Evaluating Market User Interfaces for Electric Vehicle Charging using Bid2Charge*

Sebastian Stein, Enrico H. Gerding, Adrian Nedea, Avi Rosenfeld◇, Nicholas R. Jennings•

University of Southampton, Southampton, United Kingdom
◇Jerusalem College of Technology, Jerusalem, Israel
•Imperial College, London, United Kingdom
{ss2,eg}@ecs.soton.ac.uk, adrian.nedea@gmail.com, rosenfa@gmail.com, n.jennings@imperial.ac.uk

Abstract

We consider settings where electric vehicle drivers participate in a market mechanism to charge their vehicles. Existing work typically assumes that participants are fully rational and can report their charging preferences accurately. However, this may not be reasonable in settings with non-experts. To explore this, we design a novel game called Bid2Charge and compare a fully expressive interface that covers the entire space of preferences to two restricted interfaces that offer fewer possible reports. We show that restricting the users' preferences significantly reduces deliberation times while also leading to an increase in utility by up to 70%.

1 Introduction

Market mechanisms designed for multi-agent systems hold considerable promise for addressing emerging challenges in future electricity networks, where supply will increasingly be generated from intermittent and unreliable renewables, and where demand will rise due to the electrification of transportation [Ipakchi and Albuyeh, 2009; Ramchurn *et al.*, 2012]. In particular, previous work has proposed auction-like mechanisms for scheduling the charging of electric vehicles (EVs) within supply constraints [Mets *et al.*, 2012; Robu *et al.*, 2013; Hayakawa *et al.*, 2015]. These achieve a high efficiency because they consider the preferences of drivers (i.e., availability and willingness to pay). Other work relies on real-time prices to incentivise autonomous charging agents to shift or curtail consumption when supply is low [Ramchurn *et al.*, 2011; Flath *et al.*, 2014].

However, such approaches assume that the human end-users have perfect knowledge of their preferences, i.e., they can reason accurately about the value of electricity, considering all possible, often uncertain future opportunities of using it. Such perfect rationality is often not realistic [Simon, 1972; Kahneman, 2000]. Moreover, providing the complete preferences is tedious and the associated cost could outweigh the benefits [Larson and Sandholm, 2005].

To address this challenge, there has been some work on auctions with *restricted reporting*, i.e., where bidders do not report their full preferences, but rather choose from a restricted messaging space [Milgrom, 2010; Dütting *et al.*, 2011; Bergemann *et al.*, 2012; Blumrosen and Feldman, 2013]. However, these approaches all assume rational agents and do not evaluate the auctions with real bidders.

Other work explicitly considers non-expert market participants. Specifically, research on hidden market design has looked at building simple interfaces that hide the rules and pricing mechanisms of a complex market [Seuken *et al.*, 2010]. Related to this, Seuken *et al.* (2012) investigated market user interface design, which focuses on simplifying complex market interactions by asking users to select from a discrete set of options. However, they consider significantly different domains from the EV setting and neither work provides a direct comparison to fully expressive interfaces, leaving their relative benefits unclear.

To address these limitations, we conduct the first thorough study of how to design market interfaces for the EV charging setting. Specifically, we formalise the EV charging setting and develop a novel game called Bid2Charge based on this. We then use this game to experimentally compare a fully expressive interface for reporting user preferences (as assumed by existing work) to two restricted interfaces: one that reduces the dimensionality of the reporting space (but retains infinitely many options) and one that restricts the reporting space to a discrete set of options. In two large user studies involving a total of over 300 participants, we show that the restricted interfaces significantly reduce user deliberation time. More surprisingly, they also improve performance (up to 70%). Furthermore, we show that the choice of restriction can influence the energy consumption of participants without decreasing their utility. Finally, we show that a reinforcement learning agent displays similar behaviour trends to human players, potentially enabling the optimisation of market interfaces without expensive user trials.

2 EV Charging Problem

In this section, we briefly summarise an abstract model of the EV charging problem. The aim of this model is to capture the following key challenges that are inherent in the domain. First, electricity has **no intrinsic value**. Its value instead depends on how it is utilised for journeys. Second, the problem
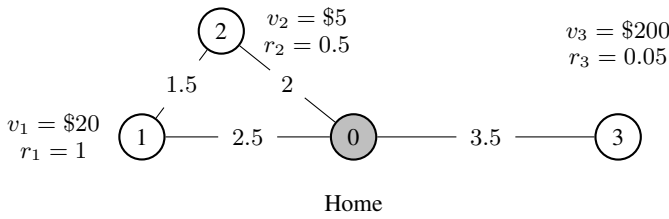

Figure 1: Illustrative example with three potential journeys.

Figure 2: Drivers's true willingness to pay.

is inherently **uncertain**. On one hand, this is a key feature of the market itself, as supply and demand fluctuate over time. On the other hand, there is uncertainty over what journeys the driver needs to complete in the future. Third, there are **complementarities**, i.e., the driver's preferences are typically highly nonlinear over the quantity of electricity. For example, a driver may require a minimum overnight charge of 10 kWh to drive to work the next day. Receiving any less has no value.

Given this, we consider a general setting where EV drivers participate in a market mechanism for charging. We abstract away from the particular market and assume that a driver simply reports her preferences to an autonomous agent at regular intervals (e.g., once a day).[1] The agent then participates in the market on its owner's behalf and procures electricity for the vehicle while it is plugged in (e.g., overnight). At the end of the charging period, the driver can use the vehicle to complete journeys and derives value from doing so.

The full formal model can be found in [Stein *et al.*, 2016; 2017]. Next, we briefly summarise it using an example.

3 Illustrative Example

Specifically, Figure 1 shows an example of the information that is available to an EV driver at the start of a given charging period — this includes a set of three potential journeys to different destinations (labelled 1, 2 and 3), which are at different distances to her home (labelled 0).

There is some uncertainty for each journey i, which is described by a realisation probability r_i (i.e., the probability that this journey will actually be available on that day), and there is a value for completing each journey, denoted by v_i. In this example, journey 1 might represent a trip to the driver's workplace, which is completely predictable, $r_1 = 1$, and carries a high value, $v_1 = \$20$. The second, journey 2, is a visit to the gym, which is of a medium value, $v_2 = \$5$, and less likely to be available, $r_2 = 0.5$ (this may be because she only goes to the gym if she can finish work early). Finally, the third, journey 3, is an emergency trip to an important customer. This is unlikely ($r_3 = 0.05$), but carries a high value ($v_3 = \$200$).

Furthermore, there are electricity costs associated with journeys, given by the edge weights. These represent the drain on the battery for each one-way trip (in kWh). For example, going to the office and back (journey 1) requires 5 kWh, while going to the gym and office requires 6 kWh.

This information represents the (typically tacit) knowledge a driver has about future journeys. As in reality, there are
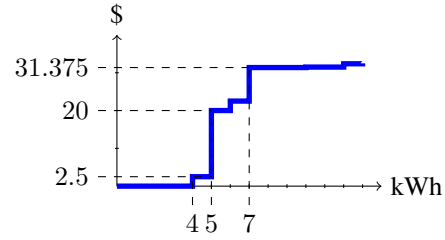
varying levels of uncertainty and values for these journeys, but we also make a number of simplifications to keep the model manageable — specifically, we assume that values and realisation probabilities are completely independent and that electricity consumption is deterministic.

Given this information, the driver next reports her maximum willingness to pay for each amount of electricity to an agent that will participate in the market on her behalf. We assume that it is optimal (in expectation) for her to report truthfully, either because the agent strategises optimally or because an incentive compatible mechanism is used, e.g., [Robu *et al.*, 2013] or [Hayakawa *et al.*, 2015].

However, doing this is challenging, as it requires the driver to reason about all possible combinations of available journeys, their costs and respective values. To illustrate this, the solution the setting above is shown in Figure 2, which relates the amount of electricity to the respective willingness to pay.[2]

Here, the driver has no value for 3 kWh or less, as this is insufficient for any journeys. She values 4 kWh at \$2.5, because that is enough to complete journey 2 for a value of \$5, but the journey's probability of being available is only $r_2 = 0.5$. However, the value for 5 kWh is significantly higher at \$20, because the driver now has enough charge to reach the office, yielding a value of \$20. The next significant increase in the driver's willingness to pay happens for 7 kWh (\$31.375), when she has sufficient charge to complete either journey 3 (if it is available) or both journeys 1 and 2.

Given these valuations, the agent now participates in the market to charge the EV and with the aim of maximising the driver's profit (i.e., the difference between the willingness to pay for the acquired electricity and the cost for this). When this phase is complete, the driver then learns how much the battery was charged and which journeys are actually available. She then chooses which journeys to complete on that day, and the setting repeats (possibly with different potential journeys). The overall aim of the driver is to maximise her actual profit (i.e., the difference between the total value derived and the total paid in the electricity market).

This example illustrates that even a simple setting can lead to complex and complementary valuations. As argued in Section 1, existing work typically assumes that these can readily be reported by drivers. Next, we describe our Bid2Charge testbed, which will allow us to test this assumption by comparing a fully expressive reporting interface with several restricted interfaces when used by non-experts.

---

[1]For example market mechanisms, see [Ramchurn *et al.*, 2011; Robu *et al.*, 2013; Hayakawa *et al.*, 2015].

[2]To keep this example brief, we assume the battery is initially empty and there is only a single day. In practice, a driver would also need to reason about potential journeys on future days.
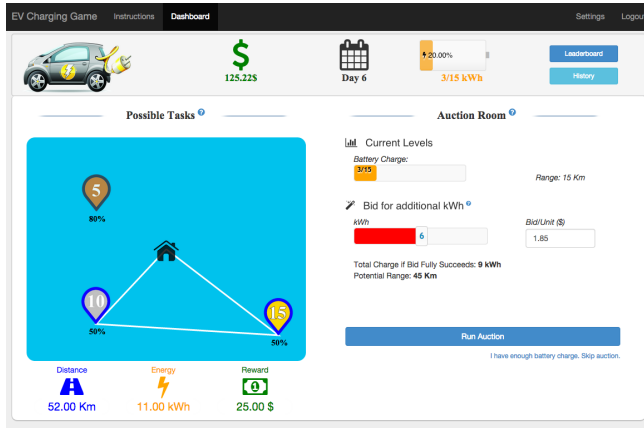
Figure 3: Main game view.

4 The Bid2Charge Testbed

To test how human non-experts interact with a variety of reporting interfaces, and to empirically determine which one works best, we designed a web-based game called Bid2Charge, which replicates the EV charging problem.[3] We frame this as a game, as this is a low-cost, controlled way of gathering data from large numbers of users in a short time.

In more detail, in Bid2Charge, the player takes the role of an EV delivery van driver. This provides an intuitive explanation of what journeys represent (they are delivery tasks resulting in a certain payment), what the objective of the game is (maximise overall profit) and what the uncertainty means (delivery tasks may not come up on a given day). We use an incentive-compatible auction based on the well-known VCG mechanism [Nisan *et al.*, 2007] as the market mechanism, and so the players' reports are framed as bids for this auction.

Figure 3 shows the main game screen. At the top, there are some general statistics, showing the player's accumulated profit, the current day and current state of charge. Below that, on the left, there is a task planning view, providing information about the tasks that are potentially available on the current day. Both the value and the realisation probability are shown for each task. Furthermore, the user can select subsets of tasks to inspect their total value and electricity cost (providing similar information as in the example in Figure 1, albeit presented slightly differently). Note that interacting with this view does not affect the game — it simply provides information about available tasks.

The auction view, which is to the right of the task planning view, allows the player to submit their bid (i.e., willingness to pay) for the current day. Importantly, this view supports three different interfaces, as shown in Figure 4:

FINITE: This is the most restricted interface, offering a finite set of options. In practice, these can be chosen by an administrator to trade off the cognitive burden on the player with the expressivity of the set. Figure 4a shows a version with three options that correspond to a marginal willingness to pay of $1, $2 and $3 per kWh, respectively.

Single Marginal Value (SMV): This allows the player to

(a) FINITE:

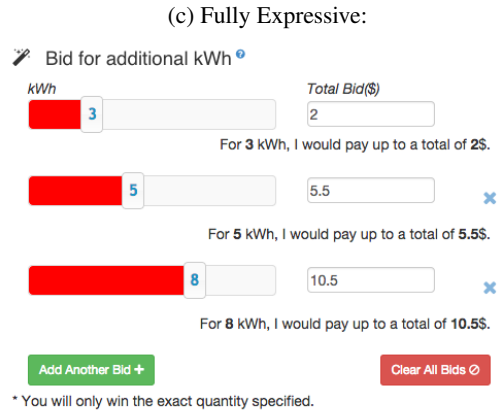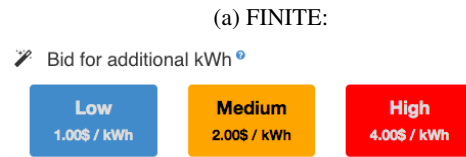(b) SMV:

(c) Fully Expressive:

Figure 4: Auction Bidding Panels.

specify a marginal willingness to pay for each unit of electricity, as well as a maximum number of units. While more expressive than FINITE, it reduces the original problem to reporting two values (rather than a function). Figure 4b shows a player entering the bid corresponding to $1.5 per kWh, but only up to a maximum of 4 kWh.

Fully expressive: Finally, this interface allows the player to enter any bid from the full space of possible bids. Here, Figure 4c shows a bid that assigns a value of $2 for 3 kWh, $5.5 for 5 kWh and $10.5 for 8 kWh.

On pressing the "Run Auction" button, they are informed of the outcome and then taken to a task view. Here, the user is presented with the tasks actually available on the current day and can select which ones to complete.

5 Experimental Evaluation

We carried out experiments with human users, in order to compare the three interfaces discussed above, and to determine whether the choice of a particular interface influences the performance of users. Our approach was to evaluate this through randomised, controlled experiments, where we allocate market interfaces randomly to users.

5.1 Experimental Setup

We recruited players from Amazon Mechanical Turk, a platform that allows requesters to advertise tasks to a large audience of online workers (http://www.mturk.com/). We paid workers for participating, as well as a variable bonus
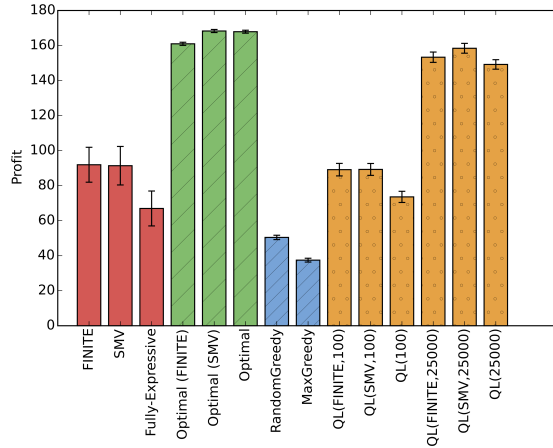
Figure 5: Average profit in first experiment.

that was aligned with a player's profit during the game. In total, we collected data from 319 workers and over two separate experiments. The first focused on general performance differences in the three interfaces (130 workers), while the second looked at learning over repeated plays (189 workers). Due to space reasons, we focus on the first experiment, but further details are in [Stein *et al.*, 2016].

The game was played for 30 simulated days with between 1–6 tasks available every day. To establish upper and lower bounds for the possible performance of human players, we compared them to a number of benchmarks: **Optimal** is the optimal strategy assuming a fully expressive interface. **Optimal (SMV)** and **Optimal (FINITE)** are two variants for the restricted interfaces. As a lower bound, **RandomGreedy** is a benchmark that places a random bid (chosen from the FINITE options) and then greedily chooses the highest-value tasks. Similarly, **MaxGreedy** places a bid that is high enough to fully charge the EV each day and then greedily chooses the highest-value tasks. Finally, in order to evaluate a strategy that might better represent how humans make decisions in this game, we also show results for reinforcement learning agents using Q-Learning [Watkins and Dayan, 1992] for all three interfaces, $QL(\lambda)$, $QL(SMV,\lambda)$ and $QL(FINITE,\lambda)$, where λ is the number of times the game was played.

5.2 Results

We first consider the overall profit achieved, as this is the players' main objective. Figure 5 shows this for the three interfaces with human participants (in red, plain), for the three optimal policies (in green, hatched), the two baseline benchmarks (in blue, finely hatched) and the reinforcement learning agent after 100 and 25000 iterations of the game, which are representative of early and late stages of learning (in orange, dotted). All results are shown with 95% confidence intervals and any significant results are confirmed with appropriate statistical tests [Stein *et al.*, 2016].

Considering the performance of human players, these are generally situated between the optimal and the baseline benchmarks. Unsurprisingly, the humans perform significantly worse than a rational agent. However, when compar-

ing the human players to the baseline benchmarks, there is a marked improvement, providing evidence that participants are putting effort into the game.

When comparing the performance of the human players using different market interfaces, we note that the choice of interface has a significant influence on performance. Counter-intuitively, the players using the fully expressive interface achieve a significantly lower performance, with an average profit of only $66.96, while players with SMV achieve an average profit of $91.38 and players with FINITE achieve an average $91.93. This constitutes an improvement of more than 35% over the fully expressive interface (in our second experiment, this reached 70% [Stein *et al.*, 2016]), which is likely due to the high complexity of the latter.

Considering the reinforcement learning agent, it is interesting to note that after 100 learning episodes, the performance of the reinforcement learning agent resembles that of human players. This indicates that it may be useful for predicting human behaviour (this is further explored in [Stein *et al.*, 2016]).

Although there is no significant difference in profit between FINITE and SMV, players using FINITE spent an average $208.44 acquiring electricity and achieved an average reward of $300.37. In contrast, players using SMV spent only an average of $148.50 and earned $239.89. Thus, while the profit for both is similar, FINITE induces a very different behaviour in the players — they spend significantly more on acquiring electricity and are then able to use this to complete a larger number of tasks. This is an interesting result, showing that significant changes in behaviour can be caused by the right choice of interface.

Finally, we note that players with the fully expressive interface spent the longest time on the auction (689 seconds on average), while those with FINITE spent the least amount of time on it (502 seconds on average), confirming the higher cognitive burden of the fully expressive interface.

6 Conclusions and Future Work

In this paper, we investigated the use of restricted market interfaces to alleviate the cognitive burden on non-expert participants in a market-based EV charging setting. Using a novel framework called Bid2Charge, we showed that such interfaces lead to lower deliberation times and higher performance than fully expressive interfaces. We also found that particular types of interfaces induce different behaviours in participants, while achieving the same utility. This could be a promising tool for nudging people towards desirable behaviours, such as energy conservation. Last, we found that a reinforcement learning agent was able to predict broad performance trends.

In future work, we will build on this framework and develop new optimised interfaces that adapt to users. In particular, we will extend the reinforcement learning agent proposed here to offer advice and act autonomously where appropriate.

Acknowledgements

References

[Bergemann *et al.*, 2012] Dirk Bergemann, Ji Shen, Yun Xu, and Edmund Yeh. Multi-dimensional mechanism design with limited information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 162–178, New York, NY, USA, 2012. ACM.

[Blumrosen and Feldman, 2013] Liad Blumrosen and Michal Feldman. Mechanism design with a restricted action space. *Games and Economic Behavior*, 82(0):424 – 443, 2013.

[Dütting *et al.*, 2011] Paul Dütting, Felix Fischer, and David C. Parkes. Simplicity-expressiveness tradeoffs in mechanism design. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, pages 341–350, New York, NY, USA, 2011. ACM.

[Flath *et al.*, 2014] Christoph M. Flath, Jens P. Ilg, Sebastian Gottwalt, Hartmut Schmeck, and Christof Weinhardt. Improving electric vehicle charging coordination through area pricing. *Transportation Science*, 48(4):619–634, 2014.

[Hayakawa *et al.*, 2015] Keiichiro Hayakawa, Enrico Gerding, Sebastian Stein, and Takahiro Shiga. Online mechanisms for charging electric vehicles in settings with varying marginal electricity costs. In *24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2610–2616, April 2015.

[Ipakchi and Albuyeh, 2009] A. Ipakchi and F. Albuyeh. Grid of the future. *Power and Energy Magazine, IEEE*, 7(2):52–62, 2009.

[Kahneman, 2000] Daniel Kahneman. A psychological point of view: Violations of rational rules as a diagnostic of mental processes (commentary on stanovich and west). *Behavioral and Brain Sciences*, 23:681–683, 2000.

[Larson and Sandholm, 2005] Kate Larson and Tuomas Sandholm. Mechanism design and deliberative agents. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and MultiAgent Systems*, pages 650–656, 2005.

[Mets *et al.*, 2012] K. Mets, R. D'hulst, and C. Develder. Comparison of intelligent charging algorithms for electric vehicles to reduce peak load and demand variability in a distribution grid. *Journal of Communications and Networks*, 14(6):672–681, Dec 2012.

[Milgrom, 2010] Paul Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70(1):62 – 70, 2010. Special Issue In Honor of Ehud Kalai.

[Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge, 2007.

[Ramchurn *et al.*, 2011] Sarvapali D Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R Jennings. Agent-based homeostatic control for green energy in the smart grid. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(4):35, 2011.

[Ramchurn *et al.*, 2012] Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R. Jennings. Putting the 'smarts' into the smart grid: A grand challenge for artificial intelligence. *Commun. ACM*, 55(4):86–97, April 2012.

[Robu *et al.*, 2013] Valentin Robu, Enrico H. Gerding, Sebastian Stein, David C. Parkes, Alex Rogers, and Nicholas R. Jennings. An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *Journal of Artificial Intelligence Research*, 48:175–230, 2013.

[Seuken *et al.*, 2010] Sven Seuken, Kamal Jain, Desney S. Tan, and Mary Czerwinski. Hidden markets: UI design for a P2P backup application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 315–324, New York, NY, USA, 2010. ACM.

[Seuken *et al.*, 2012] Sven Seuken, David C. Parkes, Eric Horvitz, Kamal Jain, Mary Czerwinski, and Desney Tan. Market user interface design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 898–915, New York, NY, USA, 2012. ACM.

[Simon, 1972] Herbert A. Simon. Theories of bounded rationality. In C. B. McGuire and R. Radner, editors, *Decision and organization: A volume in honor of Jacob Marschak*, pages 161–176. North-Holland, 1972.

[Stein *et al.*, 2016] Sebastian Stein, Enrico H Gerding, Adrian Nedea, Avi Rosenfeld, and Nicholas R Jennings. Bid2charge: Market user interface design for electric vehicle charging. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 882–890. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[Stein *et al.*, 2017] Sebastian Stein, Enrico H. Gerding, Adrian Nedea, Avi Rosenfeld, and Nicholas R. Jennings. Market interfaces for electric vehicle charging. *Journal of Artificial Intelligence Research*, 2017. (in press).

[Watkins and Dayan, 1992] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.