm. Although the Austin case appears to be exceptional given that many utilities require a subpoena for releasing records, the program hints at the growth in use we might expect as increasingly detailed consumption data becomes available. As more granular consumption data begins to flow to utilities in real time, law enforcement interest in it is quite likely to grow.

Marketing Partners

Marketing firms can potentially use behavior and appliance usage information for directed advertisements. For example, some NILM systems are powerful enough to identify specific appliance brands and might even identify malfunctioning appliances.⁸ A marketing company partnering with a utility could use this data to send customers targeted advertisements for repairs or upgrades, or more generally derive demographic data for broader advertising claims. Targeted advertising based on in-home activities transgresses the current norms of information flow and creates new privacy concerns. The exposure of in-home activities and the resulting marketing might meet with strong consumer disapproval.

Criminals

In a previous article,² our TRUST colleagues give an excellent scenario for criminal abuse of power-consumption data: criminals could tap into an intermediate AMI node or simply monitor the unencrypted traffic between it and the individual meters. They could process the data to compile lists of household appliances or determine the occupancy patterns in an entire neighborhood. Such knowledge would facilitate burglary or some other property crime, whereas appliance lists will help with choosing targets.

Quantifying Information Disclosure

The previous section showed through examples that monitoring technology's evolution creates new risks

to individual privacy by exposing data previously held within the home to various parties. However, it isn't apparent just how we can quantify these risks, especially as a function of available data. A need exists for a "privacy metric" that associates the degree of data availability with potential privacy risks. Evaluating these risks is a complex process that must necessarily take into account common industry practices regarding data privacy, the state of current jurisprudence, consumers' privacy expectations, and the relationship between utilities and interested parties. Although we feel that such an analysis isn't within our purview, we believe that we can provide a crucial technical component for it—a "disclosure metric" that associates data quality (accuracy of readings, time resolution, types of readings, and so on) from a particular source with the information that the data could reveal.

To construct this disclosure metric, we need to better understand the nature of the information various players can extract from available sensor data. Thus, we start by suggesting a formal framework for extrapolating activity, then use it to construct our metric. We also suggest various privacy-theoretic frameworks that can be used in conjunction with the disclosure metric to move toward a robust privacy metric.

Extrapolating Activity

We can think of extrapolating activity as having two stages: the first "intermediate" stage uses NILM, in combination with data from other sensors, to extract appliance usage, track an individual's position, and match particular individuals to particular observed events. During the second stage, we combine this intermediate data with demographic data, such as the number, age, or sex of individuals in the residence, tax and income records, and models of typical human behavior. We can use these data together to identify activities, behaviors, preferences, and so on. The two stages aren't cleanly separated—we might use raw data directly to estimate a parameter of interest, and the determination of some intermediate parameters might rely on contextual information. However, many parameters in the second stage rely on the same intermediate data (for example, we might extrapolate both sleeping and eating habits from tracking data).

We can more readily quantify the first stage—for a particular algorithm, parameters take definite form (such as appliance use or condition, or location tracking), and we can evaluate performance in statistical or information-theoretic terms. However, defining an absolute performance limit for the second stage is more difficult—the number of specific preferences and beliefs that we could estimate is virtually limitless. To develop a comprehensive disclosure metric, we must carefully define a list of important parameters, basing importance both on how fundamental a

parameter is (how many other parameters we can derive from it) and on home or business owners' privacy expectations. Such expectations, in turn, are partially based on previous incidents of abuse and repurposing (such as the one described in the previous section). The list of second-stage parameters can be hierarchical, with more specific parameters used to evaluate more general ones. Once we define an appropriate list and assign importance values, we can determine the sufficiency of available data based on the requirements of current and future NILM, tracking, and other relevant algorithms.

A list of important, second-stage parameters establishes the evaluation criteria. Algorithms for estimating these parameters, along with the corresponding data requirements, provide a method for evaluating the available data's sufficiency. Together, these provide a metric for how much information a particular monitoring system could reveal.

Such a metric doesn't aspire to be an absolute measure, in the sense of encompassing all existing disclosure scenarios or anticipating all future ones. However, if thoughtfully implemented, it provides a valuable tool for evaluating privacy risks associated with a particular system, as well as allowing for comparison between similar systems. We construct a sample disclosure metric later.

Using the Disclosure Metric to Assess Privacy Loss

We can use the disclosure metric in conjunction with one of several privacy theories to assess actual privacy loss. For example, we can use Helen Nissenbaum's theory of privacy as contextual integrity¹² to examine how the new system's data loss relates to context-dependent norms of information appropriateness and flow. We could also use the principles laid out in Harry Surden's "Structural Privacy Rights" essay,¹³ along with Lawrence Lessig's "What Things Regulate" chapter in Code 2.0,¹⁴ to explain how the new data flow removes a structure that once afforded privacy protection for in-home activities (that is, the walls, combined with a low level of detail about energy consumption) and replaces it with a system that breaches the home's walls and exposes real-time consumption data.

Experiment

Although we already know we can accurately estimate first-stage parameters such as appliance usage (see the performance chart elsewhere⁹), and others have explored repurposing sensor data,¹⁵ to our knowledge, our group is the first to attempt extrapolating activity from power-consumption data. We want to prove that activity extrapolation is feasible, thus lending credibility to our main claim and providing an experimental precedent that others can cite in future efforts. To do

this, we conducted a small-scale monitoring experiment on a private residence.

Experimental Setup

We conducted our experiment in a typical student residence. We used the Brultech EML energy usage monitor for data gathering and attached it to the residence's breaker panel to send real-time power-usage information to a workstation responsible for data collection. The station recorded power usage at 1- or 15-second intervals and with a 1-watt resolution. The same workstation then ran the NILM and behavior-extraction algorithms. To evaluate the system's performance, we placed a network of cameras around the residence. We elected to use the Axis 206 network camera, which we connected to a workstation using an Ethernet switch. The workstation ran the Axis Camera Station software and recorded motion events for later processing. Figure 1a shows the residence's floor plan and the camera placement; Figure 1b shows the camera control and data-gathering setups.

Experimental Protocols

We ran the experiment semi-continuously over a two-week period. This timeframe let us obtain repeated data for pattern matching while accounting for time constraints. We shut down power and camera data collection software on a semi-daily basis for archiving, maintenance, and manual video data processing.

We collected electrical data from the house breaker and passed it to our behavior-extraction algorithm through a bridge program. The Axis Camera Station collected the camera data and stored it in MPEG format at a 320 × 240 resolution at 4 frames per second. At regular intervals, we manually analyzed video data and processed it into activity logs. Upon log completion, we deleted the original video data. Activity logs had the following format:

Date/time subject activity

The subject could be any of the house's three residents or a guest. Possible activities included turning any household appliances on or off (for example, *kitchen_lamp_1_on*), entering or leaving the residence, sleeping, preparing meals, taking a bath, or having a party. Because we didn't put cameras in individual rooms, the resulting activity logs weren't fully complete. However, this arrangement respected the residents' privacy and led to more natural behavior in areas under visual observation, whereas the collected data were sufficient to estimate parameters of interest (which we detail in a later section).

After collection, we subdivided the experiment data into two sets: a smaller three-day "training" set and a larger seven-day "experimental" set. Although

Exploiting In-Home Data




Figure 1. Experimental setup. We can see (a) the residence's floor plan along with (b) the camera and electrical data-gathering setups.

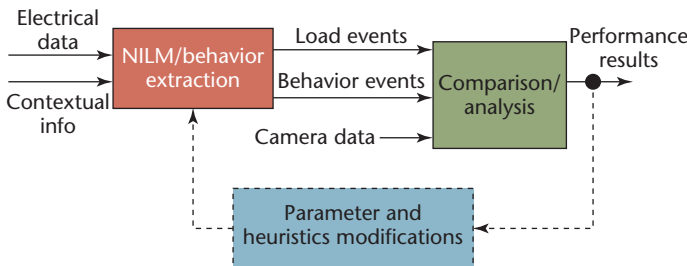


Figure 2. Flow of information between experiment components. We removed the parameter modification loop used during training for the experiment.

we actively modified our algorithms to increase performance on the training set, we kept them completely unchanged on the experimental set. Figure 2 shows the information flow between various components for

both the training and experimental stages (please refer to it as you read the subsequent sections).

Threat Model

For experimental purposes, we assume an adversary that has access to the *real power* data from a single household at a 1-watt power resolution and at least a 15-second time resolution. We further assume that the adversary has a list of appliances present inside, as well as their turn-on/turn-off profiles (not an unreasonable assumption—the Enetics NILM system⁹ has a built-in library of generic appliance profiles that you can match to unknown load signatures). Additionally, we assume that the adversary can distinguish between intermittent and periodic loads. We need not manually obtain this information—we can compile a list of intermittent appliances via reference software⁹ or through automated means,¹⁶ and we can automatically identify periodic loads using existing NILM algorithms.¹⁷

Parameters for Estimation

We chose several parameters that were both revealing and possible to estimate using our data-gathering equipment and processing algorithms:

- *presence/absence*—whether someone is present at the house;
- *appliance use*—status of the microwave, stove, water heater, TV, or miscellaneous appliances;
- *sleep/wake cycle*—when the household's occupants wake up and fall asleep; and
- *other significant events*—breakfast, dinner, showers, parties, and so on.

More formally, we begin by combining all the data into a single timeline. For each parameter, we partition this timeline into segments, with each segment assigned some value. For most parameters, the value is binary, indicating whether a person is present or absent, asleep or awake. For a specific parameter, we define the *i*th “on” interval by T_i^{on} and T_i^{off} .

Performance Metrics and Evaluation

Once we gather energy-use data and process it using behavior-extraction algorithms, we wish to compare the results against reference results obtained from camera data. To do this, we employ two classes of metrics. The first is event-based and consists of each parameter's failure-to-detect and misdetection percentages. We compute these percentages using the following procedure:

1. Define the cutoff threshold T_{thresh} , choosing it based on experimentation with training data.
2. For each parameter, examine the sequence of

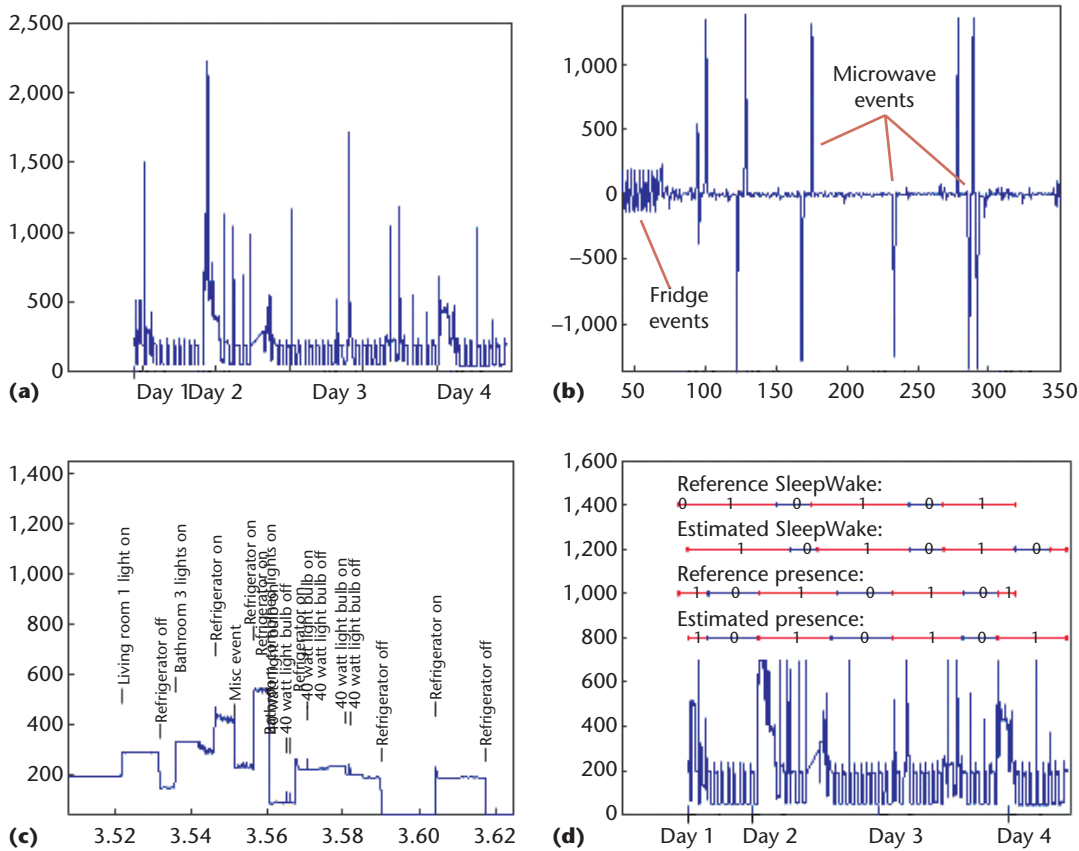


Figure 3. Behavior-extraction algorithm. We can see (a) the aggregate power-consumption data, (b) the derived switch events, (c) several identified load events, and (d) a comparison between reference and estimated intervals.

turn-on/turn-off events on both the reference and estimated intervals.

3. If a camera event occurs but a corresponding electrical event doesn't occur within T_{thresh} seconds, declare a *failure to detect*.
4. If an electrical event occurs but a corresponding camera event doesn't occur within T_{thresh} seconds, declare a *misdetection*.

The second class of metrics takes a broader perspective by computing the percentage of the reference interval that is correctly classified. This might in some cases be a better indicator of long-term performance because the algorithm could miss several short-duration events while classifying most of the interval correctly.

Together, these metrics help us get a well-rounded picture of the algorithm's performance, providing both detail and global perspective.

Behavior-Extraction Algorithms

We implemented our behavior-extraction system in Matlab; it comprises two major components: an NILM algorithm and a suite of functions that estimate the high-level parameters mentioned in the pre-

vious section. The NILM algorithm we implemented is based on an early MIT prototype.¹⁸ It analyzes the electrical data the load monitor gathers (see Figure 3a), performing edge detection, and cluster matching.

During edge detection, the algorithm computes a difference series $\Delta(t) = P(t) - P(t-1)$ from the electrical data $P(t)$. It merges adjacent $\Delta(t)$'s of the same sign and greater than a certain threshold into *switch events* (see Figure 3b).

During cluster matching, the algorithm matches switch events against a database of load signatures and classifies them as either "on" or "off" events. A load signature might be a switch event of a certain magnitude (a 40-watt light bulb has a step turn-on signature of $\Delta(t) = 40$ watts) or a series of such events (a refrigerator has a turn-on signature of $\Delta(t) = 1,100$ watts, $\Delta(t+1) = -960$ watts). The algorithm either discards unclassified events as noise or labels them with a catchall "misc. event" classifier. Figure 3c shows a sample of classified events.

Once the load events are classified, behavior-extraction routines use them to determine presence schedules, sleeping cycles, shower and bathroom use, mealtimes, and other activities. We briefly describe the most important routines:

Table 1. Algorithm performance.

	Sample size (referenced /estimated)	Reference events detected	Misdetects (%)	Interval (%)
Training data				
Presence	8/8	100	0	97.3
Sleep cycle	6/6	100	0	93.4
Microwave	8/8	50	78	43
Bathroom lights	8/8	72	44	52
Passage light	8/82	38	90	57
Living room lights	8/8	55	88	58
Experimental data				
Presence	10/10	80	20	97.4
Sleep cycle	12/10	83	0	92.3
Microwave	10/58	80	83	99
Bathroom lights	60/103	63	42	81
Passage light	8/82	38	90	57
Living room lights	19/179	21	89	52

- *Presence*. Because the refrigerator is the only load in the residence with automated turn-on/turn-off events, we assume that any nonrefrigerator event indicates presence. On the other hand, we can define absence by low power usage and a lack of events. An extended interval with low power usage during which no events occur implies that all subjects have left the residence.
- *Sleep cycle*. Intervals of inactivity occurring between late evening and early morning imply that all people are sleeping (as opposed to absent). So, we reclassify all such absence intervals as sleep intervals.

The last major component of our system is the analysis suite. Reference data derived from camera logs is automatically processed into reference intervals, which are then compared against estimated intervals using metrics described in the previous section.

Results

As mentioned, we ran our algorithms on a three-day training set and a larger seven-day experimental set. Table 1 shows the results. For each quantity, the table shows the number of events recorded, the percentage of successfully detected reference events, the mis-detection percentage, and the percentage of the reference interval correctly classified.

One important appliance Table 1 leaves out is the refrigerator, which autonomously cycles between high and low states. Unfortunately, we didn't observe these state transitions directly, which would have required a separate energy monitor exclusively for the refrigerator. However, we can comment on the algorithm's performance by manually examining the electrical data readout (a refrigerator has a distinctive operating profile—see Figure 3c). For the training data set, 101

out of roughly 104 refrigerator events (more than 97 percent) were correctly classified. The success rate was similarly high for the experimental data set.

Generally, the algorithm performed quite well in determining presence and sleep cycles. In both cases, it correctly classified more than 90 percent of the total interval length for both training and experimental data. We believe this is due to our success in identifying the refrigerator load, the small number of autonomous appliances in the residence, and the consequent simplicity of the presence and sleep-wake heuristics.

The algorithm was also relatively successful at classifying microwave and bathroom light events, supporting our hypothesis that it's possible to predict both meal times and shower times with moderate success. Living room light detection was unreliable, partly due to the presence of other 40-watt light bulbs in the residence.

It's worth noting the high percentage of misdetections. We believe this is caused in equal part by the limited capabilities of our data-gathering system and algorithms, and by our camera monitoring setup's imperfections. On the behavior-extraction side, our data logger recorded only real power and only at 15-second intervals, while our algorithm was tilted toward *false alarm* rather than *failure to detect*. On the camera side, our camera wasn't in a position to observe all loads directly, which meant that we sometimes missed turn-on/turn-off events during manual processing. Consequently, for appliances, the percentage of reference points detected is the most credible measure of algorithm performance.

Also, for appliances, *percent interval correctly classified* isn't necessarily meaningful because the *zero-performance point*, defined as the performance when the entire estimated interval is set to 0, is still 50 percent for microwave and bathroom lights.




Figure 4. Degree of disclosure. This metric measures the ability of our behavior extraction system to reveal various in-home activities by associating data quality (accuracy of readings, time resolution, types of readings, and so on) with the information that the data can reveal.

Degree of Disclosure

Figure 4 shows an implementation of the disclosure metric whose construction we discussed previously. We select eight important parameters that our NILM-based behavior-extraction algorithm could reveal, then rate the amount of information disclosure for each as *negligible*, *slight*, *fair*, *high*, or *severe*. Where possible and appropriate, we also provide a numerical measure from one to 100.

We begin with *aggregate presence* and *sleep schedules*—these are useful to a broad range of players, and their performance transfers directly from Table 1. Based on arguments made previously for balancing the short- and long-term perspectives, we average performances in columns three and five. The resulting performance numbers—more than 90 percent in both cases, seem quite good. However, we can’t track the number of people currently present or asleep, which detracts from the performance somewhat. We rate the degree of disclosure for both as high.

However, the algorithm’s tracking ability is only fair—it can localize occupants to the kitchen or bathroom, but can’t reliably localize individual movement within the house. Without the prior behavioral profiles and ancillary sensor data, the algorithm’s ability to assign events to individuals is practically nonexistent.

For appliances and appliance-derived parameters, we calculate performance by averaging all relevant entries in column three of Table 1 and subtracting an average of entries in column four weighted by .2 (because misdetects are indicative but unreliable for reasons we listed earlier).

The algorithm’s overall ability to identify appliances is only fair, due to both the fairly low detection rate and the high percentage of misdetects. However, its ability to identify large, distinctive appliances—in this case, the refrigerator and microwave—is high. Disclosure about meal times, which we obtain by windowing the microwave use detection/misdetection ratio

with likely timeframes for meals, is fair. Information about shower times, derived directly from bathroom light and fan use, is also fair. Overall disclosure is the weighted (in our case equally) average of all these parameters; we qualify it as fair.

We can now evaluate the overall threat to individual privacy by taking these results on potential disclosure from consumption data, qualifying them with the likelihood and degree of disclosure, historical precedent for repurposing, and the relationship between the data holders (utilities) and interested parties. We don’t presume to conduct such an analysis in this article,

Discussion

Our experiment shows that the algorithm can estimate presence events and sleep cycles with high confidence, at least for a household with few appliances and relatively infrequent switching events. However, although the experiment illustrates the system’s potential, it doesn’t comprehensively characterize its capabilities. First, the experiment’s scale—a week’s worth of data from a single residence—is too small to draw conclusions on the system’s limitations. Conversely, the adversarial model might also be too strong; in practice, an adversary might not have an actual list of the appliances inside a home, which would lead to less reliable presence and sleep predictions.

Despite these caveats, we believe that our results are sound and that, moreover, our approach has the potential for significant refinement. First, we note that the residence didn’t have an electric stove or a water heater—two readily identifiable loads whose “on” intervals directly correspond to meal times, laundry, and showers. Second, we used only electrical data; a behavior-extraction algorithm can combine data streams from electric, water, gas, humidity, and any other available sensors. Third, our data resolution (15 seconds in most cases) was relatively low, and our behavior-extraction algorithms were relatively unsophis-

ticated, because our aim was to prove feasibility and not to optimize performance. NILM and behavior-extraction systems of the near future will surely surpass our effort in performance, enabling person-to-event assignments and perhaps even limited tracking.

On the other hand, we believe less potent technology can still extract useful data. Various players could use hourly power averages, such as those that California's AMI system produces, to determine presence and sleep cycles (although to a coarser degree) and identify major appliances with substantial steady-state power consumption (such as heat lamps).

Note that future concerns aren't limited to these systems' performance at an individual household level. Because the algorithms are fully automated, interested parties could do analysis on extremely large scales, involving hundreds or thousands of residences. Easy access to information will inevitably generate a market for it.

Data-Handling Guidelines

A report recently submitted to the California Energy Commission¹⁹ makes several recommendations for power-data handling, including

- multiple tiers of control and oversight, both by the utilities themselves and state and federal governments;
- explicit guidelines regulating access to data for customer service, load research, and other functions;
- strong user control over information leaving the residence; and
- protocols that do most of the data processing at stations located inside the residence, as well hard prohibitions against relaying certain types of data.

One of the report's main points is that state and federal governments should carefully regulate the mining of hourly usage data. The authors advise policymakers to adopt more stringent rules on the use, release, and reuse of energy consumption data as data mining practices develop and new information in which consumers have a reasonable expectation of privacy is exposed.

This article details the sorts of conclusions that various players can readily draw from power-consumption data. Our discussion of motivations shows that the decrease in the time interval between readings of energy consumption—likely to real or near-real-time—will create new interest in repurposing consumption data. Our technology discussion and proof-of-concept demonstration show that even the simplest data mining and pattern-matching tools can convert power consumption data into information about events within “the sacred precincts of private and domestic life,”² illustrating the extent to which such players could violate residential privacy by col-

lecting and using power-consumption data. Finally, the disclosure metric we propose and implement facilitates privacy risk violation, which could let policy makers more precisely define the permitted and prohibited uses of data mining.

Discussion and advocacy efforts are already under way to address this problem. Jack Lerner and Diedre Mulligan have written an article² chronicling court opinion's evolution toward energy-data privacy and calling for its constitutional protection. They've also collaborated with the CPUC to develop a set of draft guidelines¹⁹ for a secure and privacy-preserving demand-response infrastructure. We hope that this article helps those seeking to create NG-SCADA technologies that respect and safeguard consumer privacy. □

Acknowledgments

We sincerely thank Devashree Trivedi, who provided a helpful presence and equally helpful input during every stage of the project, and who single-handedly ran data gathering during the experimental stage. We also thank Michael Baranski, whose papers helped us during project development and who provided comparison results using his own non-intrusive load-monitoring algorithm. Finally, we thank Judith Cardell, Jack Lerner, and others who were very helpful throughout the project's duration, and the US National Science Foundation for its funding of the TRUST center and our research.

References

1. “2008 Assessment of Demand Response and Advanced Metering,” staff report, Dec. 2008; www.ferc.gov/legal/staff-reports/12-08-demand-response.pdf.
2. J.I. Lerner and D.K. Mulligan, “Taking the ‘Long View’ on the Fourth Amendment: Stored Records and the Sanctity of the Home,” *Stanford Technology Law Rev.* 3, 2008; <http://certs.aeptechlab.com>.
3. *United States v. Miller*, *US Reports*, vol. 425, 1976, p. 435.
4. *Kyllo v. United States*, *US Reports*, vol. 533, 2000, p. 27.
5. A. Lee, *Privacy and the Law in Demand Response Energy Systems*, tech. report, US Nat'l Inst. for Standards and Technology, Sept. 2009; <http://csrc.nist.gov/publications/drafts/nistir-7628/draft-nistir-7628.pdf>.
6. “Assigned Commissioner and Administrative Law Judge's Joint Ruling Inviting Comments on Proposed Policies and Findings Pertaining to the Smart Grid Policies Established by the Energy Information and Security Act of 2007,” Calif. Public Utilities Commission, Sept. 2009; <http://docs.cpuc.ca.gov/efile/RULINGS/107583.pdf>.
7. “Advanced Metering Infrastructure (AMI),” Electrical Power Research Inst., Feb. 2007; www.ferc.gov/EventCalendar/Files/20070423091846-EPRI%20-%20Advanced%20Metering.pdf.
8. C. Laughman et al., “Power Signature Analysis,” *IEEE Power and Energy Magazine*, vol. 1, no. 2, 2003, pp. 56–63.
9. “Single Point End—Use Energy Disaggregation

- (SPEED) Marketing Brochure,” 2001; [www.enetics.com/downloads/SPEED Brochure.pdf](http://www.enetics.com/downloads/SPEED%20Brochure.pdf).
10. D.K. Mulligan et al., “Privacy and the Law in Demand Response Energy Systems,” *Samuelson Law, Technology and Public Policy Clinic*, 2006; www.truststc.org/pubs/36/Jones_PrivacyAndLawInDemandResponse.pdf.
 11. J. Smith, “APD Pot Hunters are Data Mining at AE,” Nov 2007, www.austinchronicle.com/gyrobase/Issue/story?oid=oid%3A561535.
 12. H. Nissenbaum, “Privacy as Contextual Integrity,” *Washington Law Rev.*, vol. 79, no. 1, 2004, pp. 119–158.
 13. H. Surden, “Structural Rights in Privacy,” *SMU Law Rev.*, vol. 60, , 2007, pp. 1605–1629.
 14. L. Lessig, *Code: And Other Laws of Cyberspace, Version 2.0*, 2nd ed., Basic Books, 2006.
 15. M.L.A.P. Jun Han and Abhishek Jain, “Don’t Sweat Your Privacy: Using Humidity to Detect Human Presence,” *Proc. 5th Int’l Workshop on Privacy in Ubi-Comp (UbiPriv 07)*, 2007; http://sparrow.ece.cmu.edu/group/pub/han_jain_luk_perrig_privacy.pdf.
 16. M. Baranski and V. Jurgen, “Genetic Algorithm for Pattern Detection in NIALM Systems,” *Proc. IEEE Int’l Conf. Systems, Man and Cybernetics*, IEEE Press, vol. 4, 2004, pp. 3462–3468.
 17. M. Baranski and V. Jurgen, “Detecting Patterns of Appliances from Total Load Data Using a Dynamic Programming Approach,” *Proc. IEEE 4th Int’l Conf. Data Mining (ICDM 04)*, IEEE Press, 2004, pp. 327–330.
 18. S. Drenker and A. Kader, “Nonintrusive Monitoring of Electric Loads,” *Proc. IEEE Computer Applications in Power*, Dec. 1999, pp. 47–51.
 19. P. Subrahmanyam et al., *Network Security Architecture for Demand Response/Sensor Networks*, tech. report, Calif. Energy Commission, Public Interest Research Group, Jan. 2008.

Mikhail Lisovich is a PhD candidate in the School of Electrical and Computer Engineering at Cornell University. His research interests include mobile wireless networks, as well as privacy and policy challenges related to upcoming pervasive wireless sensor networks. Lisovich has a BS in electrical engineering and physics from the Pennsylvania State University. He’s a member of the Team for Research in Trustworthy Systems (TRUST) Science and Technology Center and a participant in the US National Science Foundation’s Mobile Autonomous Systems and Technology (MAST) consortium. Contact him at mal86@cornell.edu.

Deirdre K. Mulligan is an assistant professor at the University of California, Berkeley, School of Information. Her current research agenda focuses on information privacy, confidentiality, and network security. Mulligan has a J.D. in law from Georgetown University. She’s the lead for law and policy for the multi-university Team for Research in Trustworthy Systems (TRUST) Science and Technology Center. Contact her at dkm@ischool.berkeley.edu.

Stephen B. Wicker is a professor of electrical and computer engineering at Cornell University. His current research interests include self-configuring wireless sensor networks and privacy-aware design for mobile computing/communication. Wicker has a PhD in electrical engineering from the University of Southern California. He’s Cornell’s principal investigator for the Team for Research in Trustworthy Systems (TRUST) Science and Technology Center, a US National Science Foundation center dedicated to the development of technologies for securing the nation’s critical infrastructure. Contact him at wicker@ece.cornell.edu.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.