Eurographics Conference on Visualization (EuroVis) 2015
H. Carr, K.-L. Ma, and G. Santucci
(Guest Editors)

Volume 34 (2015), Number 3

Visualization-assisted Insights into Home Energy Usage

Matthew Law [1], Sami Rollins[2], Nilanjan Banerjee[1], and Alark Joshi[2]

[1]University of Maryland, Baltimore County
[2]University of San Francisco

Abstract

Motivated by both cost savings and environmental concerns, managing energy in the home has become increasingly important. Though both user-driven and automated solutions have shown promise, a deeper understanding of the characteristics of energy usage behavior is necessary to inform the design of such systems. In this work, we present the results of a study that explores the insights and characteristics participants are able to derive using two well-known visualization techniques for time series data. We find that participants are able to extract a variety of relevant features and explain general and anomalous patterns of behavior. We also find that the preferred visualization is both user- and task-dependent. These findings may be used as the basis for new systems for home energy management.

Categories and Subject Descriptors (according to ACM CCS): H.1.2 [Information Systems]: User/Machine Systems—Human Information Processing

1. Introduction

Tools † and techniques for influencing the energy consumption behavior of residential users are the subject of much current research. Saving money, reducing environmental impact, and enabling broader adoption of greener energy sources are among the factors motivating the work in this area. Wider deployment of smart meters that collect fine-grained home energy usage data, smart power strips that collect energy consumption data for a collection of appliances and smart plugs that capture energy usage of individual appliances open a number of opportunities for customized energy management solutions [DWD11].

Enabling an individual user to better understand his or her energy consumption is a promising step toward influencing behavior. By making a user aware of, for example, the typical energy consumption of a particular appliance, the user may be able to detect periods of anomalous usage when the appliance may have been malfunctioning or left on unintentionally. Turning the appliance off, or even replacing it, may lead to energy savings. As another example, raising awareness of just how much energy is consumed by various devices, for example a computer monitor that is left on even when not in use, may encourage the user to modify how the device is used or when it is switched off. Technology-driven, automated solutions for reducing or shifting energy usage are also possible, however they must be carefully tuned, for example using features known to be effective or training data that has been manually labeled.

In this work, we explore the effectiveness of using data visualization to help users understand characteristics of energy usage behavior. We apply two well-known visualization techniques for time series data—a line graph and a radial layout—and conduct a study of 21 participants to explore the features they are able to identify and the insights they are able to draw from the visualizations. The participants were asked to perform two tasks—a device classification task and a pattern and anomaly detection task.

We present three main findings. First, with minimal training, a small data set, and minimal contextual information, users were able to identify relevant features, such as power consumed and frequency and duration of appliance usage, that led to successful completion of the tasks. Second, participants were able to identify anomalous behavior and general patterns of use that may be generalized to a broad range

† This material is based upon work supported by the National Science Foundation under awards CNS-1305099 and IIS-1406626, CNS-1308723, CNS-1314024, and the Microsoft SEIF Awards.

of devices. Finally, using multiple visualizations was effective as the preferred visualization was both user- and task-dependent. These insights may be used to form the basis of visualization-based recommendation and energy audit systems [BGMS10], as well as inform the design of machine learning-based techniques [KJ12].

2. Related Work

This paper builds on previous research on energy data visualization and time series visualization.

Energy visualization: Ambient appliances are the most straightforward way to increase the awareness of energy usage in the home [GG05, MAR10]. FigureEnergy [CRJ12] is an interactive system that allows residents to visualize and manually annotate their energy. Masoodian et al. [MEB+13] present a novel radial stacked bar-based technique for visualizing energy usage called the *time-pie*. Bartram et al. [BRM10,RB11] present their Adaptive Living Interface System that allows home owners to monitor their energy usage. They discuss challenges with evaluating such an interface since it is currently operational in a net zero home. Similarly, Pierce at al. [POB08] identified *evaluating the efficacy* of eco-visualizations as one of the challenges in the field. Goodwin et al. [GD12, GDJ+13] discuss an inclusive design-oriented approach that led to novel energy visualizations. Spagnolli et al. [SCG+11, GCZ+11] developed a gamification-based approach to energy conservation in the home. Janetzko et al. [JSMK14] present novel anomaly detection algorithms based on prediction as well as clustering. Our work explores the characteristics that participants are able to derive using visualizations of home energy data.

Time-series visualization: Cartesian and Radial layouts (used in this study) are widely used for visualizing time-series data [BBBD08, DBB10, AMM+08]. As per the user studies conducted by Diehl et al. [DBB10], users take more time when using radial representations [WAM01, LKL+04] but the layout was found to be more useful when visualizing two dimensions simultaneously.

3. Visualization Study

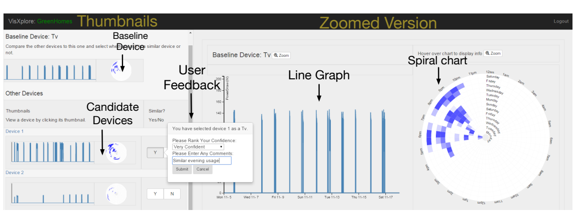

Figure 1: *Task 1 is a device classification task. The participant is shown the power consumption of a baseline labeled device (TV in the figure) and asked if the unlabeled devices are a TV.*

The goal of this preliminary study is to better understand energy usage insights. In our study, we present each participant with two visualizations of appliance energy usage and

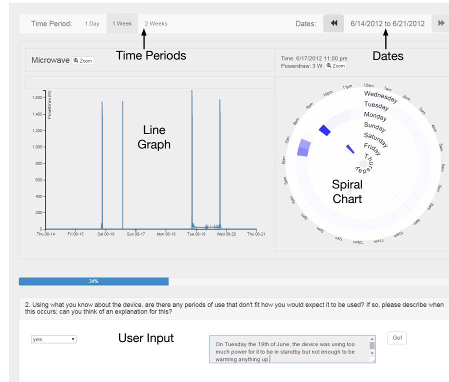

Figure 2: *Task 2 is a pattern detection task. The visualization is the energy consumption of a labeled device over a period of a day, week, and 2 weeks. The participant can also scan through different time periods using the Dates control. The participant then finds patterns in the time series and possible anomalies in the data.*

ask the participant to perform a series of tasks that involve device classification and pattern recognition.

Visualizations Used: Our study uses two types of visualizations illustrated in Figures 1 and 2: line graphs [AMM+08] and radial layouts [DBB10]. Both types are commonly used to visualize high and low dimensional data. A line graph preserves the raw energy consumption values for appliances and can show fine-grained consumption patterns. The radial layout is a two-dimensional representation of the data and can be used to infer patterns over longer time scales using a compact representation. The color shades of the slots in the radial layout represent different energy consumption values. The visualizations are implemented using D3 [BOH11] embedded in a web layout built using BootStrap.

Participant Tasks: The study asks participants to perform two tasks:

Task 1 (Device Classification): The goal of this task is to gain insight into the characteristics or features that participants consider most obvious or useful in our visualizations. Effective visualization-based device classification is useful in a direct sense insofar as manually labeling sensors in a real setup is prone to errors and changes, e.g. users moving devices. However, the possibility of exposing features that users rely on to differentiate devices in the visualizations could prove just as valuable. These features may be used to inform the design of visualization-based tools to help users better understand energy usage, which previous work has found can impact energy usage behavior [CRJ12]. These features may also be used as input to machine learning algorithms useful for tasks such as anomaly detection or disaggregation [KJ12].

Figure 1 illustrates the device classification task. We provide participants with a visualization of two weeks of data from a *baseline* device and a label describing that device, for example a TV. We then ask participants to consider visualizations of several other unlabeled devices and determine whether the unlabeled devices are in the same category as the baseline. Participants are asked to indicate the confidence of the answer provided and also share any comments to explain the reasoning behind the choice. In our study we use four baseline appliances—a TV, PC, refrigerator, and modem—and ask the participants to classify between five and seven unlabeled devices per baseline device.

Task 2 (Pattern detection): In the second task (Figure 2), we ask participants to determine whether there is a regular pattern of use that may be derived from a visualization of the device's energy consumption. Understanding what is *regular* usage can help determine anomalous usage of an appliance, for example periods of time when a device has been unintentionally left on or even devices that are malfunctioning and consuming more energy than appropriate. Determining anomalous energy usage is the basis of several energy recommendation systems [WRK10] and virtual energy audit systems [BGMS10]. A challenge for such systems, however, is defining anomalous behavior.

In this task, we provide participants with a visualization of device energy usage and ask them to select whether the device shows an hourly, daily, weekly, biweekly, or no pattern of use (see Figure 2). The visualization allows the participant to select the time period visible (one day, one week, or two weeks) as well as scroll through several weeks of data. We used four devices, a lamp, clothes washer, laptops, and a microwave, for this task. For each device, we use between four and 16 weeks of data for the visualization. As in Task 1, we ask the participant to indicate his/her confidence in the response, and we also ask the participant to indicate whether any periods of time do not fit the pattern. Finally, we ask the participant to select which of the visualizations were used for this task and rate his/her level of frustration with the task.

Study Setup: The visualizations were built using a dataset collected from a deployment of Z-Wave smart plugs in nine homes in the United States. Our dataset includes energy consumption data from five to 10 appliances per home collected over a period of two years. For this study, we used between two weeks to four months of data from a subset of appliances.

Study participants were recruited from our college campus and surrounding area using flyers and Craigslist advertisements. The study was performed in our laboratory and the participants were compensated $10/hour for their participation. Before the study began, we showed each participant a 5-minute video describing the system and tasks. We also collected demographic information, i.e., age and gender. Each session lasted for approximately 50 minutes.

We recruited 21 participants, 13 males and 8 females. The age varied from 18 to 31. It is also worth noting that the participants for the study were different from the residents of the nine homes that we collected data from for our visualizations. Answers and comments provided by users were stored in a database and parsed later to derive the insights described below.

4. Insights and Findings

In this section, we discuss the salient insights from our study. Of the 26 appliances that the participants were asked to classify, 16 were classified correctly by more than 75% of the participants. For example, 95% of participants identified one unlabeled PC correctly, and 90% were able to identify two of the unlabeled refrigerators. Devices that participants struggled with included composite devices, such as a power strip with a TV, Roku, PlayStation, DVD and a lamp that 38% of participants classified as different from a TV. For Task 2, participants were able to infer daily, weekly, or biweekly usage in appliances including laptops, a lamp, and a microwave, although a majority of participants did not find a pattern for a clothes washer, which was more sporadically used.

Insight 1: With minimal training and context and a small data set, participants were able to identify relevant features for device classification.

Participants noted several common features leading to successful device classification and pattern detection including (1) level of power consumption; (2) frequency and duration of appliance usage; (3) time and day of use; and (4) power states of a device (e.g., the device is always ON). For Task 1, for instance, 72% of parseable participant comments tied correct classifications to at least one of these four features. While participants frequently cited power levels in accurate classifications of TV and PCs, comments on refrigerator and modems tended to reference frequency and power states features, respectively. Often a decision hinged on which features a participant deemed most important given the data, hence participants could not only determine a *set* of features that might be relevant to a specific appliance but also *prioritize* features chosen based on the data. For example, four participants explicitly prioritized the time of use feature over difference in power levels to correctly classify one TV. One such participant noted:

> *The times for when it is on is similar to a TV and the powerdraw is fairly similar although device 1* (baseline device) *has a higher amount* (power consumption).

In other examples, feature selection was influenced by a participant's preconceptions about a baseline device. When comparing unlabeled devices to a refrigerator, more participants tended to consider whether an appliance was *always on* to be an important feature. For example, six participants who rejected one refrigerator that appeared to have a zero power state cited gaps, an off state, or the lack of a base power level. Overall, only nine participants correctly identified that particular device. Leveraging such prior knowledge also proved inaccurate in instances where the participant's assumptions or experience conflicted with the actual usage of the device. For example, one participant incorrectly used power states to identify a PC and noted:

This device is being used constantly, so I would think it was something other than a PC.

However, two participants who correctly identified the same PC noted:

(The device) *could be a computer that was left on most of the time*

The baseline is PC. People can have different preference in using the PC.

We note here that the participants of our study were not the original participants of the appliances that we collected data from, leading to the inaccuracy. But using a priori context would prove advantageous when a participant is analyzing his or her own device usage.

Insight 2: Participants were able to describe anomalous behavior and patterns of use using a vocabulary that can be generalized to describe such behavior in a broad range of devices.

In Task 2, our participants were able to broadly identify two categories of anomalies: (1) Anomalies derived directly based on patterns participants saw in the data; and (2) Anomalies where appliance usage differed from a priori perceived usage of the appliance. Both types of anomalies were based on features such as sudden changes; gaps or spikes in usage; atypical usage for a device type; unusual times or durations of use; and power levels. For example, a participant who identified a clothes washer as having a weekly pattern pointed out:

4/30/12 to 5/7/12 the washer was used multiples times throughout the week which deviates from the previous patterns of use.

Other pattern-based anomalies included gaps in regular usage or sporadically high power levels or anomalous time of usage. For example, a participant pointed out about the microwave:

graph shows overnight activity of microwave. door could have been left open, so the light remained on.

Participants often leveraged personal experience or device-specific assumptions about normal behavior to determine if usage was anomalous. For instance, a number of participants expressed an unwillingness to label any laptop usage as anomalous based on time of use, asserting:

Different people may prefer different time to use the laptop. A laptop is used almost all hours of the day now because of how important technology is so usage at anytime would be fine

More than 80% of the participants were able to find at least one instance of anomalous usage in the four appliances in Task 2. These anomalies can help build a participant-specific *vocabulary* and *signature* for anomalies that can be used as input to virtual energy audit systems and outlier detection algorithms.

Insight 3: The best visualization is task- and user-dependent. Having multiple visualizations for the same dataset is useful, and some visualizations require more training.

In our study, we found that the most popular visualization depends on both the task and participant. While over 72% of responses in the pattern recognition task favored the radial layout or a combination of the radial layout and line chart, the comments in the device classification task indicate high usage of the line chart. This could suggest a specific visualization was more suited for each task. The radial chart/circle visualization seemed to require more training to be useful. For instance, one participant commented:

As the study progressed I quickly used the circle graphs as my go to graphs to analyze the devices, in the beginning I used mostly the line graphs because they were familiar but after a couple tasks the circle graphs quickly became the easiest and fastest way to analyze what was going on

Another participant who preferred the line chart complained that the circles were *tiresome and untranslatable* while the lines were *easier to read...easier to understand*. Participants liked how clearly the circles showed time of day and intensity and how little space it took. Participants struggled with the unequal representation of time between the center and outer rings in the radial layout, with the transition between rings of the spiral, and close shades or faint representations of very low power level device states. Participants liked the level of detail and ease of use of the line graphs while at least one participant thought the level of fluctuation in the line graph made it harder to find patterns.

Design Recommendations The results of our study provide insight into the design of home energy management systems. Our findings suggest that, in the case of visualization-based systems, highlighting a common set of features as well as providing multiple views of the same data is recommended. Our findings could also motivate future work regarding automated, machine learning-based systems, which have yielded mixed results in some of our prior work [LBR14]. Adapting these systems to a human-in-the loop approach that considers user input regarding features and characteristics of anomalous behavior may prove effective.

5. Conclusion

In this work, we present the results of a visualization study designed to expose the characteristics and insights participants are able to identify in home energy consumption data. Our findings indicate that with minimal training, a small data set, and minimal contextual information, participants are able to extract relevant features of the data as well as describe anomalous behavior and general patterns of use. We also find that using multiple visualizations is effective as visualization preference is both user- and task-dependent. Our findings may be used to inform the design of both automated and user-driven energy management systems.

References

[AMM+08] Wolfgang Aigner, Silvia Miksch, Wolfgang Muller, Heidrun Schumann, and Christian Tominski. Visual methods for analyzing time-oriented data. *Visualization and Computer Graphics, IEEE Transactions on*, 14(1):47–60, 2008. 2

[BBBD08] Michael Burch, Felix Bott, Fabian Beck, and Stephan Diehl. Cartesian vs. radial–a comparative evaluation of two visualization tools. In *Advances in Visual Computing*, pages 151–160. Springer, 2008. 2

[BGMS10] Mario E Berges, Ethan Goldman, H Scott Matthews, and Lucio Soibelman. Enhancing electricity audits in residential buildings with nonintrusive load monitoring. *Journal of industrial ecology*, 14(5):844–858, 2010. 2, 3

[BOH11] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D^3 data-driven documents. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2301–2309, 2011. 2

[BRM10] Lyn Bartram, Johnny Rodgers, and Kevin Muise. Chasing the negawatt: visualization for sustainable living. *Computer Graphics and Applications, IEEE*, 30(3):8–14, 2010. 2

[CRJ12] Enrico Costanza, Sarvapali D Ramchurn, and Nicholas R Jennings. Understanding domestic energy consumption through interactive visualisation: a field study. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 216–225. ACM, 2012. 2

[DBB10] Stephan Diehl, Fabian Beck, and Michael Burch. Uncovering strengths and weaknesses of radial visualizations—an empirical approach. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):935–942, 2010. 2

[DWD11] Soma Shekara Sreenadh Reddy Depuru, Lingfeng Wang, and Vijay Devabhaktuni. Smart meters for power grid: Challenges, issues, advantages and status. *Renewable and Sustainable Energy Reviews*, 15(6):2736 – 2742, 2011. 1

[GCZ+11] Luciano Gamberini, Nicola Corradi, Luca Zamboni, Michela Perotti, Camilla Cadenazzi, Stefano Mandressi, Giulio Jacucci, Giovanni Tusa, Anna Spagnolli, Christoffer Björkskog, et al. Saving is fun: designing a persuasive game for power conservation. In *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology*, page 16. ACM, 2011. 2

[GD12] Sarah Goodwin and Jason Dykes. Visualising variations in household energy consumption. In *Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on*, pages 217–218. IEEE, 2012. 2

[GDJ+13] Sarah Goodwin, Jason Dykes, Sara Jones, Iain Dillingham, Graham Dove, Alison Duffy, Alexander Kachkaev, Aidan Slingsby, and Jo Wood. Creative user-centered visualization design for energy analysts and modelers. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2516–2525, 2013. 2

[GG05] Anton Gustafsson and Magnus Gyllenswärd. The power-aware cord: energy awareness through ambient information display. In *CHI'05 extended abstracts on Human factors in computing systems*, pages 1423–1426. ACM, 2005. 2

[JSMK14] Halldór Janetzko, Florian Stoffel, Sebastian Mittelstädt, and Daniel A Keim. Anomaly detection for visual analytics of power consumption data. *Computers & Graphics*, 38:27–37, 2014. 2

[KJ12] J Zico Kolter and Tommi Jaakkola. Approximate inference in additive factorial hmms with application to energy disaggregation. In *International conference on artificial intelligence and statistics*, pages 1472–1482, 2012. 2

[LBR14] D. Lachut, N. Banerjee, and S. Rollins. Predictability of energy use in homes. In *Green Computing Conference (IGCC), 2014 International*, pages 1–10, Nov 2014. 4

[LKL+04] Jessica Lin, Eamonn Keogh, Stefano Lonardi, Jeffrey P Lankford, and Donna M Nystrom. Visually mining and monitoring massive time series. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 460–469. ACM, 2004. 2

[MAR10] Paul Monigatti, Mark Apperley, and Bill Rogers. Power and energy visualization for the micro-management of household electricity consumption. In *Proceedings of the International Conference on Advanced Visual Interfaces*, pages 325–328. ACM, 2010. 2

[MEB+13] Masood Masoodian, Birgit Endrass, Rene Buhling, Pavel Ermolin, and Elisabeth André. Time-pie visualization: Providing contextual information for energy consumption data. In *Information Visualisation (IV), 2013 17th International Conference*, pages 102–107. IEEE, 2013. 2

[POB08] James Pierce, William Odom, and Eli Blevis. Energy aware dwelling: a critical survey of interaction design for eco-visualizations. In *Proceedings of the 20th Australasian Conference on Computer-Human Interaction: Designing for Habitus and Habitat*, pages 1–8. ACM, 2008. 2

[RB11] Johnny Rodgers and Lyn Bartram. Exploring ambient and artistic visualization for residential energy use feedback. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2489–2497, 2011. 2

[SCG+11] Anna Spagnolli, Nicola Corradi, Luciano Gamberini, Eve Hoggan, Giulio Jacucci, Cecilia Katzeff, Loove Broms, and Li Jönsson. Eco-feedback on the go: Motivating energy awareness. *Computer*, 44(5):38–45, 2011. 2

[WAM01] Marc Weber, Marc Alexa, and Wolfgang Müller. Visualizing time-series on spirals. In *Information Visualization, IEEE Symposium on*, pages 7–7. IEEE Computer Society, 2001. 2

[WRK10] Hendro Wicaksono, Sven Rogalski, and Enrico Kusnady. Knowledge-based intelligent energy management using building automation system. In *IPEC, 2010 Conference Proceedings*, pages 1140–1145. IEEE, 2010. 3