DOI: https://doi.org/10.51387/22-NEJSDS13

Scalable Marginalization of Correlated Latent Variables with Applications to Learning Particle Interaction Kernels

MENGYANG GU*, XUBO LIU, XINYI FANG, AND SUI TANG

Abstract

Marginalization of latent variables or nuisance parameters is a fundamental aspect of Bayesian inference and uncertainty quantification. In this work, we focus on scalable marginalization of latent variables in modeling correlated data, such as spatio-temporal or functional observations. We first introduce Gaussian processes (GPs) for modeling correlated data and highlight the computational challenge, where the computational complexity increases cubically fast along with the number of observations. We then review the connection between the state space model and GPs with Matérn covariance for temporal inputs. The Kalman filter and Rauch-Tung-Striebel smoother were introduced as a scalable marginalization technique for computing the likelihood and making predictions of GPs without approximation. We introduce recent efforts on extending the scalable marginalization idea to the linear model of coregionalization for multivariate correlated output and spatio-temporal observations. In the final part of this work, we introduce a novel marginalization technique to estimate interaction kernels and forecast particle trajectories. The computational progress lies in the sparse representation of the inverse covariance matrix of the latent variables, then applying conjugate gradient for improving predictive accuracy with large data sets. The computational advances achieved in this work outline a wide range of applications in molecular dynamic simulation, cellular migration, and agent-based models.

KEYWORDS AND PHRASES: Marginalization, Bayesian inference, Scalable computation, Gaussian process, Kalman filter, Particle interaction.

1. INTRODUCTION

Given a set of latent variables in a model, do we fit a model with a particular set of latent variables, or do we integrate out the latent variables when making predictions? Marginalization of latent variables is an iconic feature of the Bayesian analysis. The art of marginalization in statistics can at least be traced back to the De Finetti's theorem [12], which states that an infinite sequence $\{X_i\}_{i=1}^{\infty}$ is exchangeable, if and if only if there exists a random variable $\theta \in \Theta$ with probability distribution $\pi(\cdot)$, and a conditional distribution $p(\cdot \mid \theta)$, such that

$$p(x_1, \dots, x_N) = \int \left\{ \prod_{i=1}^N p(x_i \mid \theta) \right\} \pi(\theta) d\theta.$$
 (1.1)

Marginalization of nuisance parameters for models with independent observations has been comprehensively reviewed in [8]. Bayesian model selection [6, 4] and Bayesian model averaging [42], as two other examples, both rely on the marginalization of parameters in each model.

For spatially correlated data, the Jefferys prior of the covariance parameters in a Gaussian process (GP), which is proportional to the squared root of the Fisher information matrix of the likelihood, often leads to improper posteriors [7]. The posterior of the covariance parameter becomes proper if the prior is derived based on the Fisher information matrix of the marginal likelihood, after marginalizing out the mean and variance parameters. The resulting prior, after marginalization, is a reference prior, which has been studied for modeling spatially correlated data and computer model emulation [39, 46, 28, 20, 37].

Marginalization of latent variables has lately been aware by the machine learning community as well, for purposes of uncertainty quantification and propagation. In [29], for instance, the deep ensembles of models with a scoring function were proposed to assess the uncertainty in deep neural networks, and it is closely related to Bayesian model averaging with a uniform prior on parameters. This approach was further studied in [64], where the importance of marginalization is highlighted. Neural networks with infinite depth were shown to be equivalent to a GP with a particular kernel function in [38], and it was lately shown in [32] that the results of deep neural networks can be reproduced by GPs, where the latent nodes are marginalized out.

In this work, we study the marginalization of latent variables for correlated data, particularly focusing on scalable computation. Gaussian processes have been ubiquitously used for modeling spatially correlated data [3] and emulating computer experiments [50]. Computing the likelihood in

^{*}Corresponding author.

GPs and making predictions, however, cost $\mathcal{O}(N^3)$ operations, where N is the number of observations, due to finding the inverse and determinant of the covariance matrix. To overcome the computational bottleneck, various approximation approaches, such as inducing point approaches [54], fixed rank approximation [10], integrated nested Laplace approximation [48], stochastic partial differential equation representation [33], local Gaussian process approximation [15], and hierarchical nearest-neighbor Gaussian process models [11], circulant embedding [55], many of which can be summarized into the framework of Vecchia approximation [59, 27]. Scalable computation of a GP model with a multidimensional input space and a smooth covariance function is of great interest in recent years.

The exact computation of GP models with smaller computational complexity was less studied in past. To fill this knowledge gap, we will first review the stochastic differential equation representation of a GP with the Matérn covariance and one-dimensional input variable [63, 22], where the solution can be written as a dynamic linear model [62]. Kalman filter and Rauch-Tung-Striebel smoother [26, 45] can be implemented for computing the likelihood function and predictive distribution exactly, reducing the computational complexity of GP using a Matérn kernel with a halfinteger roughness parameter and 1D input from $\mathcal{O}(N^3)$ to $\mathcal{O}(N)$ operations. Here, interestingly, the latent states of a GP model are marginalized out in Kalman Filter iteratively. Thus the Kalman filter can be considered as an example of marginalization of latent variables, which leads to efficient computation. Note that the Kalman filter is not directly applicable for GP with multivariate inputs, yet GPs with some of the widely used covariance structures, such as the product or separable kernel [5] and linear model of coregionalization [3], can be written as state space models on an augmented lattice [18, 17]. Based on this connection, we introduce a few extensions of scalable marginalization for modeling incomplete matrices of correlated data.

The contributions of this work are twofold. First, the computational scalability and efficiency of marginalizing latent variables for models of correlated data and functional data are less studied. Here we discuss the marginalization of latent states in the Kalman filter in computing the likelihood and making predictions, with only $\mathcal{O}(N)$ computational operations. We discuss recent extensions on structured data with multi-dimensional input. Second, we develop new marginalization techniques to estimate interaction kernels of particles and to forecast trajectories of particles, which have wide applications in agent-based models [9], cellular migration [23], and molecular dynamic simulation [43]. The computational gain comes from the sparse representation of inverse covariance of interaction kernels, and the use of the conjugate gradient algorithm [24] for iterative computation. Specifically, we reduce the computational order from $\mathcal{O}((nMDL)^3) + \mathcal{O}(n^4L^2M^2D)$ operations in recent studies [34, 13] to $\mathcal{O}(Tn^2MDL) + \mathcal{O}(n^2MDL\log(nMDL))$ operations based on training data of M simulation runs, each containing n particles in a D dimensional space at Ltime points, with T being the number of iterations in the sparse conjugate gradient algorithm. This allows us to estimate interaction kernels of dynamic systems with many more observations. Here the sparsity comes from the use of the Matérn kernel, which is distinct from any of the approximation methods based on sparse covariance structures.

The rest of the paper is organized below. We first introduce the GP as a surrogate model for approximating computationally expensive simulations in Section 2. The state space model representation of a GP with Matérn covariance and temporal input is introduced in Section 3.1. We then review the Kalman filter as a computationally scalable technique to marginalize out latent states for computing the likelihood of a GP model and making predictions in Section 3.2. In Section 3.3, we discuss the extension of latent state marginalization in linear models of coregionaliztion for multivariate functional data, spatial and spatiotemporal data on the incomplete lattice. The new computationally scalable algorithm for estimating interaction kernel and forecasting particle trajectories is introduced in Section 4. We conclude this study and discuss a few potential research directions in Section 5. The code and data used in this paper are publicly available: https://github.com/ UncertaintyQuantification/scalable_marginalization.

2. BACKGROUND: GAUSSIAN PROCESS

We briefly introduce the GP model in this section. We focus on computer model emulation, where the GP emulator is often used as a surrogate model to approximate computer experiments [52]. Consider a real-valued unknown function $z(\cdot)$, modeled by a Gaussian stochastic process (GaSP) or Gaussian process (GP), $z(\cdot) \sim \mathcal{GP}(\mu(\cdot), \sigma^2 K(\cdot, \cdot))$, meaning that, for any inputs $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ (with \mathbf{x}_i being a $p \times 1$ vector), the marginal distribution of $\mathbf{z} = (z(\mathbf{x}_1), \dots, z(\mathbf{x}_N))^T$ follows a multivariate normal distribution,

$$\mathbf{z} \mid \boldsymbol{\beta}, \, \sigma^2, \, \boldsymbol{\gamma} \sim \mathcal{MN}(\boldsymbol{\mu}, \sigma^2 \mathbf{R}),$$
 (2.1)

where $\boldsymbol{\mu} = (\mu(\mathbf{x}_1), \dots, \mu(\mathbf{x}_N))^T$ is a vector of mean or trend parameters, σ^2 is the unknown variance and **R** is the correlation matrix with the (i, j) element modeled by a kernel $K(\mathbf{x}_i, \mathbf{x}_i)$ with parameters γ . It is common to model the mean by $\mu(\mathbf{x}) = \mathbf{h}(\mathbf{x})\boldsymbol{\beta}$, where $\mathbf{h}(\mathbf{x})$ is a $1 \times q$ row vector of basis function, and β is a $q \times 1$ vector of mean parameters.

When modeling spatially correlated data, the isotropic kernel is often used, where the input of the kernel only depends on the Euclidean distance $K(\mathbf{x}_a, \mathbf{x}_b) = K(||\mathbf{x}_a - \mathbf{x}_b||)$. In comparison, each coordinate of the latent function in computer experiments could have different physical meanings and units. Thus a product kernel is often used in constructing a GP emulator, such that correlation lengths can be different at each coordinate. For any $\mathbf{x}_a = (x_{a1}, \dots, x_{ap})$ and $\mathbf{x}_b = (x_{b1}, \dots, x_{bp})$, the kernel function can be written

as $K(\mathbf{x}_a, \mathbf{x}_b) = K_1(x_{a1}, x_{b1}) \times \cdots \times K_p(x_{ap}, x_{bp})$, where K_l is a kernel for the lth coordinate with a distinct range parameter γ_l , for $l = 1, \ldots, p$. Some frequently used kernels K_l include power exponential and Matérn kernel functions [44]. The Matérn kernel, for instance, follows

$$K_l(d_l) = \frac{1}{2^{\nu_l - 1} \Gamma(\nu_l)} \left(\frac{\sqrt{2\nu_l} d_l}{\gamma_l} \right)^{\nu_l} \mathcal{K}_{\nu_l} \left(\frac{\sqrt{2\nu_l} d_l}{\gamma_l} \right), \quad (2.2)$$

where $d_l = |x_{al} - x_{bl}|$, $\Gamma(\cdot)$ is the gamma function, $\mathcal{K}_{\nu_l}(\cdot/\gamma_l)$ is the modified Bessel function of the second kind with the range parameter and roughness parameter being γ_l and ν_l , respectively. The Matérn correlation has a closed-form expression when the roughness parameter is a half-integer, i.e. $\nu_l = 2k_l + 1/2$ with $k_l \in \mathbb{N}$. It becomes the exponential correlation and Gaussian correlation, when $k_l = 0$ and $k_l \to \infty$, respectively. The GP with Matérn kernel is $\lfloor \nu_l - 1 \rfloor$ mean square differentiable at coordinate l. This is a good property, as the differentiability of the process is directly controlled by the roughness parameter.

Denote mean basis of observations $\mathbf{H} = (\mathbf{h}^T(\mathbf{x}_1), \dots, \mathbf{h}^T(\mathbf{x}_N))^T$. The parameters in GP contain mean parameters $\boldsymbol{\beta}$, variance parameter σ^2 , and range parameters $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)$. Integrating out the mean and variance parameters with respect to reference prior $\pi(\boldsymbol{\beta}, \sigma) \propto 1/\sigma^2$, the predictive distribution of any input \mathbf{x}^* follows a student t distribution [20]:

$$z(\mathbf{x}^*) \mid \mathbf{z}, \, \boldsymbol{\gamma} \sim \mathcal{T}(\hat{z}(\mathbf{x}^*), \hat{\sigma}^2 K^{**}, N - q),$$
 (2.3)

with N-q degrees of freedom, where

$$\hat{z}(\mathbf{x}^*) = \mathbf{h}(\mathbf{x}^*)\hat{\boldsymbol{\beta}} + \mathbf{r}^T(\mathbf{x}^*)\mathbf{R}^{-1}\left(\mathbf{z} - \mathbf{H}\hat{\boldsymbol{\beta}}\right), \tag{2.4}$$

$$\hat{\sigma}^2 = (N - q)^{-1} \left(\mathbf{z} - \mathbf{H} \hat{\boldsymbol{\beta}} \right)^T \mathbf{R}^{-1} \left(\mathbf{z} - \mathbf{H} \hat{\boldsymbol{\beta}} \right), \qquad (2.5)$$

$$K^{**} = K(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{r}^T(\mathbf{x}^*) \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}^*) + \mathbf{h}^*(\mathbf{x}^*)^T$$
$$\times (\mathbf{H}^T \mathbf{R}^{-1} \mathbf{H})^{-1} \mathbf{h}^*(\mathbf{x}^*), \tag{2.6}$$

with $\hat{\boldsymbol{\beta}} = (\mathbf{H}^T \mathbf{R}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{R}^{-1} \mathbf{z}$ being the generalized least squares estimator of $\boldsymbol{\beta}$, $\mathbf{r}(\mathbf{x}^*) = (K(\mathbf{x}^*, \mathbf{x}_1), \dots, K(\mathbf{x}^*, \mathbf{x}_N))^T$ and $\mathbf{h}^*(\mathbf{x}^*) = (\mathbf{h}(\mathbf{x}^*) - \mathbf{H}^T \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}^*))$.

Direct computation of the likelihood requires $\mathcal{O}(N^3)$ operations due to computing the Cholesky decomposition of the covariane matrix for matrix inversion, and the determinant of the covariance matrix. Thus a posterior sampling algorithm, such as the Markov chain Monte Carlo (MCMC) algorithm can be slow, as it requires a large number of posterior samples. Plug-in estimators, such as the maximum likelihood estimator (MLE) were often used to estimate the range parameters γ in covariance. In [20], the maximum marginal posterior estimator (MMPE) with robust parameterizations was discussed to overcome the instability of the MLE. The MLE and MMPE of the parameters in a GP emulator with both the product kernel and the isotropic kernel are all implemented in the RobustGaSP package [19].




Figure 1: Predictions by the GP emulator of a function on 2D inputs with N=12 and N=24 observations (black circles) are shown in the left and right panels, respectively.

In some applications, we may not directly observe the latent function but a noisy realization:

$$y(\mathbf{x}) = z(\mathbf{x}) + \epsilon(\mathbf{x}),\tag{2.7}$$

where z(.) is modeled as a zero-mean GP with covariance $\sigma^2 K(.,.)$, and $\epsilon(\mathbf{x}) \sim \mathcal{N}(0,\sigma_0^2)$ follows an independent Gaussian noise. This model is typically referred to as the Gaussian process regression [44], which is suitable for scenarios containing noisy observations, such as experimental or field observations, numerical solutions of differential equations with non-negligible error, and stochastic simulations. Denote the noisy observations $\mathbf{y} = (y(\mathbf{x}_1), y(\mathbf{x}_2), \dots, y(\mathbf{x}_N))^T$ at the design input set $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ and the nugget parameter $\eta = \sigma_0^2/\sigma^2$. Both range and nugget parameters in GPR can be estimated by the plug-in estimators [19]. The predictive distribution of $f(\mathbf{x}^*)$ at any input \mathbf{x}^* can be obtained by replacing \mathbf{R} with $\tilde{\mathbf{R}} = \mathbf{R} + \eta \mathbf{I}_n$ in Equation (2.3).

Constructing a GP emulator to approximate computer simulation typically starts with a "space-filling" design, such as the Latin hypercube sampling (LHS), to fill the input space. Numerical solutions of computer models were then obtained at these design points, and the set $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ is used for training a GP emulator. For any observed input \mathbf{x}^* , the predictive mean in (2.3) is often used for predictions, and the uncertainty of observations can be obtained through the predictive distribution. In Figure 1, we plot the predictive mean of a GP emulator to approximate the Branin function [56] with N training inputs sampled from LHS, using the default setting of the RobustGaSP package [19]. When the number of observations increases from N=12 (middle panel) to N=24 (right panel), the predictive error becomes smaller.

The computational complexity of GP models increases at the order of $\mathcal{O}(N^3)$, which prohibits applications on emulating complex computer simulations, when a relatively large number of simulation runs are required. In the next section, we will introduce the state space representation of GP with Matérn covariance and one-dimensional input, where the computational order scales as $\mathcal{O}(N)$ without approximation. This method can be applied to problems with high dimensional input space discussed in Section 4.

3. MARGINALIZATION IN KALMAN FILTER

3.1 State Space Representation of GP with the Matérn Kernel

Suppose we model the observations by Equation (2.7) where the latent process z(.) is assumed to follow a GP on 1D input. For simplicity, here we assume a zero mean parameter ($\mu = 0$), and a mean function can be easily included in the analysis. It has been realized that a GP defined in 1D input space using a Matérn covariance with a half-integer roughness parameter input can be written as stochastic differential equations (SDEs) [63, 22], which can reduce the operations of computing the likelihood and making predictions from $\mathcal{O}(N^3)$ to $\mathcal{O}(N)$ operations, with the use of Kalman filter. Here we first review SDE representation and then we discuss marginalization of latent variables in the Kalman filter algorithm for scalable computation.

When the roughness parameter is $\nu = 5/2$, for instance, the Matérn kernel has the expression below

$$K(d) = \left(1 + \frac{\sqrt{5}d}{\gamma} + \frac{5d^2}{3\gamma^2}\right) \exp\left(-\frac{\sqrt{5}d}{\gamma}\right),\tag{3.1}$$

where $d = |x_a - x_b|$ is the distance between any $x_a, x_b \in \mathbb{R}$ and γ is a range parameter typically estimated by data. The output and two derivatives of the GP with the Matérn kernel in (3.1) can be written as the SDE below,

$$\frac{d\theta(x)}{dx} = \mathbf{J}\theta(x) + \mathbf{L}\epsilon(x), \tag{3.2}$$

or in the matrix form,

$$\frac{d}{dx} \begin{pmatrix} z(x) \\ z^{(1)}(x) \\ z^{(2)}(x) \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ -\lambda^3 & -3\lambda^2 & -3\lambda \end{pmatrix} \begin{pmatrix} z(x) \\ z^{(1)}(x) \\ z^{(2)}(x) \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \epsilon(x),$$

where $\epsilon(x) \sim N(0, \sigma^2)$, with $\lambda = \frac{\sqrt{2\nu}}{\gamma} = \frac{\sqrt{5}}{\gamma}$, and $z^{(l)}(\cdot)$ is the lth derivative of the process $z(\cdot)$. Denote $c = \frac{16}{3}\sigma^2\lambda^5$ and $\mathbf{F} = (1, 0, 0)$. Assume the 1D input is ordered, i.e. $x_1 \leq x_2 \leq \cdots \leq x_N$. The solution of SDE in (3.2) can be expressed as a continuous-time dynamic linear model [61],

$$y(x_i) = \mathbf{F}\boldsymbol{\theta}(x_i) + \epsilon,$$

$$\boldsymbol{\theta}(x_i) = \mathbf{G}(x_i)\boldsymbol{\theta}(x_{i-1}) + \mathbf{w}(x_i),$$
(3.3)

where $\mathbf{w}(x_i) \sim \mathcal{MN}(0, \mathbf{W}(x_i))$ for i = 2, ..., N, and the initial state follows $\boldsymbol{\theta}(x_1) \sim \mathcal{MN}(\mathbf{0}, \mathbf{W}(x_1))$. Here $\mathbf{G}(x_i) = e^{\mathbf{J}(x_i - x_{i-1})}$ and $\mathbf{W}(x_i) = \int_0^{x_i - x_{i-1}} e^{\mathbf{J}t} \mathbf{L} c \mathbf{L}^T e^{\mathbf{J}^T t} dt$ from i = 2, ..., N, and stationary distribution $\boldsymbol{\theta}(x_i) \sim \mathcal{MN}(0, \mathbf{W}(x_1))$, with $\mathbf{W}(x_1) = \int_0^{\infty} e^{\mathbf{J}t} \mathbf{L} c \mathbf{L}^T e^{\mathbf{J}^T t} dt$. Both $\mathbf{G}(x_i)$ and $\mathbf{W}(x_i)$ have closed-form expressions given in the

Appendix A.1. The joint distribution of the states follows $(\boldsymbol{\theta}^T(x_1), \dots, \boldsymbol{\theta}^T(x_N))^T \sim \mathcal{MN}(\mathbf{0}, \boldsymbol{\Lambda}^{-1})$, where the inverse covariance $\boldsymbol{\Lambda}$ is a block tri-diagonal matrix discussed in Appendix A.1.

3.2 Kalman Filter as a Scalable Marginalization Technique

For dynamic linear models in (3.3), Kalman filter and Rauch–Tung–Striebel (RTS) smoother can be used as an exact and scalable approach to compute the likelihood, and predictive distributions. The Kalman filter and RTS smoother are sometimes called the forward filtering and backward smoothing/sampling algorithm, widely used in dynamic linear models of time series. We refer the readers to [61, 41] for discussion of dynamic linear models.

Write $\mathbf{G}(x_i) = \mathbf{G}_i$, $\mathbf{W}(x_i) = \mathbf{W}_i$, $\boldsymbol{\theta}(x_i) = \boldsymbol{\theta}_i$ and $y(x_i) = y_i$ for i = 1, ..., N. In Lemma 1, we summarize Kalman filter and RTS smoother for the dynamic linear model in (3.3). Compared with $\mathcal{O}(N^3)$ computational operations and $\mathcal{O}(N^2)$ storage cost from GPs, the outcomes of Kalman filter and RTS smoother can be used for computing the likelihood and predictive distribution with $\mathcal{O}(N)$ operations and $\mathcal{O}(N)$ storage cost, summarized in Lemma 1. All the distributions in Lemma 1 and Lemma 2 are conditional distributions given parameters $(\gamma, \sigma^2, \sigma_0^2)$, which are omitted for simplicity.

Lemma 1 (Kalman Filter and RTS Smoother [26, 45]). 1. (Kalman Filter). Let $\theta_{i-1}|\mathbf{y}_{1:i-1} \sim \mathcal{MN}(\mathbf{m}_{i-1}, \mathbf{C}_{i-1})$. For i = 2, ..., N, iteratively we have,

(i) the one-step-ahead predictive distribution of $\boldsymbol{\theta}_i$ given $\mathbf{v}_{1:i-1}$.

$$\theta_i | \mathbf{y}_{1:i-1} \sim \mathcal{MN}(\mathbf{b}_i, \mathbf{B}_i),$$
 (3.4)

with $\mathbf{b}_i = \mathbf{G}_i \mathbf{m}_{i-1}$ and $\mathbf{B}_i = \mathbf{G}_i \mathbf{C}_{i-1} \mathbf{G}_i^T + \mathbf{W}_i$,

(ii) the one-step-ahead predictive distribution of Y_i given $\mathbf{y}_{1:i-1}$,

$$Y_i|\mathbf{y}_{1:i-1} \sim \mathcal{N}(f_i, Q_i), \tag{3.5}$$

with $f_i = \mathbf{F}\mathbf{b}_i$, and $Q_i = \mathbf{F}\mathbf{B}_i\mathbf{F}^T + \sigma_0^2$, (iii) the filtering distribution of $\boldsymbol{\theta}_i$ given $\mathbf{y}_{1:i}$,

$$\theta_i | \mathbf{y}_{1:i} \sim \mathcal{MN}(\mathbf{m}_i, \mathbf{C}_i),$$
 (3.6)

with
$$\mathbf{m}_i = \mathbf{b}_i + \mathbf{B}_i \mathbf{F}^T Q_i^{-1} (y_i - f_i)$$
 and $\mathbf{C}_i = \mathbf{B}_i - \mathbf{B}_i \mathbf{F}^T Q_i^{-1} \mathbf{F} \mathbf{B}_i$.

2. (RTS Smoother). Denote $\theta_{i+1}|\mathbf{y}_{1:n} \sim \mathcal{N}(s_{i+1}, S_{i+1})$, then recursively for $i = N - 1, \ldots, 1$,

$$\theta_i | \mathbf{y}_{1:N} \sim \mathcal{MN}(\mathbf{s}_i, \mathbf{S}_i),$$
 (3.7)

where
$$\mathbf{s}_i = \mathbf{m}_i + \mathbf{C}_i \mathbf{G}_{i+1}^T \mathbf{B}_{i+1}^{-1} (\mathbf{s}_{i+1} - \mathbf{b}_{i+1})$$
 and $\mathbf{S}_i = \mathbf{C}_i - \mathbf{C}_i \mathbf{G}_{i+1}^T \mathbf{B}_{i+1}^{-1} (\mathbf{B}_{i+1} - \mathbf{S}_{i+1}) \mathbf{B}_{i+1}^{-1} \mathbf{G}_{i+1} \mathbf{C}_i$.

Lemma 2 (Likelihood and predictive distribution). 1. (*Likelihood*). The likelihood follows

$$p(\mathbf{y}_{1:N} \mid \sigma^2, \sigma_0^2, \gamma) = \prod_{i=1}^{N} (2\pi Q_i)^{-\frac{1}{2}} \exp\left\{-\sum_{i=1}^{N} \frac{(y_i - f_i)^2}{2Q_i}\right\},\,$$

where f_i and Q_i are given in Kalman filter. The likelihood can be used to obtain the MLE of the parameters $(\sigma^2, \sigma_0^2, \gamma)$. 2. (**Predictive distribution**).

(i) By the last step of Kalman filter, one has $\boldsymbol{\theta}_N | \mathbf{y}_{1:N}$ and recursively by the RTS smoother; the predictive distribution of $\boldsymbol{\theta}_i$ for i = N - 1, ..., 1 follows

$$\theta_i | \mathbf{y}_{1:N} \sim \mathcal{MN}(\mathbf{s}_i, \mathbf{S}_i).$$
 (3.8)

(ii) For any x^* (W.l.o.g. let $x_i < x^* < x_{i+1}$)

$$\boldsymbol{\theta}(x^*) \mid \mathbf{y}_{1:N} \sim \mathcal{MN}\left(\hat{\boldsymbol{\theta}}(x^*), \hat{\boldsymbol{\Sigma}}(x^*)\right)$$

where $\hat{\boldsymbol{\theta}}(x^*) = \mathbf{G}_i^* \mathbf{s}_i + \mathbf{W}_i^* (\mathbf{G}_{i+1}^*)^T (\tilde{\mathbf{W}}_{i+1}^*)^{-1} (\mathbf{s}_{i+1} - \mathbf{G}_{i+1}^* \mathbf{G}_i^* \mathbf{s}_i)$ and $\hat{\boldsymbol{\Sigma}}(x^*) = ((\mathbf{W}_i^*)^{-1} + (\mathbf{G}_{i+1}^*)^T (\mathbf{W}_{i+1}^*)^{-1} \mathbf{G}_{i+1}^*)^{-1}$ with terms denoted with '*' given in the Appendix A.1.

Although we introduce the Matérn kernel with $\nu=5/2$ as an example, the likelihood and predictive distribution of GPs with the Matérn kernel of a small half-integer roughness parameter can be computed efficiently, for both equally spaced and not equally spaced 1D inputs. For the Matérn kernel with a very large roughness parameter, the dimension of the latent states becomes large, which makes efficient computation prohibitive. In practice, the Matérn kernel with a relatively large roughness parameter (e.g. with $\nu=5/2$) is found to be accurate for estimating a smooth latent function in computer experiments [20, 2]. Because of this reason, the Matérn kernel with $\nu=5/2$ is the default choice of the kernel function in some packages of GP emulators [47, 19].

For a model containing latent variables, one may proceed with two usual approaches:

- (i) sampling the latent variables $\theta(x_i)$ from the posterior distribution by the MCMC algorithm,
- (ii) optimizing the latent variables $\theta(x_i)$ to minimize a loss function.

For approach (i), the MCMC algorithm is usually much slower than the Kalman filter, as the number of the latent states is high, requiring a large number of posterior samples [17]. On the other hand, the prior correlation between states may not be taken into account directly in approach (ii), making the estimation less efficient than the Kalman filter, if data contain correlation across latent states. In comparison, the latent states in the dynamic linear model in (3.3) are iteratively marginalized out in Kalman filter, and the closed-form expression is derived in each step, which only takes $\mathcal{O}(N)$ operations and storage cost, with N being the number of observations.

In practice, when a sensible probability model or a prior of latent variables is considered, the principle is to integrate out the latent variables when making predictions. Posterior samples and optimization algorithms, on the other hand, can be very useful for approximating the marginal likelihood when closed-form expressions are not available. As




Figure 2: Comparison between the direct computation of GP, and that by the Kalman filter (KF) and RTS smoother, both having the Matérn kernel in (3.1). The left panel shows the computational cost of computing the predictive mean over different number of observations. When N=5000, direct computation takes around 20 seconds, whereas the computation by KF and RTS smoother takes 0.029 seconds. The predictive mean computed in both ways is plotted in the right panel when N=1,000, where the root of mean squared difference between two approaches is 5.98×10^{-12} .

an example, we will introduce applications that integrate the sparse covariance structure along with conjugate gradient optimization into estimating particle interaction kernels, and forecasting particle trajectories in Section 4, which integrates both marginalization and optimization to tackle a computationally challenging scenario.

In Figure 2, we compare the cost for computing the predictive mean for a nonlinear function with 1D inputs [16]. The input is uniformly sampled from [0.5, 0.25], and an independent Gaussian white noise with a standard deviation of 0.1 is added in simulating the observations. We compare two ways of computing the predictive mean. The first approach implements direct computation of the predictive mean by Equation (2.4). The second approach is computed by the likelihood function and predictive distribution from Lemma 2 based on the Kalman filter and RTS smoother. The range and nugget parameters are fixed to be 0.5 and 10^{-4} for demonstration purposes, respectively. The computational time of this simulated experiment is shown in the left panel in Figure 2. The approach based on Kalman filter and RTS smoother is much faster, as computing the likelihood and making predictions by Kalman filter and RTS smoother only require $\mathcal{O}(N)$ operations, whereas the direct computation cost $\mathcal{O}(N^3)$ operations. The right panel gives the predictive mean, latent truth, and observations, when N = 1000. The difference between the two approaches is very small, as both methods are exact.

3.3 Marginalization of Correlated Matrix Observations with Multi-dimensional Inputs

The Kalman filter is widely applied in signal processing, system control, and modeling time series. Here we introduce a few recent studies that apply Kalman filter to GP models with Matérn covariance to model spatial, spatio-temporal, and functional observations.

Let $\mathbf{y}(\mathbf{x}) = (y_1(\mathbf{x}), \dots, y_{n_1}(\mathbf{x}))^T$ be an n_1 -dimensional real-valued output vector at a p-dimensional input vector \mathbf{x} . For simplicity, assume the mean is zero. Consider the latent factor model:

$$\mathbf{y}(\mathbf{x}) = \mathbf{A}\mathbf{z}(\mathbf{x}) + \boldsymbol{\epsilon}(\mathbf{x}), \tag{3.9}$$

where $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_d]$ is an $n_1 \times d$ factor loading matrix and $\mathbf{z}(\mathbf{x}) = (z_1(\mathbf{x}), \dots, z_d(\mathbf{x}))^T$ is a d-dimensional factor processes, with $d \leq n_1$. The noise process follows $\boldsymbol{\epsilon}(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \mathbf{I}_{n_1})$. Each factor is modeled by a zero-mean Gaussian process (GP), meaning that $\mathbf{Z}_l = (z_l(\mathbf{x}_1), \dots, z_l(\mathbf{x}_{n_2}))$ follows a multivariate normal distribution $\mathbf{Z}_l^T \sim \mathcal{M}\mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_l)$, where the (i, j) entry of $\mathbf{\Sigma}_l$ is parameterized by a covariance function $\sigma_l^2 K_l(\mathbf{x}_i, \mathbf{x}_j)$ for $l = 1, \dots, d$. The model (3.9) is often known as the semiparametric latent factor model in the machine learning community [53], and it belongs to a class of linear models of coregionalization [3]. It has a wide range of applications in modeling multivariate spatially correlated data and functional observations from computer experiments [14, 25, 40].

We have the following two assumptions for model (3.9).

Assumption 1. The prior of latent processes $\mathbf{z}_i(.)$ and $\mathbf{z}_j(.)$ are independent, for any $i \neq j$.

Assumption 2. The factor loadings are orthogonal, i.e. $\mathbf{A}^T \mathbf{A} = \mathbf{I}_d$.

The first assumption is typically assumed for modeling multivariate spatially correlated data or computer experiments [3, 25]. Secondly, note that the model in (3.9) is unchanged if we replace $(\mathbf{A}, \mathbf{z}(\mathbf{x}))$ by $(\mathbf{AE}, \mathbf{E^{-1}z}(\mathbf{x}))$ for an invertible matrix \mathbf{E} , meaning that the linear subspace of \mathbf{A} can be identified if no further constraint is imposed. Furthermore, as the variance of each latent process σ_i^2 is estimated by the data, imposing the unity constraint on each column of \mathbf{A} can reduce identifiability issues. The second assumption was also assumed in other recent studies [31, 30].

Given Assumption 1 and Assumption 2, we review recent results that alleviates the computational cost. Let us first assume the observations are an $N = n_1 \times n_2$ matrix $\mathbf{Y} = [\mathbf{y}(\mathbf{x}_1), \dots, \mathbf{y}(\mathbf{x}_{n_2})]$.

Lemma 3 (Posterior independence and orthogonal projection [17]). For model (3.9) with Assumption 1 and Assumption 2, we have two properties below.

1. (Posterior Independence). For any $l \neq m$

$$Cov[\mathbf{Z}_l^T, \mathbf{Z}_m^T | \mathbf{Y}] = \mathbf{0},$$

and for each $l = 1, \ldots, d$,

$$\mathbf{Z}_{l}^{T}|\mathbf{Y} \sim \mathcal{MN}(\boldsymbol{\mu}_{Z_{l}}, \boldsymbol{\Sigma}_{z_{l}}),$$

where $\boldsymbol{\mu}_{z_l} = \boldsymbol{\Sigma}_l \tilde{\boldsymbol{\Sigma}}_l^{-1} \tilde{\mathbf{y}}_l$, $\tilde{\mathbf{y}}_l = \mathbf{Y}^T \mathbf{a}_l$ and $\boldsymbol{\Sigma}_{Z_l} = \boldsymbol{\Sigma}_l - \boldsymbol{\Sigma}_l \tilde{\boldsymbol{\Sigma}}_l^{-1} \boldsymbol{\Sigma}_l$ with $\tilde{\boldsymbol{\Sigma}}_l = \boldsymbol{\Sigma}_l + \sigma_0^2 \mathbf{I}_{n_2}$.

2. (Orthogonal Projection). After integrating $z(\cdot)$, the marginal likelihood is a product of multivariate normal densities at projected observations:

$$p(\mathbf{Y}) = \prod_{l=1}^{d} \mathcal{PN}(\tilde{\mathbf{y}}_{l}; \mathbf{0}, \tilde{\mathbf{\Sigma}}_{l}) \prod_{l=d+1}^{n_{1}} \mathcal{PN}(\tilde{\mathbf{y}}_{c,l}; \mathbf{0}, \sigma_{0}^{2} \mathbf{I}_{n_{2}}), \quad (3.10)$$

where $\tilde{\mathbf{y}}_{c,l} = \mathbf{Y}^T \mathbf{a}_{c,l}$ with $\mathbf{a}_{c,l}$ being the lth column of \mathbf{A}_c , the orthogonal component of \mathbf{A} , and \mathcal{PN} denotes the density for a multivariate normal distribution.

The properties in Lemma 3.9 lead to computationally scalable expressions of the maximum marginal likelihood estimator (MMLE) of factor loadings.

Theorem 1 (Generalized probabilistic principal component analysis [18]). Assume $\mathbf{A}^T \mathbf{A} = \mathbf{I}_d$, after marginalizing out $\mathbf{Z}_l^T \sim \mathcal{MN}(\mathbf{0}, \mathbf{\Sigma}_l)$ for l = 1, 2, ..., d, we have the results below.

• If $\Sigma_1 = \cdots = \Sigma_d = \Sigma$, the marginal likelihood is maximized when

$$\hat{\mathbf{A}} = \mathbf{US},\tag{3.11}$$

where **U** is an $n_1 \times d$ matrix of the first d principal eigenvectors of $\mathbf{G} = \mathbf{Y}(\sigma_0^2 \mathbf{\Sigma}^{-1} + \mathbf{I}_{n_2})^{-1} \mathbf{Y}^T$ and **S** is a $d \times d$ orthogonal rotation matrix;

• If the covariances of the factor processes are different, denoting $\mathbf{G}_l = \mathbf{Y}(\sigma_0^2 \mathbf{\Sigma}_l^{-1} + \mathbf{I}_{n_2})^{-1} \mathbf{Y}^T$, the MMLE of factor loadings is

$$\hat{\mathbf{A}} = \operatorname{argmax}_{\mathbf{A}} \sum_{l=1}^{d} \mathbf{a}_{l}^{T} \mathbf{G}_{l} \mathbf{a}_{l}, \quad s.t. \quad \mathbf{A}^{T} \mathbf{A} = \mathbf{I}_{d}.$$
 (3.12)

The estimator **A** in Theorem 1 is called the *generalized* probabilistic principal component analysis (GPPCA). The optimization algorithm that preserves the orthogonal constraints in (3.12) is available in [60].

In [58], the latent factor is assumed to follow independent standard normal distributions, and the authors derived the MMLE of the factor loading matrix \mathbf{A} , which was termed the probabilistic principal component analysis (PPCA). The GPPCA extends the PPCA to correlated latent factors modeled by GPs, which incorporates the prior correlation information between outputs as a function of inputs, and the latent factor processes were marginalized out when estimating the factor loading matrix and other parameters. When the input is 1D and the Matérn covariance is used for modeling latent factors, the computational order of GPPCA is $\mathcal{O}(Nd)$, which is the same as the PCA. For correlated data,

such as spatio-temporal observations and multivariate functional data, GPPCA provides a flexible and scalable approach to estimate factor loading by marginalizing out the latent factors [18].

Spatial and spatio-temporal models with a separable covariance can be written as a special case of the model in (3.9). For instance, suppose $d = n_1$ and the $n_1 \times n_2$ latent factor matrix follows $\mathbf{Z} \sim \mathcal{MN}(0, \sigma^2 \mathbf{R}_1 \otimes \mathbf{R}_2)$, where \mathbf{R}_1 and \mathbf{R}_2 are $n_1 \times n_1$ and $n_2 \times n_2$ subcovariances, respectively. Denote the eigen decomposition $\mathbf{R}_1 = \mathbf{V}_1 \mathbf{\Lambda}_1 \mathbf{V}_1^T$ with Λ_1 being a diagonal matrix with the eigenvalues λ_i , for $i = 1, ..., n_1$. Then this separable model can be written as the model in (3.9), with $\mathbf{A} = \mathbf{V}_1$, $\Sigma_l = \sigma^2 \lambda_l \mathbf{R}_2$. The connection suggests that the latent factor loading matrix can be specified as the eigenvector matrix of a covariance matrix, parameterized by a kernel function. This approach is studied in [17] for modeling incomplete lattice with irregular missing patterns, and the Kalman filter is integrated for accelerating computation on massive spatial and spatiotemporal observations.

4. SCALABLE MARGINALIZATION FOR LEARNING PARTICLE INTERACTION KERNELS FROM TRAJECTORY DATA

Collective motions with particle interactions are very common in both microscopic and macroscopic systems [35, 36]. Learning interaction kernels between particles is important for two purposes. First, physical laws are less understood for many complex systems, such as cell migration or non-equilibrium thermodynamic processes. Estimating the interaction kernels between particles from fields or experimental data is essential for learning these processes. Second, simulation of particle interactions, such as *ab initio* molecular dynamics simulation, can be very computationally expensive. Statistical learning approaches can be used for reducing the computational cost of simulations.

For demonstration purposes, we consider a simple first-order system. In [34], for a system with n interacting particles, the velocity of the ith particle at time t, $\mathbf{v}_i(t) = d\mathbf{x}_i(t)/dt$, is modeled by positions between all particles,

$$\mathbf{v}_i(t) = \sum_{j=1}^n \phi(||\mathbf{x}_j(t) - \mathbf{x}_i(t)||) \mathbf{u}_{i,j}(t), \tag{4.1}$$

where $\phi(\cdot)$ is a latent interaction kernel function between particle i and all other particles, with $||\cdot||$ being the Euclidean distance, and $\mathbf{u}_{i,j}(t) = \mathbf{x}_j(t) - \mathbf{x}_i(t)$ is a vector of differences between positions of particles i and j, for $i, j = 1, \ldots, n$. Here $\phi(\cdot)$ is a two-body interaction kernel. In molecular dynamics simulation, for instance, the Lennard-Jones potential can be related to molecular forces and the acceleration of molecules in a second-order system. The statistical learning approach can be extended to a second-order system that involves acceleration and external force terms.

The first-order system as (4.1) can be considered as an approximation of the second-order system. Furthermore, the interaction between particles is global, as any particle is affected by all other particles. Learning global interactions is more computationally challenging than local interactions [51], and approximating the global interaction by the local interaction is of interest in future studies.

One important goal is to efficiently estimate the unobservable interaction functions from the particle trajectory data, without specifying a parametric form. This goal is key for estimating the behaviors of dynamic systems in experiments and in observational studies, as the physics law in a new system may be unknown. In [13], $\phi(\cdot)$ is modeled by a zero-mean Gaussian process with a Matérn covariance:

$$\phi(\cdot) \sim \mathcal{GP}(0, \sigma^2 K(\cdot, \cdot)).$$
 (4.2)

Computing estimation of interactions of large-scale systems or more simulation runs, however, is prohibitive, as the direct inversion of the covariance matrix of observations of velocities requires $\mathcal{O}((nMDL)^3)$ operations, where M is the number of simulations or experiments, n is the number of particles, D is the dimension of each particle, L denotes the number of time points for each trial. Furthermore, constructing such covariance contains computing an $n^2LM \times n^2LM$ matrix of ϕ for a D-dimensional input space, which takes $\mathcal{O}(n^4L^2M^2D)$ operations. Thus, directly estimating interaction kernel with a GP model in (4.2) is only applicable to systems with a small number of observations [34, 13].

This work makes contributions from two different aspects for estimating dynamic systems of interacting particles. We first show the covariance of velocity observations can be obtained by operations on a few sparse matrices, after marginalizing out the latent interaction function. The sparsity of the inverse covariance of the latent interaction kernel allows us to modify the Kalman filter to efficiently compute the matrix product in this problem, and then apply a conjugate gradient (CG) algorithm [24, 21, 49] to solve this system iteratively. The computational complexity of computing the predictive mean and variance of a test point is at the order of $\mathcal{O}(TnN) + \mathcal{O}(nN\log(nN))$, for a system of n particles, N = nMDL observations, and T is the number of iterations required in the CG algorithm. We found that typically around a few hundred CG iterations can achieve high predictive accuracy for a moderately large number of observations. The algorithm leads substantial reduction of computational cost, compared to direct computation.

Second, we study the effect of experimental designs on estimating the interaction kernel function. In previous studies, it is unclear how initial positions, time steps of trajectory and the number of particles affect the accuracy in estimating interaction kernels. Compared to other conventional problems in computer model emulation, where a "space-filling" design is often used, here we cannot directly observe the

realizations of the latent function. Instead, the output velocity is a weighted average of the interaction kernel functions between particles. Besides, we cannot directly control distances between the particles moving away from the initial positions, in both simulation and experimental studies. When the distance between two particles i and j is small, the contribution $\phi(||\mathbf{x}_i(t) - \mathbf{x}_j(t)||)\mathbf{u}_{i,j}(t)$ can be small, if the repulsive force by $\phi(\cdot)$ does not increase as fast as the distance decreases. Thus we found that the estimation of interaction kernel function can be less accurate in the input domain of small distances. This problem can be alleviated by placing initial positions of more particles close to each other, providing more data with small distance pairs that improve the accuracy in estimation.

4.1 Scalable Computation by Sparse Representation of Inverse Covariance

For illustration purposes, let us first consider a simple scenario where we have M=1 simulation and L=1 time point of a system of n interacting particles at a D dimensional space. Since we only have 1 time point, we omit the notation t when there is no confusion. The algorithm for the general scenario with L>1 and M>1 is discussed in Appendix A.2. In practice, the experimental observations of velocity from multiple particle tracking algorithms or particle image velocimetry typically contain noises [1]. Even for simulation data, the numerical error could be non-negligible for large and complex systems. In these scenarios, the observed velocity $\tilde{\mathbf{v}}_i = (v_{i,1}, \ldots, v_{i,D})^T$ is a noisy realization: $\tilde{\mathbf{v}}_i = \mathbf{v}_i + \boldsymbol{\epsilon}_i$, where $\boldsymbol{\epsilon}_i \sim \mathcal{MN}(0, \sigma_0^2 \mathbf{I}_D)$ denotes a vector of Gaussian noise with variance σ_0^2 .

Assume the nD observations of velocity are $\tilde{\mathbf{v}} = (\tilde{v}_{1,1}, \dots, \tilde{v}_{n,1}, \tilde{v}_{1,2}, \dots, \tilde{v}_{n,2}, \dots, \tilde{v}_{n-1,D}, \tilde{v}_{n,D})^T$. After integrating out the latent function modeled in Equation (4.2), the marginal distribution of observations follows

$$(\tilde{\mathbf{v}} \mid \mathbf{R}_{\phi}, \sigma^2, \sigma_0^2) \sim \mathcal{MN} (\mathbf{0}, \sigma^2 \mathbf{U} \mathbf{R}_{\phi} \mathbf{U}^T + \sigma_0^2 \mathbf{I}_{nD}), \quad (4.3)$$

where **U** is an $nD \times n^2$ block diagonal matrix, with the ith $D \times n$ block in the diagonals being $(\mathbf{u}_{i,1}, \dots, \mathbf{u}_{i,n})$, and \mathbf{R}_{ϕ} is an $n^2 \times n^2$ matrix, where the (i', j') term in the (i, j)th $n \times n$ block is $K(d_{i,i'}, d_{j,j'})$ with $d_{i,i'} = ||\mathbf{x}_i - \mathbf{x}_i'||$ and $d_{j,j'} = ||\mathbf{x}_j - \mathbf{x}_j'||$ for $i, i', j, j' = 1, \dots, n$. Direct computation of the likelihood involves computing the inverse of an $nD \times nD$ covariance matrix and constructing an $n^2 \times n^2$ matrix \mathbf{R}_{ϕ} , which costs $\mathcal{O}((nD)^3) + \mathcal{O}(n^4D)$ operations. This is computationally expensive even for small systems.

Here we use an exponential kernel function, $K(d) = \exp(-d/\gamma)$ with range parameter γ , of modeling any nonnegative distance input d for illustration, where this method can be extended to include Matérn kernels with half-integer roughness parameters. Denote distance pairs $d_{i,j} = ||\mathbf{x}_i - \mathbf{x}_j||$, and there are (n-1)n/2 unique positive distance pairs. Denote the (n-1)n/2 distance pairs $\mathbf{d}_s = (d_{s,1}, \dots d_{s,(n-1)n/2})^T$ in an increasing order with the subscript s meaning 'sorted'. Here we do not need to consider

the case when $d_{i,j} = 0$, as $\mathbf{u}_{i,j} = \mathbf{0}$, leading to zero contribution to the velocity. Thus the model in (4.1) can be reduced to exclude the interaction between particle at zero distance. In reality, two particles at the same position are impractical, as there typically exists a repulsive force when two particles get very close. Hence, we can reduce the n^2 distance pairs $d_{i,j}$ for i = 1, ..., n and j = 1, ..., n, to (n-1)n/2 unique positive terms $d_{s,i}$ in an increasing order, for i = 1, ..., (n-1)n/2.

Denote the $(n-1)n/2 \times (n-1)n/2$ correlation matrix of the kernel outputs $\phi = (\phi(d_{s,1}), \dots, \phi(d_{s,(n-1)n/2}))^T$ by \mathbf{R}_s and \mathbf{U}_s is $nD \times (n-1)n/2$ sparse matrix with n-1 nonzero terms on each row, where the nonzero entries of the ith particle correspond to the distance pairs in the \mathbf{R}_s . Denote the nugget parameter $\eta = \sigma_0^2/\sigma^2$. After marginalizing out ϕ , the covariance of velocity observations follows

$$(\tilde{\mathbf{v}} \mid \gamma, \sigma^2, \eta) \sim \mathcal{MN} \left(\mathbf{0}, \sigma^2 \tilde{\mathbf{R}}_v \right),$$
 (4.4)

with

$$\tilde{\mathbf{R}}_v = (\mathbf{U}_s \mathbf{R}_s \mathbf{U}_s^T + \eta \mathbf{I}_{nD}). \tag{4.5}$$

The conditional distribution of the interaction kernel $\phi(d^*)$ at any distance d^* follows

$$(\phi(d^*) \mid \tilde{\mathbf{v}}, \gamma, \sigma^2, \eta) \sim \mathcal{N}(\hat{\phi}(d^*), \sigma^2 K^*), \tag{4.6}$$

where the predictive mean and variance follow

$$\hat{\phi}(d^*) = \mathbf{r}^T(d^*)\mathbf{U}_s^T \tilde{\mathbf{R}}_v^{-1} \tilde{\mathbf{v}}, \tag{4.7}$$

$$\sigma^2 K^* = \sigma^2 \left(K(d^*, d^*) - \mathbf{r}^T(d^*) \mathbf{U}_s^T \tilde{\mathbf{R}}_v^{-1} \mathbf{U}_s \mathbf{r}(d^*) \right), \quad (4.8)$$

with $\mathbf{r}(d^*) = (K(d^*, d_{s,1}), \dots, K(d^*, d_{s,n(n-1)/2}))^T$. After obtaining the estimated interaction kernel, one can use it to forecast trajectories of particles and understand the physical mechanism of flocking behaviors.

Our primary task is to efficiently compute the predictive distribution of interaction kernel in (4.6), where the most computationally expensive terms in the predictive mean and variance is $\tilde{\mathbf{R}}_v^{-1}\tilde{\mathbf{v}}$ and $\tilde{\mathbf{R}}_v^{-1}\mathbf{U}_s\mathbf{r}(d^*)$. Note that the \mathbf{U}_s is a sparse matrix with n(n-1)d nonzero terms and the inverse covariance matrix \mathbf{R}_{s}^{-1} is a tri-diagonal matrix. However, directly applying the CG algorithm is still computationally challenging, as neither $\tilde{\mathbf{R}}_v$ nor $\tilde{\mathbf{R}}_v^{-1}$ is sparse. To solve this problem, we extend a step in the Kalman filter to efficiently compute the matrix-vector multiplication with the use of sparsity induced by the choice of covariance matrix. Each step of the CG iteration in the new algorithm only costs $\mathcal{O}(nDT)$ operations for computing a system of n particles and D dimensions with T CG iteration steps. For most systems we explored, we found a few hundred iterations in the CG algorithm achieve high accuracy. The substantial reduction of the computational cost allows us to use more observations to improve the predictive accuracy. We term this approach the sparse conjugate gradient algorithm for




Figure 3: Estimation of particle interaction kernel by the truncated Lennard-Jones potential in [34] with D=2. The left panel shows the true interaction kernel (black curve) and estimated kernel functions (colored curves), based on three different ways of initial positions of n=200 particles. The computational time of the full GP in [13] and the sparse CG-GP approach is given in the right panel. Here the most computational intensive part of the full GP model is in constructing the correlation matrix \mathbf{R}_s of ϕ for n(n-1)/2 distance pairs in Equation (4.5), which scales as $\mathcal{O}(n^4D)$.

Gaussian processes (sparse CG-GP). The algorithm for the scenario with M simulations, each containing L time frames of n particles in a D dimensional space, is discussed in Appendix A.2.

The comparison of the computational cost between the full GP model and the proposed sparse CG-GP method is shown in the right panel in Figure 3. The most computational expensive part of the full GP model is on constructing the $n(n-1)/2 \times n(n-1)/2$ correlation matrix \mathbf{R}_s of ϕ for n(n-1)/2 distance pairs. The sparse CG-GP algorithm is much faster as we do not need to construct this covariance matrix; instead we only need to efficiently compute matrix multiplication by utilizing the sparse structure of the inverse of \mathbf{R}_s (Appendix A.2). Note the GP model with an exponential covariance naturally induces a sparse inverse covariance matrix that can be used for faster computation, which is different from imposing a sparse covariance structure for approximation.

In the left panel in Figure 3, we show the predictive mean and uncertainty assessment by the sparse CG-GP method for three different designs for sampling the initial positions of particles. From the first to the third designs, the initial value of each coordinate of the particle is sampled independently from a uniform distribution $\mathcal{U}[a_1, b_1]$, normal distribution $\mathcal{N}(a_2, b_2)$, and log uniform (reciprocal) distribution $\mathcal{L}\mathcal{U}[\log(a_3), \log(b_3)]$, respectively.

For experiments with the interaction kernel being the truncated Lennard-Jones potential given in Appendix A.2, we use $a_1 = 0$, $b_1 = 5$, $a_2 = 0$, $b_2 = 5$, $a_3 = 10^{-3}$ and $b_3 = 5$ for three designs of initial positions. Compared with the first design, the second design of initial positions, which was assumed in [34], has a larger probability mass of distributions near 0. In the third design, the distributions of the distance between particle pairs are monotonically decreasing, with

more probability mass near 0 than those in the first two designs. In all cases shown in Figure 3, we assume M=1, L=1 and the noise variance is set to be $\sigma_0=10^{-3}$ in the simulation. For demonstration purposes, the range and nugget parameters are fixed to be $\gamma = 5$ and $\eta = 10^{-5}$ respectively, when computing the predictive distribution of ϕ . The estimation of the interaction kernel on large distances is accurate for all different designs, whereas the estimation of the interaction kernel at small distances is not satisfying for the first two designs. When particles are initialized from the third design (log-uniform), the accuracy is better. as there are more particles near each other, providing more information about the particles at small values. This result is intuitive, as the small distance pairs have relatively small contributions to the velocity based on equation (4.1), and we need more particles close to each other to estimate the interaction kernel function at small distances.

The numerical comparison between different designs allows us to better understand the learning efficiency in different scenarios, which can be used to design experiments. Because of the large improvement of computational scalability compared to previous studies [13, 34], we can accurately estimate interaction kernels based on more particles and longer trajectories.

4.2 Numerical Results

Here we discuss two scenarios, where the interaction between particles follow the truncated Lennard-Jones (LJ) and opinion dynamics (OD) kernel functions. The LJ potential is widely used in MD simulations of interacting molecules [43]. First-order systems of form (4.1) have also been successfully applied in modeling opinion dynamics in social networks (see the survey [36] and references therein). The interaction function ϕ models how the opinions of pairs of people influence each other. In our numerical example, we consider heterophilious opinion interactions: each agent is more influenced by its neighbors slightly further away from its closest neighbors. As time evolves, the opinions of agents merge into clusters, with the number of clusters significantly smaller than the number of agents. This phenomenon was studied in [36] that heterophilious dynamics enhances consensus, contradicting the intuition that would suggest that the tendency to bond more with those who are different rather than with those who are similar would break connections and prevent clusters of consensus.

The details of the interaction functions are given in Appendix A.3. For each interaction, we test our method based on 12 configurations of 2 particle sizes (n=50 and n=200), 2 time lengths (L=1 and L=10), and 3 designs of initial positions (uniform, normal and log-uniform). The computational scalability of the sparse CG algorithm allows us to efficiently compute the predictions in most of these experimental settings within a few seconds. For each configuration, we repeat the experiments 10 times to average the effects of randomness in the initial positions of particles. The root of




Figure 4: Estimation of interaction function based on the sparse CG-GP method, and trajectory forecast for the truncated LJ and OD simulation. In panel (a), the colored curves are estimated interactions for different initial positions and particle sizes, all based on trajectories using only L=1 time step, whereas the black curve is the truncated LJ used in the simulation. The colored curves in panel (b) are the same as those in panel (a), but based on trajectories in L=10 time steps. The panels (d) and (e) are the same as panels (a) and (b), respectively, but the simulated data are generated by the OD interaction kernel. The shared areas are the 95% predictive interval. In panel (c), we graph the simulated trajectories of 10 out of 50 particles L=200 time steps, and the trajectory forecast based on estimated interaction function and initial positions. The arrow indicates the direction of velocities of particles at the last time step. Panel (f) is the same as panel (c), but for the OD interaction kernel.

the mean squared error in predicting the interaction kernels by averaging these 10 experiments of each configuration is given in Appendix A.4.

In Figure 4, we show the estimation of interactions kernels and forecasts of particle trajectories with different designs, particle sizes and time points. The sparse CG-GP method is relatively accurate for almost all scenarios. Among different initial positions, the estimation of trajectories for LJ interaction is the most accurate when the initial positions of the particles are sampled by the log-uniform distribution. This is because there are more small distances between particles when the initial positions follow a log-uniform distribution, providing more data to estimate the interaction kernel at small distances. Furthermore, when we have more particles or observations at larger time intervals, the estimation of the interaction kernel from all designs becomes more accurate in terms of the normalized root mean squared error with the detailed comparison given in Appendix A.4.

In panel (c) and panel (f) of Figure 4, we plot the tra-

jectory forecast of 10 particles over 200 time points for the truncated LJ kernel and OD kernel, respectively. In both simulation scenarios, interaction kernels are estimated based on trajectories of n = 50 particles across L = 20 time steps with initial positions sampled from the log-uniform design. The trajectories of only 10 particles out of 50 particles are shown for better visualization. For trajectories simulated by the truncated LJ, some particles can move very close, since the repulsive force between two particles becomes smaller as the force is proportional to the distance from Equation (4.1), and the truncation of kernel substantially reduces the repulsive force when particles move close. For the OD simulation, the particles move toward a cluster, as expected, since the particles always have attractive forces between each other. The forecast trajectories are close to the hold-out truth, indicating the high accuracy of our approach.

Compared with the results shown in previous studies [34, 13], estimating the interaction kernels and forecasting trajectories both look more accurate. The large computa-

tional reduction by the sparse CG-GP algorithm shown in Figure 3 permits the use of longer trajectories from more particles to estimate the interaction kernel, which improves the predictive accuracy. Here particle has interactions with all other particles in our simulation, making the number of distance pairs large. Yet we are able to estimate the interaction kernel and forecast the trajectories of particles within only tens of seconds in a desktop for the most time consuming scenario we considered. Since the particles typically have very small or no interaction when the distances between them are large, approximation can be made by enforcing interactions between particles within the specified radius, for further reducing the computational cost.

5. CONCLUDING REMARKS

We have introduced scalable marginalization of latent variables for correlated data. We first introduce GP models and reviewed the SDE representation of GPs with Matérn covariance and one-dimensional input. Kalman filter and RTS smoother were introduced as a scalable marginalization way to compute the likelihood function and predictive distribution, which reduces the computational complexity of GP with Matérn covariance for 1D input from $\mathcal{O}(N^3)$ to $\mathcal{O}(N)$ operations without approximation, where N is the number of observations. Recent efforts on extending scalable computation from 1D input to multi-dimensional input are discussed. In particular, we developed a new scalable algorithm for predicting particle interaction kernel and forecast trajectories of particles. The achievement is through the sparse representation of GPs in modeling interaction kernel, and then efficient computation for matrix multiplication by modifying the Kalman filter algorithm. An iterative algorithm based on CG can then be applied, which reduces the computational complexity.

There are a wide range of future topics relevant to this study. First, various models of spatio-temporal data can be written as random factor models in (3.9) with latent factors modeled as Gaussian processes for temporal inputs. It is of interest to utilize the computational advantage of the dynamic linear models of factor processes, extending the computational tools by relaxing the independence between prior factor processes in Assumption 1 or incorporating the Toeplitz covariance structure for stationary temporal processes. Second, for estimating systems of particle interactions, we can further reduce computation by only considering interactions within a radius between particles. Third, a comprehensively study the experimental design, initialization, and parameter estimation in will be helpful for estimating latent interaction functions that can be unidentifiable or weakly identifiable in certain scenarios. Furthermore, velocity directions and angle order parameters are essential for understanding the mechanism of active nematics and cell migration, which can motivate more complex models of interactions. Finally, the sparse CG algorithm developed in this work is of interest to reducing the computational complexity of GP models with multi-dimensional input and general designs.

ACKNOWLEDGEMENTS

The authors thank the editor and two referees for their comments that substantially improved the article.

FUNDING

The work is partially supported by the National Institutes of Health under Award No. R01DK130067. Gu and Liu acknowledge the partial support from National Science Foundation (NSF) under Award No. DMS-2053423. Fang acknowledges the support from the UCSB academic senate faculty research grants program. Tang is partially supported by Regents Junior Faculty fellowship, Faculty Early Career Acceleration grant, Hellman Family Faculty Fellowship sponsored by UCSB and the NSF under Award No. DMS-2111303.

APPENDIX

A.1 Closed-form Expressions of State Space Representation of GP Having Matérn Covariance with $\nu = 5/2$

Denote $\lambda = \frac{\sqrt{5}}{\gamma}$, $q = \frac{16}{3}\sigma^2\lambda^5$ and $d_i = |x_i - x_{i-1}|$. For $i = 2, \ldots, N$, \mathbf{G}_i and \mathbf{W}_i in (3.3) have the expressions below:

$$\begin{aligned} \mathbf{G}_{i} &= \frac{e^{-\lambda d_{i}}}{2} \\ &\times \begin{pmatrix} \lambda^{2}d_{i}^{2} + 2\lambda + 2 & 2(\lambda d_{i}^{2} + d_{i}) & d_{i}^{2} \\ -\lambda^{3}d_{i}^{2} & -2(\lambda^{2}d_{i}^{2} - \lambda d_{i} - 1) & 2d_{i} - \lambda d_{i}^{2} \\ \lambda^{4}d_{i}^{2} - 2\lambda^{3}d_{i} & 2(\lambda^{3}d_{i}^{2} - 3\lambda^{2}d_{i}) & \lambda^{2}d_{i}^{2} - 4\lambda d_{i} + 2 \end{pmatrix}, \\ \mathbf{W}_{i} &= \frac{4\sigma^{2}\lambda^{5}}{3} \begin{pmatrix} W_{1,1}(x_{i}) & W_{1,2}(x_{i}) & W_{1,3}(x_{i}) \\ W_{2,1}(x_{i}) & W_{2,2}(x_{i}) & W_{2,3}(x_{i}) \\ W_{3,1}(x_{i}) & W_{3,2}(x_{i}) & W_{3,3}(x_{i}) \end{pmatrix}, \end{aligned}$$

with

$$\begin{split} W_{1,1}(x_i) &= \frac{e^{-2\lambda d_i}(3+6\lambda d_i+6\lambda^2 d_i^2+4\lambda^3 d_i^3+2\lambda^4 d_i^4)-3}{-4\lambda^5},\\ W_{1,2}(x_i) &= W_{2,1}(x_i) = \frac{e^{-2\lambda d_i} d_i^4}{2},\\ W_{1,3}(x_i) &= W_{3,1}(x_i)\\ &= \frac{e^{-2\lambda d_i}(1+2\lambda d_i+2\lambda^2 d_i^2+4\lambda^3 d_i^3-2\lambda^4 d_i^4)-1}{4\lambda^3},\\ W_{2,2}(x_i) &= \frac{e^{-2\lambda d_i}(1+2\lambda d_i+2\lambda^2 d_i^2-4\lambda^3 d_i^3+2\lambda^4 d_i^4)-1}{-4\lambda^3},\\ W_{2,3}(x_i) &= W_{3,2}(x_i) = \frac{e^{-2\lambda d_i} d_i^2(4-4\lambda d_i+\lambda^2 d_i^2)}{2}, \end{split}$$

$$W_{3,3}(x_i) = \frac{e^{-2\lambda d_i}(-3 + 10\lambda^2 d_i^2 - 22\lambda^2 d_i^2 + 12\lambda^2 d_i^2 - 2\lambda^4 d_i^4) + 3}{4\lambda},$$

and the stationary covariance of θ_i , $i = 1, ..., \tilde{N}$, is

$$\mathbf{W}_1 = \begin{pmatrix} \sigma^2 & 1 & -\sigma^2 \lambda^2 / 3 \\ 0 & \sigma^2 \lambda^2 / 3 & 1 \\ -\sigma^2 \lambda^2 / 3 & 0 & \sigma^2 \lambda^4 \end{pmatrix},$$

The joint distribution of latent states follows $(\boldsymbol{\theta}^T(x_1),\dots,\boldsymbol{\theta}^T(x_N))^T \sim \mathcal{MN}(\mathbf{0},\boldsymbol{\Lambda}^{-1})$, where the $\boldsymbol{\Lambda}$ is a symmetric block tri-diagonal matrix with the ith diagonal block being $\mathbf{W}_i^{-1} + \mathbf{G}_i^T \mathbf{W}_{i+1}^{-1} \mathbf{G}_i$ for $i=1,\dots,N-1$, and the Nth diagonal block being \mathbf{W}_N^{-1} . The primary off-diagonal block of $\boldsymbol{\Lambda}$ is $-\mathbf{G}_i^T \mathbf{W}_i^{-1}$, for $i=2,\dots,N$.

Suppose $x_i < x^* < x_{i+1}$. Let $d_i^* = |x_i^* - x_i|$ and $d_{i+1}^* = |x_{i+1} - x_i^*|$. The "*" terms \mathbf{G}_i^* and \mathbf{W}_i^* can be computed by replacing d_i in \mathbf{G}_i and \mathbf{W}_i by d_i^* , whereas the "*" terms \mathbf{G}_{i+1}^* and \mathbf{W}_{i+1}^* can be computed by replacing the d_i in \mathbf{G}_i and \mathbf{W}_i by d_{i+1}^* . Furthermore, $\tilde{\mathbf{W}}_{i+1}^* = \mathbf{W}_{i+1}^* + \mathbf{G}_{i+1}^* \mathbf{W}_i^* (\mathbf{G}_{i+1}^*)^T$.

A.2 The Sparse CG-GP Algorithm for Estimating Interaction Kernels

Here we discuss the details of computing the predictive mean and variance in (4.6). The N-vector of velocity observations is denoted as $\tilde{\mathbf{v}}$, where the total number of observations is defined by N=nDML. To compute the predictive mean and variance, the most computational challenging part is to compute N-vector $\mathbf{z}=(\mathbf{U}_s\mathbf{R}_s\mathbf{U}_s^T+\eta\mathbf{I}_N)^{-1}\tilde{\mathbf{v}}$. Here \mathbf{R}_s and \mathbf{U}_s are $\tilde{N}\times\tilde{N}$ and $N\times\tilde{N}$, respectively, where $\tilde{N}=n(n-1)ML/2$ is the number of non-zero unique distance pairs. Note that both \mathbf{U}_s and \mathbf{R}_s^{-1} are sparse. Instead of directly computing the matrix inversion and the matrix-vector multiplication, we utilize the sparsity structure to accelerate the computation in the sparse CG-GP algorithm. In the iteration, we need to efficiently compute

$$\tilde{\mathbf{z}} = (\mathbf{U}_s \mathbf{R}_s \mathbf{U}_s^T + \eta \mathbf{I}_N) \mathbf{z},\tag{A.1}$$

for any real-valued N-vector z.

We have four steps to compute the quantity in (A.1) efficiently. Denote $x_{i,j,m}[t_l]$ the jth spatial coordinate of particle i at time t_l , in the mth simulation, for $i=1,\ldots,n$, $j=1,\ldots,D,\ l=1,\ldots,L$ and $m=1,\ldots,M$. In the following, we use $\mathbf{x}_{\cdot,\cdot,m}[\cdot]$ to denote a vector of all positions in the mth simulation and vice versa. Furthermore, we use z[k] to mean the kth entry of any vector \mathbf{z} , $\mathbf{A}[k,.]$ and $\mathbf{A}[.,k]$ to mean the kth row vector and kth column vector of any matrix \mathbf{A} , respectively. The rank of a particle with position $\mathbf{x}_{i,..,m}[t_l]$ is defined to be P=(m-1)Ln+(l-1)n+i.

First, we reduce the $N \times \tilde{N}$ sparse matrix \mathbf{U}_s of distance difference pairs to an $N \times n$ matrix \mathbf{U}_{re} , where 're' means reduced, with the $((m-1)LnD+(l-1)nD+(j-1)n+i_1,i_2)$ th

entry of \mathbf{U}_{re} being $(x_{i_1,j,m}[t_l]-x_{i_2,j,m}[t_l])$, for any $|i_1-i_2|=1,\ldots,n-1,\ i_1\leq n$ and $i_2\leq n$. Furthermore, we create a $\tilde{N}\times 2$ matrix \mathbf{P}_r in which the hth row records the rank of a distance pair is the hth largest in the zero-excluded sorted distance pairs \mathbf{d}_s , where $P_r[h,1]$ and $P_r[h,2]$ are the rank of rows of these distances in the matrix \mathbf{d}_{mat} , where hth column records the unordered distance pairs of the hth particle for h1, h2.

For any N-vector \mathbf{z} , the kth entry of $\mathbf{U}_s^T \mathbf{z}$ can be written as

$$(\mathbf{U}_{s}^{T}\mathbf{z})[k] = \sum_{j_{k}=0}^{D-1} \mathbf{U}_{re} \left[P_{r}[k,1] + c_{j_{k}}, P_{r}[k,2] - (m-1)nL - (l-1)n \right] \left(z[P_{r}[k,1] + c_{j_{k}}] - z[P_{r}[k,2] + c_{j_{k}}] \right), \tag{A.2}$$

where $c_{j_k} = (D-1)(m-1)nL + (D-1)(l-1)n + nj_k$ for $k = 1, ..., \tilde{N}$, if the kth largest entry of distance pair is from time frame l in the mth simulation. The output is denoted as an \tilde{N} vector $\mathbf{g_1}$, i.e. $\mathbf{g_1} = (\mathbf{U}_s^T \mathbf{z})$.

Second, since the exponential kernel is used, \mathbf{R}_s^{-1} is a tri-diagonal matrix [20]. We modify a Kalman filter step to efficiently compute the product of an upper bi-diagonal $\mathbf{g}_2 = \mathbf{L}_s^T \mathbf{g}_1$, where \mathbf{L}_s is the factor of the Cholesky decomposition $\mathbf{R}_s = \mathbf{L}_s \mathbf{L}_s^T$. Denote the Cholesky decomposition of the inverse covariance the factor $\mathbf{R}_s^{-1} = \tilde{\mathbf{L}}_s \tilde{\mathbf{L}}_s^T$, where $\tilde{\mathbf{L}}_s$ can be written as the lower bi-diagonal matrix below:

$$\tilde{\mathbf{L}}_{s} = \begin{pmatrix}
\frac{1}{\sqrt{1-\rho_{1}^{2}}} & & & \\
\frac{-\rho_{1}}{\sqrt{1-\rho_{1}^{2}}} & \frac{1}{\sqrt{1-\rho_{2}^{2}}} & & \\
& \frac{-\rho_{2}}{\sqrt{1-\rho_{2}^{2}}} & \ddots & & \\
& & \ddots & \frac{1}{\sqrt{1-\rho_{\tilde{N}-1}^{2}}} & \\
& & \frac{-\rho_{\tilde{N}-1}}{\sqrt{1-\rho_{\tilde{N}-1}^{2}}} & 1
\end{pmatrix}, (A.3)$$

where $\rho_k = \exp(-(d_{s,k+1} - d_{s,k})/\gamma)$ for $k = 1, \ldots, \tilde{N} - 1$. We modify the Thomas algorithm [57] to solve \mathbf{g}_2 from equation $(\mathbf{L}_s^T)^{-1}\mathbf{g}_2 = \mathbf{g}_1$. Here $(\mathbf{L}_s^T)^{-1}$ is an upper bi-diagonal matrix with explicit form

$$(\mathbf{L}_{s}^{T})^{-1} = \begin{pmatrix} 1 & \frac{-\rho_{1}}{\sqrt{1-\rho_{1}^{2}}} & & & & \\ & \frac{1}{\sqrt{1-\rho_{1}^{2}}} & & & & \\ & & \ddots & \ddots & & \\ & & & \frac{1}{\sqrt{1-\rho_{\tilde{N}-2}^{2}}} & \frac{-\rho_{\tilde{N}-1}}{\sqrt{1-\rho_{\tilde{N}-1}^{2}}} \\ & & & \frac{1}{\sqrt{1-\rho_{\tilde{N}-1}^{2}}} \end{pmatrix} .$$

$$(A.4)$$

Here only up to 2 entries in each row of $(\mathbf{L}_s^T)^{-1}$ are nonzero. Using a backward solver, the \mathbf{g}_2 can be obtained by the iteration below:

$$\mathbf{g}_{2}[\tilde{N}] = \mathbf{g}_{1}[\tilde{N}]\sqrt{1 - \rho_{\tilde{N}-1}^{2}},$$
 (A.5)

$$\mathbf{g}_{2}[k] = \sqrt{1 - \rho_{k-1}^{2}} \mathbf{g}_{1}[k] + \frac{\rho_{k} \mathbf{g}_{2}[k+1] \sqrt{1 - \rho_{k-1}^{2}}}{\sqrt{1 - \rho_{k}^{2}}}, \quad (A.6)$$

for $k = \tilde{N} - 1, \dots, 2, 1$. Note that the Thomas algorithm is not stable in general, but here the stability issue is greatly improved, as the matrix in the system is bi-diagonal instead of tri-diagonal.

Third, we compute $\mathbf{g}_3 = \mathbf{L}_s \mathbf{g}_2$ by solving $\mathbf{L}_s^{-1} \mathbf{g}_3 = \mathbf{g}_2$:

$$\mathbf{g}_3[1] = \mathbf{g}_2[1],\tag{A.7}$$

$$\mathbf{g}_{3}[k] = \mathbf{g}_{2}[1], \tag{A.7}$$

$$\mathbf{g}_{3}[k] = \sqrt{1 - \rho_{k-1}^{2}} \mathbf{g}_{2}[k] + \rho_{k-1} \mathbf{g}_{3}[k-1], \tag{A.8}$$

for $k = 2, ..., \tilde{N} - 1$.

Finally, we denote a $MLn \times n$ matrix \mathbf{P}_c . \mathbf{P}_c is initialized as a zero matrix. And then for $r_c = 1, \dots, MLn$, row r_c of \mathbf{P}_c stores the ranks of distances between the ith particle and other n-1 particles in \mathbf{d}_s . For instance, at the lth time step in the mth simulation, particle i has n-1 non-zero distances $||\mathbf{x}_1-\mathbf{x}_i||,\ldots,||\mathbf{x}_{i-1}-\mathbf{x}_i||,||\mathbf{x}_{i+1}-\mathbf{x}_i||$ $|\mathbf{x}_i||,\ldots,||\mathbf{x}_n-\mathbf{x}_i||$ with ranks $h_1,\ldots,h_{i-1},h_{i+1},\ldots h_n$ in \mathbf{d}_s . Then the ((m-1)Ln + (l-1)n + i)th row of \mathbf{P}_c is filled with $(h_1, \ldots, h_{i-1}, h_{i+1}, \ldots, h_n)$.

Given any N-vector \mathbf{g}_3 , the kth entry of $\mathbf{U}_s\mathbf{g}_3$ can be written as

$$(\mathbf{U}_s \mathbf{g}_3)[k] = \mathbf{U}_{re}[k, .]\mathbf{g}_3[\mathbf{P}_c[k', .]^T], \tag{A.9}$$

assuming that k satisfies k = (m-1)LnD + (l-1)nD + jn+iand k' = i + (m-1)Ln + (l-1)n for some m, l, j and i, and $k = 1, \ldots, N$. The output of this step is an N vector \mathbf{g}_4 , with the kth entry being $\mathbf{g}_4[k] := (\mathbf{U}_s^T \mathbf{g}_3)[k]$, for $k = 1, \dots, N$.

We summarize the sparse CG-GP algorithm using the following steps to compute $\tilde{\mathbf{z}}$ in (A.1) below.

- 1. Use equation (A.2) to compute $\mathbf{g}_1[k] = (\mathbf{U}_s^T \mathbf{z})[k]$, for $k=1,\ldots,N$.
- 2. Use equations (A.5) and (A.6) to solve \mathbf{g}_2 from $(\mathbf{L}_s^T)^{-1}\mathbf{g}_2 = \mathbf{g}_1$ where $(\mathbf{L}_s^T)^{-1} = (\mathbf{L}_s^{-1})^T$ with \mathbf{L}_s^{-1} given in equation (A.4).
- 3. Use equations (A.7) and (A.8) to solve \mathbf{g}_3 from $\mathbf{L}_s^{-1}\mathbf{g}_3 =$ \mathbf{g}_2 , where \mathbf{L}_s^{-1} is given in equation (A.4).
- 4. Use equation (A.9) to compute $\mathbf{g}_4[k] = (\mathbf{U}_s \mathbf{g}_3)[k]$ and let $\tilde{\mathbf{z}} = \mathbf{g}_4 + \eta \mathbf{z}$.

A.3 Interaction Kernels in Simulated Studies

Here we give the expressions of the truncated L-J and OD kernels of particle interaction in [34, 13]. The truncated LJ kernel is given by

$$\phi_{LJ}(d) = \begin{cases} c_2 \exp(-c_1 d^{12}), & d \in [0, 0.95], \\ \frac{8(d^{-4} - d^{-10})}{3}, & d \in (0.95, \infty), \end{cases}$$

$$c_1 = -\frac{1}{12} \frac{c_4}{c_3(0.95)^{11}}$$
 and $c_2 = c_3 \exp(c_1(0.95)^{12})$,

with $c_3 = \frac{8}{3}(0.95^{-4} - 0.95^{-10})$ and $c_4 = \frac{8}{3}(10(0.95)^{-11} - 4(0.05)^{-5})$

The OD kernel of particle interaction is defined as

$$\phi_{OD}(d) = \begin{cases} 0.4, & d \in [0, c_5), \\ -0.3\cos(10\pi(d - c_5)) + 0.7, & d \in [c_5, c_6), \\ 1, & d \in [c_6 \le d < 0.95), \\ 0.5\cos(10\pi(d - 0.95)) + 0.5, & d \in [0.95, 1.05), \\ 0, & d \in [1.05, \infty), \end{cases}$$

where $c_5 = \frac{1}{\sqrt{2}} - 0.05$ and $c_6 = \frac{1}{\sqrt{2}} + 0.05$.

A.4 Further Numerical Results on Estimating Interaction Kernels

We outline the numerical results of estimating the interaction functions at $N_1^* = 1000$ equally spaced distance pairs at $d \in [0,5]$ and $d \in [0,1.5]$ for the truncated LJ and OD, respectively. For each configuration, we repeat the simulation $N_2^* = 10$ times and compute the predictive error in each simulation. The total number of test points is $N^* = N_1^* N_2^* = 10^4$. For demonstration purposes, we do not add a noise into the simulated data (i.e. $\sigma_0^2 = 0$). The range and nugget parameters are fixed to be $\gamma = 5$ and $\eta = 10^{-5}$. We compute the normalized root of mean squared error (NRMSE) in estimating the interaction kernel function:

NRMSE =
$$\frac{1}{\sigma_{\phi}} \sqrt{\sum_{i=1}^{N^*} \frac{(\hat{\phi}(d_i^*) - \phi(d_i^*))^2}{N^*}},$$

where $\hat{\phi}(.)$ is the estimated interaction kernel from the velocities and positions of the particles; σ_{ϕ} is the standard deviation of the interaction function at test points.

Table 1. NRMSE of the sparse CG-GP method for estimating the truncated LJ and OD kernels of particle interaction.

Truncated LJ	n = 50	n = 200	n = 50	n = 200
	L = 1	L = 1	L = 10	L = 10
Uniform	.11	.021	.026	.0051
Normal	.037	.012	.0090	.0028
Log-uniform	.043	.0036	.0018	.00091
OD	n = 50	n = 200	n = 50	n = 200
	L = 1	L = 1	L = 10	L = 10
Uniform	.024	.0086	.0031	.0036
Normal	.13	.013	.038	.0064
Log-uniform	.076	.0045	.0018	.00081

Table 1 gives the NRMSE of the sparse CG-GP method for the truncated LJ and OD kernels at 12 configurations. Typically the estimation is the most accurate when the initial positions of the particles are sampled from the loguniform design for a given number of observations and an interaction kernel. This is because the contributions to the velocities from the kernel function are proportional to the distance of particle in (4.1), and small contributions from the interaction kernel at small distance values make the kernel hard to estimate from the trajectory data in general. When the initial positions of the particles are sampled from the log-uniform design, more particles are close to each other, which provides more information to estimate the interaction kernel at a small distance.

Furthermore, the predictive error of the interaction kernel is smaller, when the trajectories with a larger number of particle sizes or at longer time points are used in estimation, as more observations typically improve predictive accuracy. The sparse CG-GP algorithm reduces the computational cost substantially, which allows more observations to be used for making predictions.

Accepted 29 September 2022

REFERENCES

- [1] Adrian, R. J. and Westerweel, J. Particle image velocimetry 30. Cambridge university press (2011).
- [2] Anderson, K. R., Johanson, İ. A., Patrick, M. R., Gu, M., Segall, P., Poland, M. P., Montgomery-Brown, E. K. and Miklius, A. Magma reservoir failure and the onset of caldera collapse at Kīlauea Volcano in 2018. Science 366(6470) (2019).
- [3] Banerjee, S., Carlin, B. P. and Gelfand, A. E. Hierarchical modeling and analysis for spatial data. Crc Press (2014). MR3362184
- [4] BARBIERI, M. M. and BERGER, J. O. Optimal predictive model selection. The annals of statistics 32(3) 870-897 (2004). https://doi.org/10.1214/009053604000000238. MR2065192
- [5] BAYARRI, M. J., BERGER, J. O., PAULO, R., SACKS, J., CAFEO, J. A., CAVENDISH, J., LIN, Q H, C. and TU, J. A framework for validation of computer models. *Technometrics* 49(2) 138–154 (2007). https://doi.org/10.1198/004017007000000092. MR2380530
- [6] BERGER, J. O. and PERICCHI, L. R. The intrinsic Bayes factor for model selection and prediction. *Journal of the American Statistical Association* 91(433) 109–122 (1996). https://doi.org/10. 2307/2291387. MR1394065
- [7] BERGER, J. O., DE OLIVEIRA, V. and SANSÓ, B. Objective Bayesian analysis of spatially correlated data. *Journal of the American Statistical Association* 96(456) 1361–1374 (2001). https://doi.org/10.1198/016214501753382282. MR1946582
- [8] BERGER, J. O., LISEO, B. and WOLPERT, R. L. Integrated likelihood methods for eliminating nuisance parameters. Statistical science 14(1) 1–28 (1999). https://doi.org/10.1214/ss/1009211803. MR1702200
- [9] COUZIN, I. D., KRAUSE, J., FRANKS, N. R. and LEVIN, S. A. Effective leadership and decision-making in animal groups on the move. *Nature* 433(7025) 513–516 (2005).
- [10] CRESSIE, N. and JOHANNESSON, G. Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society:* Series B (Statistical Methodology) 70(1) 209–226 (2008). https://doi.org/10.1111/j.1467-9868.2007.00633.x. MR2412639
- [11] Datta, A., Banerjee, S., Finley, A. O. and Gelfand, A. E. Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association* 111(514) 800–812 (2016). https://doi.org/10.1080/01621459. 2015.1044091. MR3538706

- [12] DE FINETTI, B. La prévision: ses lois logiques, ses sources subjectives. Annales de l'institut Henri Poincaré 7. 1–68 (1937). MR1508036
- [13] Feng, J., Ren, Y. and Tang, S. Data-driven discovery of interacting particle systems using Gaussian processes (2021). arXiv preprint 2106.02735.
- [14] GELFAND, A. E., BANERJEE, S. and GAMERMAN, D. Spatial process modelling for univariate and multivariate dynamic spatial data. Environmetrics: The official journal of the International Environmetrics Society 16(5) 465–479 (2005). https://doi.org/10.1002/env.715. MR2147537
- [15] GRAMACY, R. B. and APLEY, D. W. Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics* 24(2) 561–578 (2015). https://doi.org/10.1080/10618600.2014.914442. MR3357395
- [16] GRAMACY, R. B. and LEE, H. K. Cases for the nugget in modeling computer experiments. Statistics and Computing 22(3) 713–722 (2012). https://doi.org/10.1007/s11222-010-9224-x. MR2909617
- [17] Gu, M. and Li, H. Gaussian Orthogonal Latent Factor Processes for Large Incomplete Matrices of Correlated Data. *Bayesian Analysis*. 1–26 (2022). https://doi.org/10.1214/21-BA1295
- [18] Gu, M. and Shen, W. Generalized probabilistic principal component analysis of correlated data. *Journal of Machine Learning Research* 21(13) (2020). MR4071196
- [19] Gu, M., Palomo, J. and Berger, J. O. RobustGaSP: Robust Gaussian Stochastic Process Emulation in R. The R Journal 11(1) 112–136 (2019). https://doi.org/10.1214/17-AOS1648. https://doi.org/10.32614/RJ-2019-011. MR3851764
- [20] Gu, M., Wang, X. and Berger, J. O. Robust Gaussian stochastic process emulation. Annals of Statistics 46(6A) 3038–3066 (2018). https://doi.org/10.1214/17-AOS1648. MR3851764
- [21] HACKBUSCH, W. Iterative solution of large sparse systems of equations 95. Springer (1994). https://doi.org/10.1007/ 978-1-4612-4288-8. MR1247457
- [22] HARTIKAINEN, J. and SARKKA, S. Kalman filtering and smoothing solutions to temporal Gaussian process regression models. In 2010 IEEE International Workshop on Machine Learning for Signal Processing 379–384. IEEE (2010).
- [23] HENKES, S., FILY, Y. and MARCHETTI, M. C. Active jamming: Self-propelled soft particles at high density. *Physical Review E* 84(4), 040301 (2011).
- [24] HESTENES, M. R. and STIEFEL, E. Methods of conjugate gradients for solving. *Journal of research of the National Bureau of Standards* 49(6) 409 (1952). MR0060307
- [25] HIGDON, D., GATTIKER, J., WILLIAMS, B. and RIGHTLEY, M. Computer model calibration using high-dimensional output. *Journal of the American Statistical Association* 103(482) 570–583 (2008). https://doi.org/10.1198/016214507000000888. MR2523994
- [26] KALMAN, R. E. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82(1) 35–45 (1960). MR3931993
- [27] KATZFUSS, M. and GUINNESS, J. A general framework for Vecchia approximations of Gaussian processes. Statistical Science 36(1) 124–141 (2021). https://doi.org/10.1214/19-STS755. MR4194207
- [28] KAZIANKA, H. and PILZ, J. Objective Bayesian analysis of spatial data with uncertain nugget and range parameters. Canadian Journal of Statistics 40(2) 304–327 (2012). https://doi.org/10. 1002/cjs.11132. MR2927748
- [29] LAKSHMINARAYANAN, B., PRITZEL, A. and BLUNDELL, C. Simple and scalable predictive uncertainty estimation using deep ensembles. Advances in neural information processing systems 30 (2017).
- [30] LAM, C. and YAO, Q. Factor modeling for high-dimensional time series: inference for the number of factors. The Annals of Statistics 40(2) 694–726 (2012). https://doi.org/10.1214/12-AOS970. MR2933663
- [31] LAM, C., YAO and AND BATHIA N, Q. Estimation of latent fac-

- tors for high-dimensional time series. Biometrika 98(4) 901–918 (2011). https://doi.org/10.1093/biomet/asr048. MR2860332
- [32] LEE, J., BAHRI, Y., NOVAK, R., SCHOENHOLZ, S. S., PENNINGTON, J. and Sohl-Dickstein, J. Deep neural networks as gaussian processes (2017). arXiv preprint 1711.00165.
- [33] LINDGREN, F., RUE, H. and LINDSTRÖM, J. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 73(4) 423-498 (2011). https://doi.org/10.1111/j.1467-9868.2011.00777. x. MR2853727
- [34] Lu, F., Zhong, M., Tang, S. and Maggioni, M. Nonparametric inference of interaction laws in systems of agents from trajectory data. Proc. Natl. Acad. Sci. U.S.A. 116(29) 14424–14433 (2019). https://doi.org/10.1073/pnas.1822012116. MR3984488
- [35] MARCHETTI, M. C., JOANNY, Q F, J., RAMASWAMY, S., LIVER-POOL, T. B., PROST, J., RAO, M. and SIMHA, R. A. Hydrodynamics of soft active matter. Reviews of modern physics 85(3) 1143 (2013). https://doi.org/10.1017/jfm.2012.131. MR2969140
- [36] MOTSCH, S. and TADMOR, E. Heterophilious dynamics enhances consensus. SIAM review **56**(4) 577–621 (2014). https://doi.org/ 10.1137/120901866. MR3274797
- [37] Muré, J. Propriety of the reference posterior distribution in Gaussian process modeling. The Annals of Statistics 49(4) 2356-2377 (2021). https://doi.org/10.1214/20-aos2040. MR4319254
- [38] Neal, R. M. Bayesian learning for neural networks 118. Springer (2012).
- [39] Paulo, R. Default priors for Gaussian processes. Annals of statistics 33(2) 556-582 (2005). https://doi.org/10.1214/ 009053604000001264. MR2163152
- [40] PAULO, R., GARCÍA-DONATO, G. and PALOMO, J. Calibration of computer models with multivariate output. Computational Statistics and Data Analysis 56(12) 3959–3974 (2012). https://doi.org/ 10.1016/j.csda.2012.05.023. MR2957846
- [41] Petris, G., Petrone, S. and Campagnoli, P. Dynamic linear models with. Springer (2009). https://doi.org/10.1007/b135794. MR2730074
- [42] RAFTERY, A. E., MADIGAN, D. and HOETING, J. A. Bayesian model averaging for linear regression models. Journal of the American Statistical Association 92(437) 179-191 (1997). https://doi.org/10.2307/2291462. MR1436107
- [43] Rapaport, D. C. and Rapaport, D. C. R. The art of molecular dynamics simulation. Cambridge university press (2004).
- [44] Rasmussen, C. E. Gaussian processes for machine learning. MIT Press (2006), MR2514435
- [45] RAUCH, H. E., TUNG, F. and STRIEBEL, C. T. Maximum likelihood estimates of linear dynamic systems. AIAA journal 3(8) 1445–1450 (1965). https://doi.org/10.2514/3.3166. MR0181489
- [46] Ren, C., Sun, D. and He, C. Objective Bayesian analysis for a spatial model with nugget effects. Journal of Statistical Planning and Inference 142(7) 1933-1946 (2012). https://doi.org/10.1016/ j.jspi.2012.02.034. MR2903403
- [47] ROUSTANT, O., GINSBOURGER, D. and DEVILLE, Y. DiceKriging, DiceOptim: Two R Packages for the Analysis of Computer Experiments by Kriging-Based Metamodeling and Optimization. Journal of Statistical Software **51**(1) 1–55 (2012). https://doi.org/10. 18637/jss.v051.i01
- [48] RUE, H., MARTINO, S. and CHOPIN, N. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. Journal of the royal statistical society: Series B (statistical methodology) 71(2) 319–392 (2009). https:// doi.org/10.1111/j.1467-9868.2008.00700.x. MR2649602
- [49] Saad, Y. Iterative methods for sparse linear systems. SIAM (2003).https://doi.org/10.1016/S1570-579X(01)80025-2.MR1853234

- [50] SACKS, J., WELCH, W. J., MITCHELL, T. J. and WYNN, H. P. Design and analysis of computer experiments. Statistical science 4(4) 409-423 (1989). MR1041765
- SANCHEZ-GONZALEZ, A., GODWIN, J., PFAFF, T., YING, R., LESKOVEC, J. and Battaglia, P. Learning to simulate complex physics with graph networks. In International Conference on Machine Learning 8459-8468. PMLR (2020).
- Santner, T. J., Williams, B. J. and Notz, W. I. The design and analysis of computer experiments. Springer (2003). https:// doi.org/10.1007/978-1-4757-3799-8. MR2160708
- Seeger, M., Teh, Y. Q. W. and Jordan, M. Semiparametric latent factor models. Technical Report (2005).
- SNELSON, E. and GHAHRAMANI, Z. Sparse Gaussian processes using pseudo-inputs. Advances in neural information processing systems 18 1257 (2006).
- STROUD, J. R., STEIN, M. L. and LYSEN, S. Bayesian and maximum likelihood estimation for Gaussian processes on an incomplete lattice. Journal of computational and Graphical Statistics **26**(1) 108–120 (2017). https://doi.org/10.1080/10618600.2016. 1152970. MR3610412
- [56] Surjanovic, S. and Bingham, D. Virtual Library of Simulation Experiments: Test Functions and Datasets (2017).
- Thomas, L. H. Elliptic problems in linear difference equations over a network. Watson Sci. Comput. Lab. Rept., Columbia University, New York 1-71. (1949).
- TIPPING, M. E. and BISHOP, C. M. Probabilistic principal component analysis. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 61(3) 611–622 (1999). https://doi.org/ 10.1111/1467-9868.00196. MR1707864
- [59] Vecchia, A. V. Estimation and model identification for continuous spatial processes. Journal of the Royal Statistical Society: Series B (Methodological) **50**(2) 297–312 (1988). MR0964183
- Wen, Z. and Yin, W. A feasible method for optimization with orthogonality constraints. Mathematical Programming 142(1-2) https://doi.org/10.1007/s10107-012-0584-1. (2013).MR3127080
- West, M. and Harrison, P. J. Bayesian Forecasting & Dynamic Models 2nd ed. Springer (1997). MR1482232
- [62] West, M. and Harrison, J. Bayesian forecasting and dynamic models. Springer (2006). https://doi.org/10.1007/ 978-1-4757-9365-9. MR1020301
- Whittle, P. Stochastic process in several dimensions. Bulletin of the International Statistical Institute 40(2) 974–994 (1963). MR0173287
- WILSON, A. G. and IZMAILOV, P. Bayesian deep learning and a probabilistic perspective of generalization. Advances in neural information processing systems 33 4697-4708 (2020).

Mengyang Gu. Department of Statistics and Applied Probability, University of California, Santa Barbara, USA. E-mail address: mengyang@pstat.ucsb.edu

Xubo Liu. Department of Statistics and Applied Probability, University of California, Santa Barbara, USA. E-mail address: xubo@umail.ucsb.edu

Xinyi Fang. Department of Statistics and Applied Probability, University of California, Santa Barbara, USA. E-mail address: xinyifang@umail.ucsb.edu

Sui Tang. Department of Statistics and Applied Probability, University of California, Santa Barbara, USA. E-mail address: suitang@ucsb.edu