

UCLA

UCLA Previously Published Works

Title

Graph-constrained analysis for multivariate functional data

Permalink

<https://escholarship.org/uc/item/7jq9g523>

Authors

Dey, Debangana

Banerjee, Sudipto

Lindquist, Martin A

et al.

Publication Date

2025-05-01

DOI

10.1016/j.jmva.2025.105428

Peer reviewed



Graph-constrained analysis for multivariate functional data

Debangon Dey^a, Sudipto Banerjee^b, Martin A. Lindquist^c, Abhirup Datta^{c,*}

^a National Institute of Mental Health, Bethesda, 20892, MD, USA

^b Department of Biostatistics, University of California Los Angeles, Los Angeles, 90095, CA, USA

^c Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, Baltimore, 21205, MD, USA

ARTICLE INFO

AMS 2020 subject classifications:

primary 62R10

secondary 62F30

Keywords:

Conditional independence

Functional data analysis

Gaussian processes

Graphical models

Multivariate analysis

Spatial data

ABSTRACT

The manuscript considers multivariate functional data analysis with a known graphical model among the functional variables representing their conditional relationships (e.g., brain region-level fMRI data with a prespecified connectivity graph among brain regions). Functional Gaussian graphical models (GGM) used for analyzing multivariate functional data customarily estimate an unknown graphical model, and cannot preserve knowledge of a given graph. We propose a method for multivariate functional analysis that exactly conforms to a given inter-variable graph. We first show the equivalence between partially separable functional GGM and graphical Gaussian processes (GP), proposed recently for constructing optimal multivariate covariance functions that retain a given graphical model. The theoretical connection helps to design a new algorithm that leverages Dempster's covariance selection for obtaining the maximum likelihood estimate of the covariance function for multivariate functional data under graphical constraints. We also show that the finite term truncation of functional GGM basis expansion used in practice is equivalent to a low-rank graphical GP, which is known to oversmooth marginal distributions. To remedy this, we extend our algorithm to better preserve marginal distributions while respecting the graph and retaining computational scalability. The benefits of the proposed algorithms are illustrated using empirical experiments and a neuroimaging application.

1. Introduction

Gaussian graphical models (GGM) are extensively used to represent sparsity structures emerging from conditional independence relations among a collection of Gaussian random variables. The graph specifies the zeroes in the precision matrix for multivariate observations. Conducting multivariate analysis under a known graph constraint, i.e., when the variables conform to a given graph, is of key interest. Covariance selection [11] offers a seminal result to answer this specific question. When the graph is unknown, classical approaches such as the graphical Lasso estimate the graph (from the sparse precision matrix) using L-1 regularization [17,26]. Bayesian alternatives employ Markov chain Monte Carlo (MCMC) algorithms exploiting a conjugate hyper-inverse Wishart prior for estimating the graph and the precision matrix [7,13,32].

Graphical Gaussian models are customarily applied to multivariate data sets, where the variables represented by the nodes of the graph are scalar- or vector-valued. Here we consider a collection of functions that conceptually exist in a continuum over a domain (e.g., over space or time) and are associated among themselves via a posited conditional independence graph. We consider analysis of partial realizations of these multiple functions or processes, which is always the case in practice. Examples of such multivariate functional data include physical activity or heart rate continuously measured from actigraph monitors (wearables), EEG/fMRI signals

* Corresponding author.

E-mail address: abhidatta@jhu.edu (A. Datta).

<https://doi.org/10.1016/j.jmva.2025.105428>

Received 19 September 2023; Received in revised form 14 February 2025; Accepted 17 February 2025

Available online 24 February 2025

0047-259X/© 2025 Elsevier Inc. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

from different brain regions (neuroimaging), multiple-gene expressions from cells distributed over a tissue (spatial transcriptomics), population counts of multiple biological species over a region (ecology) and multiple pollutants measured from environmental monitoring stations (environmental sciences).

Multivariate functional data are often modeled using a latent multivariate stochastic process, such as a multivariate Gaussian process (GP) [3,8,31,36]. Hence, it is natural to seek extensions of graphical models to settings where each node represents an entire function or stochastic process. Two recent approaches have extended the GGM to represent conditional dependencies among multivariate stochastic processes on continuous domains: (i) functional GGM (FGGM); and (ii) graphical Gaussian Process (GGP). The former relies on a multivariate Kosambi-Karhounen-Loève basis expansion and estimates the sparsity (graphical model) of the precision matrices of the Gaussian coefficient vectors via penalized estimation [41,42]. Functional GGMs have been primarily used for functional data analysis, where basis function representations are customarily used but have also been applied to spatial data [21]. These methods estimate the unknown graph through the aforementioned regularization techniques.

Graphical GP (GGP) [12], originally proposed for multivariate spatial analysis, offers a multivariate covariance function that exactly conforms to a specified graph among the variables. Given this undirected graph and a multivariate covariance function, a GGP is the derived optimal covariance function that satisfies the prespecified graphical constraint. This means that the processes comprising a GGP exactly satisfy process-level conditional independence specified by the graph. At the same time, GGP retains marginal process distributions from the original multivariate covariance function it is derived from, thereby facilitating the interpretation of the attributes of each individual process via inference on the respective parameters, which is often the primary inferential objective in spatial analysis. GGP is implemented via an algorithm called *stitching* that scales effectively for likelihood-based analysis of *highly multivariate* spatial data involving a large number of outcome variables by leveraging the sparsity of the graph.

We focus on multivariate functional data analysis when the inter-variable graph is known (e.g., existing brain-region networks, phylogenetic trees etc.), and needs to be preserved in the statistical model. This problem is an extension of Dempster's covariance selection from finite-dimensional covariance matrices to infinite-dimensional covariance operators. In other words, the goal is to find the best estimate of the multivariate covariance function under a graphical constraint on the variables. While having the knowledge of a graph may appear simpler than requiring the estimation of an unknown graph, current FGGM approaches that estimate the graph simultaneously with other parameters cannot preserve a given graphical constraint. Thus, akin to the typical challenges of parameter estimation under a known but complicated constraint, existing FGGM algorithms are not typically conformable to a known graphical constraint. This motivates our current developments.

Our contributions are summarized as follows. We first draw theoretical connections between functional GGM and graphical GP. These are apparently complementary methods for introducing graphical models in multivariate GP. The former relies on basis function expansions and is typically used for analyzing functional data. Graphical GP focuses on spatial data and directly uses parametric covariance functions. We bridge these two seemingly disconnected paradigms. We formally establish in [Theorem 1](#) that the class of *partially separable* functional GGM with the graphical model for each of the coefficient covariance matrices is equivalent to a graphical GP. Furthermore, using Nyström's method, we relate the stitching algorithm for constructing a graphical GP to a finite-term truncation of basis functions used in functional GGM. Specifically, we show that truncation corresponds to use of a low-rank GP in stitching that is known to oversmooth.

The theoretical connections we offer in this manuscript carry important implications for practice in multivariate functional analysis. We first leverage the equivalence between functional GGM and graphical GP to design an algorithm for analyzing multivariate functional data that exactly preserves the given functional conditional independence relationships. In particular, we provide a solution to Dempster's covariance selection for multivariate functional data by deriving the maximum likelihood estimates for a multivariate covariance function under the graphical constraint ([Theorem 2](#)). We subsequently show that the finite-term truncation of the basis function leads to over-smoothing as it corresponds to stitching using a low-rank GP. We use ideas from constructing full-rank graphical GP via stitching to offer an improved algorithm that mitigates the effects of over-smoothing by better preserving marginal distributions while retaining the graphical constraint. We show the utility of our method through numerous experiments on synthetic data. We demonstrate the application of the method for the analysis of brain region-level functional Magnetic Resonance Imaging (fMRI) data using a connectivity graph among brain regions. The data is from a study of effects of thermal pain on the brain. This application is particularly well-suited for the proposed method as it is known apriori that there are a number of brain regions that are consistently involved with pain processing whose connections together are predictive of physical pain [25,27].

2. The two paradigms of graphical functional modeling

We offer a brief review of functional Gaussian graphical models and graphical Gaussian Processes and how they differ in terms of the model forms and their applications.

2.1. Functional Gaussian graphical models

We first review the definition of conditional independence for Gaussian processes (GP) on continuous domains. For a set $\mathcal{V} = \{1, \dots, q\}$, let $w(\cdot) = (w_1(\cdot), \dots, w_q(\cdot))^T$ be a $q \times 1$ GP over a continuous domain \mathcal{D} with node j representing the univariate process $w_j(\cdot)$. Two univariate GPs $w_i(\cdot)$ and $w_j(\cdot)$ are conditionally independent, given the remaining processes $\{w_k(\cdot) \mid k \in \mathcal{V} \setminus \{i, j\}\}$ if $\text{Cov}(\epsilon_{iB}(s), \epsilon_{jB}(s')) = 0$ for all $s, s' \in \mathcal{D}$ and $B = \mathcal{V} \setminus \{i, j\}$, where $\epsilon_{kB}(s) = w_k(s) - \mathbb{E}[w_k(s) \mid \sigma(\{w_j(s') : j \in B, s' \in \mathcal{D}\})]$ [12] and $\sigma(\cdot)$ is

the σ -algebra generated by its argument. Functional GGM (FGGM) introduces this process-level conditional independence through the precision matrices of coefficient vectors in its basis expansion. Letting $\phi_{j\ell}$ be the orthonormal basis functions in the domain D , the process is represented as $w_j(s) = \sum_{\ell=1}^{\infty} \theta_{j\ell} \phi_{j\ell}(s)$. [42] show that modeling the coefficients $\theta_{j\ell}$ to be Gaussian and assuming a sparsity structure between the discrete multivariate GP $\theta^{(j)} = \{\theta_{j\ell} : \ell \in \{1, 2, \dots\}\}$ imposes the corresponding process-level conditional independence among the component processes of $w(\cdot)$. In practice, they truncate the expansion of $w_{(j)}(\cdot)$ to m_j terms and fit a GGM on the $M = \sum_{j=1}^q m_j$ dimensional vector of coefficients. The approach involves inverting an $M \times M$ matrix, thereby requiring $O(M^3)$ floating point operations (FLOPs) or $O(q^3 m^3)$ FLOPs when $m_j = m$ for all j . The computational expense being cubic in the number of basis functions m prohibits large values for m , which, in turn, restricts the richness of the function class.

The functional GGM in [21,41] introduces two further assumptions: (a) a common orthonormal basis function, i.e., $\phi_{j\ell} = \phi_\ell$ for all $j \in \{1, \dots, q\}$; and (b) the coefficients corresponding to different basis functions are independent, i.e., θ_j and θ_k are independent, where $\theta_j = (\theta_{1j}, \dots, \theta_{qj})^\top$. Due to the separability in the joint covariance structure of basis coefficients, they are formally termed as partial separable processes. The basis function expansion is truncated to m terms in practical implementation. This partial separable model only needs m inversions of $q \times q$ matrices ($\text{Cov}(\theta_j), j \in \{1, \dots, m\}$) and considerably reduces computational complexity to $O(mq^3)$, thereby allowing m to be larger. The algorithm maximizes a penalized likelihood with a graphical Lasso penalty on the estimated covariance of the coefficient vectors to induce sparsity. This results in the estimation of an unknown graphical model corresponding to the common sparsity among the precision matrices for each of the coefficients. Recently, [23,24] proposed a class of non-parametric graphical models for functional data that performs better than FGGMs when the relation between variables is non-linear or heteroskedastic in nature.

2.2. Graphical Gaussian processes

Spatial data is often modeled using GPs with parametric covariance functions. The Matérn family of covariances [34] are popular in spatial analysis because their parameters are interpreted as the variance, spatial smoothness, and range of the process. For multiple dependent outcomes $w(\cdot) = (w_1(\cdot), \dots, w_q(\cdot))^\top$, a common GP specification is $w(\cdot) \sim GP(0, C(\cdot, \cdot))$ where $C = (C_{ij})$ is the $q \times q$ multivariate covariance function. The multivariate Matérn is a popular choice for C , ensuring that the derived covariance and cross-covariance functions are also Matérn [1,18], allowing us to interpret each component process. However, the model struggles to scale up to highly multivariate (large q) settings.

In [12], the graphical Gaussian process (GGP) was introduced for analyzing highly multivariate spatial data. These are specified through multivariate covariance functions that exactly encode an inter-variable graphical model. Theorem 1 in [12] asserts that given any non-graphical stationary multivariate GP and any inter-variable graph, there is a unique and optimal GP approximating the original GP (in terms of integrated spectral Kullback–Leibler divergence) while respecting process-level conditional independence relationships as specified by the graph. This optimal *graphical GP* was shown to retain the univariate marginal distributions from the original GP exactly. For any pair of variables (i, j) included in the edge set of the graph, the cross-covariance is also retained from the original GP. We denote this optimal graphical GP derived from the original covariance function C and a graph \mathcal{G} as $GGP(C, \mathcal{G})$.

Definition 2.1 (Graphical GP). $w(\cdot) \sim GGP(C, \mathcal{G})$ is a $q \times 1$ graphical GP on D given a valid $q \times q$ cross-covariance function $C = (C_{ij})$ on $D \subseteq \mathbb{R}^d$ and a graph $\mathcal{G} = (\mathcal{V}, E)$, such that

- (i) Retains univariate marginal distributions: $w_j(\cdot) \sim GP(0, C_{jj})$ for all $1 \leq j \leq q$,
- (ii) Retains edge-specific cross-covariances: $\text{Cov}(w_i(s), w_j(t)) = C_{ij}(s, t)$, $(i, j) \in E$ and $s, t \in D$,
- (iii) Encodes process-level graphical model as in Definition 2.1: $w_i(\cdot) \perp w_j(\cdot) \mid w_{-ij}(\cdot)$, $(i, j) \notin E$.

Construction of the graphical GP used an algorithm called “stitching”. The stitching algorithm in [12] combines two algorithms, which we refer to as “stitching” and “stretching” in the current manuscript. The step of stitching is applying Dempster’s covariance selection using \mathcal{G} on the covariance matrix arising from C when considering a finite set of locations $\mathcal{L} = \{s_1, \dots, s_p\}$. This defines a graphical GP on the finite set \mathcal{L} that adheres to the conditional independences specified via the graphical model \mathcal{G} . The stretching step extends the finite-dimensional process to the continuous infinite domain by adding a residual error process, while preserving the graphical model.

2.2.1. Stitching

To perform the stitching, we need to define covariance selection for positive definite (p.d.) matrices.

Definition 2.2 (Covariance Selection, [11,33]). Given a graph $\mathcal{G} = (\mathcal{V}, E)$ and a positive definite matrix A with row and column blocks indexed by \mathcal{V} , $\text{CovSel}(A, \mathcal{G})$ is the unique p.d. matrix B from Dempster’s covariance selection on A using \mathcal{G} , i.e., $B_{ij} = A_{ij}$ for $i = j$ or $(i, j) \in E$, and $(B^{-1})_{ij} = 0$ for $(i, j) \notin E$.

Given C , \mathcal{G} , and \mathcal{L} , stitching first defines a graphical Gaussian model (GGM) on $\mathcal{L} \subset D$.

$$w_{\text{stitch}}(\mathcal{L}) \sim N(0, \tilde{C}(\mathcal{L}, \mathcal{L})) \text{ where } \tilde{C}(\mathcal{L}, \mathcal{L}) = \text{CovSel}(C(\mathcal{L}, \mathcal{L}), \mathcal{G}). \quad (1)$$

Here, $C(\mathcal{L}, \mathcal{L})$ denotes the covariance matrix obtained from evaluating the covariance function C on the reference set \mathcal{L} .

2.2.2. Stretching

As part of the stretching, the covariance matrix on a finite set is then extended to a covariance function on the infinite domain. In other words, the finite-dimensional GGM (1) is extended in [12] to an infinite-dimensional GGP on D as

$$w_{stretch,j}(s) = w_{stitch,j}^{pp}(s) + z_j(s), \text{ where } w_{stitch,j}^{pp}(s) = C_{jj}(s, \mathcal{L})C_{jj}(\mathcal{L}, \mathcal{L})^{-1}w_{stitch,j}(\mathcal{L}). \tag{2}$$

Here, $C_{jj}(s, \mathcal{L}) = (C_{jj}(s, s_1), \dots, C_{jj}(s, s_p))$ and w_{stitch}^{pp} is the (fixed-) or low-rank predictive process based on the finite-dimensional distribution $w_{stitch}(\mathcal{L})$ defined on the set of knots \mathcal{L} [4]. The processes $z_j(\cdot) \stackrel{\text{ind}}{\sim} GP(0, C_{jj, res})$ are the residual GP with $C_{jj, res}(s, t) = C_{jj}(s, t) - C_{jj}(s, \mathcal{L})C_{jj}(\mathcal{L}, \mathcal{L})^{-1}C_{jj}(\mathcal{L}, t)$. Decomposing a univariate GP into a predictive process and an orthogonal residual process is well known [4]. The stitching and stretching process introduced two innovations to this decomposition. First, the process on \mathcal{L} is endowed with the covariance $\tilde{C}(\mathcal{L}, \mathcal{L})$ instead of $C(\mathcal{L}, \mathcal{L})$. This ensures that the graph \mathcal{G} is encoded while preserving marginals on \mathcal{L} . Orthogonality of $z_j(\cdot)$ and $w_{stitch,j}^{pp}(\cdot)$ then implies that the marginals are preserved over entire D if the residual processes, $z_j(\cdot)$'s, are added to the low-rank Gaussian process $w_{stitch,j}^{pp}(s)$, i.e., $w_{stretch,j}(\cdot) \sim GP(0, C_{jj})$ for all j on the entire domain D . Second, the $z_j(\cdot)$'s are chosen to be independent component-wise (across j 's), ensuring that the GGM encoded over \mathcal{L} is extended to process-level conditional independence relations over D yielding the GGP, i.e., for all $(i, j) \notin E$, we have $w_{stretch,i}(\cdot) \perp w_{stretch,j}(\cdot) \mid w_{stretch,-ij}(\cdot)$. For a decomposable sparse graph, a GGP likelihood from this process involves considerably less parameters and FLOPs than a full GP with covariance C .

3. Theoretical connections between functional GGM and graphical GP

The GGP in [12] is constructed from a parametric multivariate covariance function C such as the multivariate Matérn. In analyzing functional data, covariance functions are often represented more generally using basis functions. Here, we demonstrate how a graphical GP can be derived from any multivariate GP that admits a partially separable basis function expansion. The resulting GGP is an FGGM that will exactly conform to a given graph, and we show, in Section 4, how this is leveraged to conduct graph-constrained analysis of multivariate functional data.

3.1. Covariance selection on partially separable Gaussian processes

Consider a $q \times 1$ partially-separable multivariate GP on a continuous domain D ,

$$w_0(s) = \sum_{\ell=1}^{\infty} \theta_{\ell} \phi_{\ell}(s), \text{ where } \theta_{\ell} = (\theta_{1\ell}, \dots, \theta_{q\ell}) \sim N_q(0, \Sigma_{\ell}). \tag{3}$$

Theorem 1 proves that the graphical GP derived from $w_0(\cdot)$ is a partially separable functional GGM that exactly preserves a given graphical model.

Theorem 1. *Let C be the cross-covariance function of a stationary partially separable process $w_0(\cdot)$ in (3). Consider a new partially separable process $w(s) = \sum_{\ell=1}^{\infty} \theta_{\ell}^* \phi_{\ell}(s)$ where $\theta_{\ell}^* = (\theta_{1\ell}^*, \dots, \theta_{q\ell}^*) \sim N_q(0, \Sigma_{\ell}^*)$. Then for any fixed graph $\mathcal{G} = (\mathcal{V}, E)$, $w(\cdot) \sim GGP(C, \mathcal{G})$ as in Definition 2.1 if and only if $\Sigma_{\ell}^* = \text{CovSel}(\Sigma_{\ell}, \mathcal{G})$ for all ℓ .*

The result shows that starting with a partially separable GP, the optimal graphical GP can be obtained by applying Dempster's covariance selection on each covariance matrix Σ_{ℓ} of the basis coefficients, and this graphical GP is a partially separable functional GGM. All proofs are provided in Appendix.

3.2. Stitching and truncation

In practice, we cannot work with infinite basis expansions and construct a process using only a finite set of orthogonal functions extracted from the graphical GP $w(\cdot)$ in Theorem 1. The following result is an immediate consequence of Theorem 1. The proof is outlined in Appendix A.

Corollary 1. *Let $w_0^m(s)$ be constructed using m basis functions of a graphical GP as*

$$w_0^m(s) = \sum_{\ell=1}^m \theta_{\ell} \phi_{\ell}(s), \theta_{\ell} = (\theta_{1\ell}, \dots, \theta_{q\ell})^T \sim N_q(0, \Sigma_{\ell}) \tag{4}$$

and C^m is the covariance function of $w_0^m(s)$. Then $w^m(\cdot) \sim GGP(C^m, \mathcal{G})$ is given by $w^m(s) = \sum_{\ell=1}^m \theta_{\ell}^ \phi_{\ell}(s)$ where $\text{Cov}(\theta_{\ell}^*) = \text{CovSel}(\Sigma_{\ell}, \mathcal{G})$ for $\ell \in \{1, \dots, m\}$.*

We discuss in Section 4 how this connection facilitates a practical estimation strategy for finite-term functional GGM that exactly conforms to a given graph. However, a finite-term truncation can come at the expense of reduced accuracy in estimating the covariance function. To explicitly demonstrate the impact of truncation, we now relate this truncated functional GGM $w^m(\cdot)$ of Corollary 1 to a low-rank stitched graphical GP (recall Section 2.2) through an approximate result using Nystrom approximation for kernel matrices [14]. Let $w_0(\cdot)$ be the partially separable GP in (3) with cross-covariance function $C = (C_{ij})$, $w(\cdot)$ be the partially separable $GGP(C, \mathcal{G})$ derived from $w_0(\cdot)$ using Theorem 1 and $w^m(\cdot)$ be the finite-term truncation of $w(\cdot)$ in Corollary 1. We show in Appendix A that $w^m(\cdot)$ is approximately equal to $w_{stitch}^{pp}(\cdot)$, where $w_{stitch}^{pp} = (w_{stitch,1}^{pp}(\cdot), \dots, w_{stitch,q}^{pp}(\cdot))$ is the low-rank predictive

process defined in (2) created as part of the stitching and stretching constructions.

Theorem 1 proves the equivalence between the theoretical constructions of functional GGM and the graphical GP. Similarly, the above approximate result establishes an analogy between the practical procedures of truncation and stitching for the two approaches. This result also highlights that truncation ignores the equivalence of the stretching step employed earlier in graphical GP (Section 2.2.2). The stitched and stretched graphical GP $w_{stretch}(s)$, as specified in (2), is the sum of the stitched low-rank predictive process $w_{stitch}^{pp}(s)$ and the independent residual processes $z(\cdot)$, i.e., $w_{stretch}(s) = w_{stitch}^{pp}(s) + z(s)$, where $z(s) = (z_1(s), \dots, z_q(s))$. The truncated functional GGM $w^m(s)$ is only equivalent to the first component $w_{stitch}^{pp}(s)$ and ignores the residual component $z(s)$. This difference has important practical implications. The marginal distributions of the process $w_{stitch}^{pp}(\cdot)$ has the same distribution as $C_{jj}(\cdot, \mathcal{L})C_{jj}(\mathcal{L}, \mathcal{L})^{-1}w_{0,j}(\mathcal{L})$ — the univariate predictive processes of [4] derived from the original process $w_0(s)$. These low-rank predictive processes have been shown, theoretically and empirically, to underestimate the marginal variances of the original process $w_0(\cdot)$, resulting in oversmoothing and inferior performance [9,35]. As $w^m(\cdot)$ is equivalent to $w_{stitch}^{pp}(\cdot)$, the truncated process $w^m(\cdot)$ also will not preserve the marginals of the original distribution $w_0(\cdot)$. Rather $w^m(\cdot)$ has the same marginal distributions as the predictive process, and will be expected to oversmooth the marginal spatial patterns. In contrast, the process $w_{stretch}(\cdot)$, on account of adding the residual processes $z(\cdot)$, preserves the marginals of the original process $w_0(\cdot)$ and would not oversmooth. Thus, we see a potential drawback of the truncation. We demonstrate the oversmoothing in our empirical results and in Section 4 we offer a solution to the issue.

4. Estimation

In scientific applications, we either know the graph apriori or have to estimate it. Functional GGM is widely deployed in the latter setting [21,41]. However, if the graph is known (e.g., a known phylogenetic tree as a graph to model multiple species distributions over a region, a known auto-regressive time-dependence structure to model spatiotemporal pollutant data, connectivity graph between brain regions estimated from a secondary dataset), current functional GGM approaches cannot explicitly encode this information and preserve the graphical model in the analysis.

In this Section, we leverage the theoretical connections developed in Section 3 to present algorithms for graph-constrained analysis of multivariate functional data.

4.1. Graph-constrained MLE for multivariate functional data

Theorem 1 and **Corollary 1** establish that the optimal approximation of a partially separable process under a graphical constraint is also a partially separable process with Dempster’s covariance selection applied to covariance matrices for each of the basis function coefficients. We now provide the data-driven analog of the result that is practically applicable when analyzing multivariate functional data occurring as partial realizations of these multivariate processes. In particular, the following theorem provides the maximum likelihood estimate (MLE) for a partially separable multivariate covariance function under a graphical constraint. The result is viewed as an extension of Dempster’s covariance selection, i.e., graph-constrained MLE of covariances, from the setting of vector-valued data with covariance matrices to functional data with covariance operators.

Theorem 2. *Let X_1, \dots, X_N denote iid realizations of q -variate functional data on p locations $\{s_1, \dots, s_p\}$, generated from a mean-zero truncated partial separable $GP(0, C^m)$ as in **Corollary 1** with known orthonormal basis functions ϕ_ℓ . Then, for a given functional graphical model \mathcal{G} , the graph-constrained MLE of the covariance function C^m is given by*

$$\widehat{C}^m = \sum_{\ell=1}^m \text{CovSel}((I_q \otimes \phi_\ell^T)S(I_q \otimes \phi_\ell), \mathcal{G}) \otimes \phi_\ell \phi_\ell^T \tag{5}$$

where $S = (1/N) \sum_{i=1}^N X_i X_i^T$ is the sample covariance matrix, $\phi_\ell^T = (\phi_\ell(s_1), \dots, \phi_\ell(s_p))$ is the basis function evaluated at p locations, and CovSel is covariance selection (see **Definition 2.2**).

Lemma 1. *Under the setting of **Theorem 2**, the graph-constrained MLE \widehat{C}^m in (5) is an asymptotically consistent estimator of the covariance operator C^m .*

Theorem 2 provides a direct approach to obtaining the MLE of a multivariate covariance function under a given graphical constraint and **Lemma 1** shows the consistency of the estimator. All the proofs are shown in Appendix. In practice, the basis functions are often chosen in a data-driven manner instead of being fixed apriori. These estimates $\hat{\phi}_\ell$ can be obtained from a functional principal components analysis. Plugging these (or any other estimate of ϕ_ℓ), we can use **Theorem 2** to calculate the profile maximum likelihood estimate (PMLE) of Σ_ℓ as $\hat{S}_\ell = (I_q \otimes \hat{\phi}_\ell^T)S(I_q \otimes \hat{\phi}_\ell)$. Subsequently the graph-constrained PMLE for C^m can be obtained as

$$\widetilde{C}^m = \sum_{\ell=1}^m \text{CovSel}(\hat{S}_\ell, \mathcal{G}) \otimes \hat{\phi}_\ell \hat{\phi}_\ell^T. \tag{6}$$

The number of basis functions m is usually chosen so that the first m eigenfunctions $\hat{\phi}_\ell$ ’s explain a given percentage v of the total variance. We formally lay out the above steps in **Algorithm 1** (see [38,39]) below and term our approach as FGGM-CovSel.

Algorithm 1: FGGM-CovSel: Graph-constrained multivariate functional data analysis

- Step 1. Fix a proportion of variance explained threshold, v , say, $v = 0.95$;
 - Step 2. Perform pooled functional principal component analysis ([38,39]) to get m common eigenfunction estimates $\hat{\phi}_\ell, \ell \in \{1, \dots, m\}$ that explain v proportion of total variance. ;
 - Step 3. Use the eigenfunctions $\hat{\phi}_\ell$ to get first m functional principal component scores $\hat{\theta}_{i\ell} = (I_q \otimes \hat{\phi}_\ell^\top)X_i$ and calculate the corresponding covariance matrix $\hat{S}_\ell = \text{Cov}_N(\hat{\theta}_{i\ell}) = (I_q \otimes \hat{\phi}_\ell^\top)S(I_q \otimes \hat{\phi}_\ell)$ for $\ell \in \{1, \dots, m\}$ where $S = \frac{1}{N} \sum X_i X_i^\top$;
 - Step 4. Using Iterative Proportional Scaling (IPS) algorithm [33] for covariance selection to obtain $\text{CovSel}(\hat{S}_\ell, \mathcal{G})$ for $\ell \in \{1, \dots, m\}$;
 - Step 5. Produce the estimate of the covariance operator \widetilde{C}^m from (6).
-

4.2. Improved estimation of marginal distributions using stretching

In Algorithm 1, using a larger v improves the retention of marginal covariance functions but requires a larger number of basis functions m . This is prohibitive computationally, as the iterative proportional scaling (IPS) algorithm [33] for covariance selection on the S_ℓ 's has a total computational complexity of $O(mq^3)$ floating point operations or FLOPs for the m terms. For large q , this mandates choosing a smaller m to reduce computations. This tradeoff leads to a worse approximation of the marginal covariances and oversmoothing as the finite-term truncated process is equivalent to a low-rank GP approximation at m locations (see the discussion in Section 3.2).

We now offer an extension of Algorithm 1 using the principles of stretching that remains computationally feasible while offering significantly better retention of the marginal distributions and, thereby, less oversmoothing. In Section 3.2, we have established that $w_{stretch}(\cdot) \approx w^m(\cdot) + z(\cdot)$. The graphical GP $w_{stretch}(\cdot)$ exactly preserves the marginals of the original process $w_0(\cdot)$, while the truncated process $w^m(\cdot)$ is equivalent to a low-rank predictive process and can oversmooth. For a univariate predictive process, the oversmoothing is mitigated by adding the residual process [16]. Stretching proceeds similarly in a multivariate graphical model by introducing the residual process $z(\cdot)$ (defined in (2)) to exactly preserve the marginals while still conforming to the graphical constraint.

We leverage this connection between truncation and stitching to remedy the oversmoothing by the process $w^m(\cdot)$ truncated at some computationally feasible value of m . We modify Algorithm 1 to emulate the full rank stitching and stretching procedure of [12] by adding component-specific residual processes. This essentially adds a block-diagonal function to the estimate of the covariance function that mitigates the bias in the marginal distributions due to truncation. To achieve this, we first run FGGM-CovSel for a small or moderate m as afforded by the computing resources. Then, we calculate the residuals and perform functional principal component analysis of the component-specific residuals, and add the diagonal matrix of the estimated residual variances to our covariance estimate \widetilde{C}^m from Algorithm 1. This yields significant improvements in retaining the univariate distributions. The Algorithm is referred to as FGGM-Stretch and is formally laid out in Algorithm 2 below.

Algorithm 2: FGGM-Stretch

- Step 1. Fix a smaller proportion of variance explained threshold, v , say, $v = 0.75$;
 - Step 2. Run Algorithm 1 to get estimates $\hat{\theta}_{i\ell}, \hat{\phi}_\ell, \ell \in \{1, \dots, m\}$ and \hat{C} ;
 - Step 3. Calculate residuals: $Z_i = X_i - \sum_{\ell=1}^m \hat{\theta}_{i\ell} \hat{\phi}_\ell$ for $i \in \{1, \dots, N\}$;
 - Step 4. Perform functional principal component analysis of the residuals of $Z^{(j)} = (Z_{1j}, \dots, Z_{Nj})$ for $j \in \{1, \dots, q\}$ independently for a larger proportion of variance explained, v' , say, $v' = 0.95$;
 - Step 5. Obtain the eigenvalues as $\hat{\lambda}_{\ell j}, \ell \in \{1, \dots, m'\}, j \in \{1, \dots, q\}$;
 - Step 6. Find the new stitched covariance operator $\widehat{C}^{m(2)} = \widehat{C}^m + \sum_{\ell=1}^{m'} \text{diag}(\hat{\lambda}_{\ell 1}, \dots, \hat{\lambda}_{\ell q})$.
-

5. Empirical results

We perform a series of simulation experiments to gauge the performance accuracy of three algorithms for graph-constrained analysis of multivariate functional data — functional Gaussian graphical model (FGGM) of [41], and our two proposed algorithms FGGM-CovSel and FGGM-Stretch. The first one cannot use the knowledge of the graph and estimate the graph, treating it as unknown while simultaneously estimating the covariance function. The latter two exactly preserve the graph in the analysis.

For the data generation, we consider $q = 10$ -variate GPs adhering to the conditional independence structure from the graph in Fig. 1. We consider two choices of covariance functions generating the data – (a) correctly specified case: partial-separable (PS) covariance, and (b) misspecified case: graphical Matérn (GM) covariance of [12] that is not partially separable. We fix the total number of basis functions at $L = 101$ for data generation. We take them to be the Fourier basis functions and evaluate them at 200 equally spaced points between (0, 1). Now to generate the basis-specific coefficients, we generate the covariance matrices as $\Sigma_\ell = a_\ell \Omega_\ell^{-1}$ for $\ell \in \{1, \dots, 101\}$, where the precision matrices Ω_ℓ are generated using the algorithm of [30] to have the exact sparsity as specified by the graph \mathcal{G} (Fig. 1). The decaying constants $a_\ell = 3^{l-1.8}$ ensure that $\text{tr}(\Sigma_\ell)$ is monotonically decreasing in ℓ . The

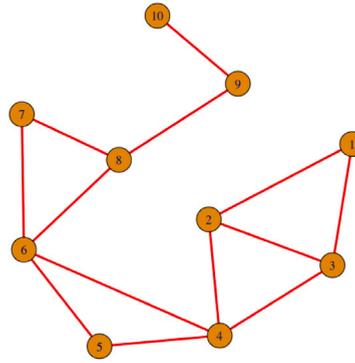


Fig. 1. The 10-variable graph used to generate $q = 10$ -variate GPs adhering to the conditional independence structure from the graph.

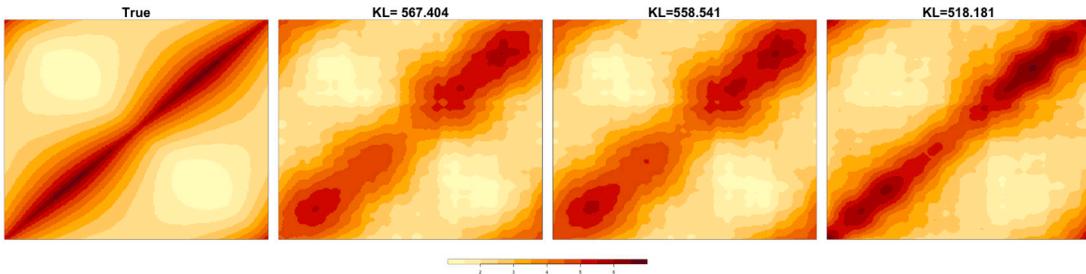
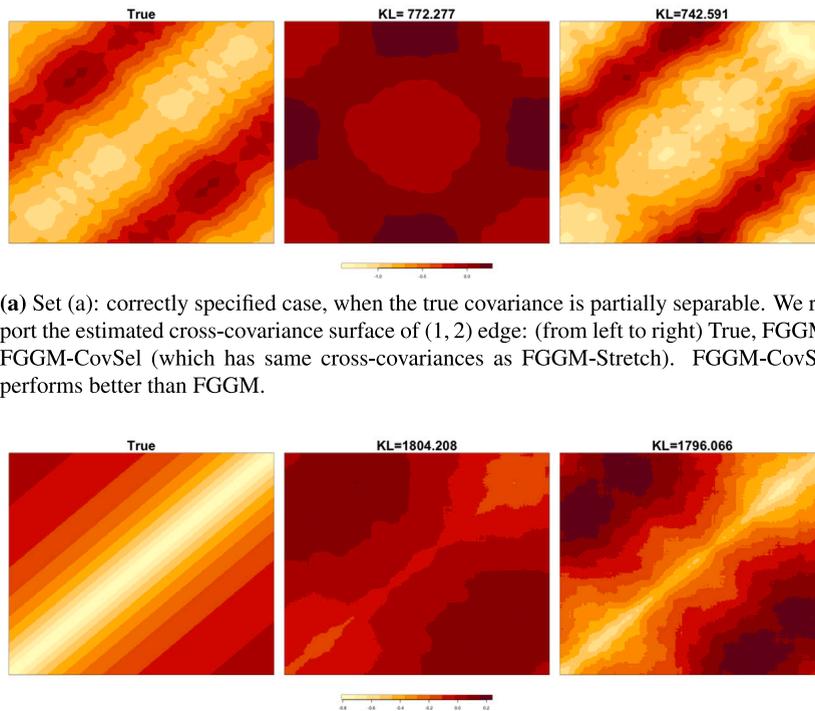


Fig. 2. Marginal covariance surfaces for variable 1: True, FGGM, FGGM-CovSel, FGGM-Stretch (from left to right). The KL-divergence from the true covariance matrix is noted at the top of the respective panels for that method. We see that FGGM-Stretch has a significantly improved estimate of the marginal covariance compared to both FGGM and FGGM-CovSel.

covariance function for the misspecified case was obtained by stitching using multivariate Matérn model [1] and the graph \mathcal{G} (Fig. 1). For every (i, j) pair of variables, the isotropic multivariate Matérn cross-covariance function on a d -dimensional domain is defined as $C_{ij}(s, s') = \sigma_{ij} H_{ij}(\|s - s'\|)$, where $H_{ij}(\cdot) = H(\cdot \mid v_{ij}, \phi_{ij})$, H being the Matérn correlation function [1]. $\sigma_{ij}, \phi_{ij}, v_{ij}$ are interpreted as the cross-covariance, cross-scale, and cross-smoothness parameters. To ensure positive-definiteness of cross-covariance matrix, we parametrize as $\sigma_{ij} = b_{ij}(\Gamma(\frac{1}{2}(v_{ii} + v_{jj} + d))\Gamma(v_{ij}))/(\phi_{ij}^{2\Delta_A + v_{ii} + v_{jj}} \Gamma(v_{ij} + \frac{d}{2}))$, where $b_{ij} = (\sigma_{ii}\sigma_{jj})^{0.5} \phi_{ii}^{v_{ii}} \phi_{jj}^{v_{jj}} r_{ij}^2 / \Gamma(v_{ij})$, and $R = (r_{ij}) > 0$, is correlation matrix. For this simulation, we take $v_{ij} = v_{ii} = v_{jj} = 0.5$, and $\Delta_A = 0$ and $\phi_{ij}^2 = (\phi_{ii}^2 + \phi_{jj}^2)/2$. The marginal range parameters ϕ_{ii} and variance parameters σ_{ii} were permutations of equispaced numbers in $(1, 5)$, while the $R = r_{ij}$'s were chosen as a random correlation matrix. For data generation locations, we take 250 points on the real line equispaced between $(0, 1)$.

We compare the estimated marginal and edge-specific cross-covariance surfaces qualitatively using some heat-maps and quantitatively using the distance between true and estimated covariance sub-matrices in terms of Kullback–Leibler (KL) divergence. The KL divergence between two matrices A and B is defined as $-\text{KL}(B \parallel A) = 0.5(\text{tr}(A^{-1}B) - \ln(\frac{|B|}{|A|}))$. In case (a), which is the truly specified case, we present the results comparing FGGM-CovSel and FGGM-Stretch with the FGGM algorithm of [41] (implemented using the R package *FGM* [28]). We report the graph estimated through FGGM and how it compares with the true graph in Fig. 7 in Appendix. The figure shows that FGGM estimates the graph better in a correctly specified setting (set (a)) than in a misspecified setting (set (b)). We first compare the true and estimated marginal covariance functions under the correctly specified setting. In Fig. 2, we plot the true marginal covariance for variable 1 and the estimated covariance from each of the three methods. We see that FGGM-Stretch has a significantly improved estimate of the marginal covariance compared to both FGGM and FGGM-CovSel. The FGGM and FGGM-CovSel rely on a low-rank truncation, which evidently leads to oversmoothing, while FGGM-Stretch mitigates this issue to a great extent. This aligns with our theoretical insights from Section 3.2. The result holds for the marginal covariance of all the variables as demonstrated in Fig. 6 in Appendix.

We then look at the estimates of the cross-covariance surfaces in Fig. 3. These will be the same for the proposed algorithms FGGM-CovSel and FGGM-Stretch, which differ only on the marginal covariances. We see that, for the correctly specified case, this estimate from our proposed algorithms performs better than the FGGM estimate (Fig. 3(a)) and estimates the true cross-covariance surface with high accuracy. This behavior is expected since the model is correctly specified for our algorithms, and the FGGM algorithms do not utilize knowledge of the graph but estimate it. Fig. 3(b) presents the cross-covariances for the mis-specified scenario. We see that FGGM-CovSel estimates the true covariance function less accurately than in the correctly specified case but does capture the broad features like the banding (decay of the Matérn cross-covariance away from the diagonal). This decrease in estimation accuracy is expected since we are trying to fit a misspecified partial separable cross-covariance to a non-partially separable GP. However, even in the misspecified case, FGGM-CovSel still significantly outperforms FGGM.



(a) Set (a): correctly specified case, when the true covariance is partially separable. We report the estimated cross-covariance surface of (1, 2) edge: (from left to right) True, FGM, FGM-CovSel (which has same cross-covariances as FGM-Stretch). FGM-CovSel performs better than FGM.

(b) Set (b): mis-specified case, when the true covariance is Graphical Matérn. We report the estimated cross-covariance surface of (1, 2) edge: (from left to right) True, FGM, FGM-CovSel (which has same cross-covariances as FGM-Stretch). FGM-CovSel estimates the true covariance function less accurately than in the correctly specified case but does capture the broad features like the banding (decay of the Matérn cross-covariance away from the diagonal).

Fig. 3. Comparison of (1,2) variable pair cross-covariance surface heatmaps under set (a) (correctly specified) and set (b) (misspecified) scenario. The KL-divergence for each method is noted at the top of the respective panel.

Overall, across all edges and in both the correctly specified and misspecified scenarios (across 25 different seeds), FGM-CovSel reports uniformly lower KL divergence between true and estimated covariance matrices than FGM (Table 1). Each row in Table 1 corresponds to the KL divergence mean (sd) across 25 replicates between edge-specific true and estimated covariance submatrices. We observe that FGM-CovSel performs according to our expectations in the correctly specified example, thus validating the theoretical basis of our approach. Moreover, in the correctly specified setting, we perform additional simulation experiments with increasing sample size (N varies from 50 to 5000) to demonstrate the asymptotical consistency of FGM-CovSel estimator. In Fig. 8 in Appendix, we see decrease in the Frobenius norm between true and estimated covariance with both increasing number of basis functions and number of samples (N). The better performance of FGM-CovSel in the misspecified example (set (b)) gives us confidence in our approach even when the partial separability assumption does not hold. This is in addition to the benefit demonstrated in Fig. 2 of using FGM-Stretch to better capture the marginal distributions.

6. Analysis of functional neuroimaging data

We illustrate the utility of the approach on multivariate functional Magnetic Resonance Imaging (fMRI) data from a study of thermal pain. The data is at the brain-region level (each variable corresponds to a region), and a graph between the variables represents the connectivity between different brain regions. In this study, the left forearm of 33 healthy right-handed participants was exposed to various levels (temperatures) of noxious heat stimuli. Participants gave their informed consent, and the Columbia University Institutional Review Board approved the study. Each participant completed seven runs within a session, consisting of 58–75 trials. During each trial, thermal stimuli were applied to the inner surface of the left forearm. Each stimulus lasted 12.5 s, including a three-second ramp-up, two-second ramp-down, and 7.5 s at the target temperature.

We collected a total of 1845 functional images ($TR=2000$ ms) from each participant using a 3T Philips Achieva TX scanner at Columbia University. Structural images were obtained using high-resolution T1 spoiled gradient recall (SPGR) imaging, while functional echo planar images (EPIs) were acquired with specific parameters; see [37] for details. The structural images were aligned with the mean functional image using an iterative mutual information-based algorithm in SPM8 and then normalized to Montreal

Table 1

The table presents the KL divergence, reported as mean (sd) across 25 replicates, which quantifies the difference between the true and estimated cross-covariance matrices for all edges in the graph. The results compare the performance of the FGGM and FGGM-CovSel methods under simulation scenarios: (a) and (b).

| Edges | Set (a) | | Set (b) | |
|---------|----------------|----------------|-----------------|-----------------|
| | FGGM | FGGM-CovSel | FGGM | FGGM-CovSel |
| (1, 2) | 787.67 (23.03) | 778.05 (22.35) | 1863.79 (36.12) | 1846.51 (35.38) |
| (1, 3) | 766.83 (12.59) | 756.66 (11.19) | 1832.1 (30.07) | 1806.42 (32.86) |
| (2, 3) | 796.32 (18.35) | 784.23 (22.14) | 1799.54 (31.76) | 1785.05 (34.07) |
| (2, 4) | 730.71 (31.39) | 722.14 (30.57) | 1840.36 (41.73) | 1833.68 (40.06) |
| (3, 4) | 730.91 (31.67) | 722.5 (30.02) | 1764.23 (35.58) | 1736.83 (38.57) |
| (4, 5) | 719.63 (22.99) | 704.54 (20.41) | 1974.63 (34.74) | 1949.14 (32.47) |
| (4, 6) | 738.21 (31.77) | 721.55 (33.27) | 1789.22 (36.9) | 1776.31 (36.65) |
| (5, 6) | 752.9 (32.11) | 742.87 (28.27) | 1793.26 (29.31) | 1781.21 (29.65) |
| (6, 7) | 781.46 (14.92) | 759.19 (20.57) | 1293.24 (32.4) | 1283.55 (31.91) |
| (6, 8) | 778.68 (19.86) | 747.99 (23.62) | 1409.48 (31.34) | 1394.82 (30.09) |
| (7, 8) | 786.81 (21.97) | 751.28 (22.99) | 1910.51 (28.87) | 1880.7 (30.58) |
| (8, 9) | 771.1 (26.52) | 751.06 (32.81) | 2004.36 (24.73) | 1982.75 (29.94) |
| (9, 10) | 771.71 (23.98) | 749.55 (24.65) | 1548.35 (36.78) | 1541.75 (36.25) |

Neurological Institute (MNI) space using SPM8's generative segment-and-normalize algorithm. Before preprocessing the functional images, the first four volumes were discarded to ensure image intensity stabilization. Outliers were identified using the Mahalanobis distance, and slice-timing differences were corrected. Motion correction was performed using SPM8, and images were transformed to SPM's normative atlas using warping parameters estimated from the aligned structural images. Finally, the functional images were smoothed with an 8 mm full width at half maximum (FWHM) Gaussian kernel and subjected to a high-pass filter with a cutoff of 180 s. For a more comprehensive description of the data acquisition and preprocessing methods, please refer to the work by [37].

We used a variant of the Yeo atlas [40] to first subdivide the brain into 286 separate non-overlapping brain regions, and thereafter organized these regions into 18 functional networks (17 cortical and one subcortical). These were further arranged into Visual (VIS), Motor (MOT), Dorsal Attention (DAN), Ventral Attention (VAN), Frontoparietal (FP), Limbic (LIM), Default Mode (DMN), and Subcortical (SC) networks. Our framework models this dataset as a 286-dimensional functional data (corresponding to the $q = 286$ brain regions) for $n = 33$ subjects. For computational efficiency, we reduce the number of time points from 1845 to 184 by binning observation within a 10 time unit interval.

Our goal is to use this dataset for a proof-of-concept analysis to judge the covariance estimation performance of our method FGGM-CovSel, which uses the knowledge of a preestimated connectivity graph among the 286 regions, compared to the FGGM method, which will re-estimate the graph while estimating the covariance function. In previous fMRI studies, partial separability has been both assumed and confirmed [41]. In a real-world application, one can assess this assumption by verifying whether the correlations among basis coefficients at different levels for different processes are close to zero. Given that our analysis serves only as a proof-of-concept, that there is past evidence supporting partial separability in fMRI data, and that our simulations demonstrated robustness of our approach to violation of the partial separability assumption, we opt not to scrutinize this assumption in this paper to maintain brevity in the discussion. We use a random subset of the data to estimate the connectivity graph. We randomly split the group of subjects into two sets of roughly equal size - (i) train (17 subjects) and (ii) test (16 subjects). On the training set, we use FGGM of [41] to learn the graph between brain regions and estimate the covariance function. Then on the test set, we use both FGGM and FGGM-CovSel to estimate covariance. The FGGM reestimates a (possibly) different graph on the test set, while FGGM-CovSel conducts graph-constrained analysis using the graph estimated from the training set. We then compare the covariance estimate of FGGM and FGGM-CovSel in the test set with the FGGM estimate obtained in the train set. We perform this random split 10 times to obtain cross-validated results.

Fig. 4 visualizes the network between brain regions from the most probable graph across 10 splits. We average the adjacency matrix across the different data splits and keep the edges in the final graph, as the pairs which have an average edge probability bigger than 0.5. Pain is known to engage multiple brain networks. In our estimated brain network, we found connections among regions within VIS, DAN, VAN, FP, LIM, and DMN. For example, the DAN may become engaged when individuals are asked to direct attention to pain-related stimuli or when they are actively engaged in pain processing tasks, and the FP network processes subjective and emotional aspects of pain and pain perception [29], the DMN modulates pain perception [2], and the limbic network mediates self-regulation of pain [37]. In addition, subcortical regions regulate pain perception and processing through autonomic responses [6] and the somatomotor network is involved in the perception and processing of pain-related sensory input [22]. In addition, connections were found between DMN, FP, LIM, DAN, and VAN, with some engagement with MOT. Particularly, we see connections between the DMN, DAN, and the FP network. In the context of pain tasks, both these relationships depend on factors such as task demands, attentional focus, and individual differences in pain processing and coping strategies. In general, these networks and connections are consistently activated by experimental pain, illustrating the method's efficacy.

We then assess how our graph-constrained analysis (FGGM-CovSel) using this graph compares with FGGM in terms of estimating the covariance function. Let \hat{C}_{train} denote the estimated covariance matrix on the training dataset, and $\hat{C}_{\text{train}}(i, j)$ denote the sub-matrix corresponding to the nodes i and j . Similarly, we denote by $\hat{C}_{\text{test}, m}$ the estimated covariance matrix on the test set by method

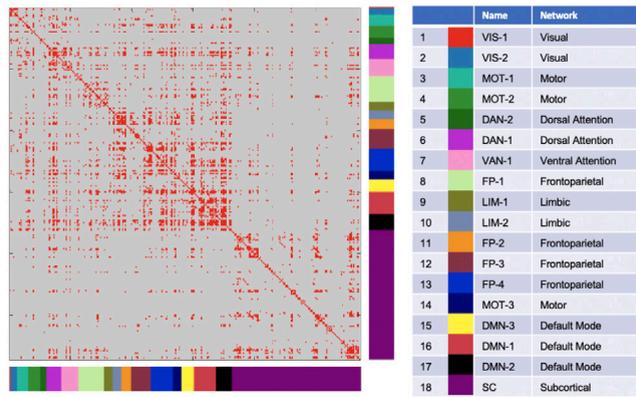


Fig. 4. The final adjacency matrix, consisting of the pairs with an average edge probability larger than 0.5. The brain is subdivided into 286 regions using a variant of the Yeo atlas [40]. These regions are organized into 18 networks. These can be further arranged into Visual (VIS), Motor (MOT), Dorsal Attention (DAN), Ventral Attention (VAN), Frontoparietal (FP), Limbic (LIM), Default Mode (DMN), and Subcortical (SC) networks. The legend to the right illustrates the colors associated with each subnetwork.

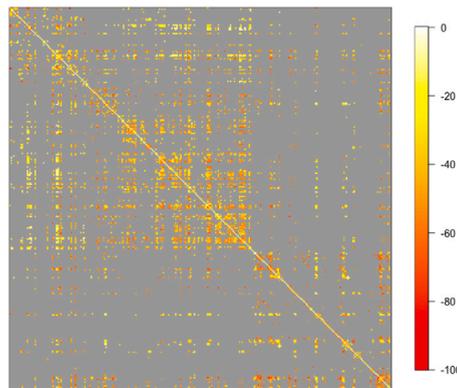


Fig. 5. The difference in KL-distances between edge-specific covariances for FGGM-CovSel (test) & FGGM (train) and FGGM (test) & FGGM (test). The difference is aggregated across 10 replicates whenever that specific edge-pair is detected. Gray cells indicate pairs that did not appear in any edge in 10 replicates.

$m, m \in \{\text{FGGM}, \text{FGGM-CovSel}\}$ and define $\hat{C}(i, j)_{\text{test}, m}$ accordingly. We then consider the KL distance $d_{KL}(\hat{C}(i, j)_{\text{train}}, \hat{C}(i, j)_{\text{test}, m})$ for all edges (i, j) in the graph. A lower KL distance implies a better estimation of the covariance. Fig. 5 plots a heatmap of the difference

$$d_{KL}(\hat{C}(i, j)_{\text{train}}, \hat{C}(i, j)_{\text{test}, \text{FGGM-CovSel}}) - d_{KL}(\hat{C}(i, j)_{\text{train}}, \hat{C}(i, j)_{\text{test}, \text{FGGM}}).$$

The result is aggregated across the replicates in which a specific pair has been detected as an edge by FGGM (train). We observe in Fig. 5 that all the values of this difference are negative, implying that the covariance estimate from FGGM-CovSel (test) has a lower KL distance with FGGM (train) than FGGM (test) has with FGGM (train). This is true across all edge-pairs that were included in the connectivity graph. Across 10 different seeds, the KL distance between FGGM-CovSel (test set) and FGGM (train) is 36.8(35, 39.9) units lower than the distance between FGGM (test) and FGGM (train) when averaged across all edge pairs. These findings show us that estimating the covariance function benefits from the knowledge of the graph and demonstrates the utility of a graph-constrained multivariate functional analysis.

7. Discussion

The current manuscript has developed graph-constrained modeling and analysis for multivariate functional data. We build upon recent developments in graphical Gaussian processes (GGP) in the context of highly multivariate spatial data analysis and draw some insightful connections of such frameworks with functional Gaussian graphical models. Our key result is an optimal partially separable functional Gaussian graphical model that is derived using the graphical Gaussian process and preserves the conditional dependencies, or lack thereof, posited by a given undirected graph. Our methodological developments are accompanied by novel theoretical insights connecting the two seemingly disparate paradigms of functional Gaussian graphical models and Graphical Gaussian processes. These theoretical developments enable us to construct a new algorithm to evaluate the maximum likelihood estimate for a partially

separable graphical Gaussian model while preserving the conditional dependence constraints imposed by the given graph. This algorithm constitutes the generalization of Dempster's covariance selection from vector-valued data to functional data. Beyond the population level theoretical connections used to derive the methodology, we also prove the consistency of the covariance operator estimator obtained from the resulting algorithm. We acknowledge that this consistency result is only established in a simplistic setting with the curves (functional data) observed at a fixed set of points and the basis functions defined on this discrete domain. Future work needs to refine the theoretical results by filling the remaining gaps between the current theory and the practice. These would need to consider basis functions on a continuous domain and the curves being sampled at a finite set of points in this domain in an increasingly dense manner, i.e., in-fill asymptotics. The use of data-driven basis functions from a functional PCA must also be accommodated. Besides consistency, asymptotic normality results would also be needed to conduct inference on the estimated covariance functions.

We note that while the manuscript focuses on the estimation of the multivariate covariance function under graphical constraints, assuming the mean of the functional processes to be zero, a non-zero mean can be easily accommodated in the framework. This will be suitable in applications where other covariates are observed in addition to the functional responses, and a mean, specified as a linear or non-linear function of these covariates, can be easily estimated as part of the algorithm. Also, our paper focuses on the analysis of multivariate functional data when the within-variable graph is known with certainty. In examples, where the graph is known with some degree of uncertainty, a Bayesian approach would be more suitable, incorporating information about graph uncertainty as a prior for the graphical model. We would then need to adapt our approach to Bayesian reversible jump MCMC algorithms [5,19] similar to what has been done in Section 4.3 of [12]. This will be an important future direction.

Our data analysis was conducted under the assumption of partial separability, which necessitates a discussion regarding the practical validity of this assumption. This assumption was motivated by the analysis of similar fMRI data in [41]. However, due to our limited sample size, we cannot explicitly test this assumption in our dataset, as [41] performed in their paper. Hence, we have undertaken an empirical assessment to evaluate the robustness of our methodology under potential violations of the partial separability assumption (Simulation set (b)). This assessment is of utmost importance as it clearly explains our approach's performance in adverse scenarios. While performance indeed declines in scenarios where the assumption is not met, our approach still outperforms the FGGM methods, which also rely on partial separability. It is important to note that when the partial separability assumption is severely violated, FGGM methods such as those proposed by [42] may be employed. However, these methods require joint estimation of the entire precision matrix of all regression coefficients, which can lead to significant computational challenges.

Given the continued interest in high-dimensional statistical inference for multivariate functional data, this work is expected to spur further research in scaling up the proposed models here to massive numbers of functional inputs. Specific examples include settings where spatial data are measured on a very large number of dependent outcomes over massive numbers of locations. In this regard, it will be possible to embed Gaussian processes, e.g., for hierarchical Nearest Neighbor Gaussian process models, [9,10,15] within the nodes of the graph to scale up to the massive numbers of locations.

While the conceptual idea of building scalable multivariate models with stochastic processes as nodes of a graph follows from graphical Gaussian processes [12] or from approaches for versatile classes of multivariate Markov random fields [20], substantial developments on precise model specifications and, perhaps even more importantly, algorithmic developments are still warranted. In this regard, we expect our current developments to generate future research pursuits that will culminate in a versatile, computationally efficient framework for graph-constrained analysis of high-dimensional dependent functional or spatial data. Applications of such a framework will span biomedical and health sciences, broader environmental sciences, social sciences, and economics.

CRedit authorship contribution statement

Debangana Dey: Methodology, Software, Formal analysis, Writing – original draft preparation. **Sudipto Banerjee:** Methodology, Writing – review & editing. **Martin A. Lindquist:** Writing – review & editing, Formal analysis. **Abhirup Datta:** Conceptualization, Methodology, Writing – original draft preparation, Funding acquisition, Supervision.

Acknowledgment

AD was partially supported by the National Science Foundation (NSF) Division of Mathematical Sciences grant DMS-1915803. AD and ML were partially supported by National Institute of Biomedical Imaging and Bioengineering award 2R01EB026549-05. SB was supported by the National Science Foundation (NSF) awards DMS-1916349 and DMS-2113778 and by the National Institute of General Medical Sciences (NIH-NIGMS) awards R01GM148761-01. DD undertook this project during his PhD at Johns Hopkins Bloomberg School of Public Health and was supported by a fellowship through a Joint Graduate Training Program, established in collaboration between the Department of Biostatistics at the Johns Hopkins Bloomberg School of Public Health and the Intramural Research Program of the National Institute of Mental Health. The authors thank Dr. Brian Caffo for helpful discussions.

Appendix A. Proofs

Proof of Theorem 1. Denote $\Sigma_\ell = (\sigma_{\ell ij})$, $\Sigma_\ell^* = (\sigma_{\ell ij}^*)$, and $\Sigma_\ell^{*-1} = \Omega_\ell = (\omega_{\ell ij}^*)$. The cross-covariance function of $w(\cdot)$ is denoted by $C^* = (C_{ij}^*)$ and the conditional cross-covariance function is $K^* = (K_{ij}^*)$. By Theorem 3 of [41],

$$K_{ij}^*(s, t) = \sum_{\ell=1}^{\infty} \frac{-\omega_{\ell ij}^*}{\omega_{\ell ii}^* \omega_{\ell jj}^* - \omega_{\ell ij}^{*2}} \phi_\ell(s) \phi_\ell(t). \tag{A.1}$$

(If part:) Suppose $\Sigma_\ell^* = \text{CovSel}(\Sigma_\ell, \mathcal{G})$ for all ℓ . To prove $w(\cdot) \sim GGP(C, \mathcal{G})$ we first show it is a graphical GP. For any $(i, j) \notin E$, by the property of covariance selection on each Σ_ℓ , we have $\omega_{\ell ij}^* = 0$ for all ℓ . Hence, from (A.1) $K_{ij}^*(s, t) = 0$ for all $s, t \in \mathcal{D}$, $(i, j) \notin E$ and $w(\cdot)$ is a graphical GP with respect to the graph \mathcal{G} . Next we show $w(\cdot)$ retains the marginals and edge-specific cross-covariances from C . This is immediate as once again for all $(i, j) \ni i = j$ or $(i, j) \in E$, we have by the property of covariance selection $\sigma_{\ell ij}^* = \sigma_{\ell ij}$ and consequently $C_{ij}^*(s, t) = \sum_{\ell=1}^{\infty} \sigma_{\ell ij}^* \phi_\ell(s) \phi_\ell(t) = \sum_{\ell=1}^{\infty} \sigma_{\ell ij} \phi_\ell(s) \phi_\ell(t) = C_{ij}(s, t)$.

(Only if part:) Let $w(s) = \sum_{\ell=1}^{\infty} \theta_\ell^* \phi_\ell(s) \sim GGP(C, \mathcal{G})$. Since it is a graphical GP, then $K_{ij}^*(s, t) = 0$ for all $s, t \in \mathcal{D}$ and $(i, j) \notin E$. By orthogonality of ϕ_ℓ 's,

$$0 = \int K_{ij}^*(s, t) \phi_\ell(s) ds = \frac{-\omega_{\ell ij}^*}{\omega_{\ell ii}^* \omega_{\ell jj}^* - \omega_{\ell ij}^{*2}} \phi_\ell(t).$$

This holds for all $t \in \mathcal{D}$ implying $\omega_{\ell ij}^* = 0$ for all ℓ . By Definition 2.1, $w(\cdot)$ retains the marginal and edge-specific cross-covariances. Hence, for $i = j$ or $(i, j) \in E$, $C_{ij}^* = C_{ij}$ implying

$$\sigma_{\ell ij} \phi_\ell(t) = \int C_{ij}(s, t) \phi_\ell(s) ds = \int C_{ij}^*(s, t) \phi_\ell(s) ds = \sigma_{\ell ij}^* \phi_\ell(t).$$

This holds for all t implying $\sigma_{\ell ij} = \sigma_{\ell ij}^*$ for all ℓ , $i = j$ or $(i, j) \in E$. Hence, Σ_ℓ^* satisfies all the conditions of covariance selection (Definition 2.2) and is the unique matrix $\text{CovSel}(\Sigma_\ell, \mathcal{G})$. \square

Proof of Corollary 1. We apply Theorem 3 of [41] to the process $w^m(s)$ and write the conditional cross-covariance function in terms of the elements of $\Sigma_\ell^{*-1} = \Omega_\ell = (\omega_{\ell ij}^*)$ and the basis functions as

$$\begin{aligned} K_{ij}^m(s, t) &= \text{Cov}(w_i^m(s), w_j^m(t) \mid w_{-(i,j)}^m) = \text{Cov} \left(\sum_{\ell=1}^m \theta_{\ell i}^* \phi_\ell(s), \sum_{\ell'=1}^m \theta_{\ell' j}^* \phi_{\ell'}(t) \mid w_{-(i,j)}^m \right) \\ &= \sum_{\ell, \ell'=1}^m \text{Cov} \left(\theta_{\ell i}^*, \theta_{\ell' j}^* \mid w_{-(i,j)}^m \right) \phi_\ell(s) \phi_{\ell'}(t) = \sum_{\ell=1}^m \text{Cov} \left(\theta_{\ell i}^*, \theta_{\ell j}^* \mid \theta_{\ell, -(i,j)}^* \right) \phi_\ell(s) \phi_\ell(t) = \sum_{\ell=1}^m \frac{-\omega_{\ell ij}^*}{\omega_{\ell ii}^* \omega_{\ell jj}^* - \omega_{\ell ij}^{*2}} \phi_\ell(s) \phi_\ell(t), \end{aligned}$$

where $w_{-(i,j)}^m$ is the full (uncountable) realization of all the component processes in $w^m(s)$ except i and j . The third equality follows from the fact that the coefficients θ_ℓ and the process $w^m(s)$ uniquely identify each other given the m basis functions. The proof now follows that of Theorem 1. \square

Here we show the approximate equivalence of stitching and low-rank truncation. We observe that $w_{stitch}^{pp}(s) = D(s)w_{stitch}(\mathcal{L})$ where $D(s) = \text{diag}(C_{11}(s, \mathcal{L})C_{11}(\mathcal{L}, \mathcal{L})^{-1}, \dots, C_{qq}(s, \mathcal{L})C_{qq}(\mathcal{L}, \mathcal{L})^{-1})$ and $w_{stitch}(\mathcal{L})$ is a random variable whose distribution is specified in (1). Next, as $w(\cdot) \sim GGP(C, \mathcal{G})$, the component processes $w_j(\cdot) \sim GP(0, C_{jj})$ retain the marginals. When the reference set \mathcal{L} is chosen to have m locations, the truncated process $w_j^m(\cdot)$ is related to the untruncated process $w_j(\cdot)$ via the following approximation (see, p. 389–390 in [3, for a derivation])

$$w_j^m(s) \approx C_{jj}(s, \mathcal{L})C_{jj}(\mathcal{L}, \mathcal{L})^{-1}w_j(\mathcal{L}) \text{ for all } s \in \mathcal{D}, j \in \mathcal{V}. \tag{A.2}$$

The approximate equivalence in (A.2) is related to the Nystrom approximation for kernel matrices [14]. Rewriting (A.2) as $w^m(s) \approx D(s)w(\mathcal{L})$ and since $w_{stitch}^{pp}(s) = D(s)w_{stitch}(\mathcal{L})$, to show $w^m(\cdot) \approx w_{stitch}^{pp}(\cdot)$, we now only need $w(\mathcal{L})$ and $w_{stitch}(\mathcal{L})$ to approximately have the same distribution. As $w(\cdot) \sim GGP(C, \mathcal{G})$, for $i = j$ or $(i, j) \in E$, we have $\text{Cov}(w_i(\mathcal{L}), w_j(\mathcal{L})) = C_{ij}(\mathcal{L}, \mathcal{L}) = \text{Cov}(w_{stitch,i}(\mathcal{L}), w_{stitch,j}(\mathcal{L}))$ by definition of covariance selection. Finally, we show that for $(i, j) \notin E$, $\text{Cov}(w_i(\mathcal{L}), w_j(\mathcal{L}) \mid w_{-(i,j)}(\mathcal{L})) \approx O = \text{Cov}(w_{stitch,i}(\mathcal{L}), w_{stitch,j}(\mathcal{L}) \mid w_{stitch,-ij}(\mathcal{L}))$ as follows:

$$\begin{aligned} &\text{Cov}(w_i(\mathcal{L}), w_j(\mathcal{L}) \mid w_{-(i,j)}(\mathcal{L})) \\ &= \text{Cov} \left(C_{ii}(\mathcal{L}, \mathcal{L})C_{ii}(\mathcal{L}, \mathcal{L})^{-1}w_i(\mathcal{L}), C_{jj}(\mathcal{L}, \mathcal{L})C_{jj}(\mathcal{L}, \mathcal{L})^{-1}w_j(\mathcal{L}) \mid C_{kk}(\mathcal{L}, \mathcal{L})C_{kk}(\mathcal{L}, \mathcal{L})^{-1}w_k(\mathcal{L}), k \in \mathcal{V} \setminus \{i, j\} \right) \\ &\approx \text{Cov}(w_i^m(\mathcal{L}), w_j^m(\mathcal{L}) \mid w_{-ij}(\mathcal{L})) \text{ by (A.2)} \\ &= \text{Cov}(w_i^m(\mathcal{L}), w_j^m(\mathcal{L}) \mid C_{kk}(s, \mathcal{L})C_{kk}(\mathcal{L}, \mathcal{L})^{-1}w_k(\mathcal{L}), s \in \mathcal{D}, k \neq i, j) \\ &\approx \text{Cov}(w_i^m(\mathcal{L}), w_j^m(\mathcal{L}) \mid w_{-ij}^m(\cdot)) \text{ again by (A.2)} \\ &= O \text{ for } (i, j) \notin E. \end{aligned}$$

The middle equality follows from the fact that for any $s \in \mathcal{L}$ and any k , we have $C_{kk}(s, \mathcal{L})C_{kk}(\mathcal{L}, \mathcal{L})^{-1}w_k(\mathcal{L}) = w_k(s)$. Hence, the σ -algebras $\sigma(w_{-ij}(\mathcal{L}))$ and $\sigma(C_{kk}(s, \mathcal{L})C_{kk}(\mathcal{L}, \mathcal{L})^{-1}w_k(\mathcal{L}), s \in \mathcal{D}, k \neq i, j)$ becomes the same since the latter is generated by the former and contains the former. Corollary 1 ensures the last equality since $w^m(\cdot)$ is a graphical GP conforming to \mathcal{G} . We have now shown that

the covariances of $w(\mathcal{L})$ and $w_{stitch}(\mathcal{L})$ agree exactly on the entries corresponding to $i = j$ or $(i, j) \in E$, and the inverse-covariances agree approximately on the entries corresponding to $(i, j) \notin E$. By uniqueness of Dempster’s covariance selection for positive definite matrices, $w(\mathcal{L})$ and $w_{stitch}(\mathcal{L})$ have approximately the same distribution. This establishes the equivalence of the m -truncated partially separable FGGM $w^m(s) = D(s)w(\mathcal{L})$ with a stitched predictive process $w_{stitch}^{pp}(s) = D(s)w_{stitch}(\mathcal{L})$ on m reference locations.

Proof of Theorem 2. We have $X_i \sim N_{pq}(0, C^m)$ for $i \in \{1, \dots, N\}$, where $C^m = \sum_{\ell=1}^m \Sigma_\ell \otimes \phi_\ell \phi_\ell^\top$ is the true multivariate covariance function of the partially separable process. We use X_{ij} to denote the p -variate vector observed for i th replicate and j th variable. Then, the sample covariance matrix $S_{(pq \times pq)}$ can be calculated as $S = \frac{1}{N} \sum_{i=1}^N (X_{i1}^\top, \dots, X_{iq}^\top)^\top (X_{i1}^\top, \dots, X_{iq}^\top) = \frac{1}{N} \sum_{i=1}^N X_i X_i^\top$. Under the constraint of a given graphical model \mathcal{G} , we want to find the maximum likelihood estimate (MLE) of the parameters Σ_ℓ ’s of the partially separable process.

First we will derive the MLE of a multivariate partially separable covariance function without the graphical constraint. Under partial separability, we can write $X_i = \sum_{\ell=1}^m \theta_{i\ell} \otimes \phi_\ell$, where $\theta_{i\ell} \sim N(0, \Sigma_\ell)$. Using orthonormality of the basis functions, we can transform the data as follows $(I_q \otimes \phi_\ell^\top) X_i = \theta_{i\ell}$. Here, the X_i ’s (jointly over i) and the $\theta_{i\ell}$ ’s (jointly over i, ℓ) are a 1-1 and invertible linear transformation of each other.

It is easy to check that $\theta_{i\ell}$ ’s are iid across i ’s and independent across ℓ ’s. For $i \neq j$, $\text{Cov}(\theta_{i\ell}, \theta_{j\ell}) = E[(I_q \otimes \phi_\ell^\top) X_i X_j^\top (I_q \otimes \phi_\ell)] = 0$, since $E[X_i X_j^\top] = 0$. Now, for $\ell \neq k$, $\text{Cov}(\theta_{i\ell}, \theta_{ik}) = E[(I_q \otimes \phi_\ell^\top) X_i X_i^\top (I_q \otimes \phi_k)] = (I_q \otimes \phi_\ell^\top) C^m (I_q \otimes \phi_k) = (I_q \otimes \phi_\ell^\top) (\Sigma_k \otimes \phi_k) = 0$. Similarly, for $i \neq j, \ell \neq k$, $\text{Cov}(\theta_{i\ell}, \theta_{jk}) = 0$.

The derivation above can help us rewrite the likelihood in following way

$$f(X_1, \dots, X_N | \Sigma_\ell; \ell \in \{1, \dots, m\}) \propto f(\theta_{i\ell}; i \in \{1, \dots, N\}, \ell \in \{1, \dots, m\} | \Sigma_\ell; \ell \in \{1, \dots, m\}) \propto \prod_{\ell=1}^m \prod_{i=1}^N f(\theta_{i\ell} | \Sigma_\ell). \tag{A.3}$$

This result (A.3) is due to the fact that the Jacobian transformation in the likelihood going from the X_i ’s to the $\theta_{i\ell}$ ’s will only depend on known ϕ_ℓ ’s. So, finding the maximum likelihood estimate (MLE) of Σ_ℓ boils down to maximizing the transformed data-likelihood of iid multivariate normal variables $\theta_{i\ell}, i \in \{1, \dots, N\}$. Hence, we get the MLE of the parameters as $\hat{\Sigma}_\ell = S_\ell = \frac{1}{n} \sum_{i=1}^N \theta_{i\ell} \theta_{i\ell}^\top$. Rewriting $\theta_{i\ell}$ in terms of X_i , we have the MLE as $\hat{\Sigma}_\ell = (I_q \otimes \phi_\ell^\top) \frac{1}{N} \sum_{i=1}^N (X_i X_i^\top) (I_q \otimes \phi_\ell) = (I_q \otimes \phi_\ell^\top) S (I_q \otimes \phi_\ell)$. Thus the MLE of a multi-variate partially separable covariance function without any graphical constraint is given by

$$\hat{C}_{unconstrained}^m = \sum_{\ell=1}^m (I_q \otimes \phi_\ell^\top) S (I_q \otimes \phi_\ell) \otimes \phi_\ell \phi_\ell^\top \tag{A.4}$$

We now derive the MLE under the graphical constraint. Based on the results of Theorem 1 and Corollary 1, a natural guess would be that the MLE under the graphical constraint can be obtained by applying covariance selection on each of the estimates $\hat{\Sigma}_\ell = S_\ell = (I_q \otimes \phi_\ell^\top) S (I_q \otimes \phi_\ell)$ in (A.4). We prove this formally below.

If we assume the variables are conforming to a graph $\mathcal{G} = (V, E)$, i.e., $\Sigma_\ell \in M^+(\mathcal{G}) = \{S \in \mathbb{R}^{q \times q} : S_{ij} = 0 \forall (i, j) \notin E\}$ for all $\ell \in \{1, \dots, m\}$, then by [11], we get graph-constrained MLE by maximizing the likelihood in (A.3) under the constraint $\{\Sigma_1, \dots, \Sigma_m : \Sigma_\ell > 0, \Sigma_\ell \in M^+(\mathcal{G}) \forall \ell \in \{1, \dots, m\}\}$. Due to the product-form (across ℓ ’s) in (A.3), this becomes equivalent to finding the MLE for each of Σ_ℓ ’s separately under the graphical constraint. In other words, the MLE is given by

$$\text{argmax}_{\Sigma_\ell > 0, \Sigma_\ell \in M^+(\mathcal{G})} \prod_{i=1}^N f(\theta_{i\ell} | \Sigma_\ell) = \text{argmin}_{\Sigma_\ell > 0, \Sigma_\ell \in M^+(\mathcal{G})} [\text{tr}(\Sigma_\ell^{-1} S_\ell) + \ln(|\Sigma_\ell|)].$$

Thus the problem has exactly reduced to Dempster’s covariance selection problem, i.e., given iid multivariate realizations $\theta_{i\ell}$ finding the MLE for the covariance matrices Σ_ℓ under the graphical constraint. The solution is given by $\hat{\Sigma}_\ell = \text{CovSel}(S_\ell, \mathcal{G})$ which yields the MLE for C^m as given in (5). \square

Proof of Lemma 1. Let us denote, for N samples, the sample covariance matrix as S_N and the basis-level covariance matrix estimate for Σ_ℓ as $\hat{\Sigma}_{\ell N} = S_{\ell N} = (I_q \otimes \phi_\ell^\top) S_N (I_q \otimes \phi_\ell)$ and the corresponding covariance selection estimator as $\hat{\Sigma}_{\ell N} = \text{CovSel}(\hat{\Sigma}_{\ell N}, \mathcal{G})$. Let us define I-divergence between two matrices A and B as $I(B \| A) = 0.5(\text{tr}(A^{-1}B) - I) - \ln(A^{-1}B) = KL(B \| A) - \frac{q}{2}$, where A and B are $q \times q$ positive-definite matrices.

Since, the X_i ’s are iid, the sample covariance S_N is an asymptotically consistent estimator for population covariance C^m . Then $\hat{\Sigma}_{\ell N} = (I_q \otimes \phi_\ell^\top) S_N (I_q \otimes \phi_\ell)$ converges to $(I_q \otimes \phi_\ell^\top) C^m (I_q \otimes \phi_\ell) = \Sigma_\ell$ in probability as $N \rightarrow \infty$. Now we will show that, $\hat{\Sigma}_{\ell N} = \text{CovSel}(\hat{\Sigma}_{\ell N}, \mathcal{G})$ converges to Σ_ℓ in probability as well.

By Lemma 1 of [33], $\hat{\Sigma}_{\ell N} = \text{argmin}_{A: A > 0, A \in M^+(\mathcal{G})} I(A \| \hat{\Sigma}_{\ell N})$. Since, $\Sigma_\ell > 0, \Sigma_\ell \in M^+(\mathcal{G})$, we have $I(\hat{\Sigma}_{\ell N} \| \hat{\Sigma}_{\ell N}) \leq I(\Sigma_\ell \| \hat{\Sigma}_{\ell N})$. Using the fact that, $\hat{\Sigma}_{\ell N}$ converges to Σ_ℓ in probability, we get $\hat{\Sigma}_{\ell N}^{-1} \Sigma_\ell \rightarrow I_q$ in probability and $I(\Sigma_\ell \| \hat{\Sigma}_{\ell N}) \rightarrow 0$. Hence, $I(\hat{\Sigma}_{\ell N} \| \hat{\Sigma}_{\ell N}) \rightarrow 0$. Using Lemma 1, part (iii) of [33], we get, $\hat{\Sigma}_{\ell N} - \hat{\Sigma}_{\ell N} \rightarrow 0$ in probability. Using Slutsky’s theorem, we get $\hat{\Sigma}_{\ell N} = \text{CovSel}(\hat{\Sigma}_{\ell N}, \mathcal{G}) = \text{CovSel}(S_{\ell N}, \mathcal{G})$ converges to Σ_ℓ in probability.

Using Slutsky’s theorem again, we get $\hat{C}^m = \sum_{\ell=1}^m \hat{\Sigma}_{\ell N} \otimes \phi_\ell \phi_\ell^\top$ converges to $\sum_{\ell=1}^m \Sigma_\ell \otimes \phi_\ell \phi_\ell^\top = C^m$ in probability. \square

Appendix B. Supplementary figures

See Figs. 6–8.

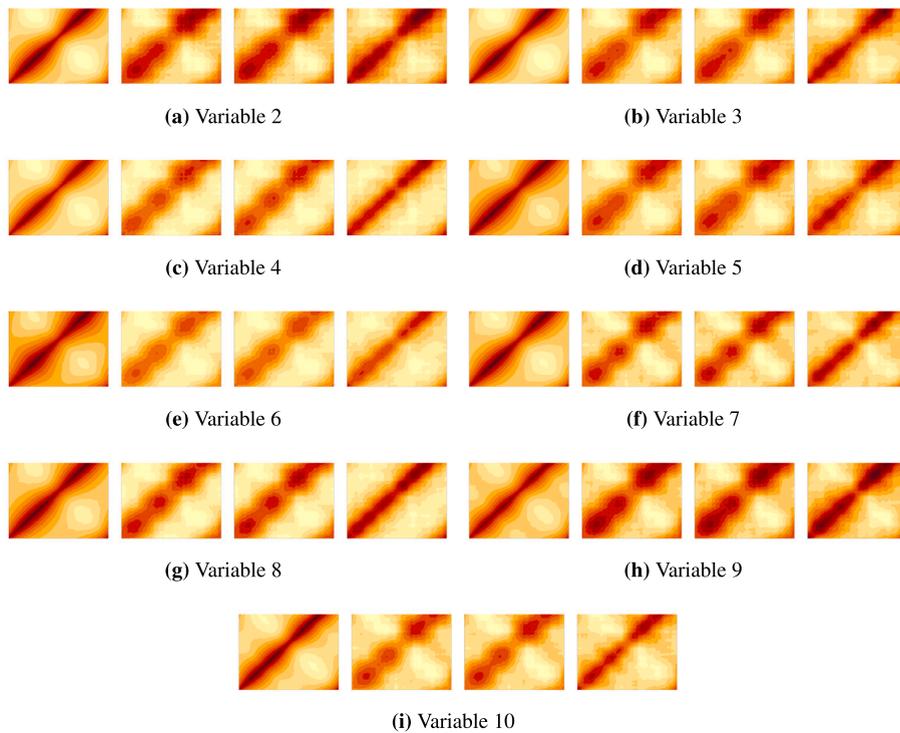


Fig. 6. Marginal covariance surfaces for variables 2 – 10 in Set (a): True, FGGM, FGGM-CovSel, FGGM-Stretch (from left to right). We see that FGGM-Stretch has a significantly improved estimate of the marginal covariance compared to both FGGM and FGGM-CovSel for variables 2 – 10.

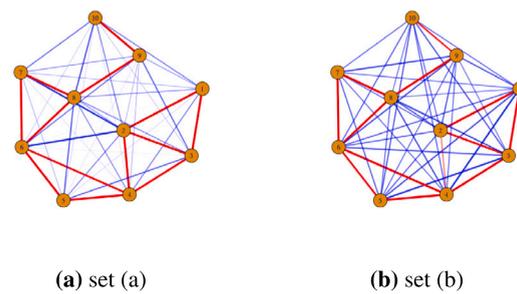


Fig. 7. Most probable graphs estimated through FGGM in set (a) (correctly specified) and set (b) (misspecified). Edge widths and opacity are proportional to edge selection probability (averaged over seeds). Edges in the true graph are colored red, otherwise the edges are colored blue. The figure shows that FGGM estimates the graph better in a correctly specified setting (set (a)) than in a misspecified setting (set (b)). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

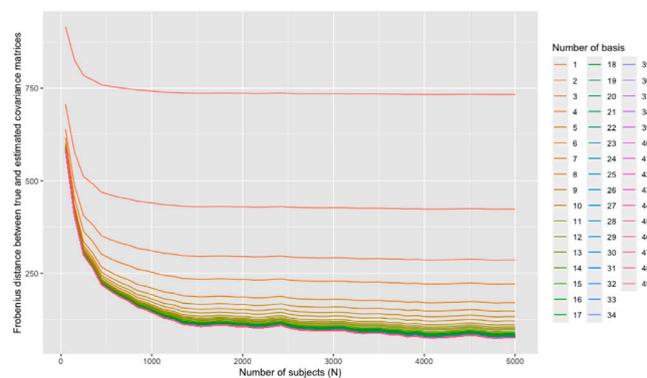


Fig. 8. The difference (in Frobenius norm) between true and FGGM-CovSel estimated covariance matrix (2000 * 2000) with increasing sample size and number of basis functions. We see a decrease in the Frobenius norm between true and estimated covariance with both increasing number of basis functions and number of samples (N).

References

- [1] T.V. Apanasovich, M.G. Genton, Y. Sun, A valid Matérn class of cross-covariance functions for multivariate random fields with any number of components, *J. Amer. Statist. Assoc.* 107 (497) (2012) 180–193.
- [2] M.N. Baliki, A.R. Mansour, A.T. Baria, A.V. Apkarian, Functional reorganization of the default mode network across chronic pain conditions, *PLoS One* 9 (9) (2014) e106133.
- [3] S. Banerjee, B.P. Carlin, A.E. Gelfand, *Hierarchical Modeling and Analysis for Spatial data*, Chapman and Hall/CRC, London, 2014.
- [4] S. Banerjee, A.E. Gelfand, A.O. Finley, H. Sang, Gaussian predictive process models for large spatial data sets, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 70 (4) (2008) 825–848.
- [5] R.J. Barker, W.A. Link, Bayesian multimodel inference by RJMCMC: A gibbs sampling approach, *Amer. Statist.* 67 (3) (2013) 150–156.
- [6] U. Bingel, M. Quante, R. Knab, B. Bromm, C. Weiller, C. Büchel, Subcortical structures involved in pain processing: evidence from single-trial fMRI, *Pain* 99 (1–2) (2002) 313–321.
- [7] C.M. Carvalho, H. Massam, M. West, Simulation of hyper-inverse wishart distributions in graphical models, *Biometrika* 94 (3) (2007) 647–659.
- [8] N.A.C. Cressie, C.K. Wikle, *Statistics for spatio-temporal data*, in: *Wiley Series in Probability and Statistics*, Wiley, Hoboken, NJ, 2011.
- [9] S. Banerjee, A.E. Gelfand, A.O. Finley, A.E. Gelfand, Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets, *J. Amer. Statist. Assoc.* 111 (514) (2016a) 800–812.
- [10] A. Datta, S. Banerjee, A.O. Finley, A.E. Gelfand, On nearest-neighbor Gaussian process models for massive spatial data, *Wiley Interdiscip. Rev. Comput. Stat. N. Y.* 8 (5) (2016b) 162–171.
- [11] A.P. Dempster, Covariance selection, *Biometrics* 28 (1) (1972) 157–175.
- [12] D. Dey, A. Datta, S. Banerjee, Graphical Gaussian process models for highly multivariate spatial data, *Biometrika* 109 (4) (2022) 993–1014.
- [13] A. Dobra, A. Lenkoski, A. Rodriguez, Bayesian inference for general Gaussian graphical models with application to multivariate lattice data, *J. Amer. Statist. Assoc.* 106 (496) (2011) 1418–1433.
- [14] P. Drineas, M.W. Mahoney, N. Cristianini, On the Nyström method for approximating a gram matrix for improved Kernel-based learning, *J. Mach. Learn. Res.* 6 (12) (2005).
- [15] A.O. Finley, A. Datta, B.D. Cook, D.C. Morton, H.E. Andersen, S. Banerjee, Efficient algorithms for Bayesian nearest neighbor Gaussian processes, *J. Comput. Graph. Statist.* 28 (2) (2019) 401–414.
- [16] A.O. Finley, H. Sang, S. Banerjee, A.E. Gelfand, Improving the performance of predictive process modeling for large datasets, *Comput. Statist. Data Anal.* 53 (8) (2009) 2873–2884.
- [17] J. Friedman, T. Hastie, R. Tibshirani, Sparse inverse covariance estimation with the graphical lasso, *Biostatistics* 9 (3) (2008) 432–441.
- [18] T. Gneiting, W. Kleiber, M. Schlather, Matérn cross-covariance functions for multivariate random fields, *J. Amer. Statist. Assoc.* 105 (491) (2010) 1167–1177.
- [19] P.J. Green, A. Thomas, Sampling decomposable graphs using a Markov chain on junction trees, *Biometrika* 100 (1) (2013) 91–110.
- [20] X. Jin, S. Banerjee, B.P. Carlin, Order-free co-regionalized areal data models with application to multiple-disease mapping, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 69 (5) (2007) 817–838.
- [21] M.L. Krock, W. Kleiber, D. Hammerling, S. Becker, Modeling massive highly multivariate nonstationary spatial data with the basis graphical lasso, *J. Comput. Graph. Statist.* (2023) 1–16.
- [22] E. Kropf, S.K. Syan, L. Minuzzi, B.N. Frey, From anatomy to function: the role of the somatosensory cortex in emotional regulation, *Braz. J. Psychiatry* 41 (2018) 261–269.
- [23] K.-Y. Lee, L. Li, B. Li, H. Zhao, Nonparametric functional graphical modeling through functional additive regression operator, *J. Amer. Statist. Assoc.* (2022) 1–15.
- [24] B. Li, E. Solea, A nonparametric graphical model for functional data with application to brain networks based on fMRI, *J. Amer. Statist. Assoc.* 113 (524) (2018) 1637–1655.
- [25] M.A. Lindquist, A. Krishnan, M. López-Solà, M. Jepma, C.-W. Woo, L. Koban, M. Roy, L.Y. Atlas, L. Schmidt, L.J. Chang, et al., Group-regularized individual prediction: theory and application to pain, *NeuroImage* 145 (2017) 274–287.
- [26] N. Meinshausen, P. Bühlmann, High-dimensional graphs and variable selection with the lasso, *Ann. Statist.* 34 (3) (2006) 1436–1462.
- [27] T. Nath, B. Caffo, T. Wager, M.A. Lindquist, A machine learning based approach towards high-dimensional mediation analysis, *NeuroImage* 268 (2023) 119843.
- [28] S.-Y. Oh, A. Petersen, FGM: Partial separability and functional Gaussian graphical models, 2019, R package version 1.0.
- [29] W.-Y. Ong, C.S. Stohler, D.R. Herr, Role of the prefrontal cortex in pain processing, *Mol. Neurobiol.* 56 (2019) 1137–1166.
- [30] J. Peng, P. Wang, N. Zhou, J. Zhu, Partial correlation estimation by joint sparse regression models, *J. Amer. Statist. Assoc.* 104 (486) (2009) 735–746.
- [31] X. Qiao, S. Guo, G.M. James, Functional graphical models, *J. Amer. Statist. Assoc.* 114 (525) (2019) 211–222.
- [32] A. Roverato, Hyper inverse wishart distribution for non-decomposable graphs and its application to Bayesian inference for Gaussian graphical models, *Scand. J. Stat.* 29 (3) (2002) 391–411.
- [33] T.P. Speed, H.T. Kiiveri, et al., Gaussian Markov distributions over finite graphs, *Ann. Statist.* 14 (1) (1986) 138–150.
- [34] M.L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*, Springer Science & Business Media, New York, 2012.
- [35] M.L. Stein, Limitations on low rank approximations for covariance matrices of spatial data, *Spat. Stat.* 8 (2014) 1–19.
- [36] H. Wackernagel, *Multivariate Geostatistics: An Introduction with Applications*, Springer Science & Business Media, New York, 2013.
- [37] C.-W. Woo, M. Roy, J.T. Buhle, T.D. Wager, Distinct brain systems mediate the effects of nociceptive input and self-regulation on pain, *PLoS Biol.* 13 (1) (2015) e1002036.
- [38] W. Yang, H.-G. Müller, U. Stadtmüller, Functional singular component analysis, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 73 (3) (2011) 303–324.
- [39] F. Yao, H.-G. Müller, J.-L. Wang, Functional data analysis for sparse longitudinal data, *J. Amer. Statist. Assoc.* 100 (470) (2005) 577–590.
- [40] B.T. Yeo, F.M. Krienen, J. Sepulcre, M.R. Sabuncu, D. Lashkari, M. Hollinshead, J.L. Roffman, J.W. Smoller, L. Zöllei, J.R. Polimeni, et al., The organization of the human cerebral cortex estimated by intrinsic functional connectivity, *J. Neurophysiol.* (2011).
- [41] J. Zapata, S.-Y. Oh, A. Petersen, Partial separability and functional graphical models for multivariate Gaussian processes, *Biometrika* 109 (3) (2022) 665–681.
- [42] H. Zhu, N. Strawn, D.B. Dunson, Bayesian graphical models for multivariate functional data, *J. Mach. Learn. Res.* 17 (204) (2016) 1–27.