

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/337729645>

A Data-driven Storage Control Framework for Dynamic Pricing

Preprint · December 2019

CITATIONS

0

READS

87

5 authors, including:



Jiaman Wu

University of California, Berkeley

15 PUBLICATIONS 76 CITATIONS

SEE PROFILE



Chenye Wu

The Chinese University of Hong Kong - Shenzhen

166 PUBLICATIONS 2,638 CITATIONS

SEE PROFILE

A Data-driven Storage Control Framework for Dynamic Pricing

Jiaman Wu, Zhiqi Wang, Chenye Wu, Kui Wang, and Yang Yu

Abstract—Dynamic pricing is both an opportunity and a challenge to the demand side. It is an opportunity as it better reflects the real time market conditions and hence enables an active demand side. However, demand’s active participation does not necessarily lead to benefits. The challenge conventionally comes from the limited flexible resources and limited intelligent devices in demand side. The decreasing cost of storage system and the widely deployed smart meters inspire us to design a data-driven storage control framework for dynamic prices. We first establish a stylized model by assuming the knowledge and structure of dynamic price distributions, and design the optimal storage control policy. Based on Gaussian Mixture Model, we propose a practical data-driven control framework, which helps relax the assumptions in the stylized model. Numerical studies illustrate the remarkable performance of the proposed data-driven framework.

Index Terms—Dynamic Price, Stochastic Control, Storage Control, Gaussian Mixture Model

I. INTRODUCTION

IN terms of economic effectiveness, dynamic pricing outperforms its rivals since it accurately represents the value of energy in real time. More importantly, it provides proper incentives to the demand side to achieve a cost effective power system [1].

Nonetheless, dynamic pricing does not necessarily benefit every end user. After implementing dynamic price in Alberta, most end users, who opted in variable electricity rates, witnessed dramatic increase in their electricity bills during winter over the past few years [2].

On the one hand, this highlights that dynamic price does reflect the true value of energy in real time. On the other hand, such price signal, due to limited intelligence in the demand side, only poses risks to most end users. To better utilize the advantage of dynamic pricing, low cost storage systems and smart meters with intelligent control capabilities are of vital importance. Together with these two necessary components, dynamic pricing can enable an active demand side, and can thus achieve the system-wide economic effectiveness.

Hence, in this paper, to design the storage control framework for demand side facing dynamic prices, we jointly utilize the storage system as well as the intelligence in smart meters. Based on a stylized model, we first identify the optimal threshold-based control policy. With the help of the Gaussian Mixture Model (GMM) [3], we design a data-driven storage control framework, which is able to further relax the

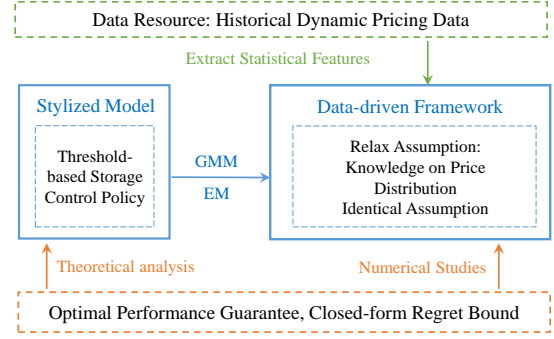


Fig. 1. Our Data-driven Storage Control Framework

assumptions in the stylized model, e.g., knowledge on the type of distribution, identical assumptions, etc. We plot the flowchart in designing the control framework in Fig. 1.

A. Literature Review

Storage system can provide the urgently needed flexibility in the power system. Hence, it has been well investigated to utilize storage system to provide various services in the electricity sector, including smoothing the dispatched wind power generation [4], maintaining system reliability [5], smart load scheduling for a single end user [6] or a networked community of cooperative users [7], grid level operation for minimum cost [8], etc.

However, the investigation on storage control policy for dynamic pricing only emerges recently. Jin *et al.* propose a heuristic algorithm using Mixed Integer Linear Programming to optimize the electric vehicle charging schedules in [9]. Oudalov *et al.* focus on conducting peak load shaving and introduce a sizing methods as well as an optimal operational scheme in [10]. Chau *et al.* assume the knowledge of future demand and the bounds of prices, and illustrate a threshold-based cost minimizing online algorithm with worst-case performance guarantee in [11]. To deal with limited information and uncertainty, Qin *et al.* introduce an online modified greedy algorithm for storage control in [12]. Vojvodic *et al.* design a simple forward thresholds algorithm to manage the operation in real-time energy market, where they decompose coupling stages using integer programming and heuristic search in [13]. Harsha *et al.* propose an optimal storage management policy based on the stochastic dynamic programming and show the policy has a dual threshold structure under mild assumptions in [14]. Different from the literature, we seek to design a *data-driven* storage control framework for dynamic pricing.

The authors are with the Institute for Interdisciplinary Information Sciences (IIIS), Tsinghua University, Beijing, China, 100084. C. Wu is the correspondence author. Email: chenye.wu@tsinghua.edu.cn.




Fig. 2. Demonstration of system model.

Another closely related research line applies data-driven methods to the electricity sector (see [15] for a systematic survey). Just to name a few, Samuelson *et al.* propose a robust storage control framework to manage the wind power fluctuations in [16]. Thanos *et al.* design a data-drive simulation framework for automated control in microgrid in [17]. In contrast to most related works, when designing the data driven framework, we provide valuable theoretical insights into the stylized model, which help us understand the value of information in storage control.

B. Our Contributions

This paper significantly extends our previous work [18] in both theoretical and practical perspectives. We provide the theoretical basis for the optimal *online* threshold based storage control policy and derive its regret bound, comparing with the *offline* optimal policy. From the practical perspectives, we take a data-driven approach and employ GMM to accurately characterize the price distribution, and use extensive numerical studies to assess the performance of our proposed framework. In summary, we highlight our principal contributions as follows:

- *Optimal Online Threshold Policy*: Based on the one-shot load decomposition technique, we propose a simple yet effective online threshold policy to minimize the consumers' expected electricity bills. We prove its optimality and derive its regret bound.
- *Data-driven Framework with GMM*: We adopt GMM to relax the assumption of knowing exact price distribution in deriving the threshold policy. With historical price data, we illustrate how to customize the Expectation-Maximization (EM) algorithm [19] for the parameter estimation in GMM, yielding our data-driven storage control framework.
- *Bridge theory and practice*: Our data-driven framework and its heuristic variants enable us to relax many assumptions on the price distributions, including the knowledge of its exact form and the *i.i.d.* assumption. Such relaxations significantly improve the practical feasibility of our framework.

The rest of our paper is organized as follows. Section II introduces the system model, and revisits the one-shot load decomposition technique. Then, in Section III, we design two optimal online threshold storage control frameworks: ETA and DETA. ETA is designed for the stylized model, while DETA and its variants relax many assumptions in the stylized




Fig. 3. Example for one-shot load decomposition.

model. In Section IV, we use numerical studies to evaluate the performance of the two frameworks with real data. Concluding remarks are given in Section V.

II. ONE-SHOT LOAD DECOMPOSITION

Consider the interaction between consumers and the grid as shown in Fig. 2. The grid operator sets dynamic price $p(t)$ at each time t . Facing such a pricing scheme, the consumer wants to satisfy its demand $d(t)$ in different ways: directly purchases energy $g(t)$ from the grid, saves energy $b(t)$ in the storage system, or uses the energy in the storage system ($c(t)$ out of $s(t)$ in the storage) to meet its demand.

Specifically, to highlight the impact of uncertainties in dynamic price on the control policy design, we assume the consumer's own demand prediction is accurate in the near future. This is reasonable as on the grid level, the impact of uncertainties in renewables is reflected through the volatility in dynamic prices. Even with this assumption, the decision making for each consumer is still quite challenging due to the uncertainty in future prices and the physical constraints (capacity constraints) coupling all storage control decisions.

Inspired by [11], in this section, we first revisit the one-shot load decomposition technique, which allows us to decouple the storage control actions across time. Then, we formally formulate the one-shot load serving problem.

A. Revisit the One-shot Load Decomposition

The one-shot load decomposition technique was first proposed by Chau *et al.* in [11] to decouple the storage capacity constrained optimization problem into a sequence of one-shot load serving problems.

For example, suppose we need serve an accumulated demand as shown by the red dashed curve in Fig. 3: serving a load of $D(t_1)$ at time t_1 , and serving a load of $D(t_2) - D(t_1)$ at time t_2 . If there were no storage devices, then we had to purchase the load at the time of serving, tolerating all the price volatility. However, with a storage of capacity B , we can reduce such risk.

Back to our example in Fig. 3, assume

$$B > D(t_1) := d_1(t_0, t_1), \quad (1)$$

then with storage, we can serve $D(t_1)$ at any time between $[t_0, t_1]$. Therefore, we define $d_1(t_0, t_1) = D(t_1)$, which

highlights its flexibility in the time span $[t_0, t_1]$. Next, we assume

$$B < D(t_2) - D(t_1), \quad (2)$$

then the demand to be met at t_2 need be decomposed into three different kinds of demand: $d_2(t_0, t_2)$, $d_3(t_1, t_2)$, and $d_4(t_2, t_2)$. Since $B > D(t_1)$, the storage has certain spare capacity to store energy and serve the load at t_2 even between $[0, t_1]$. This observation leads to the first kind of demand $d_2(t_0, t_2)$, which is reserved for load at t_2 between $[0, t_2]$ (the union of $[0, t_1]$ and $[t_1, t_2]$). The second kind of demand is due to the released capacity after serving load at t_1 , which is only flexible after t_1 . We denote it by $d_3(t_1, t_2)$. Note that since $B < D(t_2) - D(t_1)$, certain amount of load at time t_2 , $d_4(t_2, t_2)$, has to be met in real time. This leads to the last type of decomposed demand in the one-shot load decomposition technique.

Remark: The last type of decomposed demand naturally decouples the original optimization problem over a *long time* span into a set of *smaller scale* optimization problems. This implies that when we conduct the one-shot load decomposition, we only need look ahead for a couple of hours, the prediction of which can be rather accurate. This also illustrates our assumption of perfect knowledge on the near future demand is realistic.

Next, we formally introduce the construction process for the one-shot load decomposition technique as follows:

- 1) Define the accumulative demand curve $D(t)$.
- 2) Define the upward shift accumulative demand curve $D^{\text{shift}}(t)$, which is obtained by shift $D(t)$ by B .
- 3) Obtain one-shot demand $d_i(t_s^i, t_e^i)$ through sandwiched rectangle $(t_s^i - t_e^i)d_i(t_s^i, t_e^i)$ between $D(t)$ and $D^{\text{shift}}(t)$.

B. One-shot Load Serving Problem

With the one-shot load decomposition, we can focus on solving the one-shot load serving problem. More precisely, for each decomposed load, $d_i(t_s^i, t_e^i)$, which need be served between t_s^i and t_e^i , we seek to find the time t with the minimum price $p(t)$ to serve it. In fact, it suffices for us to understand the stylized one-shot load serving problem, where consumer need satisfy its *one-unit* demand between 1 and T .

Mathematically, if we were able to foresee the future prices, we could directly select:

$$t^* = \arg \min_{t \in [1, T]} p(t). \quad (3)$$

However, without the perfect knowledge of the future price, we need design an online algorithm to solve the one-shot load serving problem, based on which we can construct the control policy for the general load serving problem.

III. OPTIMAL CONTROL FRAMEWORK

To simplify our subsequent analysis with more insights, we make the following assumption:

Assumption: Dynamic price $p(t)$'s are *i.i.d* random variables.

With this assumption, we design the control framework for the following two cases:

- *Exact Distribution:* Knowing the exact distribution of the random price $p(t)$, we seek to design the optimal online control policy.
- *Data-driven Framework:* Based on the optimal control policy, we design the data-driven control framework to relax our assumptions on the price distribution.

A. Threshold Control Policy with Exact Distribution

For the one-shot load serving problem between $[1, T]$, at each time t , we only have two choices: to purchase the unit demand or not. The two choices correspond to different expected costs: $p(t)$ for purchasing and $\mathbb{E}[w_{t+1}]$ for not purchasing, where $\mathbb{E}[w_{t+1}]$ denotes the expected cost for the one-shot unit load serving between $[t+1, T]$. Note that, to calculate $\mathbb{E}[w_{t+1}]$, we assume the full knowledge of $p(t)$'s distribution.

To characterize this binary choice, we seek to design the time varying threshold $\theta(t)$ to balance the expected costs between two actions. Hence, the optimal threshold would require

$$\theta(t) = \mathbb{E}[w_{t+1}]. \quad (4)$$

If $p(t) \leq \theta(t)$, we choose to purchase the unit load at time t . Otherwise, we defer this action to later time slots.

These thresholds can be obtained in a recursive manner:

$$\theta_{t-1} = \int_0^{\theta_t} p_t f(p) dp + \int_{\theta_t}^{\infty} \theta_t f(p) dp, \quad (5)$$

where $f(p)$ denotes the probability density function of dynamic price $p(t)$, $\forall t$. We have $\theta_{T-1} = \mathbb{E}[w_T] = \mathbb{E}[p(T)]$ as the boundary condition. We term this simple threshold control policy as *expected threshold algorithm* (ETA).

B. Optimality of ETA

The optimality of ETA comes from the fact that the decision making at each time slot is a binary choice. With this fact, we can prove the following theorem:

Theorem 1: ETA is the optimal storage control policy for the *one-shot load serving problem*.

Proof: This theorem can be proved by backward induction. Note that at $\tau = T - 1$, $\theta_{T-1} = \mathbb{E}[w_T] = \mathbb{E}[p(T)]$ is the optimal choice. This constructs the induction basis. For the induction part, it suffices to identify that θ_t is the solution to the first order optimality condition of (5). ■

However, it is not obvious why the optimal control policy for the one-shot load serving problem can be used to construct the optimal policy for the original problem. This is because the one-shot load decomposition may change the feasible region of the original problem. Fortunately, we can prove that the one-shot load decomposition maintains the feasible region. Formally, we have the following Lemma.

Lemma 1: Denote the solution space of general load serving problem by \mathcal{S} . Assume this problem can be decomposed into a sequence of k one-shot load serving problems and denote the corresponding solution spaces by $\mathcal{S}_1, \dots, \mathcal{S}_K$, respectively. It holds:

$$\mathcal{S} = \bigcup_{i=1}^k \mathcal{S}_i. \quad (6)$$

Proof: Since the decomposition satisfies all the constraints in the original problem, it immediately follows that

$$\bigcup_{i=1}^k \mathcal{S}_i \subseteq \mathcal{S}. \quad (7)$$

The difficulty is to show any feasible solution in \mathcal{S} is also in $\bigcup_{i=1}^k \mathcal{S}_i$. We prove by contradiction. Suppose $\bigcup_{i=1}^k \mathcal{S}_i \subsetneq \mathcal{S}$, then there exists a sequence of storage control actions $s = \{\{b(t)\}_{t=1}^T, \{c(t)\}_{t=1}^T\}$, which satisfies the following condition:

$$s \in \mathcal{S} - \bigcup_{i=1}^k \mathcal{S}_i. \quad (8)$$

The fact that $s \in \mathcal{S}$ implies $\sum_{i=1}^t b(i) - c(i) \leq D(t) - D(t-1), \forall 0 \leq t \leq T$, while $s \notin \bigcup_{i=1}^k \mathcal{S}_i$ implies there exists a time t_0 , such that

$$\sum_{i=1}^{t_0} b(i) - c(i) > D^{\text{shift}}(t_0) - D^{\text{shift}}(t_0 - 1). \quad (9)$$

However, due to the definition of $D^{\text{shift}}(t)$, we know that

$$D^{\text{shift}}(t_0) - D^{\text{shift}}(t_0 - 1) = D(t_0) - D(t_0 - 1), \quad (10)$$

which establishes the contradiction. Hence,

$$\bigcup_{i=1}^k \mathcal{S}_i \supseteq \mathcal{S}. \quad (11)$$

Together with Eq. (7), we complete the proof. \blacksquare

With Lemma 1, the optimality of ETA for the one-shot load serving problem implies its optimality for the general load serving problem. Formally, we state this conclusion as a proposition.

Proposition 1: ETA is the optimal storage control policy for the general load serving problem.

C. Regret Bound of ETA

One important metric to evaluate the performance of any online policy is its regret, which measures the performance difference between the online policy and the offline optimal. In this section, since the regret analysis for the general load serving problem depends on the structure of specific problem, we only provide the theoretical regret analysis for the one-shot load serving problem, and conduct numerical studies to highlight ETA's performance for general load serving.

Mathematically, we define the regret of the one-shot load serving problem with shot length T by:

$$R(T) = \mathbb{E}[w]^{\text{ETA}} - \mathbb{E}[w]^{\text{Offline}}, \quad (12)$$

where $\mathbb{E}[w]^{\text{ETA}}$ and $\mathbb{E}[w]^{\text{Offline}}$ represent the expected costs obtained by ETA and offline optimal, respectively. To better characterize the regret $R(T)$, we define the following two parameters:

$$\alpha = \frac{\mathbb{E}[p_2^{\min}]}{2\mathbb{E}[p]}, \quad (13)$$

$$\beta_t = \inf\{f(p)|p \in [0, \theta_t]\}, \quad (14)$$

where $p_2^{\min} = \min\{p(1), p(2)\}$. Thus, we can characterize the upper bound of $R(T)$ as follows:

Theorem 2: Assume the dynamic price $p(t)$ is non-negative, then,

$$R(T) \leq \frac{2}{\sum_{i=2}^T \beta_i} - T\alpha^{T-1}\mathbb{E}[p]. \quad (15)$$

We provide the detailed proof in Appendix A. To highlight $R(T)$'s dependence on the moments of dynamic price, we characterize the bound of $R(T)$ in terms of mean and variance for uniform distribution as an illustration.

Assume the dynamic price $p(t)$ follows $U(a, b)$, where $a > 0$. Denote the mean and variance of the uniform distribution by μ and σ^2 . Then, standard mathematical manipulation yields

$$R(T) \leq \frac{4\sqrt{3}\sigma}{T-1} - T\mu \left(1 - \frac{\mu^2 + \sigma^2}{2\sqrt{3}\mu\sigma}\right)^{T-1}. \quad (16)$$

This bound coincides with many of our intuitions. Firstly, $R(T)$ decreases with shot length T , of order $O(T^{-1})$. This shows that $R(T)$ approaches 0 as T becomes large. Secondly, smaller variance (easier to predict) implies better performance guarantee for ETA.

D. Data-driven Control Framework

We use GMM to relax the assumption of knowing the exact distribution of dynamic price. Due to Central Limit Theorem and the mixture model in GMM, as long as the proper parameters are selected, GMM can approximate any continuous distribution accurately [3].

More precisely, GMM assumes that the data is generated from several Gaussian distributions. Hence, its log-likelihood function is of the following form:

$$L(t, K) = \sum_{i=1}^t \ln \left(\sum_{k=1}^K \pi_k \mathcal{N}(p_i | \mu_k, \sigma_k) \right), \quad (17)$$

where t is the number of samples, K is the number of Gaussian components, π_k is the weight of each component, and $\mathcal{N}(p_i | \mu_k, \sigma_k)$ is the probability density function for Gaussian distribution, given mean of μ_k , and standard deviation of σ_k .

Due to the complex structure in GMM, one primary difficulty in applying GMM for approximating any distribution is the model selection, *i.e.*, how to find the least number of parameters to best represent the data distribution. The reason for selecting the least number of parameters is not only because of Occam's Razor [20], but also because too many parameters will increase the risk of over-fitting [21].

In this paper, we choose to use the Bayesian Information Criteria (BIC) [22] to determine the optimal number of components in GMM. For a model with k parameters and t samples, its BIC value is defined as follows:

$$BIC(t, k) = k \ln t - 2 \ln L(t, k). \quad (18)$$

For the model with determined number of components, we employ the Expectation-Maximization (EM) algorithm [19] for parameter estimation. Then we use BIC to select the optimal number of components, which leads to our Data-driven Expected Threshold Algorithm (DETA). To make the algorithm description neat, we only show how to use EM algorithm to estimate the parameters in GMM in Algorithm

1¹. With the estimated GMM as price distribution, we can directly apply ETA for storage control.

Algorithm 1 DETA (Distribution Estimator)

Input: K^M : maximum number of components in GMM;
 t : current time slot;
 $p(i), i = 1, \dots, t$: historical price data;
Output: G^* : estimated GMM;

- 1: **for** $K^m = 1$ to K^M **do**
- 2: GMM($t, K^m(t)$) Training for $\pi_k, \mu_k, \sigma_k, 1 \leq k \leq K^m$:
 Set the initial value of π_k, μ_k and $\sigma_k, 1 \leq k \leq K^m$.
- 3: **repeat**
- 4: Compute posterior probability:

$$\gamma(\zeta_{ik}) = \frac{\pi_k \mathcal{N}(p_i | \mu_k, \sigma_k)}{\sum_{j=1}^{K^m} \pi_j \mathcal{N}(p_i | \mu_j, \sigma_j)}$$
- 5: Define $N_k = \sum_{i=1}^t \gamma(\zeta_{ik})$
- 6: Update parameters according to $\gamma(\zeta_{ik})$ and N_k :

$$\mu'_k = \frac{1}{N_k} \sum_{i=1}^t \gamma(\zeta_{ik}) p_i,$$

$$\sigma'_k = \frac{1}{N_k} \sum_{i=1}^t \gamma(\zeta_{ik}) (p_i - \mu'_k)(p_i - \mu'_k),$$

$$\pi'_k = \frac{N_k}{t}.$$
- 7: **until** μ_k, σ_k, π_k and $L(t, K^m)$ converge
- 8: **end for**
- 9: Model Selection : choose the model G^* as GMM($t, k^*(t)$)
 with

$$k^*(t) = \arg \min_{k \in [1, K^M]} BIC(t, k);$$

E. Heuristic Variants for Performance Improvement

Our *i.i.d.* assumption on the price distribution is for a neat analysis. In fact, in ETA, we needn't restrict ourselves to *i.i.d.* price distribution. Assume a temporally dependent price distribution $f_t(p)$ for $p(t)$ at time t . Then, we can straightforwardly generalize the threshold recursive structure:

$$\theta_{t-1} = \mathbb{E}[w_t] = \int_0^{\theta_t} p f_t(p) dp + \int_{\theta_t}^{\infty} \theta_t f_t(p) dp. \quad (19)$$

The resulting threshold policy can again guarantee the optimality for one-shot load serving problem, and hence, with Lemma 1, the optimality for general load serving. The difficulty lies in the description for price distributions for all time t . This motivates us to design the following two heuristic algorithms, both with easy descriptions for price distributions.

One immediate intuition is that the price traces show strong periodicity. The naive period can be 24 hours². To capture such periodicity, we can use 24 GMMs: one for each hour in the day. We term this approach by DETA^P, where the superscript P emphasizes the periodicity.

The major problem with *i.i.d.* assumption for our storage control is the identical assumption. Another heuristic is to divide each day into a couple of periods, and assume the price

¹In this algorithm, we simply employ the classical EM algorithm. It is possible to design a customized adaptive EM algorithm (see [23] for an example). However, a detailed discussion is beyond the scope of this paper.

²To achieve better performance, one may conduct more comprehensive hypothesis test to find the "true" period in real data. However, a detailed discussion is beyond the scope of this paper. More importantly, the "true" period may not necessarily lead to a better performance.




Fig. 4. Real price data and its distribution fittings.

distributions within each period are identical. Hence, we adopt a simple adaptive approach: by observing the training period, we first denote the mean of the observed prices by μ . And for each hour, as long as the mean of training price in this hour exceeds μ , we label this hour as peak period. We label all other hours as off-peak periods. This allows us to train only 2 GMMs: one for peak periods, the other for off-peak periods. We term this approach by DETA^I, where the superscript I emphasizes that we relax the identical assumption in the stylized model. We describe the procedure of DETA^I in detail in Algorithm 2.

Algorithm 2 DETA^I (Peak and Off-Peak Period Detecting)

Input: T : number of days for training;
 $p^i(t) : i = 1, \dots, 24, t = 1, \dots, T$: historical price data in day t at hour i .
Output: \mathcal{P} : set of peak periods;
 \mathcal{O} : set of off-peak periods.

- 1: Initialization: $\mathcal{P} = \phi, \mathcal{O} = \phi$.

$$\mu = [\sum_{i=1}^{24} \sum_{t=1}^T p^i(t)] / (24T)$$
- 2: **for** $i = 1$ to 24 **do**
- 3:
$$\hat{\mu} = \frac{\sum_{t=1}^T p^i(t)}{T}$$
- 4: **if** $\hat{\mu} > \mu$ **then**
- 5:
$$\mathcal{P} \leftarrow \mathcal{P} \cup \{i\}.$$
- 6: **end if**
- 7: **end for**
- 8:
$$\mathcal{O} \leftarrow \{1, \dots, 24\} - \mathcal{P}$$
- 9: **return** \mathcal{P} and \mathcal{O}

Remark: To achieve a better performance, one may want to select a better threshold, rather than the mean μ , to differentiate peak periods and off-peak periods. For example, in our subsequent analysis, we find that selecting a threshold of 40% percentile achieves a better performance. However, we choose not to dive into this more engineering direction. Instead, we emphasize that a simple implementation already achieves improved performance in many cases.

IV. NUMERICAL STUDIES

We use the hourly real-time price data to characterize the stochastic nature in dynamic price. The data is collected from AEP [24] during August, 2019. Figure 4 plots the histogram of the prices and its GMM approximation. Figure 5 suggests that $k = 3$ achieves the minimal BIC value (marked by *). Hence, we select a GMM with 3 components for the




Fig. 5. BIC value corresponding to different number of components.




Fig. 6. ETA performance for one-shot load serving.

distribution fitting. To align with the *i.i.d.* assumption of the price distribution in our stylized model, we first use synthetic data (generated by the fitted GMM) to evaluate ETA's performance. Then, we use real price data to evaluate the performance of DETA and its heuristic variants. As for load dataset, we randomly sample a period from the AEP users' load data, during 2019 [25].

A. Evaluate ETA for One-shot Load Serving

We first evaluate the performance of ETA for the one-shot load serving problem, and set the storage capacity to be 10% of the peak hourly energy consumption³. To better visualize the comparison between our ETA performance and offline optimal, we define regret ratio γ as follows:

$$\gamma = \frac{\text{cost}(ETA) - OPT}{OPT}, \quad (20)$$

where $\text{cost}(ETA)$ denotes the cost of ETA, and OPT denotes the offline minimal cost. Figure 6 plots the performance of ETA using the synthetic data, which displays diminishing regret ratio, and implies ETA converges to the offline optimal rather fast in the one-shot load serving problem.

B. Evaluate ETA for General Load Serving

Next, we evaluate the performance of ETA for general load serving purposes. We define the competitive ratio β as follows

$$\beta = \frac{\text{cost}(ETA)}{OPT}, \quad (21)$$

where $\text{cost}(ETA)$ denotes the *total* cost of ETA in serving load during certain period of time, while OPT denotes the corresponding offline minimal *total* cost. Figure 7 plots the competitive ratio for synthetic price data. We observe that the competitive ratio becomes stable as time goes by. Figure 7 suggests that the mean of β is bounded by 1.04.

³A detailed discussion on optimal storage sizing is provided in Appendix B, which heavily depends on the amortized cost of storage systems.






Fig. 7. ETA performance for general load serving.



(a) $B = 20\%$ of peak hourly energy consumption.



(b) $B =$ peak hourly energy consumption.

Fig. 8. Competitive Ratio of DETA.

C. Empirical Competitive Ratio

Next, we use real price data to examine the DETA's competitive ratio in the twelve months. In each month, we use the price data in the first three weeks to train the GMM parameters, and evaluate DETA's performance for the last week. Note that, although this method is not adaptive enough, it is more practical due to its minimal intelligent requirement on the local control devices. As shown in Fig. 8, the DETA performs remarkably good on all traces. When there is limited flexibility (B is only 20% of the peak hourly energy consumption), the offline optimal cannot save too much. In this case, the average stable competitive ratio is around 1.03. As there are more flexibility (B is the peak hourly energy consumption), the offline optimal achieves much more savings. In this case, the average stable competitive ratio is around 1.1.

Although remarkable, in power system control, every 1% saving may be in the order of millions of dollars. Hence, We identify three months with the worst performance in terms of competitive ratio: May, July and October. Figure 9 further visualizes the price traces for these three months, which can help explain the diverse performance of DETA's heuristic variants.




Fig. 9. Price Traces in May, July and October.

D. Performance Improvement by Heuristics

We evaluate whether the two heuristic algorithms can actually improve the DETA performance in the three months. We again use the price data in the first 3 weeks for training, and then test the performance of the three approaches: DETA, $DETA^P$, and $DETA^I$ (with 40% percentile) for the last week. Figures 10 and 11 illustrate this comparison, for different storage capacities.

It is interesting to note that $DETA^P$ improves DETA's performance in May and July, but performs poorly in October. This may be because, for example, in October, the price data does not display strong periodicity, comparing with that in the other two months. Also, $DETA^P$ requires training 24 GMMs, which significantly limits the available training data for each GMM.

In contrast, $DETA^I$ leads to remarkable improvements in all cases. This implies the potential benefits of improving the current implementation of $DETA^I$, by smartly selecting multiple thresholds.

V. CONCLUSION

In this paper, for a stylized model, we design the optimal storage control policy with theoretical guarantee. We generalize the control policy by employing GMM and propose the data-driven framework, which performs remarkably, suggested by numerical studies with real data. We design two promising heuristic algorithms to further enhance the performance of the data driven framework.

This work can be extended in many ways. It is interesting to understand the conditions under which the two heuristic algorithms can achieve the maximal improvement. We also intend to relax the assumption of perfect prediction for the future demand, and study the impact of coupled uncertainties (from both dynamic pricing and the demand) on the optimal control policy design.

REFERENCES

[1] S. Borenstein, M. Jaske, and A. Rosenfeld, "Dynamic pricing, advanced metering, and demand response in electricity markets," tech. rep., UC Berkeley: Center for the Study of Energy Markets, 2002.

[2] M. Alleckna, "Why alberta electricity bills are getting higher and what you can do about it." <https://shorturl.at/jrS34>. Accessed April 5, 2019.

[3] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.

[4] S. Teleke, M. E. Baran, A. Q. Huang, S. Bhattacharya, and L. Anderson, "Control strategies for battery energy storage for wind farm dispatching," *IEEE Trans. on Energy Conversion*, vol. 24, pp. 725–732, Sep. 2009.

[5] I. Koutsopoulos, V. Hatzi, and L. Tassiulas, "Optimal energy storage control policies for the smart power grid," in *Proc. of IEEE Smart Grid Comm*, pp. 475–480, Oct 2011.

[6] P. M. van de Ven, N. Hegde, L. Massouli, and T. Salonidis, "Optimal control of end-user energy storage," *IEEE Trans. on Smart Grid*, vol. 4, pp. 789–797, June 2013.

[7] S. Hatami and M. Pedram, "Minimizing the electricity bill of cooperative users under a quasi-dynamic pricing model," in *Proc. of IEEE Smart Grid Comm*, pp. 421–426, Oct 2010.

[8] Y. Hu and B. Defourny, "Optimal price-threshold control for battery operation with aging phenomenon: a quasiconvex optimization approach," *Annals of Oper. Res.*, May 2017.

[9] C. Jin, J. Tang, and P. Ghosh, "Optimizing electric vehicle charging with energy storage in the electricity market," *IEEE Trans. on Smart Grid*, vol. 4, pp. 311–320, March 2013.

[10] A. Oudalov, R. Cherkaoui, and A. Beguin, "Sizing and optimal operation of battery energy storage system for peak shaving application," in *Proc. of IEEE Lausanne Power Tech*, pp. 621–625, July 2007.

[11] C.-K. Chau, G. Zhang, and M. Chen, "Cost minimizing online algorithms for energy storage management with worst-case guarantee," *IEEE Trans. on Smart Grid*, vol. 7, no. 6, pp. 2691–2702, 2016.

[12] J. Qin, Y. Chow, J. Yang, and R. Rajagopal, "Online modified greedy algorithm for storage control under uncertainty," *IEEE Trans. on Power Systems*, vol. 31, no. 3, pp. 1729–1743, 2015.

[13] G. Vojvodic, A. I. Jarrah, and D. P. Morton, "Forward thresholds for operation of pumped-storage stations in the real-time energy market," *Eur. J. Oper. Res.*, vol. 254, no. 1, pp. 253–268, 2016.

[14] P. Harsha and M. Dahleh, "Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy," *IEEE Trans. on Power Systems*, vol. 30, pp. 1164–1181, May 2015.

[15] K. Zhou, C. Fu, and S. Yang, "Big data driven smart energy management: From big data to big insights," *Renewable and Sustainable Energy Reviews*, vol. 56, pp. 215 – 225, 2016.

[16] S. Samuelson and I. Yang, "Data-driven distributionally robust control of energy storage to manage wind power fluctuations," in *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pp. 199–204, Aug 2017.

[17] A. E. Thanos, M. Bastani, N. Celik, and C. Chen, "Dynamic data driven adaptive simulation framework for automated control in microgrids," *IEEE Trans. on Smart Grid*, vol. 8, pp. 209–218, Jan 2017.

[18] J. Wu, Z. Wang, Y. Yu, and C. Wu, "Optimal storage control for dynamic pricing." <https://arxiv.org/abs/1911.06963>, Nov. 2019.

[19] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–22, 1977.

[20] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth, "Occam's razor," *Information processing letters*, vol. 24, no. 6, pp. 377–380, 1987.

[21] D. M. Hawkins, "The problem of overfitting," *Journal of chemical information and computer sciences*, vol. 44, no. 1, pp. 1–12, 2004.

[22] A. A. Neath and J. E. Cavanaugh, "The bayesian information criterion: background, derivation, and applications," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 4, no. 2, pp. 199–203, 2012.

[23] J. Fan, H. Luo, and X. Lin, "Semantic video classification by integrating flexible mixture model with adaptive em algorithm," in *Proceedings of the 5th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 9–16, ACM, 2003.

[24] ENGIE, "Historical electricity pricing information." <https://www.engie.com/resources/>.

[25] PJM, "Metered load data." <https://www.pjm.com/markets-and-operations/ops-analysis/historical-load-data.aspx/>.

[26] A. Holder, "Parametric lp analysis," *Wiley Encyclopedia of Operations Research and Management Science*, 2010.




Fig. 10. Performance Improvement by Heuristics ($B = 20\%$ of peak hourly energy consumption).




Fig. 11. Performance Improvement by Heuristics ($B = \text{peak hourly energy consumption}$).

APPENDIX

A. Proof for Theorem 2

Proof To obtain the upper bound for $R(T)$, we establish the lower bound of $\mathbb{E}[w]^{\text{offline}}$ and the upper bound of $\mathbb{E}[w]^{\text{ETA}}$, respectively.

Define $p_T^{\min} = \min\{p(1), \dots, p(T)\}$. Then,

$$\mathbb{E}[w]^{\text{offline}} = \mathbb{E}[p_T^{\min}] = T \int_0^\infty p \bar{F}^{T-1}(p) f(p) dp, \quad (22)$$

where $\bar{F}(p)$ denotes the tail distribution of $p(t)$. Note that, Cauchy-Schwarz inequality guarantees:

$$\begin{aligned} \int_0^\infty p \bar{F}^T(p) f(p) dp \int_0^\infty p \bar{F}^{T-2}(p) f(p) dp \\ \geq \left(\int_0^\infty p \bar{F}^{T-1}(p) f(p) dp \right)^2. \end{aligned} \quad (23)$$

Hence, we can show that

$$\begin{aligned} \frac{\mathbb{E}[p_{T+1}^{\min}]}{\mathbb{E}[p_T^{\min}]} &= \frac{T+1}{T} \frac{\int_0^\infty p \bar{F}^T(p) f(p) dp}{\int_0^\infty p \bar{F}^{T-1}(p) f(p) dp} \\ &\geq \frac{T^2-1}{T^2} \frac{T \int_0^\infty p \bar{F}^{T-1}(p) f(p) dp}{(T-1) \int_0^\infty p \bar{F}^{T-2}(p) f(p) dp} \\ &= \frac{T^2-1}{T^2} \frac{\mathbb{E}[p_T^{\min}]}{\mathbb{E}[p_{T-1}^{\min}]}. \end{aligned} \quad (24)$$

Standard mathematical manipulation yields

$$\begin{aligned} \frac{(T+1)(T-1)}{T^2} \frac{\mathbb{E}[p_T^{\min}]}{\mathbb{E}[p_{T-1}^{\min}]} \\ \geq \frac{(T+1) \cdot 1}{T \cdot 2} \frac{\mathbb{E}[p_2^{\min}]}{\mathbb{E}[p]} = \frac{T+1}{T} \alpha, \end{aligned} \quad (25)$$

where α is defined in Eq. (13) and $\alpha < \frac{1}{2}$. Thus, we construct the exponentially decreasing lower bound for $\mathbb{E}[w]^{\text{offline}}$:

$$\mathbb{E}[w]^{\text{offline}} = \mathbb{E}[p_T^{\min}] \geq T \alpha^{T-1} \mathbb{E}[p]. \quad (26)$$

Next, we construct the upper bound for $\mathbb{E}[w]^{\text{ETA}} = \theta_{T+1}$. Using the recursive structure of threshold, we can show

$$\theta_T - \theta_{T+1} = \int_0^{\theta_T} F(p) dp \geq \frac{1}{2} \beta_T \theta_T^2, \quad (27)$$

where β_T is defined in Eq. (14). To construct the upper bound, it suffices to notice the following inequality holds:

$$\frac{1}{\theta_{T+1}} - \frac{1}{\theta_T} \geq \frac{\theta_T - \theta_{T+1}}{\theta_T^2} \geq \frac{\beta_T}{2}. \quad (28)$$

And when $T = 1$, we have to purchase for the unit demand at any price. Hence, $\theta_1 = \infty$. This yields

$$\theta_{T+1} \leq \frac{2}{\sum_{i=2}^T \beta_i}. \quad (29)$$

As long as $\sum_{i=2}^T \beta_i > 0$, then, the desirable upper bound immediately follows. ■

B. Storage Sizing Problem

To decide the optimal storage sizing, it suffices to understand the marginal benefit of storage in this problem, which can be answered by the following parametric analysis:

$$\begin{aligned} \text{MinC}(B) &:= \min \sum_{t=1}^T \mathbb{E}[(g(t) + b(t))p(t)] \\ \text{s.t. } &g(t) + c(t) = \mathbb{E}[d(t)], \quad \forall t \in [1, T], \\ &\sum_{\tau=1}^t (b(\tau) - c(\tau)) \leq B, \quad \forall t \in [1, T]. \end{aligned} \quad (30)$$

Notice that the function $\text{MinC}(B)$ is continuous, non-increasing and concave [26]. Hence, the derivative of $\text{MinC}(B)$ is also non-increasing. Denote the amortized cost for storage system over time span T by π_b . Then, the optimal storage capacity B^* is the solution to the following equation:

$$\pi_b = \text{MinC}'(B^*). \quad (31)$$

Thus, by selecting the proper amortized cost, the optimal storage capacity can be immediately obtained.