

Gaussian Graphical Models for Partially Observed Multivariate Functional Data

Marco Borriero

Department of Statistics, Computer Science and Applications,
University of Florence, Italy

Luigi Augugliaro

Department of Economics, Business and Statistics,
University of Palermo, Italy

Gianluca Sottile

Department of Economics, Business and Statistics,
University of Palermo, Italy

Veronica Vinciotti

Department of Mathematics, University of Trento, Italy

Abstract

In many applications, the variables that characterize a stochastic system are measured along a second dimension, such as time. This results in multivariate functional data and the interest is in describing the statistical dependences among these variables. It is often the case that the functional data are only partially observed. This creates additional challenges to statistical inference, since the functional principal component scores, which capture all the information from these data, cannot be computed. Under an assumption of Gaussianity and of partial separability of the covariance operator, we develop an EM-type algorithm for penalized inference of a functional graphical model from multivariate functional data which are only partially observed. A simulation study and an illustration on German electricity market data show the potential of the proposed method.

Keywords: functional Gaussian graphical models, multivariate functional data, missing data, partial separability, sparse inference

1 Introduction

In recent years, the study of data in the form of functions or curves, referred to as functional data analysis (Ramsay and Silverman, 2005; Frédéric and Vieu, 2006; Horváth and Kokoszka, 2012; Hsing and Eubank, 2015), has received much attention, with applications in a variety of fields, from medicine (Li and Solea, 2018; Zhao et al., 2024) to economics (Liebl, 2013; Goldberg et al., 2014). By their own nature, functional data are high dimensional. Therefore, dimensionality reduction techniques are often needed in order to make them more tractable. The most common is functional principal components analysis (FPCA), which uses the Karhunen-Loève theorem to represent functional data as an infinite linear combination of a given basis system, with coefficients called functional principal components scores, or simply scores. The advantage of this representation is that the scores are random variables which contain all the information about the data. Moreover, FPCA gives automatically an order of importance to the sequence of scores, with the first one containing the majority of the information and then proceeding in a decreasing fashion.

In this work, we focus on multivariate functional data in which every observation is represented by different curves. In particular, we consider the problem of modelling the dependences between the curves of the multivariate process through a functional Gaussian graphical model (Zhu et al., 2016; Lee et al., 2023; Qiao et al., 2019; Zapata et al., 2022). For this purpose, a crucial role is played by the covariance operator, which is an extension of the covariance matrix of a multivariate random variable to the functional case. However, the non invertibility of this operator makes the inference of a functional Gaussian graphical model more challenging than in the non-functional context.

Two fundamental contributions in this direction are provided by Qiao et al. (2019) and Zapata et al. (2022), who address the problem in two different ways. Qiao et al. (2019) approximate the functional data with a truncated Karhunen-Loève expansion, obtaining in this way a finite dimensional precision matrix of the scores that is used to recover

the edge set of the functional graphical model. Zapata et al. (2022) avoids making the strong assumption that the curves are exactly K -dimensional by assuming an additive hypothesis on the covariance operator, called partial separability. This gives rise to a novel version of the Karhunen-Loève expansion for multivariate functional data. Thanks to this hypothesis, the edge set of the functional Gaussian graphical model can be obtained from the precision matrices of the scores. More recently, Fici et al. (2025) have extended this second method to include a dependence of the functional curves on external covariates, while Zhao et al. (2024) have proposed a functional generalization of the neighborhood selection method (Meinshausen and Bühlmann, 2006) which avoids the definition of the precision operator and the partial separability assumption on the covariance operator but focuses primarily on graph selection.

To date, the literature on functional Gaussian graphical models has considered only the setting in which functional data are completely observed. However, functional data are often recorded with missing values. This is not the case of sparse functional data, where very few observations are available in the considered time domain (Yao et al., 2005; James et al., 2000). It is rather the case when, for instance, the technological device that is used for taking measurements stops working properly for a certain period of time or when the measures are recorded but considered as not plausible for a variety of reasons. Situations like this give rise to functional data that are partially observed. In these cases, one may be interested in reconstructing the trajectories of the curves in the missing domains and in studying the dependence structure among the curves.

There are approaches in the literature for partially observed functional data in the univariate case. In particular, Kraus (2015) propose a methodology for the functional completion of univariate functional data through a functional linear regression with the observed data as predictors, while Delaigle and Hall (2016) approximate fragmented functional data with a Markov chain model. Finally, Kneip and Liebl (2020) define a new

reconstruction operator, which belongs to a broader class than the class of linear operators and that is optimal in the sense that the reconstructed curve in the missing domain preserves the continuity at the extremes of the observed intervals.

In our work, we expand this methodology to the case of multivariate functional data. In particular, we provide a generalization of the score imputations of Kraus (2015) to the multivariate case, within the setting of a functional Gaussian graphical model and under an assumption of partial separability of the covariance operator (Zapata et al., 2022). In this way, the score imputation exploits the dependence between the functional curves. As both the scores and the precision matrices are unknown, we develop an EM-type algorithm for inference of the parameters and of the edge-set of a functional Gaussian graphical model in the presence of missing data.

The paper is organized as follows. In Section 2, we introduce multivariate functional data and functional Gaussian graphical models. In Section 3, we discuss the methodology for the inference of a functional Gaussian graphical model from partially observed multivariate data, and present the technical details of the multivariate score imputation and the proposed EM-type algorithm. Section 4 presents a simulation study, showing the performance of the method in terms of the percentage of missing data and compared to existing methods. Finally, Section 5 concludes with an illustration of the methodology on multivariate functional data describing the German electricity market (Liebl, 2019).

2 Functional Gaussian graphical models

In this section, we introduce the main notation used throughout the paper and briefly review the theory of functional Gaussian graphical models.

2.1 Notation

We denote by $L^2_{[0,1]}$ the separable Hilbert space of square-integrable real-functions defined on $[0, 1]$, which is equipped with the standard inner product $\langle f, g \rangle = \int_0^1 f(t)g(t)dt$, for any $f, g \in L^2_{[0,1]}$, and induced L_2 -norm $\|f\| = \langle f, f \rangle^{1/2}$. For any given subset of the closed interval $[0, 1]$, say S , the restrictions of the inner product and the norm are denoted by $\langle f^S, g^S \rangle = \int_S f(t)g(t)dt$ and $\|f^S\| = \langle f^S, f^S \rangle^{1/2}$, respectively. The p -fold Cartesian product of the spaces $L^2_{[0,1]}$ is denoted by $\left(L^2_{[0,1]}\right)^p$ and the inner product and norm are defined as $\langle \mathbf{f}, \mathbf{g} \rangle = \sum_{j=1}^p \langle f_j, g_j \rangle$ and $\|\mathbf{f}\| = \langle \mathbf{f}, \mathbf{f} \rangle^{1/2}$, respectively. For any given set of p subsets of $[0, 1]$, say $\mathbf{S} = \{S_j\}_{j=1}^p$, we denote by $\langle \mathbf{f}^{\mathbf{S}}, \mathbf{g}^{\mathbf{S}} \rangle = \sum_{j=1}^p \langle f_j^{S_j}, g_j^{S_j} \rangle$ and $\|\mathbf{f}^{\mathbf{S}}\|$ the restriction of the inner product and norm, respectively. Finally, given a $p \times p$ -dimensional matrix of functions, say $\mathbf{M} = \{M_{hk}\}$, with $M_{hk} \in L^2_{[0,1]}$, we denote by $\mathbf{M}_k = \{M_{hk}\} \in \left(L^2_{[0,1]}\right)^p$ the k th column of $\mathbf{M}$ and, by $\langle \mathbf{M}, \mathbf{f} \rangle \in \left(L^2_{[0,1]}\right)^p$ the p -dimensional vector whose k th function is equal to $\langle \mathbf{M}_k, \mathbf{f} \rangle$, for any given $\mathbf{f} \in \left(L^2_{[0,1]}\right)^p$. Given two matrices of functions, with dimension $p \times p$, say $\mathbf{M}$ and $\mathbf{N}$, we denote by $\langle \mathbf{M}, \mathbf{N} \rangle$ the $p \times p$ dimensional matrix with generic entry $\langle \mathbf{M}_h, \mathbf{N}_k \rangle$.

2.2 Review of the functional Gaussian graphical models

Multivariate functional data are a random sample from a multivariate process $\{\mathbf{X}(t) \in \mathbb{R}^p : t \in [0, 1]\}$, which is assumed to be zero-mean such that $\mathbf{X} \in \left(L^2_{[0,1]}\right)^p$ almost surely and $\mathbb{E}(\|\mathbf{X}\|_p) < \infty$. Functional Gaussian graphical models, introduced by Zhu et al. (2016), represent a new class of graphical models which extends the classical Gaussian graphical models (Lauritzen, 1996) to infinite dimensional spaces. In particular, $\mathbf{X}$ is assumed to follow a multivariate Gaussian process and the conditional independence structure among the p random functions is encoded by an undirected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, i.e., the pair (h, k) belongs to the edge-set $\mathcal{E}$ if and only if $X_h \perp\!\!\!\perp X_k \mid \mathbf{X}_{\mathcal{V} \setminus \{h, k\}}$.

As in the finite dimensional case, a multivariate Gaussian process is uniquely determined

by the covariance operator $\mathcal{C} : \left(L_{[0,1]}^2\right)^p \rightarrow \left(L_{[0,1]}^2\right)^p$, with $C_{hk}(s, t) = \mathbb{E}(X_h(s), X_k(s))$ the cross-covariance function between X_h and X_k (see Da Prato (2006) for more details). Following Chiou et al. (2014), the operator $\mathcal{C}$ is formally defined as the integral operator with covariance kernel $\mathbf{C}(s, t) = \{C_{hk}(s, t)\}$, i.e.,

$$\mathcal{C}\mathbf{f}(s) = \langle \mathbf{C}(s, \cdot), \mathbf{f} \rangle = \begin{pmatrix} \langle \mathbf{C}_1(s, \cdot), \mathbf{f} \rangle \\ \vdots \\ \langle \mathbf{C}_p(s, \cdot), \mathbf{f} \rangle \end{pmatrix}, \quad \forall \mathbf{f} \in \left(L_{[0,1]}^2\right)^p, \quad (1)$$

where $\langle \mathbf{C}_k(s, \cdot), \mathbf{f} \rangle = \sum_{h=1}^p \langle C_{hk}(s, \cdot), f_h \rangle$. Using the operator (1), we denote by $MGP(\mathcal{C})$ a multivariate Gaussian process and by $\{MGP(\mathcal{C}), \mathcal{G}\}$ the corresponding functional Gaussian graphical model.

As discussed in Qiao et al. (2019), although the functional setting is more complicated, some of the main properties of the standard Gaussian graphical model are inherited in an infinite dimensional setting. Specifically, letting $\mathbf{X}_{\setminus(h,k)} = \{X_m : m \in \mathcal{V} \setminus (h, k)\}$, it is possible to show that the edge-set can be defined in terms of the conditional cross-covariance function $K_{hk}(s, t) = Cov(X_h(s), X_k(t) \mid \mathbf{X}_{\setminus(h,k)})$, which represents the covariance between $X_h(s)$ and $X_k(t)$ conditional on the remaining $p - 2$ random functions. In particular, it follows that:

$$\mathcal{E} = \{(h, k) : K_{hk}(s, t) \neq 0 \text{ for some } s \text{ and } t \in [0, 1], (h, k) \in \mathcal{V}^2, h \neq k\}.$$

To estimate the edge-set from multivariate functional data, Qiao et al. (2019) propose a two-steps procedure. First, they use a finite approximation of each random function. This is formally obtained using the first L_h terms of the univariate Karhunen-Loève expansion $X_h(t) = \sum_{l=1}^{+\infty} \zeta_{hl} \varphi_{hl}(t)$, where $\{\varphi_{hl}\}_{l=1}^{+\infty}$ are the eigenfunctions and form a complete orthonormal basis system (CONS) for $L_{[0,1]}^2$. Under the Gaussianity assumption of the process, one can show that $\zeta_{hl} \sim N(0, \lambda_{hl})$, with $\lambda_{h1} \geq \lambda_{h2} \geq \dots \geq 0$. Moreover, ζ_{hl} is independent from $\zeta_{h'l'}$ for $h \neq h'$ or $l \neq l'$. Since the random vector $\boldsymbol{\zeta} = (\zeta_{11}, \dots, \zeta_{1L_1}, \dots, \zeta_{p1}, \dots, \zeta_{pL_p})^\top$ follows a multivariate Gaussian distribution whose

precision matrix has a specific block structure, the authors propose a generalization of the graphical lasso (glasso) estimator (Yuan and Lin, 2007), called functional graphical lasso, in order to encourage blockwise sparsity in the resulting precision matrix. As shown in Qiao et al. (2019), the approach can recover the edge-set $\mathcal{E}$ only under the restrictive assumption that each random function has a finite representation.

The reason for this theoretical restriction lies in the observation that the covariance operator is compact and therefore not invertible (Hsing and Eubank, 2015). Consequently, the connection between conditional independence and an inverse covariance operator is lost, as the latter does not exist. To overcome this methodological problem, Zapata et al. (2022) introduce a specific assumption on the structure of the covariance operator called partial separability. Formally, the covariance operator $\mathcal{C}$ is said to be partially separable if there exists a CONS of $L^2_{[0,1]}$, denoted by $\{\varphi_l\}_{l=1}^{+\infty}$, and if, for each $l \in \mathbb{N}$, there exists an orthonormal matrix $\mathbf{Q}_l = (\mathbf{q}_{1l} \mid \cdots \mid \mathbf{q}_{pl})$ of dimension $p \times p$ such that:

$$\mathcal{C}(\mathbf{q}_{hl}\varphi_l)(s) = \lambda_{hl}\mathbf{q}_{hl}\varphi_l(s).$$

The main advantage of this assumption is that it implies the following multivariate Karhunen-Loève expansion:

$$\mathbf{X}(t) = \sum_{l=1}^{\infty} \boldsymbol{\xi}_l \varphi_l(t), \quad (2)$$

where, under the assumption of a Gaussian process, $\boldsymbol{\xi}_l$ are mutually uncorrelated Gaussian random vectors with $\mathbb{E}(\boldsymbol{\xi}_l) = \mathbf{0}$, $\mathbb{E}(\boldsymbol{\xi}_l \boldsymbol{\xi}_l^\top) = \Sigma_l = \mathbf{Q}_l^\top \Lambda_l \mathbf{Q}_l$ and $\Lambda_l = \text{diag}(\lambda_{1l}, \dots, \lambda_{pl})$. The expansion in (2) is assumed to be ordered according to the decreasing value of $\text{tr}(\Sigma_l)$. Although the assumption of partial separability relates the basis $\{\varphi_l\}_{l=1}^{+\infty}$ to the covariance operator $\mathcal{C}$, it is possible to show that its theoretical properties, such as uniqueness and optimality, depend on the covariance operator $\mathcal{H} = \frac{1}{p} \sum_{h=1}^p \mathcal{C}_{hh}$, which admits, under the assumption of partial separability, the following spectral decomposition:

$$\mathcal{H} = \sum_{l=1}^{+\infty} \lambda_l \varphi_l \otimes \varphi_l,$$

where $\lambda_l = p^{-1}\text{tr}(\Sigma_l)$ and the tensor product is the operator $(\varphi_l \otimes \varphi_l)(\cdot) = \langle \varphi_l, \cdot \rangle \varphi_l$ on $L^2_{[0,1]}$.

The advantage of the representation (2) is related to the form of the edge set of the functional Gaussian graphical model. Indeed, letting $\Theta_l = \Sigma_l^{-1}$ and θ_{hkl} the (h, k) entry of Θ_l , if the covariance operator $\mathcal{C}$ is partially separable, then there exists a sequence of sets $\{\mathcal{E}_l\}_{l=1}^{+\infty}$, defined as $\mathcal{E}_l = \{(h, k) : \theta_{hkl} \neq 0, h \neq k\}$, such that the edge set of the functional Gaussian graphical model is given by $\mathcal{E} = \bigcup_{l=1}^{+\infty} \mathcal{E}_l$. Notice that $\mathcal{E}_l$ can be seen as the edge set of the Gaussian graphical model associated with $\boldsymbol{\xi}_l \sim N(\mathbf{0}, \Sigma_l)$. This shows that, under the hypothesis of partial separability, the edge set of the functional Gaussian graphical model can be recovered from the sequence of precision matrices $\{\Theta_l\}_{l=1}^{+\infty}$. Moreover, Zapata et al. (2022) prove that the edge selection procedure is consistent. That is, for a fixed n , there exists a $L = L_p \in \mathbb{N}$ such that $\mathcal{E} = \bigcup_{l=1}^{L_p} \mathcal{E}_l$ and L_p diverges with n only if p diverges. Since the precision matrices $\Theta_l = \Sigma_l^{-1}$ contain all the necessary information to estimate the relevant terms of a functional Gaussian graphical model, Zapata et al. (2022) propose to use an extension of the glasso estimator, called joint graphical lasso estimator (Danaher et al., 2014), which can recover, under specific assumptions, the true graph $\mathcal{G}$.

3 Inference from partially observed functional data

In this section, we first formalize the inferential problem under missing data, by defining a suitable objective function. We then embed this in an Expectation-Maximization (EM)-type algorithm. A crucial step of the algorithm is the imputation of the scores, which we discuss in the next subsection, before concluding with the computational aspects of the algorithm.

3.1 Expected log-likelihood function

Let $\mathbf{X} \sim MGP(\mathcal{C})$ be a p -dimensional vector of random functions following a multivariate Gaussian process with a partially separable covariance operator. For each function, we

consider a partition in the domain $[0, 1]$. As we consider the case of partially observed data, the domain of the j th random function X_j is split into O_j and M_j , representing the portions in which the function is observed and unobserved, respectively. In both cases, O_j and M_j might be given by the union of a finite number of sub-intervals. As in Kraus (2015), we further assume that the functional missing data is missing completely at random, which means that observed and missing domains can be thought as fixed subsets. In the remaining part of this paper, we denote by $X_j^{O_j}$ and $X_j^{M_j}$ the restriction of the random function X_j over the subdomains O_j and M_j , respectively.

When the functional data are not observed in their entire domain, the entries of the p dimensional random score $\boldsymbol{\xi}_l$ of the expansion (2) cannot be computed directly. Formally, each ξ_{jl} can be written as a sum of an observed and a missing term:

$$\xi_{jl} = \langle X_j, \varphi_l \rangle = \langle X_j^{O_j}, \varphi_l^{O_j} \rangle + \langle X_j^{M_j}, \varphi_l^{M_j} \rangle = \xi_{jl}^o + \xi_{jl}^m, \quad \forall j = 1, \dots, p, \forall l \in \mathbb{N}, \quad (3)$$

where $\varphi_l^{O_j}$ and $\varphi_l^{M_j}$ denote the restriction of φ_l on O_j and M_j , respectively. Since X_j is not observed in M_j , the term ξ_{jl}^m is missing and, consequently, some imputation is needed in order to complete the random function X_j .

To define a proper objective function by which to estimate the conditional independence structure among the p random functions, we propose to use the first L leading terms of the expansion (2) along with the identity (3), so that the log-likelihood function can be written as

$$\ell(\{\Theta\}) = \sum_{l=1}^L \left\{ \frac{1}{2} \log \det \Theta_l - \frac{1}{2} (\boldsymbol{\xi}_l^o + \boldsymbol{\xi}_l^m)^\top \Theta_l (\boldsymbol{\xi}_l^o + \boldsymbol{\xi}_l^m) \right\}, \quad (4)$$

where $\{\Theta\} = \{\Theta_1, \dots, \Theta_L\}$ denotes the set of L precision matrices.

We use function (4) to develop an algorithm that follows the rationale underlying the Expectation-Maximization (EM) algorithm (Dempster et al., 1977). In particular, instead of maximizing function (4), which cannot be computed as it depends on the missing scores, we maximize the conditional expected value of the log-likelihood function given the observed

functions $\mathbf{X}^O = \{X_j^{O_j}\}_{j=1}^p$:

$$\mathbb{E}[\ell(\{\Theta\}) \mid \mathbf{X}^O] = \sum_{l=1}^L \left\{ \frac{1}{2} \log \det \Theta_l - \frac{1}{2} \mathbb{E}[(\boldsymbol{\xi}_l^o + \boldsymbol{\xi}_l^m)^\top \Theta_l (\boldsymbol{\xi}_l^o + \boldsymbol{\xi}_l^m) \mid \mathbf{X}^O] \right\}.$$

which can be simplified by noting that:

$$\begin{aligned} \mathbb{E}[(\boldsymbol{\xi}_l^o + \boldsymbol{\xi}_l^m)^\top \Theta_l (\boldsymbol{\xi}_l^o + \boldsymbol{\xi}_l^m) \mid \mathbf{X}^O] = \\ \text{tr}\{(\boldsymbol{\xi}_l^o)(\boldsymbol{\xi}_l^o)^\top \Theta_l\} + \text{tr}\{\mathbb{V}[\boldsymbol{\xi}_l^m \mid \mathbf{X}^O] \Theta_l\} + 2\text{tr}\{\mathbb{E}[\boldsymbol{\xi}_l^m \mid \mathbf{X}^O](\boldsymbol{\xi}_l^o)^\top \Theta_l\}. \end{aligned}$$

The term $\mathbb{V}[\boldsymbol{\xi}_l^m \mid \mathbf{X}^O] = \mathbb{E}[(\boldsymbol{\xi}_l^m)(\boldsymbol{\xi}_l^m)^\top \mid \mathbf{X}^O]$, which represents the covariance matrix of $\boldsymbol{\xi}_l^m$ conditional on $\mathbf{X}^O$, can be approximated using the theoretical results given in Guo et al. (2015). In particular, under an assumption that Θ_l is sparse, it is possible to show that:

$$\mathbb{V}[\boldsymbol{\xi}_l^m \mid \mathbf{X}^O] \approx (\boldsymbol{\mu}_l^{m|o})(\boldsymbol{\mu}_l^{m|o})^\top, \quad (5)$$

where $\boldsymbol{\mu}_l^{m|o} = \mathbb{E}[\boldsymbol{\xi}_l^m \mid \mathbf{X}^O]$. This approximation was successfully used in a number of methodological papers on the inference of Gaussian graphical models from partially observed data (see for example Augugliaro et al. (2018, 2020, 2023); Sottile et al. (2024) and references therein). Using the approximation (5), the expected likelihood function can be written as follows

$$\sum_{l=1}^L \left\{ \frac{1}{2} \log \det \Theta_l - \frac{1}{2} (\boldsymbol{\xi}_l^o + \boldsymbol{\mu}_l^{m|o})^\top \Theta_l (\boldsymbol{\xi}_l^o + \boldsymbol{\mu}_l^{m|o}) \right\}.$$

Since a maximization of this function will lead to estimated precision matrices with all non-zero entries, we explore functional Gaussian graphical models with different degrees of complexity, i.e., edge-sets with varying sparsity levels, by maximizing, similarly to Zapata et al. (2022), the following penalized expected likelihood

$$\sum_{l=1}^L \left\{ \frac{1}{2} \log \det \Theta_l - \frac{1}{2} (\boldsymbol{\xi}_l^o + \boldsymbol{\mu}_l^{m|o})^\top \Theta_l (\boldsymbol{\xi}_l^o + \boldsymbol{\mu}_l^{m|o}) \right\} - P_\gamma(\{\Theta\}), \quad (6)$$

where $P_\gamma(\{\Theta\})$ is the group lasso penalty function proposed in Danaher et al. (2014)

$$P_\gamma(\{\Theta\}) = \gamma_1 \left\{ \gamma_2 \sum_{l=1}^L \sum_{h \neq k}^p |\theta_{hkl}| + (1 - \gamma_2) \sum_{h \neq k}^p \left(\sum_{l=1}^L \theta_{hkl}^2 \right)^{1/2} \right\},$$

with γ_1 a non-negative tuning parameter used to control the overall penalty level, while $\gamma_2 \in [0, 1]$ distributes the penalty between the two penalty functions.

Table 1 summarizes the main steps of the proposed EM-type algorithm for the optimization of (6). The pseudo-code reveals that a crucial step is played by the imputation of the missing scores. This aspect will be discussed in the next section.

Table 1: Pseudo-code of the proposed EM-type algorithm

Step	Description
1	compute φ_l as solution of the following identity $\langle \mathcal{H}(s, \cdot), \varphi_l \rangle = \lambda_l \varphi_l(s)$
2	choose the number of L terms to approximate expansion (2)
3	repeat
4	for $l = 1$ to L
5	let $\hat{\boldsymbol{\mu}}_l^{m o}$ be an estimate of $\boldsymbol{\mu}_l^{m o}$
6	let $\hat{\boldsymbol{\xi}}_l = \boldsymbol{\xi}_l^o + \hat{\boldsymbol{\mu}}_l^{m o}$
7	and for
8	estimate the precision matrices $\{\hat{\Theta}_\gamma\} = \arg \max \sum_{l=1}^L \left\{ \frac{1}{2} \log \det \Theta_l - \frac{1}{2} \hat{\boldsymbol{\xi}}_l^\top \Theta_l \hat{\boldsymbol{\xi}}_l \right\} - P_\gamma(\{\Theta\})$
9	until a convergence criterion is met

3.2 Multivariate imputation of the missing scores

For the imputation of the scores, we generalize the approach of Kraus (2015) to the multivariate case. In this way, the imputation of missing data exploits the dependence structure among the p random functions. In terms of the mean-squared prediction error, the j th entries of the missing score $\boldsymbol{\xi}_l^m$, denoted by ξ_{jl}^m , can be imputed using $\mathbb{E}[\xi_{jl}^m \mid \mathbf{X}^o] = \mu_{jl}^{m|o}$.

Since $\mu_{jl}^{m|o}$ could be a non-linear functional of $\mathbf{X}^O$, we propose to model $\mu_{jl}^{m|o}$ using a continuous linear functional of the observed functions. Then, by the Riesz representation theorem, we have $\mu_{jl}^{m|o} = \sum_{i=1}^p \langle \beta_{ijl}, X_i^{O_i} \rangle = \langle \beta_{jl}, \mathbf{X}^O \rangle$, where $\beta_{jl} = \{\beta_{ijl}\}_{i=1}^p$ is a p -dimensional vector of functions with $\beta_{ijl} \in L_{O_i}^2$. In matrix form, denoting by $\mathbf{B}_l = (\beta_{1l} \mid \cdots \mid \beta_{pl})$ the matrix of regression coefficient functions, the conditional expected value of ξ_l^m is modelled by

$$\mu_l^{m|o} = \langle \mathbf{B}_l, \mathbf{X}^O \rangle = \begin{pmatrix} \langle \beta_{1l}, \mathbf{X}^O \rangle \\ \vdots \\ \langle \beta_{pl}, \mathbf{X}^O \rangle \end{pmatrix}.$$

Next, we discuss how to define the infinite-dimensional optimization problem by which to estimate $\mathbf{B}_l$. Under the assumption that $\xi_l^m \sim N_p(\langle \mathbf{B}_l, \mathbf{X}^O \rangle, \Sigma_l^m)$, where Σ_l^m denotes the covariance matrix of the vector ξ_l^m , the optimal linear functional can be defined as the solution of the following minimization problem

$$\min_{\mathbf{B}_l} \frac{1}{2} \mathbb{E} \{ (\xi_l^m - \langle \mathbf{B}_l, \mathbf{X}^O \rangle)^\top \Theta_l^m (\xi_l^m - \langle \mathbf{B}_l, \mathbf{X}^O \rangle) \}, \quad (7)$$

where Θ_l^m is the inverse of Σ_l^m . The following theorem discusses the solution to this optimization problem.

Theorem 1. *Under the assumption that $\xi_l^m \sim N_p(\langle \mathbf{B}_l, \mathbf{X}^O \rangle, \Sigma_l^m)$, the optimal linear functional $\hat{\mathbf{B}}_l$, defined as the solution of the problem (7), satisfies the following matrix equation:*

$$\mathcal{C}^{OO} \hat{\mathbf{B}}_l = \mathbf{R}_l, \quad (8)$$

where $\mathcal{C}^{OO}$ is the covariance operator of the process $\mathbf{X}^O$ with kernel $\mathcal{C}^{OO} = \{C_{hk}^{O_h O_k}\}$, and

$$\mathcal{C}^{OO} \hat{\mathbf{B}}_l = (\mathcal{C}^{OO} \hat{\beta}_{1l} \mid \cdots \mid \mathcal{C}^{OO} \hat{\beta}_{pl}), \quad \mathbf{R}_l = \begin{pmatrix} \mathcal{C}_{11}^{O_1 M_1} \varphi_l^{M_1} & \cdots & \mathcal{C}_{1p}^{O_1 M_p} \varphi_l^{M_p} \\ \vdots & & \vdots \\ \mathcal{C}_{p1}^{O_p M_1} \varphi_l^{M_1} & \cdots & \mathcal{C}_{pp}^{O_p M_p} \varphi_l^{M_p} \end{pmatrix}.$$

The proof is reported in Appendix A. We notice that the system (8) is a multivariate generalization of Kraus (2015) and reveals that $\hat{\mathbf{B}}_l$ does not depend on the precision matrix

Θ_l^m . Consequently the problems of estimating $\mathbf{B}_l$ and Θ_l can be solved separately. However, the covariance operator $\mathcal{C}^{OO}$ is compact with an infinite dimensional range, which implies that the inverse operator $(\mathcal{C}^{OO})^{-1}$ is unbounded and therefore a small perturbation in $\mathbf{R}_l$ may lead to a perturbation of $\hat{\mathbf{B}}_l$. In other words, as in the univariate setting, the inverse problem (8) is ill-posed.

To solve this problem, we modify the original ill-posed inverse problem in such a way that it becomes a well-posed problem with a stable solution. Similarly to Kraus (2015), we propose to use a ridge regularization so that, instead of problem (8) we solve the problem:

$$\mathcal{C}_\alpha^{OO} \hat{\mathbf{B}}_l^\alpha = \mathbf{R}_l, \quad (9)$$

where $\mathcal{C}_\alpha^{OO} = \mathcal{C}^{OO} + \alpha \mathcal{I}^O$, α is a positive tuning parameter and $\mathcal{I}^O$ is the identity operator. Since the inverse of the operator $\mathcal{C}_\alpha^{OO}$ is bounded, the solution $\hat{\mathbf{B}}_l^\alpha = (\mathcal{C}^{OO} + \alpha \mathcal{I}^O)^{-1} \mathbf{R}_l$ is stable. In particular, the stability of the solution increases with α , though at the expense of an increased bias as the system (9) deviates more and more from the original problem. As too small values of α are not sufficient to regularize the solution, a trade-off between bias and variance is needed. This aspect will be discussed in detail in the next section.

3.3 Computational aspects

Suppose that we have a set of n independent and identically distributed realizations from the functional Gaussian graphical model $\{MGP(\mathcal{C}), \mathcal{G}\}$, where the covariance operator $\mathcal{C}$ is partially separable with basis $\{\varphi_l\}_{l=1}^p$. We also assume that a subset of the n realizations, denoted by $\mathbb{O} \subset \{1, \dots, n\}$, is completely observed. Therefore, for each $i \in \overline{\mathbb{O}}$, the j th element of the random vector $\mathbf{X}_i$ is observed in the subset O_{ij} , while it is missing in M_{ij} . To emphasize the dependence of $\mathbf{X}_i$ from these subsets, the observed and missing fragments of X_{ij} are denoted by $X_{ij}^{O_{ij}}$ and $X_{ij}^{M_{ij}}$, respectively.

As described in steps 1 and 2 of the pseudo-code in Table 1, the preliminary step of the proposed procedure is to estimate the eigenfunctions φ_l and the number of terms used

to approximate the multivariate Karhunen-Loève expansion (2). To solve these problems, we follow Kraus (2015) and estimate the covariance operator by estimating the covariance functions as follows

$$\widehat{C}_{hk}(s, t) = \frac{I_{hk}(s, t)}{\sum_{i=1}^n U_i^{hk}(s, t)} \sum_{i=1}^n U_i^{hk}(s, t) \{X_{ih}(s) - \hat{\mu}_h(s, t)\} \{X_{il}(t) - \hat{\mu}_l(s, t)\},$$

where $U_i^{hk}(s, t)$ is equal to 1 if $s \in O_{ih}$ and $t \in O_{ik}$ while it is 0 otherwise, and $I_{hk}(s, t)$ is the indicator $\mathbb{1}_{\{\sum_{i=1}^n U_i^{hk}(s, t) > 0\}}$. In the previous estimator, the mean functions are estimated by

$$\hat{\mu}_h(s, t) = \frac{I_{hk}(s, t)}{\sum_{i=1}^n U_i^{hk}(s, t)} \sum_{i=1}^n U_i^{hk}(s, t) X_{ih}(s).$$

Given the estimated covariance functions $\widehat{C}_{hh}$, one obtains the estimated covariance operators $\widehat{\mathcal{C}}_{hh}$ and, consequently, the operator $\widehat{\mathcal{H}} = p^{-1} \sum_{h=1}^p \widehat{\mathcal{C}}_{hh}$. From this the eigenfunctions are estimated by solving the empirical counterpart of the identity in Step 1 of the proposed pseudo-code, i.e.,

$$\langle \widehat{\mathcal{H}}(s, \cdot), \hat{\varphi}_l \rangle = \hat{\lambda}_l \hat{\varphi}_l(s).$$

Then, the number of L terms used to approximate expansion (2) is chosen as the minimum number of components that explain a fixed percentage of the total variance.

Given the estimated eigenfunctions $\{\hat{\varphi}_l\}_{l=1}^L$, for each $i \in \overline{\mathcal{O}}$, we impute the missing score $\boldsymbol{\xi}_{il}^m$ using $\hat{\boldsymbol{\mu}}_{il}^{m|o} = \langle \widehat{\mathbf{B}}_{il}^{\hat{\alpha}}, \mathbf{X}_i^{O_i} \rangle$, where $\widehat{\mathbf{B}}_{il}^{\hat{\alpha}} = (\mathcal{C}^{O_i O_i} + \hat{\alpha} \mathcal{I}^{O_i})^{-1} \mathbf{R}_{il}$ and $\hat{\alpha}$ is the value of the regularization parameter α that allows us to get the optimal reconstruction of the whole partially observed functions. As in Kraus (2015), we consider the minimization problem:

$$\min_{\mathcal{A}: \|\mathcal{A}\|_{\infty} < \infty} \mathbb{E} \|\mathbf{X}^M - \mathcal{A} \mathbf{X}^O\|^2,$$

where the solution is in the class of continuous (bounded) linear operators from $L_{\mathcal{O}}^2$ to L_M^2 . By Fréchet differentiation we obtain the system of normal equation $\mathcal{A} \mathcal{C}^{OO} = \mathcal{C}^{MO}$ then, analogous to what we have done for the multivariate principal scores, the solution is $\hat{\mathcal{A}}^{\alpha} = \mathcal{C}^{MO} (\mathcal{C}^{OO} + \alpha \mathcal{I}^{OO})^{-1}$ and the optimal reconstruction of $\mathbf{X}^M$ is $\hat{\mathcal{A}}^{\alpha} \mathbf{X}^O$. From

a practical point of view, the optimal α -value is selected as minimizer of the following generalized cross-validation criterion:

$$\text{gcv}_i(\alpha) = \frac{\sum_{i' \in \mathbb{O}} \|\mathbf{X}_i^{M_i} - \hat{\mathcal{A}}^\alpha \mathbf{X}_i^{O_i}\|^2}{\{1 - \text{df}_i(\alpha)/|\mathbb{O}|\}^2}, \quad (10)$$

where $\text{df}_i(\alpha)$ is the number of effective degrees of freedom, formally defined as:

$$\text{df}_i(\alpha) = \text{tr}\{(\hat{\mathcal{C}}_\alpha^{O_i O_i})^{-1} \hat{\mathcal{C}}^{O_i O_i}\} = \sum_{k=1}^{+\infty} \frac{\hat{\lambda}_k^{O_i O_i}}{\hat{\lambda}_k^{O_i O_i} + \alpha}.$$

Then, the missing scores are estimated using $\hat{\boldsymbol{\mu}}_{il}^{m|o} = \langle \hat{\mathbf{B}}_{il}^{\hat{\alpha}}, \mathbf{X}_i^{O_i} \rangle$, and $\hat{\boldsymbol{\xi}}_l$ are computed as described in Step 6 of the pseudo-code, i.e., by letting $\hat{\boldsymbol{\xi}}_{il} = \hat{\boldsymbol{\xi}}_{il}^o + \hat{\boldsymbol{\mu}}_{il}^{m|o}$, where

$$\hat{\xi}_{ijl}^o = \langle X_{ij}^{O_{ij}} - \hat{\mu}_j^{O_{ij}}, \hat{\varphi}_l^{O_{ij}} \rangle, \quad \text{and} \quad \hat{\mu}_{ij}(t) = \frac{J_j(t)}{\sum_{i=1}^n O_{ij}(t)} \sum_{i=1}^n O_{ij}(t) X_{ij}(t),$$

and $O_{ij}(t)$ is used for the indicator $\mathbb{1}_{\{t \in O_{ij}\}}$, while $J_j(t) = \mathbb{1}_{\{\sum_{i=1}^n O_{ij}(t) > 0\}}$.

Once the imputation of missing scores is complete, we estimate the precision matrices as described in Step 8 of the proposed algorithm, that is, by maximizing the following function:

$$\sum_{l=1}^L \{\log \det \Theta_l - \text{tr}(S_l \Theta_l)\} - P_\gamma(\{\Theta\}),$$

where $S_l = \sum_{i=1}^n \hat{\boldsymbol{\xi}}_{il} \hat{\boldsymbol{\xi}}_{il}^\top / n$ denotes the sample covariance matrix. The previous steps are repeated until a suitable convergence criterion is met and the edge-set $\mathcal{E}$ is estimated by $\hat{\mathcal{E}} = \bigcup_{l=1}^L \hat{\mathcal{E}}_l$, where $\hat{\mathcal{E}}_l = \{(h, k) : \hat{\theta}_{hkl} \neq 0\}$.

3.4 Selection of tuning parameters

Different values of the tuning parameters γ_1 and γ_2 correspond to different levels of sparsity in the associated graph. We select the optimal values by minimizing the extended Bayesian Information Criterion (eBIC) (Foygel and Drton, 2010). As the method we propose follows an EM-type algorithm, we consider the approach proposed by Ibrahim et al. (2008) and replace the marginal likelihood by the Q-function evaluated at the optimum

$$Q(\{\hat{\Theta}^{\gamma_1, \gamma_2}\}) = \sum_{l=1}^L \left\{ \log \det \hat{\Theta}_l^{\gamma_1, \gamma_2} - \text{tr}(S_l^{\gamma_1, \gamma_2} \hat{\Theta}_l^{\gamma_1, \gamma_2}) \right\},$$

which is readily available at convergence. Here $\{\hat{\Theta}^{\gamma_1, \gamma_2}\}$ are the precision matrices that maximize the penalized loglikelihood function, where we now explicitly emphasize the dependence on the tuning parameters γ_1 and γ_2 . Therefore, similarly to Sottile et al. (2024), the eBIC criterion that we use is given by

$$eBIC(\gamma_1, \gamma_2) = -2Q(\{\hat{\Theta}^{\gamma_1, \gamma_2}\}) + \sum_{l=1}^L \left\{ |\hat{\mathcal{E}}_l^{\gamma_1, \gamma_2}| \log n + 4\gamma_1 |\hat{\mathcal{E}}_l^{\gamma_1, \gamma_2}| \log p \right\}$$

where $\hat{\mathcal{E}}_l^{\gamma_1, \gamma_2}$ denotes the number of edges in the graph associated to $\hat{\Theta}_l^{\gamma_1, \gamma_2}$ and γ_1 is set to 0.5.

4 Simulation study

This section investigates the finite-sample performance of the proposed methodology. The data-generating mechanism is described in Section 4.1, together with the evaluation criteria. Section 4.2 presents the results of the simulations under different scenarios, including a comparison with the univariate imputation procedure of Kraus (2015).

4.1 Simulation setting and evaluation measures

We recall that an undirected graph is denoted by $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, with vertex set $\mathcal{V} = \{1, \dots, p\}$ and edge set $\mathcal{E} = \bigcup_{l=1}^L \mathcal{E}_l$, where L denotes the number of bases used to approximate the expansion (2). In this simulation study, we set $L = 3$ and $p = 15$.

Precision matrices are simulated using the following procedure. First, the precision matrices $\{\tilde{\Theta}_l\}_{l=1}^L$ are generated according to three distinct structures, displayed in Figure 1, namely a star, banded, and small-world structure, respectively. In particular, the star structure creates groups of five variables connected in a hub-and-spoke fashion, the banded structure links variables sequentially, and the small-world structure induces sparse random connections. Then, following Zapata et al. (2022), we ensure that the traces of the covariance matrices $\{\Sigma_l\}_{l=1}^L$ satisfy the decreasing ordering entailed by the expansion (2), by




Figure 1: Adjacency structures for the precision matrices Θ_l , $l = 1, 2, 3$ used in the simulation study. The same structure is imposed across all L layers.

scaling each precision matrix via a decaying factor, i.e., we let $\Theta_l = a_l^{-1} \tilde{\Theta}_l$, with $a_l = 3 \cdot l^{-1.8}$.

We set the sample size n to 100. Then, for each simulation, the p -dimensional scores $\xi_{1l}, \dots, \xi_{nl}$ are drawn from a multivariate Gaussian distribution with zero mean and covariance Σ_l . A system of Fourier basis, denoted by $\{\varphi_l\}_{l=1}^L$, is evaluated on a grid of $d = 50$ evenly spaced points, and functional data are obtained through the linear combination $X_i(t) = \sum_{l=1}^L \xi_{il} \varphi_l(t)$, for each $i = 1, \dots, n$. The functional representation is refined by performing a functional principal component analysis (FPCA) on the generated data. The number of retained components is selected so that the cumulative proportion of explained variance reaches 99.99%. This adaptive basis selection ensures that the functional representation captures nearly all relevant variability while avoiding possible redundant components.

From the complete data generated as above, we construct partially observed trajectories by randomly selecting a subset of functional data and removing a sequence of w consecutive points. We consider different settings by varying the proportion of partially observed functions (π_{po}) and the relative length of the window where the functions are missing (π_w).

In particular, we consider the various combinations resulting from $\pi_{po} \in \{0.25, 0.50, 0.75\}$ and $\pi_w \in \{0.25, 0.50, 0.75\}$.

As described in Section 3.1, the proposed estimation procedure relies, at the M-step, on the joint graphical lasso of Danaher et al. (2014). For this, we use the efficient algorithm developed in Sottile et al. (2024). Since the aim is to impose a common sparsity structure across layers, we fix $\gamma_2 = 0$ throughout. The algorithm is applied over a grid of 21 equally spaced γ_1 values between 0 and $\gamma_1^{\max}$, which is the smallest γ_1 -value such that the off-diagonal elements of all estimated precision matrices are equal to zero.

The performance of the method is assessed by three criteria. The accuracy of parameter estimation is measured through the mean squared error

$$\text{MSE}_{\Theta} = \frac{1}{L} \sum_{l=1}^L \mathbb{E} \|\Theta_l - \hat{\Theta}_l\|_F,$$

while the recovery of the sparsity pattern is quantified by the area under the ROC curve (AUC), where each point of the curve corresponds to the sensitivity and specificity of edge selection at a specific γ_1 -value. Finally, the performance in terms of functional imputation is evaluated through

$$\text{MSE}_X = \frac{1}{p} \sum_{j=1}^p \mathbb{E} \|X_j - \hat{X}_j\|_2^2,$$

which compares observed and reconstructed curves over the domain. Each configuration is replicated over 250 Monte Carlo samples.

We compare the proposed approach with the univariate imputation method of Kraus (2015). In this competing procedure, each component of the multivariate Gaussian process is imputed independently, that is disregarding cross-sectional dependences. For the method, we use the same FPCA-derived basis of our proposed estimation procedure, so that differences in performance can be attributed exclusively to the estimation step rather than to discrepancies in basis expansions. After imputation, the joint graphical lasso estimation procedure is applied on the imputed data to estimate the precision matrices, as in the M-step of our proposed method.

4.2 Evaluation and comparison of the proposed methodology

For the sake of brevity, this section focuses on the most challenging scenario, with precision matrices generated from the small-world structure. The results for the star and banded structures are deferred to Appendix B. The impact of π_{po} and π_w on the performance of the methods is examined.

Figure 2 plots the functional reconstruction error for the different simulation settings and replications, where in each case we consider its lowest value across the γ_1 sequence. The results show how the proposed method is robust to the different levels of missingness, with MSE_X consistently close to zero across all settings. On the contrary, the univariate procedure of Kraus (2015) does not exploit the correlations between the curves, leading to systematically larger errors as the proportion of missingness increases, both in terms of π_w and π_{po} .

Figure 3 focuses on the estimation of the precision matrices and recovery of the conditional independence structure. In particular, Figure 3a reports the distribution of the minimum MSE_Θ values along the γ_1 -sequence, while Figure 3b shows the distributions of the AUC values, across the 250 replications. The results show how the recovery of the precision matrices and graph structure is comparable between the two methods for low levels of missingness. As the levels of π_w and π_{po} increase, the univariate approach exhibits a marked deterioration of the performance compared to the proposed method, particularly in terms of estimation of the precision matrices. Network recovery is instead only marginally affected by the proportion of missing data.

5 An illustration on the German electricity market

This section illustrates the proposed multivariate method through an application to the German electricity market. The dataset, originally studied in Liebl (2019), contains hourly observations of four variables: electricity price, electricity demand, wind power generation,




Figure 2: Simulating 250 replications of incomplete data from a functional Gaussian graphical model, with $n = 100$, $p = 15$, a small-world graph structure and varying levels of missingness (π_w and π_{po}). Curve reconstruction error, measured by the minimum value along the γ_1 -sequence of the mean squared error, is lower for our proposed method (light grey) than the univariate method of Kraus (2015) (dark grey), particularly under high levels of missingness.

and net imports, the latter defined as the difference between import and export flows. The available data span the period from January 1, 2010, to March 27, 2012. To account for the strong weekly seasonality typical of energy markets, each week is treated as a single statistical unit. As suggested by Liebl (2019), Sundays are removed from the analysis since the dependence structure among the variables is likely to differ on non-working days. For the remaining non-working days, which do not systematically fall on the same weekday, the corresponding observations are instead treated as missing values. After preprocessing, each functional data is evaluated over $d = 144$ equally spaced time points, corresponding to 24 hourly measurements per day over six days (Monday through Saturday). The first and last incomplete weeks, as well as those consisting entirely of holidays, are excluded, resulting in $n = 114$ weekly curves. The proportion of partially observed curves ranges




Figure 3: Simulating 250 replications of incomplete data from a functional Gaussian graphical model, with $n = 100$, $p = 15$, a small-world graph structure and varying levels of missingness (π_w and π_{po}). (a) Estimation of precision matrices, measured by the minimum value of the mean squared error along the γ_1 -sequence, and (b) recovery of the graph structure, measured by the AUC, are comparable for our proposed method (light grey) and the univariate method of Kraus (2015) (dark grey) across most settings, but the proposed method has better performance under high levels of missingness.




Figure 4: Examples of observed functional curves for the four variables in the German electricity market dataset. Dark grey lines represent partially observed curves. The similar qualitative behavior across variables reflects strong interdependence.

between 70.2% and 72.8%, while the proportion of missing time points per curve is between 30.3% and 32.5%. The original time series are first smoothed using penalized cubic spline functions. The resulting functional data are then centered and scaled as suggested in Chiou et al. (2014), where mean and variance functions are estimated as proposed in Kraus (2015). This enables a coherent comparison among functions expressed in different units and characterized by heterogeneous variability.

Figure 4 shows an example of multivariate partially observed functional data. The plots highlight typical daily and weekly cycles and reveal strong correlations among variables, motivating the use of a multivariate modeling framework. The application of the functional principal component analysis (FPCA), specified using the operator $\mathcal{H}$, shows




Figure 5: Analysis of the German electricity market data. (a) eBIC values across the γ_1 sequence of values, with the vertical line indicating the selected γ_1 . (b) Estimated conditional independence graph corresponding to the optimal model.

that 29 eigenfunctions account for at least 99.99% of the total variability. Accordingly, in this study we set $L = 29$.

As in our simulation study, we use the proposed method by setting the mixing parameter γ_2 equal to 0 whereas the tuning parameter γ_1 is selected by minimizing the extended Bayesian information criterion (eBIC), as discussed in Section 3.4. Empirically, a grid of 51 equally spaced γ_1 values is considered, ranging from $\gamma_1^{\max}$, for which all estimated precision matrices are diagonal, to $\gamma_1 = 0$, corresponding to fully dense matrices. After refitting the models at each γ_1 -value through maximum likelihood estimation (Augugliaro et al., 2020, 2023), the optimal γ_1 -value is chosen as the minimizer of the eBIC. Figure 5a reports the eBIC path and the optimal value.

Considering the optimal model, Figure 6 shows four examples of imputed curves for the four variables. Figure 5b instead plots the conditional independence graph corresponding to the selected model. The graph reveals a sparse and interpretable dependence struc-




Figure 6: Example of imputed functional curves for the four variables in the German electricity market dataset. Dark grey lines denote partially observed curves; black segments indicate the imputed windows.

ture among the four variables. To gain further insight, Figure 7 displays the estimated cross-covariance functions between selected pairs of variables. The left panel highlights a clear negative association between electricity price and wind power, particularly within the same day and, to a lesser extent, across adjacent days. This relationship reflects the expected market mechanism: higher wind generation increases renewable supply, which leads to lower prices. The right panel shows a positive association between wind power and electricity imports, with an average cross-covariance of approximately 0.185, suggesting that periods of intense wind production coincide with greater import activity, possibly due to grid balancing and export–import adjustments. Finally, the pair composed of electricity demand and imports displays consistently positive cross-covariances, averaging around 0.603, indicating that higher demand levels are typically accompanied by increased reliance on external electricity sources.

Overall, the results confirm that the proposed approach captures effectively the main dependences among functional variables, even under substantial levels of missingness. The identified conditional independence structure is consistent with established economic mechanisms, such as the inverse relationship between renewable production and prices, and the joint dynamics between imports and demand. These findings indicate that the proposed methodology provides a coherent representation of the functional dependence network, offering a flexible and statistically robust framework for the analysis of high-dimensional functional data in energy systems and beyond.

6 Conclusion

We have presented a methodology for the reconstruction of partially observed multivariate curves and for the inference of the underlying functional Gaussian graphical model. The functional completion is done by the imputations of the missing part of the scores, which in turn is obtained through a functional linear regression where the predictors are represented




Figure 7: Estimated cross-covariance functions between selected pairs of variables: electricity price and wind power (left), and electricity imports and wind power (right). The color scale represents the sign and intensity of dependence across pairs of time points.

by the observed part of the curves. This constitutes a generalization of the procedure of Kraus (2015) to the multivariate case. Under the assumption of partial separability introduced in Zapata et al. (2022), and the Karhunen-Loève expansion of the multivariate functional data that results from that, the functional graphical model can be inferred using penalized likelihood approaches. Score imputation and graphical model fitting are integrated into an Expectation-Maximization algorithm. The simulation study evaluates the performance of the proposed method and shows how it is advantageous compared to existing approaches where score imputation is conducted for each variable separately, without accounting for the multivariate nature of the functional data.

Code and data availability

The R script to replicate the simulation study and the data that support the findings of this study are available from <https://github.com/gianluca-sottile/Gaussian-Graphical-Models-for-Partially-Observed-Multivariate-Functional-Data>.

Acknowledgment

Luigi Augugliaro and Gianluca Sottile gratefully acknowledge financial support from the University of Palermo (FFR2024 and FFR2025). The authors were also financially supported by the European Union - Next Generation EU - Mission 4 Component 2 - CUP: B53D23009480006.

References

L. Augugliaro, A. Abbruzzo, and V. Vinciotti. ℓ_1 -penalized censored Gaussian graphical model. *Biostatistics*, 21(2):e1–e16, 2018.

- L. Augugliaro, G. Sottile, and V. Vinciotti. The conditional censored graphical lasso estimator. *Statistics and Computing*, 30:1273–1289, 2020.
- L. Augugliaro, G. Sottile, E. C. Wit, and V. Vinciotti. cglasso: An R package for conditional graphical lasso inference with censored and missing values. *Journal of Statistical Software*, 105(1):1–58, 2023.
- J.-M. Chiou, Y.-T. Chen, and Y.-F. Yang. Multivariate functional principal component analysis: A normalization approach. *Statistica Sinica*, 24:1571–1596, 2014.
- G. Da Prato. *An Introduction to Infinite-Dimensional Analysis*. Springer Berlin, Heidelberg, 2006.
- P. Danaher, P. Wang, and D. M. Witten. The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society Series B*, 76(2):373–397, 2014.
- A. Delaigle and P. Hall. Approximating fragmented functional data by segments of Markov chains. *Biometrika*, 103(4):779–799, 2016.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B*, 39:1–38, 1977.
- R. Fici, G. Sottile, L. Augugliaro, and E. C. Wit. Functional Gaussian graphical regression models for air quality data. *Journal of the Royal Statistical Society Series C: Applied Statistics*, page qlaf042, 2025.
- R. Foygel and M. Drton. Extended Bayesian information criteria for Gaussian graphical models. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 23, Red Hook, NY, USA, 2010. Curran Associates, Inc.

- F. Frédéric and P. Vieu. *Nonparametric functional data analysis. Theory and practice.* Springer Series in Statistics. Springer, New York, 2006.
- Y. Goldberg, Y. Ritov, and A. Mandelbaum. Predicting the continuation of a function with applications to call center data. *Journal of Statistical Planning and Inference*, 147: 53–65, 2014.
- J. Guo, E. Levina, G. Michailidis, and J. Zhu. Graphical models for ordinal data. *Journal of Computational and Graphical Statistics*, 24(1):183–204, 2015.
- L. Horváth and P. Kokoszka. *Inference for functional data with applications.* Springer Series in Statistics. Springer, New York, 2012.
- T. Hsing and R. Eubank. *Theoretical foundations of functional data analysis, with an introduction to linear operators.* Wiley Series in Probability and Statistics. John Wiley & Sons, Chichester, England, 2015.
- J. Ibrahim, H. Zhu, and N. Tang. Model selection criteria for missing-data problems using the em algorithm. *Journal of the American Statistical Association*, 103(484):1648–1658, 2008. doi: 10.1198/016214508000001057.
- G. M. James, T. J. Hastie, and C. A. Sugar. Principal component models for sparse functional data. *Biometrika*, 87(3):587–602, 2000.
- A. Kneip and D. Liebl. On the optimal reconstruction of partially observed functional data. *The Annals of Statistics*, 48(3):1692 – 1717, 2020.
- D. Kraus. Components and completion of partially observed functional data. *Journal of the Royal Statistical Society Series B*, 77(4):777–801, 2015.
- S. L. Lauritzen. *Graphical Models.* Oxford University Press, Oxford, 1996.

- K.-Y. Lee, D. Ji, L. Li, T. Constable, and H. Z. and. Conditional functional graphical models. *Journal of the American Statistical Association*, 118(541):257–271, 2023.
- B. Li and E. Solea. A nonparametric graphical model for functional data with application to brain networks based on fMRI. *Journal of the American Statistical Association*, 113(524):1637–1655, 2018.
- D. Liebl. Modeling and forecasting electricity spot prices: A functional data perspective. *The Annals of Applied Statistics*, pages 1562–1592, 2013.
- D. Liebl. Nonparametric testing for differences in electricity prices: The case of the Fukushima nuclear accident. *The Annals of Applied Statistics*, 13(2):1128 – 1146, 2019.
- N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *The Annals of Statistics*, 34(3):1436 – 1462, 2006.
- X. Qiao, S. Guo, and G. M. James. Functional graphical models. *Journal of the American Statistical Association*, 114(525):211–222, 2019.
- J. O. Ramsay and B. W. Silverman. *Functional data analysis*. Springer Series in Statistics. Springer, New York, 2nd edition, 2005.
- G. Sottile, L. Augugliaro, V. Vinciotti, W. Arancio, and C. Coronello. Sparse inference of the human haematopoietic system from heterogeneous and partially observed genomic data. *Journal of the Royal Statistical Society Series C*, 74(1):204–228, 2024.
- F. Yao, H.-G. Müller, and J.-L. Wang. Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, 100(470):577–590, 2005.
- M. Yuan and Y. Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.

J. Zapata, S. Y. Oh, and A. Petersen. Partial separability and functional graphical models for multivariate Gaussian processes. *Biometrika*, 109(3):665–681, 2022.

B. Zhao, P. S. Zhai, Y. S. Wang, and M. Kolar. High-dimensional functional graphical model structure learning via neighborhood selection approach. *Electronic Journal of Statistics*, 18(1):1042 – 1129, 2024.

H. Zhu, N. Strawn, and D. B. Dunson. Bayesian Graphical Models for Multivariate Functional Data. *Journal of Machine Learning Research*, 17(204):1–27, 2016.

A Proof of Theorem 1

To obtain the system (8), we start rewriting the objective functional, denoted by $f(\mathbf{B}_k)$, in a more convenient way. First, we notice that it is equal to the sum of three specific terms:

$$\begin{aligned}
f(\mathbf{B}_l) &= \frac{1}{2} \mathbb{E} \{ (\boldsymbol{\xi}_l^m - \langle \mathbf{B}_l, \mathbf{X}^O \rangle)^\top \Theta_l^m (\boldsymbol{\xi}_l^m - \langle \mathbf{B}_l, \mathbf{X}^O \rangle) \} \\
&= \frac{1}{2} \mathbb{E} [\text{tr} \{ \Theta_l^m (\boldsymbol{\xi}_l^m - \langle \mathbf{B}_l, \mathbf{X}^O \rangle) (\boldsymbol{\xi}_l^m - \langle \mathbf{B}_l, \mathbf{X}^O \rangle)^\top \}] \\
&= \frac{1}{2} \text{tr} [\Theta_l^m \mathbb{E} \{ (\boldsymbol{\xi}_l^m) (\boldsymbol{\xi}_l^m)^\top \}] + \frac{1}{2} \text{tr} [\Theta_l^m \mathbb{E} \{ \langle \mathbf{B}_l, \mathbf{X}^O \rangle \langle \mathbf{B}_l, \mathbf{X}^O \rangle^\top \}] \\
&\quad - \text{tr} [\Theta_l^m \mathbb{E} \{ \boldsymbol{\xi}_l^m \langle \mathbf{B}_l, \mathbf{X}^O \rangle^\top \}]
\end{aligned} \tag{11}$$

The first term in (11) can be further simplified by recalling that $\xi_{hl}^m = \langle X_h^{M_h}, \varphi_l^{M_h} \rangle$, which implies identity:

$$\mathbb{E}(\xi_{hl}^m \xi_{kl}^m) = \mathbb{E}(\langle X_h^{M_h}, \varphi_l^{M_h} \rangle \langle X_k^{M_k}, \varphi_l^{M_k} \rangle) = \left\langle \left(\mathbb{E}(X_h^{M_h} X_k^{M_k}), \varphi_l^{M_k} \right), \varphi_l^{M_h} \right\rangle = \langle \mathcal{C}_{hk}^{M_h M_k} \varphi_l^{M_k}, \varphi_l^{M_h} \rangle,$$

where $\mathcal{C}_{hk}^{M_h M_k}$ denotes the cross-covariance operator with kernel $C_{hk}^{M_h M_k}(s, t) = \mathbb{E}\{X_h^M(s) X_k^M(t)\}$.

Using the previous identity we have:

$$\frac{1}{2} \text{tr} [\Theta_l^m \mathbb{E} \{ (\boldsymbol{\xi}_l^m) (\boldsymbol{\xi}_l^m)^\top \}] = \frac{1}{2} \sum_{h=1}^p \sum_{k=1}^p \theta_{hkl}^m \mathbb{E}(\xi_{hl}^m \xi_{kl}^m) = \frac{1}{2} \sum_{h=1}^p \sum_{k=1}^p \theta_{hkl}^m \langle \mathcal{C}_{hk}^{M_h M_k} \varphi_l^{M_k}, \varphi_l^{M_h} \rangle, \tag{12}$$

which shows that the first term in f is a constant with respect to $\mathbf{B}_l$, therefore it can be remove our initial problem.

To rewrite the second term in (11) in a more straightforward way, we notice that the generic term of the matrix $\mathbb{E}\{\langle \mathbf{B}_l, \mathbf{X}^O \rangle \langle \mathbf{B}_l, \mathbf{X}^O \rangle^\top\}$ is equal to:

$$\begin{aligned}
\mathbb{E}\{\langle \beta_{hl}, \mathbf{X}^O \rangle \langle \beta_{kl}, \mathbf{X}^O \rangle\} &= \sum_{i=1}^p \sum_{j=1}^p \mathbb{E}[\langle \beta_{ihl}, X_i^{O_i} \rangle \langle \beta_{jkl}, X_j^{O_j} \rangle] \\
&= \sum_{i=1}^p \langle \beta_{ihl}, \langle \mathbf{C}_i^{OO}, \beta_{kl} \rangle \rangle \\
&= \langle \beta_{hl}, \langle \mathbf{C}^{OO}, \beta_{kl} \rangle \rangle \\
&= \langle \beta_{hl}, \mathbf{C}^{OO} \beta_{kl} \rangle,
\end{aligned}$$

therefore we have:

$$\begin{aligned}
\frac{1}{2} \text{tr}[\Theta_l^m \mathbb{E}\{\langle \mathbf{B}_l, \mathbf{X}^O \rangle \langle \mathbf{B}_l, \mathbf{X}^O \rangle^\top\}] &= \frac{1}{2} \sum_{h=1}^p \sum_{k=1}^p \theta_{hkl}^m \langle \beta_{hl}, \mathbf{C}^{OO} \beta_{kl} \rangle \\
&= \frac{1}{2} \text{tr}[\Theta_l^m \langle \mathbf{B}_l, \mathbf{C}^{OO} \mathbf{B}_l \rangle], \tag{13}
\end{aligned}$$

where, with slight abuse of the notation, by $\mathbf{C} \mathbf{B}_l$ and $\langle \mathbf{B}_l, \mathbf{C}^{OO} \mathbf{B}_l \rangle$ we denote the following matrices:

$$\mathbf{C}^{OO} \mathbf{B}_l = (\mathbf{C}^{OO} \beta_{1l} \mid \dots \mid \mathbf{C}^{OO} \beta_{pl}) \quad \text{and} \quad \langle \mathbf{B}_l, \mathbf{C}^{OO} \mathbf{B}_l \rangle = \begin{pmatrix} \langle \beta_{1l}, \mathbf{C}^{OO} \beta_{1l} \rangle & \dots & \langle \beta_{1l}, \mathbf{C}^{OO} \beta_{pl} \rangle \\ \vdots & \ddots & \vdots \\ \langle \beta_{pl}, \mathbf{C}^{OO} \beta_{1l} \rangle & \dots & \langle \beta_{pl}, \mathbf{C}^{OO} \beta_{pl} \rangle \end{pmatrix}.$$

Finally, we focus on the last term in (11). Using identity $\xi_{hl}^m = \langle X_h^{M_h}, \varphi_l^{M_h} \rangle$, the generic term in $\mathbb{E}\{\xi_l^m \langle \mathbf{B}_l, \mathbf{X}^O \rangle^\top\}$ can be written as follows:

$$\begin{aligned}
\mathbb{E}\{\xi_{hl}^m \langle \beta_{kl}, \mathbf{X}^O \rangle\} &= \mathbb{E}\{\langle X_h^{M_h}, \varphi_l^{M_h} \rangle \langle \beta_{kl}, \mathbf{X}^O \rangle\} \\
&= \sum_{i=1}^p \mathbb{E}\{\langle X_h^{M_h}, \varphi_l^{M_h} \rangle \langle \beta_{ikl}, X_i^{O_i} \rangle\} \\
&= \sum_{i=1}^p \langle \beta_{ikl}, \langle C_{ih}^{O_i M_h}, \varphi_l^{M_h} \rangle \rangle \\
&= \langle \beta_{kl}, \mathbf{R}_{hl} \rangle,
\end{aligned}$$

where $\mathbf{R}_{hl}$ is the h th column of the matrix

$$\mathbf{R}_l = (\mathbf{R}_{1l} \mid \cdots \mathbf{R}_{pl}) = \begin{pmatrix} \mathcal{C}_{11}^{O_1 M_1} \varphi_l^{M_1} & \cdots & \mathcal{C}_{1p}^{O_1 M_p} \varphi_l^{M_p} \\ \vdots & & \vdots \\ \mathcal{C}_{p1}^{O_p M_1} \varphi_l^{M_1} & \cdots & \mathcal{C}_{pp}^{O_p M_p} \varphi_l^{M_p} \end{pmatrix}.$$

Using the previous identities we have that the third term in (11) can be written as follows:

$$\text{tr}[\Theta_l^m \mathbb{E}\{\boldsymbol{\xi}_k^m \langle \mathbf{B}_l, \mathbf{X}^O \rangle^\top\}] = \sum_{h=1}^p \sum_{k=1}^p \theta_{hkl}^m \langle \boldsymbol{\beta}_{kl}, \mathbf{R}_{hl} \rangle = \text{tr}[\Theta_l^m \langle \mathbf{B}_l, \mathbf{R}_l \rangle]. \quad (14)$$

Given identities (12), (13) and (14) we can write the objective functional as follows:

$$f(\mathbf{B}_l) = \frac{1}{2} \text{tr}[\Theta_l^M \langle \mathbf{B}_l, \mathcal{C}^{OO} \mathbf{B}_l \rangle] - \text{tr}[\Theta_l^m \langle \mathbf{B}_l, \mathbf{R}_l \rangle] + \text{Cost},$$

then, taking the Fréchet derivative of $f(\mathbf{B}_l)$ with respect to $\mathbf{B}_l$ and imposing it equal to zero we have:

$$\frac{\partial f(\mathbf{B}_l)}{\partial \mathbf{B}_l} = \Theta_l^m \mathcal{C}^{OO} \mathbf{B}_l - \Theta_l^m \mathbf{R}_l = \mathbf{0}.$$

By multiplying both sides on the left by $(\Theta_l^m)^{-1}$ we obtain the following system

$$\mathcal{C}^{OO} \mathbf{B}_l - \mathbf{R}_l = \mathbf{0},$$

in the unknown $\mathbf{B}_l$.

B Additional simulation results

This appendix reports additional results from the simulation study described in Section 4. Specifically, we present the same set of analyses and performance measures discussed in Section 4.2, but considering two alternative adjacency structures to generate the precision matrices, namely the star and banded configurations displayed in Figure 1.

For each structure, we show the functional reconstruction error (MSE_X), the estimation error of the precision matrices (MSE_Θ), and the recovery accuracy of the graph structure




Figure 8: Simulating 250 replications of incomplete data from a functional Gaussian graphical model, with $n = 100$, $p = 15$, a star graph structure and varying levels of missingness (π_w and π_{po}). Curve reconstruction error, measured by the minimum value along the γ_1 -sequence of the mean squared error, is lower for our proposed method (light grey) than the univariate method of Kraus (2015) (dark grey), particularly under high levels of missingness.

(AUC), across all combinations of (π_{po}) and (π_w). The same simulation design, parameter settings, and estimation procedures described in Section 4.1 are employed also here. Figures 8 and 9 report the results on curve reconstruction and dependence structure estimation, respectively, for the star structure, while Figures 10 and 11 refer to the banded graph case.

The results are consistent with those discussed in the main text for the small-world configuration. In particular, the proposed method shows a superior performance in terms of reconstruction of the curves, estimation of the precision matrices and recovery of the conditional independence structure, compared with the univariate imputation method of Kraus (2015), with differences in performance becoming more pronounced as the proportion of missing data increases.




Figure 9: Simulating 250 replications of incomplete data from a functional Gaussian graphical model, with $n = 100$, $p = 15$, a star graph structure and varying levels of missingness (π_w and π_{po}). (a) Estimation of precision matrices, measured by the minimum value of the mean squared error along the γ_1 -sequence, and (b) recovery of the graph structure, measured by the AUC, are comparable for our proposed method (light grey) and the univariate method of Kraus (2015) (dark grey) across most settings, but the proposed method has better performance under high levels of missingness.




Figure 10: Simulating 250 replications of incomplete data from a functional Gaussian graphical model, with $n = 100$, $p = 15$, a banded graph structure and varying levels of missingness (π_w and π_{po}). Curve reconstruction error, measured by the minimum value along the γ_1 -sequence of the mean squared error, is lower for our proposed method (light grey) than the univariate method of Kraus (2015) (dark grey), particularly under high levels of missingness.




Figure 11: Simulating 250 replications of incomplete data from a functional Gaussian graphical model, with $n = 100$, $p = 15$, a band graph structure and varying levels of missingness (π_w and π_{po}). (a) Estimation of precision matrices, measured by the minimum value of the mean squared error along the γ_1 -sequence, and (b) recovery of the graph structure, measured by the AUC, are comparable for our proposed method (light grey) and the univariate method of Kraus (2015) (dark grey) across most settings, but the proposed method has better performance under high levels of missingness.