Understanding Domestic Energy Consumption through Interactive Visualisation: a Field Study

Enrico Costanza, Sarvapali D. Ramchurn, Nicholas R. Jennings
Agents, Interaction and Complexity Group
School of Electronics and Computer Science
University of Southampton, UK
{ec,sdr,nrj}@ecs.soton.ac.uk

ABSTRACT

Motivated by the need to better manage energy demand in the home, in this paper we advocate the integration into Ubicomp systems of interactive energy consumption visualisations, that allow users to engage with and understand their consumption data, relating it to concrete *activities* in their life. To this end, we present the design, implementation, and evaluation of *FigureEnergy*, a novel interactive visualisation that allows users to *annotate* and *manipulate* a graphical representation of their own electricity consumption data, and therefore make sense of their past energy usage and understand when, how, and to what end, some amount of energy was used. To validate our design, we deployed FigureEnergy "in the wild" – 12 participants installed meters in their homes and used the system for a period of two weeks. The results suggest that the annotation approach is successful overall: by engaging with the data users started to relate energy consumption to *activities* rather than just to *appliances*. Moreover, they were able to discover that some appliances consume more than they expected, despite having had prior experience of using other electricity displays.

Author Keywords

Interactive visualisation, Energy management, Eco-feedback, Sustainable HCI, Consumption

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces

General Terms

Design, Human Factors

INTRODUCTION

Reducing energy consumption is recognised as one of the current major global challenges. Most of the energy we currently consume is generated from fossil fuels, which have adverse impact of the environment. Renewable Energy depends typically on weather and therefore it is not always available in plentiful supply. Even with more efficient technology it is expected that consumption habits will need to change. So there is a growing need to develop tools and techniques to empower consumers to take control of their electricity demand in order to optimise their energy usage in scenarios where energy supply is limited or variable given the intermittency of renewable sources. Moreover, reducing and controlling domestic energy consumption is important, as it accounts for a significant portion of the overall energy demand (39% in the US[1] and 28.8% in the EU [4]).

So far, commercial solutions to improve the management of energy demand have centred on the deployment of smart meters and in-home energy displays (e.g., AlertMe,[2] Current Cost,[3] The OWL[4]) that can provide whole-house real-time energy monitoring and dynamic pricing from suppliers in an attempt to motivate users to shift or reduce their energy consumption [33]. While such technology can certainly provide the *infrastructure* to deliver consumption information to home consumers, recent studies of energy displays identify that one of the main problems of current designs is the lack of context around the information provided and the consequent difficulties that users encounter in making sense of this information [15, 16, 35]. Hence, there has been considerable interest over the last few years in Ubicomp and HCI in the design of high resolution sensors [20, 30] and effective displays for energy consumption, often referred to as *eco-feedback technology* [19, 35].

In this context, our primary interest is in helping users understand their own energy consumption: since such understanding is a necessary step on the way to reduction. Hence we believe it should be a central concern of the design of eco-feedback technology. Without correct understanding, users may misplace their effort and resources, whether that is, for example, in terms of curtailment of activities or replacement of appliances. Therefore here we assume that users are independently motivated to conserve energy, be that for intrinsic (e.g. moral) or extrinsic (e.g. economic) reasons [1, 7, 11,

UbiComp '12, Sep 5-Sep 8, 2012, Pittsburgh, USA.

[1] http://www.eia.gov/energyexplained/index.cfm?page=electricity_use
[2] http://www.alertme.co.uk.
[3] http://www.currentcost.com.
[4] http://www.theowl.com.

19]. Extrinsic motivation could be, for instance, a result of fiscal policies (e.g. a carbon tax) or information campaigns. This emphasis on *understanding* rather than motivation is complementary to some prior work related to energy conservation, which has often been framed in the context of Persuasive Technology [17] – i.e. technology that motivates users to change behaviour.

Against this background, we advocate the design of *interactive* energy consumption visualisations, that allow users to engage with and understand their consumption data, relating it to concrete activities in their life. This approach is informed by the constructionist learning theory [29], the learning theory behind the Logo programming language. We believe that this approach can become foundational to effective behaviour change [17]. To substantiate this proposition, in this paper we report the design and evaluation of *FigureEnergy*, a novel interactive visualisation that allows users to *annotate* and *manipulate* a graphical representation of their own electricity consumption data, and therefore make sense of their past energy usage and understand when, how, and to what end, some amount of energy was used. The project also demonstrates the success of a medium-fidelity approach to prototyping Ubicomp systems: combining off-the-shelf networked sensors with Web technologies makes it possible to easily and economically deploy at a medium-to-large scale.

To validate our design, we deployed FigureEnergy "in the wild." As part of a field study, 12 participants installed meters in their homes and used the system for a period of two weeks. Their experience was documented through a combination of automatic recording of interaction logs and semi-structured interviews. The results suggest that the annotation approach is successful overall: by engaging with the data users started to relate energy consumption to *activities* rather than just to *appliances*. Moreover, they were able to discover that some appliances consume more than they expected, despite having had prior experience of using other electricity displays. The rest of this paper is structured as follows. In the next section we present related work. Building upon this, we present the motivation for this work and go on to describe the FigureEnergy system and its individual components. We then describe the field trial and discuss the key implications of its results. Finally, we conclude and outline future work.

RELATED WORK

The mechanisms that may support energy and, in general, resource conservation in the home setting have been the object of research in social and environmental psychology since the 1970s [1]. Reviews of this literature identify the provision of feedback about energy usage as one effective strategy for conservation [1, 10, 15]. Other reported strategies include the provision of information about energy conservation, the setting of saving goals, and the reward of savings in monetary terms (sometimes goal setting and rewards are used in combination and through the support of feedback) [1]. The display of feedback information, especially through digital and interactive displays is of particular relevance to Ubicomp & HCI as witnessed by a number of recent projects in this area. The reader will find an insightful review of these approaches in [12, 19].

Many academic and commercial projects around the visualisation of electricity consumption tend to place the main emphasis on instantaneous consumption data, often through evocative ambient displays aiming at emotional reactions [21, 22, 25, 34], or mimic stereotypical instrumentation displays, for example through the use of dials (e.g. [18]). When historical data is presented, this is often in aggregate format, for example by showing a single value for the total consumption over past days or past weeks. Notable exceptions are the work of Carlis & Konstan on visualization of serial periodic data [6] and the Energy AWARE Clock project [5]. Carlis & Konstan highlight the importance of human time accounts on the interpretation of data by using calendar days as units of analysis and a 12-months calendar as a visualisation layout. Similarly, the Energy AWARE Clock overlays the plot of energy consumption in the past 24 hours on the circular display of an analogue clock. As highlighted in [19], only a minority of proposed eco-feedback systems are evaluated, and an even smaller portion are placed in the field.

To provide more detailed feedback, the breakdown of energy consumption from different appliances can be achieved by individually monitoring each device – this strategy can, however, involve expensive hardware that may be difficult to install. Automatic methods for the disaggregation of appliance loads from whole-house metering, so called Non-Intrusive Load Monitoring (NILM), involve the use of machine learning and optimisation techniques to recognise energy signatures and is an active area of research which is, as yet, unsolved [37]. The challenge in NILM is that individual appliances have very different energy signatures that are hard to spot unless very high resolution meters are available. Moreover, even if some appliances can be accurately recognised and classified given an existing training set using machine learning techniques, such solutions perform poorly when subjected to new instances under the same category [3, 26]. Better results have been reported through high-frequency electricity sensing systems [30, 20], even though these are prohibitively expensive for the consumer market. Finally, while NILM aims at improving energy auditing methods, it is less concerned with how the information obtained is fed back to the user to improve learning. FigureEnergy is more concerned with the latter, and uses manual annotation rather than automatic disaggregation.

On related note, *tagging*, particularly as a collaborative activity, is a popular form of annotation through keywords. It has been studied in the context of online communities and user-generated content, mostly because of its application to recall and retrieval [2]. Mamykina et al. [28] investigated the use of tagging to support users assess the nutritional value of photographs of food, reporting positive results under the condition that the tagging vocabulary is pre-defined by experts. The type of annotation employed in FigureEnergy goes beyond tagging, in that users are required to *select* what to tag from a stream of information (while tagging is normally applied to pre-defined units of analysis, such as photographs or URLs). Moreover, in FigureEnergy annotations are not limited to keywords and can potentially be verbose.

In the context of information visualization, interactive visual annotations have been introduced to support synchronous and asynchronous collaboration, often for deixis [14, 24]. In these tools annotations sit on a layer separate from the data, not integrated in its processing and representation. In contrast, in our work annotations allow users to define partitions and new representations of the original data.

DESIGN RATIONALE

Recent studies of energy displays identify one of the main problems of current designs as the lack of context around the information provided and the consequent difficulties that users encounter in making sense of this information [15, 16, 23, 31, 35]. For example, information is presented numerically, but consumers are not familiar with the units of representation, and are not offered any frame of reference for comparison: "It says you've used so many kilowattevers" as a participant in [35] poignantly summarises.

In particular, Fischer [15] highlights the importance of providing a direct link between actions and consequent consumption, through the provision of feedback that is broken down for specific rooms, appliances, or times of the day. Moreover, research has highlighted how different visualisations can significantly impact consumers choices. For example, Egan et al. [13] point to the difficulties that consumers encounter in interpreting visual representations of energy consumption comparisons and the trade-offs between accuracy and clarity. In the context of visualising car consumption efficiency, Larric and Soll [27] demonstrate how inappropriate information representation can lead to systematic misjudgement, as well as the opportunities for amelioration: simply switching from one efficiency metric (miles per gallon) to another (gallons per mile) can considerably help consumers draw correct conclusions.

Against this background, we started from the observation that on the most common representation of historical energy consumption, the time-series plot, it is easy to *qualitatively* identify consumption *events* (such as the use of a washing machine, oven, or shower), that is, intervals of time where the consumption increases. On the graph, these correspond to peaks or lumps of different height and width. However, it is generally difficult to compare the *amount of energy* in different events, because the latter corresponds to the integral of the curve for the duration of the event. Therefore, we decided to use the time-series as a starting point, and design an interactive tool that allows users to select time intervals and annotate them with icons and text to isolate individual activities or events. Rather than looking into automatic solutions, we envisaged that allowing users to manually *annotate their energy consumption logs* could bring several benefits. First, we anticipated that engaging in this explicit annotation would encourage users to reflect on their daily activities and on their impact on the overall use of energy. Second, we anticipated that manual annotation could also result in richer descriptions than what could be obtained automatically. For example, while a sensor on the washing machine could easily detect "high temperature washing", users may label the same activity with more elaborate and meaningful information as "washing towels" resulting in a deeper understanding and a stronger connection between the data and the everyday activities of the users.

Furthermore, taking inspiration from the *change of representation* used in scientific and technical disciplines,[5] we switched from a time-centric representation to an energy-centric one. In this vein, we designed a visualisation, to complement the time-series plot, where events are rectangular boxes, with area proportional to the amount of energy the event consumed. The box metaphor was chosen to evoke a *materialisation* of energy consumption: a representation that helps users to relate to this intangible information in a more graspable way, as advocated, for example, in [32]. Even though the comparison of "areas" is known not to facilitate precise results [8], it still caters for approximative ones.[6] The event-boxes can be directly manipulated. The expectation – informed by the constructionist learning theory [29] – is that if users can directly manipulate the data, they would more easily reflect and understand their own consumption. This use of direct manipulation aims also at being playful, to resonate also with some of the findings by Woodruff et al. [36]: reporting a study of individuals who are particularly active in taking environmentally responsible choices, they describe that participants continuously evaluated their behaviour, and that "the decision-making process often took on a game-like or playful nature."

SYSTEM DESIGN

In this section we describe our approach to the design of FigureEnergy, an interactive electricity consumption visualisation tool. We first describe the sensing and flow of data and the measurements used in the visualisations, and then the key visualisation components.

Sensing Electricity Consumption

FigureEnergy relies on off-the-shelf digital networked electricity meters, such as AlertMe or Current Cost. These meters include a current clamp (split ferrite ring) to one of the electricity mains cables to measure the flow of electric alternating current (AC), and an RF (ZigBee) transmitter that sends measurements to a hub which, in turn, transmits this data to an Internet server (run by AlertMe or Current Cost) through a home broadband connection.[7] These meters provide information in terms of power (kW), as well as energy consumed (kWh), and real-time as well as historical data. Their reading periodicity generally falls in the range of seconds (for real-time view) to minutes (for historical data). The data is then retrieved from the AlertMe and Current Cost

[5]For example switching from a time-based to a frequency-based representation of signals through the Fourier transform, can reveal features (amplitudes of constituent frequencies of time-series) that are otherwise concealed.

[6]More precise comparison can be made by looking directly at the consumption values of events, available as text.

[7]Sometimes an in-home display is also included in the system, but this is not used by FigureEnergy, and it was not provided to participants in our study.

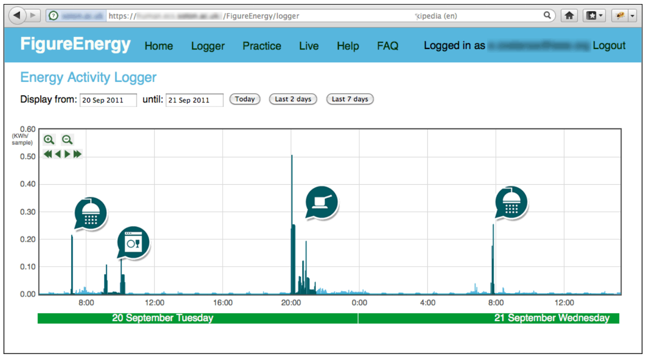

Figure 1. The FigureEnergy Logger – a time-series graph of energy consumption that can be annotated by users to help them make sense of their own data.

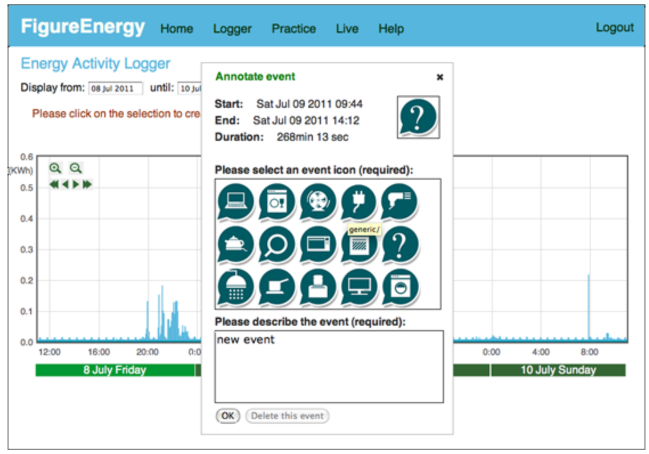

Figure 2. The user can choose one of the icons provided to associate with the annotation and provide a text description of the event.

servers to be processed and stored on our own server, and from there, presented to the users through a Web interface.

Electricity consumption related to appliances that are always plugged-in and always turned on, is referred to as *always-on consumption*. The always-on component may not be strictly constant: for example, fridges or immersion heaters operate in fairly regular cycles and can form a main part of the always-on component. We estimate the always-on consumption on a per-day basis from the raw consumption signal by calculating a smoothed (low-pass) version of the data and taking the minimum. The results were found to be comparable to the always-on value calculated by commercial meters, like AlertMe or Google PowerMeter.

In contrast to always-on consumption, we refer to the consumption related to appliances that are *explicitly* switched on and off, as *consumption events*, or simply *events*. A consumption event, can then be described in terms of start and end timestamps and the amount of energy consumed. Crucially, a consumption event can be generally associated with specific activities involving the usage of *one or more* electrical appliances. For example, a consumption event could involve running a washing machine to do laundry, or using the electric kettle and the electric hob to make dinner.

FigureEnergy also uses the concept of a *reference consumption* used for comparison with the current consumption. This value can be set to the average consumption over a fixed past period of time (like in our study – see details in the Evaluation Section), or other user-defined value, such as the user's general average, or the national or regional average of residential units of comparable size.

Interactive Visualisation

FigureEnergy is composed of two interactive views, the Logger and Practice views, and by a "Live view" designed to provide the sort of information most commonly available on other electricity meters.

Logger View

The Logger view displays a time-based plot of the average power usage in the home (Figure 1). Users can seamlessly zoom in (down to two-minute periods) and out (up to week periods) and pan backward and forward in time. We found that other displays tend not to allow such advanced functionality and provide very coarse representations for historical data (e.g., at a minimum of 15 min intervals in the case of AlertMe). The events on the graph can be annotated (Figure 2) by selecting an interval of time with the mouse and associating with it an icon and a textual description, to which we refer as the *event label*. As shown in Figure 2, a set of fifteen icons were included, their main purpose is to create a visual connection between the Logger and Practice views. These icons are not meant to cover an exhaustive range of event types, but only to give some general categories. When an icon is selected, the event label is pre-populated with a *keyword* associated to it (e.g., *dishwasher*, *kettle*, etc...); users can remove this keyword or add more text.

Practice View

The *Practice* view, shown in Figure 3, was designed to provide an *energy-centric*, or *event-centric*, representation of the same data shown in the Logger. Consumption events annotated in the Logger are represented as boxes of different sizes: the size (area) of each is proportional to the energy consumption, so if an event consumes twice as much energy as another, the former's representative box will be twice as big as the latter's. The event-boxes are marked out by the same icons used in the Logger, for immediate recognition – hovering on each box with the mouse reveals its details: textual description, amount of energy consumed, duration, date and time. The boxes are immersed in the *Practice Tub*, the main element on the page, which contains also a dark blue liquid, representing the always-on consumption, and the light blue liquid, representing the energy that was not annotated (but different from the always-on). The size of the tub represents the total amount of energy consumption over the period of time under analysis.[8]

The aim of the Practice view is to facilitate the comparison of events with each other, and to allow users to *practice* re-

[8] The visualisation refers to a specific time period, which can be set through standard date selectors. By default, the period is the same as that of the Logger view.

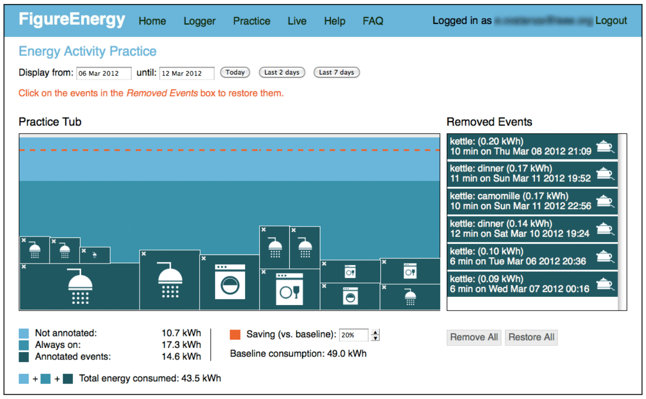

Figure 3. The Practice view allows users to play what-if scenarios by removing event boxes from the practice tub. This is equivalent to 'not-doing' an event and seeing the impact of this on the overall energy consumption.

ducing their usage through the curtailing of certain events (hence the name). When the page is first loaded, all the events sit at the bottom of the tub so that the level of the liquid fills the tub completely. Individual events can be taken out of the tub by clicking on them. By so doing, the equivalent liquid amount (i.e., energy consumption of the event) replaces the space previously occupied by the event, resulting in a decrease in the liquid level in the tub. In this way users can play *what-if scenarios* and reflect on the impact of individual events on the overall consumption. The removed events are moved to the Removed Events Box, on the right hand side of the page, and they can be restored by a simple click. The overall level of energy can be compared with the savings line, a horizontal dashed line which shows the level corresponding to an adjustable percentage of savings compared to the baseline consumption. For example, users can set the savings to 20% and then progressively remove events from the practice tub until they reach that level. A legend under the tub indicates the colour mapping and reports the amount of energy in each category as well as the total.

Live View

Finally, the Live view (see Figure 4) provides real-time information, comparable to what is available on most standard electricity meters (e.g., Current Cost or AlertMe). Information on this page includes the current power consumption of the home, in kW, the amount of energy consumed so far in the day and in the week, and the prediction of the consumption by the end of the week and the comparison to the reference consumption, as defined above. For the study, this page also displays information about the users' performance during the study. In this respect, the system calculates the savings users make against their reference consumption and displays the percentage of energy saved and the equivalent reward gained by participants for each week of the study.

Technical Implementation

The system was implemented using open source tools and open APIs. The backend was written in Python, building on the Django Web framework. The front-end uses HTML 5, in particular the canvas component, JavaScript with the jQuery, jQueryUI, and Flot libraries. This architecture was chosen to make the system more easily portable to mobile platforms, for future developments. Data is retrieved from AlertMe servers and stored locally in a MySQL database.

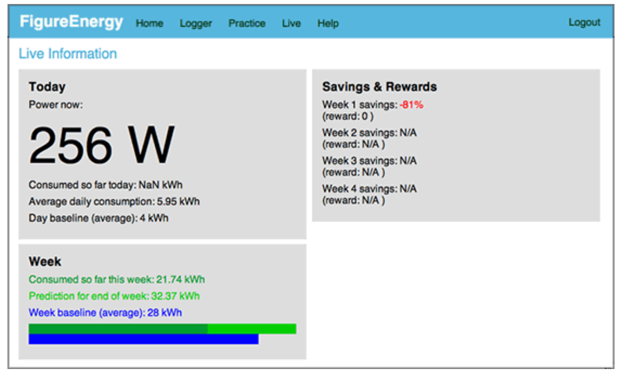

Figure 4. The Live page provides real-time measures of electricity consumption as well as estimates of consumption for the week and savings for individual weeks over which the study was run.

EVALUATION

Method

The connection between the data and the users' personal knowledge of their day-to-day activities plays a central role in FigureEnergy. Therefore a real world deployment, in the form of a *field trial*, was necessary to evaluate whether participants would be able, as we hoped, to leverage this connection. Our focus is on studying whether annotation and interactive visualisation can aid the understanding of consumption. However, for users to engage with their energy consumption logs at all, there is a need for some form of motivation. As discussed below, it turned out that many of our participants were already interested in their consumption, but to increase the likelihood of engagement for all participants for the purposes of this study, we provided a small financial incentive for energy savings. Along the lines of prior studies [1], participants were rewarded with £0.25 for each 1% of savings compared to a *reference consumption*, up to a maximum of 30% each week. They were also given extra bonus rewards of £2.5 and £5 per week if savings exceeded 10% and 20%, respectively. In other words, participants could gain up to £12.5 each week. The reference consumption was calculated from the bills from a previous period, as provided by participants.

Recruitment for the trial took place in our university cafeteria: a public space frequented by the entire university community, from academics and researchers to support staff and students. We setup a stand with posters, a monitor showing the prototype of the system and the electricity sensor used in the study, an AlertMe Energy Kit. We attended the stand for about one and a half hour at lunch time approximately every day over two weeks. As this kind of electricity meters are marketed for end-users, and can be installed without the intervention of an electrician, as part of the trial, participants were asked to install the system on their own, following the simple yet thorough instructions provided with it.

Participants

A total of 15 participants volunteered to take part in the study, all members of the university, 6 PhD students, 6 researchers and 3 members of support staff, from a variety of disciplines (including Education, Engineering, Health Sciences, Physics). Three participants (1 researcher and 2 members of support staff) were not able to start the study: two because of problems with the sensor installation and one because of compatibility issues with the web browser, as the participant was not allowed by his employer to install a different one. The rest of the analysis will focus on the 12 participants who did start.

All except one participant lived in shared accommodation: 8 couples (1 with a baby, 1 with 3 teenagers), 3 in shared accommodation, one subletting to one tenant. The age of participants ranged from 22 to 60, with the majority being in the late 20s and 30s.

At the beginning of the trial participants were asked to fill an online questionnaire addressing their attitude and familiarity with energy conservation. Answering it, all except two participants reported to be saving energy beyond the scope of the trial. The motivation for this was reported to be cost reduction by four participants, environmental concern by one and both factors by five. All except two participants (not the same ones as above) associated their own energy use with climate change or the environment. Eight participants reported to follow pro-environmental behaviours to try and reduce their carbon emissions, beyond energy consumption (e.g., limiting travel by car or plane), yet not in any radical way. The prevalence of research students and researchers in our self-selected sample indicates that energy monitors appeal to a highly educated population.

Data collection & analysis

Data was collected through a mix of techniques. Interaction logs were automatically recorded on the FigureEnergy server, documenting every time users accessed a page and created an annotation, as well as the content and details of the annotations. Semi-structured interviews were carried out with 8 participants (the others were not available), after they had used the system for 10 to 19 days. Where applicable we asked participants to compare their experience with FigureEnergy with other energy displays (4 participants). Each interview lasted approximately half an hour. Interviews were audio-recorded and later coded through open codes, then grouped in categories.

Results

Events annotation

Over the duration of the study, participants created an average of 2.56 annotations per day, for a total of 456 events. Eight participants engaged with the system actively, annotating more than one event per day. A total of 214 events labels (46% of the total) consisted exclusively of the keywords associated with the icons. The content of the remaining 242 labels was coded in the categories listed in Table 1. Note that these categories are not mutually exclusive (e.g., “kettle: luis’s tea” was coded as both *person* and *activity*).

The *other names* and *activity* categories, two of the top three in frequency, are undeniably related to specific limitations of the system: situations not covered by the icons provided, or covered by more than one icon, were described through the textual label. The prevalence of the *activity* category indicates that the annotation went beyond the mere disaggregation of appliance loads, and dealt with higher levels of abstraction such as “tom boiling water for tea” or “oven roast dinner”. This matches our expectation that, if given the opportunity to do so, users will reflect on and thus provide richer descriptions of their usage of a device.

The labels coded as “analysis” show that participants use the annotation system to record their reflections about consumption patterns, indeed the interviews (see below) revealed that FigureEnergy was sometimes used as a starting point to perform manual calculations. Finally, the “person” category is specific to users who live in shared accommodation: flatmate names were included to discriminate different events.

The distribution of event durations reveals a variety of annotation styles. Event durations ranged from a minimum of 1 minute, to a maximum of more than 20 hours. Average event duration also varied widely across active users: from about 20 minutes to almost 2 hours. Participants annotate events at most 24 hours later, while in some cases, as soon as a few minutes after the event took place.

From the interviews we learned that the choice of which events to annotate seems vary from user to user. While some annotated every “spike,” others “only logged the main things.” Factors that influenced the decision ranged from how easy the events were to remember and recognise, to how regular the corresponding activities are perceived to be, to how much energy activities are expected to consume (e.g. “the big ones”).

The annotation of events was a prominent activity within the usage of FigureEnergy, the interaction logs revealed that 58% of the page views were on the logger view, 33% on the live view and only 9% on the practice view. The limited popularity of the practice view is not surprising as the page becomes useful only after a number of events are annotated. Moreover, as it emerged from the interviews and we discuss below, participants did not use the interactive features of the practice view as much as we expected.

Understanding and Discoveries

In the interviews participants confirmed our expectation: the annotation activity and the change of representation do help in making sense of one’s own energy consumption.

Frank, a 33 year old Engineering PhD student living with his girlfriend, who had used an OWL electricity display before, nicely summarised the value of the box visualisation. Speaking about the practice view he pointed out that from

Category	*Meaning*	*Example*	*Frequency*
other names	appliance not included in the icons	“coffee filter”,“steriliser”	123
activity	use of device for an activity	“oven baking a cake”, “kettle: tea”	79
multiple	multiple activities or appliances	“tv and computer”, “oven + stove + tv”	62
analysis	comment about the energy	“fridge consumes 0.8 kWh per day”, “computer with additional peaks from...”	33
person	a person is mentioned	“shower: tom”,“kettle: luis”	27
qualifier	qualification of a keyword, without detailing the usage	“bathroom fan”,“oven: grill”	14
uncertain	the label includes a question mark or other expression of doubt	“shower?”,“fridge/freezer? will try to notice times when they come on!”	14
timing	the label includes additional information about the event time	“morning shower” or “kettle (plus microwave 09:47)”	6

Table 1. Categories of annotations, an explanation and example of each, and their frequency across the events’ labels.

the logger it is not always possible to understand whether “a big spike for a short amount of time” consumes more or less than “a low spike for a long amount of time.”

Maria, a Physics PhD student, 24 year old living with her boyfriend, not having used other energy displays before, told us that because she likes seeing things numerically, she has a preference for the Live page, because of it’s real-time and quantitative in nature:

I find the Live page probably the most useful bit, because I can just have it on while I am doing something and see how much power it takes, and I get an actual number, which means more to me than looking at the picture.

Yet, she commented:

The TV doesn’t use very much electricity, but actually in the practice box it’s a big area because it’s on for quite a long time. So, for example, the practice area is good for that.

Maria found the practice box graphical representation useful for specific events that may otherwise be difficult to notice.

Participants often reported discovering something new about their energy consumption, even to the point of being surprised. For example, John, a 32 year old Engineering researcher living with his wife and child, who had used a different electricity display before, declared:

The biggest surprise was TV watching if I was using the little cable box, which is a fairly unassuming little device, but when watching.. you can see that the TV uses a block ..I mean.. this big [indicating the size of the box with his hands on the table] but when you have the cable thing running and you are watching some on-demand thing the box for the same period is twice as big it’s just that the energy use skyrockets.. that was a little bit concerning!

The quote demonstrates two of the goals of the box visualisation. First, the visualisation allows to compare events and notice differences. Second, it can also encourage to think about energy in a *material* way, as suggested by the gesturing, especially in contrast to how *unassuming* the energy-consuming device can be.

Participants also commented on the role of the time-series visualisation in understanding how energy consumption unfolds in time, especially for multi-state appliances such as the washing machine. Other energy monitors experienced previously by our participants did not provide this visualisation. Luis, a 31 year old Engineering PhD student sharing with his girlfriend, described:

[With the logger view] you can see events. If at night you have something happening you can even see what the progress of your washing machine program cost [is].

Elaborating on the differences between FigureEnergy and a different electricity display, Frank brought up another multi-state appliance, the dishwasher. He reported another discovery:

I would have said that [the dishwasher] uses a fair bit of energy, but not huge.. because it does use a lot of energy and the fact that there is a massive spike at the beginning, as it heats the water up, and then it comes down a bit and then there is a massive spike at the end as it dries air or.. and also because the thing with the OWL [electricity display] is that you only look at the instantaneous one.. and you’re allowed to see if [an appliance] is own or off.. ..while this one, Figure Energy, gives you the overall usage so the fact is that the dishwasher uses a heck of a lot *of energy, the oven uses* tons *of energy ..the hob [stove] uses a fair bit of energy..*

Multi-state appliances are particularly hard to understand based only on instantaneous information. In contrast, the historical information presented by FigureEnergy makes it easier to understand their operation in time. The ability to integrate the event consumption over time allows for comparison with other events.

Annotations and calculations

The annotation activity in itself can raise awareness, as Luis made us notice. He had guests visiting over the week-end and noticed that they used the hair-drier a lot, he told us that he noticed it because he annotated the corresponding events. He also explained to us how he used the Logger as a measuring tool within a more complex calculation about his energy consumption:

> *I wanted the aggregate of the fridge, and I had to estimate it, I did it by hand, .. I took one small portion, then noted how many periods there are over in day, calculated the period of the fridge, ..how many hours in a day.. this is how much it burns.. so this is the fridge consumption..*

Other participants suggested us to include analysis tools in FigureEnergy, for example to show the aggregate consumption associated with each keyword, or ways to interactively define periodic events. On one hand these comments encourage the development of more automation, on the other they demonstrate engagement with the data.

Instantaneous information

Similar to what reported in other studies [16], our participants described using the instantaneous information provided by energy monitors for playful exploration. For example Frank told us about when he and his girlfriend got their OWL energy meter (before taking part in our trial):

> *We'd literally go around with pen and paper and the thing [the OWL display] and then turn off the light, see how much it drops by and then go 'ok' ..*

The FigureEnergy Live page was used in a similar fashion by Maria, who described:

> *I was running around my house and turning different things on and off and seeing how much that changed. Turning my lights on made more difference than I thought.*

The comparison of the Live page with other energy displays also revealed some of the limitations of FigureEnergy, and in particular of its Web-based nature. All four participants we interviewed who had experience with an hand-held energy display liked the ability to easily take it around their home. Response time was another limitation reported in relation to the Live page. Since in FigureEnergy, information from the electricity sensor needs to travel to two different Internet servers and back before being sent to the client browser it takes much longer than other energy displays that communicate directly with the sensor, via a local wireless network. Our participants noticed it. For example Luis mentioned:

> *The thing that I like about CurrentCost is that it's immediate: you can switch on the device and it immediately tells you how much that's consuming.*

At the same time, access to the live consumption information through the Internet was valued by some participants, and indeed the live view was the second most popular page, with 33% of the page loads. By remotely accessing the energy logs, they tried to guess what activities were going on in their homes while they were not there.

Consumption Comparison

When comparing participants' electricity consumption over the two weeks of the trial to the reference values, the results varied widely. Five participants saved between 5% and 32%, seven consumed more, from 3% to two peaks of 80% and 75% increases. During the interviews, the two participants who consumed considerably more than their reference consumption pointed out to us that was the case, explaining in one case that the increase was due to guests visiting them. These results can be explained by the short duration of our trial: domestic consumption can fluctuate considerably depending on irregular loads like usage of oven and washing machine. It is worth emphasising again that in this trial financially rewarding the savings was a strategy to draw the participants attention to their own energy consumption, and hence to the FigureEnergy system.

Limitations

The interviews also revealed some of the limitations of FigureEnergy. The most prominent one is that, in the Practice view, while all participants liked the relative size of boxes to indicate different amounts of energy consumed, they all found the representation of the always-on and not-annotated energy confusing. The metaphor of the liquids was not successful: when removing events from the practice tub, most users had the impression that the volume of the always-on energy was not constant. This may be because of the rearrangement of the boxes layout every time a box is removed or added. As a direct consequence, the adjustable *savings line* also turned out to be unusable.

The second main limitation is that FigureEnergy is not very useful to those who either consume very little (since consumption can be summarised by few events) or are already very aware of their energy consumption. One of our participants, Mike, a 25 year old Engineering PhD student, sharing a house with a relative, who had extensive experience with electricity meters, told us that because he consumes very little energy and he was already very aware of his household consumption, he did not learn much from using the system. This finding is consistent with earlier studies of energy displays [35].

Study Limitations

Our interviews were limited to one participant per household and focussed on their individual interaction with FigureEnergy. However, as most of our participants live with others, the system may influence the social dynamics in the home and play a role in the household collective decision making around energy. Further research needs to be carried out to shed light on these aspects. Moreover, recruiting from a university population introduced a bias, as all participants were educated to above average levels, so it is important to extend this work to different populations.

Implications

Interactive Displays

To the best of our knowledge, all prior designs of energy displays focused on the *presentation* of information. Users' *interaction* is reported to happen *around* the display, for example by turning devices on and off and seeing the changes on the display [16]. The engagement and discoveries that our participants described in the interviews, as well as the amount and variety of user generated annotations captured by FigureEnergy, suggest that there is a lot of potential for pulling the interaction *into the energy displays*.

The paradigms of direct manipulation and user-generated content can be applied to the design of eco-feedback technology. While our metaphor of energy consumption events (i.e., the event boxes in the Practice page) proved successful, the failure of the liquid metaphor indicates that other metaphors should be explored. The comments that our participants made about more advanced tools for analysis suggest that the potential for further work in this area is vast. Ideas brought forward in the interviews included semi-automatic event segmentation (i.e. using NILM) and various ways to aggregate consumption statistics.

Methodology

In their survey of eco-feedback technology, Froehlich et al. [19] point out that a considerable amount of research in this area employs low-to-medium fidelity prototypes evaluated through lab studies. At the same time, energy consumption is absolutely integral to everyday practices and attitudes, so lab-based studies can potentially miss important aspects of the users' experience with this kind of technology. Indeed, field evaluations of commercial energy monitors highlight the conflicts between design and real-world usage [35].

Our deployment demonstrates a middle-ground between low-to-middle fidelity solutions and the engineering of robust sensing infrastructure. A combination of web-based technologies, generally used for low-fidelity prototyping, and off-the-shelf networked electricity meters allowed us to evaluate our design in real homes without the challenges associated to rolling out our own hardware. We believe that such hybrid prototyping strategy can open the way for large-scale, even Web-based, studies of ubiquitous technology [9].

CONCLUSIONS

This paper introduced FigureEnergy, an interactive visualisation of domestic electricity consumption that leverages user-generated annotations to provide multiple views of the data. The manual annotation was designed to allow and encourage users to engage with the data, relating it to events in their day-to-day activities, and in this way help them make sense of it. We tested FigureEnergy in a field trial *in the wild*: twelve participants installed the system in their own homes and interacted with it for at least two weeks. The results of the trial show that the design was largely successful in engaging users. Even though many users had prior experience with other electricity displays, FigureEnergy allowed them to discover that some appliances consume more than they expected. The manual annotation process led users to refer to energy consumed by *activities* rather just by appliances and it made some of them more aware of the activities taking place in their household.

These results indicate that there is potential for *interactive* eco-feedback technology that engages users with their data beyond mere presentation. Moreover, our prototyping approach based on Web technologies and off-the-shelf sensors can be valuable for future real-world trials in domestic settings. Future work will explore the application of machine learning techniques to support users' annotations and provide suggestions for more efficient energy usage, possibly integrating information from other sensors as well, such as temperature and presence sensors. We are also keen to run longer term evaluations to assess possible effects on behaviour change. More in general, we hope that this work will promote more design and evaluation of Ubicomp interventions to help users understand the environmental impact of their day to day activities, relating the personal sphere to a global scale.

ACKNOWLEDGEMENTS

This work was partially supported by the EPSRC ORCHID project (EP/I011587/1) and the Adventures in Research 2011 fund from ECS, U. of Southampton. Thanks to T. Rodden, J. Fischer, N. Pantidi, Mike Hazas and the anonymous reviewers for the insightful comments.

REFERENCES

1. Abrahamse, W., Steg, L., Vlek, C., and Rothengatter, T. A review of intervention studies aimed at household energy conservation. *J. of Environmental Psychology 25*, 3 (Sept. 2005), 273–291.

2. Ames, M., and Naaman, M. Why we tag: motivations for annotation in mobile and online media. In *Proc. CHI '07*, ACM (2007), 971–980.

3. Berges, M., Goldman, E., Matthews, H., and Soibelman, L. Enhancing electricity audits in residential buildings with nonintrusive load monitoring. *J. of Industrial Ecology 14*, 5 (2010), 844–858.

4. Bertoldi, P., and Atanasiu, B. Electricity consumption and efficiency trends in the enlarged european union. Tech. rep., European Commission, Institute of Environment and Sustainability, 2007.

5. Broms, L., Ehrnberger, K., Ilstedt Hejlm, S., and Bång, M. The energy aware clock : Incorporating electricity use in the social interactions of everyday life. In *Proc. of the 6th Intl. Symposium on Environmentally Conscious Design and Inverse Manufacturing* (2009).

6. Carlis, J. V., and Konstan, J. A. Interactive Visualization of Serial Periodic Data. *Focus* (1998), 29–38.

7. Chetty, M., Tran, D., and Grinter, R. Getting to Green: Understanding Resource Consumption in the Home. In *Proc. Ubicomp '08*, ACM (2008), 242–251.

8. Cleveland, W. S., and McGill, R. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *J. American Statistical Association 79*, 387 (1984), 531–554.

9. Costanza, E., Giaccone, M., Kueng, O., Shelley, S., and Huang, J. Ubicomp to the masses: a large-scale study of two tangible interfaces for download. In *Proc. Ubicomp '10*, ACM (2010), 173–182.

10. Darby, S. The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering Billing and direct Displays April*, April (2006), 21.

11. Dillahunt, T., Mankoff, J., and Paulos, E. It's not all about green: Energy use in low-income communities. *Proc. Ubicomp '09* (2009), 255–264.

12. Disalvo, C., Sengers, P., and Brynjarsdóttir, H. Mapping the Landscape of Sustainable HCI. In *CHI '10*, vol. 30332 (2010).

13. Egan, C., Kempton, W., Eide, A., Lord, D., and Payne, C. How customers interpret and use comparative graphics of their energy use. In *Proc. 1996 ACEEE Summer Study.* (1996).

14. Ellis, S., and Groth, D. A collaborative annotation system for data visualization. In *Proc. AVI'04*, ACM (2004), 411–414.

15. Fischer, C. Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency 1*, 1 (2008), 79–104.

16. Fitzpatrick, G., and Smith, G. Technology-enabled feedback on domestic energy consumption: Articulating a set of design concerns. *IEEE Pervasive Computing 8* (January 2009), 37–44.

17. Fogg, B. J. Persuasive technology: using computers to change what we think and do. *Ubiquity 2002* (December 2002).

18. Foster, D., Blythe, M., Cairns, P., and Lawson, S. Competitive carbon counting: can social networking sites make saving energy more enjoyable? In *CHI EA '10*, ACM (2010), 4039–4044.

19. Froehlich, J., Findlater, L., and Landay, J. The design of eco-feedback technology. In *Proc. CHI '10*, ACM (2010), 1999–2008.

20. Gupta, S., Reynolds, M., and Patel, S. ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. In *Proc. Ubicomp '10*, ACM (2010), 139–148.

21. Gustafsson, A., and Gyllenswärd, M. The power-aware cord: energy awareness through ambient information display. In *CHI EA '05*, ACM (2005), 1423–1426.

22. Gyllensward, M., Gustafsson, A., and Bang, M. Visualizing energy consumption of radiators. In *Proc. PERSUASIVE '06*, Springer-Verlag (Berlin, Heidelberg, 2006), 167–170.

23. Hargreaves, T., Nye, M., and Burgess, J. Making energy visible: A qualitative field study of how householders interact with feedback from smart energy monitors. *Energy Policy 38*, 10 (2010).

24. Heer, J., Viégas, F., and Wattenberg, M. Voyagers and voyeurs: supporting asynchronous collaborative information visualization. In *Proc. CHI'07*, ACM (2007), 1029–1038.

25. Holmes, T. Eco-visualization : Combining Art and Technology to Reduce Energy Consumption. *Cognition* (2007), 153–162.

26. Kolter, J., and Johnson, M. Redd: A public data set for energy disaggregation research. In *Proc. of the SustKDD workshop on Data Mining Applications in Sustainability* (2011).

27. Larrick, R. P., and Soll, J. B. The MPG Illusion. *Science 320*, 5883 (2008), 1593–1594.

28. Mamykina, L., Miller, A. D., Grevet, C., Medynskiy, Y., Terry, M. A., Mynatt, E. D., and Davidson, P. R. Examining the impact of collaborative tagging on sensemaking in nutrition management. In *Proc. CHI '11*, ACM (2011), 657–666.

29. Papert, S. *Mindstorms: Children, computers, and powerful ideas*. Basic Books, Inc., 1980.

30. Patel, S., Robertson, T., and Kientz, J. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. *Proc. Ubicomp '07* (2007), 271–288.

31. Pierce, J., Fan, C., Lomas, D., Marcu, G., and Paulos, E. Some consideration on the (in)effectiveness of residential energy feedback systems. In *Proc. DIS '10*, ACM (2010).

32. Pierce, J., and Paulos, E. Materializing energy. In *Proc. DIS'10*, ACM (2010), 113–122.

33. Ramchurn, S., Vytelingum, P., Rogers, A., and Jennings, N. Putting the "smarts" into the smart grid: A grand challenge for artificial intelligence. *CACM* (2011).

34. Rodgers, J., and Bartram, L. Exploring ambient and artistic visualization for residential energy use feedback. *IEEE transactions on visualization and computer graphics 17*, 12 (Dec. 2011), 2489–97.

35. Strengers, Y. A. Designing eco-feedback systems for everyday life. In *Proc. CHI '11*, ACM (2011), 2135–2144.

36. Woodruff, A., Hasbrouck, J., and Augustin, S. A bright green perspective on sustainable choices. In *Proc. CHI '08*, ACM (2008), 313–322.

37. Zeifman, M., and Roth, K. Nonintrusive appliance load monitoring: Review and outlook. *Consumer Electronics, IEEE Transactions on 57*, 1 (2011), 76–84.