





gptools: Scalable Gaussian Process Inference with Stan

Till Hoffmann 
Harvard T. H. Chan
School of Public Health

Jukka-Pekka Onnela 
Harvard T. H. Chan
School of Public Health

Abstract

Gaussian processes (GPs) are sophisticated distributions to model functional data. Whilst theoretically appealing, they are computationally cumbersome except for small datasets. We implement two methods for scaling GP inference in **Stan**: First, a general sparse approximation using a directed acyclic dependency graph; second, a fast, exact method for regularly spaced data modeled by GPs with stationary kernels using the fast Fourier transform. Based on benchmark experiments, we offer guidance for practitioners to decide between different methods and parameterizations. We consider two real-world examples to illustrate the package. The implementation follows **Stan**'s design and exposes performant inference through a familiar interface. Full posterior inference for ten thousand data points is feasible on a laptop in less than 20 seconds. Details on how to get started using the popular interfaces **cmdstanpy** for Python and **cmdstanr** for R are provided.

Keywords: Gaussian process, Fourier transform, sparse approximation, Stan, Python, R.

1. Introduction

Gaussian processes (GPs) are flexible non-parametric models for functions with applications in time series analysis (Roberts, Osborne, Ebden, Reece, Gibson, and Aigrain 2013), geospatial statistics (Kriging 1951), robotics (Deisenroth, Fox, and Rasmussen 2015), and beyond. More formally, a GP is a distribution over functions $f(x)$ such that any finite set of n values $\mathbf{f} = f(\mathbf{x})$ evaluated at $\mathbf{x} = \{x_1, \dots, x_n\}$ follows a multivariate normal distribution (Rasmussen and Williams 2006). The distribution is thus fully specified by its mean $\mu(x) = \mathbb{E}(f(x))$ and covariance kernel $k(x, x') = \text{COV}(f(x), f(x'))$. A range of problem-specific kernels have been developed, such as squared exponential and Matérn kernels to model local correlations and sinusoidal kernels to capture periodic signals (Duvenaud 2014, Chapter 2). In general, evaluating the likelihood requires inverting the covariance matrix \mathbf{K} obtained by evaluating the kernel k at all pairs of observed locations \mathbf{x} . Unfortunately, the computational cost of

inverting \mathbf{K} scales as $\mathcal{O}(n^3)$, making GPs prohibitively expensive save for relatively small datasets.

Diverse schemes have been developed to approximate the likelihood for larger datasets, such as low-rank approximations of the covariance matrix (Hensman, Fusi, and Lawrence 2013), nearest-neighbor approximations (Wu, Pleiss, and Cunningham 2022), and Fourier methods (Hensman, Durrande, and Solin 2017; Greengard 2021). Numerous packages provide implementations in different programming languages, but most of them focus exclusively on GPs which makes it difficult to integrate GPs in larger models. For example, **GPflow** (Matthews *et al.* 2017) is a library for the Python (Python Software Foundation 2023) machine learning framework **TensorFlow** (Abadi *et al.* 2016). It implements many common likelihood functions, e.g., Poisson for count data, Bernoulli for classification, and Student t for robust regression, but hierarchical models cannot be easily constructed. **GPYtorch** (Gardner, Pleiss, Bindel, Weinberger, and Wilson 2018) offers similar functionality for **PyTorch** (Paszke *et al.* 2019); this implementation can be used as part of the general purpose probabilistic programming framework **Pyro** (Bingham *et al.* 2019). Likelihood gradients are readily available through the underlying machine learning frameworks, and both packages implement automatic differentiation variational inference (ADVI; Kucukelbir, Tran, Ranganath, Gelman, and Blei 2017). **GPY** (Sheffield Machine Learning Group 2012) and **george** (Ambikasaran, Foreman-Mackey, Greengard, Hogg, and O’Neil 2015) (both implemented in Python) as well as **GaussianProcesses.jl** (Fairbrother, Nemeth, Rischard, Brea, and Pinder 2022) for Julia (Bezanson, Edelman, Karpinski, and Shah 2017), **mlegp** (Dancik and Dorman 2008) and **spBayes** (Finley, Banerjee, and Carlin 2007) (both for R R Core Team 2024), and **GPML** (Rasmussen and Nickisch 2010) for MATLAB (The MathWorks Inc. 2022) implement GP inference from the ground up giving maximum flexibility. However, this approach requires error-prone manual implementation of likelihoods, their gradients, and dedicated sampling algorithms, such as Gibbs samplers. Further, users need to become intimately familiar with the codebase to develop custom models. While sampling the posterior using Hamiltonian Monte Carlo (HMC) is supported by **GPY**, **GaussianProcesses.jl**, and **GPML**, the packages implement a standard leap-frog HMC sampler, requiring extensive hyperparameter tuning for efficient sampling; **mlegp** only supports maximum-likelihood estimation of parameters. **spBayes** obtains samples efficiently using a Gibbs sampler by first marginalizing with respect to the latent Gaussian process, and **R-INLA** (Gómez-Rubio 2020) offers performant inference using the integrated nested Laplace approximation. However, both remain limited by the computational cost of inverting the covariance kernel. Building on **Stan** (Carpenter *et al.* 2017), **brms** (Bürkner 2017) offers a high-level R interface for fitting Bayesian models, including GPs, but it currently only supports squared exponential covariance kernels. **brms** implements a basis-function approximation to apply GPs to larger datasets (Riutort-Mayol, Bürkner, Andersen, Solin, and Vehtari 2022) which we discuss further in Section 8.

Despite its popularity, a library for scalable GP inference is lacking for the probabilistic programming framework **Stan**. Here, we discuss the implementation of two scalable approaches in **Stan** which can be easily integrated using the language’s `#include` directive. Building on **Stan** has distinct advantages: First, performant general purpose inference algorithms are implemented and well tested, including ADVI, penalized maximum likelihood estimation, and advanced HMC samplers with automatic hyperparameter tuning. Second, GPs can be used as components in larger hierarchical models without having to adapt or extend the library. Third, **Stan** supports automatic differentiation obviating the need for implementing gradients

manually. Finally, Stan has an engaged community whose members support one another in building statistical models, including extensive expertise in GPs.

In Section 2, we provide a brief introduction to Stan and the `cmdstanpy` (Stan Development Team 2023b) interface for Python. In Section 3, we present an implementation of GPs on directed acyclic graphs which can encode structured dependencies between observations and generalizes nearest-neighbor approximations. In Section 4, we demonstrate how to use Fourier methods to evaluate the GP likelihood exactly for observations on a regular grid in one and two dimensions. Both implementations are designed to dovetail with Stan’s design philosophy, facilitating their integration into larger models. In Section 5, we demonstrate how to use the package with a simple example in both Python and R. We consider a benchmark problem and discuss the importance of different parameterizations for performant inference in Section 6. Furthermore, we demonstrate the utility of both approaches with two examples: Inferring the density of trees in a 50 ha plot in Panama (Condit, Perez, Aguilar, Lao, Foster, and Hubbell 2019) and predicting passenger numbers on the London Underground transportation network. We summarize our contributions in Section 8 and discuss how the package can be employed to build more complex models.

As Stan does not have a package repository, we have published the library as a Python package `gptools-stan` on the Python Package Index (PyPI) at <https://pypi.org/project/gptools-stan/> and as an R package `gptoolsStan` on the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=gptoolsStan>. The packages include the Stan library code and provide utility functions to integrate with the popular Stan interfaces `cmdstanpy` (Stan Development Team 2023b) and `cmdstanr` (Gabry, Češnovar, and Johnson 2023), respectively. The library can also be obtained directly from <https://github.com/onnella-lab/gptools>. Extensive technical documentation and various examples are available at <https://gptools-stan.readthedocs.org/>.

2. Introduction to Stan

Stan is a probabilistic programming framework, comprising both a concise R-like syntax to declare probabilistic models and an efficient Hamilton Monte Carlo algorithm to draw posterior samples (Betancourt 2018). Readers familiar with Stan may skip to Section 3.

Each Stan program consists of blocks to declare inputs, parameters, and the probabilistic model. For a concrete example, consider a linear regression model with $n \times p$ design matrix \mathbf{X} , coefficient vector $\boldsymbol{\theta}$ with p elements, outcome vector \mathbf{y} with n elements, and observation noise variance σ^2 , i.e.,

$$\mathbf{y} \sim \text{Normal}(\mathbf{X}\boldsymbol{\theta}, \sigma^2).$$

The corresponding Stan program is shown below.

```
data {
  int n, p;
  matrix [n, p] X;
  vector [n] y;
}
parameters {
  vector [p] theta;
  real<lower=0> sigma;
```

```

}
model {
  theta ~ normal(0, 1);
  sigma ~ gamma(2, 2);
  y ~ normal(X * theta, sigma);
}

```

The `data` block defines inputs required to evaluate the likelihood of the model; `parameters` declares parameters of the model including any constraints, such as the noise scale being non-negative. Finally, the `model` block declares priors for parameters and the observation model of outcomes \mathbf{y} given covariates \mathbf{X} and parameters $\boldsymbol{\theta}$ and σ . To illustrate the analysis workflow in Python, we generated synthetic data using the **NumPy** package (Harris *et al.* 2020) by sampling from the prior predictive distribution with $n = 100$ observations and $p = 3$ covariates. We fix the random number generator seed for reproducibility.

```

>>> import numpy as np
>>> np.random.seed(0)
>>> n = 100
>>> p = 3
>>> X = np.random.normal(0, 1, (n, p))
>>> theta = np.random.normal(0, 1, p)
>>> sigma = np.random.gamma(2, 2)
>>> y = np.random.normal(X @ theta, sigma)
>>> print(f"coefficients: {theta}")
>>> print(f"observation noise scale: {sigma}")

```

```

coefficients: [-1.307  1.658 -0.118]
observation noise scale: 1.867

```

We used the **cmdstanpy** interface to compile the above model, draw posterior samples, and report summary statistics.

```

>>> import cmdstanpy
>>> model = cmdstanpy.CmdStanModel(stan_file = "linear.stan")
>>> fit = model.sample(data = {"n": n, "p": p, "X": X, "y": y}, seed = 0)
>>> print(fit.summary())

```

```

          5%      50%      95%    ...
theta[1]  -1.535   -1.229   -0.932  ...
theta[2]   1.439    1.754    2.062  ...
theta[3]  -0.385   -0.062    0.266  ...
sigma      1.715    1.921    2.170  ...
...

```

The 90% marginal posterior intervals for all parameters are consistent with the values used to generate the data. Having gained some intuition for **Stan** and **cmdstanpy**, we consider two approaches to scalable GP inference and their implementation in **Stan** in the following two sections.

3. Gaussian processes with structured dependencies

The joint distribution of observations \mathbf{f} may be expressed as the product of conditional distributions

$$p(\mathbf{f}) = p(f_1) \prod_{j=2}^n p(f_j | f_{j-1}, \dots, f_1). \quad (1)$$

The conditional structure in Equation 1 can be encoded by a directed acyclic graph (DAG) whose nodes represent observations such that a directed edge exists from a node j to each of its predecessors $\mathcal{P}_j = \{j-1, \dots, 1\}$; the ordering is arbitrary. If two observations do not depend on one another, the corresponding edge can be removed from the DAG to reduce the computational cost. In particular, evaluating each factor of Equation 1 requires inverting a matrix with size equal to the number of predecessors of the corresponding node—a substantial saving if the graph is sparse. For example, nearest-neighbor methods, a special case, reduce the asymptotic runtime to $\mathcal{O}(nq^3)$ by retaining only edges from each node to at most q of its nearest predecessors. This approach can yield excellent approximations provided that the neighborhoods are large enough and that the kernel only models local correlations (Wu *et al.* 2022). For example, nearest-neighbor methods are not suitable for periodic kernels but can be approximated by structured dependencies if the period is known, such as diurnal or yearly patterns. Pseudocode to approximate the likelihood of a GP realization \mathbf{f} using structured dependencies is shown in Algorithm 1 for a general kernel k . The algorithm approximates the log likelihood iteratively by evaluating the conditional mean and variance for each node i given its predecessors \mathcal{P}_i in lines 4–7; the conditional distributions are available in closed form for multivariate normal distributions (Gelman, Carlin, S, Dunson, Vehtari, and Rubin 2013, Appendix A1). The evaluation of likelihood contributions can be further accelerated by parallelizing the loop in line 3.

We implemented a custom distribution in Stan such that a GP with squared exponential kernel on a DAG embedded in a p -dimensional space can be specified as

```
f ~ gp_graph_exp_quad_cov(loc, x, sigma, length_scale, edges);
```

where vector [n] `loc` is the prior mean, and array [n] vector [p] `x` is an array of loca-

Algorithm 1 Evaluate the log likelihood of the Gaussian process realization \mathbf{f} given its mean $\boldsymbol{\mu}$, locations of observations \mathbf{x} , covariance kernel k , and the dependency graph encoded as a set of predecessors \mathcal{P} .

```

1: function GP_GRAPH_LPDF( $\mathbf{f} | \boldsymbol{\mu}, \mathbf{x}, k, \mathcal{P}$ )
2:    $\mathcal{L} \leftarrow \text{NORMAL\_LPDF}(f_1 | \mu_1, k(x_1, x_1))$    ▷ Marginal log likelihood for the first node.
3:   for  $i \in [2..n]$  do
4:      $\boldsymbol{\Sigma} \leftarrow k(\mathbf{x}_{\mathcal{P}_i}, \mathbf{x}_{\mathcal{P}_i})$    ▷ Covariance among predecessors of  $i$ .
5:      $\mathbf{s} \leftarrow k(x_i, \mathbf{x}_{\mathcal{P}_i})$    ▷ Covariance between  $i$  and its predecessors.
6:      $\nu \leftarrow \mathbf{s}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_{\mathcal{P}_i}$    ▷ Conditional mean.
7:      $\tau^2 = k(x_i, x_i) - \mathbf{s}^\top \boldsymbol{\Sigma}^{-1} \mathbf{s}$    ▷ Conditional variance.
8:      $\mathcal{L} \leftarrow \mathcal{L} + \text{NORMAL\_LPDF}(f_i | \nu, \tau^2)$    ▷ Conditional log likelihood for  $i$  given  $\mathcal{P}_i$ .
9:   end for
10:  return  $\mathcal{L}$ 
11: end function

```

tions in p dimensions for each of the n nodes of the graph. The parameters `real sigma` and `real length_scale` control the marginal scale and smoothness of the kernel which is defined as (Duvenaud 2014, Chapter 2)

$$k(x, x') = \text{sigma}^2 \times \exp\left(-\frac{|x - x'|^2}{2 \times \text{length_scale}^2}\right). \quad (2)$$

The larger the length scale the more slowly the GP varies because even points with substantial separation $|x - x'|$ remain highly correlated. The graph is encoded by the edge list `array [,] int edges`, a two-dimensional array of integer node labels. Each column represents an edge from the node in the second row to the corresponding node in the first row, i.e., edges “point up” from successors to predecessors indicating the dependence of the former on the latter. For example, the following edge list represents the directed line graph of four nodes $1 \leftarrow 2 \leftarrow 3 \leftarrow 4$:

```
array [2, 3] int edges = {
  {1, 2, 3},
  {2, 3, 4}
};
```

Following Stan’s indexing convention, node labels start at one. Similar distributions are provided for the `matern32` and `matern52` kernels.

4. Gaussian processes in Fourier space

We can use Fourier methods to evaluate the likelihood efficiently if three conditions are satisfied (Rasmussen and Williams 2006, Appendix B). First, we need to consider observation points \mathbf{x} on a regular grid to reap the computational benefits of the fast Fourier transform (FFT) (Press, Teukolsky, Vetterling, and Flannery 2007). Second, the kernel must be stationary, i.e., $k(x, x') = k(x - x')$ such that the correlation only depends on the separation between observations. Third, the kernel must be n -periodic because the FFT is subject to periodic boundary conditions, i.e., $k(x + n, x') = k(x - x')$, where n is the number of observations. These conditions may seem overly restrictive. However, in many settings, data naturally form a regular grid, e.g., financial time series with fixed sampling interval (Hoffmann, Peel, Lambiotte, and Jones 2020), resampled or binned time series (Flaxman, Wilson, Neill, Nickisch, and Smola 2015), or rasterized images (Tipping and Bishop 2002). Likewise, stationary kernels, such as squared exponential and Matérn kernels, are common choices for modeling functional data using GPs. Finally, the effect of periodic boundary conditions can be attenuated by padding the domain, as discussed in more detail in Sections 7.2 and Appendix A.

Because the Fourier transform is a linear operator and \mathbf{f} is multivariate normal, the discrete Fourier coefficients

$$\tilde{f}_\xi = \sum_{j=0}^{n-1} \exp\left(-\frac{2\pi i \xi j}{n}\right) f_j$$

are also multivariate normal, where ξ is the (discrete) frequency, f_j is the GP at the j -th grid point, and i is the imaginary number. Assuming $\mu(x) = 0$ for simplicity, the mean of Fourier

Algorithm 2 Evaluate the log likelihood of the Gaussian process realization \mathbf{f} given its mean $\boldsymbol{\mu}$ and Fourier-transformed covariance kernel $\tilde{\mathbf{k}}$. Range indexing is inclusive on the left and exclusive on the right, i.e., $\mathbf{f}_{a:b} = \{f_a, \dots, f_{b-1}\}$.

```

1: function GP_RFFT_LPDF( $\mathbf{f} \mid \boldsymbol{\mu}, \tilde{\mathbf{k}}$ )
2:    $\mathbf{z} \leftarrow |\text{RFFT}(\mathbf{f} - \boldsymbol{\mu})|$             $\triangleright$  Modulus of centered real FFT with  $\lfloor n/2 \rfloor + 1$  elements.
3:    $\mathcal{L} \leftarrow \text{NORMAL\_LPDF}(z_0 \mid 0, n\tilde{k}_0)$             $\triangleright$  Real zero-frequency term.
4:   if  $n \bmod 2 = 1$  then
5:      $m \leftarrow \frac{n+1}{2}$             $\triangleright$  Index following highest-frequency complex coefficient.
6:   else
7:      $m \leftarrow \frac{n}{2}$             $\triangleright$  Index following highest-frequency complex coefficient.
8:      $\mathcal{L} \leftarrow \text{NORMAL\_LPDF}(z_m \mid 0, n\tilde{k}_m)$             $\triangleright$  Real Nyquist-frequency term.
9:   end if
10:   $\mathcal{L} \leftarrow \mathcal{L} + 2 \times \text{NORMAL\_LPDF}(\mathbf{z}_{1:m} \mid 0, n\tilde{\mathbf{k}}_{1:m})$             $\triangleright$  Complex oscillatory terms.
11:  return  $\mathcal{L}$ 
12: end function

```

coefficients is zero and their expected complex-conjugate product at two different frequencies ξ and ξ' is

$$\begin{aligned} \mathbb{E}(\tilde{f}_\xi \overline{\tilde{f}_{\xi'}}) &= \sum_{j=0}^{n-1} \sum_{j'=0}^{n-1} \exp\left(-\frac{2\pi i}{n}(j\xi - j'\xi')\right) k(j - j') \\ &= \sum_{j'=0}^{n-1} \exp\left(-\frac{2\pi i j'}{n}(\xi - \xi')\right) \sum_{\Delta=-j'}^{n-1-j'} \exp\left(-\frac{2\pi i \Delta}{n}\right) k(\Delta), \end{aligned}$$

where we changed variables to $j = \Delta + j'$ in the second line. The argument of the inner sum is n -periodic, and we may shift the limits of summation to $[0..n - 1]$ without changing the sum. The change of limits decouples the two sums. The first is a sum-representation of the Kronecker delta $n\delta_{\xi\xi'}$; the second is the Fourier transform of the kernel $\tilde{\mathbf{k}}$. We obtain

$$\mathbb{E}(\tilde{f}_\xi \overline{\tilde{f}_{\xi'}}) = n\delta_{\xi\xi'} \tilde{k}_\xi.$$

Fourier coefficients of different frequencies are thus independent with variance $n\tilde{k}_\xi$. Subject to careful bookkeeping, we can evaluate the likelihood exactly, as illustrated in Algorithm 2. Because \mathbf{f} is real, we use the real FFT (RFFT) for efficiency. It comprises $\lfloor n/2 \rfloor + 1$ complex coefficients because just under half the coefficients are redundant (Press *et al.* 2007, Chapter 12.3). The zero-frequency term \tilde{f}_0 and, for even n , the Nyquist frequency term $\tilde{f}_{n/2}$ are real (see lines 3 and 8 of Algorithm 2, respectively). The complex coefficients contribute twice in line 10 to account for the redundant terms omitted by the RFFT.

We implemented a custom distribution in Stan such that a GP on a grid can be specified as

```
f ~ gp_rfft(loc, cov_rfft);
```

where vector $[\mathbf{n}]$ `loc` is the prior mean and vector $[\mathbf{n} \% \% 2 + 1]$ `cov_rfft` is the RFFT of the kernel evaluated on the grid (`%%` denotes floor division in Stan). Fortunately, the RFFT of common kernels, such as the squared exponential kernel and Matérn kernels, can be evaluated directly in the Fourier domain (Rasmussen and Williams 2006, Chapter 4), as

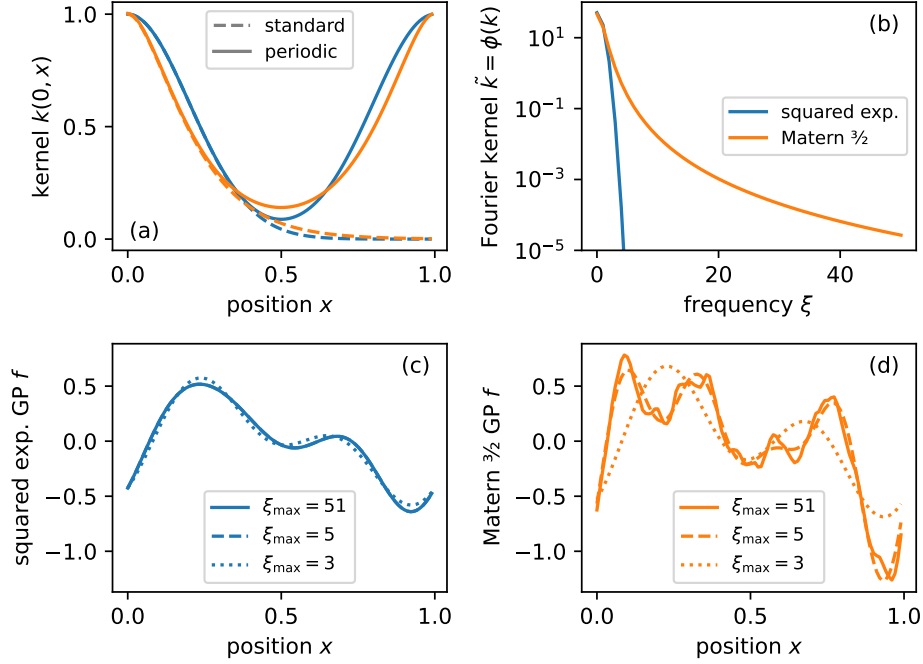


Figure 1: Gaussian processes can often be captured by a small number of Fourier modes. Panel (a) shows the periodic and non-periodic (standard) versions of the squared exponential kernel (blue) and Matérn $3/2$ kernel (orange) with $\sigma = 1$ and $\ell = 0.2$. The two versions are hardly distinguishable for $x < \ell$. Padding may need to be introduced if non-periodic signals are modeled with periodic kernels. The power spectrum of the two periodic kernels is shown in panel (b). The squared exponential kernel has negligible power for all but the first few frequencies, explaining why it is often considered too smooth to “represent natural phenomena” (Handcock and Stein 1993). Panels (c) and (d) show realizations of GPs with squared exponential and Matérn $3/2$ kernels, respectively. Different line styles correspond to approximations with a different number of Fourier modes. For the squared exponential kernel, the number of modes can be reduced by an order of magnitude without substantially affecting realizations. The Matérn kernel requires more modes due to its heavy-tailed power spectrum.

shown in Figure 1 (see Appendix B for definitions of the kernel). Evaluating the kernel in the Fourier domain also obviates the need for small “nugget” variance or “jitter” typically required for numerical stability (Neal 1997). We thus only need to evaluate one Fourier transform, that of the signal, to evaluate the likelihood. The library provides the following functions to evaluate Fourier-domain kernels:

```
gp_periodic_exp_quad_cov_rfft(n, sigma, length_scale, period)
gp_periodic_matern_cov_rfft(n, nu, sigma, length_scale, period)
```

where `period` is the size of the domain, `n` is the number of grid points with spacing `period / n`, and `nu` is the smoothness parameter of the Matérn kernel. For `nu = 0.5`, the Matérn kernel reduces to the exponential kernel $k(x, x') = \sigma^2 \exp\left(\frac{|x-x'|}{\text{length_scale}}\right)$, and, in the limit of large `nu`, it converges to the squared exponential kernel defined in Equation 2 (Rasmussen and Williams 2006, Chapter 4.2). The parameters `sigma` and `length_scale` have the same

meaning as in Section 3. Equivalent functions, which we discuss further in Section 7.2, are provided for two-dimensional grids. We implemented the Fourier-domain kernels by naively discretizing frequencies. This approach works well if the number of grid points is large and the correlation length is small compared with the size of the domain. More sophisticated methods may be required otherwise (Borovitskiy, Terenin, Mostowsky, and Deisenroth 2020).

5. Getting started

We demonstrate how to use **gptools** using a simple example: Drawing samples from a Gaussian process prior using the Fourier methods discussed in Section 4. The Stan model comprises five blocks:

```
functions {
  #include gptools/util.stan
  #include gptools/fft.stan
}
data {
  int n;
  real<lower=0> sigma, length_scale, period;
}
transformed data {
  vector [n %% 2 + 1] cov_rfft =
    gp_periodic_exp_quad_cov_rfft(n, sigma, length_scale, period) + 1e-9;
}
parameters {
  vector [n] f;
}
model {
  f ~ gp_rfft(zeros_vector(n), cov_rfft);
}
```

The first block `functions` includes the source of **gptools** in the program, and the second block `data` declares the number of grid points `n` and kernel parameters (see Section 4 for definitions). Because the kernel parameters are fixed, we can precompute the RFFT of the kernel in the `transformed data` block. Finally, we declare the vector `f` representing the Gaussian process in the `parameters` block and specify the prior in the `model` block.

5.1. Getting started in Python using `cmdstanpy`

To install **gptools** for Python, run

```
$ pip install gptools-stan
```

from the command line. This will download the Stan source and install a Python package with utility functions to compile models using the **cmdstanpy** interface (Stan Development Team 2023b). Executing the following Python script will compile and fit the model:

```

>>> import cmdstanpy
>>> from gptools.stan import get_include
>>> model = cmdstanpy.CmdStanModel(stan_file = "getting_started.stan",
...   stanc_options = {"include-paths": get_include()})
>>> fit = model.sample(
...   data = {"n": 100, "sigma": 1, "length_scale": 0.1, "period": 1},
...   chains = 1, iter_warmup = 500, iter_sampling = 50)
>>> fit.f.shape

(50, 100)

```

The function `get_include` returns the path to the Stan source files of **gptools**. The path is passed to `cmdstanpy.CmdStanModel`, which compiles the model, as `stanc_options`. The object `fit.f` is a two-dimensional array comprising 50 samples of the vector `f` each having 100 elements. If this is the first time **cmdstanpy** is used, **cmdstan** ([Stan Development Team 2023a](#)) may need to be installed by running

```
$ python -m cmdstanpy.install_cmdstan
```

from the command line before the model can be compiled and fit.

5.2. Getting started in R using `cmdstanr`

To install **gptools** for R, run

```

R> install.packages("cmdstanr",
+   repos = c("https://mc-stan.org/r-packages/", getOption("repos")))
R> install.packages("gptoolsStan")

```

from the R console. The first command installs the **cmdstanr** interface ([Gabry *et al.* 2023](#)) which is not yet available on CRAN, and the second installs **gptools** for R. Executing the following R script will compile and fit the model:

```

R> library("cmdstanr")
R> library("gptoolsStan")
R> model <- cmdstan_model(stan_file = "getting_started.stan",
+   include_paths = gptools_include_path())
R> fit <- model$sample(
+   data = list(n = 100, sigma = 1, length_scale = 0.1, period = 1),
+   chains = 1, iter_warmup = 500, iter_sampling = 50)
R> f <- fit$draws("f")
R> dim(f)

```

```
[1] 50  1 100
```

The function `gptools_include_path` returns the path to the Stan source files of **gptools**. The path is passed to `cmdstan_model` which compiles the model. The object `f` is an array of samples from the prior distribution. If this is the first time **cmdstanr** is used, **cmdstan** may need to be installed by running

```
R> install_cmdstan()
```

from the R console before the model can be compiled and fit.

6. Benchmark and the importance of parameterizations

We consider a simple benchmark problem to study the performance of different methods and compare them with standard Gaussian process inference which inverts the kernel. The model comprises a one-dimensional zero-mean Gaussian process prior with squared exponential kernel and an independent normal observation model with variance κ^2 , i.e.,

$$\begin{aligned} \mathbf{f} &\sim \text{MultivariateNormal}(0, \mathbf{K}) \\ \mathbf{y} &\sim \text{Normal}(\mathbf{f}, \kappa^2). \end{aligned} \tag{3}$$

We used a marginal kernel scale $\sigma = 1$ and unit correlation length $\ell = 1$ to evaluate the covariance matrix \mathbf{K} on an integer grid, i.e., $\mathbf{x} = \{0, \dots, n - 1\}$. Employing the `cmdstanpy` interface, we drew 100 posterior samples each from 20 independent chains after 100 warmup samples. Warmup samples are used to adapt the sampler for efficient exploration of the posterior (Hoffman and Gelman 2014). Default values were used for all other parameters. We considered different dataset sizes between $n = 2^4$ and $n = 2^{14}$ and allocated a maximum computational budget of one minute for each chain and all n , i.e., individual chains were terminated if they did not complete after 60 seconds. The mean runtime as a function of dataset size is shown in panels (a) and (b) of Figure 2 for small ($\kappa = 0.1$) and large ($\kappa = 10$) noise scales as solid lines, respectively. As expected, the runtime of the standard approach grows rapidly as n increases. We observed an empirical runtime scaling of $n^{\approx 2.5}$ for the standard approach, not dissimilar from the expected asymptotic scaling of $\mathcal{O}(n^3)$. Exploring models with more than a few hundred data points is prohibitively expensive—even for this simple setup. For the graph-based approach, we used the five nearest predecessors ($q = 5$) to construct a dependency graph. The method is comparatively slow for small datasets but outperforms the standard approach as n grows. The Fourier approach has the best performance irrespective of dataset size but is limited to observations on a grid.

The model in Equation 3 employs the natural centered parameterization (Papaspiliopoulos, Roberts, and Sköld 2007), i.e., each observation y_i is independent given the corresponding latent f_i . This parameterization works well if the data are informative (small κ) because each observation y_i constrains the corresponding latent parameter f_i . The elements of \mathbf{f} are thus relatively uncorrelated under the posterior, and the Hamiltonian sampler can explore the distribution efficiently (Hoffman and Gelman 2014).

However, if the data are weak (large κ), they cannot independently constrain each element of \mathbf{f} and the GP prior dominates the posterior. The resulting correlation among elements of \mathbf{f} frustrates the sampler, especially if the correlation length is large. We can overcome this challenge by employing a non-centered parameterization such that the parameters of the model are uncorrelated under the prior (Papaspiliopoulos *et al.* 2007). Here, we reparameterize the model in terms of a white noise vector \mathbf{z} of the same size as \mathbf{f} and obtain realizations of the GP $\mathbf{f} = \phi^{-1}(\mathbf{z}, \boldsymbol{\mu}, \mathbf{K})$ using an inverse transform ϕ^{-1} which must be selected carefully to ensure \mathbf{f} follows the desired distribution. We chose the inverse transform for consistency with the FFT: The forward transform maps to the Fourier domain, and the inverse transform

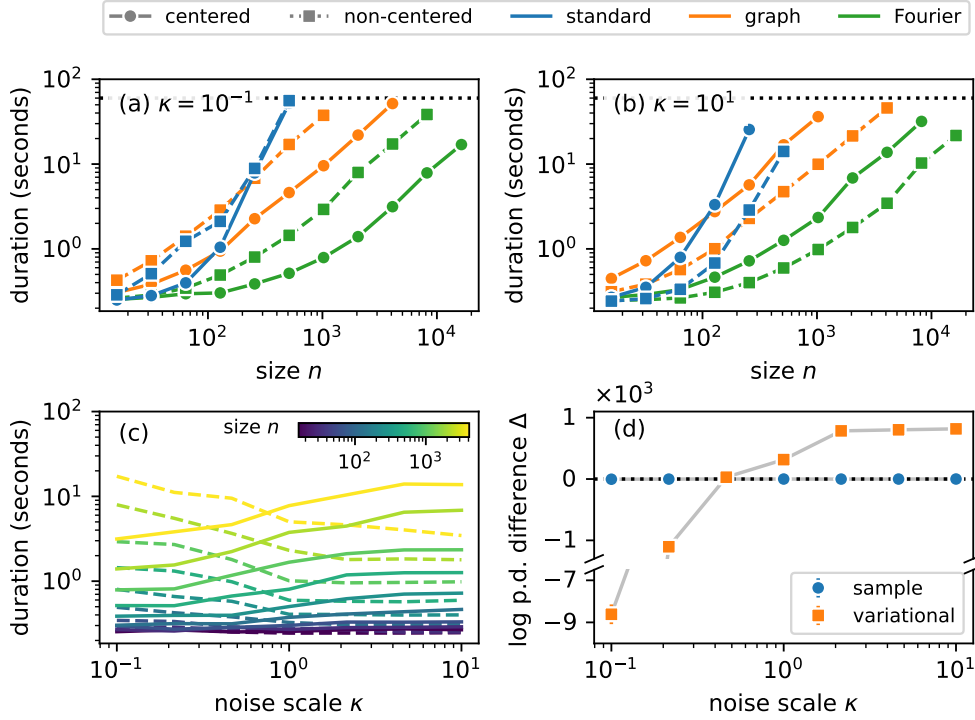


Figure 2: Different approaches, parameterizations, as well as the informativeness of the data substantially affect runtimes. Centered parameterizations are preferable when the data are strong (small observation noise scale κ), and non-centered parameterizations are superior when the data are weak (large κ), as shown in panels (a) and (b), respectively. Independent of parameterization and dataset size, Fourier methods offer the best performance when observations are regularly spaced. For datasets exceeding a few hundred observations, graph methods are faster than the standard approach which requires inversion of the covariance matrix. The dotted horizontal line represents the maximum computational budget of 60 s. As shown in panel (c) for the Fourier approach, the runtime of the centered and non-centered parameterizations increases and decreases, respectively, as the data become less informative. While different parameterizations are primarily a performance concern for drawing posterior samples, they have important consequences for variational posterior approximations. Panel (d) shows the difference Δ in log posterior density (p.d.) between the non-centered and centered parameterization on 20% held-out data for $n = 1024$ using the Fourier approach. If the model is fit with a variational mean-field approximation, the non-centered parameterization offers higher log posterior scores than the centered parameterization when the data are weak and vice versa. Bootstrapped standard errors are smaller than the size of markers in panel (d).

maps to real space. The reparameterized model is

$$\begin{aligned}
 \mathbf{z} &\sim \text{Normal}(0, 1) \\
 \mathbf{f} &= \phi^{-1}(\mathbf{z}, 0, \mathbf{K}) \\
 \mathbf{y} &\sim \text{Normal}(\mathbf{f}, \kappa^2).
 \end{aligned} \tag{4}$$

Algorithm 3 Transform white noise \mathbf{z} to a Gaussian process realization \mathbf{f} given its mean $\boldsymbol{\mu}$, locations of observations \mathbf{x} , covariance kernel k , and the dependency graph encoded as a set of predecessors \mathcal{P} .

```

1: function GP_INV_GRAPH( $\mathbf{z}, \boldsymbol{\mu}, \mathbf{x}, k, \mathcal{P}$ )
2:    $f_1 \leftarrow \mu_1 + \sqrt{k(x_1, x_1)}z_1$    ▷ Sample first observation from the marginal distribution.
3:   for  $i \in [2..n]$  do
4:     ...   ▷ Compute conditional mean  $\nu$  and variance  $\tau^2$  as in Algorithm 1.
5:      $f_i \leftarrow \nu + \tau z_i$            ▷ Sample from the conditional distribution given  $\mathcal{P}_i$ .
6:   end for
7:   return  $\mathbf{f}$ 
8: end function

```

Algorithm 4 Transform white noise \mathbf{z} to a Gaussian process realization \mathbf{f} given its mean $\boldsymbol{\mu}$ and Fourier-transformed covariance kernel $\tilde{\mathbf{k}}$. Range indexing is inclusive on the left and exclusive on the right, i.e., $\mathbf{f}_{a:b} = \{f_a, \dots, f_{b-1}\}$.

```

1: function GP_INV_RFFT( $\mathbf{z}, \boldsymbol{\mu}, \tilde{\mathbf{k}}$ )
2:    $\tau = \sqrt{n\tilde{\mathbf{k}}}$    ▷ Evaluate standard deviation of Fourier coefficients.
3:    $\tilde{f}_0 \leftarrow \tau_0 z_0$    ▷ Real zero-frequency term.
4:   if  $n \bmod 2 = 1$  then
5:      $m \leftarrow \frac{n+1}{2}$    ▷ Index following highest-frequency complex coefficient.
6:   else
7:      $m \leftarrow \frac{n}{2}$    ▷ Index following highest-frequency complex coefficient.
8:      $\tilde{f}_m \leftarrow \tau_m z_m$    ▷ Real Nyquist-frequency term.
9:   end if
10:   $\tilde{\mathbf{f}}_{1:m} \leftarrow \tau_{1:m} (\mathbf{z}_{1:m} + \mathbf{i} \times \mathbf{z}_{m+(n+1) \bmod 2:n}) / \sqrt{2}$    ▷ Complex oscillatory terms.
11:   $\mathbf{f} \leftarrow \text{INV\_RFFT}(\tilde{\mathbf{f}}, n)$    ▷ Inverse RFFT returning a vector with  $n$  elements.
12:  return  $\mathbf{f}$ 
13: end function

```

We implemented the following transforms for the graph-based and Fourier approaches:

```

 $\mathbf{f} = \text{gp\_inv\_graph\_exp\_quad\_cov}(\mathbf{z}, \text{loc}, \mathbf{x}, \text{sigma}, \text{length\_scale}, \text{edges});$ 
 $\mathbf{f} = \text{gp\_inv\_rfft}(\mathbf{z}, \text{loc}, \text{cov\_rfft});$ 

```

where vector $[\mathbf{n}] \mathbf{z}$ are the non-centered white noise parameters and all other parameters are as described previously. For the standard method, we implemented the non-centered parameterization as $\mathbf{f} = \mathbf{Lz}$ (Papaspiliopoulos *et al.* 2007), where \mathbf{L} is the Cholesky decomposition of the covariance matrix \mathbf{K} such that $\mathbf{K} = \mathbf{LL}^\top$. Algorithm 3 implements the transform for approximate GPs using structured dependencies and closely follows Algorithm 1. Instead of evaluating the log likelihood iteratively given the conditional distribution, the algorithm draws a sample sequentially by transforming white noise to the target distribution given previous samples in lines 2 and 5. Unlike Algorithm 1 the loop in line 3 cannot be parallelized because the conditional mean and variance depend on the results of previous iterations. The transformation from white noise to a GP realization using Fourier methods is illustrated in Algorithm 4. As in Algorithm 2, we account for the real zero-frequency term and, for even n , Nyquist frequency term in lines 3 and 8, respectively. Line 10 constructs the $\lfloor (n-1)/2 \rfloor$

complex Fourier coefficients from $\mathbf{z}_{1:m}$ (real part) and $\mathbf{z}_{m+(n+1) \bmod 2:n}$ (imaginary part). The $(n+1) \bmod 2$ term in the index accounts for the presence of the Nyquist frequency for even n . The omission of redundant terms in the RFFT is addressed by dividing the complex coefficients by $\sqrt{2}$.

As shown in panels (a) and (b) of Figure 2, the non-centered parameterization (dashed lines) is more performant than the centered parameterization (solid lines) if the noise scale is large and vice versa. Panel (c) further illustrates the importance of choosing the right parameterization: The runtime differs by up to a factor of five as we vary the noise scale κ . The higher-frequency terms of smooth GPs have low power, as shown in panel (b) of Figure 1. We can further improve performance of the non-centered parameterization by discarding all but the first few low-frequency terms. This approach is particularly effective for the squared exponential kernel because the power spectrum decays rapidly with increasing frequency. For the example shown in panel (c), a GP using only the first five Fourier modes is indistinguishable from the GP considering all 51 modes, reducing the dimensionality of the parameter space by an order of magnitude. GPs with Matérn kernels typically require more Fourier modes because the power spectrum has a relatively heavy tail.

While parameterization is primarily a performance concern for Hamiltonian Monte Carlo samplers, it can have a substantial impact on the predictive ability of models if variational mean-field inference is used. Variational approximations of the posterior tend to assign low probability mass to regions of the parameter space where the full posterior has low mass (Bishop 2006, Chapter 10.1). Consequently, variational approximations are too narrow if the posterior is highly correlated, and we expect predictions to be overconfident. To test this hypothesis, we sampled synthetic data with $n = 1024$ data points from the prior predictive distribution and fitted the models in Equations 3 and 4 to 80% of each synthetic dataset using variational inference (Kucukelbir *et al.* 2017). Stan’s ADVI implementation approximates the posterior by a product of independent normal distributions, one for each parameter. The variational approximation is optimized in an unconstrained space. Constrained parameters are obtained by applying a transform, e.g., an exponential transform to obtain a positive parameter such as the length scale ℓ . We repeated the analysis 20 times for each noise scale κ , terminating the algorithm, as before, if it did not complete after one minute.

We evaluated the predictive ability of the fitted models by evaluating the log posterior density on the 20% held-out GP realizations, i.e., $\log p(\mathbf{f}_{\text{test}} | \mathbf{y}_{\text{train}})$. To compare the parameterizations, we approximated the log posterior difference $\Delta = \log p(\mathbf{f}_{\text{test}} = \phi^{-1}(\mathbf{z}_{\text{test}}) | \mathbf{y}_{\text{train}}) - \log p(\mathbf{f}_{\text{test}} | \mathbf{y}_{\text{train}})$ using a Gaussian kernel density estimator (Bishop 2006, Chapter 2.5.1). The log posterior differences Δ were averaged over the 20 independent fits (or fewer if the algorithm failed to complete within the allocated computational budget). As shown in panel (d) of Figure 2, the non-centered parameterization makes better predictions than the centered parameterization when the noise scale is large and vice versa. The error bars shown represent bootstrapped standard errors, i.e., the standard deviation of Δ resampled from the population of independent chains (Rubin 1981). We also fitted the two parameterizations by drawing posterior samples using Stan’s Hamiltonian sampler and evaluated the log posterior difference. Different parameterizations did not affect the predictive performance but had a significant impact on runtime.

7. Illustrations

7.1. Passengers on the London Underground transportation network

Millions of people use the London Underground transportation network, commonly referred to as the “Tube”, to travel across the city each day (Transport for London 2019). The number of passengers using each station is affected by various factors, including how connected it is and which zone the station is in (the network comprises nine transport zones). In addition

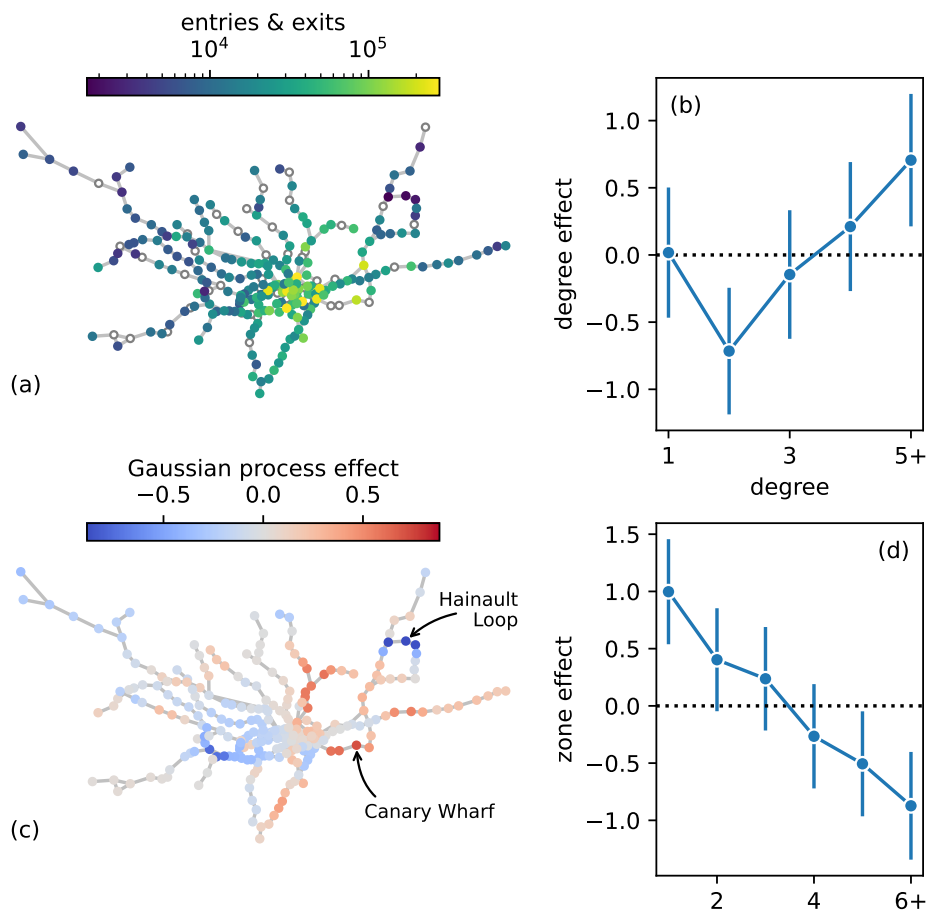


Figure 3: A Gaussian process on the London Underground transportation network identifies idiosyncrasies of transport use. Panel (a) shows the daily average number of entries and exits for each of the 267 stations in 2019 together with the transport network. Empty nodes denote held-out data. Panels (b) and (d) show the effect of degree and zone on passenger numbers, respectively. Central stations with larger degree tend to have more passengers. Termini with degree one have uncharacteristically many passengers because they serve as interchanges for longer-distance trains. Panel (c) shows the GP which explains residual variations in passenger volumes after controlling for zone and degree. On the one hand, stations in the Hainault loop have relatively few passengers because they are served infrequently compared with nearby stations. On the other hand, Canary Wharf has surprisingly many entries and exits because it is a busy financial center in London.

to these fixed effects, we also expect passenger numbers to be affected by smooth spatial effects, e.g., due to variability in population density. The spatial effect can naturally be modeled as a GP with structured dependencies induced by the transport network itself. We converted the undirected graph to a directed acyclic graph in two steps. First, we assigned an integer label to each node. Second, we added an edge from node i to node j if $i < j$ and the corresponding edge exists in the undirected graph. The order of nodes is arbitrary because the joint probability in Equation 1 can be factorized in any order. We collected network data from the Transport for London open data API (Transport for London 2022) and obtained the average daily number of entries and exits at each station in 2019 (Transport for London 2019), as shown in panel (a) of Figure 3. The model includes fixed effects for each zone and degree, i.e., the number of connections a station has. The corresponding regression coefficients were mildly regularized by half t priors with two degrees of freedom (Gelman, Jakulin, Pittau, and Su 2008). We truncated the degree and zone of each station at five and six, respectively, because only few stations exceed these values. The overall number of passengers is captured by a scalar μ , and we used the GP to explain any residual effects. A squared exponential kernel was employed for the covariance, and we used a half t prior for the marginal scale. The correlation length of the kernel is not identifiable if it is smaller than the smallest distance between stations (0.16 km) or larger than the extent of the transportation network (62 km) (Trangucci, Betancourt, and Vehtari 2016). We thus used a log-uniform prior on the interval [0.32 km, 31 km] to suppress extreme length scales. All distances were evaluated in the Ordnance Survey National Grid projection (epsg:27700). A non-centered parameterization was used because the residual effects are not strongly identified by the data after controlling for zone and degree. We used a log-normal observation model (rather than a model for count data) because passenger data are heavy-tailed and reported as daily averages.

```

functions {
  #include gptools/util.stan
  #include gptools/graph.stan
}
data {
  int num_stations, num_edges, num_zones, num_degrees;
  array [num_stations] vector [2] station_locations;
  array [num_stations] int passengers;
  array [2, num_edges] int edge_index;
  matrix [num_stations, num_zones] one_hot_zones;
  matrix [num_stations, num_degrees] one_hot_degrees;
}
parameters {
  vector [num_stations] z;
  real mu;
  real<lower=0> sigma, kappa;
  real<lower=log(0.32), upper=log(31)> log_length_scale;
  vector [num_zones] zone_effect;
  vector [num_degrees] degree_effect;
}
transformed parameters {
  real length_scale = exp(log_length_scale);

```



```

vector [num_stations] f = gp_inv_graph_exp_quad_cov(
  z, zeros_vector(num_stations), station_locations, sigma,
  length_scale, edge_index);
vector [num_stations] log_mean = mu + f + one_hot_zones
  * zone_effect + one_hot_degrees * degree_effect;
}
model {
  z ~ std_normal();
  sigma ~ student_t(2, 0, 1);
  zone_effect ~ student_t(2, 0, 1);
  degree_effect ~ student_t(2, 0, 1);
  kappa ~ student_t(2, 0, 1);
  for (i in 1:num_stations) {
    if (passengers[i] > 0) {
      log(passengers[i]) ~ normal(log_mean[i], kappa);
    }
  }
}

```

We use an implicit uniform prior for `log_length_scale` such that no corresponding statement appears in the model block.

We fitted the model to 80% of the passenger data using the below Python code, withholding 20% of the stations uniformly at random for later evaluation. Held-out data are encoded as -1 in the Stan model.

```

>>> from gptools.stan import get_include
>>> import json
>>> import numpy as np
>>> with open("tube-stan.json") as fp:
...     data = json.load(fp)
>>> train_mask = np.random.binomial(1, 0.8, data["num_stations"])
>>> data["passengers"] = np.where(train_mask, data["passengers"], -1)
>>> model = compile_model(stan_file = "tube.stan",
...     stanc_options = {"include-paths": get_include()})
>>> fit = model.sample(data)
>>> print(fit.diagnose())

```

Processing csv files: ...

Checking sampler transitions treedepth.
Treedepth satisfactory for all transitions.

Checking sampler transitions for divergences.
No divergent transitions found.

Checking E-BFMI - sampler transitions HMC potential energy.
E-BFMI satisfactory.

Effective sample size satisfactory.

Split R-hat values satisfactory all parameters.

Processing complete, no problems detected.

We employed the default configuration of **cmdstanpy** to draw posterior samples, resulting in four independent chains with 2000 samples each. The `print(fig.diagnose())` call evaluates and reports a suite of diagnostics to identify potential problems and assess convergence. For example, the split \hat{R} statistic compares samples both within and between chains to determine whether they are likely to have mixed well (Vehtari, Gelman, Simpson, Carpenter, and Bürkner 2021). Tree depth, divergence, and Bayesian fraction of missing information (BFMI) are technical diagnostics to assess whether the sampler was able to explore the posterior distribution; the effective sample size estimates the number of independent samples drawn which may be smaller than 2000 due to autocorrelation within each chain (Betancourt 2018). Here, the samples satisfied all posterior checks offered by **cmdstan**.

Panel (b) shows the effect of degree on passenger numbers on the log scale. They tend to increase with the degree of a station as they offer passengers a variety of travel options. Termini with degree one are an exception: Their passenger numbers are uncharacteristically large because they serve as stepping stones to longer-distance travel beyond the Tube network. Unsurprisingly, central stations in zones one to three tend to have more passengers than stations in the suburbs (zones four and above), as shown in panel (d). The GP captures any residuals that cannot be explained by the degree of the station or the zone it is located in, as shown in panel (c). For example, on the one hand, Canary Wharf has the largest residual effect. It is one of London’s financial centers, and the station serves tens of thousands of commuters each day despite being a station without an interchange. On the other hand, stations in the north of the Hainault loop have the largest negative residual effect because the stations are served by only three trains an hour (Transport for London 2020). Passengers divert to nearby stations that are served by twelve trains an hour (Transport for London 2020). Comparing the model with a model without GP effects using the log posterior predictive distribution on held-out data, we observe no significant difference after bootstrapping errors. Nevertheless, this example illustrates how our package can be used to easily construct GPs with structured dependencies. It may be necessary to obtain an edge list from geospatial data, e.g., for the analysis of spatially correlated outcomes (Morris, Wheeler-Martin, Simpson, Mooney, Gelman, and DiMaggio 2019). The best approach depends on the problem at hand and how many neighbors are considered for each geographical unit. Common packages for manipulating geospatial data include **shapely** (Gillies, Van der Wel, Van den Bossche, Taves, Arnott, and Ward 2023) and **geopandas** (Van den Bossche *et al.* 2023) for Python and **spdep** (Bivand 2022) for R.

7.2. Density of *T. panamensis* on a 50 ha plot in Panama

To illustrate the use of Fourier methods, we consider the density of *T. panamensis* trees during the 2015 census of the 50 ha Barro Colorado plot in Panama (Condit *et al.* 2019). The plot is divided into quadrants of 20 m side length. As shown in panel (a) of Figure 4, the data comprise the frequency of trees within each quadrant, i.e., a matrix of count data with shape (25, 50). The observed tree frequency counts \mathbf{y} were modeled by a negative-

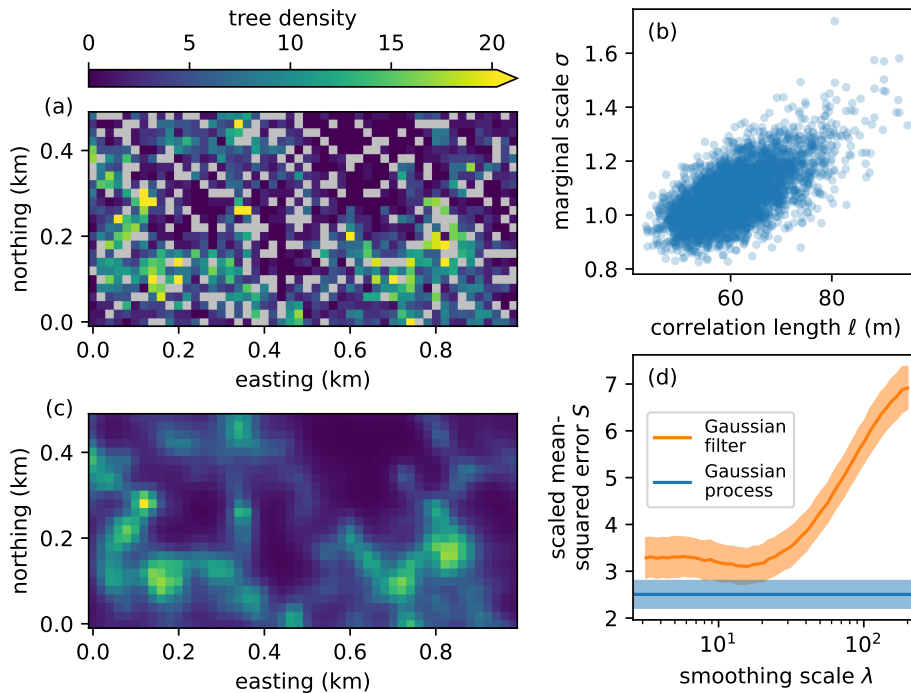


Figure 4: A Gaussian process based on the two-dimensional Fourier transform can accurately predict the frequency of trees. Panel (a) shows the number of *T. panamensis* trees per 20 m quadrant collected during the 2015 census of the 50 ha Barro Colorado plot in Panama (Condit *et al.* 2019) as a heat map. Gray quadrants indicate the 20% held-out test data. Panel (c) shows the posterior median of the expected tree frequency, recovering the held-out data and smoothing the empirical frequencies. Posterior samples of the correlation length and marginal scale of the Matérn kernel are shown in panel (b). Correlation length samples are well below the padding (200 m) introduced to mitigate the effect of periodic boundary conditions. Panel (d) shows the scaled mean squared error with bootstrapped errors on the held-out data under the Gaussian process model and a Gaussian filter with variable smoothing scale.

binomial distribution to account for possible overdispersion. We used a latent GP to model the log-mean of this distribution and capture the tree density. We employed a Matérn kernel with smoothness parameter $\nu = 3/2$, half t prior for the marginal scale and overdispersion parameter, and log-uniform prior for the correlation length as in Section 7.1. Because the quadrants are regularly spaced, the likelihood can be evaluated exactly using Fourier methods, as discussed in Section 4. However, unlike the FFT, trees are not subject to periodic boundary conditions. To mitigate this issue and reduce correlation between opposing sides of the plot, we padded the matrix with ten additional quadrants in each dimension (corresponding to 200 m) resulting in a matrix with shape (35, 60). Despite increasing the number of latent variables by almost 70%, the method is faster than the standard approach which inverts the covariance matrix.

```
functions {
  #include gptools/util.stan
  #include gptools/fft.stan
```

```

}
data {
  int num_rows, num_cols, num_rows_padded, num_cols_padded;
  array [num_rows, num_cols] int frequency;
}
parameters {
  matrix [num_rows_padded, num_cols_padded] z;
  real mu;
  real<lower=0> sigma, kappa;
  real<lower=log(2), upper=log(28)> log_length_scale;
}
transformed parameters {
  real<lower=0> length_scale = exp(log_length_scale);
  matrix [num_rows_padded, num_cols_padded %/% 2 + 1] rfft2_cov =
    gp_periodic_matern_cov_rfft2(1.5, num_rows_padded, num_cols_padded,
      sigma, [length_scale, length_scale]',
      [num_rows_padded, num_cols_padded]');
  matrix [num_rows_padded, num_cols_padded] f = gp_inv_rfft2(
    z, rep_matrix(mu, num_rows_padded, num_cols_padded), rfft2_cov);
}
model {
  to_vector(z) ~ std_normal();
  mu ~ student_t(2, 0, 1);
  sigma ~ student_t(2, 0, 1);
  kappa ~ student_t(2, 0, 1);
  for (i in 1:num_rows) {
    for (j in 1:num_cols) {
      if (frequency[i, j] >= 0) {
        frequency[i, j] ~ neg_binomial_2(exp(f[i, j]), 1 / kappa);
      }
    }
  }
}
}

```

We use an implicit uniform prior for `log_length_scale` such that no corresponding statement appears in the `model` block.

We fitted the model to 80% of the quadrants chosen uniformly at random using the below Python code, withholding the remainder for evaluation. Held-out data are encoded as -1 in the Stan model.

```

>>> from gptools.stan import get_include
>>> import numpy as np
>>> frequency = np.loadtxt("tachve.csv", delimiter = ",", dtype = int)
>>> num_rows, num_cols = frequency.shape
>>> padding = 10
>>> train_mask = np.random.binomial(1, 0.8, frequency.shape)
>>> data = {

```

```

...   "num_rows": num_rows,
...   "num_rows_padded": num_rows + padding,
...   "num_cols": num_cols,
...   "num_cols_padded": num_cols + padding,
...   "frequency": np.where(train_mask, frequency, -1),
... }
>>> model = compile_model(stan_file = "trees.stan",
...   stanc_options = {"include-paths": get_include()})
>>> fit = model.sample(data)
>>> print(fit.diagnose())

```

Processing csv files: ...

Checking sampler transitions treedepth.
Treedepth satisfactory for all transitions.

Checking sampler transitions for divergences.
No divergent transitions found.

Checking E-BFMI - sampler transitions HMC potential energy.
E-BFMI satisfactory.

Effective sample size satisfactory.

Split R-hat values satisfactory all parameters.

Processing complete, no problems detected.

Despite the noisy, masked observations, the model was able to learn a smooth estimate of the density of trees, as shown in panel (c). The posterior median of the correlation length of 60 m was well below the padding of 200 m introduced to attenuate the effect of periodic boundary conditions, as shown in panel (b). We used a scaled mean-squared error (SMSE) to evaluate the model on the held-out data and compare it with the simpler approach of smoothing the data with a two-dimensional Gaussian filter. The SMSE is

$$S(\mathbf{y}, \hat{\mathbf{y}} = \exp \hat{\mathbf{f}}) = \frac{1}{m} \sum_{j=1}^m \frac{(y_j - \exp \hat{f}_j)^2}{\max(y_j, 1)},$$

where the sum is over m test points and $\hat{\mathbf{f}}$ is the posterior median of the latent GP. We divided each term by the observed count y_j (or one if the count was zero) to ensure the measure was not dominated by large counts because the sampling variance of a Poisson count process (without overdispersion) is equal to its mean.

A simple method to estimate the number of trees in held-out quadrants is to apply a Gaussian filter to the data and compare the two methods. The Gaussian filter estimate is

$$\hat{\mathbf{y}}_\lambda = \frac{\mathbf{g}_\lambda * (\mathbf{b} \circ \mathbf{y})}{\mathbf{g}_\lambda * \mathbf{b}},$$

where $*$ denotes convolution, \circ denotes the elementwise product, \mathbf{g}_λ is a Gaussian filter with smoothing scale λ , and \mathbf{b} is the binary mask indicating which data are available for training. Gaussian filters “blur” the data locally such that adjacent elements of the smoothed signal can inform one another (Lindeberg 1990). For large λ , estimates are approximated by the sample mean, and, for small λ , they are dominated by local noise. Panel (d) of Figure 4 shows the SMSE for the Gaussian filter as a function of smoothing scale and the SMSE achieved by the GP model. The latter achieves a lower SMSE than the former for all smoothing scales, illustrating the utility of GPs for modeling spatial effects. Unlike in Section 7.1, it was not possible to use the posterior predictive distribution for evaluation because the Gaussian filter is not a generative model.

8. Discussion

We implemented two popular approaches for scaling GPs to larger datasets in *Stan*: The sparse approximation with structured dependencies discussed in Section 3 and the exact Fourier approach in Section 4 which is applicable to data on a grid. For centered parameterizations, the likelihood can be evaluated or approximated directly. For non-centered parameterizations, we sample standard normal random variables \mathbf{z} and use the inverse transform to obtain a GP sample $\mathbf{f} = \phi^{-1}(\mathbf{z})$.

Given different parameterizations and approaches, which should be used in practice? As discussed in Section 6, a non-centered parameterization is appropriate if the data are weak, and a centered parameterization is preferable if the data are strong. Choosing the right parameterization ensures parameters are relatively uncorrelated under the posterior distribution which accelerates inference. For variational mean-field approximations, choosing the right parameterization is even more important: It affects the quality of the approximation, as discussed in Section 6. Most variational approaches use a centered parameterization (Hensman *et al.* 2013; Wu *et al.* 2022), and their approximations may be improved by considering non-centered parameterizations. If the data are very strong, the benefits of GPs may be outweighed by their complexity because the likelihood dominates the GP prior. If in doubt, we suggest using a non-centered parameterization, as we have done in Section 7, because GPs are typically employed when the data are not sufficiently informative for simpler approaches to succeed. Choosing an appropriate inference approach and between the many packages discussed in Section 1 depends on the data at hand. Here, we only consider general purpose packages that can handle arbitrary likelihoods and facilitate the use of GPs as parts of larger models, as shown in Figure 5. At one extreme, posterior samples can be obtained with most inference frameworks if the data are small. At the other, if the data do not fit in memory, evaluating the likelihood repeatedly as part of a Monte Carlo sampler is not feasible—even if the likelihood can be approximated. ADVI (e.g., using **GPyTorch** and **Pyro**) is a viable approach although at the cost of considering a narrower set of posteriors (Kucukelbir *et al.* 2017). If the data are low-dimensional and form a regular grid, the Fourier methods in Section 4 are suitable. Basis function approximations (currently implemented in the R package **brms** which builds on *Stan*) may be appropriate for low-dimensional data with irregular spacing. The implementation represents the GP as a linear superposition of eigenfunctions of the Laplace operator with Dirichlet boundary conditions (Riutort-Mayol *et al.* 2022). The approach is similar to the Fourier methods presented here because Fourier modes are eigenfunctions of the Laplace operator although with periodic boundary conditions. The unique advantage of

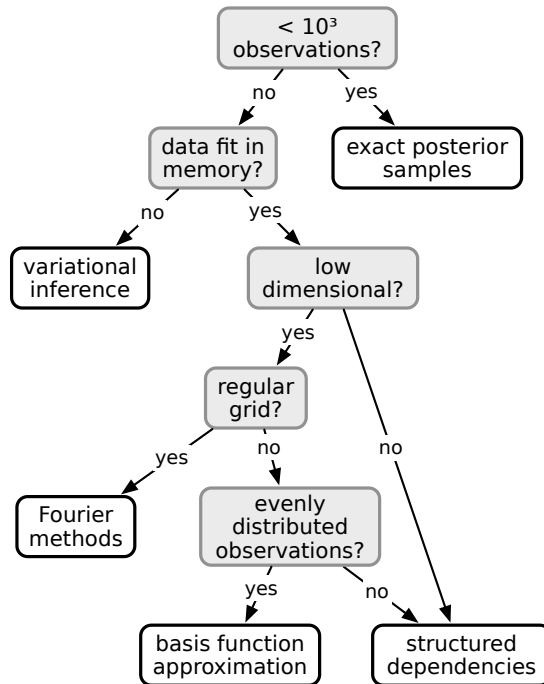


Figure 5: The optimal inference approach and package depends on the data.

Fourier methods is that the likelihood can be evaluated exactly in $\mathcal{O}(n \log n)$ if observations form a grid. In higher dimensional spaces, sparse approximations using structured dependencies can approximate the posterior, as discussed in Section 3. A meaningful performance comparison between packages is challenging because they seek to answer different questions (e.g., maximum marginal likelihood estimation or full posterior inference), employ different methods to answer the same question (e.g., ADVI, non-adaptive leapfrog sampler, Gibbs sampling, or adaptive No-U-Turn sampler), and use different programming languages with vastly different performance (e.g., compiled C++ ([International Organization for Standardization 2020](#)) or interpreted Python).

Padding may be required to attenuate the effect of periodic boundary conditions inherent to the fast Fourier transform. The necessary amount of padding depends on the kernel. We have found one to two correlation lengths ℓ to be sufficient for squared exponential and Matérn $3/2$ kernels (see Appendix A for details). However, the correlation length is often not known a priori, and finding the “right” amount of padding that appropriately balances performance and the need for non-periodic boundary conditions may be an iterative process. For example, we can start with a small amount of padding and increase it until the posterior stabilizes.

Fourier methods may also be appropriate if the density of observation points is relatively homogeneous. In particular, we may consider a latent GP \mathbf{g} on a grid and use it to predict the GP of interest \mathbf{f} at each observation point, i.e.,

$$p(\mathbf{f} | \mathbf{g}) = \prod_{j=1}^n p(f_j | \mathbf{g}).$$

This method reduces the computational cost because elements of \mathbf{f} are conditionally independent given \mathbf{g} at the regularly spaced “inducing points” ([Hensman *et al.* 2013](#)).

We hope that our library and the illustrations in Section 7 will accelerate the development of models employing GPs in Stan. Integrating GP approximations with Stan’s ecosystem, rather than developing a bespoke GP library, will allow practitioners to leverage the framework’s flexibility and the shared knowledge of the engaged Stan community.

Computational details

The results in this paper were obtained using Python 3.10.13, `cmdstanpy` 1.1.0, and `cmdstan` 2.33.0. All experiments were run on a single core of a 2020 MacBook Pro with an Apple Silicon M1 chip and 16 GB of RAM.

Acknowledgments

We thank Philip Greengard, Mike Lawrence, and Aki Vehtari for comments on the manuscript and Brian Ward for answering numerous questions about `cmdstanpy`.

References

- Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow IJ, Harp A, Irving G, Isard M, Jia Y, Józefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray DG, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker PA, Vanhoucke V, Vasudevan V, Viégas FB, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y, Zheng X (2016). “**TensorFlow**: Large-Scale Machine Learning on Heterogeneous Distributed Systems.” *arXiv 1603.04467*, arXiv.org E-Print Archive. doi:10.48550/arxiv.1603.04467.
- Ambikasaran S, Foreman-Mackey D, Greengard L, Hogg DW, O’Neil M (2015). “Fast Direct Methods for Gaussian Processes.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **38**, 252–265. doi:10.1109/tpami.2015.2448083.
- Betancourt M (2018). “A Conceptual Introduction to Hamiltonian Monte Carlo.” *arXiv 1701.02434*, arXiv.org E-Print Archive. doi:10.48550/arxiv.1701.02434.
- Bezanson J, Edelman A, Karpinski S, Shah VB (2017). “Julia: A Fresh Approach to Numerical Computing.” *SIAM Review*, **59**(1), 65–98. doi:10.1137/141000671.
- Bingham E, Chen JP, Jankowiak M, Obermeyer F, Pradhan N, Karaletsos T, Singh R, Szerlip P, Horsfall P, Goodman ND (2019). “**Pyro**: Deep Universal Probabilistic Programming.” *Journal of Machine Learning Research*, **20**(28), 1–6. URL <http://jmlr.org/papers/v20/18-403.html>.
- Bishop CM (2006). *Pattern Recognition and Machine Learning*. Springer-Verlag.
- Bivand R (2022). “R Packages for Analyzing Spatial Data: A Comparative Case Study with Areal Data.” *Geographical Analysis*, **54**(3), 488–518. doi:10.1111/gean.12319.

- Borovitskiy V, Terenin A, Mostowsky P, Deisenroth M (2020). “Matérn Gaussian Processes on Riemannian Manifolds.” In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, volume 33, pp. 12426–12437. URL <https://proceedings.neurips.cc/paper/2020/hash/92bf5e6240737e0326ea59846a83e076-Abstract.html>.
- Bürkner PC (2017). “**brms**: An R Package for Bayesian Multilevel Models Using Stan.” *Journal of Statistical Software*, **80**(1), 1–28. doi:10.18637/jss.v080.i01.
- Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li P, Riddell A (2017). “Stan: A Probabilistic Programming Language.” *Journal of Statistical Software*, **76**(1), 1–32. doi:10.18637/jss.v076.i01.
- Condit R, Perez R, Aguilar S, Lao S, Foster R, Hubbell SP (2019). “Complete Data From the Barro Colorado 50-Ha Plot: 423,617 Trees, 35 Years.” Dryad Data Platform. doi:10.15146/5xcp-0d46.
- Dancik GM, Dorman KS (2008). “**mlegp**: Statistical Analysis for Computer Models of Biological Systems Using R.” *Bioinformatics*, **24**(17), 1966–1967. doi:10.1093/bioinformatics/btn329.
- Deisenroth MP, Fox D, Rasmussen CE (2015). “Gaussian Processes for Data-Efficient Learning in Robotics and Control.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **37**(2), 408–423. doi:10.1109/tpami.2013.218.
- Duvenaud DK (2014). *Automatic Model Construction with Gaussian Processes*. Ph.D. thesis, University of Cambridge. URL <https://www.repository.cam.ac.uk/items/a54f0711-f777-4ee1-a685-3b76dae5ad91>.
- Fairbrother J, Nemeth C, Rischard M, Brea J, Pinder T (2022). “**GaussianProcesses.jl**: A Nonparametric Bayes Package for the Julia Language.” *Journal of Statistical Software*, **102**, 1–36. doi:10.18637/jss.v102.i01.
- Finley AO, Banerjee S, Carlin BP (2007). “**spBayes**: An R Package for Univariate and Multivariate Hierarchical Point-Referenced Spatial Models.” *Journal of Statistical Software*, **19**(4). doi:10.18637/jss.v019.i04.
- Flaxman S, Wilson A, Neill D, Nickisch H, Smola A (2015). “Fast Kronecker Inference in Gaussian Processes with Non-Gaussian Likelihoods.” In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pp. 607–616. URL <https://proceedings.mlr.press/v37/flaxman15.html>.
- Gabry J, Češnovar R, Johnson A (2023). **cmdstanr**: R Interface to **CmdStan**. R package version 0.8.1, URL <https://mc-stan.org/cmdstanr/>.
- Gardner JR, Pleiss G, Bindel D, Weinberger KQ, Wilson AG (2018). “**GPyTorch**: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration.” In *NIPS’18: Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 7587–7597. Curran Associates, Red Hook. URL <https://dl.acm.org/doi/10.5555/3327757.3327857>.
- Gelman A, Carlin JB, S SH, Dunson DB, Vehtari A, Rubin DB (2013). *Bayesian Data Analysis*. Chapman & Hall/CRC. doi:10.1201/b16018.

- Gelman A, Jakulin A, Pittau MG, Su YS (2008). “A Weakly Informative Default Prior Distribution for Logistic and Other Regression Models.” *Annals of Applied Statistics*, **2**(4), 1360–1383. doi:10.1214/08-aoas191.
- Gillies S, Van der Wel C, Van den Bossche J, Taves MW, Arnott J, Ward BC (2023). *Shapely, Version 2.0.2*. doi:10.5281/zenodo.5597138.
- Gómez-Rubio V (2020). *Bayesian Inference with INLA*. Chapman & Hall. doi:10.1201/9781315175584.
- Greengard P (2021). “Efficient Fourier Representations of Families of Gaussian Processes.” *arXiv 2109.14081*, arXiv.org E-Print Archive. doi:10.48550/arxiv.2109.14081.
- Handcock MS, Stein ML (1993). “A Bayesian Analysis of Kriging.” *Technometrics*, **35**(4), 403–410. doi:10.1080/00401706.1993.10485354.
- Harris CR, Millman KJ, Van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C, Oliphant TE (2020). “Array Programming with NumPy.” *Nature*, **585**(7825), 357–362. doi:10.1038/s41586-020-2649-2.
- Hensman J, Durrande N, Solin A (2017). “Variational Fourier Features for Gaussian Processes.” *Journal of Machine Learning Research*, **18**(1), 5537–5588. URL <https://jmlr.org/papers/v18/16-579.html>.
- Hensman J, Fusi N, Lawrence ND (2013). “Gaussian Processes for Big Data.” In *Uncertainty in Artificial Intelligence*, volume 29, pp. 282–290. URL <https://dl.acm.org/doi/10.5555/3023638.3023667>.
- Hoffman MD, Gelman A (2014). “The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo.” *Journal of Machine Learning Research*, **15**(1), 1593–1623. URL <http://jmlr.org/papers/v15/hoffman14a.html>.
- Hoffmann T, Peel L, Lambiotte R, Jones NS (2020). “Community Detection in Networks without Observing Edges.” *Science Advances*, **6**(4), eaav1478. doi:10.1126/sciadv.aav1478.
- International Organization for Standardization (2020). *ISO/IEC 14882:2020: Programming Languages – C++*. URL <https://www.iso.org/standard/79358.html>.
- Krige DG (1951). “A Statistical Approach to Some Basic Mine Valuation Problems on the Witwatersrand.” *Journal of the Southern African Institute of Mining and Metallurgy*, **52**(6), 119–139. URL https://hdl.handle.net/10520/AJA0038223X_4792.
- Kucukelbir A, Tran D, Ranganath R, Gelman A, Blei DM (2017). “Automatic Differentiation Variational Inference.” *Journal of Machine Learning Research*, **18**(1), 430–474. URL <https://jmlr.org/papers/v18/16-107.html>.
- Lindeberg T (1990). “Scale-Space for Discrete Signals.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **12**(3), 234–254. doi:10.1109/34.49051.

- Matthews AGDG, van der Wilk M, Nickson T, Fujii K, Boukouvalas A, León-Villagrà P, Ghahramani Z, Hensman J (2017). “**GPflow**: A Gaussian Process Library Using **TensorFlow**.” *Journal of Machine Learning Research*, **18**(40), 1–6. URL <http://jmlr.org/papers/v18/16-537.html>.
- Morris M, Wheeler-Martin K, Simpson D, Mooney SJ, Gelman A, DiMaggio C (2019). “Bayesian Hierarchical Spatial Models: Implementing the Besag York Mollié Model in Stan.” *Spatial and Spatio-Temporal Epidemiology*, **31**, 100301. doi:10.1016/j.sste.2019.100301.
- Neal RM (1997). “Monte Carlo Implementation of Gaussian Process Models for Bayesian Regression and Classification.” *arXiv 9702*, arXiv.org E-Print Archive. doi:10.48550/arXiv.physics/9701026.
- Papaspiliopoulos O, Roberts GO, Sköld M (2007). “A General Framework for the Parametrization of Hierarchical Models.” *Statistical Science*, **22**(1), 59–73. doi:10.1214/088342307000000014.
- Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Kopf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, Chintala S (2019). “**PyTorch**: An Imperative Style, High-Performance Deep Learning Library.” In *Advances in Neural Information Processing Systems*, volume 32, pp. 8024–8035. URL <https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>.
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP (2007). *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press.
- Python Software Foundation (2023). *Python Documentation, Version 3.10.13*. URL <https://docs.python.org/3.10/>.
- Rasmussen CE, Nickisch H (2010). “Gaussian Processes for Machine Learning (**GPML**) Toolbox.” *Journal of Machine Learning Research*, **11**(100), 3011–3015. URL <http://jmlr.org/papers/v11/rasmussen10a.html>.
- Rasmussen CE, Williams CKI (2006). *Gaussian Processes for Machine Learning*. MIT Press. doi:10.7551/mitpress/3206.001.0001.
- R Core Team (2024). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Riutort-Mayol G, Bürkner PC, Andersen MR, Solin A, Vehtari A (2022). “Practical Hilbert Space Approximate Bayesian Gaussian Processes for Probabilistic Programming.” *Statistics and Computing*, **33**(1), 17. doi:10.1007/s11222-022-10167-2.
- Roberts S, Osborne M, Ebdem M, Reece S, Gibson N, Aigrain S (2013). “Gaussian Processes for Time-Series Modelling.” *Philosophical Transactions of the Royal Society A*, **371**(1984), 20110550. doi:10.1098/rsta.2011.0550.
- Rubin DB (1981). “The Bayesian Bootstrap.” *Annals of Statistics*, **9**(1), 130–134. doi:10.1214/aos/1176345338.

- Sheffield Machine Learning Group (2012). **GPY**: *A Gaussian Process Framework in Python*. URL <https://github.com/SheffieldML/GPy>.
- Stan Development Team (2023a). **cmdstan** *User’s Guide, Version 2.33*. URL <https://mc-stan.org/docs/cmdstan-guide/>.
- Stan Development Team (2023b). **cmdstanpy**: *Python Interface to cmdstan*. URL <https://github.com/stan-dev/cmdstanpy>.
- Talts S, Betancourt M, Simpson D, Vehtari A, Gelman A (2018). “Validating Bayesian Inference Algorithms with Simulation-Based Calibration.” *arXiv 1804.06788*, arXiv.org E-Print Archive. doi:10.48550/arxiv.1804.06788.
- The MathWorks Inc (2022). *MATLAB – The Language of Technical Computing*. Natick. URL <https://www.mathworks.com/products/matlab/>.
- Tipping M, Bishop C (2002). “Bayesian Image Super-Resolution.” In *Advances in Neural Information Processing Systems*, volume 15. URL https://proceedings.neurips.cc/paper_files/paper/2002/file/88bfcf02e7f554f9e9ea350b699bc6a7-Paper.pdf.
- Trangucci R, Betancourt M, Vehtari A (2016). “Prior Formulation for Gaussian Process Hyperparameters.” In *Advances in Neural Information Processing Systems*. Practical Bayesian Nonparametrics Workshop.
- Transport for London (2019). “London Underground Passenger Counts Data.” URL <http://crowding.data.tfl.gov.uk/>.
- Transport for London (2020). “Central Line Working Timetable.” URL <https://content.tfl.gov.uk/cen-wtt-70.pdf>.
- Transport for London (2022). “Transport for London Unified API.” URL <https://api.tfl.gov.uk>.
- Van den Bossche J, Jordahl K, Fleischmann M, McBride J, Wasserman J, Richards M, Garcia Badaracco A, Snow AD, Tratner J, Gerard J, Ward B, Perry M, Farmer C, Hjelle GA, Taves M, ter Hoeven E, Cochran M, rraymondgh, Gillies S, Caria G, Culbertson L, Bell R, Bartos M, Eubank N, sangarshanan, Flavin J, Rey S, Gardiner J, maxalbert, Bilogur A (2023). **geopandas/geopandas**, *Version 0.14.1*. doi:10.5281/zenodo.2650956.
- Vehtari A, Gelman A, Simpson D, Carpenter B, Bürkner PC (2021). “Rank-Normalization, Folding, and Localization: An Improved \hat{R} for Assessing Convergence of MCMC.” *Bayesian Analysis*, **16**(2), 667–718. doi:10.1214/20-ba1221.
- Wu L, Pleiss G, Cunningham JP (2022). “Variational Nearest Neighbor Gaussian Process.” In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 24114–24130. URL <https://proceedings.mlr.press/v162/wu22h.html>.

A. Effect of padding for Fourier methods

We considered a simulation study to examine the effect of periodic boundary conditions inherent to Fourier methods and assess the amount of padding required to balance model misspecification and performance concerns in two steps. First, we generated $m = 100$ synthetic datasets each comprising $n = 128$ observations on an integer grid according to the benchmark model in Equation 3 with marginal kernel scale $\sigma = 1$ and observation noise $\kappa = 1$. We used a correlation length $\ell = 16$ large enough for periodic boundary conditions to have an effect on the inference. For each dataset, we fitted the standard non-centered model, i.e., the true generative model, and Fourier-based Gaussian processes with varying amounts of padding w . We evaluate both models by holding out the last data point y_n , which should be most severely affected by periodic boundary conditions, and approximating the log posterior density $\log p(f_n | y_{<n})$ of the corresponding element of the latent GP using a Gaussian kernel density estimator (Bishop 2006, Chapter 2.5.1). Even a small amount of padding, such as one correlation length, is sufficient to attenuate the effect of periodic boundary conditions,

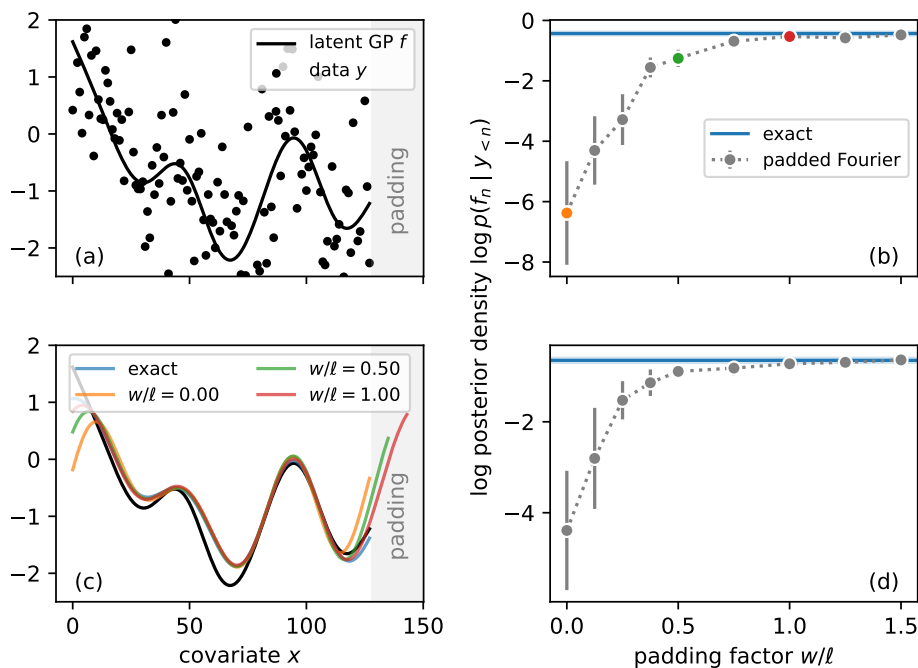


Figure 6: A small amount of padding is sufficient to attenuate the effect of periodic boundary conditions. Panel (a) shows a realization of the benchmark model in Equation 3, and panel (c) shows the true latent GP f together with posterior means for different models, including the true generative model and Fourier methods with varying amounts of padding w . The effect of periodic boundary conditions is stark for the fit without padding, but it disappears quickly with increasing padding. Panels (b) and (d) show the log posterior density evaluated at the true f_n , which corresponds to the held-out data point y_n , as a function of padding for squared exponential and Matérn $3/2$ kernels, respectively. The horizontal line corresponds to the log posterior density obtained using the true generative model. Padding with one length scale ℓ is sufficient to approximate the true model well.

as shown in panels (b) and (d) of Figure 6 for squared exponential and Matérn $3/2$ kernels, respectively.

Simulation-based calibration is a technique to validate a Bayesian inference pipeline (Talts, Betancourt, Simpson, Vehtari, and Gelman 2018). For synthetic data generated from the model, the rank of the true parameter value among posterior samples should have a uniform distribution. For each combination of the different paddings and two kernels, we evaluated the rank across m synthetic datasets and evaluated the p value of the Kolmogorov-Smirnov test by comparing with a discrete uniform reference distribution. If no padding is used, the null hypothesis that the ranks are uniform can be confidently rejected (p value $< 10^{-3}$), but the ranks are not inconsistent with a uniform distribution for $w > \ell$ (p value > 0.3 at $w = \ell$ in our simulations).

B. Kernels in the real and Fourier domains

B.1. Squared exponential kernel

The non-periodic squared exponential kernel is defined as

$$k(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2\ell^2}\right),$$

where σ is the marginal scale and ℓ is the correlation length. Its discrete power spectrum on a periodic domain of size L is

$$\tilde{k}_\xi = \frac{\sqrt{2\pi n}\sigma^2\ell}{L} \exp\left(-2\left(\frac{\pi\xi\ell}{L}\right)^2\right),$$

where n is the number of grid points and $\xi \in [0..n - 1]$ is the discrete frequency.

B.2. Matérn kernel

The non-periodic Matérn kernel is defined as

$$k_\nu(x, x') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \zeta^\nu K_\nu(\zeta),$$

where $\zeta = \frac{\sqrt{2\nu}|x - x'|}{\ell}$

is a rescaled distance, ν is a smoothness parameter, Γ denotes the gamma function, $|x - x'|$ is the Euclidean distance between x and x' , and K_ν denotes the modified Bessel function of the second kind. For $\nu = 3/2$ and $\nu = 5/2$, the kernel simplifies to

$$k_{3/2} = \sigma^2 \left(1 + \frac{\sqrt{3}|x - x'|}{\ell}\right) \exp\left(-\frac{\sqrt{3}|x - x'|}{\ell}\right)$$

$$k_{5/2} = \sigma^2 \left(1 + \frac{\sqrt{5}|x - x'|}{\ell} + \frac{5|x - x'|^2}{3\ell^2}\right) \exp\left(-\frac{\sqrt{5}|x - x'|}{\ell}\right).$$

It reduces to the Laplace kernel for $\nu = 1/2$. Its discrete power spectrum on a periodic domain of size L is

$$\tilde{k}_\xi = \sigma^2 \frac{n\ell}{L} \left(\frac{2\pi}{\nu}\right)^{p/2} \frac{\Gamma(\nu + \frac{p}{2})}{\Gamma(\nu)} \left(1 + \frac{2(\pi\ell\xi)^2}{\nu L^2}\right)^{-(\nu+p/2)},$$

where p is the dimensionality of the space, n is the number of grid points, and $\xi \in [0..n - 1]$ is the discrete frequency.

Affiliation:

Till Hoffmann, Jukka-Pekka Onnela
Harvard T. H. Chan School of Public Health
677 Huntington Avenue
Boston, MA, 02115, United States of America
E-mail: thoffmann@hsph.harvard.edu, onnela@hsph.harvard.edu
URL: <https://tillahoffmann.github.io/>,
<https://www.hsph.harvard.edu/onnela-lab/>