Managing Energy Tariffs with Agents: A Field Study of a Future Smart Energy System at Home

Alper T. Alan
Agents, Interaction and Complexity Group
University of Southampton
Southampton UK
ata1g11@ecs.soton.ac.uk

Enrico Costanza
Agents, Interaction and Complexity Group
University of Southampton
Southampton UK
ec@ecs.soton.ac.uk

Sarvapali D. Ramchurn
Agents, Interaction and Complexity Group
University of Southampton
Southampton UK
sdr@ecs.soton.ac.uk

Joel Fischer
Mixed Reality Lab
University of Nottingham
Nottingham UK
jef@cs.nott.ac.uk

Tom Rodden
Mixed Reality Lab
University of Nottingham
Nottingham UK
tar@cs.nott.ac.uk

Nicholas R. Jennings
Agents, Interaction and Complexity Group
University of Southampton
Southampton UK
nrj@ecs.soton.ac.uk

Abstract
Interactive autonomous systems are likely to be more involved in future energy systems to assist human users. Given this, we prototyped a future scenario in which householders are assisted in switching electricity tariffs by an agent-based interactive system. The system uses real-time electricity monitoring to instantiate a scenario where participants may have to make, or delegate to their agent (in a variety ways), tariff switching decisions given uncertainty about their own consumption. We carried out a field trial with 12 households for 6 weeks in order to study the notion of autonomy. The results show nuanced ways in which monitoring system performance and taking control is balanced in everyday practice. Our field study provides promising directions for future use of smart systems that help householders manage their energy.

Author Keywords
Smart internet of things, human-agent interaction, real-world deployments, energy

ACM Classification Keywords
H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous. See: http://www.acm.org/about/class/1998/ for help using the ACM Classification system.

Introduction

The increasing availability and miniaturisation of low-cost sensing, actuation and computational devices is likely to result in the wide adoption of sensor-based "smart" applications and systems that leverage large amounts of data. Researchers and practitioners need to face the challenge of designing interactive systems to help people take advantage of this data in a domestic context, often with the ambition to help us more efficiently utilise the limited resources of our planet. A significant part of this body of work has focussed on how information can be presented in the form of feedback that may help change our behaviours [5, 6], while a smaller number of projects have investigated systems that automatically respond to sensed environmental changes. Examples of the latter include "smart" thermostats that take into account our whereabouts and learn our preferences [7, 15].

We are interested in the potential of domestic interactive technology to operate autonomously, and in users' inclination, or resistance, to deliberately transfer *control* to such systems. On the one hand, autonomous operation may be desirable, or even essential, if we are to harness the capabilities offered by increasing amounts of data. On the other hand, autonomous operation may not be the best choice due to noise and biases in real world data, the limited size of training datasets, and the discrepancies between computationally feasible models and complex real-life systems, that result in the operation of these "smart" autonomous systems being, at times, sub-optimal or, in the worst case, detrimental. For example, recent work examining the real-world uptake of a smart thermostat highlighted how such errors are likely to cause users' frustration and may lead them to abandon the technology [15, 16]. It is therefore crucial that researchers and designers understand how to best design interfaces and interaction techniques that make the system status and operation clearly readable, and that allow its users to easily shift between autonomous and manual operation, a notion known as "flexible autonomy" [8].

In particular, our focus is on how autonomous interactive computing systems may mediate user interaction with future energy infrastructure: the "smart" grid. Specifically, we present a field evaluation of a novel home energy management application called TariffAgent, which monitors household energy consumption, as well as available energy tariffs, and therefore calculates the best tariff, and (optionally) automatically switches to it. The trial, which lasted for 6 weeks and involved a diverse group of 12 participants, is reported through an analysis of automatic interaction logs and qualitative analysis of post-trial interviews, aimed at uncovering orientations towards agency and smart systems. An earlier version of TariffAgent and a shorter trial (with a different group of participants) were presented elsewhere [1], here we present a longer field trial with a new version of TariffAgent, where we focus on a more realistic scenario[1], which allows us to more distinctly focus on users' orientation to smart systems, while also reducing novelty effects.

Background

An energy tariff is a pricing scheme by which consumers are charged for the energy they use. Today, there are 34 energy providers that the consumer in the UK faces when she needs to pick a tariff. The energy providers offer various tariff structures, where the most common ones are fixed tariffs. The fixed tariffs usually include two guaranteed rates: a daily standing charge that is the price for the energy service such as distribution and

[1]We removed an emphasis on renewable energy and therefore external uncertainty.

maintenance, and a unit rate that is the price per unit of energy. Given that there are many tariff options available in the energy market, finding the best energy tariff may be a daunting task for most customers. Indeed, in the UK, most customers (40%-60%) tend to stick to the same tariff and do not spend much time searching for a better tariff [9].

To assist human users within this complex and dynamic environment, automated home energy management systems that rely on autonomous software agents [11] have been proposed. Agent-based energy infrastructures make it possible to provide a wide range of services. For example, they may manage energy in off-grid homes [2]. The consequences of these services might influence people's daily lives significantly [3]. Closer to our work, researchers have introduced an agent-based recommender system that provides energy tariff suggestions based on consumption predictions and identified deferrable loads [10]. However, none of these approaches investigated in detail the human-agent interaction issues that arise when they are deployed as they either presume that consumers will embrace any schedule or tariff suggested by an agent.

In particular, we are interested in exploring the challenges raised by the notion of autonomy [12] and to observe how domestic users respond to living with such a technology for a period of time 'in the wild'. To this end, we developed a system specially designed to investigate interaction issues within the domain of automated energy tariff switching. Our system allows users to adjust the degree of autonomy of the system, and to continuously provide manual input that could improve its performance. By so doing, we aim to shed light on the challenge of balancing user control and autonomy in agent-based systems so as to make autonomous domestic systems more useful, intelligible, and trustworthy [4].

Managing Tariffs with TariffAgent

TariffAgent is developed as a combination of a web application and off-the-shelf sensors. It is inspired by the scenario depicted in a recent work [12], where autonomous agents embedded in households have the ability to switch the energy providers based on their offered rates and the user's consumption routines. In our scenario, we consider a daily electricity tariff-switching problem so as to be able to create a realistic field study (as depicted in next section). Therefore, we only concerned with the daily energy usage of a home for selecting the best tariff. In what follows, we elaborate on how this daily consumption will be billed by different tariffs and hence define the challenge of choosing between these tariffs based on the predictions of it.

Daily Tariff Switching

Let the energy consumption of a home on day d be denoted as $c_d \in \Re^+$ kWh, where $d \in \{0, \cdots, D\}$, and $D \in Z^+$. Then, let the set of tariffs provided by suppliers in the energy market be denoted as $t_1, ..., t_s \in T$ and for each tariff there exists a function $F : T \times \Re \to \Re$ 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 supplier where a customer is charged a fixed unit rate r_1 and a standing charge s_1, function F would return $c_{d+1} \times r_1 + s_1$, where c_{d+1} is actual consumption.

To create a more challenging decision environment for users of the application (and therefore incentivize them to delegate their tariff decisions to an agent), we assume there exists eight suppliers, each with their own tariffs similar to real-world tariffs. In particular, each tariff represents the best value for a particular consumption range so that it is not easy to decide which tariff is the cheapest as it may change every day unless the user is able to accurately predict her own consumption.

Software Agent
In our scenario, planning which tariff to change to and when to change is a well-suited task for a software agent since it is necessary to continuously monitor the changing consumption to predict the best tariff. Energy consumption is monitored by the agent through off-the-shelf home energy monitoring devices. These devices measure the total consumption of the household through a current clamp, and make the data available through an HTTP API.

As was shown in a previous study, applying complex machine learning techniques to predict day-ahead usage accurately is a challenging problem [14]. 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 the user's reaction to the agent when it may make mistakes. Thus, the software agent uses 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., $c'_{d+1} = c_{d-1}$). For all tariff switching suggestions and those autonomously enacted, the agent uses the functions defined in the previous subsection to determine the cheapest tariff for the day-ahead.

Applying Flexible Autonomy
To enable users to flexibly specify their relationship with our agent, we provide a number of autonomy levels for various interaction modalities that go beyond the simple notion of moving between human-controlled and fully autonomous tariff switching. These autonomy levels allow users to dynamically arrange the authority and responsibility of their own and the agent. In addition, with these levels they can decide how the human-agent interactions will be held. The three different autonomy levels designed for this study are:

- Human-guided: If the agent detects that the current tariff is different from the one predicted to be the best for the next day, it sends an SMS suggesting a tariff change. Users can accept the suggestion by replying as yes via SMS.
- Semi-autonomous: The agent automatically switches to the predicted best tariff and informs users of the change via SMS. If the users are not happy with the change they can go to the website and manually change the tariff there. This setting is semi-autonomous in that it automatically switches tariffs, but it allows users to easily regain control.
- Fully autonomous: The agent automatically switches to the predicted best tariff but does not inform the user of the change. This setting is fully autonomous in that it completely offloads users of the burden of tariff switching.

Moreover, a user, under any of the above setting, receives a daily report on her performance (with the help of the agent) for the previous day. In the next section, we present the interaction modalities that allow users to communicate with the agent.

Human-Agent Interactions
Interactions between users and the agent are supplied through two mediums: they can interact with the agent

either through SMS, 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.

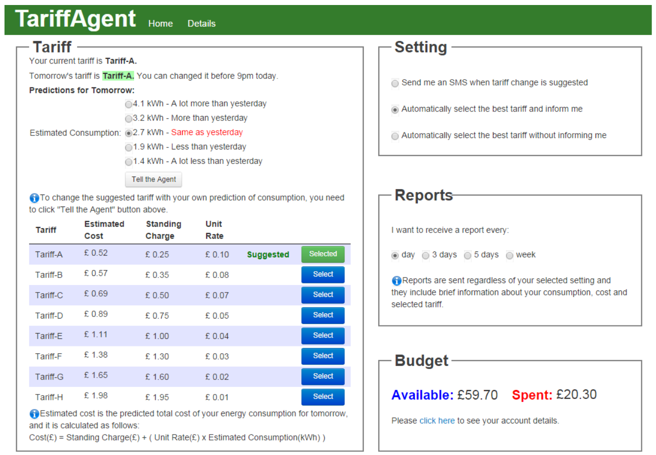

Figure 1: Home page.

The Home page, illustrated in Figure 1, comprises four components: tariff, setting, reports and budget.Through the Tariff component users can see the current tariff and manually select the tariff for the next day. The current and the next day's tariffs are displayed on the top. The next day's tariff is highlighted in green if it is the same as the one the agent suggested, otherwise in orange, to emphasise the difference. In the middle of the component, the predicted values for the user's consumption on the next day are shown. Below the predictions, the eight tariffs are listed, from the one predicted to be the cheapest to the most expensive. The suggested tariff (the cheapest) is marked as such through a text label. Users can select a tariff through a button that is placed next to each tariff.

The predicted amount of energy consumption for the next day can be modified through radio buttons. Changes in the consumption prediction are immediately reflected in the estimated costs for each tariff. The manually selected prediction can be confirmed, or 'told', to the agent by clicking a button. By so doing, users can understand how the agent uses their predicted consumption to make a better choice on their behalf and therefore inspire confidence in the system.

Setting is the second component of the home page. It allows users to select one of the three autonomy levels described in the previous section. Because the delegation of control is central to our study, we included this setting on the home page, to make it easy for participants to adjust the level of autonomy during the trial.

The third component, Reports, enables users to decide how often to receive an SMS report, where the option every day is initially selected by default. The other options are: every 3 days, every 5 days and every week. The last component on the home page is the Budget, which displays how much was spent and how much remains of the budget allocated at the beginning of the study (see the next section for more details). It also provides a link to the other web page of the system.

The Details Page provides historical information about the operation of TariffAgent, with the aim of allowing users to evaluate their performance, and make the system transparent. In particular, for each past day the predicted and actual values for energy consumption are shown, together with the suggested and actual tariff selection, the budget, the cost and the savings or loss incurred. To facilitate the understanding of the information displayed, table cells are colour-coded. Tariffs are displayed in green/red depending on whether the selection was optimal or suboptimal. Consumption predictions are shown in green when they turned out to be accurate (within 10%) and the resulting tariff suggestion is optimal. They are shown in red when they are inaccurate (outside 10%) compared to the realised values and the resulting tariff suggestion is suboptimal. They are shown in orange if the predictions are off (outside 10%), but the resulting tariff suggestion was optimal.

Agent Interacting with Human User
The agent can send three different types of notifications via SMS: reports, suggestions and confirmations. Reports provide information on how much energy was consumed, how much the cost was, which tariff was selected, and how much was saved or lost (compared to the optimal or the worst tariff). Reports were sent to all users regardless of their setting or tariff. An example report is: 'Hello, yesterday your tariff was Tariff-A, your consumption was 4.4 kWh and it cost you 0.69 pound. You saved 1.30 pound with Tariff-A.' The system sends suggestions to users who are on human-guided setting, when their tariff for the next day is predicted not to be optimal, for example: 'Hello, your tariff needs to be changed from Tariff-A to Tariff-B for tomorrow. If you confirm the change please reply as Yes.' For brevity saving assumptions are only presented in the web UI, rather than in the SMS. Confirmation messages are sent only to users on semi-autonomous setting to inform them of an automatic tariff switch, such as: 'Hello, I switched your tariff from Tariff-A to Tariff-B for tomorrow.'

Evaluation in the Wild
Our focus is to study how people interact with an autonomous system, such as TariffAgent, which can influence people's financial situation and possibly daily routines. In order to obtain meaningful results from such a study, we believe that it is crucial to offer a high degree of ecological validity. Therefore a real world deployment, in the form of a field trial, was designed and carried out to evaluate how people experienced the system by using it as part of their everyday life.

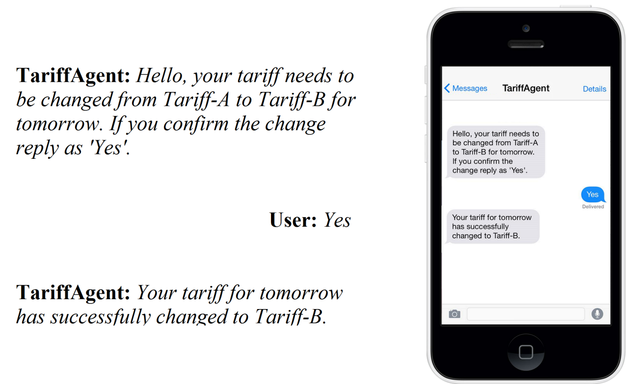

Figure 2: A text dialogue with semi-autonomy.

Participants
We recruited 12 participants (6 female, 6 male) to cover a range of lifestyles, through email lists, word of mouth and snowball sampling. The only requirement to take part in the trial was to have a broadband Internet connection and a mobile phone.

Method
The field trial involved participants installing the meters in their own homes and using the system for a period of 6 weeks (42 days). The system used the real electricity consumption data collected from the participants' homes to calculate their daily energy cost, based on the energy consumption and the selected tariff (as detailed in the previous section). To motivate participants to engage with the system and experience the use of an autonomous system, we provided monetary incentives based on their

performance. At the beginning of the study, all participants were allocated an online budget of £80, and their daily consumption cost was reduced from this budget over the period of the trial. At the end of the study, participants received payments (in the form of a gift voucher) according to the amount left on their budget. By so doing, we not only aimed to encourage participants to engage with the system, but also make saving have a real and tangible impact. This idea of using monetary incentives to mimic energy pricing was partly inspired from earlier studies around energy pricing [13], where participants were rewarded according to their study performance.

Initial Evaluation

We report a six-week deployment 'in the wild' with 12 different households. Quantitative data was collected through automatic recording of user interaction logs, including how many times users visited pages, replied SMS tariff suggestions and input their own consumption estimations. Moreover we conducted individual semi-structured exit interviews to collect qualitative data. The interviews focused on people's use, adoption and understanding of the system, and lasted between 20 and 30 minutes.

The default autonomy level at the beginning of the trial, for all participants, was human-guided: the system would send SMS suggestions about tariff switching but it would not automatically switch. This default option was chosen because it is the one that requires users the most interaction, so we wanted to see whether they would change to a less demanding one by time. Four of the participants modified the autonomy level to semi-autonomous option, where the system automatically changes to the predicted best tariff and informs the user of the change via SMS. The remaining eight users kept using the default autonomy level. No one selected full autonomy (where the system changes the tariff without informing the user).

All users except one took advantage of the web UI to provide manual estimates of their electricity consumption prediction for the following day at least once. In total this explicit input was provided 110 times during the study, and resulted in 85 times correct tariff selections. Moreover, interactions between users and agent were maintained at notable level by each user till the end of the study. Everyone either accessed the web interface or replied to SMS suggestions with some regularity that is on average at least once every 2.5 days (SD: 1.3 days). Moreover, none of the users changed the frequency of reports from daily to a less frequent option, namely opted for the agent interacting with them every day.

In the interviews, most participants appeared to hold a mental model that mirrors quite closely the actual design and implementation of TariffAgent. All participants commented that they perceived the system as helping them save money, through mostly correct suggestions, and most of them stated that they trust the system's tariff decisions. When asked about experienced mistakes in tariff suggestions or selections, they mostly considered it to be their own responsibility.

Conclusion

We present a field trial that exposed 12 participants to a prototyped future energy scenario for 6 weeks. We studied users' interaction with an interactive autonomous system designed to help in managing energy tariffs. An initial evaluation revealed that the system works in terms of engaging participants for long term, where users monitored the system's performance and took control when they considered it necessary. We believe that our study opens the route, and highlights opportunities for further research to investigate human-agent interactions in future smart systems.

Acknowledgements

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

References

[1] Alan, A., Costanza, E., Fischer, J., Ramchurn, S., Rodden, T., and Jennings, N. R. A field study of human-agent interaction for electricity tariff switching. In *Proc. AAMAS*, ACM Press (2014), 965–972.

[2] Banerjee, N., Rollins, S., and Moran, K. Automating energy management in green homes. In *Proc. SIGCOMM HomeNets*, ACM (2011), 19–24.

[3] Chetty, M., Sung, J.-Y., and Grinter, R. How smart homes learn: the evolution of the networked home and household. In *Proc. Ubicomp*, Springer (2007).

[4] Costanza, E., Fischer, J. E., Colley, J., Rodden, T., Ramchurn, S., and Jennings, N. R. Doing the laundry with agents: a field trial of a future smart energy system in the home. In *Proc. CHI*, ACM (2014), 813–822.

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

[6] Froehlich, J., Findlater, L., and Landay, J. The design of eco-feedback technology. In *Proc. CHI*, ACM (2010).

[7] Gupta, M., Intille, S. S., and Larson, K. Adding gps-control to traditional thermostats: an exploration of potential energy savings and design challenges. In *Proc. Pervasive*, Springer-Verlag (2009), 95–114.

[8] Jennings, N. R., Moreau, L., Nicholson, D., Ramchurn, S. D., Roberts, S., Rodden, T., and Rogers, A. On human-agent collectives. *CACM (In Press)*.

[9] OFGEM. The retail market review - findings and initial proposals. Tech. rep., 2011.

[10] Ramchurn, S., Osborne, M., Parson, O., Rahwan, T., Maleki, S., Reece, S., Huynh, T., Alam, M., Fischer, J., Rodden, T., Moreau, L., and Roberts, S. Agentswitch: towards smart energy tariff selection. In *Proc. AAMAS*, ACM (2013), 1401–1402.

[11] Ramchurn, S., Vytelingum, P., Rogers, A., and Jennings, N. R. Putting the smarts into the smart grid. *Communications of the ACM 55*, 4 (2012), 86–97.

[12] Rodden, T. A., Fischer, J. E., Pantidi, N., Bachour, K., and Moran, S. At home with agents: exploring attitudes towards future smart energy infrastructures. In *Proc. CHI*, ACM (2013), 1173–1182.

[13] Slavin, R. E., Wodanski, J. S., and Blackburn, B. L. A group contingency for electricity conservation in master-metered apartments. *JABA 14*, 3 (1981).

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

[15] Yang, R., and Newman, M. W. Learning from a learning thermostat: lessons for intelligent systems for the home. In *Proc. UbiComp*, ACM (2013), 93–102.

[16] Yang, R., Newman, M. W., and Forlizzi, J. Making sustainability sustainable: Challenges in the design of eco-interaction technologies. In *Proc. CHI* (2014).