Electricity Market Forecasting via Low-Rank Multi-Kernel Learning

Vassilis Kekatos, *Member, IEEE*, Yu Zhang, *Student Member, IEEE*, and Georgios B. Giannakis, *Fellow, IEEE*

***Abstract*—The smart grid vision entails advanced information technology and data analytics to enhance the efficiency, sustainability, and economics of the power grid infrastructure. Aligned to this end, modern statistical learning tools are leveraged here for electricity market inference. Day-ahead price forecasting is cast as a low-rank kernel learning problem. Uniquely exploiting the market clearing process, congestion patterns are modeled as rank-one components in the matrix of spatio-temporally varying prices. Through a novel nuclear norm-based regularization, kernels across pricing nodes and hours can be systematically selected. Even though market-wide forecasting is beneficial from a learning perspective, it involves processing high-dimensional market data. The latter becomes possible after devising a block-coordinate descent algorithm for solving the non-convex optimization problem involved. The algorithm utilizes results from block-sparse vector recovery and is guaranteed to converge to a stationary point. Numerical tests on real data from the Midwest ISO (MISO) market corroborate the prediction accuracy, computational efficiency, and the interpretative merits of the developed approach over existing alternatives.**

***Index Terms*—Block-coordinate descent, day-ahead energy prices, graph Laplacian, kernel-based learning, learning, low-rank matrix, multi-kernel learning, nuclear norm regularization.**

I. Introduction

FORECASTING electricity prices is an important decision making tool for market participants [4]. Conventional and particularly renewable asset owners plan their trading and bidding strategies according to pricing predictions. Moreover, independent system operators (ISOs) recently broadcast their own market forecasts to proactively relieve congestion [12]. At a larger geographical and time scale, electricity price analytics based solely on publicly available data rather than physical system modeling are pursued by government services to identify "national interest transmission congestion corridors" [39].

In a generic electricity market setup, an ISO collects bids submitted by generator owners and utilities [15], [24]. Compliant with network and reliability constraints, the grid is dispatched in the most economical way. Following power demand patterns, electricity prices exhibit cyclo-stationary motifs over time. More importantly and due to transmission limitations, cheap electricity cannot be delivered everywhere across the grid. Out-of-merit energy sources have to be dispatched to balance the load. Hence, congestion together with heat losses lead to spatially-varying energy prices, known as locational marginal prices (LMPs) [24], [17].

Schemes for predicting electricity prices proposed so far include time-series analysis approaches based on auto-regressive (integrated) moving average models and their generalizations [10], [14]. However, these models are confined to linear predictors, whereas markets involve generally nonlinear dependencies. To account for nonlinearities, artificial intelligence approaches, such as fuzzy systems and neural networks, have been investigated [42], [27], [40]. Hidden Markov models have been also advocated [19]. A nearest neighborhood method was suggested in [28]. Market clearance was solved as a quadratic program and forecasts were extracted based on the most probable outage combinations in [43]. Reviews on electricity price forecasting and the associated challenges can be found in [4] and [34].

Different from existing approaches where predictors are trained on a per-node basis, a framework for learning the entire market is pursued in this work. Building on collaborative filtering ideas, market forecasting is cast as a learning task over all nodes and several hours [2], [5]. Leveraging market clearing characteristics, prices are modeled as the superposition of several rank-one components, each capturing particular spatio-temporal congestion motifs. Distinct from [23], *low-rank* kernel-based learning models are developed here.

A systematic kernel selection methodology is the second contribution of this paper. Due to the postulated decomposition, different kernels must be defined over nodes and hours. Our novel analytic results extend kernel learning tools to low-rank multi-task models [30], [18], [3]. By viewing market extrapolation as learning over a graph, the commercial pricing network is surrogated here via balancing authority connections and meaningful graph Laplacian-based kernels are provided.

An efficient algorithm for solving the computationally demanding optimization involved is our third contribution. Although the problem is jointly non-convex, per block optimizations entail convex yet non-differentiable costs which are tackled via a block-coordinate descent approach. Leveraging results from (block) compressed sensing [32], the resultant algorithm boils down to univariate minimizations, exploits the Kronecker product structure, and is guaranteed to converge to a stationary point of the resultant optimization problem.

Manuscript received October 02, 2013; revised March 21, 2014; accepted July 01, 2014. Date of publication July 08, 2014; date of current version November 18, 2014. This work was supported by the Institute of Renewable Energy and the Environment (IREE) under Grant RL-0010-13, the University of Minnesota, and NSF-ECCS Grants 1202135 and 1343248. The guest editor coordinating the review of this manuscript and approving it for publication was Prof. Danilo Mandic.

The authors are with the Electrical and Computer Engineering Department, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: kekatos@umn.edu; zhan1220@umn.edu; georgios@umn.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSTSP.2014.2336611

Forecasting results on the MISO market over the summer of 2012 corroborate the accuracy, interpretative merit, and the computational efficiency of the novel learning model.

Notation: Lower- (upper-) case boldface letters denote column vectors (matrices); calligraphic letters stand for sets. Symbols $(\cdot)^\top$ and $\otimes$ denote transposition and the Kronecker product, respectively. The ℓ_2-norm of a vector is denoted by $\|\mathbf{a}\|_2$, $\|\mathbf{A}\|_F$ is the Frobenius matrix norm, and $\mathbb{S}_{++}^N$ is the set of $N \times N$ positive definite matrices. The operation $\mathrm{vec}(\mathbf{A})$ turns matrix $\mathbf{A}$ to a vector by stacking its columns, and $\mathrm{Tr}(\mathbf{A})$ denotes its trace. The property (P): $\mathrm{vec}(\mathbf{AXB}) = \left(\mathbf{B}^\top \otimes \mathbf{A}\right) \mathrm{vec}(\mathbf{X})$ will be needed throughout.

The paper outline is as follows. Electricity market forecasting is formulated in Section II, where the novel approach is presented. A block-coordinate descent algorithm is detailed in Section IV. Kernel design and forecasting results on the MISO market are in Section V. The paper is concluded in Section VI.

II. Problem Statement and Formulation

A. Preliminaries on Kernel-Based Learning

Given pairs $\{(x_n, z_n)\}_{n=1}^N$ of features x_n belonging to a measurable space $\mathcal{X}$ and target values $z_n \in \mathbb{R}$, kernel-based learning aims at finding a relationship $f : \mathcal{X} \to \mathbb{R}$ with f belonging to the linear function space

$$\mathcal{H}_\mathcal{K} := \left\{ f(x) = \sum_{n=1}^{\infty} K(x, x_n) a_n,\ a_n \in \mathbb{R} \right\} \tag{1}$$

defined by a preselected kernel (basis) $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ and corresponding coefficients a_n. When $K(\cdot, \cdot)$ is a symmetric positive definite function, then $\mathcal{H}_\mathcal{K}$ becomes a reproducing kernel Hilbert space (RKHS) whose members have a finite norm $\|f\|_\mathcal{K}^2 := \sum_{n=1}^\infty \sum_{m=1}^\infty K(x_n, x_m) a_n a_m$ [6].

Viewed either from a Bayesian estimation perspective, or as a function approximation task, learning f can be posed as the regularization problem [20], [7]

$$\hat{f} := \arg \min_{f \in \mathcal{H}_\mathcal{K}} \sum_{n=1}^{N} (z_n - f(x_n))^2 + \mu \|f\|_\mathcal{K}. \tag{2}$$

The least-squares (LS) fitting component in (2) captures the designer's reliance on data, whereas the regularizer $\|f\|_\mathcal{K}$ constraints $f \in \mathcal{H}_\mathcal{K}$ and facilitates generalization over unseen data. The two components are balanced through the parameter $\mu > 0$, which is typically tuned via cross-validation [20].

Finding $\hat{f}$ requires solving the functional optimization in (2). Fortunately, the celebrated Representer's Theorem asserts that $\hat{f}$ admits the form $\hat{f}(x) = \sum_{n=1}^N K(x, x_n) \hat{a}_n$ [20]. Hence, the sought $\hat{f}$ can be characterized by the coefficient vector $\hat{\mathbf{a}} := [\hat{a}_1 \cdots \hat{a}_N]^\top$. Upon defining the kernel matrix $\mathbf{K} \in \mathbb{S}_{++}^N$ having entries $[\mathbf{K}]_{n,m} := K(x_n, x_m)$, the vector $\mathbf{z} := [z_1 \cdots z_N]^\top$, and the norm $\|\mathbf{a}\|_\mathbf{K}^2 := \mathbf{a}^\top \mathbf{K} \mathbf{a}$; solving (2) is equivalent to the vector optimization

$$\hat{\mathbf{a}} := \arg \min_{\mathbf{a}}\ \|\mathbf{z} - \mathbf{K}\mathbf{a}\|_2^2 + \mu \|\mathbf{a}\|_\mathbf{K}. \tag{3}$$

Building on kernel-based learning, novel models pertinent to electricity market forecasting are pursued next.

B. Low-Rank Learning

Consider a wholesale electricity market over a set $\mathcal{N}$ of N commercial pricing nodes (CPNs) indexed by n. In a day-ahead market, locational marginal prices (LMPs) correspond to the cost of buying or selling electricity at each CPN and over one-hour periods for the following day [31], [17].

Viewing market forecasting as an inference problem, day-ahead LMPs are the target variables to be learned. Explanatory variables (features or regressors) can be any data available at the time of forecasting and believed to be relevant to the target variables. Due to the spatiotemporal nature of the problem, features can be either related to a CPN (nodal features), or to a specific market hour (time features).

Candidate nodal features could be the node type (generator, load, interface to another market); the generator technology (coal, natural gas, nuclear, or hydroelectric plant, wind farm); the CPN's geographical location; and the balancing authority controlling the node. Vector $\mathbf{x}_n$ collects the features related to the n-th CPN.

Vector $\mathbf{y}_t$ comprises the features related to a market period t. Candidate features include:

- same-hour LMPs from past days;
- load estimates (issued per balancing authority, region, and/or the market footprint);
- weather forecasts (e.g., temperature, humidity, wind speed, and solar radiation at selected locations);
- outage capacity (capacity of generation units closed for maintenance);
- timestamp features (hour of the day, day of the week, month of the year, holiday) to capture peak-demand hours on weekdays as well as heating and cooling patterns;
- scheduled power imports and exports to other markets.

Note that $\mathbf{y}_t$ is shared across CPNs: Weather forecasts across major cities or renewable energy sites affect several CPNs, while capacity outages, regional load estimates, and timestamps relate to the whole market. Moreover, spatially local features could not be easily related to specific CPNs, when CPN locations are unknown.

A generic approach could be to predict every single-CPN price given $\mathbf{y}_t$ and the observed LMPs. Such an approach would train N separate prediction models with identical feature variables. However, locational prices are not independent: they are determined over a transmission grid having capacity and reliability limitations [15], [22]. Leveraging this network-imposed dependence, market forecasting is uniquely interpreted here as learning over a graph; see e.g., [25]. Energy markets may change significantly due to lasting transmission and generation outages, or shifts in oil or gas markets. That is why the market is considered to be stationary only over the T most recent time periods, which together with the sought next 24 hours comprise the set $\mathcal{T}$. The market could be then thought of as a function $p : \mathcal{N} \times \mathcal{T} \to \mathbb{R}$ to be inferred.

We postulate that the price at node n and time t denoted by $p(n, t)$ belongs to the RKHS defined by the tensor product

kernel $K_{\otimes}((n,t),(n',t')) := K(n,n')G(t,t')$, where $K : \mathcal{N}\times\mathcal{N}\to\mathbb{R}$ and $G:\mathcal{T}\times\mathcal{T}\to\mathbb{R}$ are judiciously selected kernels over nodes and hours. The tensor product kernel is a valid kernel and has been used in collaborative filtering and multi-task learning [1], [2], [30], [26]. All functions in this RKHS, denoted by set $\mathcal{P}$, can be alternatively represented as [6], [2]

$$\mathcal{P}=\left\{p(n,t)=\sum_{r=1}^{R} f_r(n)g_r(t),\ f_r\in\mathcal{H}_K,\ g_r\in\mathcal{H}_G\right\} \quad (4)$$

where $\mathcal{H}_{\mathcal{K}}$ and $\mathcal{H}_{\mathcal{G}}$ are the RKHSs defined respectively by K and G, while the number of summands R is possibly infinite. Note that the decomposition in (4) is not unique [6]. Similar to (2) and upon arranging observed prices in $\mathbf{Z}\in\mathbb{R}^{N\times T}$, the market function $p(n,t)$ could be inferred via

$$\min_{p\in\mathcal{P}} \|\mathbf{Z}-\mathbf{P}\|_F^2+\mu\|p\|_{\mathcal{K}_{\otimes}} \quad (5)$$

where $\mathbf{P}\in\mathbb{R}^{N\times T}$ has entries $[\mathbf{P}]_{n,t}=p(n,t)$, $\|p\|_{\mathcal{K}_{\otimes}}$ is the norm in $\mathcal{P}$ [cf. Equation (1)], and $\mu>0$ is a regularization parameter. Notice the notational convention that when n and t are used as function arguments, the function depends on $\mathbf{x}_n$ and $\mathbf{y}_t$, respectively. In other words, $p(n,t)=p(\mathbf{x}_n,\mathbf{y}_t)$, $K(n,n')=K(\mathbf{x}_n,\mathbf{x}_{n'})$, and $G(t,t')=G(\mathbf{y}_t,\mathbf{y}_{t'})$.

The key presumption here is that $p(n,t)$ is practically the superposition of relatively few components $p_r(n,t):=f_r(n)g_r(t)$: At a specific t, usually only a few transmission lines are congested, i.e., have reached their rated power capacity [17], [15].[1] Each p_r corresponds to the pricing pattern observed whenever a specific congestion scenario occurs. Yet spatial effects are modulated by time. For example, congestion typically occurs during peak demand or high-wind periods. Moreover, due to generator ramp constraints, demand periodicities, and lasting transmission outages; pricing motifs tend to iterate over time instances with similar characteristics, e.g., the same hour of the next day or week. These specifications not only justify using the tensor product kernel $K_{\otimes}$, but they also hint at a relatively small R in (4).

To facilitate parsimonious modeling of $p(n,t)$ using a few $p_r(n,t)$ components, instead of regularizing by $\|p\|_{K_{\otimes}}$ [cf. Equation (5)], the *trace norm* $\|p\|_*$ could be used:

$$\min_{p\in\mathcal{P}} \|\mathbf{Z}-\mathbf{P}\|_F^2+\lambda\|p\|_* \quad (6)$$

for some $\lambda>0$. For the definition of trace norm see [1]. In [1], it is also shown that for every function $p\in\mathcal{P}$, its $\|p\|_*$ can be alternatively expressed as

$$\|p\|_*=\min_{\{f_r,g_r\}}\frac{1}{2}\left(\sum_{r=1}^{R}\|f_r\|_{\mathcal{K}}^2+\sum_{r=1}^{R}\|g_r\|_{\mathcal{G}}^2\right) \quad (7)$$
$$\text{s.to}\quad p=\sum_{r=1}^{R} f_r g_r,\ f_r\in\mathcal{H}_K,\ g_r\in\mathcal{H}_G.$$

Regularizing by $\|p\|_*$ is known to favor low-rank models [2], [33]. Nevertheless, in this work we advocate regularizing by the square root of $\|p\|_*$ to critically enable kernel selection (cf. Section II-C) and to derive efficient algorithms (cf. Section IV). Specifically, market inference is posed here as the regularization problem:

$$\min_{p\in\mathcal{P}} \|\mathbf{Z}-\mathbf{P}\|_F^2+\mu\sqrt{\|p\|_*} \quad (8)$$

for some $\mu>0$. The connection between (6) and (8) can be understood by the next proposition proved in Appendix A.

Proposition 1: If p_μ^* denotes a function minimizing (8) for some $\mu>0$, there exists $\lambda_\mu>0$, such that p_μ^* is also a minimizer of (6) for $\lambda=\lambda_\mu$.

Albeit Proposition 1 does not provide an analytic expression for λ_μ, it asserts that every minimizer of (8) is a minimizer for (6) too for an appropriate λ. Thus, the functions minimizing (8) are expected to be decomposable into a few p_r. Numerical tests indicate that (8) favors low-rank minimizers indeed.

Given that (8) admits low-rank minimizers anyway, its feasible set could be possibly restricted to a $\mathcal{P}$ defined by (4) but for a finite and relatively small R_0. If the p minimizing (8) over this restricted feasible set turns out to be of rank smaller than R_0, the restriction comes at no loss of optimality. Throughout the rest of the paper, (8) will be solved for a finite R. Similar approaches have been developed for low-rank matrix completion [7], collaborative filtering [2], and multi-task learning [30], [26].

To leverage the low-rank model in solving (8), the following result, proved in Appendix B, is needed:

Lemma 1: For every $p\in\mathcal{P}$, it holds $\sqrt{\|p\|_*}=h(p)$, where

$$h(p):=\min_{\{f_r,g_r\}}\frac{1}{2}\left[\left(\sum_{r=1}^{R}\|f_r\|_{\mathcal{K}}^2\right)^{1/2}+\left(\sum_{r=1}^{R}\|g_r\|_{\mathcal{G}}^2\right)^{1/2}\right] \quad (9)$$
$$\text{s.to}\ p=\sum_{r=1}^{R} f_r g_r,\ f_r\in\mathcal{H}_K,\ g_r\in\mathcal{H}_G.$$

Due to Lemma 1, the problem in (8) is reformulated, and p can be learned via the regularization

$$Q(\mathcal{K},\mathcal{G}):=\min_{p\in\mathcal{P}} Q(\mathcal{K},\mathcal{G},p) \quad (10a)$$

where

$$Q(\mathcal{K},\mathcal{G},p):=\|\mathbf{Z}-\mathbf{P}\|_F^2 \quad (10b)$$
$$+\mu\left(\sum_{r=1}^{R}\|f_r\|_{\mathcal{K}}^2\right)^{1/2}+\mu\left(\sum_{r=1}^{R}\|g_r\|_{\mathcal{G}}^2\right)^{1/2}.$$

C. Multi-Kernel Learning

Solving the inference problem in (10) assumes that μ and the kernels $\mathcal{K}$ and $\mathcal{G}$ are known. The parameter μ is typically tuned via cross-validation [20]. Choosing the appropriate kernels though is more challenging, as testified by the extensive research on *multi-kernel learning*; see the reviews [18], [3].

In this work, the multi-kernel learning approach of [30] is generalized to the function regularization in (10). Specifically, two sets of kernel function choices, $\{K_l\}_{l=1}^{L}$ and $\{G_m\}_{m=1}^{M}$, are provided for nodes and time periods, respectively. Numbers L and M are selected depending on the kernel choices and the

[1]This fact is exploited in [22] to reveal the topology of the underlying power grid by using only publicly available real-time LMPs.

computational resources available. Consider the kernel spaces constructed as the convex hulls

$$\mathcal{K} := \{K = \sum_{l=1}^{L} \theta_l K_l,\ \theta_l > 0,\ \sum_{l=1}^{L} \theta_l = 1\} \tag{11a}$$

$$\mathcal{G} := \{G = \sum_{m=1}^{M} \phi_m G_m,\ \phi_m > 0,\ \sum_{m=1}^{M} \phi_m = 1\}. \tag{11b}$$

Optimizing the outcome of the regularization problem in (10a) over $\mathcal{K}$ and $\mathcal{G}$ provides a disciplined kernel design methodology. Since all K_l and G_m are predefined, minimizing (10a) over $\mathcal{K}$ and $\mathcal{G}$, reduces to minimizing $Q(\mathcal{K}, \mathcal{G})$ over the weights $\{\theta_l\}$ and $\{\phi_m\}$. The following theorem, which is proved in Appendix C, shows how the kernel learning part can be accomplished without even finding the optimal weights.

Theorem 1: Consider the function space $\mathcal{P}$, the kernel spaces $\mathcal{K}$ and $\mathcal{G}$, and the functional $Q(\mathcal{K}, \mathcal{G}, p)$, defined in (4), (11), and (10b), respectively. Solving the regularization problem

$$\min_{\mathcal{K},\mathcal{G}} \min_{p\in\mathcal{P}}\ Q(\mathcal{K}, \mathcal{G}, p) \tag{12}$$

is equivalent to solving

$$\min_{p\in\mathcal{P}'} \|\mathbf{Z} - \mathbf{P}\|_F^2 + \mu \sum_{l=1}^{L} \sqrt{\sum_{r=1}^{R} \|f_{lr}\|_{\mathcal{K}_l}^2} + \mu \sum_{m=1}^{M} \sqrt{\sum_{r=1}^{R} \|g_{mr}\|_{\mathcal{G}_m}^2} \tag{13}$$

over $\mathcal{P}' := \Big\{p(n,t) = \sum_{r=1}^{R} f_r(n) g_r(t) : f_r = \sum_{l=1}^{L} f_{lr},\ f_{lr} \in \mathcal{H}_{\mathcal{K}_l},\ g_r = \sum_{m=1}^{M} g_{mr},\ g_{mr} \in \mathcal{H}_{\mathcal{G}_m}\Big\}$, where $\{\mathcal{H}_{\mathcal{K}_l}\}$ and $\{\mathcal{H}_{\mathcal{G}_m}\}$ are the function spaces defined by the kernels K_l and G_m, accordingly.

Theorem 1 asserts that minimizing (10b) over $f_r \in \mathcal{H}_{\mathcal{K}}$ and $g_r \in \mathcal{H}_{\mathcal{G}}$ boils down to the functional optimization in (13) where f_r and g_r are now simply decomposed as $\sum_{l=1}^{L} f_{lr}$ and $\sum_{m=1}^{M} g_{mr}$, respectively. Interestingly enough, the theorem also generalizes the multi-kernel learning results of [30] to the low-rank decomposition model of (4). After drawing some interesting connections in Section II-D, the functional inference in (13) is transformed to a matrix minimization problem in Section III.

D. Interesting Connections

Observe that when $\mathcal{N}$ and $\mathcal{T}$ are Euclidean spaces, $K(n,n') = \delta(n-n')$ and $G(t,t') = \delta(t-t')$ where $\delta(\cdot)$ is the Kronecker delta function, then $\mathcal{P}$ in (4) is the space of matrices $\mathbf{P} \in \mathbb{R}^{|\mathcal{N}|\times|\mathcal{T}|}$ having $p(n,t)$ as their (n,t)-th entry. In this case, $\|p\|_*$ is simply the nuclear norm $\|\mathbf{P}\|_*$ of matrix $\mathbf{P}$, i.e., the sum of its singular values; $R = \text{rank}(\mathbf{P})$; and (7) becomes [2], [7],

$$\|\mathbf{P}\|_* = \min_{\mathbf{F},\mathbf{G}}\ \frac{1}{2}\left(\|\mathbf{F}\|_F^2 + \|\mathbf{G}\|_F^2\right) \quad \text{s.to}\ \ \mathbf{P} = \mathbf{F}\mathbf{G}^T, \mathbf{F} \in \mathbb{R}^{|\mathcal{N}|\times R},\ \mathbf{G} \in \mathbb{R}^{|\mathcal{T}|\times R}. \tag{14}$$

The alternative representation of $\|\mathbf{P}\|_*$ in (14) has been extensively used in nuclear norm minimization [37], [33], [29]. Interestingly, the matrix counterpart of Lemma 1 reads:

Corollary 1: For $\mathbf{P} \in \mathbb{R}^{N\times T}$ with $\text{rank}(\mathbf{P}) = R$, it holds

$$\|\mathbf{P}\|_*^{1/2} = \min_{\mathbf{F},\mathbf{G}}\ \frac{1}{2}\left(\|\mathbf{F}\|_F + \|\mathbf{G}\|_F\right) \quad \text{s.to}\ \ \mathbf{P} = \mathbf{F}\mathbf{G}^T,\ \mathbf{F} \in \mathbb{R}^{N\times R},\ \mathbf{G} \in \mathbb{R}^{T\times R}. \tag{15}$$

Matrix completion aims at recovering a low-rank matrix $\mathbf{P}$ given noisy measurements for a few of its entries [13]. It can be derived from (6) after replacing $\|p\|_*$ by $\|\mathbf{P}\|_*$ [or (14)], and $\|\mathbf{Z} - \mathbf{P}\|_F^2$ by $\|(\mathbf{Z} - \mathbf{P}) \odot \mathbf{\Delta}\|_F^2$, where $\odot$ denotes element-wise multiplication and $\mathbf{\Delta}$ is a binary matrix having zeros on the missing entries. The premise is that $\mathbf{P}$ could be recovered due to its low-rank property. But recovery is impossible when entire columns or rows are missing.

For generic *yet fixed* kernels $K(n,n')$ and $G(t,t')$, low-rank kernel-based models could be similarly derived as special cases of (6); see e.g., [2], [7]. Using kernel functions other than the Kronecker delta, enables not only recovering the missing entries, but also extrapolating to unseen columns and rows. Different from matrix completion and low-rank kernel-based inference, our regularization in (13) targets to jointly learn a low-rank $p(n,t)$, together with kernels K and G.

III. Matrix Optimization

The next goal is to map the functional optimization of (13) to a vector minimization by resorting to the Representer's Theorem [20]. Observe that minimizing (13) over a specific f_{lr} is actually a functional minimization regularized by $(\|f_{lr}\|_{\mathcal{K}_l}^2 + c_{lr})^{1/2}$ for some constant $c_{lr} \geq 0$. Since the regularization is an increasing function of $\|f_{lr}\|_{\mathcal{K}_l}^2$, Representer's Theorem applies readily [20], [5].

Each one of the LR functions f_{lr} minimizing (13) can be expressed as a linear combination of the associated kernel K_l evaluated over the N training examples involved, that is

$$f_{lr}(n) = \sum_{n'=1}^{N} K_l(n,n')\beta_{lr,n'}. \tag{16}$$

Upon concatenating the unknown expansion coefficients and the function values into $\boldsymbol{\beta}_{lr} := [\beta_{lr,1}\ \cdots\ \beta_{lr,N}]^\top$ and $\mathbf{f}_{lr} := [f_{lr}(1)\ \cdots\ f_{lr}(N)]^\top$, respectively, it holds that

$$\mathbf{f}_{lr} = \mathbf{K}_l \boldsymbol{\beta}_{lr} \tag{17}$$

where $\mathbf{K}_l \in \mathbb{S}_{++}^N$ is the node kernel matrix whose (n,n')-th entry is $K_l(n,n')$. Using (17) and accounting for the decomposition $f_r = \sum_{l=1}^{L} f_{lr}$ dictated by (13), the vector collecting the values $\{f_r(n)\}_{n=1}^N$ is compactly written as

$$\mathbf{f}_r = \sum_{l=1}^{L} \mathbf{K}_l \boldsymbol{\beta}_{lr}. \tag{18}$$

Likewise, each g_{mr} minimizing (13) admits the expansion

$$g_{mr}(t) = \sum_{t'=1}^{T} G_m(t,t')\gamma_{mr,t'} \tag{19}$$

for all t. Similar to (17), the vector of function values $\mathbf{g}_{mr} := [g_{mr}(1) \ \ldots \ g_{mr}(T)]^\top$ is expressed in terms of the time kernel matrix $\mathbf{G}_m \in \mathbb{S}_{++}^T$ as

$$\mathbf{g}_{mr} = \mathbf{G}_m \boldsymbol{\gamma}_{mr} \tag{20}$$

where $\boldsymbol{\gamma}_{mr} := [\gamma_{mr,1} \ \ldots \ \gamma_{mr,T}]^\top$. Due to the decomposition $g_r = \sum_{m=1}^M g_{mr}$ in (13), the vector containing $\{g_r(t)\}_{t=1}^T$ is provided by [cf. Equation (18)]

$$\mathbf{g}_r = \sum_{m=1}^M \mathbf{G}_m \boldsymbol{\gamma}_{mr}. \tag{21}$$

So far, the functions $\{f_r(n), g_r(t)\}_{r=1}^R$ minimizing (13) have been expressed in terms of $\boldsymbol{\beta}_{lr}$'s and $\boldsymbol{\gamma}_{mr}$'s, thus enabling one to transform (13) to a minimization problem over the unknown coefficients.

Regarding the price matrix $\mathbf{P}$, the low-rank model $p(n,t) = \sum_{r=1}^R f_r(n) g_r(t)$ implies that

$$\mathbf{P} = \sum_{r=1}^R \mathbf{f}_r \mathbf{g}_r^\top. \tag{22}$$

Plugging (18) and (21) into (22), yields

$$\mathbf{P} = \sum_{l=1}^L \sum_{m=1}^M \mathbf{K}_l \mathbf{B}_l \boldsymbol{\Gamma}_m^\top \mathbf{G}_m \tag{23}$$

where $\mathbf{B}_l := [\boldsymbol{\beta}_{l1} \ \cdots \boldsymbol{\beta}_{lR}] \in \mathbb{R}^{N \times R}$ and $\boldsymbol{\Gamma}_m := [\boldsymbol{\gamma}_{m1} \ \cdots \boldsymbol{\gamma}_{mR}] \in \mathbb{R}^{T \times R}$ for all l and m.

Consider now the regularization terms in (13). Due to (16) and (19), the associated norms can be written as $\|f_{lr}\|_{\mathcal{K}_l}^2 = \boldsymbol{\beta}_{lr}^\top \mathbf{K}_l \boldsymbol{\beta}_{lr}$ and $\|g_{mr}\|_{\mathcal{G}_m}^2 = \boldsymbol{\gamma}_{mr}^\top \mathbf{G}_m \boldsymbol{\gamma}_{mr}$ [cf. Equation (1)–(5)]. Using the properties of the trace operator, it can be shown that

$$\sum_{r=1}^R \|f_{lr}\|_{\mathcal{K}_l}^2 = \mathrm{Tr}(\mathbf{B}_l^\top \mathbf{K}_l \mathbf{B}_l) \tag{24a}$$

$$\sum_{r=1}^R \|g_{mr}\|_{\mathcal{G}_m}^2 = \mathrm{Tr}(\boldsymbol{\Gamma}_m^\top \mathbf{G}_m \boldsymbol{\Gamma}_m). \tag{24b}$$

The right-hand sides in (24) can be identified as the norms $\|\mathbf{B}_l\|_{\mathbf{K}_l}^2 := \mathrm{Tr}(\mathbf{B}_l^\top \mathbf{K}_l \mathbf{B}_l)$ and $\|\boldsymbol{\Gamma}_m\|_{\mathbf{G}_m}^2 := \mathrm{Tr}(\boldsymbol{\Gamma}_m^\top \mathbf{G}_m \boldsymbol{\Gamma}_m)$. By using (23)–(24), the functional optimization in (13) can be compactly expressed as the matrix optimization problem

$$\min_{\mathbf{P}, \{\mathbf{B}_l\}, \{\boldsymbol{\Gamma}_m\}} \|\mathbf{Z} - \mathbf{P}\|_F^2 + \mu \sum_{l=1}^L \|\mathbf{B}_l\|_{\mathbf{K}_l} + \mu \sum_{m=1}^M \|\boldsymbol{\Gamma}_m\|_{\mathbf{G}_m} \tag{25}$$

$$\text{s.to} \quad \mathbf{P} = \sum_{l=1}^L \sum_{m=1}^M \mathbf{K}_l \mathbf{B}_l \boldsymbol{\Gamma}_m^\top \mathbf{G}_m.$$

Solving (25) faces two challenges. Even though optimizing separately over $\{\mathbf{B}_l\}$ or $\{\boldsymbol{\Gamma}_m\}$ entails a convex cost, the joint minimization is non-convex. Secondly, solving (25) involves multiple high-dimensional matrices, which raises computational concerns. The algorithm developed in the next section scales well with the problem dimensions, and converges to a stationary point of (25).

Price Forecasting: Having found all $\hat{\mathbf{B}}_l$ and $\hat{\boldsymbol{\Gamma}}_m$, the electricity prices over the training period can be reconstructed via (22). Of course, the ultimate learning goal is inferring future prices. Based on the modeling approach in Section II-B, the price $p(n_0, t_0)$ for an unseen pair (n_0, t_0) can be predicted simply as

$$\hat{p}(n_0, t_0) = \sum_{r=1}^R \sum_{l=1}^L \sum_{m=1}^M \hat{f}_{lr}(n_0) \hat{g}_{mr}(t_0) \tag{26}$$

where $\hat{f}_{lr}(n_0) = \sum_{n=1}^N K_l(n_0, n) \hat{\beta}_{lr,n}$ and $\hat{g}_{mr}(t_0) = \sum_{t=1}^T G_m(t_0, t) \hat{\gamma}_{mr,t}$ [cf. Equation (16), (19)]. In essence, extrapolation to (n_0, t_0) is viable conditioned on availability of the kernel values involved.

If network-wide forecasts are needed over a future interval $\mathcal{T}'$ and over the node set $\mathcal{N}'$, the predicted values can be stored in the $|\mathcal{N}'| \times |\mathcal{T}'|$ matrix $\hat{\mathbf{P}}'$. According to (26), matrix $\hat{\mathbf{P}}'$ is compactly expressed as

$$\hat{\mathbf{P}}' = \sum_{m=1}^M \sum_{l=1}^L \mathbf{K}_l' \hat{\mathbf{B}}_l \hat{\boldsymbol{\Gamma}}_m^\top \mathbf{G}_m' \tag{27}$$

where $\mathbf{K}_l' \in \mathbb{R}^{N \times |\mathcal{N}'|}$ and $\mathbf{G}_m' \in \mathbb{R}^{T \times |\mathcal{T}'|}$ are the kernel matrices between the training and the forecast points, i.e., having entries $[\mathbf{K}_l']_{n,n'} = K_l(n, n')$ and $[\mathbf{G}_m']_{t,t'} = G_m(t, t')$. Important remarks are now in order.

Remark 1: Price forecasts are not confined to future t_0's (essentially unseen feature vectors $\mathbf{x}_{t_0}$); they can be issued even for a new node $n_0 \notin \mathcal{N}$. This is an important feature when dealing with electricity markets having seasonal pricing models. For example, MISO updates its commercial grid quarterly by adding, removing, merging, and redefining CPNs, to accommodate transmission grid updates and market participants leaving or entering the market.

Remark 2: In addition to extrapolation (prediction), the proposed approach is general enough to encompass imputation of missing entries. Similar to matrix completion [cf. Section II-D], that would be possible upon substituting $\|\mathbf{Z} - \mathbf{P}\|_F^2$ in (25) by $\|(\mathbf{Z} - \mathbf{P}) \odot \boldsymbol{\Delta}\|_F^2$.

Remark 3: As justified in Section IV, (25) promotes *block-sparse solutions*. In particular, some of the $\{\hat{\mathbf{B}}_l\}_{l=1}^L$ and $\{\hat{\boldsymbol{\Gamma}}_m\}_{m=1}^M$ may be driven to zero. The latter indicates that the corresponding K_l or G_m are not influential in price clearing. Since experimentation with kernels defined over different feature subsets can be highly interpretative, the proposed approach becomes a systematic prediction and kernel selection tool.

IV. Block-Coordinate Descent Algorithm

A block-coordinate descent (BCD) algorithm is developed here for solving (25). According to the BCD methodology, the initial optimization variable is partitioned into blocks. Per block minimizations having the remaining variables fixed are then iterated cyclically over blocks.

Solving (25) in particular, variable blocks are selected in the order $\{\mathbf{B}_1, \ldots, \mathbf{B}_L, \boldsymbol{\Gamma}_1, \ldots, \boldsymbol{\Gamma}_M\}$. The per block minimizations

Algorithm 1 Minimizing the canonical form **(30)**

function SOLVECANONICAL $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mu)$
  if $\|\mathbf{B}^{1/2}\mathbf{A}\mathbf{C}\|_F \leq \mu/2$ **then** $\hat{\mathbf{X}} = \mathbf{0}$
  else
    $(\mathbf{U}_B, \{\lambda_i\}) =$ EIGENDECOMPOSITION $(\mathbf{B})$
    $(\mathbf{U}_C, \{\mu_j\}) =$ EIGENDECOMPOSITION $(\mathbf{C}\mathbf{C}^{\mathrm{T}})$
    Define $\mathbf{W} = \mathbf{U}_B^\top \mathbf{A} \mathbf{U}_C$
    Initialize $w^0 = 0$ and $t = 0$
    repeat
      Evaluate $s'(w^t)$ via **(33)**
      Update $w^{t+1} = \max\{0, w^t - c \cdot s'(w^t)\}$
      $t = t + 1$
    until $\left|s(w^t) - s(w^{t-1})\right| < \epsilon_{\mathrm{c}}$
    Set $\hat{w} = w^t$
    Obtain $\hat{\mathbf{X}}$ by solving the Sylvester equation **(31)**
  end if
end function

involved are detailed next. Consider minimizing (25) over a specific $\mathbf{B}_l$, while all other variables are maintained to their most recent values $\{\hat{\mathbf{B}}_{l'}\}_{l' \neq l}$ and $\{\hat{\mathbf{\Gamma}}_m\}_{m=1}^M$. Upon rearranging terms in (25), block $\mathbf{B}_l$ can be updated as

$$\hat{\mathbf{B}}_l = \arg\min_{\mathbf{B}_l} \|\mathbf{Z}_l^B - \mathbf{K}_l \mathbf{B}_l \mathbf{H}^\top\|_F^2 + \mu \|\mathbf{B}_l\|_{\mathbf{K}_l} \tag{28}$$

where $\mathbf{H} := \sum_{m=1}^M \mathbf{G}_m \hat{\mathbf{\Gamma}}_m$ is the contribution of all $\hat{\mathbf{\Gamma}}_m$, and $\mathbf{Z}_l^B := \mathbf{Z} - \sum_{l' \neq l} \mathbf{K}_{l'} \hat{\mathbf{B}}_{l'} \mathbf{H}^\top$.

Similarly, updating a particular $\mathbf{\Gamma}_m$ entails finding

$$\hat{\mathbf{\Gamma}}_m = \arg\min_{\mathbf{\Gamma}_m} \|\mathbf{Z}_m^\Gamma - \mathbf{F}\mathbf{\Gamma}_m^\top \mathbf{G}_m\|_F^2 + \mu \|\mathbf{\Gamma}_m\|_{\mathbf{G}_m} \tag{29}$$

where $\mathbf{F} := \sum_{l=1}^L \mathbf{K}_l \hat{\mathbf{B}}_l$ is the contribution of all $\hat{\mathbf{B}}_l$, and $\mathbf{Z}_m^\Gamma := \mathbf{Z} - \sum_{m' \neq m} \mathbf{F}\mathbf{\Gamma}_{m'}^\top \mathbf{G}_{m'}$.

Problems (28) and (29) are convex, yet not differentiable, and exhibit the same canonical form. This form can be efficiently solved according to the following lemma that is proved in Appendix D.

Lemma 2: Let $\mathbf{A} \in \mathbb{R}^{d_1 \times d_3}$, $\mathbf{B} \in \mathbb{S}_{++}^{d_1}$, $\mathbf{C} \in \mathbb{R}^{d_3 \times d_2}$, and $\mu > 0$. The convex optimization problem

$$\min_{\mathbf{X}} \ \|\mathbf{A} - \mathbf{B}\mathbf{X}\mathbf{C}^\top\|_F^2 + \mu \|\mathbf{X}\|_{\mathbf{B}} \tag{30}$$

has a unique minimizer $\hat{\mathbf{X}}$ provided by the solution of

$$\mathbf{B}\hat{\mathbf{X}}\mathbf{C}^\top\mathbf{C} + \frac{\mu^2}{4\hat{w}}\hat{\mathbf{X}} = \mathbf{A}\mathbf{C} \tag{31}$$

if $\|\mathbf{B}^{1/2}\mathbf{A}\mathbf{C}\|_F > \mu/2$; or, $\hat{\mathbf{X}} = \mathbf{0}$, otherwise. The scalar $\hat{w} > 0$ in (31) is the minimizer of the convex problem

$$\hat{w} := \arg\min_{w \geq 0} \ w - \sum_{i=1}^{d_1} \sum_{j=1}^{d_2} \frac{[\mathbf{W}]_{ij}^2 \lambda_i \mu_j w}{\lambda_i \mu_j w + \frac{\mu^2}{4}} \tag{32}$$

Algorithm 2 BCD algorithm for solving **(25)**

Input: $\mathbf{Z}$, $\{\mathbf{K}_l\}_{l=1}^L$, $\{\mathbf{G}_m\}_{m=1}^M$, R, μ
Randomly initialize $\{\hat{\mathbf{B}}_l\}_{l=1}^L$ and $\{\hat{\mathbf{\Gamma}}_m\}_{m=1}^M$
Compute $\mathbf{F} = \sum_{l=1}^L \mathbf{K}_l \hat{\mathbf{B}}_l$ and $\mathbf{H} = \sum_{m=1}^M \mathbf{G}_m \hat{\mathbf{\Gamma}}_m$
Store $\{\hat{\mathbf{B}}_l^{\mathrm{old}} = \hat{\mathbf{B}}_l\}_{l=1}^L$ and $\{\hat{\mathbf{\Gamma}}_m^{\mathrm{old}} = \hat{\mathbf{\Gamma}}_m\}_{m=1}^M$
repeat
  for $l = 1 \to L$ **do**
    Update $\mathbf{F} = \mathbf{F} - \mathbf{K}_l \hat{\mathbf{B}}_l$
    Define $\mathbf{Z}_l^B = \mathbf{Z} - \mathbf{F}\mathbf{H}^\top$
    $\hat{\mathbf{B}} =$ SOLVECANONICAL $(\mathbf{Z}_l^B, \mathbf{K}_l, \mathbf{H}, \mu)$
    Update $\mathbf{F} = \mathbf{F} + \mathbf{K}_l \hat{\mathbf{B}}_l$
  end for
  for $m = 1 \to M$ **do**
    Update $\mathbf{H} = \mathbf{H} - \mathbf{G}_m \hat{\mathbf{\Gamma}}_m$
    Define $\mathbf{Z}_m^\Gamma = \mathbf{Z} - \mathbf{F}\mathbf{H}^\top$
    $\hat{\mathbf{\Gamma}}_m =$ SOLVECANONICAL $((\mathbf{Z}_m^\Gamma)^\top, \mathbf{G}_m, \mathbf{F}, \mu)$
    Update $\mathbf{H} = \mathbf{H} + \mathbf{G}_m \hat{\mathbf{\Gamma}}_m$
  end for
until $\left|\frac{(f(\{\hat{\mathbf{B}}_l\}, \{\hat{\mathbf{\Gamma}}_m\})}{f(\{\hat{\mathbf{B}}_l^{\mathrm{old}}\}, \{\hat{\mathbf{\Gamma}}_m^{\mathrm{old}}\}))} - 1\right| < \epsilon_{\mathrm{BCD}}$: $f(\cdot)$ is the cost in **(25)**
Output: $\{\hat{\mathbf{B}}_l\}_{l=1}^L$, $\{\hat{\mathbf{\Gamma}}_m\}_{m=1}^M$

where $\mathbf{W} := \mathbf{U}_B^\top \mathbf{A} \mathbf{U}_C$; $(\mathbf{U}_B, \{\lambda_i\}_{i=1}^{d_1})$ are the eigenpairs of $\mathbf{B}$; and $(\mathbf{U}_C, \{\mu_j\}_{j=1}^{d_2})$ the non-zero eigenpairs of $\mathbf{C}\mathbf{C}^\top$.

Lemma 2 provides valuable insights for solving (30). It reveals that by simply calculating $\|\mathbf{B}^{1/2}\mathbf{A}\mathbf{C}\|_F$, the sought $\hat{\mathbf{X}}$ may be directly set to zero. Hence, (30) admits block-zero minimizers depending on the value of μ. This property critically implies that some of the $\{\hat{\mathbf{B}}_l\}$ and $\{\hat{\mathbf{\Gamma}}_m\}$ minimizing (25) will be zero, thus, effecting kernel selection.

Back to Lemma 2, if $\|\mathbf{B}^{1/2}\mathbf{A}\mathbf{C}\|_F > \mu/2$, a non-zero solution emerges. The univariate optimization in (32) and the linear matrix equations in (31) can be efficiently tackled as described next. First, the constrained convex problem in (32) can be solved by a projected gradient algorithm. If $s(w)$ denotes the cost function in (32), its derivative is

$$s'(w) = 1 - \sum_{i=1}^{d_1} \sum_{j=1}^{d_2} \frac{\mu^2 [\mathbf{W}]_{ij}^2 \lambda_i \mu_j}{4\left(\lambda_i \mu_j w + \frac{\mu^2}{4}\right)^2}. \tag{33}$$

The iterates $w^{t+1} = \max\{0, w^t - c \cdot s'(w^t)\}$ are guaranteed to converge to the global minimum $\hat{w}$ for a sufficiently small step size $c > 0$; see [8] for details. Each iterate costs $\mathcal{O}(d_1 d_2)$ operations. Secondly, concerning (31), it can be rewritten as a Sylvester equation as advocated also in [23], [35]. Hence, $\hat{\mathbf{X}}$ can be found in $\mathcal{O}(d_1^3 + d_2^3)$ numerical operations using the Bartels-Stewart algorithm [16, Alg. 7.6.2], instead of the $\mathcal{O}(d_1^3 d_2^3)$ complexity of a generic linear system solver. The steps for solving the canonical problem (30) are tabulated as Alg. 1, whose overall worst-case complexity is $\mathcal{O}(d_1^3 + d_2 d_3^2)$.

Proceeding with the BCD steps (28) and (29), those can be efficiently performed after carefully updating $\mathbf{H}$ and $\mathbf{F}$. The

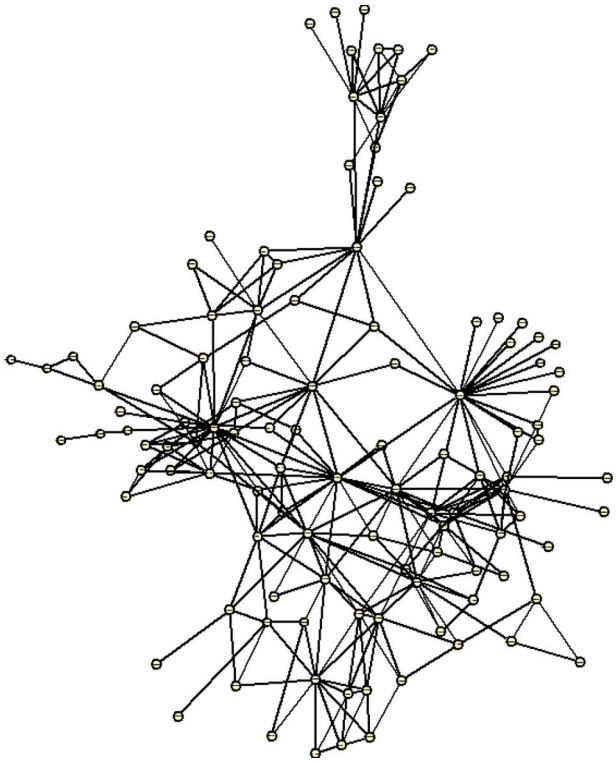

Fig. 1. Graph of the LBAs involved in the MISO market.

final steps for solving (25) are listed as Alg. 2. Due to the separability of the non-differentiable cost over the chosen variable blocks, the BCD algorithm is guaranteed to converge to a stationary point of (25) [38]. The BCD iterates are terminated when the relative cost value error becomes smaller than some threshold e.g., $\epsilon_{\text{BCD}} = 10^{-3}$. The eigendecomposition of all kernel matrices can be computed once. Algorithm 2 has complexity $\mathcal{O}\left(L(N^3+RT^2)+M(T^3+RN^2)\right)$ per iteration. In the numerical experiments of Section V, and depending on the value of μ, 5–15 BCD iterations were sufficient.

V. Numerical Tests

The derived low-rank multi-kernel learning approach was tested using real data from the Midwest ISO (MISO) electricity market. Day-ahead hourly LMPs were collected across N =1,732 nodes for the period June 1 to August 31, 2012, yielding a total of 92 days or 2,208 hours.

A pool of $L = 5$ nodal and $M = 5$ time kernels was selected as detailed next. Starting with the nodal ones, when learning over a graph, the corresponding graph Laplacian matrix is oftentimes used to design meaningful kernels [25]. CPNs are considered here as vertices of a similarity graph, connected with edges having non-negative weights proportional to the similarity between incident CPNs. Nonetheless, lacking any other type of geographical or electrical distance, the local balancing authority (LBA) each CPN belongs to was adopted here as a topology surrogate. The presumption is that nodes of the same LBA experience similar prices. Further, nodes controlled by neighboring authorities are expected to have prices correlated more than nodes under non-adjacent ones. The connectivity graph of 131 LBAs involved in MISO was constructed based on publicly available data found on MISO's website; cf. Fig. 1.

Kernel matrices $\mathbf{K}_1, \mathbf{K}_2 \in \mathbb{S}_{++}^N$ were built based on this LBA connectivity graph as follows. Edges between CPNs of the same LBA were assigned unit weights; edges across CPNs from different LBAs received weight 0.5; and all other edges were set to zero. If weight values are stored in the adjacency matrix $\mathbf{A}_{\mathcal{N}}$, the normalized Laplacian matrix of a graph is defined as $\mathbf{L}_{\mathcal{N}} := \mathbf{I}_N - \mathbf{D}_{\mathcal{N}}^{-1/2}\mathbf{A}_{\mathcal{N}}\mathbf{D}_{\mathcal{N}}^{-1/2}$, where $\mathbf{D}_{\mathcal{N}}$ is a diagonal matrix with diagonal entries the row sums of $\mathbf{A}_{\mathcal{N}}$ [25]. Then, $\mathbf{K}_1$ was selected as the regularized Laplacian $\mathbf{K}_1 := (\mathbf{L}_{\mathcal{N}} + \mathbf{I}_N)^{-1}$, and $\mathbf{K}_2$ as the diffusion Laplacian $\mathbf{K}_2 := \exp(-3\mathbf{L}_{\mathcal{N}})$ [36].

Kernel $\mathbf{K}_3$ utilized information that could be inferred from CPN names. Specifically, the prefix of every CPN name in MISO denotes its LBA, while some CPNs have similar names. For example, nodes ALTE.COLUMBAL1 and ALTE.COLUMBAL2 belong to the LBA named ALTE, and they are assumed to be geographically colocated. Every CPN is classified in the MISO market as generator, load, interface, or hub. The LBA, the name similarity, and the CPN type, were all used as binary coded categorical features. The vectors obtained were then used as arguments of a Gaussian kernel. The kernel bandwidth was fixed to the median of all pairwise squared Euclidean vector distances.

To capture potential independence across nodes, kernel $\mathbf{K}_4$ was chosen to be the identity matrix. The last nodal kernel $\mathbf{K}_5$ was selected as the covariance matrix of market prices empirically estimated using the training data.

Regarding temporal kernels $\{\mathbf{G}_m\}_{m=1}^5$, the following publicly available features were used:

1) Yesterday's day-ahead LMPs for the same hour.
2) Load forecasts for the north, south, and central regions of MISO footprint.
3) Generation capacity outage publicized by MISO.
4) MISO forecast for market-wide wind energy generation.
5) Hourly temperature and humidity in major cities across the MISO footprint (Bismarck, Des Moines, Detroit, Kansas City, Milwaukee, Minneapolis). Instead of predicted values, the actual values recorded by the National Oceanic and Atmospheric Administration (NOOA) were used.
6) Binary encoded categorical features for the hour of the day, the day of the week, and a holiday indicator.

For all but the categorical features, their one-hour delayed and one-hour advanced values were also considered. For example, the market forecast for 3 pm depended on temperature forecasts for 2 pm, 3 pm, and 4 pm. The reason was to model wind power and weather volatility, as well as time coupling across hours introduced by unit commitment as exemplified next. Having a high temperature forecast for 4 pm increases the load demand at 4 pm and 5 pm. Additionally, industrial consumers aware of the weather forecast may start their cooling systems at 3 pm or even earlier to save money and achieve space cooling by 4 pm. Secondly, weather forecasts are characterized by delay uncertainties: a 24-hour ahead weather model predicts quite accurately that high winds or a cold wave will be coming say in the afternoon, yet the exact hour is not precisely known. Third, many generation units have physical constraints: e.g., once they are started, they should remain on for at least a specific number of

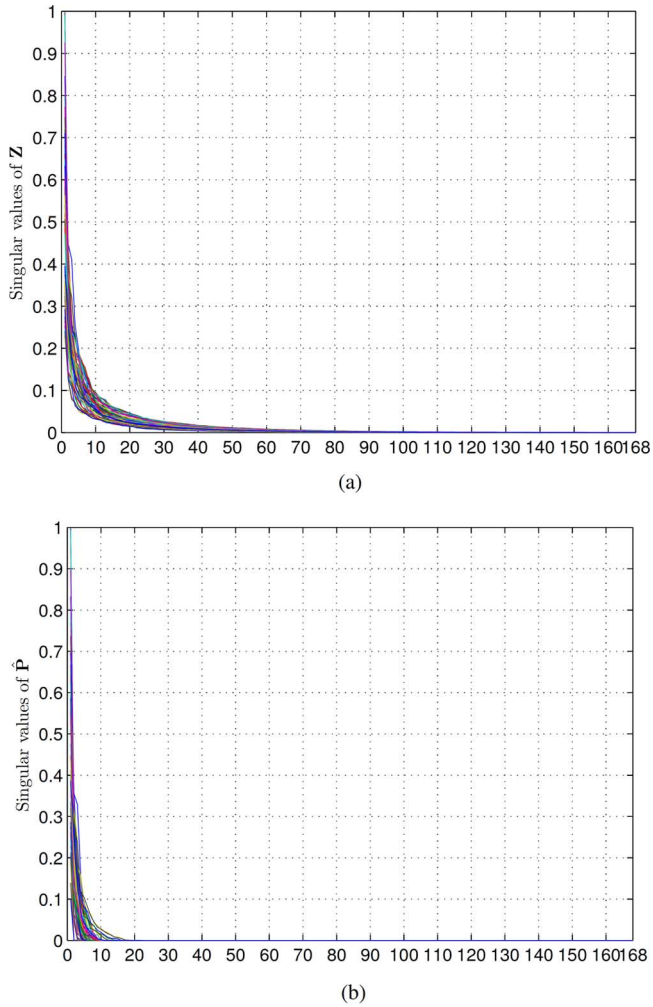

Fig. 2. Empirical distribution for the sorted singular values of price matrices: (a) for actual price matrices $\mathbf{Z} \in \mathbb{R}^{1732\times 168}$; and (b) for predicted price matrices $\hat{\mathbf{P}}$ as obtained by (25) for $R = 25$. (a) Singular values for actual price matrices $\mathbf{Z}$. (b) Singular values for predicted price matrices $\hat{\mathbf{P}}$.

hours; see e.g., [15]. Such constraints introduce time-coupling across power generation ranges and hence prices.

Temporal kernels $\mathbf{G}_1$ to $\mathbf{G}_3$ were designed by plugging the aforementioned features into Gaussian kernels of bandwidths 1, 430 (the median of all pairwise Euclidean feature distances), and 10^4, respectively. Kernel $\mathbf{G}_4$ was the Gaussian kernel obtained from all but the time-shifted features, and with its bandwidth set to the median of all pairwise Euclidean feature distances. Finally, $\mathbf{G}_5$ was selected as the linear kernel. As a standard preprocessing step, both nodal and temporal features were centered and standardized, while all $\mathbf{K}_l$'s and $\mathbf{G}_m$'s were normalized to unit diagonal elements.

Market data are cyclo-stationary: the market-wide price mean fluctuates hourly, yet with a period of one day. To cope with cyclo-stationarity, market prices in $\mathbf{Z}$ were centered upon subtracting the per-hour sample mean. The developed predictor will hence forecast the mean-compensated prices, and not the actual ones. It is important to mention though that usually the *price differences* across CPNs, rather than absolute nodal prices, are of interest. This is because bilateral transactions and power transfer contracts depend on exactly such nodal differentials [11]. In such cases, our price forecasts can be readily used. Otherwise, a simple market-wide price mean predictor could be easily trained. Several factors not captured by the publicly available features used here (e.g., transmission and generation outages) can severely affect the market. Due to this source of non-stationarity, the designed day-ahead predictors depend on market data only from the previous week. Hence, the dimension T of $\mathbf{Z}$ and $\mathbf{P}$ in (25) is 168 (hours).

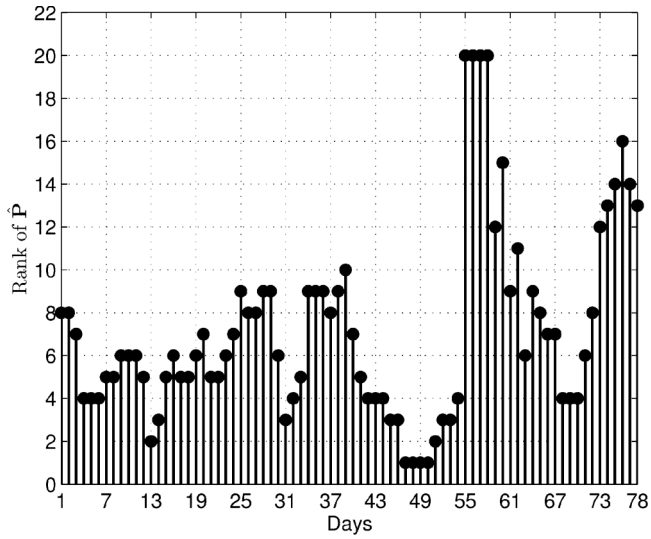

Fig. 3. Rank for predicted price matrices $\hat{\mathbf{P}}$ as obtained by (25) for $R = 25$.

Tuning the regularization parameter μ was based on market data from the first 14 days. The causal nature of the market did not allow shuffling data across time, as it is typically done in cross-validation. Instead, days 1–7 were used to predict day 8, days 2–8 for day 9, and the process was repeated up to day 14. The value of μ attaining the lowest prediction root mean square error (RMSE) over a grid of values was fixed when predicting all the remaining 78 evaluation days.

Fig. 2(a) depicts the singular values of 78 price matrices $\mathbf{Z} \in \mathbb{R}^{1732\times 168}$. The figure shows that singular values decay quickly, and retaining the top 25 could possibly express most of the information in market data. Such an observation not only justifies the trace norm regularization in (8), but also hints at fixing R to 25 for a good complexity-performance tradeoff. Fig. 2(b) shows the singular values of matrices $\hat{\mathbf{P}} \in \mathbb{R}^{1732\times 168}$ as obtained by solving (25). In addition, the rank of $\hat{\mathbf{P}}$'s is depicted in Fig. 3. Interestingly, the rank is at most 20 in all 78 matrices, which again justifies the prescribed choice of $R = 25$.

Fig. 4 shows the kernel selection capability of the novel multi-kernel learning approach. Checking whether the $\{\|\mathbf{B}_l\|_{\mathbf{K}_l}\}_{l=1}^{L}$ and $\{\|\mathbf{\Gamma}_m\|_{\mathbf{G}_m}\}_{m=1}^{M}$ obtained by Alg. 2 are zero or not, indicates whether the corresponding kernels $\{\mathbf{K}_l\}$ and $\{\mathbf{G}_m\}$ have been eliminated. A black (white) square in Fig. 4 indicates that the respective kernel has been selected (eliminated) while forecasting that specific day. Regarding nodal kernels, note that the identity kernel $\mathbf{K}_4 = \mathbf{I}_{1732}$ has been eliminated almost consistently; hence, providing experimental evidence that coupling price forecasting across CPNs is beneficial. On the other hand, kernel $\mathbf{K}_5$ computed as the sample nodal covariance seems to capture rich information of CPN pair similarities and is always selected. As far as time kernels are concerned, note that the

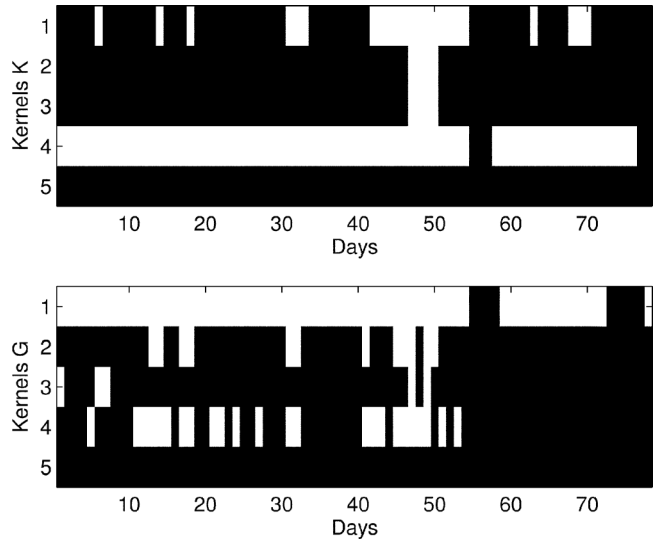

Fig. 4. Kernel selection: a black (white) square indicates that the respective kernel has been selected (eliminated) while forecasting that specific day.

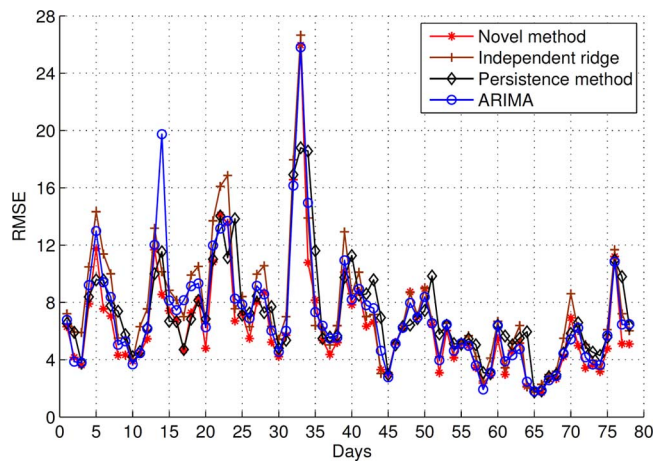

Fig. 5. RMSE comparison of forecasting methods.

bandwidth for the Gaussian kernel $\mathbf{G}_1$ turns out to be inappropriate, while the linear kernel $\mathbf{G}_5$ is always activated.

Finally, the forecasting performance of the novel method is provided in Figs. 5 and 6. Specifically, four methods were tested: (i) the novel multi-kernel learning method; (ii) the ridge regression forecast where each CPN predictor is independently obtained by solving $\min_{\mathbf{a}} \|\mathbf{z} - \mathbf{G_1 a}\|_2^2 + \mu \mathbf{a}^T \mathbf{G_1 a}$ for the Gaussian kernel $\mathbf{G}_1$; (iii) the persistence method which simply repeats yesterday's prices; and (iv) the autoregressive integrated moving average (ARIMA) approach. Regarding the last one, an ARIMA model is first estimated to fit the prices from the previous week, and it is then utilized to forecast the prices of the next 24 hours. The functions `auto.arima` and `forecast` in the R package "forecast" are used for model estimation and price forecasting, respectively, while model selection was based on the Akaike information criterion (AIC) [21], [9]. Two forecasting errors were evaluated and are listed in Table I: the root mean-square error (RMSE) $\|\hat{\mathbf{P}}' - \mathbf{P}'\|_F/\sqrt{24N}$, and the mean-absolute errors (MAEs) $\sum_{i,j} |[\hat{\mathbf{P}}']_{ij} - [\mathbf{P}']_{ij}|/(24N)$, both averaged over the 78-day evaluation period. The derived low-rank multi-kernel forecast attains the lowest RMSE and MAE.

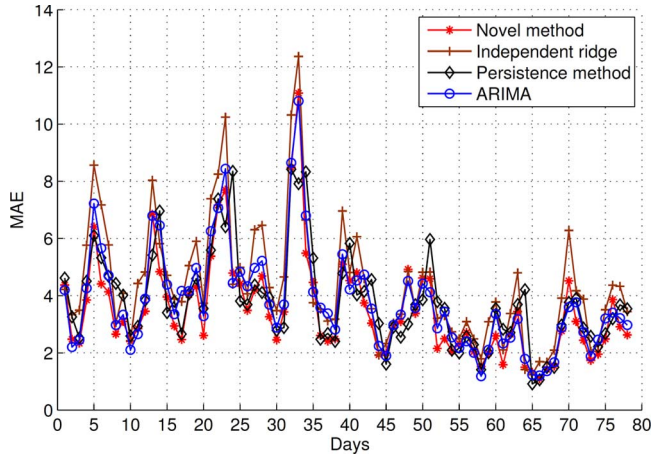

Fig. 6. MAE comparison of forecasting methods.

TABLE I
FORECASTING ERRORS [$/MWH].

	Novel	Ridge	Persistence	ARIMA
RMSE	6.395	7.550	7.197	7.062
MAE	3.514	4.395	3.810	3.798

VI. CONCLUSIONS

A novel learning approach was developed here for electricity market inference. The congestion mechanisms causing the variations in wholesale electricity prices were specifically accounted for. After viewing prices across CPNs and hours as entries of a matrix, a pertinent low-rank model was postulated. Its factors were selected from a set of candidate kernels by solving a non-convex optimization problem. Stationary points of this problem can be attained using a computationally attractive block-coordinate descent algorithm. The block-sparse properties of the per-coordinate minimizations facilitate kernel selection. Meaningful nodal kernels were built upon utilizing the related LBA connectivity graph. Applying the novel approach to MISO market data demonstrated its low-rank and kernel selection features. Even though the devised market predictor was based only on publicly available data which may not fully characterize the market outcome, it outperforms standard per-CPN predictors. The developed kernel selection methodology is sufficiently generic. It can be utilized in any low-rank collaborative filtering setup where kernels need to be selected across two types of features. Extensions to low-rank tensor scenarios where kernels are chosen over three or more feature types is an interesting research direction too. Focusing on applications for smart grids, kernel learning for low-rank models could be further used to predict load demand, as well as solar and wind energy, across nodes and time periods.

APPENDIX

A. Proof of Proposition 1

Proof of Proposition 1: The proof follows the Pareto efficient argument of [43, App. A]. Let $\mathcal{S}_\lambda$ and $\mathcal{S}_\mu$ be the sets of functions minimizing (6) and (8) for all $\lambda \geq 0$ and $\mu \geq 0$, respectively. Since (6) is a convex problem, the set $\mathcal{S}_\lambda$ coincides

with the set of *weakly efficient* functions $\mathcal{S}_p$[41]: A function p^* belongs to $\mathcal{S}_p$ if at least one of the following conditions hold:

1) $p^* \in \arg\min_{p\in\mathcal{P}} \|\mathbf{Z}-\mathbf{P}\|_F^2$;
2) $p^* \in \arg\min_{p\in\mathcal{P}} \|p\|_*$;
3) p^* is Pareto efficient, i.e., there is no $p' \in \mathcal{P}$ such that $\|\mathbf{Z}-\mathbf{P}'\|_F^2 \leq \|\mathbf{Z}-\mathbf{P}\|_F^2$ and $\|p'\|_* \leq \|p\|_*$ with at least one strict inequality.

Observe next that if p_μ^* minimizes (8) for some $\mu \geq 0$, then it is also weakly efficient. Hence, $\mathcal{S}_\mu \subseteq \mathcal{S}_p = \mathcal{S}_\lambda$, which proves the claim. ∎

B. Proof of Lemma 1

Proving Lemma 1, requires the following result.

Lemma 3: If $\{f_r^*, g_r^*\}_{r=1}^R$ are the minimizers of (9), it holds that $\sum_{r=1}^R \|f_r^*\|_{\mathcal{K}}^2 = \sum_{r=1}^R \|g_r^*\|_{\mathcal{G}}^2$.

Proof of Lemma 3: Arguing by contradiction, suppose there exist $\{f_r^0, g_r^0\}_{r=1}^R$ minimizing (9) with $\sum_{r=1}^R \|f_r^0\|_{\mathcal{K}}^2 \neq \sum_{r=1}^R \|g_r^0\|_{\mathcal{G}}^2$. Without loss of generality, assume $\sum_{r=1}^R \|f_r^0\|_{\mathcal{K}}^2 = (1+\epsilon)^2 \sum_{r=1}^R \|g_r^0\|_{\mathcal{G}}^2$ for some $\epsilon > 0$. The minimum value attained in (9) is $(2+\epsilon)\sqrt{\sum_{r=1}^R \|g_r^0\|_{\mathcal{G}}^2}/2$.

Consider now the functions $\{(1+\epsilon/2)^{-1} f_r^0\}_{r=1}^R$ and $\{(1+\epsilon/2) g_r^0\}_{r=1}^R$ which are feasible for (9), yielding a cost of $\left(\frac{1+\epsilon}{1+\epsilon/2)} + 1 + \frac{\epsilon}{2}\right)\sqrt{\sum_{r=1}^R \|g_r^0\|_{\mathcal{G}}^2}/2$. The fact that $\frac{1+\epsilon}{1+\epsilon/2} + 1 + \frac{\epsilon}{2} < 2+\epsilon$ for all $\epsilon > 0$ contradicts the assumed optimality of $\{f_r^0, g_r^0\}$. ∎

Proof of Lemma 1: Every $p \in \mathcal{P}$ admits a spectral factorization $p(n,t) = \sum_{r=1}^\infty \sigma_r u_r(n) v_r(t)$, where $\{\sigma_r\}$ is a non-negative sequence converging to zero, and $\{u_r(n)\}$ and $\{v_r(t)\}$ are orthonormal functions in $\mathcal{N}$ and $\mathcal{T}$, accordingly. The trace norm of p is then defined as $\|p\|_* := \sum_{r=1}^\infty \sigma_r$[2].

To show that $h(p) \leq \sqrt{\|p\|_*}$, consider the spectral decomposition of $p = \sum_{r=1}^R \sigma_r u_r v_r$. Choose $f_r = \sqrt{\sigma_r} u_r$ and $g_r = \sqrt{\sigma_r} v_r$ for $r = 1, \ldots, R$. Since $\{f_r, g_r\}$ are feasible for (9) and attain a cost of $\sqrt{\|p\|_*}$, it follows that $h(p) \leq \sqrt{\|p\|_*}$.

It is next shown that $\sqrt{\|p\|_*} \leq h(p)$. Because the square root is strictly increasing, it can be applied on (7) to yield

$$\|p\|_*^{1/2} = \min_{\{f_r, g_r\}} \left\{ \sqrt{\frac{1}{2}\sum_{r=1}^R \|f_r\|_{\mathcal{K}}^2 + \|g_r\|_{\mathcal{G}}^2} : p = \sum_{r=1}^R f_r g_r \right\}. \tag{34}$$

Let $\{f_r^*, g_r^*\}_{r=1}^R$ be minimizers of (9). By Lemma 3, they yield a minimum of $h(p) = \sqrt{\sum_{r=1}^R \|g_r^*\|_{\mathcal{G}}^2}$. These minimizers are also feasible for (34), while attaining a cost of $\sqrt{\sum_{r=1}^R \|g_r^*\|_{\mathcal{G}}^2}$; hence, $\sqrt{\|p\|_*} \leq \sqrt{\sum_{r=1}^R \|g_r^*\|_{\mathcal{G}}^2} = h(p)$. ∎

C. Proof of Theorem 1

Theorem 1 builds upon the key result of [6, p. 352–53]:

Theorem 2 (Aronszajn, 1950): If K_l is the kernel of the function family $\mathcal{H}_{\mathcal{K}_l}$ having norm $\|\cdot\|_{\mathcal{K}_l}$, then $K = \sum_{l=1}^L \theta_l K_l$ for any $L \geq 2$ and $\theta_l > 0$, is the reproducing kernel of the function family $f = \sum_{l=1}^L f_l$ with $f_l \in \mathcal{H}_{\mathcal{K}_l}$, having the norm $\|f\|_{\mathcal{K}}^2 = \min\left\{\sum_{l=1}^L (\|f_l\|_{\mathcal{K}_l}^2/\theta_l) : f = \sum_{l=1}^L f_l, f_l \in \mathcal{H}_{\mathcal{K}_l}\right\}$.

Proof of Theorem 1: Theorem 2 asserts that a conic combination of kernels defines a function family whose members can be alternatively represented as a sum of functions defined by the constituent kernels. Applying this result to the convex combinations of (11), allows replacing (12) with

$$\min_{\mathcal{K},\mathcal{G}} \min_{p\in\mathcal{P}'} Q(\mathcal{K}, \mathcal{G}, p) \tag{35}$$

where $\mathcal{P}'$ has been defined in (13). Upon exchanging the order of minimizations in (35), consider solving the inner one, that is $\min_{\mathcal{K},\mathcal{G}} Q(\mathcal{K}, \mathcal{G}, p)$. The LS term is constant for a fixed $p \in \mathcal{P}'$, while the two regularization terms can be separately minimized over $\mathcal{K}$ and $\mathcal{G}$, respectively.

Focus now on solving $\min_{\mathcal{K}} \left(\sum_{r=1}^R \|f_r\|_{\mathcal{K}}^2\right)^{1/2}$. By Theorem 2, for fixed $f_r \in \mathcal{H}_{\mathcal{K}}$, there exist $\{f_{lr} \in \mathcal{H}_{\mathcal{K}_l}\}$ such that

$$\|f_r\|_{\mathcal{K}}^2 = \sum_{l=1}^L \frac{\|f_{lr}\|_{\mathcal{K}_l}^2}{\theta_l}. \tag{36}$$

Summing (36) over r and defining $\alpha_l^2 := \sum_{r=1}^R \|f_{lr}\|_{\mathcal{K}_l}^2$ yields

$$\sum_{r=1}^R \|f_r\|_{\mathcal{K}}^2 = \sum_{r=1}^R \sum_{l=1}^L \frac{\|f_{lr}\|_{\mathcal{K}_l}^2}{\theta_l} = \sum_{l=1}^L \frac{\alpha_l^2}{\theta_l}. \tag{37}$$

Recall that minimizing over $\mathcal{K}$ amounts to finding the optimum $\{\theta_l\}_{l=1}^L$. By applying the Cauchy-Schwarz inequality, it can be shown that [30, Lemma 26]

$$\min_{\{\theta_l\}_{l=1}^L} \left\{ \sqrt{\sum_{l=1}^L \frac{\alpha_l^2}{\theta_l}} : \theta_l > 0, \sum_{l=1}^L \theta_l = 1 \right\} = \sum_{l=1}^L \alpha_l. \tag{38}$$

Utilizing (38) to minimize the square root of (37), and replicating the analysis for $\{g_r\}_{r=1}^R$ completes the proof. ∎

D. Proof of Lemma 2

Lemma 2 generalizes [32, Corollary 2] to matrix variables.

Lemma 4 ([32]): The solution to the ℓ_2-penalized LS problem

$$\hat{\boldsymbol{\theta}} := \arg\min_{\boldsymbol{\theta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\theta}\|_2^2 + \mu\|\boldsymbol{\theta}\|_2$$

is $\hat{\boldsymbol{\theta}} = \left(\mathbf{X}^\top\mathbf{X} + (\mu^2/4\hat{w})\mathbf{I}\right)^{-1}\mathbf{X}^\top\mathbf{y}$ when $\|\mathbf{X}^\top\mathbf{y}\|_2 > \mu/2$; and $\mathbf{0}$, otherwise. The scalar $\hat{w} > 0$ minimizes the convex problem

$$\min_{w\geq 0} w - \mathbf{y}^\top\mathbf{X}\left(\mathbf{X}^\top\mathbf{X} + \frac{\mu^2}{4w}\mathbf{I}\right)^{-1}\mathbf{X}^\top\mathbf{y}. \tag{39}$$

Proof of Lemma 2: Since $\mathbf{B} \succ \mathbf{0}$, the problem in (30) can be equivalently expressed in terms of $\mathbf{X}' := \mathbf{B}^{1/2}\mathbf{X}$ as

$$\min_{\mathbf{X}'} \|\mathbf{A} - \mathbf{B}^{1/2}\mathbf{X}'\mathbf{C}^\top\|_F^2 + \mu\|\mathbf{X}'\|_F. \tag{40}$$

Upon defining $\mathbf{a} := \text{vec}(\mathbf{A})$ and using property (P), (40) can be expressed in terms of $\mathbf{x}' := \text{vec}(\mathbf{X}')$ as

$$\min_{\mathbf{x}'} \ \|\mathbf{a} - (\mathbf{C} \otimes \mathbf{B}^{1/2})\mathbf{x}'\|_2^2 + \mu\|\mathbf{x}'\|_2. \tag{41}$$

By Lemma 4, the minimizer of (41) is the solution of

$$\left(\mathbf{C}^\top\mathbf{C} \otimes \mathbf{B} + \frac{\mu^2}{4w}\mathbf{I}\right)\hat{\mathbf{x}}' = (\mathbf{C}^\top \otimes \mathbf{B}^{1/2})\mathbf{a} \tag{42}$$

when $\|(\mathbf{C}^\top \otimes \mathbf{B}^{1/2})\mathbf{a}\|_2 > \mu/2$; or $\hat{\mathbf{x}}' = \mathbf{0}$, otherwise. Using property (P) and if $\hat{\mathbf{x}}' = \text{vec}(\hat{\mathbf{X}}')$, then $\hat{\mathbf{X}}'$ satisfies $\mathbf{B}\hat{\mathbf{X}}'\mathbf{C}^\top\mathbf{C} + \mu^2/(4w)\hat{\mathbf{X}}' = \mathbf{B}^{1/2}\mathbf{A}\mathbf{C}$ when $\|\mathbf{B}^{1/2}\mathbf{A}\mathbf{C}\|_F > \mu/2$; otherwise, $\hat{\mathbf{X}}' = \mathbf{0}$. Transforming back to the sought $\hat{\mathbf{X}} = \mathbf{B}^{-1/2}\hat{\mathbf{X}}'$, yields finally (31).

The scalar $\hat{w}$ in (31) is the minimizer of the optimization problem obtained upon replacing $\mathbf{X}$ and $\mathbf{y}$ in (39) by $\mathbf{C} \otimes \mathbf{B}^{1/2}$ and $\mathbf{a}$, respectively. Given the singular value decompositions $\mathbf{C} = \mathbf{U}_C\mathbf{\Sigma}_C\mathbf{V}_C^\top$ and $\mathbf{B}^{1/2} = \mathbf{U}_B\mathbf{\Sigma}_B\mathbf{V}_B^\top$, and after algebraic manipulations, $\hat{w}$ can be shown to be the minimizer of

$$\min_{w>0} \ w - \mathbf{w}^\top\left(\mathbf{\Sigma}_C^2 \otimes \mathbf{\Sigma}_B^2\right)\left(\mathbf{\Sigma}_C^2 \otimes \mathbf{\Sigma}_B^2 + \frac{\mu^2}{4w}\mathbf{I}\right)^{-1}\mathbf{w} \tag{43}$$

where $\mathbf{w} := (\mathbf{U}_C^\top \otimes \mathbf{U}_B^\top)\mathbf{a}$. Recognizing that the matrices in (43) are diagonal and that the $d_1 \times d_2$ matrix version of $\mathbf{w}$ is $\mathbf{W} = \mathbf{U}_B^\top\mathbf{A}\mathbf{U}_C$, yields (32) thus completing the proof. ■

REFERENCES

[1] J. Abernethy, F. Bach, T. Evgeniou, and J.-P. Vert, "Low-rank matrix factorization with attributes," Ecole des Mines de Paris, Tech. Rep. N24/06/MM, Sep. 2006.

[2] J. Abernethy, F. Bach, T. Evgeniou, and J.-P. Vert, "A new approach to collaborative filtering: Operator estimation with spectral regularization," *J. Mach. Learn. Res.*, vol. 10, pp. 803–826, 2009.

[3] M. A. Alvarez, L. Rosasco, and N. D. Lawrence, "Kernels for vector-valued functions: A review," *Foundat. Trends Mach. Learn.*, vol. 4, no. 3, pp. 195–266, 2012.

[4] N. Amjady and M. Hemmati, "Energy price forecasting—problems and proposals for such predictions," *IEEE Power Energy Mag.*, vol. 4, no. 2, pp. 20–29, Mar./Apr. 2006.

[5] A. Argyriou, C. A. Michelli, and M. Pontil, "When is there a representer theorem? Vector versus matrix regularizers," *J. Mach. Learn. Res.*, vol. 10, pp. 2507–2529, 2009.

[6] N. Aronszajn, "Theory of reproducing kernels," *Trans. Amer. Math. Soc.*, vol. 68, no. 3, pp. 337–404, May 1950.

[7] J. A. Bazerque and G. B. Giannakis, "Nonparametric basis pursuit via sparse kernel-based learning," *IEEE Signal Process.*, vol. 12, no. 7, pp. 112–125, Jul. 2013.

[8] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA, USA: Athena Scientific, 1999.

[9] P. J. Brockwell and R. A. Davis, *Time Series: Theory and methods*, 2nd ed. New York, NY, USA: Springer, 1991.

[10] J. Contreras, R. Espinola, F. J. Nogales, and A. J. Conejo, "ARIMA models to predict next-day electricity prices," *IEEE Trans. Power Syst.*, vol. 18, no. 3, pp. 1014–1020, Aug. 2003.

[11] S. J. Deng and S. S. Oren, "Electricity derivatives and risk management," *Energy*, vol. 31, no. 6, pp. 940–953, 2006.

[12] Electric Reliability Council of Texas, ERCOT launches wholesale pricing forecast tool, 2012, [Online]. Available: http://www.ercot.com/news/press_releases/show/26244

[13] M. Fazel, "Matrix rank minimization with applications," Ph.D. dissertation, Stanford Univ., Stanford, CA, USA, 2002.

[14] R. C. Garcia, J. Contreras, M. v. Akkeren, and J. B. C. Garcia, "A GARCH forecasting model to predict day-ahead electricity prices," *IEEE Trans. Power Syst.*, vol. 20, no. 2, pp. 867–874, May 2005.

[15] G. B. Giannakis, V. Kekatos, N. Gatsis, S.-J. Kim, H. Zhu, and B. Wollenberg, "Monitoring and optimization for power grids: A signal processing perspective," *IEEE Signal Process. Mag.*, vol. 30, no. 5, pp. 107–128, Sep. 2013.

[16] G. H. Golub and C. F. v. Loan, *Matrix Computations*. Baltimore, MD, USA: John Hopkins Univ. Press, 1996.

[17] *Electric Energy Systems, Analysis and Operation*, A. Gómez-Expósito, A. Conejo, and C. Canizares, Eds. Boca Raton, FL, USA: CRC, 2009.

[18] M. Gonen and E. Alpaydin, "Multiple kernel learning algorithms," *J. Mach. Learn. Res.*, vol. 12, pp. 2211–2268, Sep. 2011.

[19] A. M. Gonzalez, A. M. S. Roque, and J. G. Gonzalez, "Modeling and forecasting electricity prices with input/output hidden Markov models," *IEEE Trans. Power Syst.*, vol. 20, no. 1, pp. 13–24, Feb. 2005.

[20] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York, NY, USA: Springer Series in Statistics, 2009.

[21] R. J. Hyndman, Forecast: Forecasting functions for time series and linear models, Feb. 2014, [Online]. Available: http://cran.r-project.org/web/packages/forecast/index.html

[22] V. Kekatos, G. B. Giannakis, and R. Baldick, "Grid topology identification using electricity prices," in *Proc. IEEE PES Soci. General Meeting*, Washington, DC, USA, Jul. 2014, pp. 91–96.

[23] V. Kekatos, S. Veeramachaneni, M. Light, and G. B. Giannakis, "Day-ahead electricity market forecasting," in *Proc. IEEE PES Innovative Smart Grid Technol.*, Washington, DC, USA, Feb. 2013, pp. 1–5.

[24] D. Kirschen and G. Strbac, *Power System Economics*. West Sussex, U.K.: Wiley, 2010.

[25] E. D. Kolaczyk, *Statistical Analysis of Network Data, Methods and Models*. New York, NY, USA: Springer, 2010.

[26] V. Koltchinskii and M. Yuan, "Sparsity in multiple kernel learning," *Ann. Statist.*, vol. 38, no. 6, pp. 3660–3695, 2010.

[27] G. Li, C.-C. Liu, C. Mattson, and J. Lawarree, "Day-ahead electricity price forecasting in a grid environment," *IEEE Trans. Power Syst.*, vol. 22, no. 1, pp. 266–274, Feb. 2007.

[28] A. T. Lora, J. M. R. Santos, A. G. Exposito, J. L. M. Ramos, and J. C. R. Santos, "Electricity market price forecasting based on weighted nearest neighbors techniques," *IEEE Trans. Power Syst.*, vol. 22, no. 3, pp. 1294–1301, Aug. 2007.

[29] M. Mardani, G. Mateos, and G. Giannakis, "Decentralized sparsity-regularized rank minimization: Algorithms and applications," *IEEE Trans. Signal Process.*, vol. 61, no. 21, pp. 5374–5388, Nov. 2013.

[30] C. Michelli and M. Pontil, "Learning the kernel function via regularization," *J. Mach. Learn. Res.*, vol. 6, pp. 1099–1125, Sep. 2005.

[31] A. L. Ott, "Experience with PJM market operation, system design, and implementation," *IEEE Trans. Power Syst.*, vol. 18, no. 2, pp. 528–534, May 2003.

[32] A. T. Puig, A. Wiesel, G. Fleury, and A. H. Hero, "Multidimensional shrinkage-thresholding operator and group LASSO penalties," *IEEE Signal Process. Lett.*, vol. 18, no. 6, pp. 363–366, Jun. 2011.

[33] B. Recht, M. Fazel, and P. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010.

[34] M. Shahidehpour, H. Yamin, and Z. Li, *Market Operations in Electric Power Systems: Forecasting, Scheduling, and Risk Management*. New York, NY, USA: IEEE-Wiley Interscience, 2002.

[35] V. Sindhwani, H. Q. Minh, and A. C. Lozano, "Scalable matrix-valued kernel learning for high-dimensional nonlinear multivariate regression and granger causality," in *Proc. Uncertainty Artif. Intell.*, Bellevue, WA, USA, Jul. 2013.

[36] A. J. Smola and R. Kondor, B. Schölkopf and M. Warmuth, Eds., "Kernels and regularization on graphs," in *Proc. Ann. Conf. Comput. Learn. Theory Kernel Workshop, Ser. Lecture Notes in Comput. Sci. Springer*, 2003.

[37] N. Srebro and A. Shraibman, "Rank, trace-norm and max-norm," in *Ser. Learning Theory, Ser. Lecture Notes in Comput. Sci*, P. Auer and R. Meir, Eds. Berlin, Germany: Springer, 2005, vol. 3559, pp. 545–560.

[38] P. Tseng, "Convergence of block coordinate descent method for non-differentiable minimization," *J. Optimiz. Theory Applicat.*, vol. 109, pp. 475–494, Jun. 2001.

[39] U.S. Dept. of Energy, National electric transmission congestion study, 2012, [Online]. Available: http://energy.gov/oe/services/electricity-policy-coordination-and-implementation/transmission-planning/2012-national

[40] L. Wu and M. Shahidehpour, "A hybrid model for day-ahead price forecasting," *IEEE Trans. Power Syst.*, vol. 25, no. 3, pp. 1519–1530, Aug. 2010.

[41] H. Xu, C. Caramanis, and S. Mannor, "Robust regression and LASSO," *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3561–3574, Jul. 2010.

[42] L. Zhang, P. B. Luh, and K. Kasiviswanathan, "Energy clearing price prediction and confidence interval estimation with cascaded neural network," *IEEE Trans. Power Syst.*, vol. 18, no. 1, pp. 99–105, Feb. 2003.

[43] Q. Zhou, L. Tesfatsion, and C.-C. Liu, "Short-term congestion forecasting in wholesale power markets," *IEEE Trans. Power Syst.*, vol. 26, no. 4, pp. 2185–2196, Nov. 2011.

Vassilis Kekatos (M'10) obtained his Diploma, M.Sc., and Ph.D. in computer engineering and informatics from the University of Patras, Greece, in 2001, 2003, and 2007, respectively. He is currently a postdoctoral associate with the Dept. of Electrical and Computer Engineering of the University of Minnesota. In 2009, he received a Marie Curie fellowship. During the summer of 2012, he worked as a consultant for Windlogics Inc. His current interests lie in the areas of signal processing, optimization, and statistical learning for smart power grids.

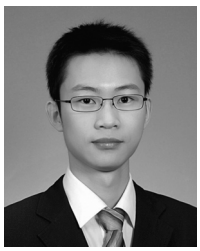

Yu Zhang (S'11) received his B.Eng. and M.Sc. degrees (both with highest honors) in electrical engineering from Wuhan University of Technology, Wuhan, China, and from Shanghai Jiao Tong University, Shanghai, China, in 2006 and 2010, respectively. Since September 2010, he has been working towards the Ph.D. degree with the Department of Electrical and Computer Engineering (ECE) at the University of Minnesota (UMN). During the summer of 2014, he was a research intern at ABB US Corporate Research Center, Raleigh, NC. His research interests span the areas of smart power grids, machine learning, and wireless communications. Mr. Zhang received the Huawei Scholarship, the Infineon Scholarship in Shanghai, 2009, and the UMN ECE Dept. Fellowship in Minneapolis, 2010.

Georgios B. Giannakis (F'97) received his Diploma in Electrical Engineering from the National Technical University of Athens, Greece, 1981. From 1982 to 1986 he was with the University of Southern California (USC), where he received his M.Sc. in electrical engineering, 1983, the M.Sc. in mathematics in 1986, and the Ph.D. in electrical engineering in 1986. Since 1999 he has been a professor with the University of Minnesota, where he now holds an ADC Chair in Wireless Telecommunications in the ECE Department, and serves as director of the Digital Technology Center.

His general interests span the areas of communications, networking and statistical signal processing — subjects on which he has published more than 370 journal papers, 630 conference papers, 20 book chapters, two edited books and two research monographs (h-index 107). Current research focuses on sparsity and big data analytics, wireless cognitive radios, mobile ad hoc networks, renewable energy, power grid, gene-regulatory, and social networks. He is the (co-) inventor of 21 patents issued, and the (co-) recipient of 8 best paper awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize Paper Award in Wireless Communications. He also received Technical Achievement Awards from the SP Society (2000), from EURASIP (2005), a Young Faculty Teaching Award, the G. W. Taylor Award for Distinguished Research from the University of Minnesota, and the IEEE Fourier Technical Field Award (2014). He is a Fellow of EURASIP, and has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SP Society.