A Field Study of Human-Agent Interaction for Electricity Tariff Switching

Alper Alan,[1] Enrico Costanza,[1] Joel Fischer,[2] Sarvapali D. Ramchurn,[1] Tom Rodden,[2] Nicholas R. Jennings[1]
[1]Agents, Interaction, and Complexity Group
University of Southampton
Southampton, UK
{ata1g1,ec,sdr,nrj}@ecs.soton.ac.uk
[2]Mixed Reality Lab
University of Nottingham
Nottingham, UK
{jef,tar}@cs.nott.ac.uk

ABSTRACT

Recently, many algorithms have been developed for autonomous agents to manage home energy use on behalf of their human owners. By so doing, it is expected that agents will be more efficient at, for example, choosing the best energy tariff to switch to when dynamically priced tariffs come about. However, to date, there has been no validation of such technologies in any field trial. In particular, it has not been shown whether users prefer fully autonomous agents as opposed to controlling their preferences manually. Hence, in this paper we describe a novel platform, called TariffAgent, to study notions of *flexible autonomy* in the context of tariff switching. TariffAgent uses real-world datasets and real-time electricity monitoring to instantiate a scenario where human participants may have to make, or delegate to their agent (in different ways), tariff switching decisions given uncertainties about their own consumption and tariff prices. We carried out a field trial with 10 participants and, from both quantitative and qualitative results, formulate novel design guidelines for systems that implement flexible autonomy.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Intelligent Agents

Keywords

Human-Agent Interaction; Autonomous Agents; Flexible Autonomy; Energy; Smart Grid

1. INTRODUCTION

The vision of the Smart Grid includes technologies that enable the efficient integration of intermittent renewable energy sources (such as wind or solar energy) and electric vehicles, and will reduce inefficiencies by allowing consumers to better manage how electricity is used, stored, and delivered [17]. One of the key underpinnings of this endeavour is the concept of the *smart meter* that aims to help users manage energy consumption in the home to minimise inefficiencies in usage and maximise the user's savings. One particular use of the smart meter would allow suppliers to offer different energy tariffs (i.e., energy pricing methods) every day, depending on fluctuations in demand and supply on the grid. However, such efficiencies will only be accrued if users pay attention to, understand, and react to these tariffs.

In this context, a number of agent-based solutions have recently been proposed to minimise the need for users to monitor real-time electricity prices (i.e., dynamic pricing) for different times of the day (via their smart meter) in order to react to changes in the price [9, 18]. However, most of these approaches have only been tested in simulation, and therefore, tend to simplify the problem significantly in a number of ways. First, most agent-based approaches assume users have predictable energy consumption profiles [1, 9]. However, studies have shown that most users' consumption may not at all be predictable (e.g., at half-hourly or even daily intervals) [16]. Second, they assume users will trust the agent to make the right decisions a day-ahead when the prices for the next day are not completely predictable and may change in real-time. Thus, the agent may make mistakes that may lead the user to distrust the system [4]. Third, they assume that users will take the trouble to understand and react to price changes regularly [8]. However, as shown in a recent report by US Department of Energy [17], consumers are reported to typically, spend at most two hours per year optimising their energy settings and many of them struggle to comprehend energy tariffs.[1] Thus, the traditional assumptions made in the design of agent-based systems for home energy management may render them impractical when deployed in the real-world, especially if one takes into account that such systems have the potential to disrupt users' comfort and have financial implications. In fact, to date, there are still no design guidelines derived from a real-world deployment to implement agents for home energy management systems.

Against this background, in this paper we present a novel approach to building and testing agent-based systems for real-world home energy management applications. In particular, we focus on a tariff-switching problem that is likely to become reality in the future:[2] users have to choose an electricity tariff that is most tailored to their consumption profile for the next day with the help of an autonomous software agent. This is a challenging problem because the users' consumption and tariffs may change every day and some tariffs may vary more the next day than others.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

[1]`http://www.bbc.co.uk/news/uk-22745110.`

[2]Day-ahead and real-time tariffs have been proposed both in Europe and US [3, 17].

Hence, we develop a platform to evaluate human-agent interactions in order to define the design guidelines for *flexible autonomy* (i.e., an agent's autonomy may be monitored and adjusted by its owner) in the context of tariff switching. The platform, called TariffAgent, allows users to interact with their tariff-switching agent in order to choose the best tariff for the next day. In particular, we advance the state of the art in the following ways:

1. We develop a novel web application that simulates dynamic pricing tariffs, and, more importantly, a tariff switching agent and interfaces to interact with the agent in order to help it choose the best tariff for the day-ahead for its owner. Our system allows us to evaluate how people perceive the benefit of the agent even if it may choose the wrong tariff at times (due to uncertainty in predictions).

2. We carried out a field trial of the platform over a period of two weeks and we demonstrate that users find the system useful and generally trustworthy (in spite of the savings made by the agent being minimal).

3. We present new design guidelines for agent-based home energy management systems.

In general, our work introduces a methodology, borrowed from the discipline of Human-Computer Interaction (HCI) [11], to build and test agent-based systems that goes beyond the simulation-based approaches typical to most previous work.

The rest of this paper is structured as follows. Section 2 provides some background on agent technologies for the energy management as well as eco-feedback technologies that have been studied in HCI. Section 3 describes our agent-based tariff switching model and the real-time pricing tariff we implement. Then, Section 4 presents our implementation of the agent and tariff schemes. The evaluation of the platform in a pilot study is then provided in Section 5 and the discussion of the key results is given in Section 6. The design guidelines derived from the results is presented in Section 7. Finally, Section 8 concludes.

2. BACKGROUND

Here we present previous work relevant to the design of autonomous agents for home energy management. In particular, we survey the multi-agent systems literature as well as the HCI literature as the work we present in this paper lies at the boundary of these two research areas.

Previous work in multi-agent systems within the energy domain has focused on algorithms to manage, among others, micro-storage [18], load deferral [9], energy exchange [1], or energy use within non-domestic buildings [6]. Moreover, closer to our work, Ramchurn et al. [8] presented a deployed recommender agent that provides energy tariff suggestions based on accurate predictions of energy consumption and the identification of deferrable loads. However, none of these approaches investigated the human-agent interaction issues that arise when they are deployed as they either assume that agent owners have well-defined linear utility functions or will adopt any schedule or tariff suggested by an agent.

Instead, HCI research within the energy domain has focused on studies of "visualisations" of energy monitoring [7]. While these studies have shown that energy monitoring alone can help consumers save from 5 to 15 per cent, they also show that only environmentally-motivated users are likely to exhibit conservation behaviours [2, 5]. These behaviours may fade over time as most of these technologies were missing social, cultural, technical, and institutional aspects of the energy use [13]. However, apart from just one study by Rodden et al. [10], HCI research has, so far, typically ignored autonomous agent technologies described above and what interactional issues these will result in when users are faced with automated load deferral or tariff switching for example [7]. In contrast to the lab-based approach taken by Rodden et al. [10], our work is based on a real-world study.

In particular, we develop a system specially designed to investigate interactional issues within the domain of automated tariff-switching [8]. Our work builds upon the seminal works that defined the notion of adjustable autonomy, whereby its expert human owner can adjust the autonomy of an agent [14, 15]. However, while they focused on *algorithms for transfer-of-control* (that, in short, define whether the human owner will take control of the system at certain fixed points in the operation of the agent), we instead do not constrain *when or at what points* such transfers of control should happen. Rather, as we show in Section 4, we design the interactions with TariffAgent to be flexible to the (non-expert) user's preferences to control the agent at any point in time and fall back to automatic and default choices when user input is not available. Hence, we cast our approach within the paradigm of *flexible autonomy* and define novel design guidelines for agent-based systems that espouse this paradigm.

3. THE TARIFF SWITCHING PROBLEM

Here we present a model of the tariff switching problem which we then use to formulate the agent-based strategies to switch the user to different tariffs (in Section 4). In our scenario, we consider a daily tariff switching problem in order to be able to create a realistic user study (as depicted in Section 5).[3] Hence, we are only concerned with the *daily energy usage* of a home in terms of choosing the best tariff to go for. Thus, let the energy use of a given home on day d be denoted as $c_d \in \Re^+$ kWh, where $d \in \{0, \cdots, D\}$, and $D \in \mathbb{Z}^+$. In what follows, we elaborate on how this daily consumption (and predictions of it) will be billed by different types of tariffs and hence define the challenge of choosing between these tariffs.

3.1 Tariff Specifications

Let the set of tariffs provided by suppliers in the retail energy market be denoted as $t_1, ..., t_s \in T$ and for each tariff, there exists a function $F : T \times \Re \rightarrow \Re$ that takes the predicted energy consumption for the next day c_{d+1} of the home and returns the predicted cost for that tariff. For example, given a standard tariff from typical UK retailers where a customer is charged a fixed rate r_1, function F would return $c_{d+1} \times r_1$.

The rate r_1 is typically chosen by the supplier to mirror the risk it incurs in the wholesale energy market (e.g., when consumption increases unexpectedly, it may have to buy at a high price, and if consumption is lower than expected, it sells at a lower price than it paid for the energy a day-ahead) as well as standard costs for metering the customer [18]. In future, however, we expect rates to change on a daily basis as suppliers increasingly rely on renewable energy generators. This is because, the output from renewable energy generators will be significantly dependent on the weather (e.g., solar or wind), and therefore, the amount available to power all the homes contracted to a given supplier may not be sufficient. To meet shortfalls in real-time, the supplier will, in turn, have to buy energy at a higher price from other (possibly non-renewable) producers in the energy market and therefore pass such costs to its consumers.

[3] This scenario could be extended to consider hourly tariff changes but this is beyond the scope of this paper.

To account for such a scenario, we therefore consider that the home can choose its tariff only from two types of suppliers P_{std} and P_{wind} with different characteristics as follows:

1. P_{std} suppliers sell electricity at a standard rate r_{std} that is constant across all days $d \in D$. They are able to do so because they buy energy from coal-fired power stations that can ramp up or down their production as and when needed. Thus the cost to the customer for any day d is:
$$cost_{fixed} = c_d \times r_{std} \tag{1}$$
2. P_{wind} suppliers sell electricity at a variable rate that may change every day and the rate for day $d+1$ is set on day d. This type of supplier buys its day-ahead energy from a wind generator that has different outputs every day. To meet possible shortfalls in supplier P_{wind} also needs to buy energy at extra cost in the real-time market and therefore passes such costs to its consumers. Thus, P_{wind} suppliers advertise two rates r_{wind} and r_{nowind} where $r_{wind} << r_{nowind}$ and r_{wind} is the rate charged for any $c_{d+1} \leq c'_{d+1}$ and r_{nowind} is charged for any demand $c_{d+1} - c'_{d+1} > 0$ respectively (we express this in detail in the next section). The value c'_{d+1} reflects the amount a given supplier can afford to supply to each of its customers. However, both c_{d+1} and c'_{d+1} can only be estimated to some degree of accuracy on day d by the user and supplier and hence we denote them as $\hat{c}_{d+1}$ and $\hat{c}'_{d+1}$ respectively.

The fact that $\hat{c}'_{d+1}$ may not be accurately predicted poses a problem for the user as it may mean she will incur much higher costs if $\hat{c}'_{d+1}$ is poorly predicted to be high. We therefore detail this challenge in the next section.

3.2 Choosing a Tariff

The user is faced with a tariff comparison problem involving the two types of tariffs detailed above. First, we need to compute $cost_{fixed}$ as per Equation (1) and then compute a *predicted* cost for the wind based tariff from P_{wind}. To this end, the user would need to compute an expected demand for the next day, $\hat{c}_{d+1}$ and use the predicted consumption threshold from the supplier, $\hat{c}'_{d+1}$, in order to compute the *expected cost* from the wind based rates as follows:

$$\begin{aligned} cost_{variable} = {} & min(\hat{c}_{d+1}, \hat{c}'_{d+1}) \times r_{wind} \\ & + max(0, \hat{c}_{d+1} - \hat{c}'_{d+1}) \times r_{nowind} \end{aligned} \tag{2}$$

Using $cost_{fixed}$ and $cost_{variable}$, the *predicted* cheapest tariff can be identified for the day-ahead. However, as can be seen from Equation (2) the decision to go with the variable tariff is fraught with uncertainties. In particular, there are two dimensions to this uncertainty:

1. Personal uncertainty: in predicting the user's own consumption for the next day. Apart from days when the user knows she will not be at home (and therefore consume only minimally), other days may involve activities, which are unpredictable (e.g., visits by friends for dinner, running out of clean clothes and therefore needing to do a wash).
2. Environmental uncertainty: the availability of wind energy is weather dependent and hence, while hourly predictions may be reasonably good, day-ahead predictions are likely to be inaccurate.

Hence, to alleviate the need for the user to make sense of these uncertainties, compute the cost of each tariff for the next day, and decide on which one to switch to, we develop a simple tariff switching agent called TariffAgent that takes over this burden from the user.[4] In particular, we are interested in how users will use such a tariff switching agent. The three key questions we wish to investigate are: (i) are users comfortable with an agent making decisions with financial consequences on their behalf? (ii) are users trusting of an agent autonomously predicting their consumption or do they want to control its predictions and choices? (iii) how do we engineer a tariff switching agent that is acceptable to users and does not intrude upon users' daily activities? We believe that answering these questions is central to the development of design guidelines for systems involving human-agent interactions, particularly when such interactions take place within such intimate settings (i.e., relevant to the users' own energy use) and lead to decisions with financial consequences (i.e., where penalties may be incurred if the agent makes the wrong choice). Hence, in the next section we describe the architecture of TariffAgent and then go on to evaluate it empirically in a field trial in Section 5.

4. TARIFFAGENT

TariffAgent is developed as a web application using current web technologies (django/python for the backend, mysql database, and HTML/Javascript). The choice of a web-based agent is motivated by the fact that future tariff switching scenarios will rely on users (through their smart meters) accessing tariff information and interacting with suppliers over the Internet. Moreover, this setup will allow agents to gather weather forecast data to predict the output of renewable energy sources. However, for the purposes of this work, TariffAgent incorporates both live features (in terms of real-time energy monitoring and tariff switching interactions) and simulated features (in terms of wind energy prediction) in order to embed the users into a realistic test environment. Thus TariffAgent involves the following key modules:

1. Usage Monitoring and Prediction: this module collects data from the live energy monitoring device installed in the home.
2. Wind-Energy Predictor: this module retrieves actual and predicted (i.e., c'_d and $\hat{c}'_{d+1}$) wind energy values from a real-world dataset.
3. Tariff Switching Interfaces: this is a set of web pages and notification mechanisms that give the user information and control over the actions of TariffAgent.

In the remaining subsections, we detail each of the above modules.

4.1 Usage Monitoring and Prediction

In order to monitor the energy use of a home we rely on off-the-shelf home energy monitoring devices. In particular we use kits by AlertMe[5] since they provide the following useful features. First, a non-expert can easily install these kits by simply connecting a current clamp to one of the mains electricity cables. This clamp contains a transmitter that provides readings at regular intervals (seconds). Second, the data from the clamp is sent to AlertMe data servers using an ethernet hub connected to the broadband router of the home. Third, the collection of the data from the AlertMe server is facilitated through an open API that allows TariffAgent to pull both power and energy readings (for a given period) at regular

[4] Instead of a complex tariff switching agent such as AgentSwitch [8], we simplify the computations here to make it easier for users to model the reasoning of the agent.

[5] `http://www.alertme.co.uk.`

intervals. Given these readings, TariffAgent can then predict the user's consumption for the day-ahead.

As was shown in a previous study, applying complex machine learning techniques to predict day-ahead usage accurately is a challenging problem [16]. While some houses may be very predictable (e.g., a family home with both parents working during daytime or a single elderly person with a regular pattern of behaviour), other homes may be more liable to variations on a daily basis (e.g., houses of sharers or users who work at different hours of the day). Moreover, even if some residents have very regular lifestyles there will always be days when high-power devices will be used unpredictably (e.g., heavy washes or the use of the oven). Here we do not aim to test the user's trust in the accuracy of predicted consumption (since this is likely to be low when using state-of-the-art algorithms in any case) and instead focus on trust in the agent when it may make mistakes. Thus, in TariffAgent we use a very simple prediction algorithm to predict the day-ahead consumption on day d that simply uses the previous day's consumption as a prediction for the next day's consumption (i.e., $\hat{c}_{d+1} = c_{d-1}$). To supplement this prediction, we also build in a mechanism (as we show in Section 4.3) to let the user *fine-tune the prediction* made by TariffAgent (and this was shown to help users understand their energy consumption better as we show later in Section 5).

4.2 Wind Energy Prediction Simulation

The goal of TariffAgent is to simulate the availability of wind energy from suppliers of type P_{wind} both for the current day (i.e., actual wind energy) and for the day-ahead (i.e., predicted values $\hat{c}'_{d+1}$). To this end, we collected wind speed data from a third-party weather forecast service[6] for 28 days from regions where wind turbines are located in the UK and are used by Ecotricity (a renewable energy provider).[7] For each of these days, we collected the day's hourly wind speed data as well as the hourly predictions of wind speed for the day-ahead. These wind speeds are then aggregated and used to create a simulated wind energy values c'_d and $\hat{c}'_{d+1}$ that is calibrated for individual users. This is important because each individual user of TariffAgent may have different daily electricity consumption levels, so we need to calibrate the wind energy values to ensure users' consumption may, at times, be higher than the available wind energy (predicted or actual), and, at other times, be within the available wind energy (predicted or actual). Our aim is to introduce uncertainty in the system, so that the agent would at times provide correct suggestions, but at times incorrect ones because of incorrect wind predictions and mismatch between available wind energy and users' consumption. Thus, given the average consumption of a home c_{avg}, we compute the available wind energy as follows:

$$\hat{c}'_{d+1} = \frac{\hat{w}_{d+1}}{\delta} - \frac{k \times c_{avg} - 1}{2} \tag{3}$$

where $\hat{w}_{d+1}$ is the average wind speed predicted for the next day, $k = 0.9$ and $\delta = 20$ are factors that were empirically determined to align the wind and energy consumption values. Essentially we shift the wind probability distribution so its average and range roughly match the probability distribution of the energy consumption. This may make it particularly difficult for the user to perceive the benefit of the agent as the savings are likely to be small.

At the start of the trial, TariffAgent is set to start from day zero of the wind energy prediction computations. It is also important to note that, even though *realised* wind energy values are stored on the same platform, these are never passed on to the tariff switching algorithm used by TariffAgent.

4.3 Tariff Switching Interactions

The tariff switching interactions are at the centre of our study of flexible autonomy. In particular, we provide a number of interaction modalities that allow us to capture different levels of autonomy that go beyond the simple notion of moving between human controlled and fully autonomous tariff switching. We designed the system to offer different levels, noted as $S1$ to $S3$, of information provisioning and control as follows:

- S1 (human-guided): TariffAgent informs the user, only once a day, using SMS reminders when it believes the tariff set for the next day should be switched. This suggestion is only made if the user has not set the best tariff (as computed by TariffAgent using predictions of wind energy and electricity consumption) for the next day. The tariff changes are made by users *manually*. This setting is the fully human-guided system but with an agent autonomously making predictions of the user's consumption (as specified in Section 4.1). The user can authorise the switch using the SMS or manually action it. Moreover, in this setting, the user cannot disable such notifications.
- S2 (semi-autonomous): TariffAgent automatically switches the user to the best tariff and notifies the user of the change by SMS. Then the user has to manually go to the web application to change it if she wishes to. This setting can be seen as *semi-autonomous* in that it automates the tariff switching process (and usage prediction) but also allows the user to change the tariff chosen by first notifying her of the change.
- S3 (fully autonomous): Same as S2 but TariffAgent never notifies the user. This is the fully autonomous setting in our case as the user completely offloads the burden of tariff switching to TariffAgent.

For all tariff switching suggestions and those autonomously enacted, TariffAgent uses the functions defined in Section 3.2 to determine the cheapest tariff for the day-ahead. Moreover, a user, under any of the above setting, receives a daily report on her performance (with the help of TariffAgent) for the previous day. This report includes information on how much energy she consumed, its cost, and how much was saved (by choosing the cheapest tariff) or lost (by choosing a tariff other than the cheapest one). Finally, if the user has chosen a tariff (for the next day) from a supplier of type P_{wind} (i.e., a wind-dependent tariff), she will be sent a notification if her predicted consumption for the next day goes above 80% of the predicted wind energy available (i.e., $\hat{c}_{d+1} > 0.8 \times \hat{c}'_{d+1}$). This notification is only sent once a day and helps make the user aware of the uncertainties in her own consumption and the predicted wind energy.

Through the users' interactions with the above settings (as we show later), we expect to understand to what degree users trust the system to act autonomoustly. In what follows, we detail the tariff rates that make the tariff switching problem challenging enough to require users to use TariffAgent and then describe the interfaces that allow users to interact with TariffAgent to either set the above autonomy levels, gather information about the system, and also interact with the system to choose the best tariff suited to their needs.

4.3.1 *Tariffs*

To create challenging tariffs for users of TariffAgent (and therefore incentivise them to delegate their tariff switching problem to

[6] http://www.weatherunderground.com.
[7] http://www.ecotricity.com.

Table 1: Tariffs in p/kWh.

Tariff	r_{wind}	r_{nowind}	Risk Level
Variable-A	3 p/kWh	23 p/kWh	High
Variable-B	8 p/kWh	18 p/kWh	Medium
Variable-C	10 p/kWh	16 p/kWh	Low

an agent), we assume there exists one supplier of type P_{std} that charges the user a fixed 15 p/kWh (pence per kilowatt-hour) every day and three different suppliers A, B, and C of type P_{wind}, each with their own wind-based tariffs as described in Table 1.

The only difference between each variable tariff is in terms of r_{wind} and r_{nowind} where combinations of these values represent different level of risks. In other words, the gap between the minimum rate and maximum rates offered varies for each variable tariff. Tariff Variable-A is a high risk tariff because the user may incur a very high cost (23 p/kWh compared to 3 p/kWh – an increase of 20p) in case $\hat{c}'_{d+1}$ is lower than predicted while Variable-C (10p/kWh presents a relatively low risk in such circumstances (16 p/kWh compared to 10 p/kWh – an increase of 6p). Variable-B is considered medium risk in comparison to the other two tariffs since the extra cost of a poor prediction lies in between the two (i.e., 10p).

As can be seen, deciding between these three tariffs is not easy unless the user is able to accurately predict her own consumption and wind energy levels for the next day. Hence, in the next section, we present the interaction modalities that allow users to transfer control to TariffAgent to make such complex decisions.

4.4 User Interface

Users can interact with TariffAgent either through SMS, as described in the previous section, or through a Web Site that represents the main "face" of the system. The site includes two pages: *Home* and *Details*, which are described in what follows.

4.4.1 Home View

The Home view is the first page that users encounter when they log into the system (Figure 1). It is composed of three key elements (each visually enclosed in a box):

1. Tariff (selection): shows the user's selected tariffs for the current and the next day, the agent's prediction of the user's consumption, and the agent's suggested tariff out of the four that are available (based on the predicted consumption and wind energy prediction). The tariff for the current day cannot be changed but the tariff for the next day can be chosen out of any of the four tariffs listed at the bottom of the box. This section also contains a checkbox to allow the user to let the agent know about their predicted consumption for the next day, so that the user can help TariffAgent in its choice of the best tariff.

2. Setting (autonomy levels): gives the options to the user to switch between any of the autonomy levels presented in the previous section, that is S1, S2, or S3.

3. Budget (status): shows how much remains of the budget allocated at the beginning of the study (see Section 5 for more details). It also provides a link to a page that has many more details on the historical data collected by the system for those users interested in analysing their past data.

As can be seen, the Home view allows the user to *manually control* TariffAgent (by setting the tariff and consumption predictions).

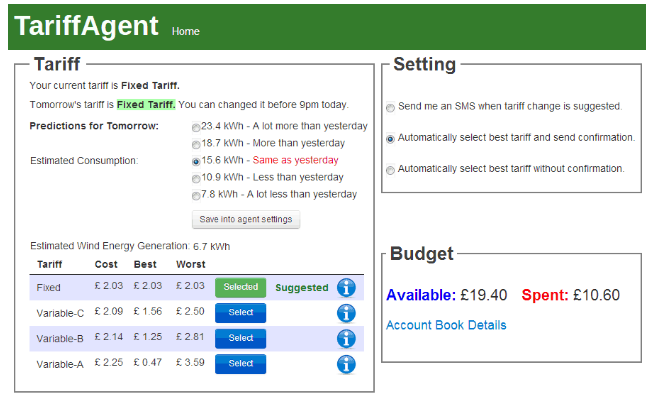

Figure 1: Home View

More importantly, the user can get direct feedback on her best suggested tariff given her energy consumption prediction. By so doing, the user can start understanding how the agent is using her predicted consumption to make a better choice on her behalf and therefore inspire confidence in the system. In the next section we present the other view in the system that allows users to probe even deeper into their operation of the system (with or without human control).

4.4.2 Details View

The Details view (Figure 2) is designed to provide daily historical energy cost information. This allows users to evaluate the TariffAgent's or their own success at choosing the right tariff in the past. The table sorts the detailed information by date, placing the most recent one on top. A daily energy cost information includes estimated consumption and estimated wind energy values, actual consumption and actual wind energy data, the suggested and selected tariff names, total energy price and the information of monetary saving or loss. Users can also check the accuracy of the agent's predictions of consumption and the external wind energy provided. In order to facilitate the user's understanding of the information provided to her, table cells are colour-coded.

5. EVALUATION

Our aim was to study how people interact with a system, such as TariffAgent, that is based on an autonomous agent, that can affect their financial savings or losses and has potential to disrupt their daily routines. We believe that for results of such a study to be meaningful it is important to provide a high level of realism, or in more formal terms a high degree of *ecological validity*.

We therefore decided to deploy a fully functioning prototype of TariffAgent in the field. The trial involved participants installing the system in their own homes and using it for a period of 14 days. The system used the real electricity consumption data from the participants' homes, recorded and transmitted through the AlertMe system, and this data was used to calculate their daily energy cost, based on the wind data and the selected tariff (as detailed in the previous section). To let participants experience the situation of an autonomous system affecting their money, the study included monetary incentives based on performance according to the following setup. Participants were allocated an online budget of £30 at the beginning of the study from which their consumption cost is taken over the duration of the trial. At the end of the study, participants were rewarded with the amount of money left in their budget (in the form of a shopping voucher), both as an incentive to engage

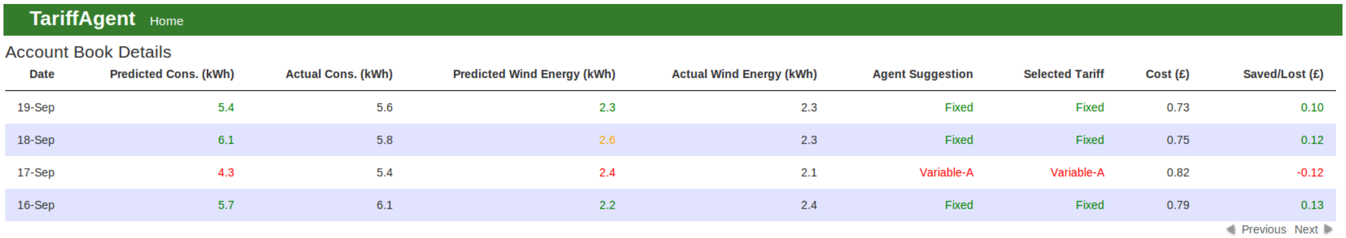
TariffAgent Home

Account Book Details

Date	Predicted Cons. (kWh)	Actual Cons. (kWh)	Predicted Wind Energy (kWh)	Actual Wind Energy (kWh)	Agent Suggestion	Selected Tariff	Cost (£)	Saved/Lost (£)
19-Sep	5.4	5.6	2.3	2.3	Fixed	Fixed	0.73	0.10
18-Sep	6.1	5.8	2.6	2.3	Fixed	Fixed	0.75	0.12
17-Sep	4.3	5.4	2.4	2.1	Variable-A	Variable-A	0.82	-0.12
16-Sep	5.7	6.1	2.2	2.4	Fixed	Fixed	0.79	0.13

Previous Next

Figure 2: Details View

with the system, and to make the saving have an actual, tangible impact on participants. The idea of using monetary incentives to simulate dynamic energy pricing is in part based on an early study in which participants received payments of the value of electricity saved [12].

Data was collected through different methods, both quantitative and qualitative. System logs were automatically collected from TariffAgent, documenting (i) when participants accessed the website, (ii) when they provided consumption information to the agent, (iii) when they manually selected a tariff, (iv) when they changed the agency settings and (v) when they replied to SMS messages. This data allows us to observe, in a basic way, the interaction of participants with the system (e.g., to see whether any of the settings is more popular than the others, or how often participants looked at a specific page). However, such data alone does not provide enough information for a more general analysis of human-agent interaction, for example information about *why* participants would use the system in a certain way.

Therefore, to try and understand how participants perceived the system and whether and how they adapted their daily routines to the novel technology we complemented the quantitative data collection with semi-structured interviews, conducted at the end of the trial, after each participant had interacted with TariffAgent for 14 days. Semi-structured interviews are part of a methodology commonly used in HCI and originally borrowed from Social Sciences [11]. These are interviews that aim to cover a set of topics and questions, but that do not rigidly follow a pre-defined scheme, resembling, from a participant's point of view, more the flow of a conversation. Our interviews addressed participants' usage of the system, and any possible changes in their use of electricity during the study.

It is worth emphasising that rather than asking questions like "did you like the system" or "what do you think about the system", normally in semi-structured interviews the aim is to induce participants to reveal information in a more indirect way, with the expectation to obtain richer data. For example, by asking questions like "how would you describe to someone else what TariffAgent does" or "who do you think the system works for (e.g., for you, for the energy company, and for the environment) and why?" we aim to uncover the participants' perception and understanding of the system, as it is exposed through the UI (both web and SMS). It is important to underline how this method relies on the fact that participants used a working prototype of the system, and therefore it aims to expose their actual experience, rather than abstract opinions.

5.1 Participants

Participants were recruited through personal contacts and snowball sampling. The recruitment criteria were for participants to have a broadband internet connection, basic knowledge of internet use, not to have been involved in energy-related research before and to live in flats or houses where the AlertMe monitoring kit could be installed without the intervention of a certified electrician.

A total of 10 volunteers, all adults and living in Southampton, took part in the study. Half of the participants were members of the university, 4 PhD students (Law, Education, 2 from CS) and 1 postdoc researcher. The other half of the participants were local professionals working in different sectors (Law, HR, administration, procurement and supply). None of the participants had previous experience with energy monitoring devices or applications.

6. RESULTS

To analyse the qualitative data collected through the interviews we adopted an approach inspired by grounded theory, a qualitative analysis method characterised by an absence of predefined assumption [11]. The interviews were documented through audio-recording (later partially transcribed) and notes; analysis started by categorising the material at the sentence level through open codes. Initially 41 open codes were used, later grouped in 6 broader categories: "usage patterns", "system perception: agency & control", "energy & environment", "usability", "study aspects" and "long term". The most relevant ones in the context of this paper are described in the following subsections, in some cases together with the relevant quantitative data.

6.1 Usage Patterns

All participants reported using the SMS messages to keep track of the system. Everyone liked the daily SMS notifications, and these were often described as the preferred way to keep track of the system. Nobody found them intrusive or too frequent. In addition everyone accessed the web interface with some regularity, at least once every 5 days and on average once every 2.2 days (see details in Table 2). Most participants accessed the TariffAgent web interface through desktop computers, only 3 accessed the system from smart phones, from the interviews this appears to be due to most participants not owning a smart phone connected to the Internet during the trial. The *Home page* of TariffAgent was accessed more frequently, with 277 page loads over the entire trial (on average 27.7 times per participant), while the *Details page* was loaded overall 210 times (on average 21.0 times per participant), still accounting for about 43% of the page views.

Table 2: Frequency of interactions with TariffAgent.

Action	Frequency
Loading home view	277
Loading details view	210
Saving consumption level	26
Selecting tariff from website	20
Changing system settings	18
Replying to SMS as Yes	10
Not replying to SMS	5

Table 3: Frequency of user activities during the trial.

User	Manual corrections to consumption prediction	Manual tariff selections	Changes to autonomy setting
Ender	10	3	4
Maria	2	3	5
Ivan	8	5	5
Omar	0	1	0
Mehmet	0	0	0
Louisa	1	1	1
Greta	0	0	1
Claudia	2	3	0
Alisa	3	4	0
Sinha	0	0	0

Settings & Tariffs.
Only 5 participants modified the system settings and used the S2 setting (where the agent automatically selects the predicted optimal tariff and sends confirmation via SMS) for a maximum of 14 days, and 29 days in total. The remaining 5 participants kept the default settings S1, where the agent sends suggestions via SMS, but does not automatically change the selected tariff. None of the participants in our trial selected settings S3 where the agent automatically changes the tariff without sending any confirmation to users. In terms of tariff selection, out of the 8 participants[8] who received tariff switching suggestions from the agent (while on setting S1), 5 accepted them by replying 'yes' via SMS at least once, the other 3 participant never accepted any tariff switching suggestion. Overall, 15 suggestions were sent from the system to the users, with a 66% acceptance rate. These values indicate that participants generally had trust in the system, despite the uncertainty and incorrect suggestions caused by the predictions. This trust was confirmed in the interviews, by all except one participant. All 5 participants who used setting S2 took advantage of the web UI to provide manual estimates of their electricity consumption prediction for the following day. In total this explicit input was provided 26 times over the course of the study. Overall our participants spent £134.5 over the entire study, on average £13.45 per participant.

6.2 System Perception: Agency & Control

All participants reported that they perceived TariffAgent to be working *for them*. Sometimes this was seen to be in contrast to the system working for an energy company, for example they told us:

> *I think mainly what I thought it was for my benefit, so I can actually see what is going on. So if it is telling me I can save money in terms of my bills then it is definitely not directly the company's advantage, 'cause they would want me to pay more, they want to increase their sales... [Omar]*

> *It does not seem to work for a company, 'cause it is offering higher flexibility [in choosing a tariff], it offers me to choose more options so it seems that it tries to help me [Ivan]*

Other times the system was reported to be working for the end user *as well as* for energy companies, and also for the environment:

> *..[it] saves me money, gives me the opportunity to save money, really [...] because this one takes into account renewable sources of energy available, I mean it works for everybody as it is. For the company that provides the energy, for me 'cause I get to choose the tariff so it is a good deal... in the end company sets the price right not me [Louisa]*

> *I think it is quite beneficial for the user because obviously you know and you have the control over how much you use and you can monitor how much you use. It also benefits the company 'cause it is almost like a self-service. [...] and definitely the environment 'cause it sort of raises the awareness [...] benefits environment as well as it increases awareness to reduce energy use. [Greta]*

It is interesting to note that, even though autonomy is one of the dominant features of TariffAgent, the utility provided by the system is often reported in terms of *control* ("gives me the opportunity to save money"). However, sometimes, as in the second quote above, the system is also described in terms of automation ("self-service").

This feeling of control was further illustrated by the fact that participants, when prompted to tell us about how cost was calculated in TariffAgent, talked mostly about consumption, rather than wind or tariff selection:

> *..it tracks how much electricity we use during a day,.. and then suggests whether I should keep that the tariff for the next day, or change the different tariff, mmm so that is pretty much all I 've got [laughing].. I guess it suggests sort of trying to see a pattern.., everyday we are gonna use similar amount, I imagine it goes on the day before more.. yeah I do not know 'cause I do not see how it suggested moving on somedays why did it suggest it 'cause it is not like being huge dramatically different each day [consumption].. [Claudia]*

More interestingly, participants' interviews suggest that while correct decisions made by the agent clearly had a positive impact on the user's trust in the agent, participants were tolerant to incorrect decisions:

> *It suggested me fixed tariff and I did not want to choose that and I lost money. In that case I learnt that I should trust the system. [Ender]*

Energy consumption, and its declared prediction, is what participants can take action on (they have no control over wind), what they have control over. At the same time, autonomous agency was sometimes explicitly reported as a benefit, for example:

> *..it is an easy way to change tariff without me having to put a lot of attention. because sometimes I think I need to revise the contract of the energy, but it is so difficult sometimes to understand if it is better, that sometimes I give up and I just continue with the current one. so it is useful to do that [Maria]*

The same participant described to us how she took advantage of the opportunity to instruct the agent, overriding the consumption prediction:

> *..this weekend it was obvious, 'cause I was not here.. I decided to change consumption to be much less and ..mmm.. also I changed on friday 'cause on saturday I use more energy, usually than during the weekdays, 'cause it is the day put the washing machine and I stay more hours at home.. [Maria]*

This quote demonstrates that being aware of agency is not at odds with keeping in control.

[8] Two participants never received switching suggestions because they happened to be already on an optimal tariff.

7. DESIGN GUIDELINES

Based on the results of our evaluation, we list here design guidelines for human-agent interaction:

1. Provide an easy way for users to receive updates about the status and operation of the agent.
2. Enable users to instruct the agent by offering them opportunities to declare their plans[9] and integrate these plans into the agent's operations.
3. Leave the system open to transfer of control by allowing users to adjust the system's level of autonomy (i.e., when to release or retain agency).

The first guideline is based on the observation that in our field trial participants were very keen on keeping track of the agent operations. None of them disabled the SMS notification, instead everyone reported to find them useful. Moreover, the web access logs illustrate that participants visited the detailed information page quite frequently, about 43% of all page views were on it, illustrating a desire to monitor in detail what the system is doing.

The second guideline builds on the perception of being *in control* that our participants reported in the trial. This feeling of control was related to them inputting into TariffAgent their predicted consumption for the following day, but also to some adjustments of the participants schedule. Indeed in some cases this action took the form almost of "booking" their activities into the system. This is in contrast to existing trends in mixed initiative systems [15], where it is part of the system's functionality to decide when to attempt and transfer control to users. In contrast we suggest that systems involving humans and agents, or so-called *human-agent collectives*, should be left open enough so that users can decide when to intervene. It should not be necessary to express this operation as "removing" or "diminishing" agency from the system, indeed in TariffAgent the optional user input provides more information for the autonomous system to help them to save money.

The third guideline is based on the fact that half of our participants used the autonomous system setting of TariffAgent (S2), but half of them kept the manual confirmation. Moreover, of those who took advantage of automation, two reverted to manual mode (S1) after some time. It should be emphasised that this manual setting does not correspond to simply turning the system agency off completely, indeed the system still continuously monitors consumption and it autonomously offers suggestions for when to change tariff. However, the user needs to explicitly accept such suggestions, before they are turned into practice.

8. CONCLUSIONS

In this paper we have applied, for the first time, an existing HCI methodology to study human-agent interactions and to investigate the concept of flexible autonomy within the domain of home energy management. In particular, we developed the TariffAgent platform that embeds participants in a tariff switching problem scenario where they need to interact with an agent to choose the best tariff based on predictions of their own (real) energy consumption and (simulated) wind energy available. The results of our pilot study run over two weeks with ten participants show that users are ready to give away control to TariffAgent to switch their tariff but were keen to monitor the performance of the agent. From our analysis, we formulate a number of design guidelines for flexible autonomy in terms of easy of use, control over agency and transfer of control.

Moreover, we find it surprising and intriguing that participants chose to monitor and manually confirm the agent's action around energy tariffs. In fact, this finding is in contrast to what reported in prior literature: that people want to know as little as possible about energy tariffs [17], which led us to expect that at least some participants would turn off the SMS notifications and simply rely on the agent. Perhaps participants felt that the extra effort they were required to put in monitoring the agent and feeding it data was a good investment taking into account the extra work that the agent does for them. On the other hand it could be argued for this result to be an artefact due to the study setup (in particular the limited duration and the introduction of a novel system). However, we believe, at the very least, it entices further research in this area.

[9] e.g., how much energy they are going to consume the next day

Acknowledgements

This work was partially supported by the EPSRC ORCHID project (EP/I011587/1) and through a PhD scholarship from the Turkish Government.

9. REFERENCES

[1] M. Alam, S. Ramchurn, and A. Rogers. Cooperative energy exchange for the efficient use of energy and resources in remote communities. In *Proc. AAMAS*. ACM, 2013.

[2] S. Darby. *The effectiveness of feedback on energy consumption: A review for DEFRA of the literature on metering, billing and direct displays.* Environmental Change Institute, University of Oxford, 2006.

[3] DECC. Smarter grids: The opportunity. Technical report, Department of Energy and Climate Change (DECC), 2009.

[4] R. R. Hoffman, M. Johnson, and J. M. Bradshaw. Trust in automation. *IEEE Intelligent Systems*, 28(1), 2013.

[5] T. G. Holmes. Eco-visualization: combining art and technology to reduce energy consumption. In *Proc. Creativity & Cognition 2007*. ACM, 2007.

[6] J. Kwak, P. Varakantham, R. Maheswaran, Y. Chang, M. Tambe, B. Becerik-Gerber, and W. Wood. Tesla: An extended study of an energy-saving agent that leverages schedule flexibility. *Journal of Autonomous Agents and Multiagent Systems*, 2013.

[7] J. Pierce and E. Paulos. Beyond energy monitors: interaction, energy, and emerging energy systems. In *Proc. CHI*. ACM, 2012.

[8] S. D. Ramchurn, M. A. Osborne, O. Parson, T. Rahwan, S. Maleki, S. Reece, T. D. Huynh, M. , J. E. Fischer, T. Rodden, L. Moreau, and S. Roberts. Agentswitch: Towards smart energy tariff selection. In *Proc. AAMAS*. ACM, 2013.

[9] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *AAMAS*. ACM, 2011.

[10] T. A. Rodden, J. E. Fischer, N. Pantidi, K. Bachour, and S. Moran. At home with agents: Exploring attitudes towards future smart energy infrastructures. In *Proc. CHI*. ACM, 2013.

[11] Y. Rogers, H. Sharp, and J. Preece. *Interaction Design: Beyond Human-Computer Interaction*. Wiley, 2011.

[12] R. E. Slavin, J. S. Wodanski, and B. L. Blackburn. A group contingency for electricity conservation in master-metered apartments. *Journal of Applied Behavior Analysis*, 14(3), 1981.

[13] Y. Strengers. Designing eco-feedback systems for everyday life. In *Proc. CHI*. ACM, 2011.

[14] M. Tambe, E. Bowring, J. P. Pearce, P. Varakantham, P. Scerri, and D. V. Pynadath. Electric elves: What went wrong and why. *AI Magazine*, 29(2), 2008.

[15] M. Tambe, P. Scerri, and D. V. Pynadath. Adjustable autonomy for the real world. *Journal of Artificial Intelligence Research*, 17(1), 2002.

[16] N. C. Truong, J. McInerney, L. Tran-Thanh, E. Costanza, and S. D. Ramchurn. Forecasting multi-appliance usage for smart home energy management. In *Proc. IJCAI*. , 2013.

[17] US Department of Energy. Grid 2030: A national vision for electricity's second 100 years, 2013.

[18] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid. *Journal of Artificial Intelligence Research*, 42, 2010.