Recommending Energy Tariffs and Load Shifting Based on Smart Household Usage Profiling

Joel E. Fischer[*], Sarvapali D. Ramchurn[+], Michael A. Osborne[~], O. Parson[+], T.D. Huynh[+], M. Alam[+], N. Pantidi[*], S. Moran[*], K. Bachour[*], S. Reece[~], E. Costanza[+], T. Rodden[*], N.R. Jennings[+]

*University of Nottingham
Nottingham, UK, NG8 1BB
joel.fischer@nottingham.ac.uk

[+]University of Southampton
{sdr,op106,tdh,moody,ec,nrj}@
ecs.soton.ac.uk

[~]University of Oxford
{mosb,reece}@
robots.ox.ac.uk

ABSTRACT

We present a system and study of personalized energy-related recommendation. AgentSwitch utilizes electricity usage data collected from users' households over a period of time to realize a range of smart energy-related recommendations on energy tariffs, load detection and usage shifting. The web service is driven by a third party real-time energy tariff API (uSwitch), an energy data store, a set of algorithms for usage prediction, and appliance-level load disaggregation. We present the system design and user evaluation consisting of interviews and interface walkthroughs. We recruited participants from a previous study during which three months of their household's energy use was recorded to evaluate personalized recommendations in AgentSwitch. Our contributions are a) a systems architecture for personalized energy services; and b) findings from the evaluation that reveal challenges in designing energy-related recommender systems. In response to the challenges we formulate design recommendations to mitigate barriers to switching tariffs, to incentivize load shifting, and to automate energy management.

Author Keywords

Energy tariffs; recommender systems; load shifting; personalization; smart grid; demand response.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Human Factors; Design.

INTRODUCTION

In the face of dwindling fossil fuel resources, energy security and efficiency has emerged as one of the key global concerns of our time. Many countries have implemented political, societal and technological initiatives to respond to the challenges of how to become more energy efficient [8]. For example, initiatives in the UK include the roll out of *smart meters,* a key technology associated with the idea of the *smart grid* [9]. Essentially, the smart grid will serve as an "information technology backbone" [15]. Two-way information flow between suppliers and consumers is essential for smart grid technologies aimed at making the grid more energy efficient, such as *demand side management*. For example, 'dynamic' or 'time-of-use' pricing have been suggested as techniques to reduce peak demand by encouraging *load shifting* of demand through higher prices at peak times to off-peak periods [31]. These techniques raise major challenges for interactive and intelligent systems design to support the consumers' understanding and control of their energy usage in this complex environment. For example, recommender systems have been proposed [18] in order to empower users to play an active role in demand side management.

Research into demand side technologies has been embraced by HCI and adjacent computing disciplines, and domestic energy consumption has emerged as a key application domain. In particular, efforts have been focused on persuasive technologies that provide feedback on consumption to raise awareness or to promote conservation behavior [5,13,24]. However, research has also shown that despite raised awareness, comfort and everyday habits can prevent true change [30]. For most people in the 'developed world', the fact that energy (still) is a readily available 'commodity' often appears to fail to engage people sufficiently to care. In spite of predictions of annual energy bills exceeding £1,500 for the average UK household [12], a recent study has shown that people spend as little as two hours per year to shop for energy tariffs [32].

Against this background, we present a system and evaluation that takes on the challenges of people's limited interest in energy tariffs and explores reactions to load shifting advice. The premise behind AgentSwitch is that the future proliferation of monitoring devices such as smart meters connected to online data stores will lead to intelligent services that process, analyze and reason based on energy usage data. AgentSwitch presents a novel prototype of such a service, recommending energy tariffs and load shifting based on energy usage profiles (EUP). Our research interest focuses on how user interaction with such systems might best be supported. We evaluate the system through a user study consisting of task-based walkthroughs and interviews. Our findings highlight perceived barriers to switch to

IUI'13, March 18–22, 2013, Santa Monica, California, USA.

cheaper tariffs. In particular, for recommendations of shifting appliance usage to different times, it appears that more benefits than monetary savings are required to incentivize changing behavior in this way.

MANAGING ENERGY ON THE DEMAND SIDE

HCI systems to encourage conservation often provide multiple views of users' energy usage at various levels of temporal granularity, and may for example combine these with persuasive messages sent to the user's inbox [17]. Recently, this focus on *consumption feedback* has been critiqued for its 'one size fits all' approach that glosses over people's differences in motivations [14], and can reduce energy issues to an optimization problem [5]. An inherent rationalistic orientation that feedback leads to reduction does also not take into account that everyday practices are often driven by comfort and convenience rather than by rational calculations [30]. A recent review of the field has also called for more engagement with the opportunities smart grid systems hold to promote energy efficiency and reductions in costs for consumers, producers and the environment [24].

Smart grid and smart home energy management

A key question for our research is how can smart grid technologies support consumers in becoming more energy efficient without impeding too much on their comfort and convenience. As the US Department of Energy writes in their report introducing the smart grid:

"Consumers are not interested in sitting around for an hour a day to change how their house uses energy; what they will do is spend two hours per year to set their comfort, price and environmental preferences – enabling collaboration with the grid to occur automatically on their behalf and saving money each time." [13: 20].

More automated approaches to home energy management have been proposed that rely on machine learning and multi-agent systems [1], e.g., to enable automatic appliance control [26], home heating based on occupancy [29], or off-grid home energy management [2].

The work presented here engages with these proposals. Firstly, AgentSwitch is designed as a service that consumers could use only once or twice a year to optimize their energy use. Secondly, we are interested in users' reactions to personalized recommendations on how much they can save by shifting loads to off-peak times, under current tariffs and assuming no automation. Will the potential savings be perceived as sufficient to motivate behavior change? In this light, what are people's attitudes towards automation to support load shifting?

Recommender Systems and Energy

At the time of writing, the consumer in the UK can choose between 24 energy providers that offer plans with many different payment options, energy sources, fixed price and tiered tariff structures. Currently, 'Economy7' is the most common tiered tariff type that offers a cheaper rate for a seven-hour period during the night. Consumers also need a special meter installed in their homes to switch to this tariff. Conceivably, consumers may find choosing the right energy tariff a daunting task. Hence, a recommender systems approach proposed to cope with the problem of information overload [21] appears to be worth investigating.

Generally speaking, recommender systems can be distinguished by whether they employ a content-based or a collaborative strategy (or a mix of both) [21]. They have been applied to the energy domain, for example to support building automation [18]. Context-aware and model-based recommender systems have also taken into account behavior [19] and user state in time [16]; approaches that are more closely aligned with the energy usage profiling our work relies on. Price comparison websites also employ profiling to support selecting appropriate products such as insurances, mobile phone contracts and energy tariffs. Such services usually require the user to manually provide parameters of preference to narrow the choices. In contrast, our approach relies on the analysis of pre-recorded data to enable *personalized* price comparison based on energy usage. Potential privacy issues related to access to this kind of fine-grained data have been highlighted [28], and will also be explored in the evaluation.

We draw on a user-centric evaluation framework for recommender systems [25] to guide the evaluation of our system. In the context of giving energy advice, literature shows that prior domain knowledge and commitment to save energy are important factors that influence the intention to follow advice [7,22]. To evaluate our system, the intention to follow recommendations is an important indicator of usefulness, in addition to factors relating to the predicted information, such as confidence and intelligibility, as well as potential issues with privacy.

AGENTSWITCH SYSTEM ARCHITECTURE

AgentSwitch relies on energy data collected from household-level monitoring to compute personalized recommendations. The web service is implemented as a RESTful Django web application,[1] consisting of a number of core modules (see figure 1). The system setup is designed for a deregulated electricity market (e.g., the consumer can chose between various suppliers and tariffs) as found in a number of countries across the world. Its modular design should make it relatively straightforward to plug in other tariff data sources and metering equipment available than the ones used in our deployment.

Household monitoring

The evaluation presented in this paper is based on data collected during a three-month deployment in 18 households in the UK. The technical setup consisted of a commercially available current sensor (CT-clamp) wirelessly transmitting Watt readings to an off-the-shelf In-home display (IHD). A low-power (< 8W), small form-factor computer connected

[1] http://hac.ecs.soton.ac.uk/agentswitch

to the display aggregates the readings and pushes them to the data store via the home's broadband every 5 seconds.

Data store

AgentSwitch retrieves the user's personal energy data from the data store that exposes an API for authentication and data requests. The user logs in to authorize AgentSwitch to retrieve their data, which is explained in the UI. The data is only handled for the duration of the session; it is not stored persistently by AgentSwitch. The data store is part of a bigger vision in which a multitude of web and mobile applications provide personalized energy-specific services. The user authorizes the applications to use their energy data that is stored in a single secure online location.

Energy usage profiling

Once the actual energy data is retrieved from the data store, an EUP for the whole year is predicted based on the actual data. Annual consumption estimates are a required input parameter for the energy tariff API. We use a Gaussian process (GP) to model power consumption as a function of time, using scaled national average consumption to provide a mean function. The latter embeds seasonal variations within our model, meaning that we can provide accurate power consumption prediction on yearlong scales even with sparse data. We can then employ this GP model to estimate the integral of power consumption over a year, giving an estimate of a household's total annual energy consumption. This technique is known as Bayesian Quadrature (BQ), a model-based means of numerical integration [10]. In this work, we adapt the technique for quasi-periodic energy consumption signals emerging from the weekly cycles of typical domestic consumers.

Evaluating BQ

We evaluated the efficacy of this approach with the data collected during the deployment: we divided the data into a set of training data (one third of the collected data) and test data; the goal was to accurately predict the total energy usage given only training data. Results reported in more detail in [27] show that the approach generates more accurate predictions than simpler alternatives.

A key challenge for the prediction module is computation time: users are usually unwilling to wait more than a few seconds to receive recommendations. The Bayesian quadrature algorithm requires fitting a GP to the data, an expensive operation that scales as $O(N^3)$ in the length of the data, N. To tackle this problem, we aggregated our data, therefore reducing its size. Rather than using all minute-by-minute observations, we supplied to the model only observations of the total energy consumed in an hour. This was found to have negligible effect on predictive accuracy due to the redundancy in the finer-grained observations. Training the hyperparameters of the GP is a particularly computationally demanding process. We therefore trained most hyperparameters on a fixed set of training data off-line, reducing the required processing for the on-line module to about 0.02 seconds for N=5 and 0.4 seconds for N=500, for example.

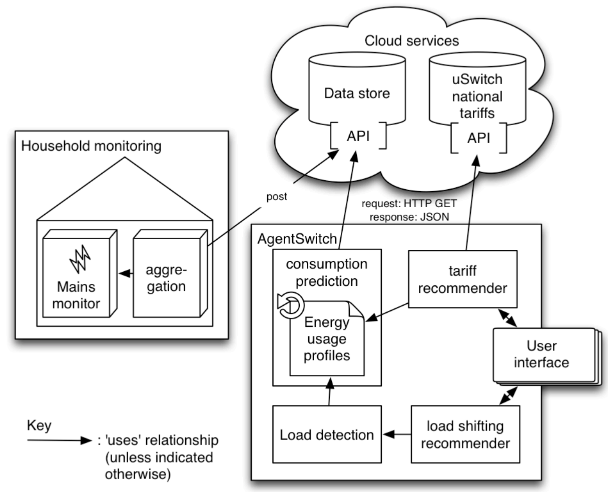

Figure 1. AgentSwitch system and infrastructure.

uSwitch API wrapper

AgentSwitch makes essential use of national tariff data provided by uSwitch.com, a price comparison website that provides its up-to-date energy tariff information through an API. In AgentSwitch the user selects their tariff details from a set of locally available tariffs retrieved from uSwitch (e.g., supplier and plan details). The recommender module then uses the details (e.g., unit rates, thresholds) to compute annual cost estimates based on personal usage profiles.

Load detection

Deferrable loads (non-overlapping uses of the washing machine, clothes dryer or dishwasher) are detected from the aggregate household-level data. Loads from these appliances share similar profiles, and therefore a template can be constructed to match the behavior of all of them. To detect deferrable loads disaggregation is performed using a variant of the hidden Markov model approach described by Parson et al. [23]. The disaggregation module takes a list of time-stamped energy readings and returns a list of deferrable loads. However, due to the lack of training data and low data granularity, it is very hard to determine which of the three appliances generated an individual load. Whether the accuracy of such appliance-level disaggregation is sufficient from the user's perspective is investigated in the evaluation.

User interface

The GUI (e.g., see figure 3, 4 and 5) has been iteratively developed from wireframe sketches through several rounds of expert walkthroughs. The following key considerations have driven the development.

Minimize manual user input

A key feature of AgentSwitch compared to existing commercial price comparison websites is that the user does not have to provide parameters detailing their energy consumption. Arguably, manual input is error prone, can be tedious, and important details like the proportion of daytime vs. nighttime consumption are difficult to estimate. In AgentSwitch, after the user logs on they only have to provide a postcode (to select the right set of locally available tariffs). Their usage profile is computed from their actual energy data fetched from the data store. Comparison of their current tariff to the list of available tariffs is optional, which requires the user to manually provide parameters detailing their current tariff to the system.

Explanation driven by data provenance

Throughout the AgentSwitch UI, care is taken not to overload views with information. Instead, more information is provided on demand through hover-overs and links. Availability of explanations of system actions was deemed important in concordance with findings in the literature on recommender systems [6]. In order to support understanding and confidence in the recommendations, explanations are constructed using the *provenance* of data used in computations. Provenance information is modeled using the W3C PROV Data Model [20], which captures the various relationships between (data) entities, activities (e.g. consumption prediction, load disaggregation), and agents (e.g. the user, uSwitch, the data store). For example, using such traces, users may be able to identify incorrect predictions of yearly consumption, or justify changes to daily routines to make significant savings.

In particular, due to its potential to raise concerns related to privacy and confidence in the accuracy of the presented information, the following explanations are provided on demand:

- Explanations justifying why personal data is required and how personal energy data is handled.
- Explanations describing how the presented information was computed, including its provenance (i.e., its origin).

Familiar concepts

Throughout the application, interface concepts were deployed that we hoped most people would be familiar with. In addition to visual layout inspired by existing price comparison websites, we adopted the concept of *authorizing an application to use your data*. Users should be familiar with this for example from logging on to OAuth-enabled social networking application clients. It is conceivable that future third party energy applications will be able to request usage of users' energy data stored in a private online repository in the same way, much like initiatives such as Green Button[2] already advocate.

EVALUATION

The key objective for AgentSwitch is to provide a useful and trustworthy personalized service that allows the user to make an informed decision on how to save by switching tariffs or shifting energy usage. As outlined in the previous section, this is realized through various modules that compute information surfaced by the UI. Drawing on [25], the focus of the evaluation is on how users with genuine personal energy data perceive AgentSwitch with regards to

Participant	Colleague	Home occupancy	Age, gender	Occupation
P1	Yes	2	29, m	PhD student
P2	No	1	32, f	Marketing manager
P3	No	4	48, m	University manager
P4	Yes	2	34, m	Researcher
P5	No	4	34, f	Artist
P6	No	2	35, f	Translator
P7	No	3	49, f	Housewife
P8	Yes	2	25, m	PhD student
P9	Yes	2	27, m	Researcher
P10	Yes	3	55, m	PhD student

Table 1. Participant details.

- **usefulness** and use cases of provided functionality,
- **provenance** (Where did this data/figure come from?),
- **comprehensiveness** (Is all the desired information present?),
- **confidence** (Do users trust that the information is accurate?),
- **intelligibility** (Is it possible to understand the information, and the way it was computed?),
- **actionability** (Will users do anything as a result of this information?) of the information, and
- **privacy** issues and emerging attitudes towards the use of personal energy data to drive the service.

We chose to employ a task-based interface walkthrough and also asked our participants to 'think aloud' while completing the tasks. After the walkthrough, a researcher conducted a semi-structured interview to investigate their experience of using AgentSwitch in more detail. We recorded video, audio and screencasts during the procedure. For qualitative analysis, the statements made during the walkthrough and the interview were coded and grouped according to emerging broader themes, of which the quotes present in this paper are representative [4]. Before we present the results from the study, we briefly describe participant recruitment, procedure, walkthrough tasks, and interview structure.

Participants

To evaluate AgentSwitch's core functionality, it is essential that there is a sufficient amount of energy data available for each user. Hence, we recruited 10 participants that took part in a study approximately one year earlier during which we recorded three months worth of their home's electricity usage data aggregated to 5min intervals (in Watt-hours). In the earlier study, participants had been exposed to real-time and historic electricity feedback delivered via various devices such as an In-home display (IHD), a web app, and text messages. Each participant was from a different household with varying occupancy (see table 1). Due to the

[2] www.greenbuttondata.org

sensitive nature of conducting a long-term study in their homes, participants had been recruited among colleagues that were not involved with the research (5) and friends of friends (5). The only requirement we had is that they frequently handled energy bill payment in their household.

Procedure

We invited our 10 participants to come to one of three sessions in a meeting space in the local University. We asked them to have their current energy tariff details to hand as they were required for the study. Two of the sessions had three participants; the other session had four.

Upon arrival we briefed participants on the study procedure and gained informed consent. We then demonstrated uSwitch to them, a commercial price comparison website. The site offers energy tariff comparison based on annual usage data users have to input manually. The purpose of the demo was to acknowledge an underlying service we built on, but also to give participants the opportunity to compare AgentSwitch to an existing solution.

One researcher accompanied one participant to a laptop set up in a separate and quiet part of the lab. The browser already displayed the landing page of AgentSwitch. Participants received printed task instructions for the walkthrough. The researcher then remained with the participant throughout the walkthrough, encouraged them to think aloud and clarified or helped where necessary. Finally, interviews were conducted. The entire procedure usually lasted about 45 minutes.

Walkthrough tasks

To evaluate AgentSwitch's core functionality in use, we designed a set of walkthrough tasks to be completed by each participant. Participants were tasked to answer a set of questions that required them to explore most of AgentSwitch's UI. We wanted to ensure participants engaged with most of the functionality in a systematic and comparable way, to enable later discussion about the actual figures AgentSwitch presented.

Tasks were ordered in increasing complexity and in keeping with the order of interface transitions. Participants had to find information by engaging with the following features:

- Estimation of annual cost for their current tariff (figure 3).
- Breakdown of annual cost into daytime and nighttime (if they had an Economy7 tariff) (see figure 3).
- Cheapest recommended standard and Economy7 tariff, and estimated cost and savings compared to their current tariff (see figure 4).
- Potential savings by shifting 20% of daytime usage to nighttime (based on their own Economy 7 tariff, or else based on the cheapest Economy7 tariff).
- Percentage of their overall use accounted for by detected deferrable loads (washing machine, dish washer, or tumble dryer) (see figure 5).

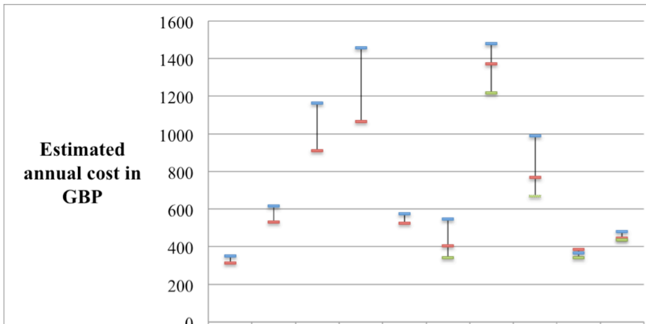

	P1	P2	P6	P7	P9	P3	P4	P5	P8	P10
current tariff	352	615	1164	1456	574	546	1479	989	368	482
cheapest available tariff	313	531	912	1065	525	406	1373	768	385	447
shifting 20% usage to nighttime						342	1218	669	342	437

Figure 2. Estimated annual cost for current tariff, cheapest available tariff, and when shifting 20% of usage from daytime to nighttime (for participants on Economy 7 tariffs).

- Potential savings by shifting all of the detected deferrable loads from daytime to nighttime (see figure 5).

Interviews

A researcher conducted a semi-structured interview per participant after the walkthrough. Aside from allowing for first impressions and overall comments or concerns, questions were designed for participants to elaborate on perceived usefulness and whether and for what reasons they would use AgentSwitch again in the future. Questions also probed their understanding of the provenance of the information, their confidence in the accuracy, and whether they would consider switching tariffs or shifting usage as a result of using the service. Lastly, questions probed potential privacy issues as well as attitudes towards more autonomous systems that perform actions like switching tariffs automatically on the user's behalf.

RESULTS

Participants found AgentSwitch relatively easy to use and easy to make sense of. The fact that all of the participants said they would use the site again speaks to its perceived usefulness. Confirming initial motivations, participants' comments suggest that the domain of energy tariffs only attracts usage once or twice per year, or less often.

We begin by giving an overview of the personalized figures AgentSwitch presented, and how they were typically made sense of and perceived by our participants. We then move on to present qualitative findings from a thematic analysis of statements participants made during the walkthrough and the interview.

Personalized recommendations

AgentSwitch found cheaper tariffs for 9 of our 10 participants, with estimated annual savings between £35 and £391 (M=£132) if they switched to the cheapest tariff. In addition, AgentSwitch calculated that participants with an Economy 7 tariff could save an additional £26-£110 if they shifted 20% of their overall usage to the cheaper night rate of the two-tiered tariff (see figure 2).

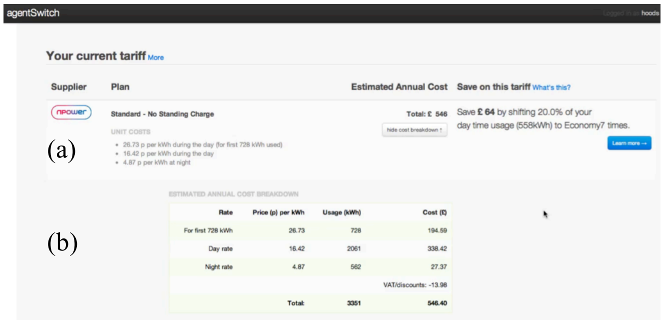

Figure 3. Screenshot of annual cost estimates (a) and breakdown (b) for P3's current tariff.

Making sense of the recommendations

The following transcript from the walkthrough illustrates how a participant is using and making sense of the presented information, discovering discrepancies in the predictions compared to his provider (see figure 3). This participant has brought paperwork detailing his power company's estimations of their annual cost.

[Reading the results page] *Estimated Annual cost. That must be using the data you hold. You reckon it'll cost us, electricity wise 540 pounds a year -* [looks down at paperwork] *which…. They* [the supplier] *think is a lot more. By 200 pounds worth. Because - they reckon I'm gonna use a lot. I'm thinking the majority is at 16p during the day. So… let's see how you break it down first of all.* [clicks on 'show cost breakdown' – reading breakdown] *Day rate -* [looks down] *yea. that's the difference. It's 200 pounds - it's virtually 170 pounds difference in the day rate. The night rate is virtually the same -so, it's in the day rate, that it's a big difference.* [P3]

Perceived accuracy of the predictions

AgentSwitch estimates annual consumption cost of the current tariff on the basis of a three-month data sample. For the purpose of this paper, more important than the accuracy of the predictions is how participants perceived the accuracy of the predicted figures. We anticipated that a certain level of perceived accuracy was important to further engage with and trust information provided by the system.

P1 and P7 perceived the estimate of the annual bill as about right, while P4, P5 and P6 felt that is was higher than expected; and P2, P8, P9 felt it was lower than expected (P10 does not comment). This discrepancy in the data does not necessarily lead to a distrust of the system, as P3 puts it,

*I mean, **the thing that struck me was the difference first of all between what you estimate my cost should be. And what they're*** **[the power company]** ***estimating it should be**. I mean, the interesting thing is why are they estimating an awful lot more than you are?* [P3]

This statement illustrates how confidence in the presented information extends beyond immediate system usage:

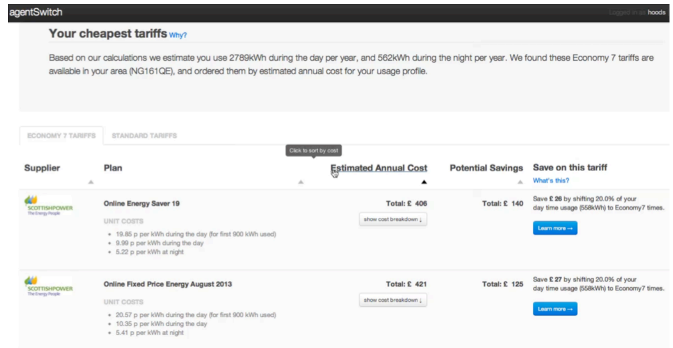

Figure 4. 2 of the recommended cheapest available tariffs for P3.

*I'd want to make sure that information is accurate. I don't know how I would... I mean I supposed I would have to look up my... go back through my bills to double check. I think I only feel that way because I know that this system is new. And I know also that the data that was collected was over a certain period of time. Whereas if let's say I was involved in a year-long, or even six months actually... **If I'd have been involved in a six month project, where over six months my data is collected, and I knew that this technology, the glitches and issues, were ironed out, then I would be quite happy to trust the information** I was being given and I would switch, definitely*. [P5]

Barriers to following advice on switching

Despite that cheaper tariffs than the one they were currently on were presented to 9 out of 10 participants, they presented many accounts of the barriers that made the prospect of switching tariffs less desirable.

Inconvenience

Despite efforts of consumer-oriented agencies, switching tariffs and particularly providers is perceived by all of our participants as an inconvenience.

*I get the impression that you could just keep moving. Your whole life, and **it's a bit of a pain every time**… you get the whole welcome pack and things. It's not something I want to do all the time.* [P1]

*One of the reasons why we haven't changed our energy provider was because I did uSwitch and I found it was such **a hassle to go through the process of changing***. [P5]

Switching tariffs […] *I mean, it's a **necessarily horrifying and complicated process** anyway, ahem, because there are so many different suppliers*. [P4]

(Not) knowing the provider

Another stated barrier for switching is unfamiliarity with the provider that offered the cheapest tariff. P3 states that they prefer to be with a supplier that actually owned the infrastructure as well and were not just providing a service. This could be described as a 'reputational effect'.

Disagreeing with the provider's policy

Prior experience and attitudes towards the recommended provider play a role in why switching is rejected. P7 says they would not switch to the provider who offered the cheapest tariff *…because* ***I object to their selling techniques, I object to them on a moral [basis]****…* [P7]

Environmental vs. monetary motivation

'Green' motives can outweigh monetary incentives to switch. P7 goes on to highlight that they are on their current tariff for environmental reasons.

The reason we're with the tariff is because, ***it's the green side of it****, I know we're using electricity, we all are but I kind of justify it by thinking well if that's being offset slightly by them* [current supplier] *investing in renewable energy, so I know its not the cheapest, but that's not why I use them.* [P7]

Satisfied with the current provider's service

There can be a myriad of other reasons why users might prefer to stay on their current tariff, or with their current supplier. For example P1 states,

[current supplier] ***has a really cool website*** *which lets me see what I'm using… I think if that need was filled by some kind of personal energy monitoring then… the appeal of* [current supplier] *might go away a little bit. And then maybe I'd think about changing then.* [P1]

Cost of switching

The current supplier can sometimes penalize switching providers before the end of a contract (for example on a fixed term tariff). In addition, the procedure of signing up with a new supplier can also be associated with costs, monetary or in terms of time spent.

… ***what kind of penalties there might be for changing****, if that information is available, I don't know; and perhaps in the new tariffs, what kind of procedure is expected to be followed, do I need to do anything with it (...) even if it does give you a nice button, sometimes you are not really sure where is that going to take you to; it requires a lot of information.* [P2]

Incompleteness of AgentSwitch

The prematurity of the prototype meant that energy tariffs could not be covered comprehensively. For example, combined gas and electricity tariffs ('dual fuel') were excluded from the system as AgentSwitch relied solely on electricity usage data.

…the only slight problem is that because I have a dual plan, I don't know if this… ***I don't think this takes it into account****. Because obviously you're not monitoring gas use. I suppose if the system existed that would allow me to do that I would* [consider switching]. [P4]

However, consulting other sources can probably mitigate problems with incompleteness:

I will probably go into for example the [supplier] *website and check the details again. I think this is very useful in terms of,* ***it gives you the idea*** *of what is the cheapest and what is the second cheapest****, so I will probably go into the cheapest and the second cheapest and check their websites and decide which one to go for****.* [P6]

Thresholds for switching

The perhaps most obvious reason not to switch is that savings are not perceived to be high enough to incentivize switching.

It's ***not actually as much savings as I thought****. There's a threshold… it's not worth changing supplier for 2 pounds a year. It might be worth it for 40 pounds a year… maybe… I'm more inclined now after having a look at this I'd say… to.. to think about changing. But I probably won't. I'm not on a tragic tariff.* [P1]

In fact, participants often talk about the threshold that potential savings have to exceed to motivate them to investigate switching further.

I did do a price comparison on a website once, I think it was after I had been in the flat for a year so I had an annual figure and I think it said ***the most I would be able to save was 60 pounds over a year so I decided it was not worth moving****. But with these kinds of savings I think it would be worthwhile* [P2]

We're still talking of ***2-300 pound difference, which - is… worth exploring****. So I'd be like… right OK, tell you what,* [clicks Economy 7 tariff tab, see figure 4] *give me a screen print of that and I'll take it home, start to have a look, dig into it.* [P3]

I think it was like 60 quid ***and for a saving of 60 pounds over a whole year it just felt like it was more hassle*** *to change […]* ***whereas with this it's actually telling me that I could save almost 300 pounds*** *by a combination of changing my behavior and changing my tariff. So it's much more clear* [P6].

Interestingly, automation of switching could perhaps lower this threshold, as P6 elaborates.

If it's just 50 quid for a year's worth of energy, to me it's not worth the hassle cause I've had the experience before of switching providers. Whereas ***if there was a system that could do it for you****, if that system worked fairly effectively,* ***then the saving of 50 pounds or even 10 pounds is worth it****. Because there's little effort in doing that.* [P6]

Autonomous switching

Whether this is seen as a blessing or a curse seems to depend on the perceived accuracy of the system as experienced in the walkthrough

But ***it's discredited itself by giving me incorrect savings amounts****!* [P1]

as well as whether a host of conditions could be met by such an autonomous system:

Yes, I'd be quite happy to, ***as long as I had some notification and control over it*** *I'd be quite happy for an agent to change my tariff between energy companies. … It*

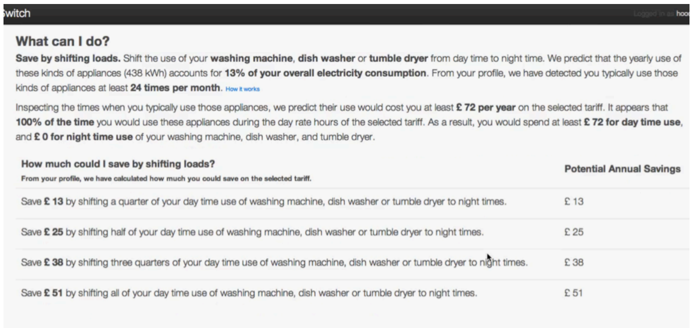

Figure 5. Participant P3's load shifting advice.

*would have to be set up sufficiently cleverly that **it would notify me when it was gonna do it and ask me whether I was willing to do it - for each time**. The reason why is that I'd want to know what's coming out of my bank account in any given month. So, as long as it has some **limiting factors**, like, I want to be on a **monthly direct debit and I don't want to spend more than X**, or something like that, then, I'd be quite happy for it to flick around.* [P4]

Load shifting

The load detection module made more specific recommendations possible, enabling detailed advice of how many loads users typically run, what their footprint is, how much they can save by shifting them and so on (see figure 5). People have their own ways of making sense of how the system is able to present this kind of information to them

You can tell the heavy loads they would have come on during - well all the time you're monitoring - say you've looked at where the heavy loads are and say - right ok, if they moved into the night time tariff then they would have done this, I can understand where the information could be gleaned from. Because of the information you collected. It's nice to see it all together. [P3]

Aside from praise, this participant also flags up the risk involved with running appliances during night time, which

*… we are reluctant to do … because of **risk of them going up in flames** … as we know a couple of people that had those issues. So we tend not to …* [P3]

Furthermore, the actual savings by shifting loads are perhaps not high enough to outweigh the cost of shifting.

*The **shifting the loads is not going to do us a great deal**. Behavior wise-we're **more comfortable with it being run during the day even if it costs us an extra 30 pounds a year**. That isn't a problem.* [P3]

In addition to noting that the potential for load shifting is limited to certain appliances, it was also noted that accommodating for shifting these loads comes with an extra cost.

*That would probably mean **buying a timer, or indeed a washing machine with a timer on it**. So washing I could do at night, and the dishwasher could go on at night… cooking as I say, I might struggle…. 'Wake up everybody, time for dinner'. I **mean paying off a timer switch, I've no idea what one costs, but say it's 20 pounds**, that's stupid... that's a complete waste of time, for the hassle factor.* [P7]

Understanding 'saving by shifting'

Whilst the concept of 'saving by switching' was familiar to all of our participants, the potential to save costs by shifting loads to different times of day with cheaper rates was novel to most of them. Hence, the information related to potential additional savings through shifting is sometimes misunderstood, and usually takes users more effort to make sense of.

*I find it difficult to imagine an interface which is clearer… but **it is a bit confusing at times**. You have to **stop and think about what you're looking at** and what the savings actually mean.* [P1]

People struggle perhaps especially because there are two kinds of recommendations related to savings (switching and shifting), which could be separated more clearly.

I mentioned before about the potential savings and then the savings on this tariff there isn't a clear indication of how those things are different and how they relate to one another. [P5]

However, once people make sense of the information, they appreciate its value.

*I do think the breakdown of **how much you can save by moving things - by moving energy usage to different times** is really interesting. Not so much, just saying, you know, you'll get some savings. But saying, you know, **here are actual, genuine numbers**, 'if you move THIS percentage of your usage.' Which - if you're talking about, erm, tumble dryers, dish washers, these are things than can easily be moved. You don't want to move your kettle, toaster, you know. It's really nice, it kinda just reminds you of the stuff. And that's a whole, kind of, area I hadn't thought about, since the last study… moving stuff to doing them at night, particularly.* [P4]

Provenance of information

When asked whether it was clear where the data came from and how the information was computed, none of our participants explicitly mention any of the features of the UI that support explanations. Instead, participants orient to the previous study in their responses.

*I know they're based on real usage - **because I watched it being recorded**, and I reviewed what it was saying* [P4]

Moreover, through this orientation participants may have gained a more favorable view of AgentSwitch.

*Maybe because I went through the previous process and I was able to see my data and how the data was collected and so on, so it might be **because I'm aware of how the data was collected that I feel that I can trust to a certain degree, I have a better understanding of what is actually there**…**if***

that wasn't the case and I wasn't aware perhaps I wouldn't be quite so quick to trust it. [P5]

Privacy

Records of energy data – especially if recorded continuously over a period of time – are sensitive pieces of information as people's behavior can be inferred from it to a certain degree. Two of our participants point out the ability to infer whether people are not at home from (near-) real-time energy consumption data. While these concern relate to energy data being publicly available, participants also voice concerns regarding sharing of energy data with energy suppliers.

If you're not in touch with any of the companies*, if the companies can't see it…****then that's not a problem****. I would only want my provider to know what tariff I'm on. I wouldn't want other people to say 'Ah, you're on this tariff, we can save you money', so that's none of their business.* [P2]

Other comments suggest that the personal benefit derived from recording and analyzing the information may mitigate concerns about sharing the data with the supplier.

The live information that's being produced by my activity and my family's activity, I'm not completely comfortable about organizations having access to that data and potentially selling that data to one another. Without it actually benefiting me. Whereas ***this I think can demonstrate that there is more benefit than what I've seen before***. [P5]

A couple of participants also point out that the voluntary nature of use is essential, which implies consent to the data being used to provide a service.

Well ***you don't have to use it****, so I'm gonna guess that if you were concerned about privacy in that way, then you wouldn't be using it at all.* [P7]

This participant puts it in more technical terms, and raises an issue of control over data and autonomous systems.

But this is kind of interesting because ***it is a pull service rather than a push service****. So I'm asking it to use my data to tell me something interesting, rather than IT is using my data to tell me something interesting. That distinction is what's important to me in terms of privacy.* [P4]

DISCUSSION

The findings reveal a set of barriers to switching despite that 9 out of 10 of our participants could save by switching. With regards to shifting loads, aside from some difficulties understanding the concept due to novelty, participants furthermore largely struggled to see the benefits in shifting. We now discuss the findings and suggest ways to overcome the barriers, and relate our findings back to broader concerns in the literature.

Making advice easier to follow: barrier-free switching

The study has revealed people's manifold barriers to switching. What can be done to address these barriers? Fundamentally, energy policy and regulation needs to be implemented to enable interaction designers to build better systems. For example, to make switching more convenient (e.g., enabling one-click switching), to commit energy providers to be more transparent about the provenance of the energy sold to consumers (e.g., to enable choosing tariffs based on the proportions of the energy-mix), and so on.

Some barriers are also more directly attributable to the system design, which can be addressed through improvements. Trust is a core issue for recommender systems [21]. To be perceived as trustworthy by users is important for these kinds of intelligent and complex systems [6], especially when they involve personal data [28]. Therefore, we suggest the following design recommendations:

Have a complete product base

Focus on the determining factors that help people make a decision. In the case of energy tariffs, in addition to pricing structure we found information on the provider (e.g., their policies, practices, area of operation etc.), and contractual details (e.g., on payment, duration, and penalties for premature cancelation), potential additional services of the provider (e.g., web tools and apps), and gas and dual fuel tariffs to be decisive factors that were not covered sufficiently by the current prototype of AgentSwitch.

Know the product recommended to be replaced

To help users make an informed decision on whether to switch to another supplier, the system needs to take into account the details of the current tariff. Does the current contract allow switching, are there any penalties associated with leaving the provider, how do other aspects of the current provider compare to the one to switch to?

Support different levels of motivation

People have different attitudes, beliefs and values. In addition to rationalistic price comparisons, people may want to base their decisions on prior attitudes, for example on environmental impact, or emotional persuasion, for example through positive reinforcement [14]. In terms of AgentSwitch, information on the environmental impact of the tariff could have helped to make an informed decision in this vein (e.g., the actual provenance of the electricity distributed to the home).

Make following recommendations as convenient as possible

Even if regulation means 'one-click' switching is not possible, improvements to the system that makes switching more convenient might be as simple as populating an email template with the desired tariff details, or making the relevant provider's website available with one click.

Reveal provenance and certainty of predictions on demand

People did not explicitly mention that they used the explanations to support intelligibility of the system (echoing recent findings [6]), but their orientation towards the previous study in which their data was recorded suggests that an understanding of the provenance of the predictions was important overall. In real-life settings (e.g., without the presence of explaining researchers), availability of explanations how the predictions were arrived at in terms of provenance of the raw data used is likely more important to be perceived as trustworthy.

In addition, statistical methods of predictions yield probability values that indicate the level of certainty associated with the prediction. It may also be of value to take into account the level of certainty when presenting the predictions. Whether and how to present uncertainty in intelligent UIs is a key question for future work.

Incentivizing load shifting

The study revealed participants' manifold cost/benefit trade-offs with regards to incentivizing switching tariffs and shifting loads. The analysis highlights that these barriers differ contextually by person, activity, provider, perceived reward, and so on.

Particularly, the responses suggest that the monetary benefits of load shifting under current tariff structure (e.g., Economy7) did not sufficiently motivate our participants to change their behavior. Incentives to shift loads may be more attractive in the future with rising prices. However, even small savings may be important for some; fuel poverty is a core problem in low-income communities [11].

For load shifting in particular, the onus is on governmental policy and energy providers to create the infrastructure and services to incentivize load shifting. For example, the anticipated roll out of smart grid technologies may lead to the introduction of dynamically priced tariffs to incentivize load shifting. The price structure of these dynamic, multi-tier tariffs can be expected to provide a larger monetary incentive.

In addition, once the smart grid is in place, interactive systems can then be realized that aim to motivate load shifting in similar ways to current prevailing research and system design to motivate reductions in energy consumption [13,17]. One approach might be to convey to consumers the environmental benefits of load shifting (e.g., avoiding peak demand and thus reducing carbon emissions). New kinds of 'smart' persuasive technologies could draw on grid data to dissuade use in high-demand periods, and reward consent not just with money, but also with game-like rewards or environmental praise. In future work we will investigate providing such incentives in addition to monetary incentives especially where the latter are not sufficiently high.

Towards autonomous energy management

The study has also revealed people's attitude towards systems that automate switching and shifting. The anticipated increase in convenience seems to lower the (monetary) threshold in savings required to make switching worthwhile. The results support recent findings that people are not interested in spending a great deal of time optimizing their home energy management [32]. This finding supports the potential of investing in research towards (semi-) autonomous home energy management.

However, participants' comments suggest that accuracy of the predictions is essential in whether automation would be trusted and desired. With increasing autonomy, the requirements to being able to monitor and control the systems become more important. Challenges for the intelligent interface community include how the balance between user control and autonomy can be achieved in a flexible way without overwhelming the user with requests, and without risking undesired system actions.

Further, questions of ownership and stakeholders implicated with such an autonomous system arose during the study. Who owns the system and the data it generates and stores? Who has access to the data and the recommendations? Who is responsible if the system takes an undesired action? The user should be able to answer all these questions by accessing information through the UI, or by being notified by the system.

The insight still holds that, perhaps above all, intelligent systems need to be *intelligible* and *accountable* to be perceived as useful and trustworthy [3], particularly if the aim is that the system is adopted for use in everyday life.

CONCLUSIONS

This paper presented a system and evaluation of personalized energy-related recommendations based on household usage profiling. Based on user profiles, AgentSwitch provides comparisons of the user's current energy tariffs to the tariffs available on the market, and provides advice on how much the user can save by shifting detected deferrable loads (e.g., washing machine or tumble dryer) to off-peak times. We described the system architecture that drives the web service, including energy usage profiling based on time series of electricity consumption monitored in users' homes, deferrable load detection, and usage of up-to-date national energy tariff information to compute the recommendations.

The evaluation based on task-driven walkthroughs with 10 users with three months worth of monitored consumption data showed that AgentSwitch found cheaper tariffs for most of them (9/10). Despite potential annual savings of between £35-£390 interviews revealed a host of barriers to actually switching tariffs. As a result, we discussed interface issues that can be applied more broadly to overcome barriers to following recommendations. Moreover, anticipated regulatory as well as technological changes are needed, especially to create an environment in which load shifting can be encouraged, for example by devising persuasive systems not only to convey monetary incentives, but also to convey environmental benefits more effectively. Lastly, our evaluation supports the potential for semi-autonomous systems and we discuss some considerations relevant to intelligent UI designers, such as the need to balance user control and autonomy.

ACKNOWLEDGEMENTS

This work is supported by EPSRC grant EP/I011587/1.

REFERENCES

1. Baker, O.F., Khalil, I., and Kotsis, G. Agents for energy efficiency in ubiquitous environments. *Proc. iiWAS '09*, ACM Press (2009), 674-677.

2. Banerjee, N., Rollins, S., and Moran, K. Automating energy management in green homes. *Proc. SIGCOMM HomeNets '11*, ACM Press (2011), 19-24.

3. Bellotti, V. and Edwards, K. Intelligibility and accountability: human considerations in context-aware systems. *Human-Comp. Interact.* 16, 2 (2001), 193–212.

4. Braun, V. and Clark, V. Using thematic analysis in psychology. *Qual. Res. in Psych.* 3, (2006), 77–102.

5. Brynjarsdottir, H., Håkansson, M., Pierce, J., Baumer, E., DiSalvo, C., and Sengers, P. Sustainably unpersuaded: how persuasion narrows our vision of sustainability. *Proc. CHI '12*, ACM Press (2012), 947-956.

6. Bunt, A., Lount, M., and Lauzon, C. Are explanations always important?: a study of deployed, low-cost intelligent interactive systems. *Proc. IUI '12*, ACM Press (2012), 169-178.

7. Darby, S. Making sense of energy advice. *Proc. of ECEEE 2003*, (2003), 1217–1226.

8. MacKay, D.J.C. Sustainable energy — without the hot air. UIT Cambridge, Ltd, Cambridge, 2009.

9. Department of Energy and Climate Change (DECC) and the Office of Gas and Electricity Markets (OfGEM). *Smart Metering Implementation Programme*. 2011.

10. Diaconis, P. Bayesian numerical analysis. In S.G.J. Berger, ed., *Statistical Decision. Theory and Related Topics IV*. 1988, 163–175.

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

12. Frankcom, N. Household energy bills unaffordable in less than three years. 2012. http://bit.ly/LSE175. Accessed 09/01/2013.

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

14. He, H.A., Greenberg, S., and Huang, E.M. One size does not fit all: applying the transtheoretical model to energy feedback technology design. *Proc. CHI '10*, ACM Press (2010), 927-936.

15. Hledick, R. How green is the smart grid? *Electricity* 22, 3 (2009), 29-41.

16. Ji, T., Tuo-Yu, P., Zhen-Min, Z., and Kai, L. A personalized recommendation model based on the user-state awareness. *Proc. 1st IEEE Symposium on Web Society*, IEEE (2009), 10–13.

17. Kjeldskov, J., Skov, M.B., Paay, J., and Pathmanathan, R. Using mobile phones to support sustainability: a field study of residential electricity consumption. *Proc. CHI '12*. ACM Press (2012), 2347-2356.

18. LeMay, M., Haas, J.J., and Grunter, C.A. Collaborative recommender systems for building automation. *Proc. 42nd Hawaii International Conference on System Sciences*, IEEE (2009), 1–10.

19. Lima, J.C.D., Rocha, C.C., Augustin, I., and Dantas, M´.A.R. A Context-aware recommendation system to behavioral based authentication in mobile and pervasive environments. *Proc. 9th International Conference on Embedded and Ubiquitous Computing*, IEEE (2011), 312–319.

20. Moreau, L. and Missier, P. PROV-DM: The PROV data model. 2012. http://www.w3.org/TR/prov-dm/ Accessed 09/01/2013.

21. O'Donovan, J. and Smyth, B. Trust in recommender systems. *Proc. IUI '05*, ACM Press (2005), 167-174.

22. Parnell, R. and Larsen, O.P. Informing the development of domestic energy efficiency initiatives. *Env. and Behavior* 37, 6 (2005) 787-807.

23. Parson, O., Ghosh, S., Weal, M., and Rogers, A. Non-intrusive load monitoring using prior models of general appliance types. *Proc. of the 26th AAAI Conference on Artificial Intelligence*, AAAI (2012), 356-362.

24. Pierce, J. and Paulos, E. Beyond energy monitors: interaction, energy, and emerging energy systems. *Proc. CHI '12*, ACM Press (2012), 665-674.

25. Pu, P., Chen, L., and Hu, R. A user-centric evaluation framework for recommender systems. *Proc. RecSys '11*, ACM Press (2011), 157-164.

26. Ramchurn, S.D., Vytelingum, P., Rogers, A., and Jennings, N.R. Agent-based homeostatic control for green energy in the smart grid. *ACM Trans. Intell. Syst. Technol.* 2, 4 Article 35 (2011).

27. Ramchurn, S.D., Osborne, M.A., Parson, O., Maleki, S., Rahwan, T., Huynh, T.D., Reece, S., Alam, M., Fischer, J.E., Hines, G., Costanza, E., Moreau, L., Rodden, T., and Roberts, S. AgentSwitch: smart energy tariff selection using agent technology. *Proc. AAMAS '13*, IFAAMAS (2013).

28. Rial, A. and Danezis, G. Privacy-preserving smart metering. *Proc. WPES '11*, ACM Press (2011), 49-60.

29. Scott, J., Brush, A.J.B., Krumm, J., Meyers, B., Hazas, M., Hodges, S. and Villar, N. PreHeat: controlling

home heating using occupancy prediction. *Proc. UbiComp '11,* Springer (2011), 281-290.

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

31. US Department of Energy. Benefits of demand response in electricity markets and recommendations for achieving them. 2006.

32. US Department of Energy. The smart grid: an introduction. 2008.