

Online Learning of Load Elasticity for Electric Vehicle Charging

Nasim Yahya Soltani, Seung-Jun Kim, and Georgios B. Giannakis

Dept. of ECE, University of Minnesota, 200 Union St. SE, Minneapolis, MN 55455, USA

E-mail:{yahya015, seungjun, georgios}@umn.edu

Abstract—While electric vehicles (EVs) are expected to provide environmental and economical benefits, judicious coordination of EV charging may be necessary to prevent overloading of the distribution grid. Leveraging the smart grid infrastructure, the utility company can adjust the electricity price intelligently for individual customers to elicit desirable load curves. In this context, the present paper addresses the problem of predicting the EV charging behavior of the consumers at different prices, which is a prerequisite for the price adjustment. The dependencies on price responsiveness among neighbouring consumers are captured by adopting a conditional random field (CRF) model. To account for temporal dynamics even in an adversarial setting, the framework of online convex optimization is adopted to develop an efficient online algorithm for estimating the CRF parameters. Numerical tests verify the proposed approach.

I. INTRODUCTION

The smart grid vision aims at capitalizing on information technology for efficient use of energy resources in grid operation without sacrificing user satisfaction. Recently, there has been growing interest in electric vehicles (EVs), which are expected to be widely deployed by 2050 [1]. EV penetration may contribute to alleviating dependence on fossil fuel and cutting greenhouse gas emissions. EV owners may also benefit from lower energy cost in the face of spiking gasoline prices.

Although EVs can undoubtedly effect environmental and economical benefits, simultaneous charging of a large number of EVs in residential distribution grids can significantly overload the infrastructure. Also, uncoordinated EV charging can aggravate load peaks due to, e.g., concentrated charging demand before commuting hours, resulting in higher generation cost for the utility company. Therefore, coordinated charging of EVs is a crucial task for efficient grid operation.

In order to elicit desirable electricity consumption patterns, various time-of-use (ToU) pricing schemes have been proposed. By allowing the electricity price to vary over different hours of a day, consumers are encouraged to shift inessential loads to the periods of low prices, which generally correspond to off-peak hours [2]. There is extensive literature on scheduling the time and rate of EV charging [3], [4], [5]. Generation and EV charging costs were minimized in [3] under power flow constraints. With the goal of shifting EV loads to fill the overnight demand valley, a distributed algorithm for day-ahead charging rate schedules was proposed in [4]. The EV

This work is supported by NSF grant 1202135, and by the Inst. of Renewable Energy and the Environment (IREE) under grant No. RL-0010-13.

charging schedule was optimized in [5] by minimizing load variance and maximizing load factor.

Prerequisite to the demand coordination task is to obtain reliable analytics. In addition to the essential load and price forecasting, it is instrumental also to learn the consumers' behavioral patterns. In [6], smart grid consumers' price elasticity was estimated using a linear regression model with price changes as regressors and the corresponding shift in total demand as the response. The price responsiveness is useful, for instance, when one desires to set the prices optimally with various objectives such as minimizing the generation cost or maximizing revenue of the utility company.

However, existing techniques for acquiring price elasticity fall short of capturing the following important aspects. First, the dynamics of consumer preferences over time have been not accounted for. In practice, price elasticity might change abruptly and even in an adversarial manner, as the consumers may also react so as to maximize their own profit. Secondly, the spatial dependencies of consumer behaviors, e.g., the correlations existent in the behaviors of consumers in geographic proximity, have not been exploited. In terms of algorithm implementation, online algorithms are preferred over batch alternatives for real-time streaming analytics.

The present work adopts a model that captures the dependency of EV consumers' charging decisions on the announced electricity price. As the charging decisions can be best described by discrete values (e.g., “charging” or “not charging”), a logistic regression-type framework is employed. To capture spatial dependence of the behaviors, the dependency structure is encoded as a graph, and the overall model corresponds to a conditional random field (CRF) [7]. Subsequently, an online algorithm is developed to estimate the relevant model parameters, based on an online learning framework, which provides performance guarantees with minimal assumptions on the structure of temporal dynamics [8].

The rest of the paper is organized as follows. Sec. II states the problem and recaps the CRF modeling framework. Sec. III develops an online algorithm for estimating the model parameters. Preliminary results of numerical tests are presented in Sec. IV, and conclusions are provided in Sec. V.

II. PROBLEM STATEMENT

A. System model

Consider N EV owners, who want to charge their EVs via the distribution grid. Let $b_i^t \in \mathcal{S} := \{0, 1\}$ indicate

the charging behavior of consumer i at time t ; i.e., $b_i^t = 1$ when consumer i is charging her EV at time t , and $b_i^t = 0$, otherwise.¹ Similarly, ρ_i^t denotes the electricity price during time slot t for consumer i . To capture spatial dependence (e.g., behavioral dependence of consumers living in the same neighborhood, or having similar income levels), an undirected graph $G = (V, E)$ is introduced, where the vertex set $V := \{1, 2, \dots, N\}$ corresponds to the consumers, and edges $(i, j) \in E$ capture the dependence between consumers i and j .

It is assumed that the customer premises are equipped with smart meters so that bi-directional communication between the utility and the consumers is feasible. Leveraging such an infrastructure, the utility announces prices $\{\rho_i^t\}$ for all consumers $i \in V$ at the beginning of time slot $t = 1, 2, \dots$. Consequently, the charging decisions $\{b_i^t\}$ of the consumers are reported back to the utility at the end of time slot t .

In this context, the following problem is of interest: estimate the probability with which each consumer $i \in V$ will charge her EV at time $t + 1$ paying price $\{\rho_i^{t+1}\}$, given past prices $\{\rho_i^\tau, i \in V, \tau = 1, \dots, t\}$, and the corresponding observed behaviors $\{b_i^\tau, i \in V, \tau = 1, \dots, t\}$, while accounting for possible spatial dependencies in $\{b_i^t\}$ captured by G .

B. CRF model for EV charging behavior

To solve the aforementioned problem, the framework of conditional random fields (CRFs) is adopted [7]. Collect in vectors \mathbf{b}^t and $\boldsymbol{\rho}^t$ variables $\{b_i^t\}_{i=1}^N$ and $\{\rho_i^t\}_{i=1}^N$, respectively. The CRF models the conditional probability distribution function (pdf) $p(\mathbf{b}^t | \boldsymbol{\rho}^t)$. In short, $p(\mathbf{b}^t | \boldsymbol{\rho}^t)$ is a CRF with respect to G if it obeys the Markov property for every $\boldsymbol{\rho}^t$. This means that conditioned on $\boldsymbol{\rho}^t$, for any $i, j \in V$, behavior b_i^t is independent of b_j^t given the neighbors $\{b_k^t : (i, k) \in E\}$. Intuitively, this means that given $\boldsymbol{\rho}^t$, the behavior of the neighbors of b_i^t contains all the information needed for predicting b_i^t , and other variables are redundant.

Let $\psi_{i,j}(b_i^t, b_j^t)$ denote feature functions quantifying the dependency in charging behavior of consumers i and j . In addition, functions $\phi_i(b_i^t, \rho_i^t)$ model the dependency of b_i^t on price ρ_i^t . Parameters $\boldsymbol{\theta}_i^t$ and $\boldsymbol{\theta}_{i,j}^t$ are introduced for ϕ_i and $\psi_{i,j}$, respectively, to capture the strengths of these dependencies. With $\boldsymbol{\theta}^t := [\boldsymbol{\theta}_i^t, \boldsymbol{\theta}_{i,j}^t]$, the price-conditional behavior pdf can thus be modeled as

$$p_{\boldsymbol{\theta}^t}(\mathbf{b}^t | \boldsymbol{\rho}^t) = \frac{1}{Z(\boldsymbol{\rho}^t)} \prod_{i \in V} e^{\langle \boldsymbol{\theta}_i^t, \phi_i(b_i^t, \rho_i^t) \rangle} \prod_{(i,j) \in E} e^{\langle \boldsymbol{\theta}_{i,j}^t, \psi_{i,j}(b_i^t, b_j^t) \rangle} \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, and

$$Z(\boldsymbol{\rho}^t) := \sum_{\mathbf{b}^t} \prod_{i \in V} e^{\langle \boldsymbol{\theta}_i^t, \phi_i(b_i^t, \rho_i^t) \rangle} \prod_{(i,j) \in E} e^{\langle \boldsymbol{\theta}_{i,j}^t, \psi_{i,j}(b_i^t, b_j^t) \rangle} \quad (2)$$

is a normalization factor, also known as the partition function. Motivated by the CRF model pdf involved with logistic

¹Multiple charging rates can be accommodated straightforwardly by increasing the number of labels in \mathcal{S} .

regression, we adopt a scalar function $\phi_i(b_i^t, \rho_i^t) = b_i^t \rho_i^t$. Furthermore, inspired by the Ising model for modeling dependencies of binary random variables, we choose $\psi_{i,j}(b_i^t, b_j^t) = b_i^t b_j^t$ [9]. Now the problem of finding $p_{\boldsymbol{\theta}^{t+1}}(\mathbf{b}^{t+1} | \boldsymbol{\rho}^{t+1})$ given $\{\rho_i^\tau, b_i^\tau, i \in V\}_{\tau=1}^t$ translates to estimating $\boldsymbol{\theta}^{t+1}$ at each time t , given the past prices and charging decisions observed.

III. ONLINE LEARNING OF LOAD ELASTICITY

Compared to batch algorithms that process the entire collection of data to obtain the desired estimates, online algorithms feature the capability to process data one by one in a recursive fashion. To develop an online algorithm for estimating $\boldsymbol{\theta}^t$, the approach here utilizes online convex optimization, which requires minimal assumptions on the temporal dynamics of $\boldsymbol{\theta}^t$, and can provide provable performance guarantees even in adversarial settings [8]. These guarantees in the “adversarial” setting are meaningful because the consumers can also act strategically to maximize their own benefit (and mend their price responsiveness accordingly). Next, the online convex programming framework is outlined first.

A. Online convex programming

Online convex programming can be viewed as a multi-round game with a forecaster and an adversary. The loss functions $l^t(\cdot)$ associated with the forecasts for $t = 1, 2, \dots, T$ and the feasible set Θ are assumed convex. In round t , the forecaster chooses $\hat{\boldsymbol{\theta}}^t \in \Theta$, after which the adversary reveals $l^t(\cdot)$, incurring loss $l^t(\hat{\boldsymbol{\theta}}^t)$ for the t -th round. Performance of the online actions $\{\hat{\boldsymbol{\theta}}^t\}$ is assessed through the so-termed *regret* that is given by

$$R_T = \sum_{t=1}^T l^t(\hat{\boldsymbol{\theta}}^t) - \min_{\boldsymbol{\theta} \in \Theta} \sum_{t=1}^T l^t(\boldsymbol{\theta}) \quad (3)$$

and represents the relative cumulative loss of the *online* forecaster after T rounds, compared to an optimal *offline* minimizer, which has the advantage of hindsight. Online convex programming algorithms provide ways to generate $\hat{\boldsymbol{\theta}}^t$ to achieve a regret that is sublinear in T .

B. Dynamic learning of load elasticity

In our context, the forecaster is the utility company and the adversaries are the EV owners. The loss represented by the negative log-likelihood function $l^t(\boldsymbol{\theta}^t) := -\log p_{\boldsymbol{\theta}^t}(\mathbf{b}^t | \boldsymbol{\rho}^t)$ is not revealed to the utility company until the utility predicts $\hat{\boldsymbol{\theta}}^t$ and announces $\boldsymbol{\rho}^t$ based on $\hat{\boldsymbol{\theta}}^t$, since only then can the consumers report their charging decisions \mathbf{b}^t . Note that, the chosen loss $l^t(\boldsymbol{\theta}^t)$ with $p_{\boldsymbol{\theta}^t}(\mathbf{b}^t | \boldsymbol{\rho}^t)$ as in (1) is convex [10].

A popular online convex programming algorithm relies on the online mirror descent (OMD) iteration, which is a projected subgradient method with the Bregman divergence used as a proximal term. It yields an efficient first-order algorithm with sublinear convergence rate [11]. Vector $\hat{\boldsymbol{\theta}}^{t+1}$ is obtained in OMD as

$$\hat{\boldsymbol{\theta}}^{t+1} = \arg \min_{\boldsymbol{\theta}} \langle \nabla l^t(\hat{\boldsymbol{\theta}}^t), \boldsymbol{\theta} \rangle + \frac{1}{\mu_t} D(\hat{\boldsymbol{\theta}}^t \| \boldsymbol{\theta}) \quad (4)$$

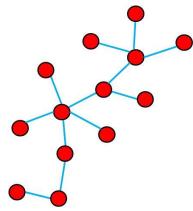


Fig. 1: Spatial dependency graph.

```

1: Initialize  $\hat{\theta}^1$ 
2: For  $t = 1, 2, \dots$  do
3:   Set prices  $\rho^t$  based on  $\hat{\theta}^t$ .
4:   Collect EV charging decisions  $\mathbf{b}^t$ .
5:   Compute  $\nabla l^t(\hat{\theta}^t)$  using BP.
6:   Update  $\hat{\theta}^{t+1} = \hat{\theta}^t - \mu_t \nabla l^t(\hat{\theta}^t)$ 
7: Next  $t$ 

```

TABLE I: Overall online algorithm.

where μ_t denotes a step size, and $D(\cdot||\cdot)$ represents the Bregman divergence. The Bregman divergence associated with the ℓ_2 -norm is simply given by $D(\hat{\theta}^t||\theta) = \frac{1}{2}\|\theta - \hat{\theta}^t\|^2$. Upon substituting this into (4), the OMD update boils down to an online gradient descent given by

$$\hat{\theta}^{t+1} = \hat{\theta}^t - \mu_t \nabla \ell^t(\hat{\theta}^t). \quad (5)$$

To evaluate the gradient, note that

$$\frac{\partial \ell^t}{\partial \theta_i^t} = \mathbb{E}[b_i^t \rho_i^t | \rho^t] - b_i^t \rho_i^t, \quad \forall i \in V \quad (6a)$$

$$\frac{\partial \ell^t}{\partial \theta_{i,j}^t} = \mathbb{E}[b_i^t b_j^t | \rho^t] - b_i^t b_j^t, \quad \forall (i, j) \in E \quad (6b)$$

where the expectation is with respect to $p_{\hat{\theta}^t}(\mathbf{b}^t | \rho^t)$. It is straightforward to see that $\mathbb{E}[b_i^t \rho_i^t | \rho^t] = \rho_i^t p_{\hat{\theta}^t}(b_i^t = 1 | \rho^t)$ and $\mathbb{E}[b_i^t b_j^t | \rho^t] = p_{\hat{\theta}^t}(b_i^t = 1, b_j^t = 1 | \rho^t)$. These marginal conditional probabilities can be efficiently evaluated by employing the belief propagation (BP) algorithm, which converges to the correct parameters when the underlying graph is a tree [12].

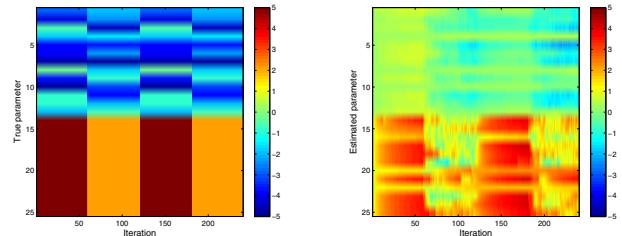
The overall online algorithm for estimating θ^{t+1} is summarized in Table I. Steps 5 and 6 confirm that $\hat{\theta}^{t+1}$ is dependent on past values of $\{\theta^\tau\}_{\tau=1}^t$, and thus on $\{\mathbf{b}^\tau, \rho^\tau\}_{\tau=1}^t$ as well. Therefore, (5) makes use of the information in the entire input history. Regarding step 3 of the algorithm, the utility can optimally determine price ρ^t using stochastic optimization to minimize a certain power system cost while accounting for the consumer behavior through $p_{\hat{\theta}^t}(\mathbf{b}^t | \rho^t)$. Specifics of this price setting step will be included in the journal version of this work.

Clearly, the consumers' charging decisions are correlated across time in practice. Compared to stochastic approximation alternatives [13], the novel algorithm based on online convex programming requires minimal assumptions on the structure of temporal correlation of data (charging decisions). In addition, the framework accommodates strategic actions of the consumers as well.

C. Performance analysis

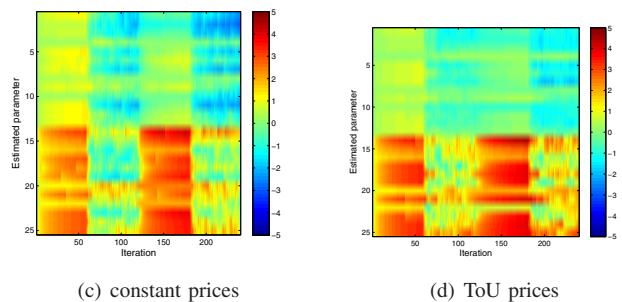
The algorithm in Table I yields a regret bound that is sublinear in T , as described in the following proposition, whose proof is omitted due to space limitation.

Proposition 1: Let $N_e := |E|$ denote the number of edges in E . If $\max\{|\theta_i^t|, |\theta_{i,j}^t|\} \leq \theta_0$ and $|\rho_i^t| \leq \rho_0$ for all $t, i \in V$, and $(i, j) \in E$, then for $\{\hat{\theta}^t\}$ obtained from the algorithm in



(a) true parameters

(b) random prices



(c) constant prices

(d) ToU prices

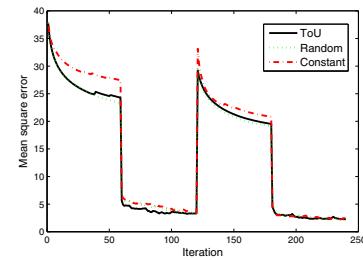
Fig. 2: True and estimated θ^t .

Fig. 3: Average squared prediction error for CRF parameters.

Table I, it holds that

$$R_T \leq \theta_0 \sqrt{2(N\rho_0^2 + N_e)T} = O(\sqrt{T}). \quad (7)$$

It is natural to assume that the price is bounded. In addition, to account for consumers that do not respond to price changes (corresponding to $\theta_i^t = \pm\infty$), one can use a sufficiently large bound for $\max\{|\theta_i^t|, |\theta_{i,j}^t|\}$ in practice. Thus, the conditions in the proposition can be readily satisfied.

D. Logistic regression benchmark

If one neglects the spatial dependencies by setting $\theta_{i,j} = 0$ for $(i, j) \in E$, the CRF model reduces to N parallel logistic regression models; that is [cf. (1)]

$$p_{\theta^t}(\mathbf{b}^t | \rho^t) = \frac{1}{Z(\rho^t)} e^{\sum_{i \in V} \theta_i^t b_i^t \rho_i^t} = \prod_{i=1}^N \frac{e^{\theta_i^t b_i^t \rho_i^t}}{1 + e^{\theta_i^t \rho_i^t}}. \quad (8)$$

To obtain online estimates of $\{\theta_i^t\}$, the algorithm in Table I is again applicable, but the gradient evaluation can be performed without using BP, simply as

$$\frac{\partial \ell^t}{\partial \theta_i^t} = \frac{\rho_i^t e^{\theta_i^t \rho_i^t}}{1 + e^{\theta_i^t \rho_i^t}} - b_i^t \rho_i^t, \quad i \in V. \quad (9)$$

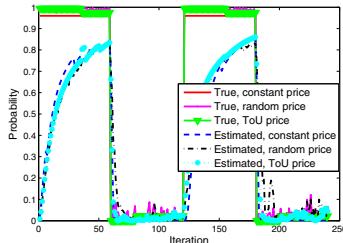


Fig. 4: Joint probability of consumer charging decisions.

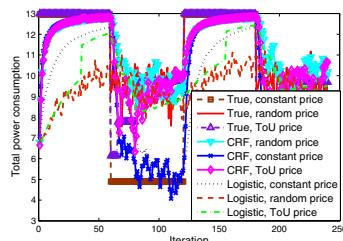


Fig. 5: Predicted total load.

IV. NUMERICAL TESTS

The performance of the proposed algorithm was verified via numerical tests. A set of 13 EV owners was considered, whose spatial dependencies were captured by the graph G of Fig. 1. Given the true CRF parameters θ^{*t} , the charging decisions b^t were deviates from the CRF model in (8), with ρ^t chosen as explained next. The values of θ^{*t} were changed occasionally but otherwise were kept fixed over time; see Fig. 2(a), which depicts θ^{*t} . As price setting based on the estimated model is not addressed in this work, three different strategies for selecting ρ^t independent of estimated parameters were considered for the simulations: s1) constant prices; s2) ToU prices; and s3) random prices. For the constant pricing setup, ρ^t was set to 1 for all consumers and held constant across time. For ToU pricing, the following was used:

$$\rho_i^t = \begin{cases} \$0.508, & 7\text{am}-2\text{pm} \\ \$0.880, & 2\text{pm}-8\text{pm} \\ \$0.722, & 8\text{pm}-11\text{pm} \\ \$0.508, & 11\text{pm}-7\text{am} \end{cases} \quad \forall i \in V. \quad (10)$$

In random pricing, the prices were randomly selected from a uniform distribution over $[0, 1]$. Prices were updated every 12 minutes, which corresponds to the duration of one time slot. A step size $\mu_t = 0.72$ was used.

Fig. 2 shows true and estimated parameters for the different pricing strategies. It can be seen that the parameters are tracked approximately.

Fig. 3 depicts the average squared prediction error of the parameters, averaged over $N + N_e = 25$ parameters. For all price setting mechanisms, the errors tend to decrease as iterations progress, while sharp changes in the prediction error result whenever the parameter values are changed abruptly.

A comparison of $p_{\hat{\theta}^t}(b_i^t | \rho^t)$ with $p_{\theta^{*t}}(b_i^t | \rho^t)$ per iteration is shown in Fig. 4. It can be seen that the joint probability of the charging decisions is tracked very well.

Once the joint probabilities $p(b^t | \rho^t)$ have been estimated, the expected value of the total load can be readily found as

$$P_{tot}^t := P_{EV} \sum_{i=1}^N p_{\hat{\theta}^t}(b_i^t = 1 | \rho^t) \quad (11)$$

where P_{EV} is the load due to charging of a single EV. The predicted total loads are depicted in Fig. 5, where the CRF-based estimates and the logistic regression-based ones are plotted together for comparison. The CRF-based algorithm achieves the performance gain by incorporating spatial dependencies.

V. CONCLUSIONS

An algorithm to predict the elasticity of individual EV charging loads was developed. Such information is essential for the utility company to set electricity prices in real time and thus be able to coordinate EV charging. The probabilities that EV consumers charge their vehicles when presented with arbitrary ToU prices were obtained based on a CRF model, in which the spatial dependency of the customers' behavior was captured. Without explicit models for temporal dynamics, an online algorithm to estimate the CRF parameters was derived in the framework of online convex optimization. Performance of the proposed algorithm was verified by numerical tests.

REFERENCES

- [1] International Energy Agency, "Technology roadmap electric and plug-in hybrid electric vehicles," http://www.iea.org/publications/freepublications/publication/EV_PHEV_Roadmap.pdf, 2011.
- [2] K. Parks, P. Denholm, and T. Markel, "Costs and emissions associated with plug-in hybrid electric vehicle charging in the XCEL energy Colorado service territory," Ntl. Renewable Energy Lab. (NREL), Tech. Rep. NREL/TP-640-41410, 2007.
- [3] S. Sojoudi and S. H. Low, "Optimal charging of plug-in hybrid electric vehicles in smart grids," in *Proc. of IEEE PES General Meeting*, Detroit, Michigan, Jul. 2011.
- [4] L. Gan, U. Topcu, and S. H. Low, "Optimal decentralized protocol for electric vehicle charging," in *Proc. of 50th IEEE Conf. on Decision and Control and European Control Conf.*, Orlando, FL, Dec. 2011.
- [5] E. Sortomme, M. M. Hindi, S. D. J. MacPherson, and S. S. Venkata, "Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses," *IEEE Trans. Smart Grid*, vol. 2, no. 1, pp. 198–205, Mar. 2011.
- [6] V. Gomez, M. Chertkov, S. Backhaus, and H. Kappen, "Learning price-elasticity of smart consumers in power distribution systems," in *Proc. of IEEE SmartGridComm*, Tainan City, Taiwan, Nov. 2012.
- [7] J. Lafferty, A. McCallum, and F. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *Proc. of Intl. Conf. on Mach. Learning*, Williamstown, MA, 2001.
- [8] S. Shalev-Shwartz, "Online learning and online convex optimization." *Foundations and Trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2011.
- [9] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Foundations and Trends in Machine Learning*, vol. 1, no. 1–2, pp. 1–305, 2008.
- [10] M. Raginsky, R. F. Marcia, S. Jorge, and R. Willett, "Sequential probability assignment via online convex programming using exponential families," in *Proc. of IEEE Intl. Symp. Info. Theory*, Seoul, Korea, Jun.-Jul. 2009.
- [11] N. Srebro, K. Sridharan, and A. Tewari, "On the universality of online mirror descent," in *Proc. of Advances in Neural Information Processing Systems*, Granada, Spain, Dec. 2011.
- [12] J. S. Yedidia, W. T. Freeman, and Y. Weiss, "Bethe free energy, Kikuchi approximations, and belief propagation algorithms," Mitsubishi Electric Research Laboratories, Tech. Rep. TR-2001-16, May 2001.
- [13] H. Kushner and G. Yin, *Stochastic Approximation Algorithms and Applications*, 2nd ed. Springer-Verlag, 2003.