



HAL
open science

Sensitivity to forecast errors in energy storage arbitrage for residential consumers

Diego Kiedanski, Umar Hashmi, Ana Bušić, Daniel Kofman

► **To cite this version:**

Diego Kiedanski, Umar Hashmi, Ana Bušić, Daniel Kofman. Sensitivity to forecast errors in energy storage arbitrage for residential consumers. SmartGridComm 2019 - IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids, Oct 2019, Beijing, China. hal-02163114v2

HAL Id: hal-02163114

<https://telecom-paris.hal.science/hal-02163114v2>

Submitted on 14 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sensitivity to Forecast Errors in Energy Storage Arbitrage for Residential Consumers

Diego Kiedanski*, Md Umar Hashmi[†], Ana Bušić[†], Daniel Kofman*

*Telecom Paristech, 23 Avenue d'Italie, Paris, France
Emails: {diego.kiedanski, daniel.kofman}@telecom-paristech.fr

[‡] INRIA, DI ENS, École Normale Supérieure, CNRS, PSL Research University, Paris, France
Email: : md-umar.hashmi@inria.fr

Abstract—With the massive deployment of distributed energy resources, there has been an increase in the number of end consumers that own photovoltaic panels and storage systems. The optimal use of such storage when facing Time of Use (ToU) prices is directly related to the quality of the load and generation forecasts as well as the algorithm that controls the battery. The sensitivity of such control to different forecast techniques is studied in this paper. It is shown that good and bad forecasts can result in losses in, particularly bad days. Nevertheless, it is observed that performing Model Predictive Control (MPC) with a simple forecast that is representative of the pasts can be profitable under different price and battery scenarios. We observe that performing MPC at a faster sampling time with a receding optimization horizon makes arbitrage less sensitive to uncertainties in forecasting. We use real data from Pecan Street and ToU price levels with different buying and selling price for the numerical experiments.

Index Terms—Sensitivity, Arbitrage, Smart Grids, Battery, Load Forecast, Model Predictive Control

I. INTRODUCTION

With new technological advances and an increase in Distributed Energy Resources (DER), exploiting end customer's flexibility becomes a feasible and important task. A major policy change in this direction was allowing end customers to sell their energy surplus and inject it to the grid.

In the face of such changes, a model that has grown in popularity is that of a household with a photovoltaic panel, a battery (that can be shared between neighbours [1]) and possibly an electric vehicle (EV) [2], [3]. The model requires an algorithm to operate the battery in the most efficient way possible. This usually means maximizing self-consumption.

The above model is similar to a pure arbitrage problem (where the profits depend on the fluctuations of the market's price) with two differences. First, the household has an inflexible load that has to be satisfied at all moments. Second, residential households are rarely faced with real time prices. Moreover, the types of tariffs faced by these customers consist, most of the time, of a selling price that is lower than the buying price at all moments (that is the case for the Parisian tariff for buying electricity, *heures creuses*, which is always greater than the selling price).

This setting, in which an end customer owns a battery and faces a Time of Use (ToU) tariff has been studied in [4], where they use Model Predictive Control (MPC) paired with a persistence forecast to control the battery. The same setting was used in [5], but the Reinforcement Learning technique used by the authors requires only a forecast for the next time slot, which they model as the true value plus a Gaussian noise. Variations of the problem when price changes are dynamic have already been studied [6]–[9].

If the household is subject to a Time of Use, the objective of arbitrage is to shift the consumption from the expensive period into the low price one [10]. As in the arbitrage model, stochasticity might prevent from optimally utilizing the battery. While in the pure arbitrage case the uncertainties are in the prices, for households they originate from renewable generation and consumption. This is an important point because unlike market price prediction which has received a lot of attention from many disciplines or forecasting the aggregation of several loads [11]–[15]; forecasting the load of a single household is a difficult problem [16], [17]. Not only it is hard, but also expensive: unlike a market forecast that can be used by many, the forecast of a residential household has to be done independently for each house. Nevertheless, the literature in forecast of residential household is vast, see for example [18]–[20].

Sensitivity analysis for storage systems when the variability comes from the price (real time market prices) but ignore the load have been studied thoroughly: [21]–[24]. Authors in [25], include a forecast of the load in addition to prices, but do not consider its accuracy in the overall cost during their analysis, they mostly focus on the length of the forecast horizon and the price forecasts. Furthermore, it has already been shown that if the price of buying and selling are the same, then the load variation does not affect at all the operation of a battery [26]. For cases where selling price is not equal to buying price, the storage decisions are governed by price and load variations.

Studies regarding sensitivity in the operation of a battery with respect to its different characteristics have been carried in the past. In [27], the authors perform a sensitivity analysis on the profitability of PV-storage systems for households but do not consider the impact of forecasts, which is the main goal in

this work. A sensitivity analysis in the battery characteristics together with the evaluation of different optimization techniques to solve the problem has also been studied in [28]. In this work we consider energy storage arbitrage under ToU price and Feed in Tariff for excess production. Under such a case, storage operation is affected by load variations. This work aims to answer how important is the forecast of the load when the tariff is a Time of Use. We use data from Pecan Street, real ToU prices and forecast of varying degree of precision in order to obtain an estimate of the importance of such forecasts. We provide important insight on the relationship of the predictions and battery usage that will help future practitioners decide the accuracy of the model they are willing to work with.

The rest of the paper is organized as follows: in Section II the mathematical model of the problem is formulated. Section III describes the experiment setup which includes the forecasting techniques, the evaluation metrics and the price structure used. Numerical results with the appropriate discussion are presented in Section IV. Finally, Section V includes closing remarks.

II. MODEL DESCRIPTION

This study focuses on households that own a battery. Such battery will be modeled by the difference equation (1) in discrete time:

$$b_{t+1} = b_t + x_t, \quad (1)$$

where b_i is the State of Charge (SoC) of the battery at time i and b_0 , the initial SoC, is given. In Equation (1), $x_t \in \mathcal{X}(b)$ denotes the energy charging (positive) or discharging (negative) the battery. The set $\mathcal{X}(b)$ of feasible actions depends on the SoC and is defined in Equation (2) where $\delta_{\max} > 0$ and $\delta_{\min} < 0$ denote the maximum charging and discharging rate (in power), h is the length of each time interval and b_{\max} and b_{\min} are the maximum and minimum SoC of the battery.

$$\mathcal{X}(b) = [\max\{\delta_{\min}h, b_{\min} - b\}, \min\{\delta_{\max}h, b_{\max} - b\}] \quad (2)$$

Because we consider efficiency losses, the actual energy required to charge the battery an amount x_t is $\frac{x_t}{\eta_c}$ ($\eta_c \in (0, 1]$) while the energy seen from outside the battery when it discharges x_t is $x_t\eta_d$ ($\eta_d \in (0, 1]$). In what follows, we will denote the energy seen from the outside of the battery s_t .

Each time slot t , consumers are faced with a buying price p_t^b and a selling price p_t^s . The inflexible load of the consumer at time t is denoted z_t (a positive value represents consumption while a negative value represents surplus of generation).

For a given and finite time horizon T , the consumer is interested in solving the optimization problem given as

$$P) \underset{x_0, \dots, x_{T-1}}{\text{minimize}} \quad f(x) = \sum_{i=0}^{T-1} p_i^b [s_i + z_i]^+ - p_i^s [s_i + z_i]^- \quad (3a)$$

subject to

$$b_{j+1} = b_j + x_j \quad j = 0, \dots, T-1, \quad (3b)$$

$$x_j \in \mathcal{X}(b_j) \quad j = 0, \dots, T-1 \quad (3c)$$

where $[\cdot]^+ = \max\{\cdot, 0\}$ and $[\cdot]^- = \max\{-\cdot, 0\}$.

Problem P) was studied in [26], [29] where the authors proved that the optimal policy has a threshold structure and propose an algorithm to solve the optimization problem efficiently.

The solution of the optimization problems depends on the battery characteristics, the price of electricity (both of which are known) and on the electricity load profile z . Denote $x^* = x^*(z)$ the optimal solution of the optimization problem when the load profile z is used as input and $f^* = f(x^*(z))$ (f is the electricity cost incurred by a household in a day). If instead of the original profile z only an approximation of it is available (\hat{z}), we seek to understand how x^* and $\hat{x} = x^*(\hat{z})$ are related, as well as f^* and $\hat{f} = f^*(x^*(\hat{z}))$.

A. Model Predictive Control

A natural idea in this kind of iterative problems where information becomes available as time goes by, is to use a receding horizon and solve the problem iteratively, as new information arrives. In the next sections we will make use of Model Predictive Control (MPC). Here we briefly describe some of the conventions used. Let $j \in \{0, 1, \dots, T-1\}$ index the time slots of a given day. For $j = 0$, problem P) is solved using $[z_0 \hat{z}_{1:T-1}]$ as the load profile (where $[\cdot]$ denotes the concatenation of the two vectors). We denote the optimal solution x^{0*} . In general, at timeslot j , problem P) is solved using $[z_{0:j} \hat{z}_{j+1:T-1}]$ as the load profile and also imposing that the first j coordinates of x^{*j} are the decisions already taken, i.e., $x_{0:j-1}^{*j} = x_{0:j-1}^{*j-1}$. We refer to this as MPC(1).

In this work, we also consider MPC with updates that can be made only at time steps $j \in \{0, d, 2d, \dots, T-d\}$, where d is a divisor of T . We will denote this version of MPC with decreased sampling time as MPC(d).

Because we will sometimes need the error of the forecast used while doing MPC(X), we will adopt the following convention: the error associated with MPC(X) is the weighted average of all the errors of the forecasts used, where each weight corresponds to the number of decisions taken with it.

III. EXPERIMENT SETUP

In this section we will describe the experiment design and the parameters used. This involves the tariff structures, the battery parameters, the consumption profiles, the forecasting techniques and the algorithms to solve the minimization problem. Whenever it is possible, we try to use real data as to get "practical" results.

To begin with, the Parisian two-period ToU *heures creuses* was used. During the cheap period *heures creuses* (hours 23:00 to 7:00), the price of energy is 12.3c/KWh and during the expensive period *heures pleines* (hours 7:00 to 23:00) the price is 15.8c/KWh. Moreover, photovoltaic energy can be sold back to the grid at the *tarif rachat* or Feed in Tariff (FIT), which as of March 2019 is 10c/KWh. In order to understand how different tariffs affect the profitability of arbitrage we considered four additional tariffs summarized in Table I, based on the Parisian ToU. In all cases, the low price period is from 23:00 to 7:00.

The battery parameters are modeled after Tesla's Powerwall 2 [30], arguably the most well known storage product in the market in the Spring of 2019. The parameters of the battery are

TABLE I
TARIFFS USED IN THE EXPERIMENT

Name	Price Low	Price High	Selling Price (FIT)
P1	12.3	15.8	10
P2	12.3	15.8	5
P3	12.3	15.8	1
P4	12.3	20	10
P5	12.3	20	5

described in Table II. It has been shown that the performance of a battery doing arbitrage with real time prices depends on the speed of charging and discharging with respect to battery charge capacity [31]. To take this phenomenon into account, we also used a modified versions of the Powewall 2 with $B_{\max} = 25$ kWh and $\delta_{\max} = -\delta_{\min} = 100$ kW. We will refer to this new battery as **Fastbat**.

TABLE II
POWERWALL PARAMETERS

δ_{max}	δ_{min}	b_{max}	b_{min}	η_c	η_d
5 kW	-5 kW	13.5 kWh	0	0.95	0.95

For the user profiles, data was obtained from the PecanStreet project [32]. Data from two users: one with solar generation (id 5403) and another one without it (id 821) was used. The granularity of the data used was 15 minutes and for the experiments, only weekdays from July 2015 were considered (all weekends were removed from the dataset). In what follows, when referring to a previous day, we mean the previous day in the dataset, i.e., the previous day of Monday is Friday.

To run the experiments we solved problem P) using Model Predictive Control with 4 different sampling times: every 15 minutes (MPC(1)), every hour (MPC(4)), every 12 hours (MPC(48)) and only once every day (MPC(96)). We expected to see a natural trade-off between accuracy and increased computational time.

To generate forecasts, the four following methods were used:

- 1) $\hat{z}(d) = z(d-1)$
- 2) $\hat{z}(d) = z(d) + w$; $z, w \in \mathbb{R}^T$; $w_i \sim \mathcal{N}(0, \Delta)$ i.i.d
- 3) $\hat{z}(d) = \frac{1}{|S|} \sum_{j \in S} z(s)$, $S = \{j: j \equiv d \pmod{7}, j < d\}$
- 4) $\hat{z}(d) = \text{SARIMA}(2, 1, 1)(0, 1, 1, 96)[z(d-1), \dots, z(d-4)]$

where the notation $z(d)$ means the real profile of the day d ¹ and Δ is exactly half of the the 3rd quantile of the user consumption (in absolute value). Method 1 is the standard *persistence forecast* or also known as day ahead. It assumes that the profile is going to be exactly as the real profile of the day before. Method 2 (*Gauss*) corresponds to adding additive noise equally distributed and independent to all coordinates. We emphasize that we do not consider this to be a realistic forecast (as it requires knowledge of the future), but it acts as a placeholder for more accurate forecasts. Method 3 (*AvgPast*), consists on an average of all past days in the dataset that fall in the same weekday, i.e., the average of all consumption on the past Mondays, Tuesdays, etc. Finally, for the last method we trained a Seasonal ARIMA (*SARIMA*) model. Each time a forecast had to be issued, the previous four

¹We enumerated all days in the dataset starting form 0.

days in the dataset were used to fit the parameters of the model and an out-of-sample forecast was produced.

To evaluate the quality of the forecasts, we used two methods usually encountered in the literature: the Mean Absolute Difference / Mean (MAD/MEAN) ratio, an adjusted version of the Mean Absolute Percentage Error (MAPE) that is useful when the actual values might be zero or close to it (defined in Equation 4, [33]) and the Normalized Root Mean Squared Error (NRMSE) (defined in Equation 5), where \bar{z} denotes the average of the coordinates of $z_{T-k+1:T}$. The advantage of using the normalised version of RMSE is that it is comparable among datasets.

$$\text{MAD/MEAN}(\hat{z}_{T-k+1:T}, z_{T-k+1:T}) = \frac{100\%}{k} \sum_{i=T-k+1}^T \left| \frac{\hat{z}_i - z_i}{\bar{z}} \right| \quad (4)$$

$$\text{NRMSE}(\hat{z}_{T-k+1:T}, z_{T-k+1:T}) = \sqrt{\frac{1}{k} \sum_{i=T-k+1}^T \left(\frac{\hat{z}_i - z_i}{\bar{z}} \right)^2} \quad (5)$$

The NRMSE of the forecasts used in the dataset is illustrated in Figure 1. For each of the 22 days considered in the experiment, we include 16 forecasts, one for each combination of forecasting technique and sampling time. To obtain the forecast error for a sampling time lower than 96, the weighted average as described in Section II.A was used.




Fig. 1. NRMSE for the forecasts used in the study

As a benchmark, we compare the results obtained with [16], where the authors show that their forecasts have a NRMSE of approximately 2.

IV. NUMERICAL RESULTS

In this section the numerical results obtained are presented. First, we will look at the opportunity of arbitrage with perfect information. Second, we will show how the different sampling and forecasting techniques affect the earnings obtained as well as the optimal battery actions. Third, the relationship between the forecast error metrics and the arbitrage gains is studied. Finally, a profitability study is included.

A. Arbitrage opportunity with perfect information

The arbitrage opportunity with perfect information (AO) can be defined as the difference in electricity cost in a single day between a scenario without battery (CwB) and a scenario with a

battery and perfect information (CBPI). This metric will enable us to better understand the results obtained by performing arbitrage. Moreover, this metric is sufficient to determine whether the investment in a battery is profitable. As the AO is the optimal case, if a particular storage is not profitable under these circumstances, it will not be at all.




Fig. 2. AO using a Powerwall 2 for households: 5403 on top, 821 below.

Figure 2 shows the AO for the two household (5403 with generation on top and 821 without generation on the bottom). For the household with id 5403, it can be seen that the AO increases as the selling price decreases ($P1 \rightarrow P2 \rightarrow P3$). This effect can be explained as follows: when the FIT is high enough, the advantage between storing for later consumption and selling to the grid and re-buying later is less compared to the scenario when the selling price is very low. On the contrary, for a consumer without generation, a high selling price helps mitigate forecast errors: if energy was bought but it was not needed, it can be resold again for a small loss.




Fig. 3. AO using **Fastbat** for households: house id 5403 on top, 821 below.

Figure 3 shows the equivalent results when **Fastbat** is used instead. It can readily be seen that while the overall gains increase, the trend persists. Also, the difference between scenarios P4 and P5 with the other three is easily explained. In those cases, the gap between low buying price and the high buying price is bigger and so the AO is bigger as well. Recall that for a ToU tariff, arbitrage is mostly equivalent to switching consumption from the expensive period to the cheap one.




Fig. 4. AOF for house id 5403 for some days under tariff P1. Each row corresponds to a different forecast technique. From top to bottom: persistence, Gaussian, AvgPast and SARIMA.




Fig. 5. AOF for house id 821 during some days under tariff P1. Each row corresponds to a different forecast technique. From top to bottom: Persistence, Gaussian, AvgPast and SARIMA.

B. Arbitrage opportunity without perfect information

Having assessed the potential of arbitrage for the different prices and households, we turn our attention on how those benefits are affected when forecasts are used instead. For household 5403, these results while using a Powerwall 2 are shown in Figure 4. In it, each bar represents the cost incurred using a forecast (CF) minus the cost of having perfect information (CBPI) while the black line depicts the cost of not having a battery (CwB) minus the CBPI. In this way, all results between the x-axis and the black curve imply that profit was obtained when using the forecast while bars above the black line stand for days in which **doing nothing would have been a better choice**. As it can be seen in Figure 4, July 24 th is a particularly bad day. Only with a very accurate forecast losses could have been avoided.

In Figure 4, each subplot shows the result for a different forecasting technique. For top to bottom: persistence, Gaussian, AvgPast and SARIMA. Moreover, each color stands for the frequency of sampling used for running MPC. The blue bars (1) are standard MPC taking actions every timeslot while the purple bars (96) are cases when all actions were decided at the beginning of the day. As it was to be expected, the most

frequent the sampling time is while running MPC, the lower the incurred cost is. There is a natural trade-off between CPU usage and accuracy. It can be observed that using a more frequent sampling is more important than using a good forecast.

In contrast with household 5403 who is a producer, household 821 can operate closer to optimality using less accurate forecasts. This is evidenced in Figure 5 where the same information as in Figure 4 is plotted for household 821. MPC(1) performs remarkably well.

Turning our attention to the optimal actions, for household 5403 and the forecasting technique AvgPast, Figures 6 and 7 plot the different usages of the battery for days 1st and 24th of July respectively.




Fig. 6. Differences in optimal battery operation during the 1st of July using the AvgPast forecast for household 5403

In Figure 6, only MPC(96) differs strongly from the real data at the beginning of the day. During July 1st all sampling frequencies were profitable using the AvgPast forecast. In Figure 7 however, it can be seen that all techniques made the wrong decision at the beginning and near the end of the day, resulting in an overall loss. This can happen if the day in question is very different than the previous ones.




Fig. 7. Differences in optimal battery operation during the 24th of July using the AvgPast forecast for household 5403

C. Sensitivity of the gains to NRMSE and MAD/MEAN




Fig. 8. Average savings as a function of the forecasting error metrics.

In a sensitivity analysis, it is natural to consider how does the output vary in response to the input. In Figure 8, we investigate the arbitrage gains as a function of the precision of the forecast. In the x-axis the two error metrics are plotted: MAD/MEAN in the left and NRMSE in the right. The y-axis represents the linear mapping defined as

$$G(\hat{z}) = \frac{\hat{f} - f^*}{CwB(z) - f^*} \quad (6)$$

It is 0, when the forecast performs as having perfect information, 1 when it performs as not having a battery and values greater than 1 are worst than doing nothing. Observe that although there is not a clear correlation between both axes, it is possible to find a threshold (vertical bar) for which all forecast with less error than the threshold (to the left of it) perform better than not having a battery (below the horizontal line).




Fig. 9. Savings using the Powerwall 2 battery for household 5403 and price scenario P1

D. Profitability analysis for Powerwall 2

We conclude the analysis with the evaluation of the average performance of each forecasting technique and each sampling




Fig. 10. Savings using the **Fastbat** for household 5403 and price scenario P1




Fig. 11. Arbitrage opportunity for the extended dataset

TABLE III

NUMBER OF YEARS THAT IT WOULD TAKE HOUSEHOLD 5403 TO RECOVER ITS INVESTMENT IN A POWERWALL 2 FOR THE DIFFERENT PRICE SCENARIOS AND FORECASTING TECHNIQUES

	Persistence	Gaussian	AvgPast	SARIMA	Perfect
P1	89.0	58.0	83.0	94.0	54.0
P2	51.0	31.0	45.0	63.0	29.0
P3	39.0	23.0	33.0	49.0	21.0
P4	26.0	19.0	27.0	26.0	18.0
P5	21.0	15.0	21.0	24.0	14.0

frequency using MPC for household 5403. Figure 9 and 10 depict the savings obtained by having a battery with respect to not having one in a given day for Powerwall 2 and Fastbat, respectively. As expected, gains decrease with less sampling, but interestingly for a simple forecast as AvgPast, it would be possible to use a less coarse sampling time and still obtain profits. Regarding both batteries, it can be seen that while SARIMA and the persistence forecast are not profitable with the Powerwall 2, they are with the **Fastbat**. Finally Table III shows the number of years it would take consumer 5403 using MPC(1), for the different forecasting techniques and price scenarios, to recover the investment in a Powerwall 2². Unfortunately, it does not seem likely that under any scenario the investment will be profitable. With perfect information and the best price structure (P5) it will take 14 years to get the Return of Investment while with a good but imperfect forecast the time to get the ROI increases to 15 or 21 years.

E. Extended Dataset

An extended dataset was considered consisting in the 15-minute energy usage of users with ids: 26, 59, 86, 93, 101, 114, 171, 186, 821, 5403, 5450, 3204 between the 1st February 2019 and the 31th May 2019. The first three months were used to fit the AvgPast forecast, but only the month of May was used to run the experiments. Given SARIMA's bad performance in Figures 9 & 10, it was excluded from the analysis. In the following experiments, weekends were not removed.

²The cost of a Powerwall 2 is 7000 euros in France.

Figure 11 depicts the value of the linear map defined in Equation (6), for each sampling frequency of MPC and each of the forecasting techniques used. Recall that a y-value of 0 is equivalent to having perfect information, a value of 1 is equivalent to not having a battery and for values greater than 1, the performance is worst than doing nothing. Qualitatively, results preserve the structure observed before: for all the forecasting techniques, MPC(1) (the highest sampling) is enough to guarantee a profit for the battery. Moreover, with less sampling, AvgPast performs better than Persistence.

V. CONCLUSIONS

Consumers facing a Time of Use can perform arbitrage by shifting their consumption in the expensive period to the cheap period. For a consumer without generation, a higher selling price will mitigate unnecessary energy bought because of a forecast error. On the contrary, for a consumer with generation, a high selling price will translate in large lost opportunities whenever energy is not sold due to a forecasting error. Even then, as prices do not spike, the margins are low. Such low margins imply that for some days, even with very accurate forecast it will be very hard not to lose money. That is the case of household 5403 during the 24th of July. It is found that the sampling time greatly impacts the profit and only very good forecasts can avoid losses with a sampling frequency greater than one. In spite of such negative results, on average, running MPC(1) with the AvgPast forecast seems to be sufficient to guarantee profit. These conclusions were validated with our extended dataset. We conclude that forecasting does have a big impact on the profits obtained facing a Time of Use. However, since such profits are low, the gap between the AvgPast and very accurate Gaussian approximation of forecast do not yield a very big difference.

REFERENCES

- [1] D. Kiedanski, A. Orda, and D. Kofman, "The effect of ramp constraints on coalitional storage games," in *Proceedings of the Tenth ACM International Conference on Future Energy Systems, e-Energy '19*, (New York, NY, USA), pp. 226–238, ACM, 2019.
- [2] E. L. Ratnam, S. R. Weller, and C. M. Kellett, "An optimization-based approach to scheduling residential battery storage with solar pv: Assessing customer benefit," *Renewable Energy*, vol. 75, pp. 123 – 134, 2015.

- [3] S. J. Chiang, K. T. Chang, and C. Y. Yen, "Residential photovoltaic energy storage system," *IEEE Transactions on Industrial Electronics*, vol. 45, pp. 385–394, June 1998.
- [4] J. Horta, D. Kofman, D. Menga, and A. Silva, "Novel market approach for locally balancing renewable energy production and flexible demand," in *2017 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 533–539, Oct 2017.
- [5] D. Kiedanski, D. Kofman, J. Horta, and D. Menga, "Strategy-proof local energy market with sequential stochastic decision process for battery control," in *IEEE Innovative Smart Grid Technologies 2019 NA*, (Washington DC, United States), Feb. 2019.
- [6] D. Thomas, O. Deblecker, and C. S. Ioakimidis, "Optimal operation of an energy management system for a grid-connected smart building considering photovoltaics uncertainty and stochastic electric vehicles driving schedule," *Applied Energy*, vol. 210, pp. 1188 – 1206, 2018.
- [7] Q. Wei, G. Shi, R. Song, and Y. Liu, "Adaptive dynamic programming-based optimal control scheme for energy storage systems with solar renewable energy," *IEEE Transactions on Industrial Electronics*, vol. 64, pp. 5468–5478, July 2017.
- [8] X. Wu, X. Hu, S. Moura, X. Yin, and V. Pickert, "Stochastic control of smart home energy management with plug-in electric vehicle battery energy storage and photovoltaic array," *Journal of Power Sources*, vol. 333, pp. 203 – 212, 2016.
- [9] R. Cruise, L. Flatley, R. Gibbens, and S. Zachary, "Control of energy storage with market impact: Lagrangian approach and horizons," *Operations Research*, vol. 67, pp. 1–9, 1 2019.
- [10] C. Byrne and G. Verbic, "Feasibility of residential battery storage for energy arbitrage," in *2013 Australasian Universities Power Engineering Conference (AUPEC)*, pp. 1–7, Sep. 2013.
- [11] T. Senjyu, H. Takara, K. Uezato, and T. Funabashi, "One-hour-ahead load forecasting using neural network," *IEEE Transactions on Power Systems*, vol. 17, pp. 113–118, Feb 2002.
- [12] D. C. Park, M. El-Sharkawi, R. Marks, L. Atlas, and M. Damborg, "Electric load forecasting using an artificial neural network," *IEEE transactions on Power Systems*, vol. 6, no. 2, pp. 442–449, 1991.
- [13] H. S. Hippert, C. E. Pedreira, and R. C. Souza, "Neural networks for short-term load forecasting: A review and evaluation," *IEEE Transactions on power systems*, vol. 16, no. 1, pp. 44–55, 2001.
- [14] H. K. Alfares and M. Nazeeruddin, "Electric load forecasting: literature survey and classification of methods," *International journal of systems science*, vol. 33, no. 1, pp. 23–34, 2002.
- [15] M. Y. Cho, J. C. Hwang, and C. S. Chen, "Customer short term load forecasting by using arima transfer function model," in *Proceedings 1995 International Conference on Energy Management and Power Delivery EMPD '95*, vol. 1, pp. 317–322 vol.1, Nov 1995.
- [16] P. Lulis, K. R. Khalilpour, L. Andrew, and A. Liebman, "Short-term residential load forecasting: Impact of calendar effects and forecast granularity," *Applied Energy*, vol. 205, pp. 654 – 669, 2017.
- [24] G. Bathurst and G. Strbac, "Value of combining energy storage and wind in short-term energy and balancing markets," *Electric Power Systems Research*, vol. 67, no. 1, pp. 1 – 8, 2003.
- [17] S. Humeau, T. K. Wijaya, M. Vasirani, and K. Aberer, "Electricity load forecasting for residential customers: Exploiting aggregation and correlation between households," in *2013 Sustainable Internet and ICT for Sustainability (SustainIT)*, pp. 1–6, Oct 2013.
- [18] W. Kong, Z. Y. Dong, D. J. Hill, F. Luo, and Y. Xu, "Short-term residential load forecasting based on resident behaviour learning," *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 1087–1088, 2017.
- [19] S. Humeau, T. K. Wijaya, M. Vasirani, and K. Aberer, "Electricity load forecasting for residential customers: Exploiting aggregation and correlation between households," in *2013 Sustainable Internet and ICT for Sustainability (SustainIT)*, pp. 1–6, Oct 2013.
- [20] M. Ghofrani, M. Hassanzadeh, M. Etezadi-Amoli, and M. S. Fadali, "Smart meter based short-term load forecasting for residential customers," in *2011 North American Power Symposium*, pp. 1–5, Aug 2011.
- [21] A. Dunbar, F. Tagliaferri, I. M. Viola, and G. P. Harrison, "The impact of electricity price forecast accuracy on the optimality of storage revenue," in *3rd Renewable Power Generation Conference (RPG 2014)*, pp. 1–6, Sep. 2014.
- [22] D. Krishnamurthy, C. Uckun, Z. Zhou, P. R. Thimmapuram, and A. Botterud, "Energy storage arbitrage under day-ahead and real-time price uncertainty," *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 84–93, 2018.
- [23] K. Bradbury, L. Pratson, and D. Patio-Echeverri, "Economic viability of energy storage systems based on price arbitrage potential in real-time u.s. electricity markets," *Applied Energy*, vol. 114, pp. 512 – 519, 2014.
- [25] K.-H. Ahlert and C. Block, "Assessing the impact of price forecast errors on the economics of distributed storage systems," in *2010 43rd Hawaii International Conference on System Sciences*, pp. 1–10, IEEE, 2010.
- [26] M. U. Hashmi, A. Mukhopadhyay, A. Bušić, and J. Elias, "Optimal control of storage under time varying electricity prices," in *2017 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 134–140, IEEE, 2017.
- [27] G. Mulder, D. Six, B. Claessens, T. Broes, N. Omar, and J. V. Mierlo, "The dimensioning of pv-battery systems depending on the incentive and selling price conditions," *Applied Energy*, vol. 111, pp. 1126 – 1135, 2013.
- [28] P. Mokrian, M. Stephen, *et al.*, "A stochastic programming framework for the valuation of electricity storage," in *26th USAEE/IAEE North American Conference*, pp. 24–27, 2006.
- [29] M. Hashmi, A. Mukhopadhyay, A. Busic, and J. Elias, "Storage optimal control under net metering policies," *to be submitted IEEE Transactions on Smart Grid*.
- [30] "Powerwall 2 datasheet." Accessed: 2019-01-29.
- [31] Y. Chen, M. U. Hashmi, D. Deka, and M. Chertkov, "Stochastic battery operations using deep neural networks," in *IEEE ISGT, NA Washington DC*, 2019.
- [32] P. Street Inc., "Dataport." "<https://www.pecanstreet.org/dataport/>", 2019.
- [33] S. Kolassa and W. Schtz, "Advantages of the mad/mean ratio over the mape," *Foresight: The International Journal of Applied Forecasting*, vol. 6, pp. 40–43, 01 2007.