Optimal Vehicle to Grid Regulation Service Scheduling

Christian Osorio

Introduction

With the growing popularity and market share of electric vehicles comes several opportunities for electric power utilities, vehicle owners, and vehicle manufacturers. Vehicles have a tremendous capacity to provide services to better the electric grid. These services are the focus of so-called Vehicle to Grid (V2G) technologies. V2G represents the idea that electric vehicles can not only charge from the grid, but also send power back to the grid when it is advantageous to do so.

An area of focus for V2G technologies is the ability of fleets of electric vehicles to provide ancillary services to the grid. Because power must be consumed at the time it is produced, ancillary services make up the difference between the scheduled amount of power generated, and the actual amount of power consumed. In order to maintain a 60 Hz grid frequency, regulation services are employed. Regulation service providers are called upon by utility companies to increase or decrease grid power in order to regulate the grid frequency. Power sources that can be switched quickly are valuable for ancillary services. Electric vehicles are well suited for this purpose.

V2G implementation poses several problems. Contracts for regulation services must be scheduled one day ahead of time. As electric vehicles are not static generators; they will not always be available to provide ancillary services. Also, the vehicles need to be available when needed by their drivers. Because of this, I propose a machine learning algorithm solution to schedule both charging and regulation services for a fleet of distributed, consumer-owned electric vehicles.

Scheduling electric vehicle charging poses several problems. The aim is to predict when vehicles will be plugged into their home chargers, and optimize the schedules to charge when electricity is least expensive, and provide regulation services when it is most lucrative, all while charging the vehicle to an acceptable level by the time the driver needs to use it.

Ideally, one would like to be able to predict the ancillary service (AS) clearing prices, electricity prices, and vehicle availabilities, then schedule the charging and regulation services as a constrained linear optimization problem. Unfortunately, there are several difficulties with this approach. First, it is not clear that all of the required features are predictable. In fact, it seems highly unlikely that the AS clearing prices can be predicted. Power utilities schedule electricity production to track actual consumption by using algorithms to predict the consumption. These algorithms take into account a variety of features, and are very good at scheduling electricity production. The AS clearing prices are a response to the error of their predictions. For this reason, the AS clearing prices do not appear to be correlated with any of the features. If the AS prices were predictable, the profit margins of the utility could be increased by raising or lowing electricity production based on the predicted AS prices. Another problem with scheduling via linear programming is that there will be error in the predictions. These errors could lead to cars being charged below an acceptable level, and failure to meet ancillary service contracts. Charging and regulation services should be scheduled taking safety margins on the predictions into account.

Method

In order to schedule the charging and regulation services of electric vehicles, I propose a bottom up approach. First, the availability of each individual vehicle is predicted. The hourly electricity prices are also predicted. Reinforcement learning is used to schedule charging and regulation services for each vehicle. After creating the individual schedules, these schedules will be aggregated at a regional level. It will also be necessary to develop a control system to fulfill regulation service contracts and manage the charging of vehicles. For this project, I focus on the first three steps: predicting the availability of vehicles, predicting electricity prices, and using reinforcement learning to create individual schedules.

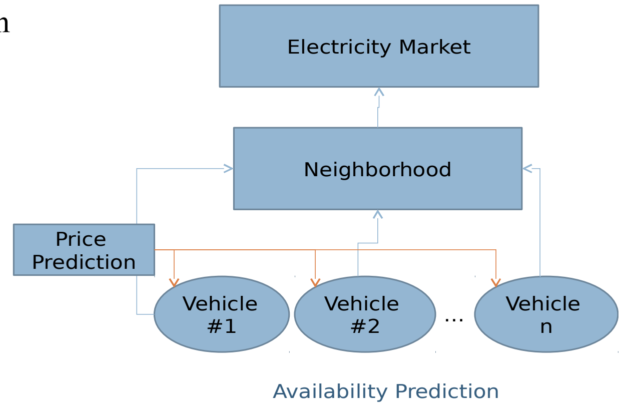

Electricity Price Prediction

Electricity price prediction is a commonly researched problem. In a paper by Raquel Gareta et. al[1], hourly electricity prices are forecasted using neural networks. Gareta uses features including the day of the month, a flag to indicate weekdays or weekends, and prices from the previous days. In one of their most basic models, they use a neural network with 40 hidden nodes.

I decided to use different features than the Gareta paper suggests. Instead, I used one year of hourly temperatures and dew points obtained from Weather Underground[2]. The weather data was collected at the San Francisco Airport. I also created a binary matrix indicating the days of the week and the hour of the day. The elements combined to form 33-dimensional input vectors .The neural networks have a single output node for the hourly electricity price. I obtained one year of Locational Marginal electricity prices (LMP) for a grid node near the San Francisco Airport from the California Independent System Operators (CAISO) website [3]. MATLAB was used to create and train neural networks with various numbers of hidden nodes. The data set was randomly split into training, validation and test sets. The training data represented 70% of the data points, the validation 15%, and the test data 15%. Each network was trained three times, and the average and minimum mean squared errors of the network versus the training set were collected. These values are plotted below. From the plot we see that there isn't significant improvement in performance of the neural network after 30 neurons.

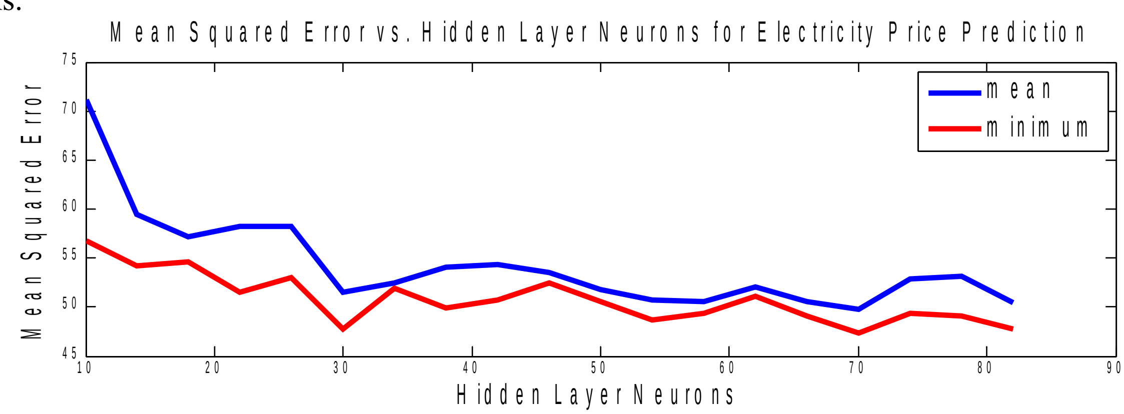

Because of the lack of improvement past 30 neurons, I selected a neural network with 40 hidden layer neurons. After retraining a few times, the network achieved an R value of .81 and a mean squared error of 53. this neural network is .81, and the mean squared error is 60.2.

Strangely, the mean squared error of the network seems unusually high. It seems that due to the size of the data set, MATLAB has some problems checking the mean of the data. However, by plotting a fit to the data, we can see that the neural network seems to be working well. Below, I've plotted the neural network prediction of the LMP for the first week of the year, as well as the actual prices. In the future, I hope to obtain better metrics to measure the performance of the neural network.

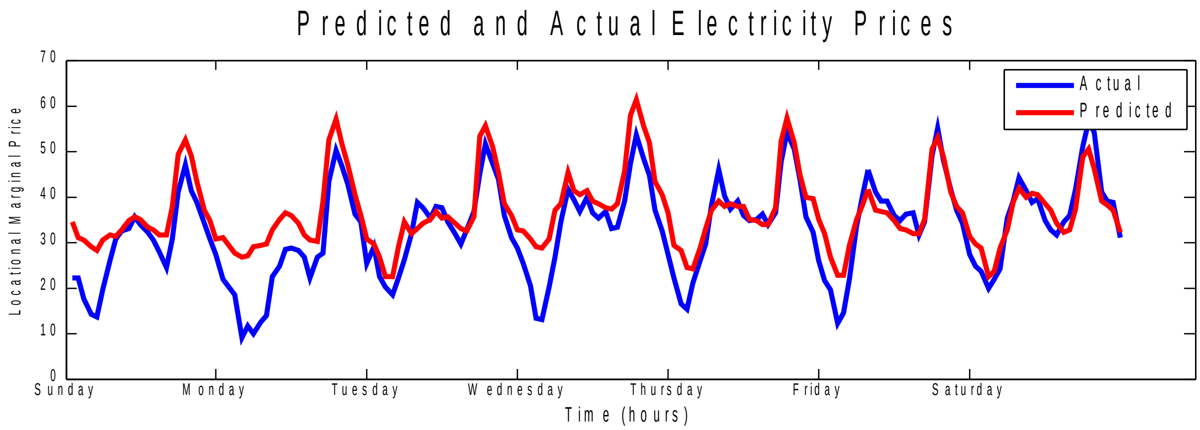

Availability Predictions

After predicting electricity prices, the next step towards planning a vehicle-to-grid schedule is predicting when vehicles are available. For this problem, I was provided a data set of charging histories of ten electric cars from BMW. The data set contains logs of charging events for the vehicles for a five month period. The data indicates when a vehicle is plugged in, when it is disconnected. It also indicates the mileage of the vehicle and the state of charge of the battery at each charging event. This data is used to create a binary vector indicating whether the vehicle is plugged-in or not each hour. Because some data is missing, the missing hours are skipped.

In order to predict when vehicles will be plugged in, the prediction is treated as a classification problem. Unfortunately, the data is anonymized, so that I know nothing about the customer. Because of this, I am unable to use any data based on the location of the vehicle. The features used to classify the availability of the vehicle are represented by a binary matrix indicating the day of the week and the hour of the day. For this reason, we have a 31 dimensional feature space.

I first attempted to classify the data using support vector machines (SVM). I trained an SVM with a linear kernel using the first 1000 data points from the first vehicle, and achieved a classification error of 15%. However, the process was very slow. The data did not separate very well in the 31 dimensional feature space, and out of 1000 data points, there were 806 support vectors. Because the application is made to be real time, as time goes on, the data set will grow. When attempting to train the SVM on more data points, the computer ran out of memory. Because of this, an SVM would not make a good candidate for the final implementation of the program.

Instead of the SVM, a naive Bayes classifier was used to classify the data. The naive Bayes algorithm took much less memory to calculate, and was able to be trained much faster. For the first vehicle, trained on 1000 data points, the naive Bayes classifier achieved a classification error of 19%. However, when trained on a 3000 data points, the classifier achieved a classification error of 13%.

Unfortunately, some drivers are more predictable than others. While the first vehicle becomes fairly predictable with more data, other drivers are not. The Naive Bayes classifier for another driver in the data set misclassifies the availability 35% of the time. The unpredictable nature of some drivers may indicate that they are poor candidates for vehicle to grid services.

Scheduling with Reinforcement Learning

In order to accommodate error in the predictions of the availabilities of the vehicles, a reinforcement learning algorithm will be developed to optimize the margin of safety of scheduling the charging and regulation services. The schedules for regulation services and vehicle charging need to be decided upon at least one day ahead of time for a 24 hour period.

In order to frame the problem, the data is separated into 24 hour periods. During a period, charging periods are selected by the availability predictor. Actions are decided upon for each charging period. The charging period will have an expected length, which will be used to determine the action taken.

For each day, the outcomes of all actions are evaluated. A value matrix keeps track of the expected rewards of a each action based on the expected length of the charge period. If the vehicle is charged to a target value by the time the vehicle is unplugged, it is assumed that the vehicle owner is satisfied with the action taken during the charge period. For this, the reward will be the sum of a bonus for satisfying the customer, and the net revenue during the charge period. If the customer is unsatisfied, the reward function is just the net revenue.

Because the space of possible actions is so large, the actions will take the form of different charging policies adjusted for different safety margins. An example of a policy is maximizing electricity purchased when electricity is projected to be the least expensive. Within an expected charge period, we can choose to limit the charge period of the policy as a built in safety margin. The value matrix is parametrized by the actions taken, and the safety margin. Actions are decided upon by the highest expected reward for a given length of an expected charge period.

The proposed reinforcement learning algorithm will be my next undertaking in the project. Unfortunately, I was not able to implement the algorithm in time for this paper. However, I believe that the framework will produce a viable reinforcement learning problem.

Conclusion

Scheduling electric vehicle charging for a day-ahead market is a complex problem with many moving parts. The steps taken in this project will help to push the problem closer to a solution. Previous studies had already shown that electricity prices can be accurately predicted using neural networks. I implemented a neural networks predictor in hopes that I could use in in the future with a reinforcement learning based scheduler. The Naive Bayes classifier worked fairly well for predicting when drivers are available. It shows that some drivers may be better candidates than others for day ahead scheduling. In implementing vehicle to grid technologies, many vehicles will have to be aggregated to provide enough power capacity to make a difference in the market. Hopefully, the aggregated predictions will be even more accurate. While I was not able to finish implementing the reinforcement learning based scheduler, I am well on the way, and hope to complete this soon. By demonstrating these methods, I hope to demonstrate that vehicle-to-grid technologies can be profitable and useful.

References

1.) Raquel Gareta, Luis M. Romeo, Antonia Gil, Forecasting of electricity prices with neural networks, Energy Conversion and Management, Volume 47, Issues 13–14, August 2006, Pages 1770-1778, ISSN 0196-8904, 10.1016/j.enconman.2005.10.010.

2.) Weather Underground. www.wunderground.com

3.) California Independent System Operator. www.caiso.com.

Thanks

Thank you to the BMW group for providing me with data from their test fleet, and to the CS 229 course staff for a great quarter.