Smarter Electricity through Argumentation

Menelaos Makriyiannis, Tudor Lung, Robert Craven, Francesca Toni and Jack Kelly
Department of Computing,
Imperial College London,
London SW7 2AZ
{mm2810,rac101,ft,dk3810}@doc.ic.ac.uk, tudor.lung@gmail.com

Abstract—The current, widespread introduction of smart electricity meters is resulting in large datasets' becoming available, but there is as yet little in the way of advanced data analytics and visualization tools, or recommendation software for changes in contracts or user behaviour, which use this data. In this paper we present an integrated tool which combines the use of abstract argumentation theory with linear optimization algorithms, to achieve some of these ends.

Keywords-Smart electricity; Argumentation; Optimization.

I. INTRODUCTION

Recent research [1] has shown that UK energy prices are rising at eight times the rate of average earnings, and the industry's trade body, Energy UK, recently warned that "household bills could rise by another 50% over the next six years" [2]. Partly as a response to this trend, and partly also in an attempt to improve the sustainability of current forms of energy usage in at least the medium term, the UK Deptartment of Energy and Climate Change has required a gradual introduction of 'smart meters', to be rolled out to all UK homes by 2020 [3]. EU legislation from 2012 states that "80% of all electricity meters in the EU have to be replaced by smart meters by 2020" [4].

Smart meters record the utility (e.g., electricity, gas, water) consumption within a household—our focus here is on electricity. After a regular, and often short interval, the energy consumed is recorded, and can be stored or sent to a local database for future analysis. This can result in a substantial amount of data's being available for analysis. However, there is currently no complete software package available which makes an innovative use of the data, to properly empower individual users to control their electricity bills by providing detailed analysis of existing usage, and intelligent recommendations on how to adjust usages or contract providers, in order best to meet the users' demands.

In this paper we present preliminary work to that end. Specifically, we make use of abstract argumentation (AA) theory [5], a branch of logical AI, to construct theoretically-underpinned arguments for how a user might change electricity contracts and behaviour in order to minimize their electricty bill. The use of AA is combined with linear optimization algorithms, and both are applied to a data output from smart meters currently available. We use two large, real-world data sets for testing and evaluation. The resulting web-based implementation is available online, as a prototype, at http://smartelectricity.io.[1]

The paper is organized as follows. In §II we present background on abstract argumentation, and describe the data sets used. In §III we describe central parts of the functionality of our tool. In §IV we present results from the detailed evaluation, and comparison, made. We conclude in §V.

II. BACKGROUND

Abstract argumentation [5] is used to represent the reason about the relations of conflict between opposing arguments, without delving into the internal logical structure of those arguments.

Definition 1. An *abstract argumentation framework* is a tuple $(Args, \rightsquigarrow)$, where

- *Args* is a set of *arguments*;
- $\rightsquigarrow \subseteq Args \times Args$ is the *attack relation*. ⌐

For example, the set *Args* might represent arguments for changing electricity contracts to a different provider based on current or projected electricity usage patterns, with the attack relation $\rightsquigarrow$ being between arguments for different contracts.

A total set of 'winning' arguments is determined in several alternative ways, based on the attack relation over *Args*, as follows.

Definition 2. Where $(Args, \rightsquigarrow)$ is an AA-framework:

- $A \subseteq Args$ is *conflict-free* if there are no $a, b \in A$ such that $a \rightsquigarrow b$;
- $A \subseteq Args$ is *admissible* if (i) it is conflict-free, and (ii) for any $a \in A$ such that $b \rightsquigarrow a$, there is $a' \in A$ such that $a' \rightsquigarrow b$;
- $A \subseteq Args$ is *preferred* if it is maximally (w.r.t. $\subseteq$) admissible. ⌐

Admissibility of a set of arguments thus ensures for them a form of internal consistency, and means they collectively defend themselves from external attack; preferredness requires this in a $\subseteq$-maximal way. (Other semantics have been widely studied. See [6] for details.)

We were able to make use of two large electricity smart-meter datasets in testing and evaluating our implementation.

First, the *UK Power Data* (UKPD) dataset [7][2] contains detailed power consumption data from four London houses

---

recorded over several months. The data were recorded using two types of sensor, Current Transformers[3] and EDF Transmitter Plugs,[4] for individual appliance measurements. The data contains entries for individual appliances, as well as for the entire house consumption.

Secondly, the *Household Electricity Survey* (HES) dataset [8] includes data from 250 owner-occupied households from England, from 2010 to 2011. 26 households were monitored for a full year, and the remaining 224 for one month, on a rolling basis throughout the trial. This data is disaggregated—that is, appliance-specific—with no aggregates provided per house.

In addition to the two datasets, we collected and represented electricity contract data. All widely-used contracts from the two main UK electricity suppliers (British Gas[5] and EDF Energy[6]) were used.

III. FUNCTIONALITY

A. *Data representations*

Contract data was standardized to be represented in a UML format of the form
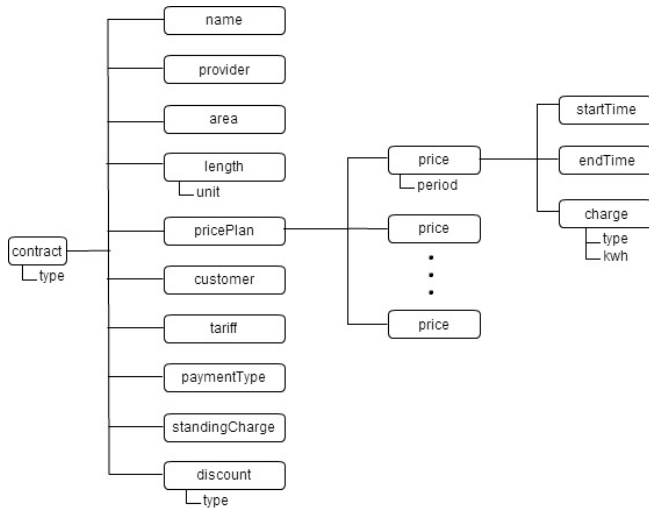
A configurable, stand-alone application was written to perform data-standardization for the smart electricity readings from the two datasets. This standardization had several aspects. The data in each of the two sets represents the electricity usage over a given interval of n seconds; but the value of n varies, both between datasets and also within specifically the HES dataset—as determined by the particular smart meters in use. The UKPD data uses an interval of 6 seconds, and the HES data has an interval of 2 minutes for the majority of houses, and 10 minutes for the others. In the case of the UKPD set, this would mean 14,400 data points per house per day—a substantial storage cost. As our implementation is intended to be the prototype for a web-based application, data calls with

such amounts of data would hinder the user experience and also make computation much more expensive, both on the server-side and on the client-side. Accordingly, we selected a time of 1 hour to be the standard length of interval, and aggregated both datasets appropriately.—Other aspects of the standardization were less significant. We chose the DateTime format[7] for the time-stamp, and since the HES dataset only exists in a disaggregated form (values per appliance), we also summed the data to produce total series of values per household—values which are used in some of the functions of our implementation.

B. *Data visualization*

For the data visualization components of the implementation, we used the HighchartsJS[8] framework, which allows for a wide variety of visualization types, and is supplied with extensive documentation. In the present subsection we describe three of the data visualizations we developed.

First, the **aggregated consumption and cost chart**. (For a screenshot, see Figure 1.) This chart represents two types of data: a line series depicts the consumption of electricity over a given period, with points at each hourly interval as determined by the data standardization discussed in §III-A. Secondly, at each point of the aggregated consumption data line, the respective cost of that consumption sample under one or more contracts is visualised in the form of column bars. This may be the cost under the contract the user is currently under or any alternative contracts, or a combination thereof. The user may choose the total period for which the data is displayed, either using presets or by sliding a bar to fix the period precisely; and the contracts whose pricing is displayed can be toggled on or off in the legend. Finally, as the user hovers over points on the graph, the precise values of the nearest data point are displayed. This form of visualization enables the user to see, of their aggregated electricity consumption, how a selection of contracts compare in a fine-grained way.

The aggregated consumption and cost chart may prove useful to a user if the user wants to examine their usage as a whole, with a view to changing contract. However, the user may wish, instead, to alter their behaviour whilst remaining with their current provider and on their current contract. For this, visualizations using the disaggregated datasets are more appropriate. Thus, secondly, we implemented a **disaggregated consumption chart** (see Figure 2). As with the aggregated consumption and cost chart, the user may select a period for which the data is visualized. A number of appliances are selected, and the visualization then presents the changing total consumption, per hourly interval, over the period selected, as well as the electricity consumption per individual appliance. This enables the user to see frequent spikes in the power consumed by particular appliances—thus allowing them to consider altering their behaviour with respect to those devices which use the most power.
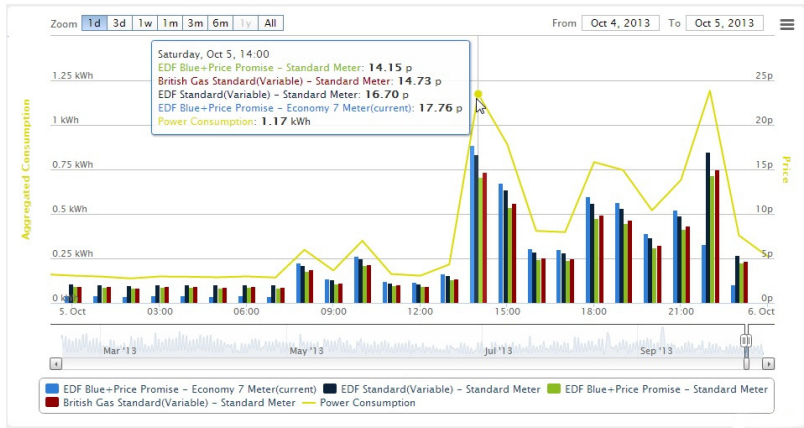
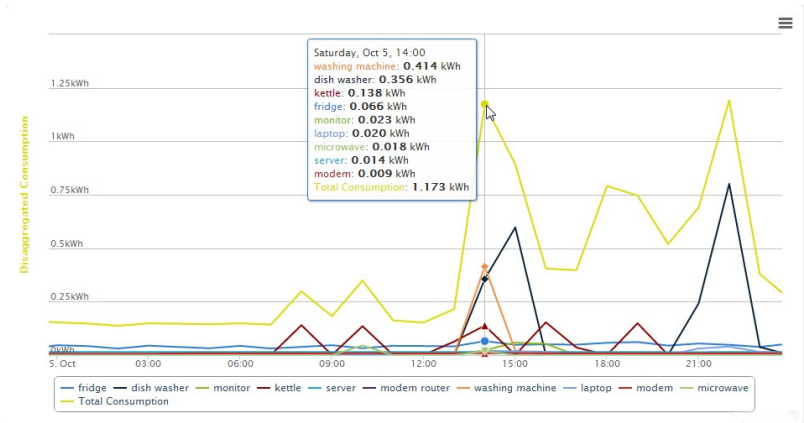Fig. 1.   Aggregated consumption and cost chart.
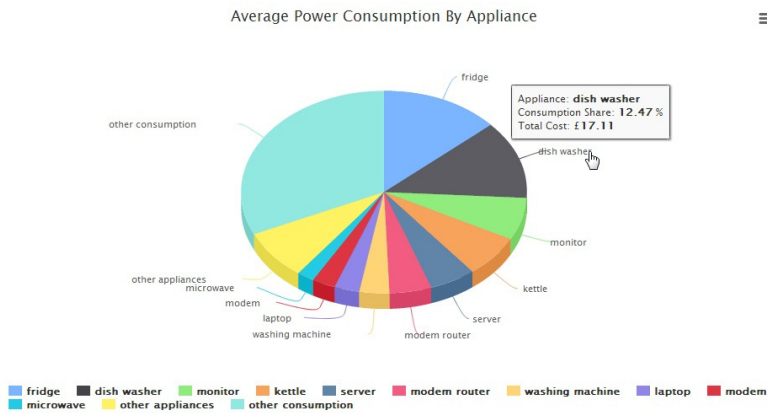

Fig. 2.   Disaggregated consumption chart.

Fig. 3.   Average consumption appliance chart.

Contract Name	Selected Interval (£)	Standing Charges (£)	Entire Consumption (£)	Past Year (£)	Past 6 Months (£)	Past 3 Months (£)	Past Month (£)	Switching Savings (£)
EDF Blue+Price Promise - Economy 7 Meter	260.14	45.56	410.16	399.88	204.67	111.14	38.61	37.88
EDF Standard(Variable) - Economy 7 Meter	290.70	45.56	458.52	447.01	228.86	124.49	43.28	7.32
British Gas Standard(Variable) - Economy 7 Meter (current)	298.02	62.68	469.57	457.86	234.32	126.67	43.91	0.00
EDF Blue+Price Freeeeze - Economy 7 Meter	310.83	45.56	490.40	478.06	244.81	133.29	46.35	-12.82
British Gas Fix and Control - Economy 7 Meter	317.48	77.99	499.87	487.46	249.32	134.32	46.50	-19.47

Fig. 4.   Contract comparison chart.

Related to this fine-grained depiction of applicance electricity usage, our **average consumption appliance chart** (shown in Figure 3) depicts the average electricity consumption over the entirety of a selected data interval, at a single glance. The consumption of the ten appliances which use the most electricity in the given period is shown, together with a section ('other applicances') for those appliances whose electricity consumption is known. Since the datasets we use sometimes show a discrepancy between the sum of the values for the dis-aggregated appliances, and the aggregated consumption value (since not all appliances are included in the disaggregation), there is also a sector in the pie chart for 'other consumption', representing the other appliance usage not included in the disaggregated data.

C. Contract comparison

In order to give a summary of comparisons between different contracts of electricity providers, we also allow a detailed tabular comparison between the prices a user would have been charged, on a set of different contracts, for selectable, different periods of their historical power usage (see Figure 4). The **Contract Summary Table** provides a breakdown of how a user's electricity would have been charged according to different contracts, over a number of different periods; the entries are ordered according to the saving, or loss, to be accrued for that period, with the row for the current contract highlighted. For example, for the user whose electricity is being analyzed in Figure 4, a switch to EDF's 'Blue + Price Promise' Economy 7 contract is indicated: this would involve a saving of £37.88 for the selected period. Switching contracts to either EDF's 'Blue + Price Freeeze' Economy 7 contract, or British Gas's 'Fix and Control' Economy 7, would have resulted in the user being £12.82 or £19.47 out of pocket, respectively.

Although the Contract Summary Table is useful for visualising contracts and giving aggregate cost figures calculated on the basis of the user's consumption, it does not explicitly state comparison conclusions and justification. Further, they are passive features and their purpose is to complement comparisons rather than being the source of them. Thus, we also supplement the Contract Summary Table with **Contract comparison recommendations**, which give detailed, argumentation-theoretic justifications for the results summarized in the table. Let X and

Y be two contracts to be compared. Each contract splits the day into a number of intervals $X_1, \ldots, X_m$ and $Y_1, \ldots, Y_n$. These may not overlap, so that there is a set $\mathsf{T}_{X,Y}$ of $k \leqslant m \times n$ *tarif stretches*:

$$\mathsf{T}_{X,Y} = \{X_i \cap Y_j \mid X_i \in \{X_1, \ldots, X_m\},$$
$$Y_j \in \{Y_1, \ldots, Y_n\}, X_i \cap Y_j \neq \emptyset\}$$

Each tarif-stretch represents a period of time over which each contract's tarif is constant, and for each tarif-stretch $T \in \mathsf{T}_{X,Y}$ either the total charge for contract X, or Y, is greater, or they are equal. Let $X(T)$ be the amount charged by contract X for tarif-stretch $T \in \mathsf{T}_{X,Y}$, and $Y(T)$ the amount charged by contract Y; since contracts also have standing charges, let $X(S)$ be the standing charge for X, and $Y(S)$ that for Y. Further, let $X(\{T_1, \ldots, T_n\})$ be $X(T_1) + \cdots + X(T_n)$, etc. Define $w : \mathsf{T}_{X,Y} \cup \{S\} \to \{X, Y, \{X, Y\}\}$ so that:

$$w(T) = \begin{cases} X & \text{if } X(T) < Y(T), \\ Y & \text{if } X(T) > Y(T), \\ \{X, Y\} & \text{if } X(T) = Y(T). \end{cases}$$

so that $w(T)$ gives the 'winner' between contracts X and Y, in the sense of the contract which charges least for that stretch (or for the standing charge, in the case of $w(S)$).

So, given two contracts X and Y, the resulting abstract argumentation framework $(Args_{X,Y}, \rightsquigarrow_{X,Y})$ has:

$$Args_{X,Y} = \{(\{T_1, \ldots, T_k\}, C) \mid T_1, \ldots, T_k \in \mathsf{T}_{X,Y} \cup \{S\},$$
$$w(T_1) = \cdots = w(T_k) = C, (C = X \vee C = Y)\}$$
$$\rightsquigarrow_{X,Y} = \{((\mathsf{T}_1, C_1), (\mathsf{T}_2, C_2)) \mid C_1 \neq C_2,$$
$$C_1(\mathsf{T}_1) \leqslant C_2(\mathsf{T}_2)\}$$

We then find preferred extensions of the resulting argumentation framework.

Example 1. Consider two contracts, X and Y:
- X charges electricity at a rate of 15p per kWh between 8am and 9pm, and 13.54p per kWh between 9pm and 8am; and there is a standing charge of £40;
- Y charges electricity at a rate of 16.54p per kWh between 9am and 10pm, and 13.54p per kWH between 10pm and 9am, with a standing charge of £35.

This gives $\mathsf{T}_{X,Y}$ and standing charges (S) as follows (shown with the kWh used for a sample period, and the resultant prices for X and Y):

	p/kWh (X/Y)	kWh	Price (X/Y), £	w
T_1, 08–09	15 / 13.54	166	24.9 / 22.48	Y
T_2, 09–21	15 / 16.54	260	39 / 43	X
T_3, 21–22	14.54 / 16.54	300	43.62 / 49.62	X
T_4, 22–08	14.54 / 13.54	400	58.16 / 54.16	Y
S	n/a	n/a	40 / 35	Y

This gives the AA-framework shown below. (The arguments have been annotated with the £ value of the saving made on the winning contract; so the $(\{T_1\}, Y) : 2.4236$ represents that over tarif stretch T_1, contract Y is the cheaper contract by £2.4236.)
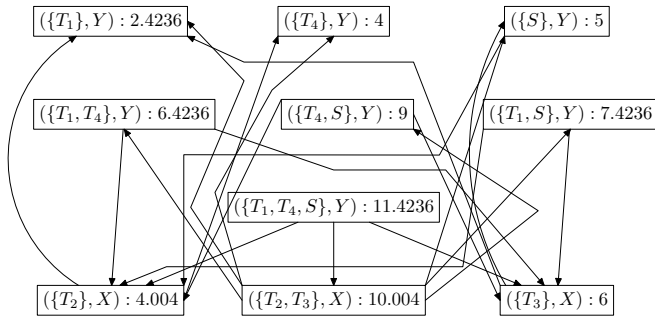
The preferred extension has all and only arguments for Y.

The argumentation framework exhibits a number of features. If there is a 'winning' contract, then there will be at least one argument in $Args_{X,Y}$ supporting that contract which is unattacked by any other argument; in fact, any unattacked argument in the framework represents a decisive argument, based on cost, for the contract it supports. Further, there may be two preferred extensions; this represents the case where, although for particular tarif stretches one contract may be better than another, overall, the contracts would require equal payments for the period and usage under consideration.

The output from the argumentation component is processed in order to give a prose output to the user. For example, w.r.t. the example above:

> Even though in the stretch 09:00–22:00 contract X is less costly by £35.60, contract Y in the interval 22:00–09:00 is less costly by £23.62 and has £15 less standing charges, which is enough to make it less costly overall by £3.02, over the selected date interval.

Whilst the sum-total of the difference between the charges for the two contracts might easily be computed without any argumentation theory, the use of an argumentation-theoretic underpinning brings the logical relations between the structure of charges to the surface, and provides a stepping-stone which lets the natural-language rendition be easily computed.

D. Recommendation generation

Recommendation generation is a central feature of our approach. The recommendations we implemented are divided into two broad categories, of **General** and **Appliance** recommendations. In the present subsection we select some of the many recommendation capacities that our suite of tools provides.

General recommendations concern a user's aggregated consumption (without paying attention to any appliance-specific, disaggregated data). We have already mentioned, in §III-C, contract comparisons: if all contracts available are compared, these can be used to provide a **Contract switch recommendation** based in a straightforward fashion on the user's history of electricity usage. If the user wishes to stay with her or his current electricity provider and contract, then a **Consumption behaviour recommender** can provide advice on how the cost of the user's bill might have been lowered by shifting portions of the aggregate power consumption between

what, in §III-C, were called 'tarif stretches'. The Consumption behaviour recommender finds portions of the power usage which occur near the boundaries between tarif stretches, and optimizes the price of the bill by allocating this power to a different stretch. The tool does this by looking at the total usage for a given period (again, selectable by the user).

A **Reduced usage estimator** allows the user to see what percentage of the various tarif stretches' usage needs to be cut if a desired reduction in the price of a bill is to be achieved. One sliding bar represents the cost of electricity, and more sliding bars (one each per tarif stretch) shows the percentage of usage allocated to that stretch. When the user manipulates the slider on any one bar, the others compensate in real time. If the total bill cost is reduced, aggregate electricity usage is shifted between tarif stretches to indicate how the user must change behaviour; and the process also works in reverse, as the user alters the amount of power usage allocated to any particular stretch.

Finally, **Shifted usage recommendations**, perhaps the most useful of the general recommendations provided, combine the Contract switch recommendations with Consumption behaviour recommendations. The recommendations examine minimal changes to the user's aggregate electricity usage patterns (up to a threshold which is easily configurable) in combination with all possible contracts with all possible electricity providers, to give overall advice on the best bill price possible (given historical usage patterns), if the user can switch contract and provider, and alter a small portion of their power consumption to times slightly earlier or later.

Appliance recommendations use disaggregated data to achieve more fine-grained recommendations. A basic **Appliance consumption table** shows the average hourly consumption per appliance, together with the total price per appliance over a user-selectable interval. An **Appliance recommendation table** performs a similar function to the Consumption behaviour recommender described above, but for individual pieces of equipment. This tool calculates the n appliances, a portion of whose usage might mostly profitably be switched from one tarif stretch to another, displays this information to the user, and finds the resultant monetary benefit. An **Appliance savings estimator** (Figure 5) allows the user to see the effect of different reductions in the usage of individual appliances on total bill cost whilst remaining with the current provider and contract, *or* to set a desired bill reduction and see how this can best be balanced by reducing appliances' use. Thresholds can be set—a useful constraint when a subset of appliances may only be reduced to a certain degree. Recommendations are displayed as percentage values of current usage. **Additional recommendations** for appliances (Figure 6) converts these percentages into cardinal values. This works by estimating, on by an analysis of appliance power usage values, the number of times a given appliance is used per day on average, and calculating how many times fewer the appliances should be used to give the desired cost reduction.

Most of the recommendation functionality is implemented as the solution of linear optimization problems solved in real

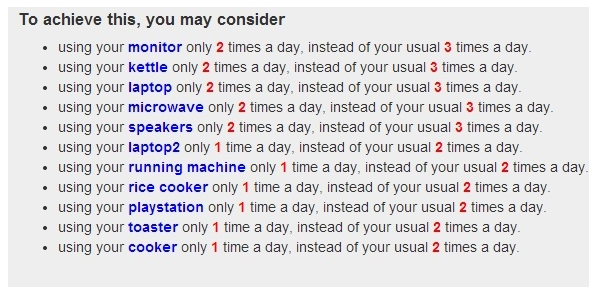Fig. 5. Appliance savings estimator.

Fig. 6. Additional appliance recommendations.

time; we omit the details, which are straightforward.

IV. EVALUATION

We wanted to provide some basic testing of the various forms of recommendation our tool offers. The ideal way to do this would have been to roll out the tool as a beta to users, who could have collected disaggregated data for a year, then followed different recommendations the tool gave. A form of validation could then be found if the users' bills, over the following year, were low relative to all possible combinations of provider and contract.

Unfortunately, this sort of testing was markedly infeasible: we did not have the time to conduct a study of such length, nor the number of users, and did not want to become immersed in technicalities of data collection when our primary intention was with recommendations and the *use* of data. Accordingly, we made use of the existing datasets. These were split into a set from which recommendations were made, and a set against which, once the recommendations had been followed, the effect was measured (call the former *training*, the latter *testing* data). For example, for **Contract switch recommendations**, we split the data in two ways: first, into an 80%/20% division of training/testing; and second, into a 50%/50% division (different splits were used to ensure that any results found

were not unduly sensitive to any arbitrariness in the split). Our tool was used to find a recommendation for the best contract to switch to, given the user's power usage for the training set. We then tested whether, if that recommendation were followed and the user switched contracts, the bill would be the lowest possible. Of the 30 datasets we used, our recommendation tool gave the optimal contract switch in 87% of the cases for the 80/20 split, and 90% of the cases in the 50/50 split. (The difference indicates a shift in usage pattern.) Other tests indicate a comparable rate of success for the various recommendation components. Detailed results are omitted owing to lack of space.

V. CONCLUSION

With the advent of ubiquitous electricity smart-meter presence in households, there is a need for advanced software tools which can enable the end-user to understand his or her electricity usage. The availability of large amounts of data should enable users to make informed decisions about how best to choose providers and contracts, and how to alter their own electricity usage to lower their bills, thus saving both money, and also, potentially, energy.

In the current paper we presented a prototype, web-based implementation underpinned by the use of argumentation theory and linear optimization, as a step towards this end. The tool combines a large number of visualization and recommendation components, and tests indicate it may be highly successful in reducing power bills. Abstract argumentation was used to bring rational structures underpinning recommendations to the user.

Future work will study automatic methods for the disaggregation of appliances from aggregated data.

REFERENCES

[1] The Observer, "Energy bills rise by 37% in three years," November 2013. [Online]. Available: http://www.theguardian.com/money/2013/nov/16/energy-prices-rise
[2] The Guardian, "Energy uk warns that household bills could rise by 50% over six years," November 2013. [Online]. Available: http://www.theguardian.com/money/2013/nov/12/energy-uk-warns-50-percent-rise-household-bills
[3] ——, "Smart energy meters in every uk home by 2020," May 2009. [Online]. Available: http://www.theguardian.com/environment/2009/may/11/smart-meters-energy-efficiency
[4] The European Commission, "Energy: Commission paves the way for massive roll-out of smart metering systems," January 2012. [Online]. Available: http://ec.europa.eu/energy/gas_electricity/smartgrids/doc/20120309_smart_grids_press_release.pdf
[5] P. M. Dung, "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games," *Artificial Intelligence*, vol. 77, no. 2, pp. 321–358, 1995.
[6] P. Baroni and M. Giacomin, "Semantics of Abstract Argumentation Systems," in *Argumentation in Artificial Intelligence*, I. Rahwan and G. R. Simari, Eds. Springer, 2009, pp. 25–44.
[7] J. Kelly and W. Knottenbelt, "UK-DALE: A dataset recording UK Domestic Appliance-Level Electricity demand and whole-house demand," *ArXiv e-prints*, April 2014.
[8] Department for Environment, Food and Rural Affairs, "Household electricity study, ev0702," 2009–2012. [Online]. Available: http://randd.defra.gov.uk/Default.aspx?Module=More&ProjectID=17359