

Matrix Completion and Extrapolation via Kernel Regression

Pere Giménez-Febrer , Alba Pagès-Zamora , and Georgios B. Giannakis 

Abstract—Matrix completion and extrapolation (MCEX) are dealt with here over reproducing kernel Hilbert spaces (RKHSs) in order to account for prior information present in the available data. Aiming at a fast and low-complexity solver, the task is formulated as one of kernel ridge regression. The resultant MCEX algorithm can also afford online implementation, while the class of kernel functions also encompasses several existing approaches to MC with prior information. Numerical tests on synthetic and real datasets show that the novel approach is faster than widespread methods such as alternating least-squares (ALS) or stochastic gradient descent (SGD), and that the recovery error is reduced, especially when dealing with noisy data.

Index Terms—Matrix completion, extrapolation, RKHS, kernel ridge regression, graphs, online learning.

I. INTRODUCTION

WITH only a subset of its entries available, matrix completion (MC) amounts to recovering the unavailable entries by leveraging just the low-rank attribute of the matrix itself [1]. The relevant task arises in applications as diverse as image restoration [2], sensor networks [3], and recommender systems [4]. To save power for instance, only a fraction of sensors may collect and transmit measurements to a fusion center, where the available spatio-temporal data can be organized in a matrix format, and the unavailable ones can be eventually interpolated via MC [3]. Similarly, collaborative filtering of ratings given by users to a small number of items are stored in a sparse matrix, and the objective is to predict their ratings for the rest of the items [4].

Existing MC approaches rely on some form of rank minimization or low-rank matrix factorization. Specifically, [1] proves that when MC is formulated as the minimization of the nuclear norm subject to the constraint that the observed entries remain unchanged, exact recovery is possible under mild assumptions;

Manuscript received July 31, 2018; revised March 18, 2019; accepted July 24, 2019. Date of publication August 2, 2019; date of current version August 23, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jarvis Haupt. This work was supported in part by the Ministerio de Economía y Competitividad of the Spanish Government and ERDF funds under Grants TEC2016-75067-C4-2-R and TEC2015-51569648-REDC; in part by Catalan Government funds under Grant 2017 SGR 578 AGAUR; and in part by NSF under Grants 1500713, 1514056, 1711471, and 1509040. (Corresponding author: Pere Giménez-Febrer.)

P. Giménez-Febrer and A. Pagès-Zamora are with the SPCOM Group, Universitat Politècnica de Catalunya-Barcelona Tech, Barcelona 08034, Spain (e-mail: p.gimenez@upc.edu; alba.pages@upc.edu).

G. B. Giannakis is with the Department of ECE and Digital Technology Center, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: georgios@umn.edu).

Digital Object Identifier 10.1109/TSP.2019.2932875

see also [5] where reliable recovery from a few observations is established even in the presence of additive noise. Alternatively, [4] replaces the nuclear norm by two low-rank factor matrices that are identified in order to recover the complete matrix.

While the low-rank assumption can be sufficient for reliable recovery, prior information about the unknown matrix can be also accounted to improve the completion outcome. Forms of prior information can include sparsity [3], local smoothness [6], and interdependencies encoded by graphs [7]–[10]. These approaches exploit the available similarity information or prior knowledge of the bases spanning the column or row spaces of the unknown matrix. In this regard, reproducing kernel Hilbert spaces (RKHSs) constitute a powerful tool for leveraging available prior information thanks to the kernel functions, which measure the similarity between pairs of points in an input space. Prompted by this, [11]–[14] postulate that columns of the factor matrices belong to a pair of RKHSs spanned by their respective kernels. In doing so, a given structure or similarity between rows or columns is effected on the recovered matrix. Upon choosing a suitable kernel function, [3] as well as [6]–[10] can be cast into the RKHS framework. In addition to improving MC performance, kernel-based approaches also enable extrapolation of rows and columns, even when all their entries are missing – a task impossible by the standard MC approaches in e.g. [1] and [4].

One major hurdle in MC is the computational cost as the matrix size grows. In its formulation as a rank minimization task, MC can be solved via semidefinite programming [1], or proximal gradient minimization [8], [15]–[17], which entails a singular value decomposition of the recovered matrix per iteration. Instead, algorithms with lower computational cost are available for the bi-convex formulation based on matrix factorization [4]. These commonly rely on iterative minimization schemes such as alternating least-squares (ALS) [18], [19] or stochastic gradient descent (SGD) [13], [20]. With regard to kernel-based MC, the corresponding algorithms rely on alternating convex minimization and semidefinite programming [11], block coordinate descent [12], and SGD [13]. However, algorithms based on alternating minimization only converge to the minimum after infinite iterations. In addition, existing kernel-based algorithms adopt a specific sampling pattern or do not effectively make use of the Representer Theorem for RKHSs that will turn out to be valuable in further reducing the complexity, especially when the number of observed entries is small.

The present contribution offers an RKHS-based approach to MCEX that also unifies and broadens the scope of MC

approaches, while offering reduced complexity algorithms that scale well with the data size. Specifically, we develop a novel MC solver via kernel ridge regression as a convex alternative to the nonconvex factorization-based formulation that offers a closed-form solution. Through an explicit sampling matrix, the proposed method offers an encompassing sampling pattern, which further enables the derivation of upper bounds on the mean-square error. Moreover, an approximate solution to our MCEX regression formulation is developed that also enables online implementation using SGD. Finally, means of incorporating prior information through kernels is discussed in the RKHS framework.

The rest of the paper is organized as follows. Section II outlines the RKHS formulation and the kernel regression task. Section III unifies the existing methods for MC under the RKHS umbrella, while Section IV introduces our proposed Kronecker kernel MCEX (KKMCEX) approach. Section V develops our ridge regression MCEX (RRMCEX) algorithm, an accelerated version of KKMCEX, and its online variant. Section VI deals with the construction of kernel matrices. Finally, Section VII presents numerical tests, and Section VIII concludes the paper.

Notation: Boldface lower case fonts denote column vectors, and boldface uppercase fonts denote matrices. The (i, j) th entry of matrix \mathbf{A} is $\mathbf{A}_{i,j}$, and the i th entry of vector \mathbf{a} is \mathbf{a}_i . Superscripts T and \dagger denote transpose and pseudoinverse, respectively; while hat $\hat{\cdot}$ is used for estimates. Matrix $\mathbf{F} \in \mathcal{H}$ means that its columns belong to a vector space \mathcal{H} . The symbols \mathbf{I} and $\mathbf{1}$ stand for the identity matrix and the all-ones vector of appropriate size, specified by the context. The trace operator is $\text{Tr}(\cdot)$, the function $\text{eig}(\mathbf{A})$ returns the diagonal eigenvalue matrix of \mathbf{A} ordered in ascending order, and $\lambda_k(\mathbf{A})$ denotes the k th eigenvalue of \mathbf{A} with $\lambda_k(\mathbf{A}) \leq \lambda_{k+1}(\mathbf{A})$.

II. PRELIMINARIES

Consider a set of N input-measurement pairs $\{(x_i, m_i)\}_{i=1}^N$ in $\mathcal{X} \times \mathbb{R}$, where $\mathcal{X} := \{x_1, \dots, x_N\}$ is the input space, \mathbb{R} denotes the set of real numbers, and measurements obey the model

$$m_i = f(x_i) + e_i \quad (1)$$

where $f : \mathcal{X} \rightarrow \mathbb{R}$ is an unknown function and $e_i \in \mathbb{R}$ is noise. We assume this function belongs to an RKHS

$$\mathcal{H}_x := \left\{ f : f(x_i) = \sum_{j=1}^N \alpha_j \kappa_x(x_i, x_j), \quad \alpha_j \in \mathbb{R} \right\} \quad (2)$$

where $\kappa_x : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is the kernel function that spans \mathcal{H}_x , and $\{\alpha_i\}_{i=1}^N$ are weight coefficients. An RKHS is a complete linear space endowed with an inner product that satisfies the reproducing property [21]. If $\langle f, f' \rangle_{\mathcal{H}_x}$ denotes the inner product in \mathcal{H}_x between functions f and f' , the reproducing property states that $f(x_i) = \langle f, \kappa_x(\cdot, x_i) \rangle_{\mathcal{H}_x}$; that is, f in \mathcal{H}_x can be evaluated at x_i by taking the inner product between f and $\kappa_x(\cdot, x_i)$. With $\{\alpha_i\}_{i=1}^N$ and $\{\alpha'_i\}_{i=1}^N$ denoting the coefficients of f and f' in (2) respectively, we have $\langle f, f' \rangle_{\mathcal{H}_x} := \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha'_j \kappa_x(x_i, x_j)$; that is,

$$\langle f, f' \rangle_{\mathcal{H}_x} = \boldsymbol{\alpha}^T \mathbf{K}_x \boldsymbol{\alpha}' \quad (3)$$

where $\boldsymbol{\alpha} := [\alpha_1, \dots, \alpha_N]^T$, $\boldsymbol{\alpha}' := [\alpha'_1, \dots, \alpha'_N]^T$ and $(\mathbf{K}_x)_{i,j} := \kappa_x(x_i, x_j)$. In order for $\langle \cdot, \cdot \rangle_{\mathcal{H}_x}$ in (3) to be an inner product, κ_x must be symmetric and semipositive definite, meaning $\langle f, f \rangle_{\mathcal{H}_x} \geq 0 \forall f \in \mathcal{H}_x$. As a consequence, \mathbf{K}_x will be symmetric positive semidefinite since $\boldsymbol{\alpha}^T \mathbf{K}_x \boldsymbol{\alpha} \geq 0 \forall \boldsymbol{\alpha} \in \mathbb{R}^N$.

While κ_x is usually interpreted as a measure of similarity between two elements in \mathcal{X} , it can also be seen as the inner product of two corresponding elements in feature space \mathcal{F} to which \mathcal{X} can be mapped using function $\phi_x : \mathcal{X} \rightarrow \mathcal{F}$. Formally, we write

$$\kappa_x(x_i, x_j) = \langle \phi_x(x_i), \phi_x(x_j) \rangle_{\mathcal{F}} \quad (4)$$

Function ϕ_x is referred to as feature map, and its choice depends on the application. For an input space of text files, for example, the files could be mapped to a feature vector that tracks the number of words, lines, and blank spaces in the file. Since \mathcal{F} can potentially have infinite dimension, evaluating the kernel using (4) might be prohibitively expensive. This motivates specifying the kernel through a similarity function in \mathcal{X} , which bypasses the explicit computation of the inner product in \mathcal{F} . Typical examples include the Gaussian kernel $\kappa_x(x_i, x_j) = \exp\{-\|x_i - x_j\|_2^2 / (2\eta)\}$ with η being a free parameter, and the polynomial kernel [22]. In certain cases however, it is difficult to obtain the kernel similarity function on the input space. Such cases include metric input spaces with misses (as in MC), and non-metric spaces. The alternative to both is deriving the kernel from prior information. For instance, if we have a graph connecting the points in \mathcal{X} , a kernel can be obtained from the graph Laplacian [23].

Having introduced the basics of RKHS, we proceed with the kernel regression task, where given $\{m_i\}_{i=1}^N$ we seek to obtain

$$\hat{f} = \arg \min_{f \in \mathcal{H}_x} \frac{1}{N} \sum_{i=1}^N l(m_i, f(x_i)) + \mu' \|f\|_{\mathcal{H}_x}^2 \quad (5)$$

with $l(\cdot)$ denoting the loss, $\mu' \in \mathbb{R}^+$ the regularization parameter, and $\|f\|_{\mathcal{H}_x} := \langle f, f \rangle_{\mathcal{H}_x}$ the norm induced by the inner product in (3). We will henceforth focus on the square loss $l(m_i, f(x_i)) := (m_i - f(x_i))^2$. Using \mathbf{K}_x , consider without loss of generality expressing the vector $\mathbf{f} := [f(x_1), \dots, f(x_N)]^T$ as $\mathbf{f} = \mathbf{K}_x \boldsymbol{\alpha}$, where $\boldsymbol{\alpha} := [\alpha_1, \dots, \alpha_N]^T$. Using the latter in the square loss, (5) boils down to a kernel ridge regression (KRR) problem that can be solved by estimating $\boldsymbol{\alpha}$ as

$$\hat{\boldsymbol{\alpha}} = \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^N} \|\mathbf{m} - \mathbf{K}_x \boldsymbol{\alpha}\|_2^2 + \mu \boldsymbol{\alpha}^T \mathbf{K}_x \boldsymbol{\alpha} \quad (6)$$

where $\mathbf{m} := [m_1, \dots, m_N]^T$ and $\mu = N\mu'$. The weights can be found in closed form as

$$\hat{\boldsymbol{\alpha}} = (\mathbf{K}_x + \mu \mathbf{I})^{-1} \mathbf{m} \quad (7)$$

and the estimate of the sought function is obtained as $\hat{f} = \mathbf{K}_x \hat{\boldsymbol{\alpha}}$.

III. KERNEL-BASED MCEX

Matrix completion considers $\mathbf{F} \in \mathbb{R}^{N \times L}$ of rank r observed through a $N \times L$ matrix of noisy observations

$$\mathbf{M} = P_{\Omega}(\mathbf{F} + \mathbf{E}) \quad (8)$$

where $\Omega \subseteq \{1, \dots, N\} \times \{1, \dots, L\}$ is the sampling set of cardinality $S = |\Omega|$ containing the indices of the observed entries; $P_\Omega(\cdot)$ is a projection operator that sets to zero the entries with index $(i, j) \notin \Omega$ and leaves the rest unchanged; and, $\mathbf{E} \in \mathbb{R}^{N \times L}$ is a noise matrix. According to [5], one can recover \mathbf{F} from \mathbf{M} with an error proportional to the magnitude of $\|\mathbf{E}\|_F^2$ by solving the following convex optimization problem:

$$\begin{aligned} \min_{\mathbf{F} \in \mathbb{R}^{N \times L}} \quad & \text{rank}(\mathbf{F}) \\ \text{subject to} \quad & \|P_\Omega(\mathbf{F} - \mathbf{M})\|_F^2 \leq \delta \end{aligned} \quad (9)$$

where $\|\cdot\|_F$ is the Frobenius norm, and we assume $\|P_\Omega(\mathbf{E})\|_F^2 \leq \delta$ for some $\delta > 0$. Since solving (9) is NP-hard, the nuclear norm $\|\mathbf{F}\|_* := \text{Tr}(\sqrt{\mathbf{F}^T \mathbf{F}})$ can be used to replace the rank to obtain the convex problem [8], [15]

$$\min_{\mathbf{F} \in \mathbb{R}^{N \times L}} \|P_\Omega(\mathbf{M} - \mathbf{F})\|_F^2 + \mu \|\mathbf{F}\|_*. \quad (10)$$

Because \mathbf{F} is low rank, it is always possible to factorize it as $\mathbf{F} = \mathbf{W}\mathbf{H}^T$, where $\mathbf{W} \in \mathbb{R}^{N \times p}$ and $\mathbf{H} \in \mathbb{R}^{L \times p}$ are the latent factor matrices with $p \geq r$. This factorization allows expressing the nuclear norm as [24] $\|\mathbf{F}\|_* = \min_{\mathbf{F}=\mathbf{W}\mathbf{H}^T} \frac{1}{2}(\|\mathbf{W}\|_F^2 + \|\mathbf{H}\|_F^2)$ which allows reformulating (10) as

$$\begin{aligned} \{\hat{\mathbf{W}}, \hat{\mathbf{H}}\} = \arg \min_{\substack{\mathbf{W} \in \mathbb{R}^{N \times p} \\ \mathbf{H} \in \mathbb{R}^{L \times p}}} & \|P_\Omega(\mathbf{M} - \mathbf{W}\mathbf{H}^T)\|_F^2 \\ & + \mu \left(\|\mathbf{W}\|_F^2 + \|\mathbf{H}\|_F^2 \right) \end{aligned} \quad (11)$$

and yields $\hat{\mathbf{F}} = \hat{\mathbf{W}}\hat{\mathbf{H}}^T$. While the solutions to (10) and (11) are equivalent when the rank of the matrix minimizing (10) is smaller than p [18], solving (10) can be costlier since it involves the computation of the singular values of the matrix. On the other hand, since (11) is bi-convex it can be solved by alternately optimizing \mathbf{W} and \mathbf{H} , e.g. via ALS [19] or SGD iterations [20]. Moreover, leveraging the structure of (11), it is also possible to optimize one row from each factor matrix at a time instead of updating the full factor matrices, which enables faster and also online and distributed implementations [25].

Aiming at a kernel-based MCEX, we model the columns and rows of \mathbf{F} as functions that belong to two different RKHSs. To this end, consider the input spaces $\mathcal{X} := \{x_1, \dots, x_N\}$ and $\mathcal{Y} := \{y_1, \dots, y_L\}$ for the column and row functions, respectively. In the user-movie ratings paradigm, \mathcal{X} could be the set of users, and \mathcal{Y} the set of movies. Then $\mathbf{F} := [\mathbf{f}_1, \dots, \mathbf{f}_L]$ is formed with columns $\mathbf{f}_l := [f_l(x_1), \dots, f_l(x_N)]^T$ with $f_l: \mathcal{X} \rightarrow \mathbb{R}$. Likewise, we rewrite $\mathbf{F} := [\mathbf{g}_1, \dots, \mathbf{g}_N]^T$, with rows $\mathbf{g}_n^T := [g_n(y_1), \dots, g_n(y_L)]$ and $g_n: \mathcal{Y} \rightarrow \mathbb{R}$. We further assume that $f_l \in \mathcal{H}_x \forall l = 1, \dots, L$ and $g_n \in \mathcal{H}_y \forall n = 1, \dots, N$, where

$$\mathcal{H}_x := \left\{ f : f(x_i) = \sum_{j=1}^N \alpha_j \kappa_x(x_i, x_j), \quad \alpha_j \in \mathbb{R} \right\} \quad (12)$$

$$\mathcal{H}_y := \left\{ g : g(y_i) = \sum_{j=1}^L \beta_j \kappa_y(y_i, y_j), \quad \beta_j \in \mathbb{R} \right\} \quad (13)$$

and $\kappa_x: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $\kappa_y: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ are the kernels forming $\mathbf{K}_x \in \mathbb{R}^{N \times N}$ and $\mathbf{K}_y \in \mathbb{R}^{L \times L}$, respectively.

Since \mathbf{W} and \mathbf{H} span the column and row spaces of \mathbf{F} , their columns belong to \mathcal{H}_x and \mathcal{H}_y as well. Thus, the m th column of \mathbf{W} is

$$\mathbf{w}_m := [w_m(x_1), \dots, w_m(x_N)]^T \quad (14)$$

where $w_m: \mathcal{X} \rightarrow \mathbb{R}$ and $w_m \in \mathcal{H}_x \forall m = 1, \dots, p$, and the m th column of \mathbf{H} is

$$\mathbf{h}_m := [h_m(y_1), \dots, h_m(y_L)]^T \quad (15)$$

where $h_m: \mathcal{Y} \rightarrow \mathbb{R}$ and $h_m \in \mathcal{H}_y \forall m = 1, \dots, p$. Hence, instead of simply promoting a small Frobenius norm for the factor matrices as in (11), we can also promote smoothness on their respective RKHS. The kernel-based formulation in [12] estimates the factor matrices by solving

$$\begin{aligned} \{\hat{\mathbf{W}}, \hat{\mathbf{H}}\} = \arg \min_{\substack{\mathbf{W} \in \mathcal{H}_x \\ \mathbf{H} \in \mathcal{H}_y}} & \|P_\Omega(\mathbf{M} - \mathbf{W}\mathbf{H}^T)\|_F^2 \\ & + \mu \text{Tr}(\mathbf{W}^T \mathbf{K}_x^{-1} \mathbf{W}) + \mu \text{Tr}(\mathbf{H}^T \mathbf{K}_y^{-1} \mathbf{H}). \end{aligned} \quad (16)$$

Note that (16) is equivalent to (11) for $\mathbf{K}_x = \mathbf{I}$ and $\mathbf{K}_y = \mathbf{I}$. Since the constraints $\mathbf{W} \in \mathcal{H}_x$ and $\mathbf{H} \in \mathcal{H}_y$ can be challenging to account for when solving (16), we can instead find the coefficients that generate \mathbf{W} and \mathbf{H} in their respective RKHSs in order to satisfy such constraints. Thus, if we expand $\mathbf{W} = \mathbf{K}_x \mathbf{B}$ and $\mathbf{H} = \mathbf{K}_y \mathbf{C}$, where $\mathbf{B} \in \mathbb{R}^{N \times p}$ and $\mathbf{C} \in \mathbb{R}^{L \times p}$ are coefficient matrices, (16) becomes

$$\begin{aligned} \{\hat{\mathbf{B}}, \hat{\mathbf{C}}\} = \arg \min_{\substack{\mathbf{B} \in \mathbb{R}^{N \times p} \\ \mathbf{C} \in \mathbb{R}^{L \times p}}} & \|P_\Omega(\mathbf{M} - \mathbf{K}_x \mathbf{B} \mathbf{C}^T \mathbf{K}_y)\|_F^2 \\ & + \mu \text{Tr}(\mathbf{B}^T \mathbf{K}_x \mathbf{B}) + \mu \text{Tr}(\mathbf{C}^T \mathbf{K}_y \mathbf{C}). \end{aligned} \quad (17)$$

Nevertheless, with nonsingular kernel matrices, \mathbf{B} and \mathbf{C} can be found by solving (16) and substituting $\hat{\mathbf{B}} = \mathbf{K}_x^{-1} \hat{\mathbf{W}}$ and $\hat{\mathbf{C}} = \mathbf{K}_y^{-1} \hat{\mathbf{H}}$ [12].

Alternating minimization schemes that solve the bi-linear MC formulation (11) tends to the solution to the convex problem (10) in the limit [19], thus convergence to the global optimum is not guaranteed unless the number of iterations is infinite. Since algorithms for kernel-based MC [12] solving (16) rely on such alternating minimization schemes, they lack convergence guarantees given finite iterations as well. In addition to that, their computational cost scales with the size of \mathbf{F} . On the other hand, online implementations have a lower cost [13], but only guarantee convergence to a stationary point [26]. In the ensuing section we develop a convex kernel-based reformulation of MCEX that enables a closed-form solver which purely exploits the extrapolation facilitated by the kernels. By casting aside the low-rank constraints, the computational complexity of our solver scales only with the number of observations while, according to our numerical tests, providing better performance. Moreover, we derive an online implementation that can be seamlessly extended to distributed operation.

IV. KRONECKER KERNEL MCEX

In the previous section, we viewed the columns and rows of \mathbf{F} as functions evaluated at the points of the input spaces \mathcal{X} and \mathcal{Y} in order to unify the state-of-the-art on MC using RKHSs. Instead, we now postulate entries of \mathbf{F} as the output of a function lying on an RKHS evaluated at a tuple $(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}$. Given the spaces \mathcal{X} and \mathcal{Y} , consider the space $\mathcal{Z} := \mathcal{X} \times \mathcal{Y}$ with cardinality $|\mathcal{Z}| = NL$ along with the two-dimensional function $v : \mathcal{Z} \rightarrow \mathbb{R}$ as $v(x_i, y_j) = f_j(x_i)$, which belongs to the RKHS

$$\mathcal{H}_z := \left\{ v : v(x_i, y_j) = \sum_{n=1}^N \sum_{l=1}^L \gamma_{n,l} \kappa_z((x_i, y_j), (x_n, y_l)), \quad \gamma_{n,l} \in \mathbb{R} \right\} \quad (18)$$

with $\kappa_z : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$. While one may choose any kernel to span \mathcal{H}_z , we will construct one adhering to our bilinear factorization $\mathbf{F} = \mathbf{W}\mathbf{H}^T$ whose (i, j) th entry yields

$$\mathbf{F}_{ij} = v(x_i, y_j) = \sum_{m=1}^p w_m(x_i) h_m(y_j) \quad (19)$$

with w_m and h_m functions capturing m th column vector of \mathbf{W} and \mathbf{H} as in (14) and (15). Since $w \in \mathcal{H}_x$ and $h \in \mathcal{H}_y$, we can write $w_m(x_i) = \sum_{n=1}^N b_{n,m} \kappa_x(x_i, x_n)$ and $h_m(y_j) = \sum_{l=1}^L c_{l,m} \kappa_y(y_j, y_l)$, where $b_{n,m}$ and $c_{l,m}$ are the entries at (n, m) and (l, m) of the factor matrices \mathbf{B} and \mathbf{C} from (17), respectively. Therefore, (19) can be rewritten as

$$\begin{aligned} v(x_i, y_j) &= \sum_{m=1}^p \sum_{n=1}^N b_{n,m} \kappa_x(x_i, x_n) \sum_{l=1}^L c_{l,m} \kappa_y(y_j, y_l) \\ &= \sum_{n=1}^N \sum_{l=1}^L \left(\sum_{m=1}^p b_{n,m} c_{l,m} \right) \kappa_x(x_i, x_n) \kappa_y(y_j, y_l) \\ &= \sum_{n=1}^N \sum_{l=1}^L \gamma_{n,l} \kappa_z((x_i, y_j), (x_n, y_l)) \end{aligned} \quad (20)$$

where $\gamma_{n,l} = \sum_{m=1}^p b_{n,m} c_{l,m}$, and $\kappa_z((x_i, y_j), (x_n, y_l)) = \kappa_x(x_i, x_n) \kappa_y(y_j, y_l)$ since a product of kernels is itself a kernel [22]. Using the latter, (20) can be written compactly as

$$v(x_i, y_j) = \mathbf{k}_{i,j}^T \boldsymbol{\gamma} \quad (21)$$

where $\boldsymbol{\gamma} := [\gamma_{1,1}, \gamma_{2,1}, \dots, \gamma_{N,1}, \gamma_{1,2}, \gamma_{2,2}, \dots, \gamma_{N,L}]^T$, and correspondingly,

$$\begin{aligned} \mathbf{k}_{i,j} &= [\kappa_x(x_i, x_1) \kappa_y(y_j, y_1), \dots, \kappa_x(x_i, x_N) \kappa_y(y_j, y_1), \\ &\quad \kappa_x(x_i, x_1) \kappa_y(y_j, y_2), \dots, \kappa_x(x_i, x_N) \kappa_y(y_j, y_L)]^T \\ &= (\mathbf{K}_y)_{:,j} \otimes (\mathbf{K}_x)_{:,i} \end{aligned} \quad (22)$$

where a subscript $(:, j)$ denotes the j th column of a matrix, and we have used that \mathbf{K}_x and \mathbf{K}_y are symmetric matrices. In accordance with (22), the kernel matrix of \mathcal{H}_z in (18) is

$$\mathbf{K}_z = \mathbf{K}_y \otimes \mathbf{K}_x. \quad (23)$$

Clearly, $\mathbf{k}_{i,j}$ in (22) can also be expressed as $\mathbf{k}_{i,j} = (\mathbf{K}_z)_{:(j-1)N+i}$. This together with (21) implies that

$$\mathbf{v} = [v(x_1, y_1), v(x_2, y_1), \dots, v(x_N, y_1), v(x_1, y_2), \quad (24)$$

$$v(x_2, y_2), \dots, v(x_N, y_N)]^T \quad (25)$$

can be expressed in matrix-vector form as

$$\mathbf{v} = \mathbf{K}_z \boldsymbol{\gamma} \quad (26)$$

or, equivalently, $\mathbf{v} = \text{vec}(\mathbf{F})$. Note that entries of the kernel matrix are $(\mathbf{K}_z)_{i',j'} = \kappa_x(x_{i'}, x_n) \kappa_y(y_{j'}, y_l)$, where $n = j' \bmod N, i = i' \bmod N, l = \lceil \frac{j'}{N} \rceil$, and $j = \lceil \frac{i'}{N} \rceil$.

Since the eigenvalues of \mathbf{K}_z are the product of eigenvalues of \mathbf{K}_y and \mathbf{K}_x , it follows that \mathbf{K}_z is positive semidefinite and thus a valid kernel matrix. With the definition of the function v and its vector form we have transformed the matrix of functions specifying \mathbf{F} into a function that lies on the RKHS \mathcal{H}_z . Hence, we are ready to formulate MCEX as a kernel regression task for recovering \mathbf{v} from the observed entries of $\mathbf{m} = \text{vec}(\mathbf{M})$.

Given $\{(x_i, y_j), m_{i,j}\}_{(i,j) \in \Omega}$ in $\mathcal{Z} \times \mathbb{R}$, our goal is to recover the function v as

$$\hat{v} = \arg \min_{v \in \mathcal{H}_z} \sum_{(i,j) \in \Omega} (m_{i,j} - v(x_i, y_j))^2 + \mu \|v\|_{\mathcal{H}_z}^2 \quad (27)$$

where $\|v\|_{\mathcal{H}_z}^2 := \boldsymbol{\gamma}^T \mathbf{K}_z \boldsymbol{\gamma}$. Define next $\mathbf{e} := \text{vec}(\mathbf{E})$ and $\bar{\mathbf{m}} = \mathbf{S}\mathbf{m}$, where \mathbf{S} is an $S \times NL$ binary sampling matrix also used to specify the sampled noise vector $\bar{\mathbf{e}} = \mathbf{S}\mathbf{e}$. With these definitions and (26), the model in (8) becomes

$$\bar{\mathbf{m}} = \mathbf{S}\mathbf{v} + \mathbf{S}\mathbf{e} = \mathbf{S}\mathbf{K}_z \boldsymbol{\gamma} + \bar{\mathbf{e}} \quad (28)$$

which can be solved to obtain

$$\hat{\boldsymbol{\gamma}} = \arg \min_{\boldsymbol{\gamma} \in \mathbb{R}^{NL}} \|\bar{\mathbf{m}} - \mathbf{S}\mathbf{K}_z \boldsymbol{\gamma}\|_2^2 + \mu \boldsymbol{\gamma}^T \mathbf{K}_z \boldsymbol{\gamma} \quad (29)$$

in closed form

$$\hat{\boldsymbol{\gamma}} = (\mathbf{S}^T \mathbf{S} \mathbf{K}_z + \mu \mathbf{I})^{-1} \mathbf{S}^T \bar{\mathbf{m}}. \quad (30)$$

Since the size of \mathbf{K}_z is $NL \times NL$, the inversion in (30) can be very computationally intensive. To alleviate this, we will leverage the Representer Theorem (see [27] for a formal proof), which allows us to reduce the number of degrees of freedom of the regression problem. In our setup, this theorem is as follows.

Theorem 1 (Representer Theorem): Given the set of input-observations pairs $\{(x_i, y_j), m_{i,j}\}_{(i,j) \in \Omega}$ in $\mathcal{Z} \times \mathbb{R}$ and the function v as in (20), the solution to

$$\arg \min_{v \in \mathcal{H}_z} \sum_{(i,j) \in \Omega} (m_{i,j} - v(x_i, y_j))^2 + \mu \|v\|_{\mathcal{H}_z}^2 \quad (31)$$

is an estimate \hat{v} that satisfies

$$\hat{v} = \sum_{(n,l) \in \Omega} \tau_{n,l} \kappa_z((\cdot, \cdot), (x_n, y_l)) \quad (32)$$

for some coefficients $\tau_{n,l} \in \mathbb{R}, \forall (n,l) \in \Omega$.

Theorem 1 asserts that $\hat{\boldsymbol{\gamma}}$ in (29) satisfies $\hat{\gamma}_{n,l} = 0 \forall (n,l) \notin \Omega$. Therefore, we only need to optimize $\{\gamma_{n,l} : (n,l) \in \Omega\}$

which correspond to the observed pairs. In fact, for our vector-based formulation, the Representer Theorem boils down to applying on (30) the matrix inversion lemma (MIL), which asserts the following.

Lemma 1 (MIL [28]): Given matrices \mathbf{A}, \mathbf{U} and \mathbf{V} of conformal dimensions, with \mathbf{A} invertible, it holds that

$$(\mathbf{U}\mathbf{V} + \mathbf{A})^{-1}\mathbf{U} = \mathbf{A}^{-1}\mathbf{U}(\mathbf{V}\mathbf{A}^{-1}\mathbf{U} + \mathbf{I}). \quad (33)$$

With (30) $\mathbf{A} = \mu\mathbf{I}, \mathbf{U} = \mathbf{S}^T$ and $\mathbf{V} = \mathbf{S}\mathbf{K}_z$, application of (33) to (30) yields

$$\hat{\boldsymbol{\gamma}} = \mathbf{S}^T(\mathbf{S}\mathbf{K}_z\mathbf{S}^T + \mu\mathbf{I})^{-1}\bar{\mathbf{m}}. \quad (34)$$

Subsequently, we reconstruct \mathbf{v} as

$$\hat{\mathbf{v}}_K = \mathbf{K}_z\mathbf{S}^T(\mathbf{S}\mathbf{K}_z\mathbf{S}^T + \mu\mathbf{I})^{-1}\bar{\mathbf{m}} \quad (35)$$

and we will henceforth refer to as the *Kronecker kernel MCEX* (KKMCEX) estimate of \mathbf{v} . Regarding the computational cost incurred by (34), inversion costs $\mathcal{O}(S^3)$, since the size of the matrix to be inverted is reduced from NL to S . Clearly, there is no need to compute $\mathbf{K}_z = \mathbf{K}_y \otimes \mathbf{K}_x$. As \mathbf{S} has binary entries, $\mathbf{S}\mathbf{K}_z\mathbf{S}^T$ is just a selection of S^2 entries in \mathbf{K}_z ; and, given that $\kappa_z((x_i, y_j), (x_n, y_l)) = \kappa_x(x_i, x_n)\kappa_y(y_j, y_l)$, it is obtained at cost $\mathcal{O}(S^2)$. Overall, the cost incurred by (34) is $\mathcal{O}(S^3)$. Compared to the MC approach in (16), the KKMCEX method is easier to implement since it only involves a matrix inversion. Moreover, since it admits a closed-form solution, it facilitates deriving bounds on the estimation error of $\hat{\mathbf{v}}_K$.

Remark 1: Matrices built via the Kronecker product have been used in regression for different purposes. Related to MC, [9] leverages Kronecker product structures to efficiently solve the Sylvester equations that arise in alternating minimization iterations to find $\{\hat{\mathbf{W}}, \hat{\mathbf{H}}\}$ in (16). On the other hand, [14], [29] propose a Kronecker kernel ridge regression method that can be used to extrapolate missing entries in a matrix. However, the methods in [14], [29] assume a complete training set and Kronecker structure for the regression matrix; this implies that the observed entries in \mathbf{M} can be permuted to form a full submatrix. In our formulation, we introduce \mathbf{S} which encompasses any sampling pattern in Ω . Thus, the properties of the Kronecker product used in [9], [14], [29] cannot be applied to solve (35) since $\mathbf{S}\mathbf{K}_z\mathbf{S}^T$ is not necessarily the Kronecker product of two smaller matrices.

Remark 2: The KKMCEX solution in (35), differs from that obtained as the solution of (17). On the one hand, the loss in (29) can be derived from the factorization-based one by using the Kronecker product kernel $\mathbf{K}_y \otimes \mathbf{K}_x$ and $\boldsymbol{\gamma} = \text{vec}(\mathbf{B}\mathbf{C}^T)$ to arrive at

$$\begin{aligned} & \left\| P_\Omega(\mathbf{M} - \mathbf{K}_x\mathbf{B}\mathbf{C}^T\mathbf{K}_y) \right\|_{\text{F}}^2 \\ &= \left\| \bar{\mathbf{m}} - \mathbf{S}(\mathbf{K}_y \otimes \mathbf{K}_x)\text{vec}(\mathbf{B}\mathbf{C}^T) \right\|_2^2. \end{aligned} \quad (36)$$

One difference between the two loss functions is that (29) does not explicitly limit the rank of the recovered matrix $\hat{\mathbf{F}} = \text{unvec}(\hat{\mathbf{v}}_K)$ since it has NL degrees of freedom through $\hat{\boldsymbol{\gamma}}$, while in (17) the rank of $\hat{\mathbf{F}}$ cannot exceed p since \mathbf{B} and \mathbf{C} are of rank p at most. In fact, the low-rank property is indirectly promoted in (29) through the kernel matrices. Since

$\text{rank}(\mathbf{F}) \leq \min(\text{rank}(\mathbf{K}_x), \text{rank}(\mathbf{K}_y))$, we can limit the rank of $\hat{\mathbf{F}}$ by selecting rank deficient kernels. On the other hand, the regularization terms in (17) and (29) play a different role in each formulation. The regularization in (17) promotes smoothness on the columns of the estimated factor matrices $\{\hat{\mathbf{W}}, \hat{\mathbf{H}}\}$; or, in other words, similarity between the rows of $\{\hat{\mathbf{W}}, \hat{\mathbf{H}}\}$ as measured by κ_x and κ_y . On the contrary, the regularization in (29) promotes smoothness on $\hat{\mathbf{v}}$, which is tantamount to promoting similarity between the entries of $\hat{\mathbf{F}}$ in accordance with κ_z .

A. KKMCEX Error Analysis

In order to assess the performance of KKMCEX we will rely on the mean-square error

$$\text{MSE} := \mathbb{E}_e\{\|\mathbf{v} - \hat{\mathbf{v}}_K\|_2^2\} \quad (37)$$

where $\mathbb{E}_e\{\cdot\}$ denotes the expectation with respect to \mathbf{e} . Before we proceed, we will outline Nyström's approximation.

Definition 1: Given a kernel matrix \mathbf{K} and a binary sampling matrix \mathbf{S} of appropriate dimensions, the Nyström approximation [30] of \mathbf{K} is $\mathbf{T} = \mathbf{K}\mathbf{S}^T(\mathbf{S}\mathbf{K}\mathbf{S}^T)^\dagger\mathbf{S}\mathbf{K}$, and the regularized Nyström approximation is

$$\tilde{\mathbf{T}} = \mathbf{K}\mathbf{S}^T(\mathbf{S}\mathbf{K}\mathbf{S}^T + \mu\mathbf{I})^{-1}\mathbf{S}\mathbf{K}. \quad (38)$$

Nyström's approximation is employed to reduce the complexity of standard kernel regression problems such as the one in (6). Instead of \mathbf{K} , the low-rank approximation \mathbf{T} is used to reduce the cost of inverting large-size matrices using the MIL [31]. While it is known that the best low-rank approximation to a matrix is obtained from its top eigenvectors, Nyström's approximation is cheaper. Using Def. 1, the following lemma provides the bias and variance of the KKMCEX estimator in (35):

Lemma 2: Given the kernel matrix \mathbf{K}_z and its regularized Nyström approximation $\tilde{\mathbf{T}}_z$ with $\mu > 0$, the MSE of the KKMCEX estimator is

$$\text{MSE} = \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z)\boldsymbol{\gamma} \right\|_2^2 + \mathbb{E}_e \left\{ \frac{1}{\mu^2} \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z)\mathbf{S}^T\bar{\mathbf{e}} \right\|_2^2 \right\} \quad (39)$$

where the first term accounts for the bias and the second term accounts for the variance.

Lemma 2 shows that the MSE of the KKMCEX can be expressed in terms of $\tilde{\mathbf{T}}_z$; see proof in the Appendix. Knowing that the 2-norm satisfies $\|\mathbf{A}\|_2 \leq \|\mathbf{A}\|_{\text{F}}$, we have

$$\begin{aligned} & \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z)\boldsymbol{\gamma} \right\|_2^2 + \mathbb{E}_e \left\{ \frac{1}{\mu^2} \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z)\mathbf{S}^T\bar{\mathbf{e}} \right\|_2^2 \right\} \\ & \leq \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z) \right\|_{\text{F}}^2 \left(\|\boldsymbol{\gamma}\|_2^2 + \mathbb{E}_e \left\{ \frac{1}{\mu^2} \left\| \mathbf{S}^T\bar{\mathbf{e}} \right\|_2^2 \right\} \right). \end{aligned} \quad (40)$$

Consequently, the upper bound on the MSE is proportional to the approximation error of $\tilde{\mathbf{T}}_z$ to \mathbf{K}_z . This suggests selecting $\{m_{i,j}\}_{(i,j) \in \Omega}$ so that this approximation error is minimized; see also [31] where Ω is chosen according to the so-called leverage scores of \mathbf{K}_z in order to minimize the regression error. The next theorem uses Lemma 2 to upper bound the MSE in (37); see the Appendix for its proof.

Theorem 2: Let σ_{NL} be the maximum eigenvalue of a non-singular \mathbf{K}_z , and $\tilde{\gamma} := \mathbf{L}^T \gamma$, where \mathbf{L} is the eigenvector matrix of $\mathbf{K}_z - \tilde{\mathbf{T}}_z$. If \mathbf{e} is a zero-mean vector of iid Gaussian random variables with covariance matrix $\nu^2 \mathbf{I}$, the MSE of the KKMCEX estimator is bounded as

$$\text{MSE} \leq \frac{\mu^2 \sigma_{NL}^2}{(\sigma_{NL} + \mu)^2} \sum_{i=1}^S \tilde{\gamma}_i^2 + \sigma_{NL}^2 \sum_{i=S+1}^{NL} \tilde{\gamma}_i^2 + \frac{S \nu^2 \sigma_{NL}^2}{\mu^2}. \quad (41)$$

Considering the right-hand side of (41), the first two terms correspond to the bias, while the last term is related to the variance. We observe that when \mathbf{M} is fully observed, that is, $S = NL$, the bias can be made arbitrarily small by having $\mu \rightarrow 0$. It is also of interest to assess how the MSE bound behaves as S increases. Considering $\mu = S\mu'$ and fixed values in $(0, \infty)$ for μ' , $\|\tilde{\gamma}_i\|^2$ and σ_{NL}^1 , the bias term reduces to

$$\frac{S^2 \mu'^2 \sigma_{NL}^2}{(\sigma_{NL} + S\mu')^2} \sum_{i=1}^S \tilde{\gamma}_i^2 + \sigma_{NL}^2 \sum_{i=S+1}^{NL} \tilde{\gamma}_i^2. \quad (42)$$

We observe in (42) that as S increases, terms move from the second summation to the first. Therefore, whether the bias term grows or diminishes depends on the multiplication factors in front of the two summations. Since $\frac{S^2 \mu'^2}{(\sigma_{NL} + S\mu')^2} \leq 1$ the bias term in (42) decreases with S . On the other hand, the variance term becomes $\frac{\nu^2 \sigma_{NL}^2}{S \mu'^2}$ and decays with S as well. As a result, the MSE bound in Theorem 2 decays up until $S = NL$.

V. RIDGE REGRESSION MCEX

Although the KKMCEX algorithm is fast when S is small, the size of the matrix to be inverted in (34) grows with S , hence increasing the computational cost. Available approaches to reducing the computational cost of kernel regression methods are centered around the idea of approximating the kernel matrix. For instance, [31] uses Nyström's approximation, that our performance analysis in Section IV was based on, whereas [32] relies on a sketch of \mathbf{K}_z formed by a subset of its columns, hence reducing the number of regression coefficients; see also [33], where the kernel function is approximated by the inner product of random finite-dimensional feature maps, which also speeds up the matrix inversion. In this section, we reformulate the KKMCEX of Section IV to incorporate a low-rank approximation of \mathbf{K}_z in order to obtain a reduced complexity estimate for \mathbf{v} . Moreover, we also develop an online method based on this reformulation.

Recall from Eq. (4) that a kernel can be viewed as the inner product of vectors mapped to a feature space \mathcal{F}_z , namely $\kappa_z((x_i, y_j), (x_n, y_l)) = \langle \phi_z(x_i, y_j), \phi_z(x_n, y_l) \rangle_{\mathcal{F}_z}$. Let $\tilde{\phi}_z : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ be a feature map approximating κ_z so that

$$\kappa_z((x_i, y_j), (x_n, y_l)) \simeq \langle \tilde{\phi}_z(x_i, y_j), \tilde{\phi}_z(x_n, y_l) \rangle. \quad (43)$$

Then, we define the $NL \times d$ feature matrix $\tilde{\mathbf{\Phi}}_z := [\tilde{\phi}_z(x_1, y_1), \tilde{\phi}_z(x_2, y_1), \dots, \tilde{\phi}_z(x_N, y_L)]^T$ and form $\tilde{\mathbf{K}}_z = \tilde{\mathbf{\Phi}}_z \tilde{\mathbf{\Phi}}_z^T$.

¹Note that $\|\tilde{\gamma}_i\|^2$ and σ_{NL} depend on the selected kernel \mathbf{K}_z and matrix \mathbf{F} , and do not depend on \mathbf{S} .

Note that $\tilde{\mathbf{K}}_z$ is a rank- d approximation of \mathbf{K}_z , and that the equality $\mathbf{K}_z = \tilde{\mathbf{K}}_z$ is only feasible when $\text{rank}(\mathbf{K}_z) \leq d$. Consider $\tilde{\mathbf{\Phi}}_x = [\tilde{\phi}_x(x_1), \dots, \tilde{\phi}_x(x_N)]$ and $\tilde{\mathbf{\Phi}}_y = [\tilde{\phi}_y(y_1), \dots, \tilde{\phi}_y(y_L)]$, where $\tilde{\phi}_x : \mathcal{X} \rightarrow \mathbb{R}^{d_x}$ and $\tilde{\phi}_y : \mathcal{Y} \rightarrow \mathbb{R}^{d_y}$, as the feature matrices forming low-rank approximations to \mathbf{K}_x and \mathbf{K}_y , respectively. Since $\mathbf{K}_z = \mathbf{K}_y \otimes \mathbf{K}_x$ in KKMCEX, a prudent choice is $\tilde{\mathbf{\Phi}}_z = \tilde{\mathbf{\Phi}}_y \otimes \tilde{\mathbf{\Phi}}_x$. In the next section we will present means of constructing $\{\tilde{\mathbf{\Phi}}_x, \tilde{\mathbf{\Phi}}_y, \tilde{\mathbf{\Phi}}_z\}$ maps.

Since $\tilde{\mathbf{K}}_z$ is a valid kernel matrix, upon replacing \mathbf{K}_z in (28) with $\tilde{\mathbf{K}}_z$, the observation model reduces to

$$\tilde{\mathbf{m}} = \mathbf{S} \tilde{\mathbf{\Phi}}_z \tilde{\mathbf{\Phi}}_z^T \gamma + \tilde{\mathbf{e}}, \quad (44)$$

where $\tilde{\mathbf{e}} = \mathbf{e} + \mathbf{S}(\mathbf{K}_z - \tilde{\mathbf{K}}_z)\gamma$. With this model, the weights in (29) are obtained as

$$\hat{\gamma} = \arg \min_{\gamma \in \mathbb{R}^{NL}} \left\| \tilde{\mathbf{m}} - \mathbf{S} \tilde{\mathbf{\Phi}}_z \tilde{\mathbf{\Phi}}_z^T \gamma \right\|_2^2 + \mu \gamma^T \tilde{\mathbf{\Phi}}_z \tilde{\mathbf{\Phi}}_z^T \gamma. \quad (45)$$

Letting $\xi := \tilde{\mathbf{\Phi}}_z^T \gamma$ and substituting into (45), we arrive at

$$\hat{\xi} = \arg \min_{\xi \in \mathbb{R}^d} \left\| \tilde{\mathbf{m}} - \mathbf{S} \tilde{\mathbf{\Phi}}_z \xi \right\|_2^2 + \mu \|\xi\|_2^2 \quad (46)$$

which admits the closed-form solution

$$\hat{\xi} = (\tilde{\mathbf{\Phi}}_z^T \mathbf{S}^T \mathbf{S} \tilde{\mathbf{\Phi}}_z + \mu \mathbf{I})^{-1} \tilde{\mathbf{\Phi}}_z^T \mathbf{S}^T \tilde{\mathbf{m}}. \quad (47)$$

Using $\hat{\xi}$, we obtain $\hat{\mathbf{v}}_R = \tilde{\mathbf{\Phi}}_z \hat{\xi}$ as the *ridge regression MCEX* (RRMCEX) estimate. Using the MIL (33) on (47), it follows that

$$\hat{\xi} = \tilde{\mathbf{\Phi}}_z^T \mathbf{S}^T (\mathbf{S} \tilde{\mathbf{\Phi}}_z \tilde{\mathbf{\Phi}}_z^T \mathbf{S}^T + \mu \mathbf{I})^{-1} \tilde{\mathbf{m}} \quad (48)$$

and thus,

$$\hat{\mathbf{v}}_R = \tilde{\mathbf{\Phi}}_z \hat{\xi} = \tilde{\mathbf{K}}_z \mathbf{S}^T (\mathbf{S} \tilde{\mathbf{K}}_z^T \mathbf{S}^T + \mu \mathbf{I})^{-1} \tilde{\mathbf{m}}. \quad (49)$$

Therefore, (49) shows that $\hat{\mathbf{v}}_R$ is equivalent to the KKMCEX solution $\hat{\mathbf{v}}_K$ in (35) after replacing \mathbf{K}_z by its low-rank approximation $\tilde{\mathbf{K}}_z$. For error-free approximation, $\mathbf{K}_z = \tilde{\mathbf{\Phi}}_z \tilde{\mathbf{\Phi}}_z^T$, while $\hat{\xi}$ in (47) can be viewed as the primal solution to the optimization problem in (46), and $\hat{\gamma}$ in (34) as its dual [21]. Still, obtaining $\hat{\xi}$ requires multiplying two $d \times S$ matrices and inverting a $d \times d$ matrix, which incurs computational cost $\mathcal{O}(d^2 S)$ when $S \geq d$, and $\mathbf{S} \tilde{\mathbf{\Phi}}_z$ is obtained at cost $\mathcal{O}(dS)$. Thus, the cost of RRMCEX grows linearly with S in contrast to KKMCEX that increases with S^3 .

By choosing an appropriate feature map so that $d \ll S$, it is possible to control the computational cost of calculating $\hat{\xi}$. However, reduced computational cost by selecting a small d might come at the price of an approximation error to \mathbf{K}_z , which correspondingly increases the estimation error of $\hat{\mathbf{v}}_R$. The selection of a feature matrix to minimize this error and further elaboration on the computational cost are given in Section VI.

A. Online RRMCEX

Online methods learn a model by processing one datum at a time. An online algorithm often results when the objective can be separated into several subfunctions, each depending on one or multiple data. In the context of MC, online implementation

updates $\hat{\mathbf{F}}$ every time a new entry $M_{i,j}$ becomes available. If we were to solve (34) each time a new observation was becoming available, inverting an $S \times S$ matrix per iteration would result in an overall prohibitively high computational cost. Still, the cost of obtaining an updated solution per observation can stay manageable using online kernel regression solvers that fall into three categories [34]: dictionary learning, recursive regression and stochastic gradient descent based. Akin to [35], [36], we will pursue here the SGD.

Consider rewriting (46) entrywise as

$$\hat{\xi} = \arg \min_{\xi \in \mathbb{R}^d} \sum_{(i,j) \in \Omega} \left[m_{i,j} - \tilde{\phi}_z^T(x_i, y_j) \xi \right]^2 + \mu \|\xi\|_2^2. \quad (50)$$

With n denoting each scalar observation, SGD iterations form a sequence of estimates

$$\hat{\xi}^n = \hat{\xi}^{n-1} - t_n \left[\tilde{\phi}_z(x_i, y_j) (\tilde{\phi}_z^T(x_i, y_j) \hat{\xi}^{n-1} - m_{i,j}) + \mu \hat{\xi}^{n-1} \right] \quad (51)$$

where t_n is the step size, $n = 1, \dots, S$ and the tuple (i, j) denotes the indices of the entry revealed at iteration n . With properly selecting t_n , the sequence $\hat{\xi}^n$ will converge to (50) at per iteration cost $\mathcal{O}(d)$ [37]. Apart from updating all entries in the matrix, (51) can also afford a simple distributed implementation using e.g., the algorithms in [38].

Remark 3: Online algorithms for MC can be designed to solve the factorization-based formulation from (11) rewritten as

$$\begin{aligned} \arg \min_{\substack{\mathbf{W} \in \mathbb{R}^{N \times p} \\ \mathbf{H} \in \mathbb{R}^{L \times p}} \sum_{(i,j) \in \Omega} & \left((m_{i,j} - \mathbf{w}_i^T \mathbf{h}_j)^2 \right. \\ & \left. + \frac{\mu}{|\Omega_i^w|} \|\mathbf{w}_i\|_2^2 + \frac{\mu}{|\Omega_j^h|} \|\mathbf{h}_j\|_2^2 \right) \end{aligned} \quad (52)$$

where \mathbf{w}_i^T and \mathbf{h}_j^T denote the i th and j th rows of \mathbf{H} and \mathbf{W} respectively, $\Omega_i^w = \{j : (i, j) \in \Omega\}$, and $\Omega_j^h = \{i : (i, j) \in \Omega\}$. When $m_{i,j}$ becomes available, algorithms such as SGD and online ALS update the rows $\{\mathbf{w}_i^T, \mathbf{h}_j^T\}$ of the coefficient matrices. This procedure can also be applied to the kernel MCEX formulation in (16), that solves for \mathbf{W} and \mathbf{H} , although the rows $\{\mathbf{w}_i^T, \mathbf{h}_j^T\}$ cannot be updated independently due to the involvement of the kernel matrices [13]. Then, all entries in the i th row and j th column of $\hat{\mathbf{F}}$ are also updated per iteration, as opposed to our method which updates the whole matrix.

VI. CHOOSING THE KERNEL MATRICES

In this section, we provide pointers on how to build matrices \mathbf{K}_z for KKMCEX and $\tilde{\Phi}_z$ for RRMCEX when prior information about either the matrix \mathbf{F} , or the input spaces \mathcal{X} and \mathcal{Y} , is available.

A. Kernels Based on the Graph Laplacian

Suppose that the columns and rows of \mathbf{F} lie on a graph, that is, each entry of a column or row vector is associated with a node on a graph that encodes the interdependencies with entries in the same vector. Specifically, we define an undirected weighted

graph $\mathcal{G}_x = (\mathcal{X}, \mathcal{E}_x, \mathbf{A}_x)$ for the columns of \mathbf{F} , where \mathcal{X} is the set of vertices with $|\mathcal{X}| = N$, $\mathcal{E}_x \subseteq \mathcal{X} \times \mathcal{X}$ is the set of edges connecting the vertices, and $\mathbf{A}_x \in \mathbb{R}^{N \times N}$ is a weighted adjacency matrix. Then, functions $\{f_l : \mathcal{X} \rightarrow \mathbb{R}\}_{l=1}^L$ are what is recently referred to as a graph signal [39], that is, a map from the set \mathcal{X} of vertices into the set of real numbers. Likewise, we define a graph $\mathcal{G}_y = (\mathcal{Y}, \mathcal{E}_y, \mathbf{A}_y)$ for the rows of \mathbf{F} , i.e., $\{g_n : \mathcal{Y} \rightarrow \mathbb{R}\}_{n=1}^N$, which are also graph signals. In a matrix of user-movie ratings for instance, we would have two graphs: one for the users and one for the movies. The graphs associated with the columns and rows yield the underlying structure of \mathbf{F} that can be used to generate a pair of kernels.

Using \mathbf{A}_x and \mathbf{A}_y , we can form the corresponding graph Laplacian as $\mathbf{L}_x := \text{diag}(\mathbf{A}_x \mathbf{1}) - \mathbf{A}_x$ and likewise for \mathbf{L}_y , that can serve as kernels. A family of graphical kernels results using a monotonic inverse function $r^\dagger(\cdot)$ on the Laplacian eigendecomposition as [40]

$$\mathbf{K} = \mathbf{Q} r^\dagger(\boldsymbol{\Lambda}) \mathbf{Q}^T. \quad (53)$$

A possible choice of $r(\cdot)$ is the Gaussian radial basis function, which generates the diffusion kernel $r(\lambda_i) = e^{\eta \lambda_i}$, where λ_i is the i th eigenvalue of \mathbf{L} , and η a weight parameter. Alternatively, one can choose just the linear function $r(\lambda_i) = 1 + \eta \lambda_i$, which results in the regularized Laplacian kernel. By applying different weighting functions to the eigenvalues of \mathbf{L}_x and \mathbf{L}_y , we promote smoother or more rapidly changing functions for the columns and rows of $\hat{\mathbf{F}}$ [41]. While \mathbf{K}_x and \mathbf{K}_y are chosen as Laplacian kernels, this would not be the case for $\mathbf{K}_z = \mathbf{K}_y \otimes \mathbf{K}_x$ used in our KKMCEX context since it does not result from applying $r^\dagger(\cdot)$ to a Laplacian matrix. However, since $\mathbf{K}_x = \mathbf{Q}_x \boldsymbol{\Sigma}_x \mathbf{Q}_x^T$ and $\mathbf{K}_y = \mathbf{Q}_y \boldsymbol{\Sigma}_y \mathbf{Q}_y^T$, the eigendecomposition of \mathbf{K}_z is $\mathbf{K}_z = (\mathbf{Q}_y \otimes \mathbf{Q}_x) (\boldsymbol{\Sigma}_y \otimes \boldsymbol{\Sigma}_x) (\mathbf{Q}_y^T \otimes \mathbf{Q}_x^T)$, and the notions of frequency and smoothness carry over. In other words, we are still promoting similarity among entries that are connected on the row and columns graphs through \mathbf{K}_z .

A key attribute in graph signal processing is that of ‘‘graph bandlimitedness’’, which arises when a signal can be generated as a linear combination of a few eigenvectors of the Laplacian matrix. Therefore, a bandlimited graph signal belongs to an RKHS that is spanned by a bandlimited kernel [23] that suppresses some of the frequencies of the graph. A bandlimited kernel is derived from the Laplacian matrix of a graph as in (53), using

$$r(\lambda_i) = 0 \quad \forall i \notin \Psi, \quad (54)$$

where $\Psi \subseteq \mathbb{N}$ contains the indices of frequencies not to be suppressed. As mentioned earlier, we define a graph for the columns and a graph for the rows of \mathbf{F} . Therefore, graph signals contained in the columns and rows may be bandlimited with different bandwidths. In order to form \mathbf{K}_z our KKMCEX approach, we will need to apply different weighting functions akin to the one in (54) for kernel matrices \mathbf{K}_x and \mathbf{K}_y .

B. Kernels From Known Basis or Features

In some applications the basis that spans the columns or rows of the unobserved matrix is assumed known, although this basis

matrix needs not be a kernel. In order to be able to include such basis into the kernel framework, we need to generate kernel functions that span the same spaces as the columns and rows of \mathbf{F} .

Consider the input sets $\{\mathcal{X}, \mathcal{Y}\}$ whose entries can be mapped into an Euclidean space through feature extraction functions $\theta_x : \mathcal{X} \rightarrow \mathbb{R}^{t_x}$ and $\theta_y : \mathcal{Y} \rightarrow \mathbb{R}^{t_y}$ such that $\theta_x(x_i) := \mathbf{x}_i$ and $\theta_y(y_j) := \mathbf{y}_j$. For instance, in a movie recommender system where the users are represented in \mathcal{X} and the movies in \mathcal{Y} , each coordinate of \mathbf{y}_j could denote the amount of action, drama and nudity in the movie, and \mathbf{x}_i would contain weights according to the user's preference for each attribute. We may then use the feature vectors to determine the similarities among entries in \mathcal{X} and \mathcal{Y} by means of kernel functions.

Let $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_N]^T$ and $\mathbf{Y} := [\mathbf{y}_1, \dots, \mathbf{y}_L]^T$. If $\text{span}(\mathbf{F}) \subseteq \text{span}(\mathbf{X})$ and $\text{span}(\mathbf{F}^T) \subseteq \text{span}(\mathbf{Y})$, we may conveniently resort to the linear kernel. The linear kernel amounts to the dot product in Euclidean spaces, which we use to define the pair $\kappa_x(x_i, x_j) = \mathbf{x}_i^T \mathbf{x}_j$ and $\kappa_y(y_i, y_j) = \mathbf{y}_i^T \mathbf{y}_j$. This leads to a straightforward construction of the kernel matrices for KKMCEX as $\mathbf{K}_x = \mathbf{X}\mathbf{X}^T$ and $\mathbf{K}_y = \mathbf{Y}\mathbf{Y}^T$.

Besides the linear kernel, it is often necessary to use a different kernel class for each κ_x and κ_y chosen to better fit the spaces spanned by the rows and columns of \mathbf{F} . For instance, the Gaussian kernel defined as $\kappa_x(x_i, x_j) = \exp\{-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / (2\eta)\}$, is a widely used alternative in the regression of smooth functions.

C. Feature Maps for RRMCEX

Aiming to construct $\tilde{\Phi}_z$ that approximates \mathbf{K}_z at reduced complexity, we choose $\tilde{\phi}_z$ with $d \ll S$. To approximate linear kernels, let $\tilde{\phi}_x(x_i) = \mathbf{x}_i$ and $\tilde{\phi}_y(y_j) = \mathbf{y}_j$ so that we can set $\tilde{\phi}_z(x_i, y_j) = \tilde{\phi}_y(y_j) \otimes \tilde{\phi}_x(x_i)$ and $\tilde{\Phi}_z = \mathbf{Y} \otimes \mathbf{X}$. Note that in this case $\tilde{\Phi}_z \tilde{\Phi}_z^T$ yields a zero-error approximation to $\mathbf{K}_z = (\mathbf{Y} \otimes \mathbf{X})(\mathbf{Y} \otimes \mathbf{X})^T$, which renders the KKMCEX and RRMCEX solutions equivalent.

On occasion, \mathbf{X} and \mathbf{Y} may have large column dimension, thus rendering $\mathbf{Y} \otimes \mathbf{X}$ undesirable as a feature matrix in RRMCEX. In order to overcome this hurdle, we build an approximation to the column space of $\mathbf{Y} \otimes \mathbf{X}$ from the SVD of \mathbf{X} and \mathbf{Y} . Consider the SVDs of matrices $\mathbf{X} = \mathbf{U}_x \mathbf{D}_x \mathbf{V}_x^T$ and $\mathbf{Y} = \mathbf{U}_y \mathbf{D}_y \mathbf{V}_y^T$, to obtain $\mathbf{Y} \otimes \mathbf{X} = (\mathbf{U}_y \otimes \mathbf{U}_x)(\mathbf{D}_y \otimes \mathbf{D}_x)(\mathbf{V}_y^T \otimes \mathbf{V}_x^T)$. Let $\tilde{\Phi}_z = \mathbf{U}_d \mathbf{D}_d$, where \mathbf{U}_d and \mathbf{D}_d respectively hold the top d singular vectors and singular values of $\mathbf{Y} \otimes \mathbf{X}$. The SVD has cost $\mathcal{O}(Nt_x^2)$ for \mathbf{X} and $\mathcal{O}(Lt_y^2)$ for \mathbf{Y} . Comparatively, the cost of building \mathbf{K}_x and \mathbf{K}_y for the linear kernel is $\mathcal{O}(N^2t_x)$ and $\mathcal{O}(L^2t_y)$, respectively. Therefore, choosing RRMCEX over KKMCEX in this case incurs no extra overhead.

When a function other than the linear kernel is selected, obtaining an approximation is more complex. To approximate a Gaussian kernel on $\mathcal{X} \times \mathcal{X}$, the vectors $\{\tilde{\phi}_x(\mathbf{x}_i)\}_{i=1}^N$ can be obtained by means of Taylor series expansion [42] or random Fourier features [33], which can also approximate Laplacian, Cauchy and polynomial kernels [33], [43]. Therefore, the maps $\tilde{\phi}_x$ and $\tilde{\phi}_y$ must be designed according to the chosen kernels.

In some instances, such as when dealing with Laplacian kernels, \mathbf{X} and \mathbf{Y} are not available and we are only given \mathbf{K}_x and \mathbf{K}_y . We are then unable to derive approximations to the kernel matrices by means of maps $\tilde{\phi}_x$ and $\tilde{\phi}_y$. Nevertheless, we can still derive an adequate $\tilde{\Phi}_z$ to approximate \mathbf{K}_z . Indeed, Mercer's Theorem asserts that there are eigenfunctions $\{q_n\}_{n=1}^{NL}$ in \mathcal{H}_z along with a sequence of nonnegative real numbers $\{\sigma_n\}_{n=1}^{NL}$, such that

$$\kappa_z((x_i, y_j), (x_n, y_l)) = \sum_{n=1}^{NL} \sigma_n q_n(x_i, y_j) q_n(x_n, y_l). \quad (55)$$

We can find (55) from the eigendecomposition $\mathbf{K}_z = \mathbf{Q}_z \Sigma_z \mathbf{Q}_z^T$, where q_n is the n th eigenvector in \mathbf{Q}_z and σ_n the n th eigenvalue in Σ_z . If \mathbf{K}_z is low rank, we can construct $\tilde{\Phi}_z = \mathbf{Q}_d \Sigma_d^{\frac{1}{2}}$, where \mathbf{Q}_d and Σ_d respectively hold the top d eigenvectors and eigenvalues of \mathbf{K}_z . Note that, since $\mathbf{K}_z = (\mathbf{Q}_y \otimes \mathbf{Q}_x)(\Sigma_y \otimes \Sigma_x)(\mathbf{Q}_y^T \otimes \mathbf{Q}_x^T)$, we only need to eigendecompose smaller matrices \mathbf{K}_x and \mathbf{K}_y at complexity $\mathcal{O}(N^3 + L^3)$. In some cases however, such as when using Laplacian kernels, the eigendecompositions are readily available, and $\tilde{\Phi}_z$ can be built at a markedly reduced cost.

VII. NUMERICAL TESTS

In this section, we test the performance of the KKMCEX, RRMCEX and online (o)RRMCEX algorithms developed in Sections IV, V and V-A, respectively; and further compare them to the solution of (16) obtained with ALS [19] and SGD [13]. We run the tests on synthetic and real datasets, with and without noise, and measure the signal-to-noise-ratio (SNR) as $\frac{\|\mathbf{F}\|_F^2}{\|\mathbf{E}\|_F^2}$. The algorithms are run until convergence over $N_r = 50$ realizations with different percentages of observed entries, denoted by $P_s = 100S/(NL)$, which are taken uniformly at random per realization. As figure of merit, we use

$$\text{NMSE} = \frac{1}{N_r} \sum_{i=1}^{N_r} \frac{\|\hat{\mathbf{F}}_i - \mathbf{F}\|_F^2}{\|\mathbf{F}\|_F^2} \quad (56)$$

where $\hat{\mathbf{F}}_i$ is the estimate at realization i . We show results for the optimal combination of regularization and kernel parameters, found via grid search. Finally, ALS and SGD are initialized by a product of two random factor matrices.

A. Synthetic Data

We first test the algorithms on synthetic data. The 250×250 data matrix is generated as $\mathbf{F} = \mathbf{K}_x \mathbf{\Gamma} \mathbf{K}_y$, where $\mathbf{\Gamma}$ is a 250×250 matrix of Gaussian random deviates. For \mathbf{K}_x and \mathbf{K}_y we use Laplacian diffusion kernels with $\eta = 1$ based on Erdős-Rényi graphs, whose binary adjacency matrices are unweighted and any two vertices are connected with probability 0.03. The resulting \mathbf{F} is approximately low-rank, with the sum of the first 10 eigenvalues accounting for 96% of the total eigenvalue sum. Therefore, we set the rank bound p to 10 for the ALS and SGD algorithms. Whether \mathbf{F} is approximately low-rank or exactly

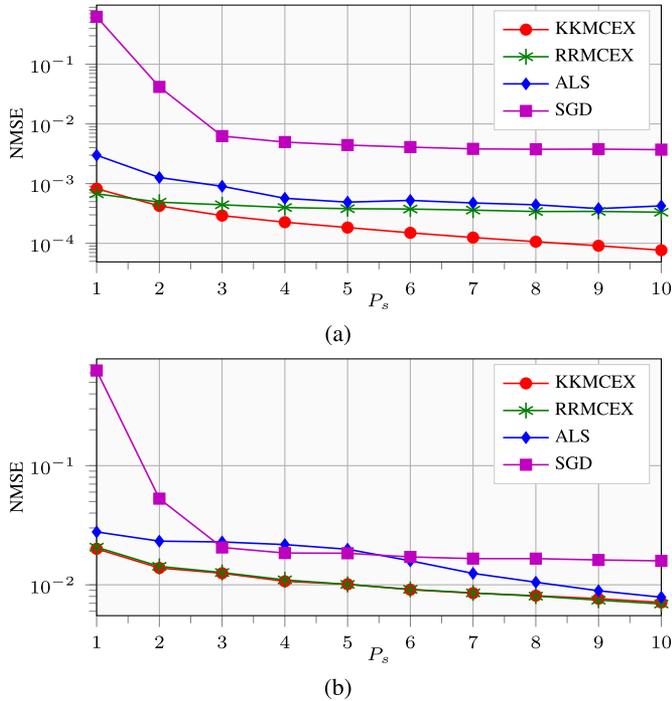


Fig. 1. NMSE vs. P_s for (a) synthetic noiseless matrix; and (b) synthetic noisy matrix.

low-rank did not affect our results, as they were similar for matrices with an exact rank of 10. For KKMCEX, $\mathbf{K}_z = \mathbf{K}_y \otimes \mathbf{K}_x$, and for RRMCEX $\tilde{\Phi}_z = \mathbf{Q}_d \Sigma_d^{\frac{1}{2}}$, where \mathbf{Q}_d contains the top 250 eigenvectors of \mathbf{K}_z , and Σ_d the corresponding top 250 eigenvalues.

Fig. 1 shows the simulated NMSE when \mathbf{M} is noiseless (a) or noisy (b). We deduce from Fig. 1a that all algorithms except SGD achieve a very small NMSE, below 0.003 at $P_s = 1\%$ that falls to 0.007 at $P_s = 10\%$. Of the three algorithms, KKMCEX has the smallest error except at $P_s = 1\%$, where RRMCEX performs best. Although the error drops below 0.005 for SGD at $P_s > 4\%$, it is outperformed by the other algorithms by an order of magnitude. Fig. 1b shows the same results when Gaussian noise is added to \mathbf{F} at $\text{snr} = 1$. We observe that KKMCEX and RRMCEX are matched and attain the lowest error, whereas ALS and SGD have larger errors across P_s . This corroborates that thanks to the regularization term that smoothes over all the entries instead of row or column-wise, the noise effect is reduced. Interestingly, RRMCEX is able to reduce the noise effect despite the bias it suffers because it only uses the top 250 eigenvalues of \mathbf{K}_z from a total of 62,500. This is mainly due to the additive noise being evenly distributed across the eigenspace of \mathbf{K}_z . Therefore, by keeping only the eigenvectors associated with the top 250 eigenvalues in \mathbf{K}_z , we are discarding those dimensions in which the SNR is lower.

Fig. 2 depicts the time needed for the algorithms to perform the simulations reported in Fig. 1. We observe in Fig. 2a that RRMCEX has an almost constant computation time, whereas the time for KKMCEX grows with P_s as expected since the size of the matrix to be inverted increases with S . On the other

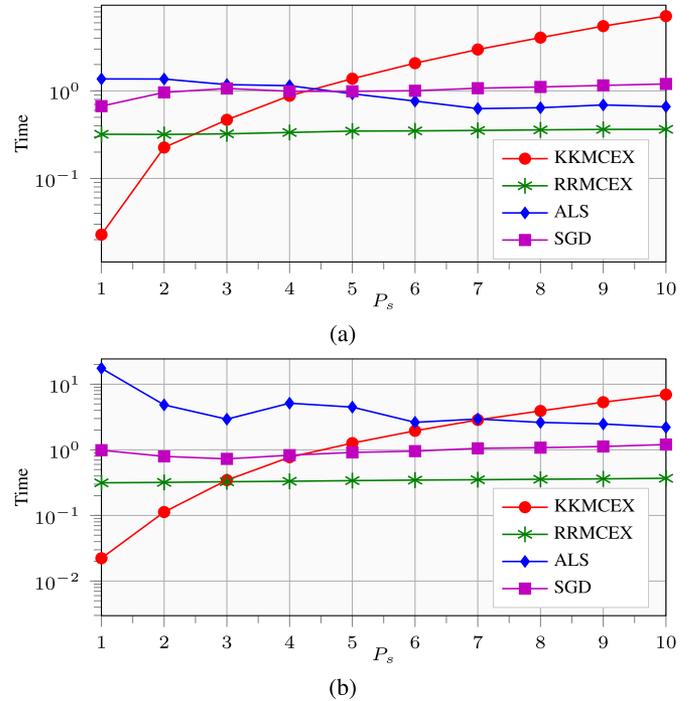


Fig. 2. Time vs. P_s for (a) synthetic noiseless matrix; and (b) synthetic noisy matrix.

hand, ALS and SGD require less time than KKMCEX for the larger values of P_s , but are always outperformed by RRMCEX. Moreover, the ALS time is reduced as P_s increases because the number of iterations required to converge to the minimum is smaller. Fig. 2b suggests that the noise only impacts ALS, which has its computation time rise considerably across all P_s . Overall, Figs. 1 and 2 illustrate that RRMCEX has the best performance for the synthetic matrix both in terms of NMSE and computational cost.

B. Temperature Measurements

In this case, \mathbf{F} has size 150×365 comprising temperature readings taken by 150 stations over 365 days in 2002 in the United States.² The columns and rows of \mathbf{F} are modeled as graph signals with \mathbf{A}_x and \mathbf{A}_y for the graphs formed by the stations and the days of the year, respectively. We use Laplacian diffusion kernels both for \mathbf{K}_x and \mathbf{K}_y , while \mathbf{K}_z and $\tilde{\Phi}_z$ are obtained as in the tests on synthetic data, except that $\tilde{\Phi}_z$ is constructed with the top 150 eigenvectors of \mathbf{K}_z . The matrix \mathbf{A}_x is obtained as in [8], where a graph \mathcal{G}_x with unweighted adjacency matrix \mathbf{P} is generated for the stations, and each station is a vertex connected to the 8 geographically closest stations. Next, we obtain the undirected graph \mathcal{G}'_x with symmetric adjacency matrix $\mathbf{P}' = \text{sign}(\mathbf{P}^T + \mathbf{P})$. Finally, the entries of \mathbf{A}_x are constructed as $(\mathbf{A}_x)_{i,j} = \exp(-\frac{N^2 d_{i,j}}{\sum_{i,j} d_{i,j}})$, where $\{d_{i,j}\}$ are geodesic distances on \mathcal{G}'_x . We adopt a graph on which each day is a vertex and each day is connected to the 10 past and future days to form \mathbf{A}_y .

²<http://earthpy.org/ulmo.html>

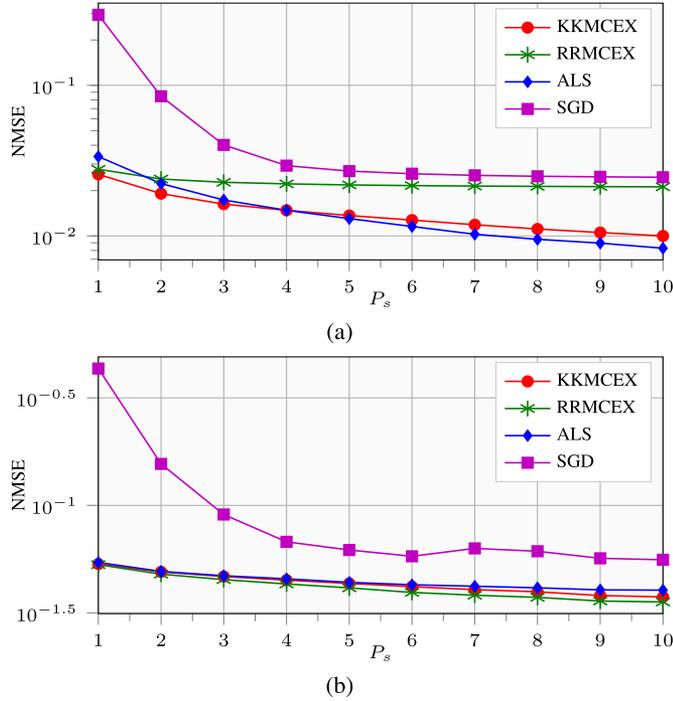


Fig. 3. NMSE vs. P_s for the matrix of temperature measurements (a) without noise, and (b) with noise.

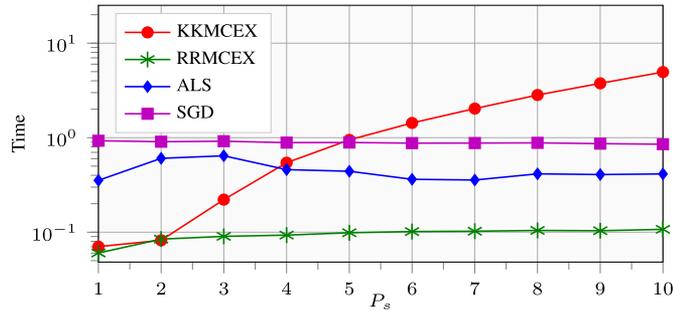


Fig. 4. Time vs. P_s for the noiseless matrix of temperature measurements.

Fig. 3 shows the simulated tests for (a) the matrix of temperature readings, and (b) the same matrix with additive Gaussian noise at $\text{snr} = 1$. Fig. 3a demonstrates that KKMCEX achieves the lowest error for the first three data points, while afterwards ALS has a slight edge over KKMCEX. The real data matrix \mathbf{F} is approximately low rank, since the sum of the first 10 singular values accounts for 75% of the total sum. This explains why RRMCEX fares worse than KKMCEX. Because it only contains the top 150 eigenvectors of \mathbf{K}_z , which is full rank, the vectorized data \mathbf{m} lies in part outside the space spanned by $\tilde{\Phi}_z$. Indeed, increasing the number of eigenvectors in $\tilde{\Phi}_z$ results in a lower error, although the computational cost increases accordingly. Fig. 3b further demonstrates that the addition of noise has the least impact on RRMCEX, which attains the lowest error slightly below KKMCEX. On the other hand, ALS has a marginally higher error whereas the gap between SGD and the other three methods remains. Fig. 4 depicts the computational time for the

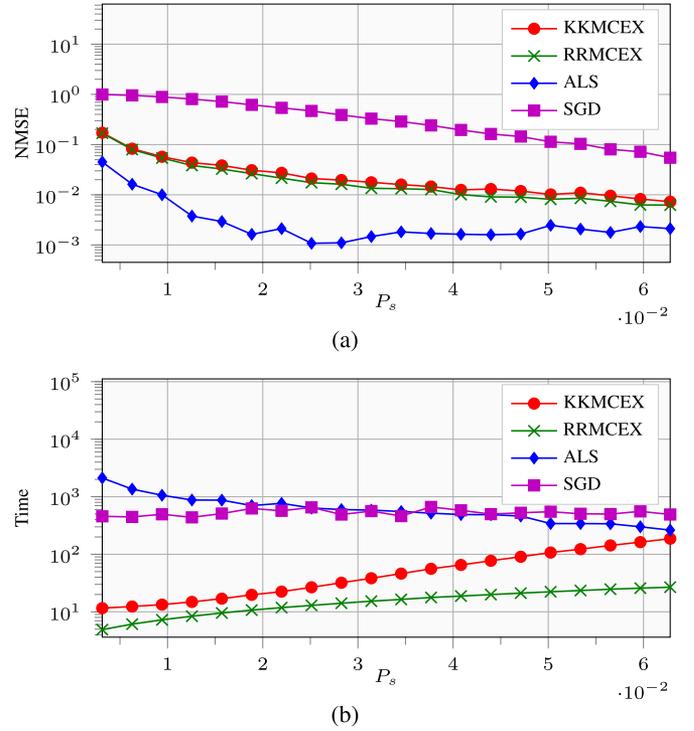


Fig. 5. Results for the mushroom adjacency matrix as (a) NMSE vs. P_s , and (b) time vs. P_s .

results in Fig. 3a, which are similar to those obtained for the synthetic dataset.

C. Mushroom Dataset

The Mushroom dataset³ comprises 8,124 labels and as many feature vectors. Each label indicates whether a sample is edible or poisonous, and each vector has 22 entries describing the shape, color, etc. of the mushroom sample. After removing items with missing features, we are left with 5,643 labels and feature vectors. Here, we solve a clustering problem in which \mathbf{F} is a $5,643 \times 5,643$ adjacency matrix, where $\mathbf{F}_{i,j} = 1$ if the i th and the j th mushroom samples belong to the same class (poisonous or edible), and $\mathbf{F}_{i,j} = -1$ otherwise.

We encode the matrix stacking the feature vectors via one-hot encoding to produce a $5,643 \times 98$ binary feature matrix analogous to \mathbf{X} in Section VI-B. We build the kernel matrix \mathbf{K}_x from the Pearson correlation coefficients of the rows of \mathbf{X} , and let $\mathbf{K}_y = \mathbf{K}_x$. The feature matrix $\tilde{\Phi}_z$ for RRMCEX is built using the top 3,000 left singular vectors of $\mathbf{X} \otimes \mathbf{X}$.

Fig. 5a shows the test results on the mushroom adjacency matrix from $S = 2,000$ ($P_s = 0.006\%$) to $S = 20,000$ ($P_s = 0.036\%$) in steps of 1,000 observations. KKMCEX and RRMCEX achieve similar NMSE, while SGD has an error one order of magnitude higher, and ALS outperforms both by around one order of magnitude. This difference with ALS is because regression-based methods restrict the solution to belong to the

³<http://archive.ics.uci.edu/ml>

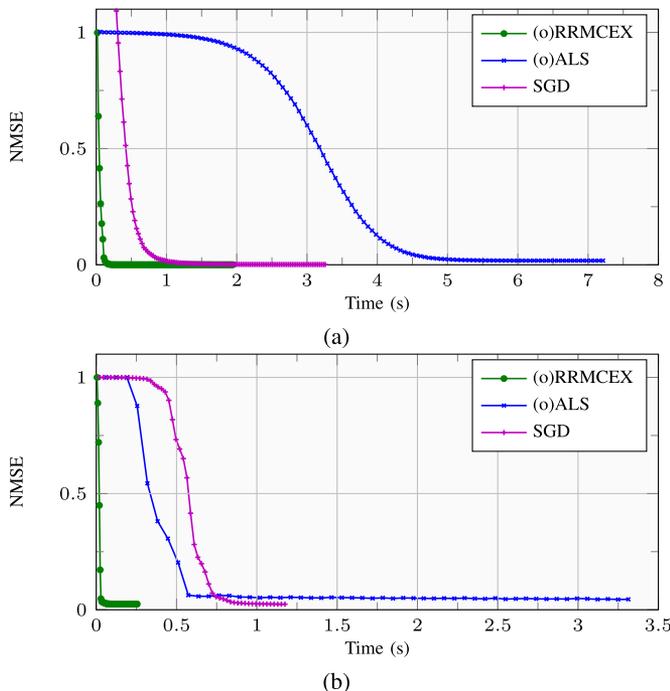


Fig. 6. NMSE vs time for the online algorithms on the (a) synthetic noiseless matrix; and (b) matrix of noiseless temperature measurements. Each mark denotes 1000 iterations have passed.

space spanned by the basis matrix. On the other hand, when solving (16), we do not enforce the constraints $\mathbf{W} \in \mathcal{H}_x$, $\mathbf{H} \in \mathcal{H}_y$ [12], [13]. Therefore, when the prior information encoded in the kernel matrices is imperfect, ALS might be able to find a factorization that fits better the data at the cost of having $\hat{\mathbf{W}} \notin \mathcal{H}_x$ and $\hat{\mathbf{H}} \notin \mathcal{H}_y$. However, in Fig. 5b we see that the computational cost for ALS and SGD is much higher than for KKMCEX and RRMCEX for the smaller P_s . On the other hand, the time for ALS decreases with S due to requiring less iterations to converge, whereas for KRRMCEX and RRMCEX it increases with S .

D. Online MC

In the online scenario, we compare the (o)RRMCEX algorithm with online (o)ALS and SGD. One observation is revealed per iteration at random, and all three algorithms process a single observation per iteration in a circular fashion. Per realization, we run tests on both synthetic and temperature matrices with $P_s = 10\%$, that is, $S = 6,250$ and $S = 5,475$ observations for the synthetic and temperature matrices, respectively, for a single realization.

Fig. 6a depicts the tests for the noiseless synthetic matrix. Clearly, (o)RRMCEX converges much faster than SGD and (o)ALS. Indeed, as opposed to SGD and (o)ALS, which require several passes over the data, (o)RRMCEX approaches the minimum in around 6,000 iterations. Moreover, it achieves the smallest NMSE of 0.0004, which is slightly below the 0.0011 obtained by SGD. Fig. 6b shows the results for the temperature matrix without noise. Again, we observe that (o)RRMCEX

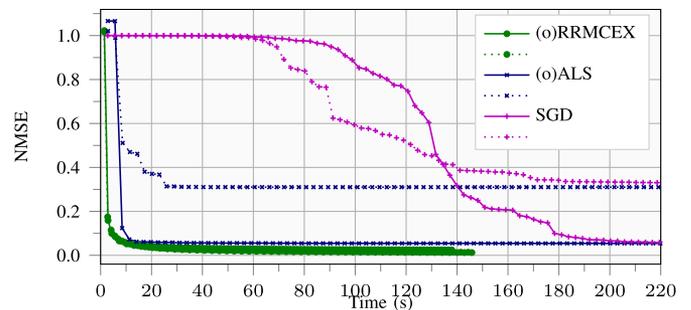


Fig. 7. NMSE vs time for the mushroom adjacency matrix. Solid lines denote $S = 20,000$, and dotted lines denote $S = 10,000$. Each mark denotes 10000 iterations have passed.

converges the fastest to the minimum, whereas SGD requires many passes through the data before it starts descending, while (o)ALS converges much faster than with the synthetic data. Regarding the NMSE, (o)RRMCEX and SGD achieve the same minimum value.

The tests on the Mushroom dataset are run with $S = 10,000$ ($P_s = 0.033\%$) and $S = 20,000$ ($P_s = 0.036\%$) observations following the same procedure as with the synthetic and temperature datasets. Fig. 7 shows results for the Mushroom adjacency matrix with the error for $S = 20,000$ plotted in solid lines, and for $S = 10,000$ in dotted lines. We observe that for $S = 20,000$, (o)RRMCEX crosses the minimum of (o)ALS and SGD in 7 seconds, whereas (o)ALS and SGD converge to this minimum in 12 and 200 seconds, respectively. Afterwards, the line for (o)RRMCEX keeps descending until an error of 0.012 is reached. For $S = 10,000$ the convergence time of (o)RRMCEX and SGD remains almost unchanged, whereas for (o)ALS it increases to 26 seconds. Moreover, the error of both (o)ALS and SGD grows much larger, whereas (o)RRMCEX exhibits just a small increase.

VIII. CONCLUSION

In this paper, we have taken a comprehensive look at MC under the framework of RKHS. We have viewed columns and rows of the data matrix as functions from an RKHS, and leveraged kernel theory to account for the available prior information on the contents of the sought matrix. Moreover, we have developed two estimation algorithms that offer simplicity and speed as their main advantages. When the number of observed data is small, KKMCEX obtains the full matrix estimate by inverting a reduced-size matrix thanks to the Representer Theorem. On the other hand, when the number of observations is too large for KKMCEX to handle, RRMCEX can be employed instead in order to lower the computational cost with no impact on the recovery error when noise is present. In addition, RRMCEX can be easily turned into an online method implemented via SGD iterations. Compared to mainstream methods designed for the factorization-based formulation, namely ALS and SGD, our RRMCEX exhibited improved performance in simulated and real data sets.

Our future research agenda includes improving both KKMCEX and RRMCEX through parallel and accelerated regression methods, as well as designing robust sampling strategies for MCEX formulated as a kernel regression.

APPENDIX

A. Proof of Lemma 1

For the KKMCEX estimator (35), the MSE is given as

$$\text{MSE} := \mathbb{E}_e \left\{ \|\mathbf{v} - \hat{\mathbf{v}}_K\|_2^2 \right\} = \mathbb{E}_e \left\{ \|\mathbf{v} - \mathbf{K}_z \hat{\boldsymbol{\gamma}}\|_2^2 \right\}. \quad (57)$$

Plugging the estimator from (34) into (57) yields

$$\begin{aligned} \text{MSE} &= \mathbb{E}_e \left\{ \left\| \mathbf{v} - \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} (\mathbf{S} \mathbf{v} + \bar{\mathbf{e}}) \right\|_2^2 \right\} \\ &= \left\| (\mathbf{I} - \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} \mathbf{S}) \mathbf{v} \right\|_2^2 \\ &\quad + \mathbb{E}_e \left\{ \left\| \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} \bar{\mathbf{e}} \right\|_2^2 \right\} \end{aligned} \quad (58)$$

where we have used that $\mathbb{E}\{\bar{\mathbf{e}}\} = \mathbf{0}$. Further, the first and second terms in (58) are the bias and variance of the KKMCEX estimator, respectively. If we substitute $\mathbf{v} = \mathbf{K}_z \boldsymbol{\gamma}$ into the first term of (58), we obtain

$$\begin{aligned} \text{bias} &= \left\| (\mathbf{I} - \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} \mathbf{S}) \mathbf{K}_z \boldsymbol{\gamma} \right\|_2^2 \\ &= \left\| (\mathbf{K}_z - \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} \mathbf{S} \mathbf{K}_z) \boldsymbol{\gamma} \right\|_2^2 \\ &= \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z) \boldsymbol{\gamma} \right\|_2^2 \end{aligned} \quad (59)$$

where $\tilde{\mathbf{T}}_z$ is the regularized Nyström approximation of \mathbf{K}_z in (38). On the other hand, the variance term is

$$\begin{aligned} \text{var} &= \mathbb{E}_e \left\{ \left\| \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} \bar{\mathbf{e}} \right\|_2^2 \right\} \\ &= \mathbb{E} \left\{ \frac{1}{\mu^2} \left\| \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} \right. \right. \\ &\quad \left. \left. \times (\mu \mathbf{I} + \mathbf{S} \mathbf{K}_z \mathbf{S}^T - \mathbf{S} \mathbf{K}_z \mathbf{S}^T) \bar{\mathbf{e}} \right\|_2^2 \right\} \\ &= \mathbb{E}_e \left\{ \frac{1}{\mu^2} \left\| \mathbf{K}_z \mathbf{S}^T - \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T \right. \right. \\ &\quad \left. \left. + \mu \mathbf{I})^{-1} \mathbf{S} \mathbf{K}_z \mathbf{S}^T \bar{\mathbf{e}} \right\|_2^2 \right\} \\ &= \mathbb{E}_e \left\{ \frac{1}{\mu^2} \left\| (\mathbf{K}_z - \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T \right. \right. \\ &\quad \left. \left. + \mu \mathbf{I})^{-1} \mathbf{S} \mathbf{K}_z) \mathbf{S}^T \bar{\mathbf{e}} \right\|_2^2 \right\} \\ &= \mathbb{E}_e \left\{ \frac{1}{\mu^2} \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z) \mathbf{S}^T \bar{\mathbf{e}} \right\|_2^2 \right\}. \end{aligned} \quad (60)$$

Adding the two terms in (59) and (60), we obtain the MSE in (39).

B. Proof of Theorem 2

Since $\mathbf{K}_z - \tilde{\mathbf{T}}_z$ appears in the bias and variance terms in Lemma 1, we will first derive an upper bound on its eigenvalues

that will eventually lead us to a bound on the MSE. To this end, we will need a couple of lemmas.

Lemma 3: Given a symmetric matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and a symmetric nonsingular matrix $\mathbf{B} \in \mathbb{R}^{N \times N}$, it holds that $\lambda_k(\mathbf{A}\mathbf{B}) = \lambda_k(\mathbf{B}^{\frac{1}{2}} \mathbf{A} \mathbf{B}^{\frac{1}{2}})$; and also $\lambda_k(\mathbf{A}\mathbf{B}) \leq \lambda_k(\mathbf{A}) \lambda_N(\mathbf{B})$.

Proof: Since \mathbf{B} is invertible and symmetric, we can write $\mathbf{A}\mathbf{B} = \mathbf{B}^{-\frac{1}{2}} (\mathbf{B}^{\frac{1}{2}} \mathbf{A} \mathbf{B}^{\frac{1}{2}}) \mathbf{B}^{\frac{1}{2}}$. Therefore, $\mathbf{A}\mathbf{B}$ is similar to $\mathbf{B}^{\frac{1}{2}} \mathbf{A} \mathbf{B}^{\frac{1}{2}}$, and they both share the same eigenvalues. Let $\mathcal{U} \subset \mathbb{R}^N \setminus \{\mathbf{0}\}$. From the min-max theorem [44], the k^{th} eigenvalue of \mathbf{A} satisfies

$$\lambda_k(\mathbf{A}) = \min_{\mathcal{U}} \left\{ \max_{\mathbf{x} \in \mathcal{U}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \mid \dim(\mathcal{U}) = k \right\}. \quad (61)$$

Therefore, we have

$$\begin{aligned} \lambda_k(\mathbf{A}\mathbf{B}) &= \lambda_k(\mathbf{B}^{\frac{1}{2}} \mathbf{A} \mathbf{B}^{\frac{1}{2}}) \\ &= \min_{\mathcal{U}} \left\{ \max_{\mathbf{x} \in \mathcal{U}} \frac{\mathbf{x}^T \mathbf{B}^{\frac{1}{2}} \mathbf{A} \mathbf{B}^{\frac{1}{2}} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \mid \dim(\mathcal{U}) = k \right\} \\ &= \min_{\mathcal{U}} \left\{ \max_{\mathbf{x} \in \mathcal{U}} \frac{\mathbf{x}^T \mathbf{B}^{\frac{1}{2}} \mathbf{A} \mathbf{B}^{\frac{1}{2}} \mathbf{x}}{\mathbf{x}^T \mathbf{B}^{\frac{1}{2}} \mathbf{B}^{\frac{1}{2}} \mathbf{x}} \frac{\mathbf{x}^T \mathbf{B} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \mid \dim(\mathcal{U}) = k \right\} \\ &\leq \min_{\mathcal{U}} \left\{ \max_{\mathbf{x} \in \mathcal{U}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \mid \dim(\mathcal{U}) = k \right\} \lambda_N(\mathbf{B}) \\ &= \lambda_k(\mathbf{A}) \lambda_N(\mathbf{B}). \end{aligned} \quad (62)$$

The following lemma bounds the eigenvalues of $\mathbf{K}_z - \tilde{\mathbf{T}}_z$, and the regularized Nystrom approximation $\tilde{\mathbf{T}}_z$ in (38). ■

Lemma 4: With \mathbf{K}_z as in (23) and $\tilde{\mathbf{T}}_z$ as in (38), the eigenvalues of $\mathbf{K}_z - \tilde{\mathbf{T}}_z$ are bounded as

$$\mathbf{K}_z - \tilde{\mathbf{T}}_z \preceq \frac{\mu \sigma_{NL}}{\sigma_{NL} + \mu} \mathbf{I}'_S + \sigma_{NL} \mathbf{I}_S \quad (63)$$

where σ_{NL} is the largest eigenvalue of \mathbf{K}_z , $\mathbf{I}_S := \text{diag}([0, 0, \dots, 1, 1])$ has S zeros on its diagonal, and $\mathbf{I}'_S := \mathbf{I} - \mathbf{I}_S$.

Proof: Using the eigendecomposition $\mathbf{K}_z = \mathbf{Q}_z \boldsymbol{\Sigma}_z \mathbf{Q}_z^T$, we can write

$$\begin{aligned} \mathbf{K}_z - \tilde{\mathbf{T}}_z &= \mathbf{K}_z - \mathbf{K}_z \mathbf{S}^T (\mathbf{S} \mathbf{K}_z \mathbf{S}^T + \mu \mathbf{I})^{-1} \mathbf{S} \mathbf{K}_z \\ &= \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} \left[\mathbf{I} - \boldsymbol{\Sigma}_z^{-\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T (\mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T \right. \\ &\quad \left. + \mu \mathbf{I})^{-1} \mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} \right] \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T. \end{aligned} \quad (64)$$

Applying the MIL to the matrix inside the square brackets of (64), we arrive at

$$\begin{aligned} \mathbf{I} - \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T (\mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T + \mu \mathbf{I})^{-1} \mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} \\ = (\mathbf{I} + \frac{1}{\mu} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T \mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}})^{-1}. \end{aligned} \quad (65)$$

That in turn implies

$$\begin{aligned} \mathbf{K}_z - \tilde{\mathbf{T}}_z &= \mu \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} (\mu \mathbf{I} + \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T \mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}})^{-1} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \\ &= \mu \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} (\boldsymbol{\Sigma}_z + \mu \mathbf{I} - \boldsymbol{\Sigma}_z + \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T \mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}})^{-1} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \end{aligned}$$

$$\begin{aligned}
&= \mu \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} \left[(\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{\frac{1}{2}} \left(\mathbf{I} - (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \boldsymbol{\Sigma}_z (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \right. \right. \\
&\quad \left. \left. + (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T \mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \right) \right] \\
&\quad \times (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{\frac{1}{2}} \left. \right]^{-1} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \\
&= \mu \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} (\mathbf{I} - \mathbf{P})^{-1} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \quad (66)
\end{aligned}$$

where

$$\begin{aligned}
\mathbf{P} &= \boldsymbol{\Sigma}_z (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-1} \\
&\quad - (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T \mathbf{S}^T \mathbf{S} \mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}}. \quad (67)
\end{aligned}$$

Regarding the eigenvalues of $\mathbf{K}_z - \tilde{\mathbf{T}}_z$ in (66), we can bound them as

$$\begin{aligned}
&\lambda(\mathbf{K}_z - \tilde{\mathbf{T}}_z) \\
&= \mu \lambda(\mathbf{Q}_z \boldsymbol{\Sigma}_z^{\frac{1}{2}} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} (\mathbf{I} - \mathbf{P})^{-1} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \mathbf{Q}_z^T) \\
&= \mu \lambda(\boldsymbol{\Sigma}_z^{\frac{1}{2}} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} (\mathbf{I} - \mathbf{P})^{-1} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}}) \\
&= \mu \lambda((\mathbf{I} - \mathbf{P})^{-1} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-1} \boldsymbol{\Sigma}_z) \\
&\leq \frac{\mu \sigma_{NL}}{\sigma_{NL} + \mu} \lambda((\mathbf{I} - \mathbf{P})^{-1}) \quad (68)
\end{aligned}$$

where $\lambda(\cdot)$ denotes the eigenvalues of a matrix, and we have applied Lemma 3 on the third equality and the last inequality. Knowing that $\lambda(\mathbf{I} - \mathbf{P}) = \mathbf{I} - \lambda(\mathbf{P})$ we can now bound the eigenvalues of \mathbf{P} as

$$\begin{aligned}
\lambda(\mathbf{P}) &= \lambda \left(\boldsymbol{\Sigma}_z^{\frac{1}{2}} (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \mathbf{Q}_z^T (\mathbf{I} - \mathbf{S}^T \mathbf{S}) \right. \\
&\quad \left. \times \mathbf{Q}_z (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-\frac{1}{2}} \boldsymbol{\Sigma}_z^{\frac{1}{2}} \right) \\
&= \lambda(\mathbf{Q}_z^T (\mathbf{I} - \mathbf{S}^T \mathbf{S}) \mathbf{Q}_z (\boldsymbol{\Sigma}_z + \mu \mathbf{I})^{-1} \boldsymbol{\Sigma}_z) \\
&\leq \frac{\sigma_{NL}}{\sigma_{NL} + \mu} \lambda(\mathbf{Q}_z^T (\mathbf{I} - \mathbf{S}^T \mathbf{S}) \mathbf{Q}_z) \\
&= \frac{\sigma_{NL}}{\sigma_{NL} + \mu} \lambda(\mathbf{I}_S) \quad (69)
\end{aligned}$$

where we have applied Lemma 3 on the second and third inequalities, and $\mathbf{I}_S := \text{diag}[0, 0, \dots, 1, 1]$ has S zeros on its diagonal. Next, we have that $\mathbf{I} - \mathbf{P} \succeq \mathbf{I} - \frac{\sigma_{NL}}{\sigma_{NL} + \mu} \mathbf{I}_S$, and thus

$$(\mathbf{I} - \mathbf{P})^{-1} \preceq \mathbf{I}'_S + \frac{\sigma_{NL} + \mu}{\mu} \mathbf{I}_S \quad (70)$$

where $\mathbf{I}'_S := \mathbf{I} - \mathbf{I}_S$. Finally, combining (70) with (68) yields

$$\mathbf{K}_z - \tilde{\mathbf{T}}_z \preceq \frac{\mu \sigma_{NL}}{\sigma_{NL} + \mu} \mathbf{I}'_S + \sigma_{NL} \mathbf{I}_S \quad (71)$$

which concludes the proof. \blacksquare

Using Lemmas 3 and 4, we can proceed to establish a bound on the bias and variance. Considering the eigendecomposition $\mathbf{K}_z - \tilde{\mathbf{T}}_z = \mathbf{L} \boldsymbol{\Lambda} \mathbf{L}^T$, we can write the bias in (59) as

$$\text{bias} = \|\mathbf{L} \boldsymbol{\Lambda} \mathbf{L}^T \boldsymbol{\gamma}\|_2. \quad (72)$$

With $\tilde{\boldsymbol{\gamma}} := \mathbf{L}^T \boldsymbol{\gamma}$, and using Lemma 4 the bias is bounded as

$$\begin{aligned}
\text{bias} &= \|\mathbf{L} \boldsymbol{\Lambda} \tilde{\boldsymbol{\gamma}}\|_2^2 = \tilde{\boldsymbol{\gamma}}^T \boldsymbol{\Lambda}^2 \tilde{\boldsymbol{\gamma}} \\
&\leq \frac{\mu^2 \sigma_{NL}^2}{(\sigma_{NL} + \mu)^2} \tilde{\boldsymbol{\gamma}}^T \mathbf{I}'_S \tilde{\boldsymbol{\gamma}} + \sigma_{NL}^2 \tilde{\boldsymbol{\gamma}}^T \mathbf{I}_S \tilde{\boldsymbol{\gamma}} \\
&= \frac{\mu^2 \sigma_{NL}^2}{(\sigma_{NL} + \mu)^2} \sum_{i=1}^S \tilde{\gamma}_i^2 + \sigma_{NL}^2 \sum_{i=S+1}^{NL} \tilde{\gamma}_i^2. \quad (73)
\end{aligned}$$

To bound the variance in (60), recall that \mathbf{e} is a Gaussian random vector with covariance matrix $\nu^2 \mathbf{I}$, while $\bar{\mathbf{e}}$ has covariance matrix $\nu^2 \mathbf{S} \mathbf{S}^T$. Then (60) is a quadratic form in \mathbf{e} , whose the variance becomes

$$\begin{aligned}
\text{var} &= \mathbb{E}_{\mathbf{e}} \left\{ \frac{1}{\mu^2} \left\| (\mathbf{K}_z - \tilde{\mathbf{T}}_z) \mathbf{S}^T \bar{\mathbf{e}} \right\|_2^2 \right\} \\
&= \frac{\nu^2}{\mu^2} \text{Tr}(\mathbf{S} (\mathbf{K}_z - \tilde{\mathbf{T}}_z)^2 \mathbf{S}^T) \\
&= \frac{\nu^2}{\mu^2} \text{Tr}((\mathbf{K}_z - \tilde{\mathbf{T}}_z)^2 \mathbf{S}^T \mathbf{S}). \quad (74)
\end{aligned}$$

The matrix inside the trace in (74) has $NL - S$ zero entries in its diagonal. Lemma 4, on the other hand, implies that diagonal entries of $\mathbf{K}_z - \tilde{\mathbf{T}}_z$ are smaller than its largest eigenvalue; that is, $[\mathbf{K}_z - \tilde{\mathbf{T}}_z]_{i,i} \leq \sigma_{NL}$. Coupling this with (74) yields

$$\text{var} \leq \frac{S \nu^2 \sigma_{NL}^2}{\mu^2}. \quad (75)$$

Finally, combining the bias bound in (73) with the variance bound in (75), yields the bound for the MSE as

$$\text{MSE} \leq \frac{\mu^2 \sigma_{NL}^2}{(\sigma_{NL} + \mu)^2} \sum_{i=1}^S \tilde{\gamma}_i^2 + \sigma_{NL}^2 \sum_{i=S+1}^{NL} \tilde{\gamma}_i^2 + \frac{S \nu^2 \sigma_{NL}^2}{\mu^2}. \quad (76)$$

REFERENCES

- [1] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Found. Comput. Math.*, vol. 9, no. 6, pp. 717–772, Dec. 2009.
- [2] H. Ji, C. Liu, Z. Shen, and Y. Xu, "Robust video denoising using low rank matrix completion," in *Proc. Comput. Vision Pattern Recognit. Conf.*, San Francisco, CA, USA, Jun. 2010, pp. 1791–1798.
- [3] K. Yi, J. Wan, T. Bao, and L. Yao, "A DCT regularized matrix completion algorithm for energy efficient data gathering in wireless sensor networks," *Int. J. Distrib. Sensor Netw.*, vol. 11, no. 7, Jul. 2015, Art. no. 272761.
- [4] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, Aug. 2009.
- [5] E. J. Candès and Y. Plan, "Matrix completion with noise," *IEEE*, vol. 98, no. 6, pp. 925–936, Jun. 2010.
- [6] J. Cheng, Q. Ye, H. Jiang, D. Wang, and C. Wang, "STCDG: An efficient data gathering algorithm based on matrix completion for wireless sensor networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 2, pp. 850–861, Feb. 2013.
- [7] V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst, "Matrix completion on graphs," in *Proc. Neural Inf. Process. Syst. Workshop "Out of the Box: Robustness in High Dimension"*, Montreal, QC, Canada, Dec. 2014.
- [8] S. Chen, A. Sandryhaila, J. M. Moura, and J. Kovacević, "Signal recovery on graphs: Variation minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 17, pp. 4609–4624, Sep. 2015.
- [9] N. Rao, H.-F. Yu, P. K. Ravikumar, and I. S. Dhillon, "Collaborative filtering with graph information: Consistency and scalable methods," in *Proc. Adv. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2015, pp. 2107–2115.

- [10] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King, "Recommender systems with social regularization," in *Proc. ACM Int. Conf. Web Search Data Mining*, Hong Kong, Feb. 2011, pp. 287–296.
- [11] J. Abernethy, F. Bach, T. Evgeniou, and J.-P. Vert, "Low-rank matrix factorization with attributes," Ecole des mines de Paris, Paris, France, Tech. Rep. N-24/06/MM, Sep. 2006.
- [12] J. A. Bazerque and G. B. Giannakis, "Nonparametric basis pursuit via sparse kernel-based learning: A unifying view with advances in blind methods," *IEEE Signal Process. Mag.*, vol. 30, no. 4, pp. 112–125, Jul. 2013.
- [13] T. Zhou, H. Shan, A. Banerjee, and G. Sapiro, "Kernelized probabilistic matrix factorization: Exploiting graphs and side information," in *Proc. SIAM Int. Conf. Data Mining*, Minneapolis, MN, USA, Jul. 2012, pp. 403–414.
- [14] M. Stock, T. Pahikkala, A. Airola, B. de Baets, and W. Waegeman, "A comparative study of pairwise learning methods based on kernel ridge regression," *Neural Comput.*, vol. 30, no. 8, pp. 2245–2283, 2018.
- [15] S. Ma, D. Goldfarb, and L. Chen, "Fixed point and Bregman iterative methods for matrix rank minimization," *Math. Program.*, vol. 128, no. 1–2, pp. 321–353, Jun. 2011.
- [16] J. F. Cai, E. J. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. Optim.*, vol. 20, no. 4, pp. 1956–1982, Jan. 2010.
- [17] P. Giménez-Febrer and A. Pagès-Zamora, "Matrix completion of noisy graph signals via proximal gradient minimization," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, New Orleans, LA, USA, Mar. 2017, pp. 4441–4445.
- [18] T. Hastie, R. Mazumder, J. D. Lee, and R. Zadeh, "Matrix completion and low-rank SVD via fast alternating least squares," *J. Mach. Learn. Res.*, vol. 16, pp. 3367–3402, Jan. 2015.
- [19] P. Jain, P. Netrapalli, and S. Sanghavi, "Low-rank matrix completion using alternating minimization," in *Proc. ACM Symp. Theory Comput.*, Palo Alto, CA, USA, Jun. 2013, pp. 665–674.
- [20] R. Gemulla, E. Nijkamp, P. J. Haas, and Y. Sismanis, "Large-scale matrix factorization with distributed stochastic gradient descent," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, San Diego, CA, USA, Aug. 2011, pp. 69–77.
- [21] J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [22] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning* (Springer Series in Statistics). New York, NY, USA: Springer-Verlag 2001.
- [23] D. Romero, M. Ma, and G. B. Giannakis, "Kernel-based reconstruction of graph signals," *IEEE Trans. Signal Process.*, vol. 65, no. 3, pp. 764–778, Feb. 2017.
- [24] N. Srebro, J. Rennie, and T. S. Jaakkola, "Maximum-margin matrix factorization," in *Proc. Adv. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2005, pp. 1329–1336.
- [25] C. Teflioudi, F. Makari, and R. Gemulla, "Distributed matrix completion," in *Proc. Int. Conf. Data Mining*, Brussels, Belgium, Dec. 2012, pp. 655–664.
- [26] M. Mardani, G. Mateos, and G. B. Giannakis, "Subspace learning and imputation for streaming big data matrices and tensors," *IEEE Trans. Signal Process.*, vol. 63, no. 10, pp. 2663–2677, May 2015.
- [27] B. Schölkopf, R. Herbrich, and A. J. Smola, "A generalized representer theorem," in *Proc. Int. Conf. Comput. Theory*, 2001, pp. 416–426.
- [28] H. V. Henderson and S. R. Searle, "On deriving the inverse of a sum of matrices," *SIAM Rev.*, vol. 23, no. 1, pp. 53–60, 1981.
- [29] T. Pahikkala, M. Stock, A. Airola, T. Aittokallio, B. de Baets, and W. Waegeman, "A two-step learning approach for solving full and almost full cold start problems in dyadic prediction," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*, 2014, pp. 517–532.
- [30] P. Drineas and M. W. Mahoney, "On the Nyström method for approximating a Gram matrix for improved kernel-based learning," *J. Mach. Learn. Res.*, vol. 6, pp. 2153–2175, Dec. 2005.
- [31] A. Alaoui and M. W. Mahoney, "Fast randomized kernel ridge regression with statistical guarantees," in *Proc. Adv. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2015, pp. 775–783.
- [32] Y. Yang, M. Pilanci, and M. J. Wainwright, "Randomized sketches for kernels: Fast and optimal non-parametric regression," *Ann. Statist.*, vol. 45, no. 3, pp. 991–1023, Jun. 2017.
- [33] H. Avron, K. L. Clarkson, and D. P. Woodruff, "Faster kernel ridge regression using sketching and preconditioning," *SIAM J. Matrix Anal. Appl.*, vol. 38, no. 4, pp. 1116–1138, Jan. 2017.
- [34] S. Van Vaerenbergh and I. Santamaría, "Online regression with kernels," in *Regularization, Optimization, Kernels, and Support Vector Machines*. London, U.K.: Chapman & Hall, 2014, ch. 21, pp. 477–501.
- [35] J. Lu, S. C. Hoi, J. Wang, P. Zhao, and Z.-Y. Liu, "Large scale online kernel learning," *J. Mach. Learn. Res.*, vol. 17, no. 47, pp. 1–43, Jan. 2016.
- [36] F. Sheikhholeslami, D. Berberidis, and G. B. Giannakis, "Large-scale kernel-based feature extraction via low-rank subspace tracking on a budget," *IEEE Trans. Signal Process.*, vol. 66, no. 8, pp. 1967–1981, Apr. 2018.
- [37] L. Bottou, "Stochastic gradient descent tricks," in *Neural Networks: Tricks of the Trade*. New York, NY, USA: Springer-Verlag, 2012, pp. 421–436.
- [38] I. D. Schizas, G. Mateos, and G. B. Giannakis, "Distributed LMS for consensus-based in-network adaptive processing," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2365–2382, Jun. 2009.
- [39] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 83–98, May 2013.
- [40] A. J. Smola and R. Kondor, "Kernels and regularization on graphs," in *Learning Theory and Kernel Machines*. New York, NY, USA: Springer-Verlag, 2003, pp. 144–158.
- [41] D. Romero, V. N. Ioannidis, and G. B. Giannakis, "Kernel-based reconstruction of space-time functions on dynamic graphs," *IEEE J. Sel. Topics Signal Process.*, vol. 11, no. 6, pp. 856–869, Sep. 2017.
- [42] A. Cotter, J. Keshet, and N. Srebro, "Explicit approximations of the gaussian kernel," 2011, arXiv:1109.4603.
- [43] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," in *Proc. Adv. Neural Inf. Process. Syst.*, 2008, pp. 1177–1184.
- [44] P. Lax, *Linear Algebra and Its Applications*. Hoboken, NJ, USA: Wiley, 2007.