

# Multi-group Gaussian Processes

Didong Li<sup>1,2</sup>, Andrew Jones<sup>1</sup>, Sudipto Banerjee<sup>2</sup> and Barbara E. Engelhardt<sup>1,3</sup>

*Department of Computer Science, Princeton University<sup>1</sup>*

*Department of Biostatistics, University of California, Los Angeles<sup>2</sup>*

*Gladstone Institutes<sup>3</sup>*

## Abstract

Gaussian processes (GPs) are pervasive in functional data analysis, machine learning, and spatial statistics for modeling complex dependencies. Modern scientific data sets are typically heterogeneous and often contain multiple known discrete subgroups of samples. For example, in genomics applications samples may be grouped according to tissue type or drug exposure. In the modeling process it is desirable to leverage the similarity among groups while accounting for differences between them. While a substantial literature exists for GPs over Euclidean domains  $\mathbb{R}^p$ , GPs on domains suitable for multi-group data remain less explored. Here, we develop a multi-group Gaussian process (MGGP), which we define on  $\mathbb{R}^p \times \mathcal{C}$ , where  $\mathcal{C}$  is a finite set representing the group label. We provide general methods to construct valid (positive definite) covariance functions on this domain, and we describe algorithms for inference, estimation, and prediction. We perform simulation experiments and apply MGGP to gene expression data to illustrate the behavior and advantages of the MGGP in the joint modeling of continuous and categorical variables.

## 1 Introduction

Gaussian processes (GPs) are widely used in functional data analysis, machine learning, and spatial statistics due to their flexibility and expressiveness in modeling complex dependent data (Rasmussen and Williams, 2005; Stein, 1999; Gelfand et al., 2010; Cressie and Wikle, 2011; Banerjee et al., 2014). For example, in nonparametric regression models, GPs are commonly used to model unknown arbitrary functions; in Bayesian contexts, they act as priors over functions (Ghosal and Van der Vaart, 2017). In spatial statistics, GPs are widely used to model spatial dependencies in geostatistical models and to perform spatial prediction or interpolation (“kriging”) (Matheron, 1963). The GP framework may also be adapted to deal with discrete outcomes and applied to classification problems (Bernardo et al., 1998). GPs are also being increasingly used in deep learning and reinforcement learning applications (Damianou and Lawrence, 2013; Deisenroth et al., 2013).

---

Code for the model and experiments is available at <https://github.com/andrewcharlesjones/multi-group-GP>.

A GP is determined by its covariance function, also known as the covariogram or Mercer kernel, which for statistical modeling purposes is specified to be a real-valued positive definite (PD) function  $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ . Most commonly, GPs are defined over the Euclidean domain  $\mathcal{X} = \mathbb{R}^p$ , where  $p$  is the dimension of the input samples. However, in many cases, data sets admit a “multi-group” structure, where each sample belongs to one of  $k$  known groups. For example, biological measurements might come from multiple tissue types or cell types (Consortium et al., 2020; Regev et al., 2017); geospatial data might be collected from multiple locations defined by discrete demarcations, such as state or country borders (Pan et al., 2020); and census data comes from people of different races, ethnicities, and genders (U.S. Census Bureau, 2020). Models built on Euclidean domains may be unsuitable for such data because they are not designed to capture the discrete nature of the groups. Although inference for GPs on non-Euclidean manifolds and graphs has recently attracted attention in statistics (Niu et al., 2019; Dunson et al., 2020; Li et al., 2021) and machine learning (Borovitskiy et al., 2020, 2021), the multi-group setting remains largely unaddressed.

As a motivating example (which we will explore in depth in this paper), consider the data collected by the Genotype-Tissue Expression (GTEx) consortium (Consortium et al., 2020). The GTEx data contain measurements from thousands of biological samples across hundreds of human donors, and these samples are collected from up to 52 tissue types for each donor. The measurements include gene expression and genotype profiles, along with a variety of additional metadata, including demographic variables and tissue health measurements. In this dataset, each tissue type constitutes a “group.” While each tissue type has a unique genomic environment, some tissue types are expected to be similar to one another. For example, a subset of the tissue types correspond to subregions of the brain, and another subset of tissue types correspond to subregions of the heart. In these cases, it may be beneficial to share information across similar tissue types when analyzing dependencies in the data.

For flexibly modeling  $k$ -group data, we seek a stochastic process over  $\mathbb{R}^p \times \mathcal{C}$  to drive the inference, where  $\mathcal{C} = \{c_1, \dots, c_k\}$  is a finite set representing the group labels. To begin, we identify three existing approaches for applying GPs to multi-group data: the Separated GP (SGP), the Union GP (UGP), and the Hierarchical GP (HGP)—all of which are defined over Euclidean domains. The SGPs and UGPs are straightforward applications of GPs that effectively ignore the group structure and either model the groups separately (SGP) or combine samples across groups and ignore the known structure (UGP). These “separated” and “union” modeling approaches are often used in practice for their simplicity and extend beyond the context of GPs (Consortium et al., 2020; Tsherniak et al., 2017). We briefly describe these three models below.

1. The Separated GP (SGP) assumes independence across groups. Thus, the across-group correlation is set to zero:  $K((x, c_i), (x', c_j)) = 0$  if  $i \neq j$ . The SGP is equivalent to modeling each group with a separate, independent GP.
2. The Union GP (UGP) assumes the same dependencies within and across groups so the covariance function does not depend on the members of  $\mathcal{C}$ , i.e.,  $K((x, c_i), (x', c_j)) = K_0(x, x')$ . It is equivalent to modeling all groups jointly with a single GP.
3. The Hierarchical GP (HGP) accommodates both across- and within-group dependencies, where all within-group and across-group dependencies are assumed to be identical. Here,  $K((x, c_i), (x', c_j)) = K_0(x, x') + 1_{\{c_i=c_j\}}K_1(x, x')$ , where  $K_0$  and  $K_1$  are real-valued positive definite functions (Park and Choi, 2010; Hensman et al., 2013).

Each of the above models build covariance kernels based upon a standard GP over a Euclidean domain. While each of the above models offers a reasonable modeling solution for multi-group data, they offer limited flexibility. Specifically, the conditions implicit in the above models are restrictive and fail to model heterogeneity in the between-group dependencies. As a consequence, when the model is misspecified, the performance of these GPs will be unsatisfactory, especially when the overall sample size is small or the groups are unbalanced in terms of sample size.

In this manuscript we introduce a class of Multi-Group Gaussian Process (MGGP) models as a flexible tool to model complex dependent grouped data. A specific contribution here is the construction of positive definite covariance functions defined over  $\mathcal{X} := \mathbb{R}^p \times \mathcal{C}$  that will provide a valid specification for the MGGP over  $\mathbb{R}^p \times \mathcal{C}$ . Specifically, our proposed MGGP framework will offer the following benefits compared to existing models:

1. The MGGP flexibly models multi-group heterogeneity, leveraging varying levels of similarity between groups;
2. The MGGP is robust to model misspecification;
3. The MGGP can make use of domain/prior/expert knowledge about which groups are expected to be similar to one another;
4. The MGGP outperforms existing GPs, especially when the overall sample size is limited or when group-specific sample sizes are unbalanced;
5. The MGGP does not depend on a sophisticated model; instead, the multi-group structure is encoded in the covariance function, which enables the seamless use of existing methods for

estimation and inference in GPs;

6. Given that the MGGP is a valid Gaussian process, it is compatible with the growing set of existing scalable Bayesian inferential methods for GPs (see, e.g., Banerjee, 2017, and references therein for a review).

Much work in Bayesian statistics has focused on the challenge of joint modeling of continuous and categorical variables. One approach builds hierarchical models to partition the data based on each category and share strength across observations through the hierarchical structure (Dunson et al., 2003; Dunson, 2000; Teimourian et al., 2015; Ru et al., 2020). Other approaches use additive effects for each category (Schulam et al., 2015). More recent work uses Dirichlet process mixture models with both multinomial and Gaussian data likelihoods, which includes both additive mean effects and mixture components to effectively combine the above two strategies to model dependencies between continuous and categorical features (Murray and Reiter, 2016). Our MGGP model adds to this literature by avoiding additive, hierarchical, and mixture models entirely, and instead explicitly modeling the dependencies within and among the two types of variables through a non-Euclidean covariance function with a flexible Gaussian process prior. The flexibility of the Gaussian process allows straightforward conditioning on multiple categorical partitions, offering a closed form conditional posterior. This is especially useful when some groups have extremely small sample sizes, and the posterior distribution can exploit dependencies between groups.

The paper is organized as follows. Section 2 presents a model-based inferential framework for MGGP regression models. We turn to the construction of the MGGP using valid covariance functions on  $\mathbb{R}^p \times \mathcal{C}$  in Section 3 considering separable and inseparable functions in Sections 3.1 and 3.2, respectively, and show that the three aforementioned GPs for multi-group data are special cases in Section 3.3. In Section 4 we demonstrate our proposed method through simulation experiments and an application to a large gene expression data set. Section 5 concludes the article with a brief discussion. Mathematical details, including all proofs, are presented in the Appendix.

## 2 Multi-Group Gaussian Process Regression Models

We envision a Multi-Group Gaussian Process (MGGP) regression model, where our dependent variable  $y(x; c_j)$  is generated from a latent stochastic process over  $\mathbb{R}^p \times \mathcal{C}$  for inputs  $x \in \mathbb{R}^p$  and group  $j$  as

$$y(x; c_j) = \mu(x; c_j) + Z(x; c_j) + \epsilon(x; c_j), \quad \epsilon(x; c_j) \stackrel{ind}{\sim} N(0, \tau_j^2), \quad (1)$$

where  $\mu(x; c_j)$  is a mean function,  $Z(x; c_j)$  is a zero-centered latent process and  $\epsilon(x; c_j)$  is a zero-centered white-noise process capturing measurement error or fine-scale variation with group-specific variances. The mean function can be further modeled, if deemed appropriate, as  $\mu(x; c_j) = f_j(x)^\top \beta_j$ , where  $f_j(x)$  is a  $q_j \times 1$  vector of design variables possibly, but not necessarily, depending on  $x$ , and the  $\beta_j$  variables are  $q_j \times 1$  vectors of group-specific regression coefficients. This specification accommodates predictors or other explanatory variables that need neither be continuous nor reside within  $\mathbb{R}^p$ .

The model in Equation (1) can be regarded as a semiparametric regression model with a parametric specification offered by the mean function and a nonparametric specification offered by the latent process. Our focus in this paper, however, is not so much on modeling  $\mu(x; c_j)$ , which can be built from customary parametric linear model specifications, as it is on the latent stochastic process  $Z(x; c_j) : \chi \rightarrow \mathbb{R}$ , where  $\chi = \mathbb{R}^p \times \mathcal{C}$ . We will specify  $Z(\cdot; \cdot)$  to be a GP with zero mean and a covariance kernel  $K((x; c_j), (x'; c_{j'})) : \chi \times \chi \rightarrow \mathbb{R}$  so that  $K((x; c_j), (x'; c_{j'})) = \text{cov}(Z(x; c_j), Z(x'; c_{j'}))$  is a positive-definite covariance function. In general, we consider settings where data arise over a finite, possibly unbalanced, set of points  $\{(x_i; c_j)\}$  for  $i = 1, 2, \dots, n_j$  and  $j = 1, 2, \dots, k$ . Thus, each group can have a different number,  $n_j$ , of inputs. Given the covariance kernel, the realizations of the process over the finite set of points is the  $n \times 1$  vector  $Z = (Z_1^\top, \dots, Z_k^\top)^\top$ , where  $n = \sum_{j=1}^k n_j$  and  $Z_j = (Z(x_1; c_j), \dots, Z(x_{n_j}; c_j))^\top$  follows a multivariate Gaussian distribution with an  $n \times 1$  zero vector as mean and an  $n \times n$  covariance matrix  $K$ , whose  $(j, j')$ th block is given by the  $n_j \times n_{j'}$  matrix  $K_{jj'}$  whose  $(i, i')$  element is given by  $K((x_i; c_j), (x_{i'}; c_{j'}))$  for  $i = 1, 2, \dots, n_j$  and  $i' = 1, 2, \dots, n_{j'}$ .

Equation (1) leads to likelihood-based inference and can be extended to a Bayesian framework. Assuming, for elucidation purposes only, that  $\mu(x; c_j) = f_j(x)^\top \beta_j$ , our Bayesian MGGP model specifies the joint distribution of the parameters and the data as

$$p(\{\tau_j^2\}, \theta, \{\beta_j\}) \times N(Z | 0, K_\theta) \times \prod_{j=1}^k \prod_{i=1}^{n_j} N(y(x_i; c_j) | f_j(x_i)^\top \beta_j + Z(x_i; c_j), \tau_j^2), \quad (2)$$

where we now index the process covariance matrix by  $\theta$  to denote parametric specifications for the covariance kernel, and  $p(\{\tau_j^2\}, \theta, \{\beta_j\})$  is the prior on the model parameters. Inference on these parameters and the latent process proceeds by drawing samples from the posterior distribution  $p(\{\tau_j^2\}, \theta, \{\beta_j\}, Z | \{y(x_i; c_j)\}, \{f_j(x_i)\})$ , which is proportional to Equation (2).

Sampling from the joint posterior distribution including the process realizations  $Z$  will be challenging due to the dimension of  $Z$ . Exploiting the Gaussian likelihood, we work with the collapsed

likelihood after integrating out  $Z$  from Equation (2), which yields the posterior distribution

$$p(\tau, \theta, \beta | y, F) \propto p(\tau, \theta, \beta) \times N(y | F\beta, K_\theta + D_\tau), \quad (3)$$

where  $y$  is the  $n \times 1$  vector of observations,  $y(x_i; c_j)$ , constructed analogous to  $Z$ ,  $F$  is an  $n \times q$  block-diagonal matrix,  $q = \sum_{j=1}^k q_j$ , with  $n_j \times q_j$  blocks  $F_j = (f_j(x_1), \dots, f_j(x_{n_j}))^\top$ ,  $\beta = (\beta_1^\top, \dots, \beta_k^\top)^\top$  is the  $q \times 1$  vector of stacked regression coefficients,  $\tau = \{\tau_j^2\}$  is the collection of measurement error variances and  $D_\tau$  is the  $n \times n$  diagonal matrix with  $\tau_j^2 I_{n_j}$  as  $n_j \times n_j$  diagonal blocks. Markov chain Monte Carlo (MCMC) algorithms operate much more efficiently in sampling from Equation (3) because of the reduced parameter space than from Equation (2). Once posterior samples of  $\{\tau, \theta, \beta\}$  are collected, we recover the posterior samples for the latent process from  $p(Z | y, F) = \mathbb{E}[p(Z | \{\tau, \theta, \beta\}, y, F)]$ , where the expectation  $\mathbb{E}[\cdot]$  is taken with respect to the posterior distribution in Equation (3); Monte Carlo sampling will draw one  $Z \sim p(Z | \{\tau, \theta, \beta\}, y, F)$  for each posterior draw of  $\{\tau, \theta, \beta\}$ . This is straightforward because  $p(Z | \{\tau, \theta, \beta\}, y, F)$  is of the form  $N(Mm, M)$ , where  $M^{-1} = K_\theta^{-1} + D_\tau^{-1}$  and  $m = y - F\beta$ , and the draws need to be made using only the post-convergence samples of  $\{\tau, \theta, \beta\}$ .

To estimate the latent process at an unobserved input  $x_0 \in \mathbb{R}$  for a given group  $c_j \in \mathcal{C}$ , we evaluate the Bayesian posterior predictive distribution

$$p(Z(x_0; c_j) | \{y(x_i; c_j)\}, \{f_j(x_i)\}) \propto \int p(Z(x_0; c_j) | Z, \theta) \times p(Z, \{\tau, \theta, \beta\} | y, F) dZ d\{\tau, \theta, \beta\}, \quad (4)$$

where we have used the conditional independence  $p(Z(x_0; c_j) | Z, \{\tau, \theta, \beta\}, y, F) = p(Z(x_0; c_j) | Z, \theta)$  derived from the hierarchical model in Equation (2). We can sample from Equation (4) by drawing one  $Z(x_0; c_j) \sim p(Z(x_0; c_j) | Z, \theta)$  for each drawn posterior sample of  $Z$  and  $\theta$ , where  $p(Z(x_0; c_j) | Z, \theta)$  is Gaussian with mean  $K_\theta((x_0; c_j); \cdot)^\top K_\theta^{-1} Z$ , where  $K_\theta((x_0; c_j); \cdot)$  is the  $n \times 1$  vector with elements  $K_\theta((x_0; c_j), (x_i, c_{j'}))$  for  $j' = 1, \dots, k$  and  $i = 1, 2, \dots, n_{j'}$ , and variance  $K_\theta((x_0; c_j), (x_0; c_j)) - K_\theta((x_0; c_j); \cdot)^\top K_\theta^{-1} K_\theta((x_0; c_j); \cdot)$ . Finally, to predict the response at  $(x_0; c_j)$ , we can sample from the predictive distribution  $p(Y(x_0; c_j) | y, F)$  by drawing one  $Y(x_0; c_j) \sim N(f_j(x_0)^\top \beta_j + Z(x_0; c_j), \tau_j^2)$  for each posterior sample of  $\{\beta_j, \tau_j^2\}$  and  $Z(x_0; c_j)$ , found as above.

Our theoretical contribution includes the construction of valid positive-definite functions to serve as  $K_\theta((x; c_j), (x'; c_{j'}))$ . This is crucial for the above inferential framework as it ensures that the matrix  $K_\theta$  in Equation (2) will be positive definite for any finite set of distinct elements, observed or unobserved, in  $\mathbb{R} \times \mathcal{C}$ . A key advantage of driving the inference through a latent process is the convenience of predictive inference for the process and the response at new inputs.

A specific contribution of the proposed fully process-based framework is that it allows us to carry out predictive inference even for new unobserved groups. For example, if  $c_j \in \mathcal{C}$  for some  $j$  is a new group with no observations, sampling from the posterior predictive distributions of  $Z(x_0; c_j)$  and  $Y(x_0; c_j)$  can be executed as described above, possibly with appropriate modeling adjustments on the prior for  $\tau_j^2$ . Hence, we turn to the construction of valid MGGP covariance kernels.

It is worth pointing out the computational bottleneck arising from the dimension of  $K_\theta$  in GP models for large data sets. There is, by now, a substantial literature on various approaches that build scalable models for massive data sets by building low-rank or sparsity-inducing processes (see, e.g. Wikle, 2010; Banerjee, 2017; Heaton et al., 2019, for expository treatments of such processes) out of any valid covariance kernel. While scalable processes is not the focus of the current manuscript, it is important to point out that our approach of constructing MGGPs using valid covariance kernels renders the resulting processes as “scalable-ready” since low-rank or sparsity-inducing MGGPs can be easily derived using current methods.

### 3 Multi-Group Gaussian Processes

We focus upon building MGGPs using positive-definite covariance kernels. We start with simpler separable specifications and proceed to derive richer and more flexible alternatives.

#### 3.1 Separable multi-group GP

We start with a simple case where the covariance function over  $\mathbb{R}^p \times \mathcal{C}$  is separable.

**Definition 1.**  *$K$  is said to be separable if  $K((x, c_i), (x', c_j)) = K_{\mathbb{R}^p}(x, x')K_{\mathcal{C}}(c_i, c_j)$ , where  $K_{\mathbb{R}^p}$  and  $K_{\mathcal{C}}$  are over  $\mathbb{R}^p$  and  $\mathcal{C}$ .*

First, observe that  $K$  is positive definite if and only if both  $K_{\mathbb{R}^p}$  and  $K_{\mathcal{C}}$  are positive definite. Since GPs over  $\mathbb{R}^p$  have been thoroughly studied, we focus on covariance functions over  $K_{\mathcal{C}}$ , or GPs over a categorical set. Also,  $\mathcal{C}$  being finite, any function on  $\mathcal{C} \times \mathcal{C}$  is completely determined by  $C \in \mathbb{R}^{k \times k} : C_{ij} = K_{\mathcal{C}}(c_i, c_j)$ .

**Proposition 1.**  *$K_{\mathcal{C}}$  is positive definite if and only if  $C$  is a positive definite matrix.*

The above result implies that it will suffice to find a positive definite function on  $\mathbb{R}^p$  and a positive definite matrix  $C \in \mathbb{R}^{k \times k}$  to construct a separable positive definite function on  $\mathcal{X}$ . Homogeneous kernels arise as a special case.

**Definition 2.** A function  $K_{\mathcal{C}} : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}$  is said to be homogeneous if  $K_{\mathcal{C}}(c_i, c_j) = K_0(1_{\{c_i \neq c_j\}})$  for some function  $K_0$  on  $\{0, 1\}$ .

A homogeneous GP is completely determined by two scalars:  $a := K_{\mathcal{C}}(c_i, c_i)$ ,  $b := K_{\mathcal{C}}(c_i, c_j)$  where  $i \neq j$ , which represent the within-group correlation and across-group correlation, respectively. Without loss of generality, we assume  $a = 1$ ; otherwise, we can rescale  $K_{\mathcal{C}}$ . In this case, all within-group and between-group correlations are the same. A homogeneous GP is appropriate if we only need to distinguish pairs of observations that are in the same group from those in different groups, while the specific group identities are irrelevant. Equivalently,  $K$  is homogeneous if  $K$  is isotropic with respect to the discrete metric  $d(c_i, c_j) = 1_{\{c_i \neq c_j\}}$ .

**Corollary 1.** Let  $K_{\mathcal{C}}$  be homogeneous, then  $K_{\mathcal{C}}$  is positive definite if and only if  $-\frac{1}{k-1} \leq b \leq 1$ , where  $b = K_{\mathcal{C}}(c_i, c_j)$  with  $i \neq j$ .

The inequality  $-\frac{1}{k-1} \leq b \leq 1$  implies the across-group correlation should not dominate the within-group correlation, which is intuitively reasonable.

Separable models provide computational benefits because the resulting covariance matrix for  $Z$  can be expressed as a Kronecker product of  $K_{\mathbb{R}^p}$  and  $K_{\mathcal{C}}$ . However, separable covariance functions tend to have “ridges” or discontinuities (Stein, 2005) that can lead to poorer inference. They also assume that the same covariance structure ( $K_{\mathbb{R}^p}$ ) is retained for all groups, which seems unduly restrictive in terms of accommodating associations of the latent process for different pairs of inputs in  $\mathbb{R}^p \times \mathcal{C}$ . Hence, we turn to inseparable cases in the next section.

## 3.2 Inseparable multi-group GP

### 3.2.1 Isotropic MGGP

In order to discuss “isotropic” GPs on  $\mathbb{R}^p \times \mathcal{C}$ , we need to endow  $\mathcal{C}$  with more structure. To minimize assumptions, we introduce a metric  $d$  on  $\mathcal{C}$  so that  $(\mathcal{C}, d)$  is a metric space.

**Definition 3.** Given two metric spaces  $(\mathcal{Y}, d)$  and  $(\mathcal{Y}', d')$ , a GP on  $\mathcal{Y} \times \mathcal{Y}'$  is said to be semi-isotropic if  $K((x_1, x'_1), (x_2, x'_2)) = K_0(d(x_1, x_2), d'(x'_1, x'_2))$ .

Intuitively, a semi-isotropic GP is isotropic in each component  $\mathbb{R}^d$  and  $\mathcal{C}$  separately. Isotropic functions imply semi-isotropic, but the other direction does not hold in general. In practice,  $d$  is usually obtained by domain knowledge, including prior, exterior, or expert knowledge. For example, if  $\mathcal{C} = \{\text{North Carolina, New Jersey, California}\}$ , then the distance can be the geographical distance between the center of these states. As another example, if  $\mathcal{C}$  contains a set of human tissue

types, prior biomedical knowledge might lead to candidates for  $d$ ; two tissue types from different parts of the same organ, say the brain, might be expected to be more similar to each other than a brain tissue and liver tissue. When  $\mathcal{C}$  is a weighted graph, the graph distance serves as a valid metric (Bouttier et al., 2003). If no domain knowledge is available, we suggest using a default noninformative distance:  $d_{ij} = 1 - \delta_{ij}$ , i.e., all groups are equidistant.

Recall that the restriction of a GP to a subset of the original domain is again a GP (Rasmussen and Williams, 2005). Motivated by this fact, if we can isometrically embed  $\mathcal{C}$  to an Euclidean space  $\mathbb{R}^{p'}$  by some mapping  $\iota : \mathcal{C} \rightarrow \mathbb{R}^{p'}$  such that

$$d_{ij} := d(c_i, c_j) = \|\iota(c_i) - \iota(c_j)\|,$$

then an isotropic GP on  $\mathbb{R}^{p'}$  induces a GP on the image  $\iota(\mathcal{C})$ . Our next result creates a large family of semi-isotropic covariance functions on  $\mathcal{X} = \mathbb{R}^p \times \mathcal{C}$ .

**Theorem 1.** *Let  $\iota$  be an isometric embedding from  $(\mathcal{C}, d)$  to  $\mathbb{R}^{p'}$ . Then, if  $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}$  is a completely monotone function and  $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  is a positive function with a completely monotone derivative, then:*

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2}{(\psi(\|\iota(c_i) - \iota(c_j)\|^2))^{p/2}} \varphi\left(\frac{\|x - x'\|^2}{\psi(\|\iota(c_i) - \iota(c_j)\|^2)}\right). \quad (5)$$

*is a valid covariogram, where  $\sigma^2 > 0$  is the spatial variance. In particular, if  $d(c_i, c_j) = 1 - \delta_{ij}$  is the discrete metric, and if  $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}$  is a completely monotone function, then*

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2}{\alpha^{\frac{p}{2}(1-\delta_{ij})}} \varphi\left(\frac{\|x - x'\|^2}{\alpha^{1-\delta_{ij}}}\right) \quad (6)$$

*is a valid covariogram, where  $\sigma^2 \in (0, 1]$  is the spatial variance and  $\alpha > 0$  controls the interaction between  $\mathbb{R}^p$  and  $\mathcal{C}$ .*

There are two remaining questions to fully define the covariance function: How do we determine  $p'$ ? and how do we find the embedding  $\iota$ ? In fact, the explicit formula of  $\iota$  is not necessary to define the covariance function since only  $d$  is involved in Equation (5). Therefore, as long as  $d$  is known, the above covariance functions are well-defined. If one is interested in the embedding  $\iota$  itself, the following lemma answers the first question:

**Lemma 1** (Maehara (2013)). *Let  $G \in \mathbb{R}^{k \times k}$  be the Gram matrix of  $d$ , that is,*

$$G_{ij} = \frac{1}{2} (d(c_1, c_i) + d(c_1, c_j) - d(c_i, c_j)) .$$

*Then there exists an isotropic embedding  $\iota : \mathcal{C} \rightarrow \mathbb{R}^{p'}$  if and only if  $G$  is positive semi-definite with rank at most  $p'$ .*

For simplicity, we choose  $p' = \text{rank}(G)$  to reduce the dimension. For the second question, such an  $\iota$  can be obtained using existing methods in closed-form, including a power transform (Maehara, 2013; Hopkins, 2015). If  $\mathcal{C}$  has a graph structure, then graph embedding algorithms apply here, see Deza and Laurent (2009); Frankl et al. (2020) for more details.

The solution for  $\iota$  admits a simple form when  $d$  is the discrete metric, that is,  $d(c_i, c_j) = 1 - \delta_{ij}$ .

For the discrete metric, the Gram matrix is  $B = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & \frac{1}{2} & \cdots & \frac{1}{2} \\ 0 & \frac{1}{2} & 1 & \cdots & \frac{1}{2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \frac{1}{2} & \frac{1}{2} & \cdots & 1 \end{bmatrix} = \begin{bmatrix} 0 & \mathbf{0}_{1 \times (k-1)} \\ \mathbf{0}_{(k-1) \times 1} & \tilde{B} \end{bmatrix},$

where  $\tilde{B} = \frac{1}{2} \mathbf{I}_{k-1} + \frac{1}{2} \mathbf{1}_{(k-1) \times (k-1)}$  is full rank. As a result,  $\text{rank } B = k - 1 = p'$ , so we can isometrically embed  $\mathcal{C}$  to  $\mathbb{R}^{k-1}$ . A natural embedding  $\iota$  is to map each  $c_i$  to the vertex of the  $k - 1$  standard simplex  $\Delta_{k-1} \subset \mathbb{R}^{k-1}$ . Then the embedding  $\iota$  is given by

$$\iota : \mathcal{C} \rightarrow \Delta_{k-1}, c_i \mapsto \frac{1}{\sqrt{2}} \mathbf{e}_i - \frac{1}{(k-1)\sqrt{2}} \left( 1 + \frac{1}{\sqrt{k}} \right) \mathbf{1}_{k-1}, \quad 1 \leq i \leq k-1, \quad c_k \mapsto -\frac{1}{\sqrt{2k}} \mathbf{1}_{k-1}.$$

### 3.2.2 Candidate covariance functions for the MGGP

Some candidates for completely monotone functions  $\phi$  and positive functions with completely monotone derivatives  $\psi$  are in Table 3.2.2 (Gneiting, 2002):

$\phi(t)$	$\psi(t)$
$\exp(-ct^\gamma)$	$(at^\alpha + 1)^\beta$
$(2^{\nu-1} \Gamma(\nu))^{-1} (ct^{1/2})^\nu K_\nu(ct^{1/2})$	$\log(at^\alpha + b) / \log b$
$(1 + ct^\gamma)^{-\nu}$	$(at^\alpha + \beta) / (\beta(at^\alpha + 1))$
$2^\nu (\exp(ct^{1/2}) + \exp(-ct^{1/2}))^{-\nu}$	

Table 1: Candidate functions for completely monotone functions  $\phi$  and positive functions with completely monotone derivatives  $\psi$ . Here,  $a, c, \nu > 0$ ,  $b > 1$ ,  $0 < \alpha, \beta, \gamma \leq 1$ .

Selecting a  $\phi$  and a  $\psi$  function from Table 3.2.2, we obtain the following semi-isotropic covariance functions on  $\mathcal{X}$ , for more covariance functions, see Appendix 6.7:

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2}{(a^2 d_{ij}^2 + 1)^{p/2}} \exp \left\{ -\frac{b^2 \|x - x'\|^2}{a^2 d_{ij}^2 + 1} \right\}, \quad (7)$$

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2}{(a d_{ij} + 1)^{p/2}} \exp \left\{ -\frac{b^2 \|x - x'\|^2}{a d_{ij} + 1} \right\}, \quad (8)$$

$$K((x, c_i), (x', c_j)) = \begin{cases} \frac{\sigma^2 2c^{p/2}}{(a^2 d_{ij}^2 + 1)^\nu (a^2 d_{ij}^2 + c)^{p/2} \Gamma(\nu)} \left\{ \frac{b}{2} \left( \frac{a^2 d_{ij}^2 + 1}{a^2 d_{ij}^2 + c} \right)^{1/2} \|x - x'\| \right\}^\nu \\ \quad \times K_\nu \left( b \left( \frac{a^2 d_{ij}^2 + 1}{a^2 d_{ij}^2 + c} \right)^{1/2} \|x - x'\| \right) & x \neq x', \\ \frac{\sigma^2 c^{p/2}}{(a^2 d_{ij}^2 + 1)^\nu (a^2 d_{ij}^2 + c)^{p/2}} & x = x' \end{cases} \quad (9)$$

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2 c^{p/2}}{(a^2 d_{ij}^2 + 1)^{1/2} (a^2 d_{ij}^2 + c)^{p/2}} \exp \left\{ -b \left( \frac{a^2 d_{ij}^2 + 1}{a^2 d_{ij}^2 + c} \right)^{1/2} \|x - x'\| \right\}. \quad (10)$$

In the equations above,  $\sigma^2 > 0$  is the spatial variance,  $a \geq 0$  is the group similarity scale,  $b \geq 0$  is the feature scale,  $c \geq 0$  is the dependency scale, and  $\nu > 0$  is a smoothness parameter. The covariance functions in Equation (7) and Equation (8) are analogues of the squared exponential or radial basis function (RBF). The covariance function in Equation (9) is the analogue of the Matérn covariance function. In particular, the covariance function in Equation (10) is the special case of Equation (9) when  $\nu = 1/2$ , so we can call it the exponential covariance function. Equation (9) becomes separable when  $c = 1$ .

### 3.2.3 Stationary MGGP with $k = 2$

We now aim to weaken the isotropic condition described above. Since  $(\mathcal{C}, d)$  does not admit a natural group structure, we start with the simple case where  $k = 2$ , which appears frequently in practice, including in data sets where the two groups are male/female, adults/children, treatment/control, etc. Note that by rescaling, we can always embed a binary metric space within  $\mathbb{R}$ :  $\iota(c_1) = 0$ ,  $\iota(c_2) = 1$ . Then any spatial-temporal covariogram induces an semi-isotropic covariance function for  $\mathbb{R}^d \times \mathcal{C}$ .

For the non-isotropic case observe that  $\mathcal{C}$  can be identified with  $\mathbb{Z}_2$  when  $k = 2$ , an Abelian group. In this situation,  $K$  is said to be stationary if  $K((x, d), (x', l)) = K_0(x - x', d - l)$ . We will use  $K$  instead of  $K_0$  for simplicity, where  $K$  is characterized by two covariance functions:  $K_w = K(\cdot, 0)$ , the within-group covariance function, and  $K_c = K(\cdot, 1)$ , the cross-group covariance function. It is clear that any covariance function on  $\mathbb{R}^p \times \mathbb{Z}_2$  determines two covariance functions on  $\mathbb{R}^p$ . On the

other hand, not all pairs of covariance functions on  $\mathbb{R}^p$  define a valid (PD) covariance function on  $\mathbb{R}^p \times \mathbb{Z}_2$ . In order to construct a valid covariance function on  $\mathbb{R}^p \times \mathbb{Z}_2$ , we need to find a sufficient condition for  $K$  being PD.

**Theorem 2.** *Let  $K_w$  and  $K_c$  be two PD functions on  $\mathbb{R}^p$  with spectral densities  $\rho_w$  and  $\rho_c$  such that*

$$K(x, 0) = K_w(x) = \int_{\mathbb{R}^p} e^{-2\pi i \omega x} \rho_w(\omega) d\omega, \quad K(x, 1) = K_c(x) = \int_{\mathbb{R}^p} e^{-2\pi i \omega x} \rho_c(\omega) d\omega .$$

Then,  $K(x, l) = \begin{cases} K_w(x) & l = 0 \\ K_c(x) & l = 1 \end{cases}$  is PD on  $\mathbb{R}^p \times \mathbb{Z}_2$  if and only if  $\rho_w \geq \rho_c$ .

**Example 1.** *Recall the RBF covariance function in Equation (7):*

$$K(x, l) = \frac{\sigma^2}{(a^2 l^2 + 1)^{p/2}} \exp \left\{ -\frac{b^2 \|x\|^2}{a^2 l^2 + 1} \right\} .$$

The two spectral densities are given by

$$\begin{aligned} \rho_w(\omega) &= \sigma^2 \left( \frac{\pi}{b^2} \right)^{\frac{p}{2}} \exp \left\{ -\frac{\pi^2 \|\omega\|^2}{b^2} \right\}, \\ \rho_c(\omega) &= \sigma^2 \left( \frac{\pi}{b^2} \right)^{\frac{p}{2}} \exp \left\{ -\frac{\pi^2 (a^2 + 1) \|\omega\|^2}{b^2} \right\}, \end{aligned}$$

where  $\rho_w \geq \rho_c$ .

The stationary MGGP assumes homogeneity in the within-group correlation; however, for certain applications, heterogeneity exists. To account for this heterogeneity, we consider a weaker version, which we call the semi-stationary MGGP. A semi-stationary MGGP is stationary in  $\mathbb{R}^p$  but not in  $\mathcal{C}$ .

**Definition 4.**  *$K$  is said to be semi-stationary if  $K((x, c_i), (x', c_j)) = K_0(x - x', c_i, c_j)$  where  $K_0$  is defined on  $\mathbb{R}^p \times \mathcal{C} \times \mathcal{C}$ .*

This semi-stationary assumption is suitable for applications where groups are expected to have different within-group correlations, but the GP is stationary once the group is fixed. For semi-stationary MGGPs,  $K$  is determined by  $K_0(x) = K(x, 0, 0)$ ,  $K_c(x) = K(x, 0, 1) = K(x, 1, 0)$  and  $K_1 = K(x, 1, 1)$ , where  $K_0 \neq K_1$  in general, otherwise  $K$  becomes stationary.

**Theorem 3.** Let  $K_0$ ,  $K_c$  and  $K_1$  be PD functions on  $\mathbb{R}^p$  with spectral densities  $\rho_0$ ,  $\rho_c$ , and  $\rho_1$ .

$$\text{Then } K(x, l, l') = \begin{cases} K_0(x) & l = l' = 0 \\ K_c(x) & l + l' = 1 \text{ is PD on } \mathbb{R}^p \times \mathbb{Z}_2 \text{ if and only if } \rho_0 \rho_1 \geq \rho_c^2. \\ K_1(x) & l = l' = 1 \end{cases}$$

Data sets with more than two groups are ubiquitous in scientific applications. For example, in a biomedical setting, data may be collected from multiple human tissue types; from people of different race and ethnicity; and from varying age groups. Hence, we generalize the above theory to  $k > 2$  groups. The main difficulty here is that  $\mathcal{C}$  does not admit a natural group structure for  $k > 2$ . A straightforward solution would be to identify  $\mathcal{C}$  with  $Z_k$ , but the modular structure of  $Z_k$ , that is,  $1 - 0 = 2 - 1 = \dots k - 1 - (k - 1 - 1) \neq k - (k - 1)$ , is not satisfied in practice. Hence, Bochner's Theorem, which fully characterizes positive definite functions on locally compact Abelian groups, is not applicable.

Motivated by the proof of Theorem 3, we observe that this result is formally identical to an application of Cramér's Theorem to bivariate GPs.

**Lemma 2** (Cramér (1940)). Let  $\tilde{K} : \mathbb{R}^p \rightarrow \mathbb{R}^{2 \times 2}$  be a function with spectral density  $\tilde{\rho}_{ij}$ . Then,  $\tilde{K}$  is positive definite, hence defines a bivariate GP on  $\mathbb{R}^p$ , if and only if  $\tilde{\rho}(\omega) = \{\tilde{\rho}(\omega)\}_{i,j=1}^2$  is positive semi-definite for any  $\omega \in \mathbb{R}^p$ .

That is, Theorem 3 draws an equivalence between Bochner's Theorem on  $\mathbb{R}^p \times \mathbb{Z}_2$  and Cramér's Theorem on  $\mathbb{R}^p$  with bivariate outputs. This insight is crucial, as it enables us to leverage the large literature on bivariate GPs to construct two-group GPs. Furthermore, given that Cramér's Theorem holds for a general  $k$ -variate GP, this observation suggests a general theory for an arbitrary number of groups with  $k > 2$ , which we explore below.

### 3.2.4 Stationary MGGP with arbitrary $k$

We first draw similarities between the MGGP with  $k$  groups and  $k$ -variate GPs. Recall that a Gaussian  $k$ -variate random field  $\tilde{Z}$  on  $\mathcal{Y}$  is characterized by its cross-covariance function  $\tilde{K} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^{k \times k}$ :

$$\text{Cov}(\tilde{Z}(x), \tilde{Z}(x')) = \tilde{K}(x, x').$$

**Theorem 4.** Let  $\mathcal{G}$  be the space of all Gaussian random fields on  $\mathcal{Y} \times \mathcal{C}$ , where  $\mathcal{C} = \{c_1, \dots, c_k\}$  and  $\mathcal{V}$  is the space of all Gaussian  $k$ -variate random fields on  $\mathcal{Y}$ . Then

$$\Phi : \mathcal{G} \rightarrow \mathcal{V}, (\Phi(Z))_i(x) := Z(x, c_i), \forall Z \in \mathcal{G}$$

is a bijection, and its inverse  $\Phi^{-1}$  is given by

$$\Phi^{-1} : \mathcal{V} \rightarrow \mathcal{G}, (\Phi^{-1}(\tilde{Z}))(x, c_i) = \tilde{Z}_i(x), \forall \tilde{Z} \in \mathcal{V}.$$

The correspondence between the covariance function of  $Z$  and the cross-covariance function of  $\tilde{Z}$  is given by

$$K((x, c_i), (x', c_j)) = \tilde{K}(x, x')_{ij}.$$

This shows that in order to construct a  $k$ -group MGGP it will suffice to construct a  $k$ -variate GP, and vice versa. For our purposes, existing constructions of multi-variate GPs can be applied (see, e.g., Gneiting et al., 2010; Apanasovich and Genton, 2010; Genton and Kleiber, 2015, for possible ideas.)

The following Lemma is often referred to as Cramér’s Theorem, which is a multivariate generalization of Bochner’s Theorem.

**Lemma 3** (Cramér (1940)). *Let  $\tilde{K} : \mathbb{R}^p \rightarrow \mathbb{R}^{k \times k}$  be a matrix-valued function with elements  $(\tilde{K}(x))_{ij} = \tilde{K}_{ij}(x)$  and spectral density  $\tilde{\rho}_{ij}$ . Then,  $\tilde{K}$  is positive definite, hence defines a  $k$ -variate GP on  $\mathbb{R}^p$ , if and only if  $\tilde{\rho}(\omega) = \{\tilde{\rho}(\omega)\}_{i,j=1}^k$  is positive semi-definite for any  $\omega \in \mathbb{R}^p$ .*

We prove the following related result.

**Theorem 5.** *Let  $K : \mathbb{R}^p \times \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}$  be a function with  $K_{ij} = K(\cdot, c_i, c_j)$  being stationary on  $\mathbb{R}^p$  and spectral densities  $\rho_{ij}$ . Then,  $K$  is positive definite, hence defines a semi-stationary GP on  $\mathbb{R}^p \times \mathcal{C}$ , if and only if  $\rho(\omega) = \{\rho(\omega)\}_{i,j=1}^k$  is positive semi-definite for any  $\omega \in \mathbb{R}^p$ .*

Note that Theorem 3 is a special case of Theorem 5 when  $k = 2$  but can be proved differently. As a result, the connection between Bochner’s Theorem on  $\mathbb{R}^p \times \mathcal{C}$  and Cramér’s Theorem on  $\mathbb{R}^p$  is analogous to the relationship between multi-group GPs and multi-variate GPs.

### 3.3 Special cases of the MGGP

We discuss the relationship between the MGGP and several commonly used GPs for multi-group data. We find that each of these models is a special case of the MGGP, which demonstrates the generality of our model. The first simple model we consider is the separated GP, which models each group separately with  $k$  independent GPs on  $\mathbb{R}^p$ . Let  $K_1, \dots, K_k$  be the covariance functions of these  $k$  GPs.

**Proposition 2.** *The SGP described above is equivalent to MGGP with covariance function*

$$K((x, c_i), (x', c_j)) = \begin{cases} K_i(x, x') & i = j \\ 0 & i \neq j. \end{cases}$$

*That is, the SGP is a special case of the MGGP with zero between-group correlation.*

Another simple model ignores the group structure entirely, Thus, we treat each observation  $(x, c)$  as simply  $x$  by dropping the label  $c \in \mathcal{C}$ . Then, we can consider a GP on the union of all observations across all groups. Assume the covariance function of the UGP is  $K_{\cup}$ , which is a GP over  $\mathbb{R}^p$ . We obtain the following.

**Proposition 3.**  *$K_{\cup}$  is equivalent to MGGP with covariance function*

$$K((x, c_i), (x', c_j)) = K_{\cup}(x, x').$$

*That is, the within group and between-group correlations are all equal.*

Finally, we consider the hierarchical GP:

$$\mu_i \sim GP(0, K_g) ; \quad Z | c_i \sim GP(\mu_i, K_z).$$

where  $K_g$  is the covariance function at the group level, while  $K_z$  is the covariance function at the covariate level.

**Proposition 4.** *The HGP is equivalent to MGGP with covariance function*

$$K((x, c_i), (x', c_j)) = K_g(x, x') + 1_{\{i=j\}}K_z(x, x').$$

*That is, the HGP is a special case of the MGGP with identical within-group dependencies and identical across-group dependencies.*

These results reveal that the MGGP is a non-trivial generalization of existing GPs that allows heterogeneity through a variety of flexible covariance functions that it accommodates.

## 4 Experiments

### 4.1 Simulations

#### 4.1.1 Demonstrating the MGGP’s relationship to related models

We designed and conducted an experiment to ascertain the MGGP’s ability to recover the Separated GP, Union GP, and Hierarchical GP as special cases. We generated synthetic data from each of these models using the likelihood function corresponding to Equation (1) using  $k = 2$  groups. We specified a zero mean, i.e.,  $\mu(x; c_j) = 0$  for both groups, and specified the latent process through the covariance function corresponding to the three models. We set  $b = \sigma^2 = a = 1$  in

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2}{(a^2 d_{ij}^2 + 1)^{p/2}} \exp \left\{ -\frac{b^2 \|x - x'\|^2}{a^2 d_{ij}^2 + 1} \right\}. \quad (11)$$

(Note that  $a$  is only used in the generation of data from the MGGP). We also assumed  $\tau_1^2 = \tau_2^2 = \tau^2$  in Equation (1) and used  $\tau^2 = 0.1$  to generate our data. Using these settings, we generated  $n_1 = n_2 = 100$  measurements for each group.

We computed the log marginal likelihood of the data, i.e.,  $N(y | 0, K_\theta + D_\tau)$  under each model for each data set. For the SGP, UGP, and HGP, we used the RBF covariance function,  $K(x, x') = \sigma^2 \exp\{-b^2 \|x - x'\|^2\}$ . For the MGGP, we used the “multi-group” version of the RBF (Equation (11)). When computing the likelihood under each model, we fix  $b, \sigma^2$ , and  $\tau^2$  to their true values, for  $a$  we use a grid of values,  $a = 10^{-5}, 10^{-4}, \dots, 10^2$  and specified  $D_\tau = \tau^2 I_n$

We find that the MGGP performs at par with the UGP, SGP, and HGP in the expected regimes (Figure 1). In particular, the MGGP recovers the performance (as measured by the log marginal likelihood) of the SGP when  $a$  is large, and it recovers the performance of the UGP as  $a \rightarrow 0$ . For data generated from the MGGP, we find that the likelihood peaks at the true value of  $a$  and is higher than all other models at this value. These results i) serve as a demonstration of the role of the  $a$  parameter; ii) confirm numerically that the MGGP recovers these models in certain regimes; and iii) suggest that the MGGP is a viable generalization of these models.

#### 4.1.2 Estimation and inference for the MGGP

In our previous experiment, we used the multi-group covariance function in Equation (7) and fixed the value of  $a$ . However, in practice we will need to estimate this and all other covariance parameters from the data. Thus, we next evaluate our ability to fit the parameters of the MGGP using both maximum likelihood estimation and fully Bayesian posterior inference.

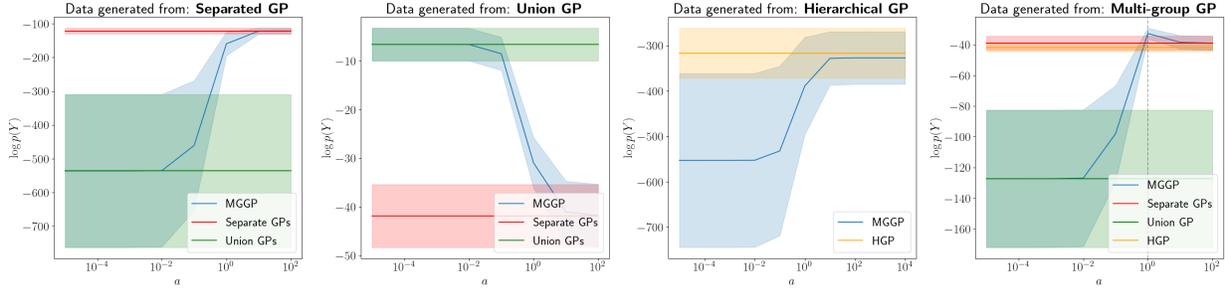


Figure 1: **Comparison between the MGGP, Separated GP, Union GP, and Hierarchical GP.** Using two-group data generated from each of the four models, we computed the log marginal likelihood of the data under each model. For the MGGP, we used the covariance function in Equation (7) and used a range of different values the parameter  $a$ . In the rightmost plot, the dashed vertical line indicates the true value of  $a$  used for data generation. We use an RBF covariance function for the *Separated GP* and *Union GP*, which does not have an  $a$  parameter. We repeated this experiment 20 times, and the bands in each plot represent 95% confidence intervals.

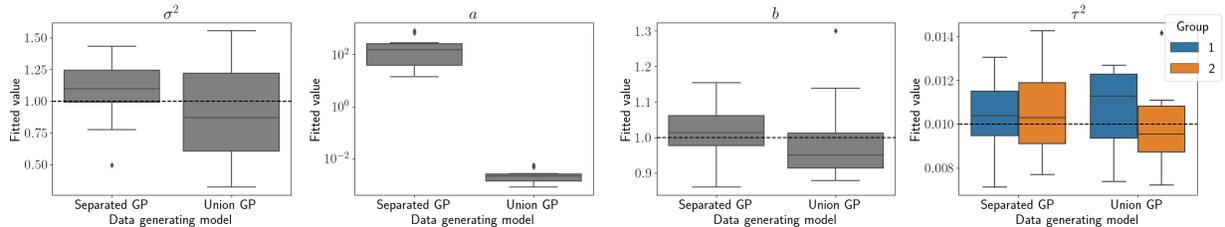


Figure 2: **Covariance function parameter estimation.** Using data generated from the SGP and UGP, we fit the MGGP by finding the maximum likelihood estimate for the true parameters of the kernel function in Equation (7).

**Maximum likelihood estimation.** We first conducted an experiment where we generated data from the SGP and UGP models as in the previous section. We maximize the collapsed or marginalized likelihood corresponding to Equation (1), i.e.,  $\mathcal{N}(y | 0, K_\theta + \tau^2 I_n)$  with respect to  $\theta = \{a, b, \sigma^2\}$  and a common measurement error variance  $\tau^2$ , where  $\theta$  corresponds to the three parameters in the multi-group RBF covariance function in Equation (7). We used a conjugate gradient ascent algorithm (Nocedal and Wright, 2006) to obtain the joint estimates of  $\{\theta, \tau^2\}$ . We implemented the algorithm in Python using the JAX software framework (Bradbury et al., 2018), which is designed for fast computation, compilation, and automatic differentiation. Experiments were run on an internal computing cluster using a 320 NVIDIA P100 GPU. We found that the MLE for  $a$  was consistently high for data generated from the SGP and low for the UGP data, as expected (Figure 2, middle panel). Additionally, the estimation was able to recover the true values for  $\sigma^2$  and  $b$  (Figure 2, left and right panels).

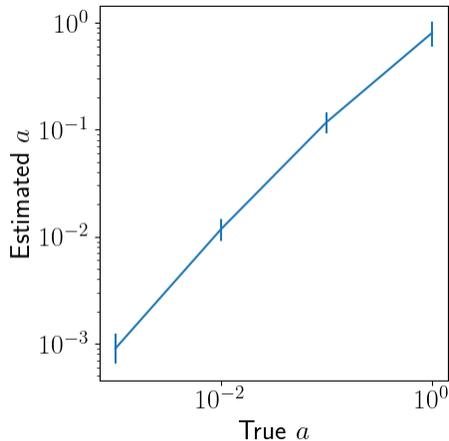


Figure 3: **Recovering  $a$  with ML estimation.** We generated synthetic data from the MGGP at different values for  $a$  and fit the MGGP to these data sets. Here, we fix all other parameters to their true values. We find that we are able to recover the true value of  $a$  to some extent.

Next, we evaluated whether the MGGP could recover the true value of  $a$  from synthetic data sampled from the MGGP model. Here, we generated four datasets, each with a different value of  $a$  for  $a \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0\}$ . We again optimized all parameters jointly by maximizing the marginal likelihood, and we examined the estimated value of  $a$  for each. We repeated this experiment ten times. We found that we could consistently estimate a reasonable value of  $a$  (Figure 3). While the estimated values do not exactly coincide with the true values, they still showed a monotone relationship. Together, these results show that likelihood-based parameter estimation is feasible in the MGGP model and that existing estimation and optimization techniques can be successfully applied.

**Bayesian Inference.** Next, we used a simulated dataset generated in the same manner as above to we conduct a full Bayesian analysis for the MGGP model as described in Section 2. Here, we also include a group-specific intercept for the  $k = 2$  groups given by  $\beta = (\beta_1, \beta_2)^\top$ , with the data generated according to  $\beta_1 = 1, \beta_2 = 2$ . We form the  $(n_1 + n_2) \times 2$  binary design matrix  $F$  appropriately in order to apply the group-specific intercept in the model. With  $\theta = \{a, b, \sigma^2\}$  the prior distribution in Equation (3) is specified as

$$p(\theta, \tau^2, \beta) = p(a, b, \sigma^2, \tau^2, \beta) = IG(a | \alpha_a, \alpha'_a)IG(b | \alpha_b, \alpha'_b)IG(\sigma^2 | \alpha_\sigma, \alpha'_\sigma)IG(\tau^2 | \alpha_\tau, \alpha'_\tau)N(\mu_\beta, V_\beta) ,$$

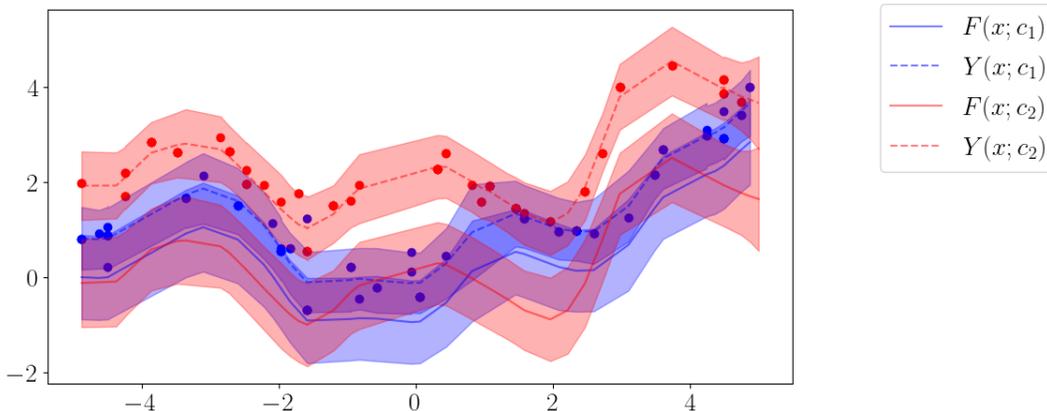


Figure 4: **Posterior predictive distribution under the MGGP.** The points represent training data; the solid lines show the means of the latent processes  $F(x; c)$ ; the dashed lines represent the predictive means of  $Y(x; c)$ ; and the shaded areas around the lines are twice the standard deviation of the PPD at the corresponding input points.

where we set  $\alpha_a = \alpha'_a = \alpha_b = \alpha'_b = \alpha_\sigma = \alpha'_\sigma = \alpha_\tau = \alpha'_\tau = 1$ ,  $\mu_\beta = 0$  and  $V_\beta^{-1} = I$ . We sample from the posterior distribution in Equation (3) using a Hamiltonian No U-Turn Sampling (Hoffman et al., 2014) algorithm as implemented in the `Stan` programming environment (Stan Development Team, 2020; Riddell et al., 2021).

We ran four chains with dispersed initial values for 400 iterations each. Convergence was diagnosed after 200 iterations using visual inspection of autocorrelation plots (Supplementary Figure 2) and computation of Gelman-Rubin R-hat and Monte Carlo standard errors, and the subsequent 800 samples were retained for posterior inference. The posterior median and 95% credible intervals are presented in Supplementary Figure 3, which shows that the posterior samples of the covariance function parameters center around their true values (Supplementary Figure 3). We also sample from the posterior predictive distribution (PPD),  $p(Y(x_0; c_j) | y, F)$  (see Section 2), for a collection of new inputs or test cases. These are presented in Figure 4. Notably, because all of the MGGP assumptions are encoded in the covariance function, any appropriate method for estimation and inference in standard GPs can be applied.

**Prediction** As a final experiment with simulated data, we evaluated the MGGP in its ability to predict held-out values in a Gaussian process regression task. We generated data from a GP regression model (Equation (1)), using the Separated GP, Union GP, HGP, and MGGP. We then fit these models on each of the data sets, using 50% of the data for fitting and computing predictions for the other 50%. We use the posterior mean as a point prediction for each of the  $n^*$  held-out

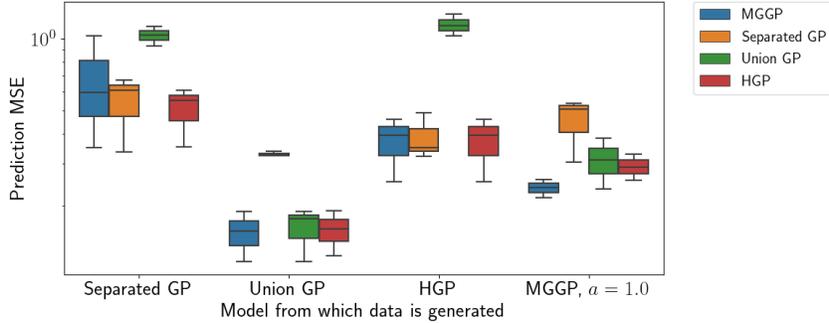


Figure 5: **GP Prediction with simulated data.** We generate data from each of the four models — the SGP, UGP, HGP, and MGGP — and fit each of these models too all datasets. Prediction error was computed on a hold-out dataset.

samples:

$$\mu^* = K_{X^*X} K_{XX}^{-1} y, \quad (12)$$

where  $K_{X^*X}$  is the  $n^* \times n$  matrix of covariance function evaluations for each pair of test and train samples, and  $K_{XX}$  is the  $n \times n$  matrix of covariance function evaluations for each pair of train samples. We mean-center the data for each group. To measure the quality of the predictions, we compute the mean squared error of the predictions,

$$E = \frac{1}{n^*} \sum_{i=1}^{n^*} (y_i - \mu_i^*)^2 \quad (13)$$

We found that the MGGP emulates the performance of the SGP, UGP, and HGP on their respective simulated datasets (Figure 5), while with data generated from the MGGP itself, the MGGP clearly outperformed the other models. While the Separated GP can be expected to perform well when each group has a large sample size, a primary benefit of the MGGP is its ability to share information across similar groups when the sample size is limited. Thus, we expect the performance of the MGGP to improve relative to the SGP in settings in which some groups have a small number of samples but are closely related to other, larger groups.

To test this hypothesis, we conducted another multi-group GP regression experiment in which we sought to predict the held-out values for one group that contained few samples. Specifically, we generated a synthetic dataset containing three groups, where group 1 and group 2 are similar to one another, and group 3 is dissimilar from the other two. We then generated a series of datasets, varying the number of samples in group 1 for each. Finally, we computed the prediction error for the SGP, UGP, HGP, and MGGP. We found that the MGGP outperformed the other models,

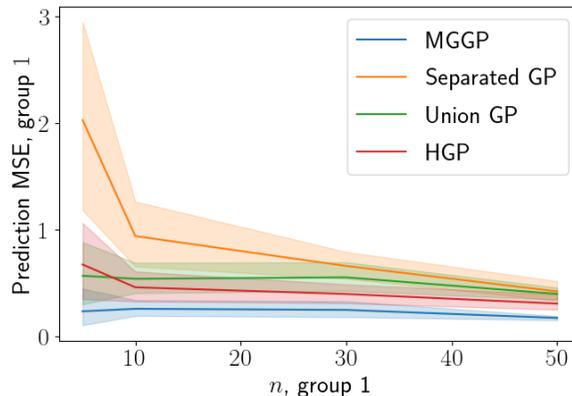


Figure 6: **Prediction using simulated data with imbalanced groups.** We perform a prediction experiment with  $k = 3$  groups. To generate a series of datasets, we fix the sample size of groups  $c_2$  and  $c_3$  to be 50, and vary the sample size of group  $c_1$ .

especially when the sample size for group 1 was small (Figure 6), confirming our hypothesis. This result shows that the MGGP thrives relative to other setups when the sample size for some groups is limited, as it can most effectively leverage information from similar groups.

## 4.2 Application to GTEx tissue samples

We next applied the MGGP to a large gene expression data set collected by the Genotype-Tissue Expression (GTEx) project (Consortium et al., 2020). The GTEx v8 data contain measurements from 17382 samples that span 52 tissue types collected from 838 human donors; see Appendix 6.9.1 for a full list of tissue types and the sample size for each tissue. Along with gene expression profiling, a variety of additional metadata characteristics were collected, including demographic variables and tissue health measurements.

In these experiments, we use Gaussian process regression models to analyze the relationship between a sample’s gene expression profile and its ischemic time (the duration of time between death and tissue collection). Previous work has shown a robust relationship between gene expression and ischemic time (Musella et al., 2013; Ferreira et al., 2018); however, whether this relationship exhibits tissue-specific patterns remains largely unknown. In these experiments, the groups correspond to tissue types.

### 4.2.1 Exploring group similarities with the MGGP

As an initial test with the GTEx data, we applied the MGGP to just two tissue types at a time. These experiments aim to validate that the MGGP can appropriately model known associations

between similar groups.

**Examining model fit via the log marginal likelihood** In a preliminary experiment, we examined three tissue types: tibial artery ( $n = 657$ ), coronary artery ( $n = 238$ ), and breast ( $n = 456$ ). First, for each of the three pairs of tissues, we fit the MGGP with maximum likelihood estimation, as described in Section 4.1.2, using the multi-group RBF covariance function in Equation (11). In this experiment, we fixed  $a$  to one value in a preset range, and found the maximum likelihood estimates of the remaining parameters. Using these maximum likelihood estimates and the fixed  $a$ , we then computed the log marginal likelihood of the data,

$$\log p(y|X, a, \hat{b}, \hat{\sigma}^2, \hat{\tau}^2) = -\frac{k}{2}2\pi - \frac{1}{2} \det(K_{XX} + \hat{\tau}^2 I) - \frac{1}{2} y^\top (K_{XX} + \hat{\tau}^2 I)^{-1} y,$$

where  $K_{XX}$  is the  $(n_1 + n_2) \times (n_1 + n_2)$  matrix of covariance function evaluations for each pair of samples. We also fit the Separated GP and Union GP for each pair of tissues using the standard RBF kernel, and computed the log marginal likelihood of the data under these models. Examining the log marginal likelihood across varying values of  $a$  (Figure 7), we found that two tissue types that are expected to be similar to one another — tibial and coronary artery — showed a higher marginal likelihood under small values of  $a$  ( $a \lesssim 0.01$ ), while tissues that have unique expression patterns — tibial artery and breast — showed a higher marginal likelihood under large values of  $a$  ( $a \gtrsim 10$ ). In both cases, the MGGP gracefully recovered the Separated GP and Union GP marginal likelihoods for  $a \rightarrow \infty$  and  $a \rightarrow 0$ , respectively. This result implies that the MGGP could be a viable strategy not only for sharing information across groups, but also for learning the group relationships themselves.

**Recovering group similarities** Next, we conducted a similar experiment in which we also estimated  $a$  from the data itself (along with all other model and covariance function parameters), again using maximum likelihood estimation. In this experiment, we applied the model to all 52 tissue types. We fit the MGGP for every pair of tissues, and extracted  $\hat{a}_{MLE}$  for each pair. This experiment yields  $\frac{1}{2}(52 \times 51) = 1326$  estimated values of  $a$  (one for each pair of tissue types).

We found that the estimated values of  $a$  reflect many of the expected relationships between the tissue types (Figure 8). Most notably, we found that 11 subtypes of brain tissue yield low values for  $a$ , suggesting that gene expression changes in these tissue types in a similar manner as ischemic time changes.

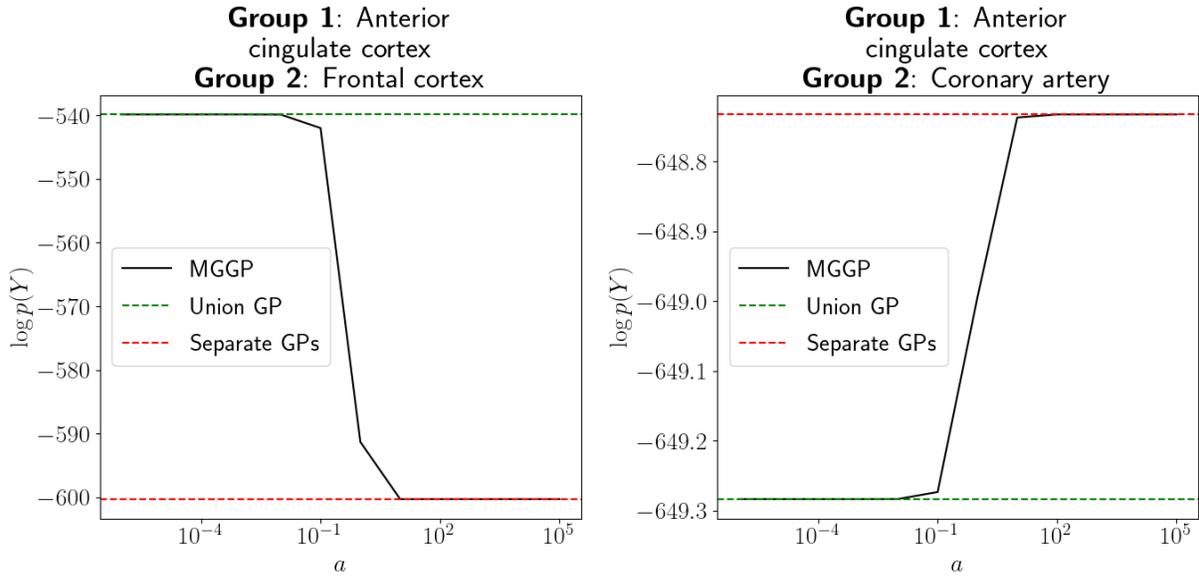


Figure 7: **Likelihood of GTEx gene expression data under the MGGP.** For two pairs of tissues, we computed the log marginal likelihood of the data under the MGGP with  $a$  set to be a range of values. Similar tissue types (anterior cingulate cortex and frontal cortex) prefer low values of  $a$ , while more dissimilar tissues (anterior cingulate cortex and coronary artery) prefer high values of  $a$ .

**Recovering group similarities with group-specific errors** As a final test of the MGGP in the two-group setting, we performed a similar pairwise analysis as in the previous paragraph, but here we allowed each group to have its own noise variance  $\tau_j$  as in Equation 1. We apply this model to a subset of the tissue types: Brain Cortex, Brain Hypothalamus, Esophagus Gastroesophageal Junction, and Esophagus Mucosa. We fit the MGGP to each of the six pairs of tissues and extracted the estimated value for  $a$ . We found that the resulting values of  $a$  recovered similarities between tissue types: the MGGP obtained a small estimated value for the two brain tissues and for the two esophagus tissues, while all other values were quite large (Supplementary Figure 4). This small-scale experiment demonstrates that the MGGP behaves similarly when group-specific noise variances are included. More generally, these experiments indicate that, in the pairwise case, the estimated value of  $a$  reflects prior biological knowledge about similarity of tissue types.

#### 4.2.2 Leveraging group relationships for prediction

We next conducted a series of experiments with the GTEx data to test whether the MGGP’s model of group relationships could improve downstream predictive performance relative to standard GPs. In each of these experiments, we randomly split the data into 50% training samples and 50% testing



Figure 8: **Estimation of  $a$  for each pair of GTEx tissue types.** Cell  $ij$  in the heatmap represents  $\log_{10}(a_{ij})$ , where  $a_{ij}$  is the maximum likelihood estimate of  $a$  when fitting the MGGP using tissues  $i$  and  $j$ . Lower values of  $a$  (red) indicate higher similarity, while higher values of  $a$  (black) indicate lower similarity.

samples, and mean-centered each dataset within each group. We fit the MGGP on the training set using maximum likelihood estimation, using the multi-group RBF kernel in all experiments here (Equation (11)). Then, we computed the predicted ischemic time for each test sample using the predictive mean (Equation (12)) for the ischemic time for each gene expression sample and computed the mean squared error relative to the true ischemic time (Equation (13)). For comparison, we also fit the Union and Separated GPs using the standard RBF covariance function and the Hierarchical GP using the standard RBF as both the within- and between-group covariance functions.

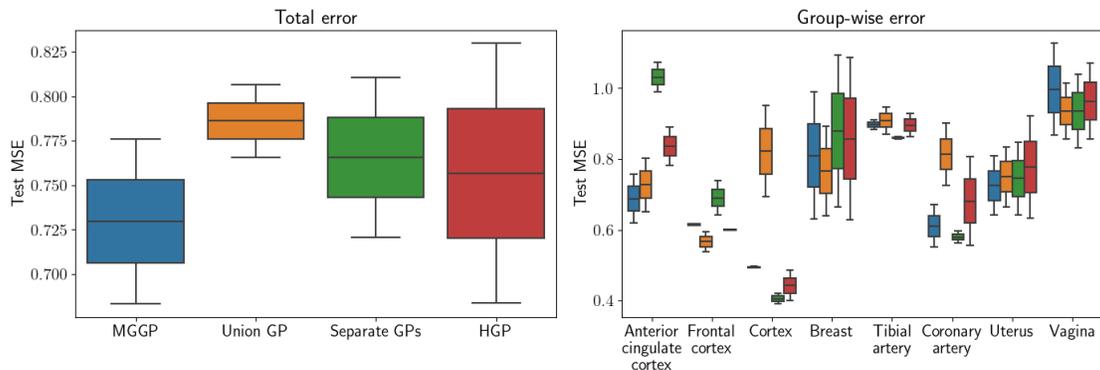


Figure 9: **Prediction with GTEx data.** We fit each of the four models in a regression setting, where the covariates are gene expression profiles of samples from various tissue types, and the response is each sample’s ischemic time. We compute the error on a held-out set of data.

**Benchmarking the MGPP, UGP, SGP, and HGP** We first performed a prediction exercise using eight GTEx tissue types: Anterior Cingulate Cortex, Frontal Cortex, Cortex, Breast, Tibial Artery, Coronary Artery, Uterus, and Vagina. Given *a priori* similarities between tissue types from different body parts, we specify the “distance” between two similar tissue types ( $d_{ij}$ ) to be small, while the distance between dissimilar tissue types is large. Here, we identify sets of similar tissues as  $\{\text{Anterior Cingulate Cortex, Frontal Cortex, Cortex}\}$ ,  $\{\text{Breast}\}$ ,  $\{\text{Tibial Artery, Coronary Artery}\}$ , and  $\{\text{Uterus, Vagina}\}$ . For each pair of tissues within each of these subsets, we specify the distance to be  $d_{ij} = 10$ , and for each pair of tissues between subsets, we specify the distance as  $d_{ij} = 0.1$ . The distance from a group to itself is zero,  $d_{ij} = 0$  for  $i = j$ . Note that, for the multi-group RBF covariance function (Equation (11)), only the *relative* distances between groups are relevant, as the  $a$  parameter as a scaling factor on these relative distances, and we estimate  $a$  from the data. Thus, our distance settings imply an assumption that tissue types within each of these subsets are ten-fold more similar than tissues between the subsets.

Examining the prediction error for each of the four models, we find that MGPP outperforms UGP, SGP and HGP in terms of overall prediction error, although the HGP performs somewhat comparably (Figure 9, left panel). Further examining the error for each tissue, we find that the performance varies substantially across tissue types (Figure 9, right panel).

**Fine-tuning group distances** While we observe that the MGPP performs well in a prediction setting, the *a priori* distances between tissues were chosen somewhat heuristically as a coarse approximation of prior biological knowledge. However, one may wonder whether these distances can be chosen more precisely, or perhaps even be learned from the data itself. To test this, we performed

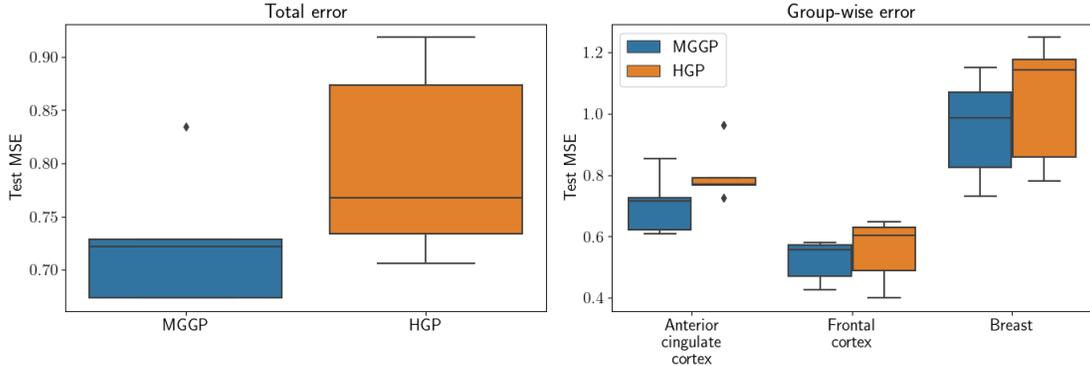


Figure 10: **Fine-tuning group distances and comparison of the MGGP and HGP via prediction with GTEx data.** We use just three tissue types and carefully tune the pairwise distances between the groups using a data-driven approach.

another prediction experiment with three tissue types: Anterior Cingulate Cortex, Frontal Cortex, and Breast. Instead of hard-coding the distances, we used a data-driven approach that uses two steps. In the first step, we fit a two-group MGGP on each of the three pairs of tissues, and used the resulting MLE for  $a$  as the distance between two tissues,  $d_{ij} = \hat{a}_{ij}$ . For the two cortex tissues, we find  $a \approx 10^{-4}$ , and for each of the cortex and breast pairs, we find  $a \approx 1$ . In the second step, we fit the full three-group MGGP using these learned distances. We also fit the HGP for comparison.

We find that after using this data-driven approach to select the group-wise distances, the MGGP performs better than before, both in terms of overall prediction error and group-wise error (Figure 10). Moreover, we find that the MGGP more clearly outperforms the HGP in this case. This difference is likely due to the nature of the HGP’s model for group relationships: the HGP treats all groups symmetrically, in the sense that the only relevant quantity is whether two samples are in the same group or different groups. In contrast, the MGGP relies on the specific group identities of any two samples, as well as the distance between those groups. This result implies that carefully defining the group relationships can be important for the MGGP, and that the demonstrated data-driven approach is a viable strategy for doing so.

**Assessing univariate relationships** The MGGP regression framework also allows for assessing relationships between the outcome variable and individual explanatory variables. To demonstrate this, we again leveraged data for three GTEx tissue types: Anterior Cingulate Cortex, Frontal Cortex, and Breast. In this experiment, we fit the MGGP to these data using just one gene at a time, finding a one-dimensional regression function describing each gene’s relationship to

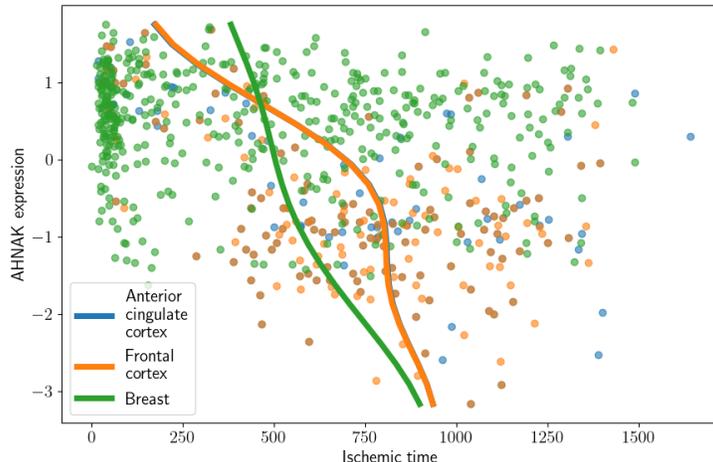


Figure 11: **Examining univariate relationships with the MGGP.** For three tissue types, each sample’s expression of the gene *AHNAK* (x-axis) and its ischemic time (y-axis) are shown by the scatter points. The solid lines show the MGGP’s posterior predictive mean when fit on this single gene in these three tissues.

ischemic time. We fit the model using maximum likelihood estimation and the multi-group RBF covariance function (Equation (11)), and we used the between-group distances found in the previous experiment. We then computed the predicted mean (Equation (12)) for a dense grid of input points, and we examined the resulting regression functions.

We find that the MGGP identifies several trends in the relationship between gene expression and ischemic time that show group-level trends. For instance, we find that expression of the gene *AHNAK*, which codes for a neuroblast differentiation-associated protein, decays similarly for the two cortex tissues, while this decay pattern follows a different trend in breast tissue (Figure 11). Moreover, we find that this relationship is robust to the choice of group distances (Supplementary Figure 5). This experiment demonstrates that the MGGP can be used for more specifically characterizing the group-wise relationships present in a dataset.

## 5 Discussion

In this work, we present the multi-group Gaussian process (MGGP), a flexible approach for modeling complex dependencies in datasets with subgroup structure. We present several options for constructing valid covariance functions on  $\mathbb{R}^p \times \mathcal{C}$ , and we show that this structure generalizes several existing GP models. We demonstrate the behavior of the MGGP through a series of simulation experiments and an application to a dataset of gene expression measurements.

Several future directions remain to be explored. First, this paper lays the groundwork for developing new positive definite covariance functions on  $\mathbb{R}^p \times \mathcal{C}$ . A particularly interesting direction is to construct covariance functions whose within-group and between-group correlations exhibit fundamentally different structure (e.g., the within-group correlation may be Matérn-like, while the between group correlation may be RBF-like). Second, as briefly mentioned in Section 2, recent advances in scalable GPs could be applied to the MGGP. For example, sparsity-inducing GPs have received much attention in the recent past (see, e.g., Datta et al., 2016; Katzfuss and Guinness, 2021; Peruzzi et al., 2020) and such methods can be applied to the class of MGGP models presented here. Finally, although we primarily focused on GP regression in this paper, the MGGP can be applied in any GP-based model. Thus, there could be benefits from using the MGGP in classification models, latent variable models, and other applications. We envision the MGGP being a flexible tool in a variety of contexts.

## 6 Appendix

### 6.1 Proof of Proposition 1

*Proof.* First, assume  $K_2$  is positive definite, and let the observation set be  $\mathcal{C} = \{c_1, \dots, c_k\}$ . Then the covariance matrix is exactly  $C$ , which is positive definite as well.

We assume  $C$  is positive definite. Given any observation set, we can rearrange the observations to be  $\{\underbrace{c_1, \dots, c_1}_{n_1}, \dots, \underbrace{c_k, \dots, c_k}_{n_k}\}$ . Then the corresponding covariance matrix  $\Sigma$  can be written as the following block matrix:

$$\Sigma = \begin{bmatrix} C_{11}\mathbf{1}_{n_1 \times n_1} & C_{12}\mathbf{1}_{n_1 \times n_2} & \cdots & C_{1k}\mathbf{1}_{n_1 \times n_k} \\ C_{21}\mathbf{1}_{n_2 \times n_1} & C_{22}\mathbf{1}_{n_2 \times n_2} & \cdots & C_{2k}\mathbf{1}_{n_2 \times n_k} \\ \vdots & \ddots & & \vdots \\ C_{k1}\mathbf{1}_{n_k \times n_1} & C_{k2}\mathbf{1}_{n_k \times n_2} & \cdots & C_{kk}\mathbf{1}_{n_k \times n_k} \end{bmatrix},$$

where  $\mathbf{1}$  is the matrix with all entries equal to 1. Let  $P_i \in \mathbb{R}^{n_i \times n_i}$  be rotation matrices such that  $P_i\mathbf{1}_{n_i} = \sqrt{n_i}\mathbf{e}_{n_i}$ , where  $\mathbf{e}_{n_i} = [1, 0, \dots, 0]^\top \in \mathbb{R}^{n_i}$ ,  $i = 1, \dots, k$ , and let  $P := \text{diag}\{P_1, \dots, P_k\} \in$

$\mathbb{R}^{\sum_i n_i \times \sum_i n_i}$ . Then,

$$\begin{aligned}
P\Sigma P^\top &= \begin{bmatrix} P_1 & & \\ & \ddots & \\ & & P_k \end{bmatrix} \begin{bmatrix} C_{11}\mathbf{1}_{n_1 \times n_1} & C_{12}\mathbf{1}_{n_1 \times n_2} & \cdots & C_{1k}\mathbf{1}_{n_1 \times n_k} \\ C_{21}\mathbf{1}_{n_2 \times n_1} & C_{22}\mathbf{1}_{n_2 \times n_2} & \cdots & C_{2k}\mathbf{1}_{n_2 \times n_k} \\ \vdots & \ddots & & \vdots \\ C_{k1}\mathbf{1}_{n_k \times n_1} & C_{k2}\mathbf{1}_{n_k \times n_2} & \cdots & C_{kk}\mathbf{1}_{n_k \times n_k} \end{bmatrix} \begin{bmatrix} P_1^\top & & \\ & \ddots & \\ & & P_k^\top \end{bmatrix} \\
&= \begin{bmatrix} C_{11}P_1\mathbf{1}_{n_1 \times n_1}P_1^\top & \cdots & C_{1k}P_1\mathbf{1}_{n_1 \times n_k}P_k^\top \\ \vdots & \ddots & \vdots \\ C_{k1}P_k\mathbf{1}_{n_k \times n_1}P_1^\top & \cdots & C_{kk}P_k\mathbf{1}_{n_k \times n_k}P_k^\top \end{bmatrix} \\
&= \begin{bmatrix} C_{11}(P_1\mathbf{1}_{n_1})(P_1\mathbf{1}_{n_1})^\top & \cdots & C_{1k}(P_1\mathbf{1}_{n_1})(P_k\mathbf{1}_{n_k})^\top \\ \vdots & \ddots & \vdots \\ C_{k1}(P_k\mathbf{1}_{n_k})(P_1\mathbf{1}_{n_1})^\top & \cdots & C_{kk}(P_k\mathbf{1}_{n_k})(P_k\mathbf{1}_{n_k})^\top \end{bmatrix} \\
&= \begin{bmatrix} C_{11}n_1\mathbf{e}_{n_1}\mathbf{e}_{n_1}^\top & \cdots & C_{1k}\sqrt{n_1n_k}\mathbf{e}_{n_1}\mathbf{e}_{n_k}^\top \\ \vdots & \ddots & \vdots \\ C_{k1}\sqrt{n_kn_1}\mathbf{e}_{n_k}\mathbf{e}_{n_1}^\top & \cdots & C_{kk}n_k\mathbf{e}_{n_k}\mathbf{e}_{n_k}^\top \end{bmatrix}
\end{aligned}$$

As a result, eigenvalues of  $\Sigma$  are the same as eigenvalues of  $\tilde{\Sigma} = \begin{bmatrix} C_{11}n_1 & \cdots & C_{1k}\sqrt{n_1n_k} \\ \vdots & \ddots & \vdots \\ C_{k1}\sqrt{n_kn_1} & \cdots & C_{kk}n_k \end{bmatrix}$ .

Let  $v = [\sqrt{n_1}, \dots, \sqrt{n_k}]^\top$  then

$$\tilde{\Sigma} = C \odot vv^\top.$$

Then by assumption and Schur product theorem (Styan, 1973),  $\tilde{\Sigma}$  is positive semi-definite, which finishes the proof.  $\square$

## 6.2 Proof of Corollary 1

*Proof.* By Proposition 1,  $K_2$  is positive definite if and only if  $C$  is positive definite. Under the

assumption of homogeneity,  $C = \begin{bmatrix} 1 & b & \cdots & b \\ b & 1 & \cdots & b \\ \vdots & \vdots & \ddots & \vdots \\ b & b & \cdots & 1 \end{bmatrix}$ . We can rewrite  $C = (1-b)I_k + b\mathbf{1}_{k \times k}$ . Recall

that  $\mathbf{1}_{k \times k}$  has eigenvalue  $p$  with multiplicity 1, and the corresponding eigenvector is  $\mathbf{1}_k$ . Moreover, 0 is another eigenvalue with multiplicity  $k-1$ . As a result, eigenvalues of  $\Sigma$  are  $1-b+bp$  with

multiplicity 1 and  $1 - b$  with multiplicity  $-1$ . We can conclude that  $C$  is positive definite if and only if  $-\frac{1}{k-1} \leq b \leq 1$ .  $\square$

### 6.3 Proof of Theorem 1

*Proof.* By the isometric embedding, finding  $K$  is equivalent to finding a semi-isotropic covariance function  $K_0$  in  $\mathbb{R}^p \times \mathbb{R}^{p'}$ . For any completely monotone function  $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}$  and positive function  $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  with a completely monotone derivative, Gneiting (2002) proved that  $K_0((x_1, x'_1), (x_2, x'_2)) = \frac{\sigma^2}{(\psi(\|x'_1 - x'_2\|^2))^{p/2}} \varphi\left(\frac{\|x_1 - x_2\|^2}{\psi(\|x'_1 - x'_2\|^2)}\right)$  is positive definite for any  $\sigma^2 > 0$ . As a result,

$$K((x, c_i), (x', c_j)) := K_0((x, \iota(c_i)), (x', \iota(c_j)))$$

is positive definite.  $\square$

### 6.4 Proof of Theorem 2

Recall the general form of Bochner's Theorem for a locally compact Abelian group:

**Lemma 4** (Bochner's Theorem). *Let  $G$  be a locally compact Abelian group and  $\widehat{G}$  be its dual group, then for any continuous positive-definite function  $K$  on  $G$ , there exists a unique positive measure  $\mu$  on  $\widehat{G}$  such that*

$$K(g) = \int_{\widehat{G}} \xi(g) d\mu(\xi).$$

Note that  $G = \mathbb{R}^p \times \mathbb{Z}_2$  is a locally compact Abelian group and the dual group is  $\widehat{G} = \mathbb{R}^p \times U_2$ , where  $U_2$  is the group of second roots of unity, that is,  $U_2 = \{1, -1\}$ .  $\widehat{G}$  acts on  $G$  as

$$(\omega, z)((x, l)) = e^{-2\pi i \omega x} z^l.$$

The spectral measure of  $K$ , denoted by  $\mu$  on  $\mathbb{R}^p \times U_2$  splits as  $\mu_1 \times \mu_2$  where  $\mu_1$  is a measure on  $\mathbb{R}^p$  and  $\mu_2$  is a measure on  $U_2$ . Then we claim that

$$K(x, l) = \sum_{z \in U_2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} z^l \rho(\omega, z) d\omega,$$

where  $\rho(\omega, 1) = \frac{1}{2}\rho_w(\omega) + \frac{1}{2}\rho_c(\omega)$  and  $\rho(\omega, -1) = \frac{1}{2}\rho_w(\omega) - \frac{1}{2}\rho_c(\omega)$ . We derive from the right hand

side to the left hand side. First consider the case when  $l = 0$ :

$$\begin{aligned}
& \sum_{z \in U_2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} z^0 \rho(\omega, z) d\omega \\
&= \frac{1}{2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} (\rho_w(\omega) + \rho_c(\omega)) d\omega + \frac{1}{2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} (\rho_w(\omega) - \rho_c(\omega)) d\omega \\
&= \int_{\mathbb{R}^p} e^{-2\pi i \omega x} \rho_w(\omega) d\omega \\
&= K_w(x) = K(x, 0).
\end{aligned}$$

Similarly, when  $l = 1$ ,

$$\begin{aligned}
& \sum_{z \in U_2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} z^1 \rho(\omega, z) d\omega \\
&= \frac{1}{2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} (\rho_w(\omega) + \rho_c(\omega)) d\omega - \frac{1}{2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} (\rho_w(\omega) - \rho_c(\omega)) d\omega \\
&= \int_{\mathbb{R}^p} e^{-2\pi i \omega x} \rho_c(\omega) d\omega \\
&= K_c(x) = K(x, 1).
\end{aligned}$$

As a result,  $K$  is PD  $\iff \rho$  is a positive measure  $\iff \rho_w \geq \rho_c$ , and the Theorem follows.

## 6.5 Proof of Theorem 3

The generalized spectral measure of  $K$ , denoted by  $\mu$  (with density  $\rho$ ) on  $\mathbb{R}^p \times U_2 \times U_2$ , splits as  $\mu_1 \times \mu_2$  where  $\mu_1$  is a measure on  $\mathbb{R}^p$  and  $\mu_2$  is a positive semi-definite measure on  $U_2 \times U_2$ . Then

$$K(x, l, l') = \sum_{z \in U_2} \sum_{z' \in U_2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} z^l \overline{z'}^{l'} \rho(\omega, z, z') d\omega.$$

We claim that  $\rho(\omega, z, z') = \begin{cases} \frac{1}{4} (\rho_0(\omega) + 2\rho_c(\omega) + \rho_1(\omega)) & z = z' = 1 \\ \frac{1}{4} (\rho_0(\omega) - \rho_1(\omega)) & zz' = -1 \\ \frac{1}{4} (\rho_0(\omega) - 2\rho_c(\omega) + \rho_1(\omega)) & z = z' = -1. \end{cases}$

We derive from the right hand side to the left hand side. First consider the case when  $l = l' = 0$ :

$$\begin{aligned}
& \sum_{z \in U_2} \sum_{z' \in U_2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} z^0 \overline{z'^0} \rho(\omega, z, z') d\omega \\
&= \frac{1}{4} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} (\rho_0(\omega) + 2\rho_c(\omega) + \rho_1(\omega) + \rho_0(\omega) - \rho_1(\omega) + \rho_0(\omega) - \rho_1(\omega) + \rho_0(\omega) - 2\rho_c(\omega) + \rho_1(\omega)) d\omega \\
&= \int_{\mathbb{R}^p} e^{-2\pi i \omega x} \rho_0(\omega) d\omega \\
&= K_0(x) = K(x, 0, 0).
\end{aligned}$$

Similarly, when  $l = 0, l' = 1$ ,

$$\begin{aligned}
& \sum_{z \in U_2} \sum_{z' \in U_2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} z^0 \overline{z'^1} \rho(\omega, z, z') d\omega \\
&= \frac{1}{4} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} (\rho_0(\omega) + 2\rho_c(\omega) + \rho_1(\omega) - \rho_0(\omega) + \rho_1(\omega) + \rho_0(\omega) - \rho_1(\omega) - \rho_0(\omega) + 2\rho_c(\omega) - \rho_1(\omega)) d\omega \\
&= \int_{\mathbb{R}^p} e^{-2\pi i \omega x} \rho_c(\omega) d\omega \\
&= K_c(x) = K(x, 0, 1).
\end{aligned}$$

When  $l = l' = 1$ ,

$$\begin{aligned}
& \sum_{z \in U_2} \sum_{z' \in U_2} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} z^1 \overline{z'^1} \rho(\omega, z, z') d\omega \\
&= \frac{1}{4} \int_{\mathbb{R}^p} e^{-2\pi i \omega x} (\rho_0(\omega) + 2\rho_c(\omega) + \rho_1(\omega) - \rho_0(\omega) + \rho_1(\omega) - \rho_0(\omega) + \rho_1(\omega) + \rho_0(\omega) - 2\rho_c(\omega) + \rho_1(\omega)) d\omega \\
&= \int_{\mathbb{R}^p} e^{-2\pi i \omega x} \rho_1(\omega) d\omega \\
&= K_1(x) = K(x, 1, 1).
\end{aligned}$$

As a result,  $K$  is PD  $\iff \rho$  is a positive semi-definite measure  $\iff (\rho_0 + 2\rho_c + \rho_1)(\rho_0 - 2\rho_c + \rho_1) - (\rho_0 - \rho_1)^2 = 4(\rho_0\rho_1 - \rho_c^2) \geq 0$ , and the Theorem follows.

## 6.6 Proof of Theorem 4

*Proof.* Given a Gaussian random field  $Z$  on  $\mathcal{Y} \times \mathcal{C}$  with covariance function  $K$ , we prove that  $\tilde{Z} := \Phi(Z)$  is a Gaussian  $k$ -variate random field on  $\mathcal{Y}$  with cross-covariance function  $\tilde{K}$ . It suffices

to check that  $\text{Cov}(\tilde{Z}(x), \tilde{Z}(x')) = \tilde{K}(x, x')$ , for any  $x, x' \in \mathcal{Y}$ .

$$\begin{aligned}
\text{Cov}(\tilde{Z}(x), \tilde{Z}(x')) &= \text{Cov}([Z(x, c_1), \dots, Z(x, c_k)]^\top, [Z(x', c_1), \dots, Z(x', c_k)]^\top) \\
&= \begin{bmatrix} K((x, c_1), (x', c_1)) & K((x, c_1), (x', c_2)) & \cdots & K((x, c_1), (x', c_k)) \\ K((x, c_2), (x', c_1)) & K((x, c_2), (x', c_2)) & \cdots & K((x, c_2), (x', c_k)) \\ \vdots & \vdots & \ddots & \vdots \\ K((x, c_k), (x', c_1)) & K((x, c_k), (x', c_2)) & \cdots & K((x, c_k), (x', c_k)) \end{bmatrix} \\
&= \begin{bmatrix} \tilde{K}_{11} & \tilde{K}_{12} & \cdots & \tilde{K}_{1k} \\ \tilde{K}_{21} & \tilde{K}_{22} & \cdots & \tilde{K}_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{K}_{k1} & \tilde{K}_{k2} & \cdots & \tilde{K}_{kk} \end{bmatrix} \\
&= \tilde{K}(x, x').
\end{aligned}$$

Then assume  $\tilde{Z}$  is a Gaussian  $k$ -variate random field on  $\mathcal{Y}$ . We prove that  $Z: \Phi^{-1}(\tilde{Z})$  is a Gaussian random field on  $\mathcal{Y} \times \mathcal{C}$  with covariance function  $K$ . It suffices to check that  $\text{Cov}(Z(x, c_i), Z(x', c_j)) = K((x, c_i), (x', c_j))$  for any  $x, x' \in \mathcal{Y}$ ,  $i, j = 1, \dots, k$ . Then,

$$\begin{aligned}
\text{Cov}(Z(x, c_i), Z(x', c_j)) &= \text{Cov}(\tilde{Z}_i(x), \tilde{Z}_i(x')) \\
&= \tilde{K}(x, x')_{ij} \\
&= K((x, c_i), (x', c_j)).
\end{aligned}$$

□

## 6.7 More covariance functions

We provide more semi-isotropic covariance functions below.

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2(a^2 d_{ij}^2 + 1)}{[(a^2 d_{ij}^2 + 1)^2 + b^2 \|x - x'\|^2]^{\frac{p+1}{2}}}, \quad (14)$$

$$K((x, c_i), (x', c_j)) = \frac{\sigma^2(ad_{ij} + 1)}{[(ad_{ij} + 1)^2 + b^2 \|x - x'\|^2]^{\frac{p+1}{2}}}, \quad (15)$$

$$K((x, c_i), (x', c_j)) = \sigma^2 \exp \{-a^2 d_{ij}^2 - b^2 \|x - x'\|^2 - cd_{ij}^2 \|x - x'\|^2\}, \quad (16)$$

$$K((x, c_i), (x', c_j)) = \sigma^2 \exp \{-ad_{ij} - b^2 \|x - x'\|^2 - cd_{ij} \|x - x'\|^2\}. \quad (17)$$

## 6.8 Code

All code for the models and experiments can be found in our GitHub repository: <https://github.com/andrewcharlesjones/multi-group-GP>. Within this repository, we provide an installable Python package for model fitting, covariance functions, prediction, and analysis.

## 6.9 Data

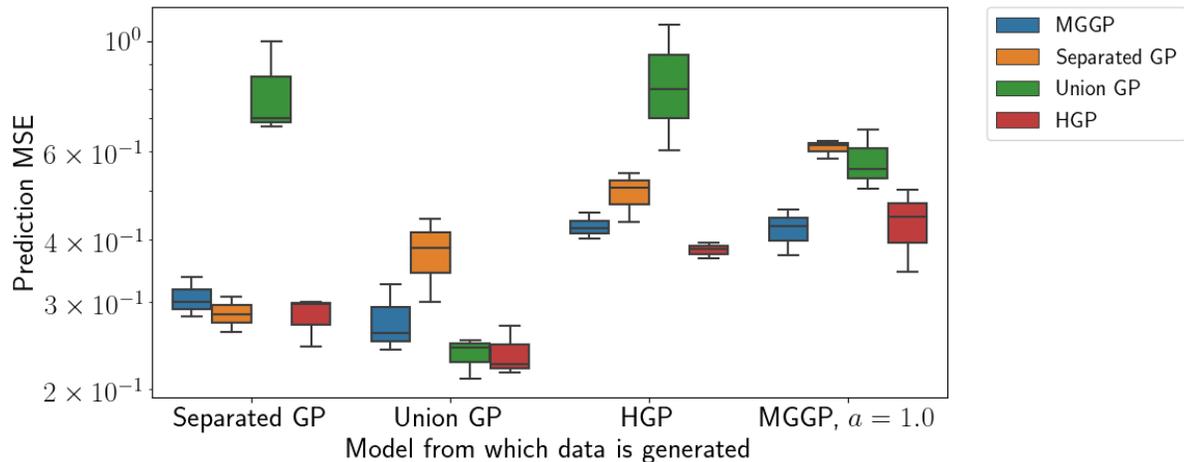
### 6.9.1 GTEx

The GTEx data can be downloaded from the GTEx portal: <https://gtexportal.org/home/datasets>. We use 52 tissue types, listed below, although some experiments use a subset of these tissue types. The sample size for each tissue type is shown in square brackets. Adipose Subcutaneous [644], Adipose Visceral (Omentum) [539], Adrenal Gland [254], Artery Aorta [424], Artery Coronary [238], Artery Tibial [657], Bladder [21], Brain Amygdala [147], Brain Anterior cingulate cortex (BA24) [172], Brain Caudate (basal ganglia) [230], Brain Cerebellar Hemisphere [208], Brain Cerebellum [241], Brain Cortex [255], Brain Frontal Cortex (BA9) [200], Brain Hippocampus [188], Brain Hypothalamus [193], Brain Nucleus accumbens (basal ganglia) [232], Brain Putamen (basal ganglia) [194], Brain Spinal cord (cervical c-1) [155], Brain Substantia nigra [133], Breast Mammary Tissue [456], Cells Cultured fibroblasts [444], Cells EBV-transformed lymphocytes [174], Cervix Ectocervix [9], Cervix Endocervix [10], Colon Sigmoid [373], Colon Transverse [406], Esophagus Gastroesophageal Junction [375], Esophagus Mucosa [551], Esophagus Muscularis [515], Fallopian Tube [9], Heart Atrial Appendage [422], Heart Left Ventricle [428], Kidney Cortex [85], Kidney Medulla [4], Liver [224], Lung [573], Minor Salivary Gland [162], Nerve Tibial [615], Ovary [180], Pancreas [37], Pituitary [283], Prostate [245], Skin Not Sun Exposed (Suprapubic) [595], Skin Sun Exposed (Lower leg) [665], Small Intestine Terminal Ileum [187], Spleen [233], Stomach [359], Testis [359], Uterus [142], Vagina [156], Whole Blood [139].

### 6.10 Supplementary figures

## References

- Apanasovich, T. V. and Genton, M. G. (2010). Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika*, 97(1):15–30.
- Banerjee, S. (2017). High-Dimensional Bayesian Geostatistics. *Bayesian Analysis*, 12(2):583 – 614.



Supplementary Figure 1: **Prediction experiment with synthetic data using a Matérn covariance function.**

Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2014). *Hierarchical modeling and analysis for spatial data*. CRC Press, Boca Raton, FL.

Bernardo, J., Berger, J., Dawid, A., Smith, A., et al. (1998). Regression and classification using Gaussian process priors. *Bayesian Statistics*, 6:475.

Borovitskiy, V., Azangulov, I., Terenin, A., Mostowsky, P., Deisenroth, M., and Durrande, N. (2021). Matérn Gaussian processes on graphs. In *International Conference on Artificial Intelligence and Statistics*, pages 2593–2601. PMLR.

Borovitskiy, V., Terenin, A., Mostowsky, P., and Deisenroth, M. (2020). Matérn Gaussian processes on Riemannian manifolds. In *Advances in Neural Information Processing Systems*, volume 33, pages 12426–12437.

Bouttier, J., Di Francesco, P., and Guitter, E. (2003). Geodesic distance in planar graphs. *Nuclear physics B*, 663(3):535–567.

Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: composable transformations of Python+NumPy programs.

Consortium, G. et al. (2020). The GTEx Consortium atlas of genetic regulatory effects across human tissues. *Science*, 369(6509):1318–1330.



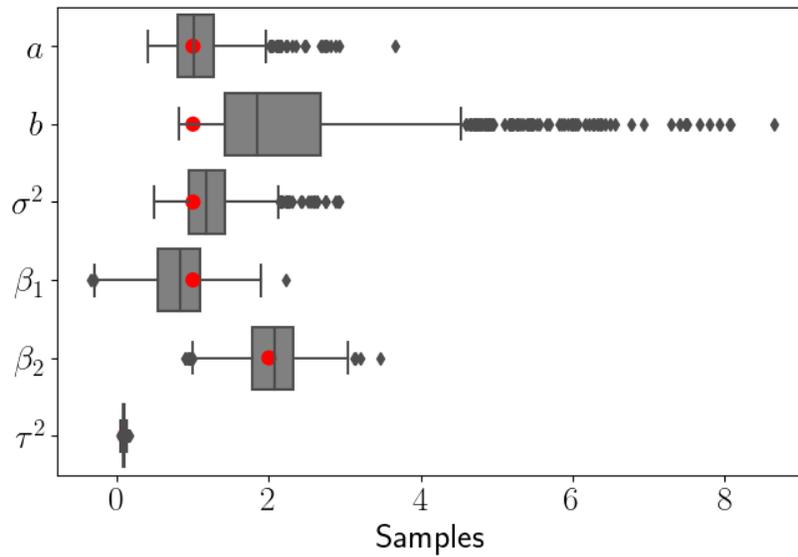
Supplementary Figure 2: **Autocorrelation plots for MCMC samples from MGGP posterior.**

Cramér, H. (1940). On the theory of stationary random processes. *Annals of Mathematics*, pages 215–230.

Cressie, N. and Wikle, C. K. (2011). *Statistics for spatio-temporal data*. John Wiley & Sons, Hoboken, NJ.

Damianou, A. and Lawrence, N. D. (2013). Deep Gaussian processes. In *Artificial Intelligence and Statistics*, pages 207–215. PMLR.

Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016). Hierarchical nearest-neighbor



Supplementary Figure 3: **Posterior samples of covariance function and model parameters.** Boxplots show the 95% confidence intervals for each set of samples. Red points indicate the parameter values used to generate the data.

Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111(514):800–812.

Deisenroth, M. P., Fox, D., and Rasmussen, C. E. (2013). Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423.

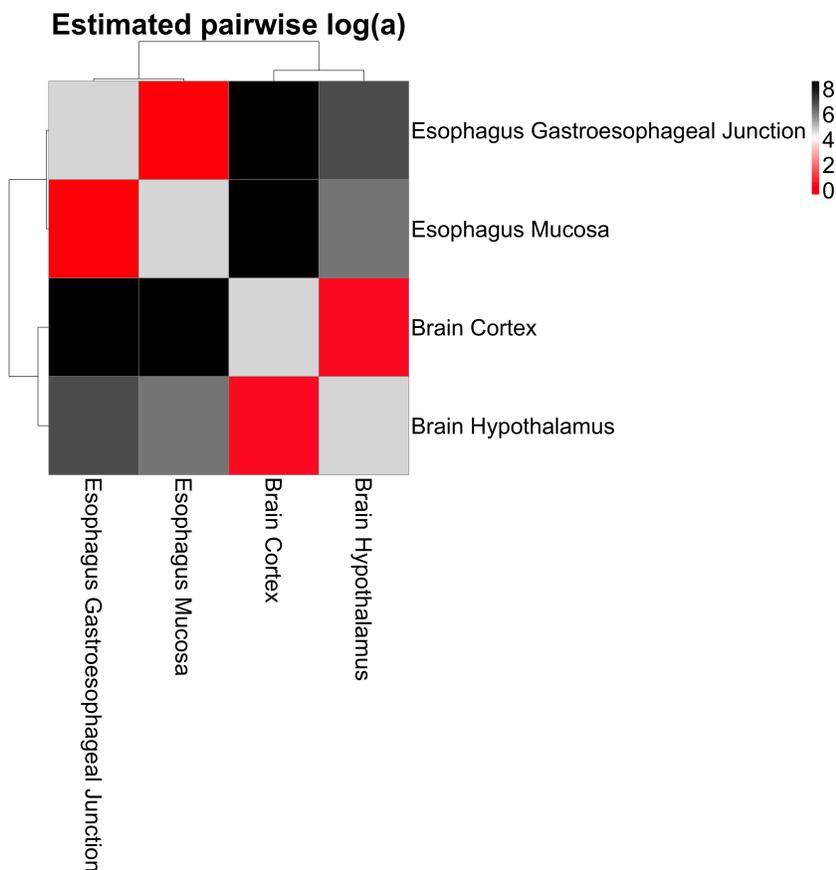
Deza, M. M. and Laurent, M. (2009). *Geometry of cuts and metrics*, volume 15. Springer.

Dunson, D. B. (2000). Bayesian latent variable models for clustered mixed outcomes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(2):355–366.

Dunson, D. B., Chen, Z., and Harry, J. (2003). A bayesian approach for joint modeling of cluster size and subunit-specific outcomes. *Biometrics*, 59(3):521–530.

Dunson, D. B., Wu, H.-T., and Wu, N. (2020). Diffusion based Gaussian processes on restricted domains. *arXiv preprint arXiv:2010.07242*.

Ferreira, P. G., Muñoz-Aguirre, M., Reverter, F., Godinho, C. P. S., Sousa, A., Amadoz, A., Sodaei, R., Hidalgo, M. R., Pervouchine, D., Carbonell-Caballero, J., et al. (2018). The effects of



Supplementary Figure 4: **Estimation of  $a$  for each pair of GTEx tissue types.** Cell  $ij$  in the heatmap represents  $\log_{10}(a_{ij})$ , where  $a_{ij}$  is the maximum likelihood estimate of  $a$  when fitting the MGGP using tissues  $i$  and  $j$ . Lower values of  $a$  (red) indicate higher similarity, while higher values of  $a$  (black) indicate lower similarity. Here, we allow each group its own noise variance  $\tau_j^2$ .

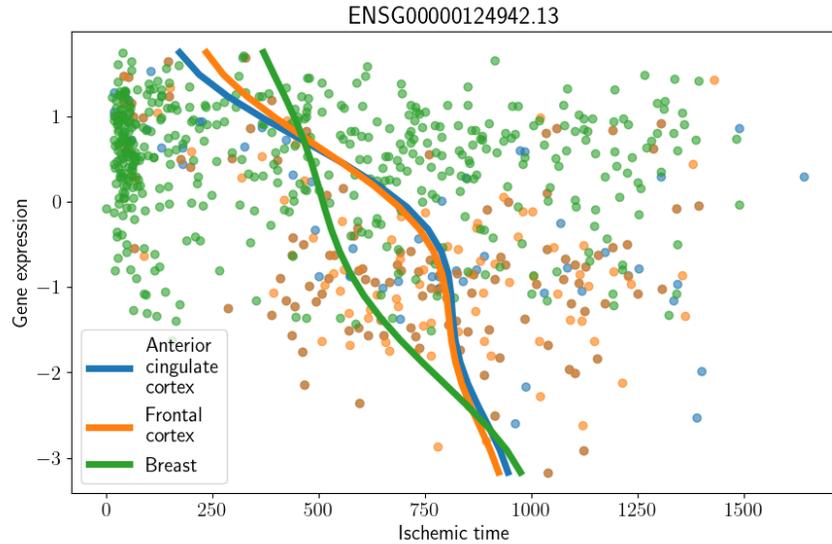
death and post-mortem cold ischemia on human tissue transcriptomes. *Nature communications*, 9(1):1–15.

Frankl, N., Kupavskii, A., and Swanepoel, K. J. (2020). Embedding graphs in Euclidean space. *Journal of Combinatorial Theory, Series A*, 171:105146.

Gelfand, A. E., Diggle, P., Guttorp, P., and Fuentes, M. (2010). *Handbook of spatial statistics*. CRC press, Boca Raton, FL.

Genton, M. G. and Kleiber, W. (2015). Cross-covariance functions for multivariate geostatistics. *Statistical Science*, 30(2):147–163.

Ghosal, S. and Van der Vaart, A. (2017). *Fundamentals of nonparametric Bayesian inference*, volume 44. Cambridge University Press.



Supplementary Figure 5: **Examining univariate relationships with the MGGP.** We reran the experiment shown in Figure 11, but here we used a noninformative distance matrix such that all three groups were equidistant.

Gneiting, T. (2002). Nonseparable, stationary covariance functions for space–time data. *Journal of the American Statistical Association*, 97(458):590–600.

Gneiting, T., Kleiber, W., and Schlather, M. (2010). Matérn cross-covariance functions for multivariate random fields. *Journal of the American Statistical Association*, 105(491):1167–1177.

Heaton, M., Datta, A., Finley, A., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R., Hammerling, D., Katzfuss, M., Lindgren, F., Nychka, D., Sun, F., and Zammit-Mangion, A. (2019). Methods for analyzing large spatial data: A review and comparison. *Journal of Agricultural, Biological and Environmental Statistics*, 24(3):398–425.

Hensman, J., Lawrence, N. D., and Rattray, M. (2013). Hierarchical Bayesian modelling of gene expression time series across irregularly sampled replicates and clusters. *BMC bioinformatics*, 14(1):1–12.

Hoffman, M. D., Gelman, A., et al. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623.

Hopkins, R. (2015). Finite metric spaces and their embedding into Lebesgue spaces.

Katzfuss, M. and Guinness, J. (2021). A General Framework for Vecchia Approximations of Gaussian Processes. *Statistical Science*, 36(1):124 – 141.

- Li, D., Tang, W., and Banerjee, S. (2021). Fixed-domain inference for Gaussian processes with Matérn covariogram on compact Riemannian manifolds. *arXiv preprint arXiv:2104.03529*.
- Maehara, H. (2013). Euclidean embeddings of finite metric spaces. *Discrete Mathematics*, 313(23):2848–2856.
- Matheron, G. (1963). Principles of geostatistics. *Economic geology*, 58(8):1246–1266.
- Murray, J. S. and Reiter, J. P. (2016). Multiple imputation of missing categorical and continuous values via bayesian mixture models with local dependence. *Journal of the American Statistical Association*, 111(516):1466–1479.
- Musella, V., Verderio, P., Reid, J. F., Pizzamiglio, S., Gariboldi, M., Callari, M., Massimo, M., De Cecco, L., Veneroni, S., Pierotti, M. A., et al. (2013). Effects of warm ischemic time on gene expression profiling in colorectal cancer tissues and normal mucosa. *PloS one*, 8(1):e53406.
- Niu, M., Cheung, P., Lin, L., Dai, Z., Lawrence, N., and Dunson, D. (2019). Intrinsic Gaussian processes on complex constrained domains. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(3):603–627.
- Nocedal, J. and Wright, S. (2006). *Numerical optimization*. Springer Science & Business Media.
- Pan, T., Hu, G., and Shen, W. (2020). Identifying latent groups in spatial panel data using a Markov random field constrained product partition model. *arXiv preprint arXiv:2012.10541*.
- Park, S. and Choi, S. (2010). Hierarchical gaussian process regression. In *Proceedings of 2nd Asian conference on machine learning*, pages 95–110. JMLR Workshop and Conference Proceedings.
- Peruzzi, M., Banerjee, S., and Finley, A. O. (2020). Highly scalable Bayesian geostatistical modeling via meshed Gaussian processes on partitioned domains. *Journal of the American Statistical Association*. in press.
- Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, MA, first edition.
- Regev, A., Teichmann, S. A., Lander, E. S., Amit, I., Benoist, C., Birney, E., Bodenmiller, B., Campbell, P., Carninci, P., Clatworthy, M., et al. (2017). Science forum: the human cell atlas. *elife*, 6:e27041.
- Riddell, A., Hartikainen, A., and Carter, M. (2021). pystan (3.0.0). PyPI.

- Ru, B., Alvi, A., Nguyen, V., Osborne, M. A., and Roberts, S. (2020). Bayesian optimisation over multiple continuous and categorical inputs. In *International Conference on Machine Learning*, pages 8276–8285. PMLR.
- Schulam, P., Wigley, F., and Saria, S. (2015). Clustering longitudinal clinical marker trajectories from electronic health data: Applications to phenotyping and endotype discovery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.
- Stan Development Team (2020). Stan modeling language users guide and reference manual.
- Stein, M. L. (1999). *Interpolation of spatial data: Some theory for kriging*. Springer Science & Business Media.
- Stein, M. L. (2005). Space–time covariance functions. *Journal of the American Statistical Association*, 100(469):310–321.
- Styan, G. P. (1973). Hadamard products and multivariate statistical analysis. *Linear algebra and its applications*, 6:217–240.
- Teimourian, M., Baghfalaki, T., Ganjali, M., and Berridge, D. (2015). Joint modeling of mixed skewed continuous and ordinal longitudinal responses: a bayesian approach. *Journal of Applied Statistics*, 42(10):2233–2256.
- Tsherniak, A., Vazquez, F., Montgomery, P. G., Weir, B. A., Kryukov, G., Cowley, G. S., Gill, S., Harrington, W. F., Pantel, S., Krill-Burger, J. M., et al. (2017). Defining a cancer dependency map. *Cell*, 170(3):564–576.
- U.S. Census Bureau, . (2020). 2020 census. U.S. Department of Commerce.
- Wikle, C. K. (2010). Low-rank representations for spatial processes. *Handbook of Spatial Statistics*, pages 107–118. Gelfand, A. E., Diggle, P., Fuentes, M. and Guttorp, P., editors, Chapman and Hall/CRC, pp. 107-118.