AgentSwitch: Towards Smart Energy Tariff Selection (Demonstration)

Sarvapali D. Ramchurn,[1] Michael A. Osborne,[2] Oliver Parson,[1] Talal Rahwan,[1]
Sasan Maleki,[1] Steve Reece,[2] Trung D. Huynh,[1] Mudasser Alam,[1] Joel E. Fischer,[3]
Tom Rodden,[3] Luc Moreau,[1] Stephen Roberts[2]

[1]Electronics and Computer Science, University of Southampton, United Kingdom
{sdr,op106,tr,sm9g09,tdh,ma606,lavm}@ecs.soton.ac.uk

[2]Pattern Recognition Lab, University of Oxford, United Kingdom
{mosb,sjrob}@oxford.ac.uk

[3]Mixed Reality Lab, University of Nottingham, United Kingdom
{Joel.Fischer,Tom.Rodden}@nottingham.ac.uk

ABSTRACT

We present AgentSwitch, a prototype agent-based platform to solve the tariff selection problem for homeowners. AgentSwitch incorporates novel algorithms that work on the coarse data provided by smart meters to make predictions of hourly energy usage as well as detect (and suggest to the user) deferrable loads that could be shifted to off-peak times to maximise savings. Our demo will allow users to interact with AgentSwitch and explore test user accounts in order to understand the impact of different usage profiles and appliance loads.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence*

Keywords

Electricity, Smart Grid, Optimisation, Group Buying, Provenance, Recommender Systems

1. INTRODUCTION

Energy poverty is a rapidly growing issue across the world due to the signficant rise in energy costs over the last few years. Such increments are due to the unprecedented growth in energy demand (e.g., world energy usage is set to grow by more than 50% by 2030) coupled with dwindling fossil fuels and the high costs of (and resentment against) constructing large renewable energy generation facilities. In the UK energy market, we note that energy security issues are exacerbated by a misalignment of incentives whereby energy companies have no motivation to innovate nor to adopt cleaner sources of energy, given that they can easily pass on any rising costs they face directly to their consumers. In addition, most consumers (40%-60% in the UK) tend to 'stick' to the same energy supplier year in year out and do not spend much time looking for a cheaper deal. This reduces competition in the energy retail market and does not help drive down prices [4]. Indeed, research by the U.S. Dept. of Energy found that most people in the U.S. are likely to spend *no more than two hours a year* setting their preferences for comfort, tariffs, and environmental impact [4].

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

Now, consumers cannot be completely blamed for not finding the best deal, given that energy tariffs are often made explicitly complex and, at times, confusing. For example, tariffs may have multiple tiers (e.g., the first tier may be priced at 20p per kWh, and beyond this the cost drops to 5p/kWh), implement time-of-use pricing (e.g., 5p/kWh between 11pm and 7am), and may include additional one-off discounts (often only for a limited period). To help consumers combat this complexity, a number of third-party online services exist to help consumers submit simple estimates of their yearly consumption and obtain the cheapest tariff (e.g., uswitch.com and moneysupermarket.com). Some other services also claim to help consumers come together as a collective in order to access group discounts from retailers (e.g., which.com, incahoot.com). Crucially, however, these services rely on consumers being able to make a reasonable estimate of their yearly consumption (taking into account varying usage over different seasons and usage at on- and off-peak times) and being able to understand how to take advantage of the various tiers or time-of-use tariffs they offer (e.g., by shifting appliance usage to off-peak times). Moreover, in existing collective purchasing systems, all members of the collective tend to obtain the same contract without considering whether the discounts are fairly distributed across all the members of the collective (e.g., those with unpredictably peaky consumption profiles should be charged more than those with predictably flat profiles, as they tend to cause higher penalties in the balancing market).

Against this background, in this demo, we present a fully working prototype of an agent-based platform, called AgentSwitch, that integrates state-of-the art techniques and mechanisms to address the challenging issue of energy tariff selection. AgentSwitch builds upon the data provided by off-the-shelf energy monitoring devices and applies a number of machine learning, optimisation, and coalition formation algorithms in order to solve the energy tariff selection problem. In more detail, this demo will show how state-of-the-art intelligent algorithms as detailed in [3] can be used to solve the problem of tariff selection. These include (i) novel extensions to Bayesian Quadrature (a machine learning technique), in order to generate predictions of yearly consumption at hourly level and help select the best tariff traditionally available from energy retailers, (ii) a novel mechanism for collective energy purchasing,[1] (iii) a new non-intrusive appliance load monitoring (NIALM) algorithm that works on coarse energy data (at five-minute level rather than

[1]The group buying elements are more speculative as there are currently no group buying tariffs on the market.

second-level as is traditionally the case in this field) in order to detect deferrable loads that might benefit from being shifted to off-peak times, (iv) a novel provenance service that allows the tracking of data throughout the system in order to provide accountability for its recommendations. The demo will walk viewers through the key algorithms and the user interfaces used to provide advice to users. A video is provided (see: `http://bit.ly/11OeghI`) and describes the rationale behind AgentSwitch and the mechanics of its main components focusing on the prediction, disaggregation, and provenance tracking elements of the prototype.

The rest of this paper is structured as follows. Section 2 details the system architecture underlying AgentSwitch and discusses what will be presented in our demo.

2. AGENTSWITCH ARCHITECTURE

AgentSwitch uses a data store that collects aggregate (whole-house) energy usage data from the energy monitoring device installed in the home (see Figure 1) and uses a number of AI modules to provide recommendations as follows:

Annual Load Prediction: once the data (power in kW averaged over 5 min intervals) is retrieved from the data store, the energy usage for the whole year is predicted based on the actual data. 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 year-long 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 [2]. In this work, we adapt the technique to allow for quasi-periodic energy consumption signals emerging from the weekly cycles of typical domestic consumers.

Group Buying: this module identifies and forms coalitions of consumers to take advantage of group discounts from retailers. Given a group of consumers and their estimates, the value of a coalition (group of energy consumers signed up to AgentSwitch) is the expected cost of the aggregated consumption of all its members given the group discount applied to the collective. In particular, we note that, given the workings of the forward energy market and the penalties applying for unexpected consumption in the real-time (balancing) market, flat and predictable (low variance) load profiles will obtain the highest discounts. Given this we construct the cost function for a coalition based on the expected usage profile over a given period (e.g., day, month, or year) as well as the variances (capturing the uncertainty in demand) of all coalion members and use this cost function in computing the Shapley value for all agents using a novel approximation algorithm that scales to thousands of agents and returns solutions in minutes.

Load disaggregation: AgentSwitch focuses on the identification of appliances (e.g., washing machine or dishwasher) which both consume a large amount of energy and can be deferred to another time of day (i.e., deferrable loads) with minimal inconvenience to the household occupants. Since these deferrable loads have a high energy consumption, they can be easily disaggregated from the remainder of a household's energy consumption, despite the low data sampling rate. Given this, we first construct appliance models from individually metered appliances from houses other than those in which disaggregation will be performed, which we refer to as the training phase. Second, the appliance models are used to identify appliance signatures within the aggregate electricity data, which we refer to as the disaggregation phase. The training phase consists of operation detection, followed by feature extraction and model construction, while the disaggregation phase consists of operation detection, followed by feature extraction and operation classification (i.e., identifying the appliance run).

Provenance: as recommendations from AgentSwitch might potentially lead to financial gains or losses, it is crucial to be able to identify the origin of error once a recommendation is deemed to be inaccurate. This, however, is challenging due to the multiple paths of (data) dependencies inherent to such a complex system. In order to address this issue, the chains of dependencies that lead to a recommendation, i.e., its provenance [1] were fully tracked in AgentSwitch. Such information enables a systematic approach to pinpointing the sources or agents responsible for the errors and to auditing the data produced within the system.

Now, in order to provide tariff recommendations to users who sign up to use the service (i.e., allow AgentSwitch to analyse their data to provide recommendations), AgentSwitch needs to access at least[2] two key sources, namely consumers' electricity consumption readings (to be kept in a database) and live electricity tariff specifications from all retailers (for AgentSwitch to match consumption predictions or group consumption against the best tariffs). While the former can be obtained from users' off-the-shelf energy monitors or smart meters that provide average power readings at different levels of granularity, the latter can be obtained from online third-party providers and suppliers (in our case, we used live tariff data from `uswitch.com` with their permission).

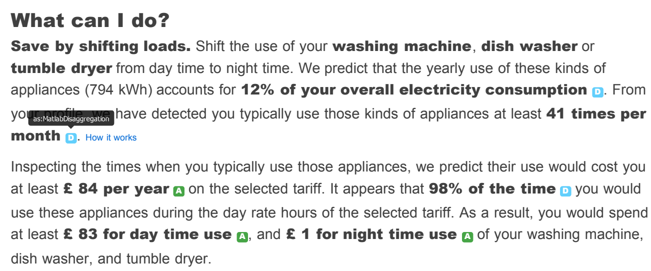

Figure 1: Suggestions based on disaggregated energy consumption and provenance data.

The demo will consist of user walk-throughs of the system and will run through the challenges in generating recommendations for different homes that exhibit different consumption patterns. A number of energy monitoring devices will be on display to explain the operation of AgentSwitch. Finally, the demo will be used to gather feedback from users in order to help build the first publicly deployed version of AgentSwitch.

Acknowledgements

This work is supported by the EPSRC-Funded ORCHID Project EP/I011587/1. We also thank uSwitch for access to their API.

3. REFERENCES

[1] L. Moreau and P. Missier. PROV-DM: The PROV Data Model. Technical report, World Wide Web Consortium, 2012.

[2] M. Osborne, R. Garnett, S. Roberts, C. Hart, S. Aigrain, N. Gibson, and S. Aigrain. Bayesian quadrature for ratios. *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics (AISTATS 2012)*, 2012.

[3] S. D. Ramchurn, M. Osborne, O. Parson, T. Rahwan, S. Maleki, S. Reece, T. D. Huynh, M. Alam, J. Fischer, L. Moreau, S. Roberts, and T. Rodden. Agentswitch: Towards smart energy tariff selection. In *The 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012) – to appear*, May 2013.

[4] US Department of Energy. The Smart Grid: An Introduction, 2007.

[2]Other sources of data such as the users' schedule and travel plans for the year ahead, the users' home types and contents, would all be useful to improve predictions and load disaggregation.