

Learning from Summarized Data: Gaussian Process Regression with Sample Quasi-Likelihood

Yuta Shikuri

Tokio Marine Holdings, Inc.
Tokyo, Japan
shikuriyuta@gmail.com

Abstract

Gaussian process regression is a powerful Bayesian nonlinear regression method. Recent research has enabled the capture of many types of observations using non-Gaussian likelihoods. To deal with various tasks in spatial modeling, we benefit from this development. Difficulties still arise when we can only access summarized data consisting of representative features, summary statistics, and data point counts. Such situations frequently occur primarily due to concerns about confidentiality and management costs associated with spatial data. This study tackles learning and inference using only summarized data within the framework of Gaussian process regression. To address this challenge, we analyze the approximation errors in the marginal likelihood and posterior distribution that arise from utilizing representative features. We also introduce the concept of sample quasi-likelihood, which facilitates learning and inference using only summarized data. Non-Gaussian likelihoods satisfying certain assumptions can be captured by specifying a variance function that characterizes a sample quasi-likelihood function. Theoretical and experimental results demonstrate that the approximation performance is influenced by the granularity of summarized data relative to the length scale of covariance functions. Experiments on a real-world dataset highlight the practicality of our method for spatial modeling.

1 Introduction

Gaussian process regression is a Bayesian nonlinear regression method that can handle many types of observation data (Rasmussen and Williams 2006). A key application of Gaussian process regression is spatial modeling (Cressie 1993) across various fields, such as geology, agriculture, marketing, and public health. For example, it is used to estimate the distribution of soil moisture and regional medical expenses. Specifying likelihood functions according to these tasks impacts prediction performance. While a conventional choice is a Gaussian likelihood that is conjugate with a Gaussian process prior, recent research has focused extensively on applying Gaussian process regression to non-Gaussian likelihoods. The challenge of specifying non-Gaussian likelihoods lies in posterior distributions that lack closed-form expressions. Gaussian approximations of the posterior distribution are commonly used to overcome this challenge.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Laplace approximation represents the posterior distribution as a Gaussian centered at the maximum a posteriori estimate (Williams and Barber 1998). Expectation propagation finds the parameters in a Gaussian that approximate the posterior distribution through an iterative scheme of matching moments (Minka 2001). Variational inference maximizes a lower bound on the marginal likelihood via a Gaussian approximation (Sheth, Wang, and Kharon 2015). These methods for specifying non-Gaussian likelihoods allow us to handle various types of observations.

Difficulties have been encountered when only summarized data is available. This situation arises particularly in the context of spatial modeling. Due to privacy concerns associated with location-specific observation information, spatial data is often summarized to include representative locations, summary statistics, and counts. Too fine locations might identify individuals even if direct identifiers are removed. Hence, the granularity of summarized data tends to be intentionally coarse to ensure that personal information (e.g., financial assets, interests, purchase activity, and medical history) is not exposed. In addition to protecting individual privacy, management costs necessitate aggregating data by specific units (e.g., cities, branches, hospitals, and schools). To analyze regional trends, the latitude and longitude pair for each unit is linked to the corresponding summary statistics. Consequently, techniques are required for spatial modeling given only summarized data.

In this study, we address learning and inference from summarized data in Gaussian process regression. Figure 1 illustrates the overview of our approach. For input approximation using representative features, we demonstrate the theoretical errors in the marginal likelihood and posterior distribution within the Gaussian process regression framework. We also propose the concept of sample quasi-likelihood. This approach enables the computation of the marginal likelihood and posterior distribution using only summarized data, offering straightforward implementation and low computational complexity. The sample quasi-likelihood is defined by a variance function. Specifying this variance function allows for capturing non-Gaussian likelihoods that satisfy certain conditions. Theoretical and experimental results indicate that the approximation performance strongly depends on the granularity of summarized data relative to the length scale of covariance functions. Experiments using real-world

Complete Data		Summarized Data		
Input	Output	Centroid	Statistic	Count
\mathbf{x}_1	y_1			
\mathbf{x}_2	y_2	\mathbf{z}_1	\bar{y}_1	n_1
\mathbf{x}_3	y_3			
\mathbf{x}_4	y_4	\mathbf{z}_2	\bar{y}_2	n_2
\mathbf{x}_5	y_5			

Learning from $(\mathbf{z}_i, \bar{y}_i, n_i)_{i=1}^2$




Figure 1: Overview of our approach. We address Gaussian process regression using only summarized data, which includes representative features, summary statistics, and data point counts. Our approach involves the input approximation and the introduction of sample quasi-likelihood under specific assumptions. Specifying a variance function that characterizes a sample quasi-likelihood allows us to capture non-Gaussian likelihoods.

spatial data demonstrate the usability of our approach for supervised learning tasks when only summarized data is available.

2 Related Work

There are existing methods capable of handling aggregated outputs. Composite likelihood is derived by multiplying a collection of component likelihoods (Besag 1975; Varin, Reid, and Firth 2011). This approach is advantageous for dealing with aggregated outputs while retaining some properties of the full likelihood. Additionally, synthetic likelihood (Wood 2010; Price et al. 2018) and approximate Bayesian computation (Beaumont 2019) are also promising. These methods approximate the posterior distribution by comparing summary statistics with simulated data generated from a tractable probability model. Outputs sometimes take distributional forms, such as random lists, intervals, and histograms. Beranger, Lin, and Sisson 2022 constructed likelihood functions for these forms. Considering that a Gaussian process has parameters corresponding to its inputs, handling outputs not directly linked to inputs becomes crucial. Some studies (Law et al. 2018; Tanaka et al. 2019) have provided methods for learning and inference from aggregated outputs in Gaussian process regression, assuming that input data points or their distribution are available.

3 Preliminaries

3.1 Notations

Appendix A provides the list of symbols. Let $\mathcal{X} \subset \mathbb{R}^d$ denote the domain of inputs, and \mathcal{Y} denote the domain of outputs. The symbol of \mathbf{K}_{\cdot} is a Gram matrix. The (i, j) -th element of the Gram matrix for (\mathbf{A}, \mathbf{B}) is the return value of a covariance function when the inputs are the i -th row of \mathbf{A} and the j -th row of \mathbf{B} . The matrix with the opposite order of indices is the transposed matrix (e.g., $\mathbf{K}_{*\mathbf{f}} = \mathbf{K}_{\mathbf{f}*}^\top$). Let

\mathcal{N} denote the probability density function of Gaussian distribution, the vertical line with equality denote substitution (e.g., $p(\mathbf{y} | \mathbf{f})|_{\mathbf{f}=\mathbf{w}_{\text{fu}}\mathbf{u}}$), O denote the Landau symbol, and F denote the upper cumulative distribution function of the chi-square distribution given by

$$F(a, b) \equiv \left(\int_0^\infty t^{\frac{a}{2}-1} \exp(-t) dt \right)^{-1} \int_{\frac{b}{2}}^\infty t^{\frac{a}{2}-1} \exp(-t) dt$$

for any $a \in \mathbb{N}$ and $b \in [0, \infty)$. The Landau symbol in matrix operations is applied to each matrix element. We assume that symmetric Gram matrices are positive-definite, that likelihood functions and their gradients are differentiable and bounded, and that covariance functions are symmetric and continuous.

3.2 Gaussian Process Regression

In this subsection, we describe the basic framework of Gaussian process regression. A more detailed introduction is presented in (Rasmussen and Williams 2006).

A Gaussian process $f \sim \mathcal{GP}(\tau(\cdot), k(\cdot, \cdot))$ is a distribution over functions characterized by a mean function $\tau : \mathcal{X} \rightarrow \mathbb{R}$ and a covariance function $k : \mathcal{X} \times \mathcal{X} \rightarrow (0, \infty)$. A stochastic process $\{f(\mathbf{x}) | \mathbf{x} \in \mathcal{X}\}$ is a Gaussian process if and only if the random variables $\{f(\mathbf{x}) | \mathbf{x} \in \mathcal{X}'\}$ for any finite set $\mathcal{X}' \subseteq \mathcal{X}$ follow a multivariate normal distribution. For simplicity, we take $\tau(\cdot) = 0$. Given inputs $\mathbf{X} \equiv (\mathbf{x}_i)_{i=1}^n$ with $\mathbf{x}_i \in \mathcal{X}$ and outputs $\mathbf{y} \equiv (y_i)_{i=1}^n$ with $y_i \in \mathcal{Y}$, the hyperparameters of the covariance function are learned to maximize the log marginal likelihood defined as

$$\mathcal{L} \equiv \log \int_{\mathbb{R}^n} p(\mathbf{y} | \mathbf{f}) p(\mathbf{f}) d\mathbf{f}, \quad (1)$$

where $\mathbf{f} \equiv (f(\mathbf{x}_i))_{i=1}^n$, $p(\mathbf{f}) \equiv \mathcal{N}(\mathbf{f}; \mathbf{0}, \mathbf{K}_{\mathbf{ff}})$, $p(\mathbf{y} | \mathbf{f})$ is a likelihood function of \mathbf{f} , and $\mathbf{K}_{\mathbf{ff}}$ is the Gram matrix of (\mathbf{X}, \mathbf{X}) . The complexity of optimizing \mathcal{L} given a Gaussian

Function Name	$k(\mathbf{x}, \mathbf{x}')$	$\zeta_1(\alpha, \mathbf{z}, \mathbf{z}')$	$\zeta_2(\alpha, \mathbf{z}, \mathbf{x}^*)$
Laplacian	$\exp(- \mathbf{x} - \mathbf{x}')$	$1 - \exp(-2\alpha)$	$1 - \exp(-\alpha)$
Gaussian	$\exp(-\frac{1}{2} \mathbf{x} - \mathbf{x}' ^2)$	$1 - \exp(-2\alpha(\mathbf{z} - \mathbf{z}' + \alpha))$	$1 - \exp(-\alpha(\mathbf{z} - \mathbf{x}^* + \frac{3}{2}\alpha))$

Table 1: Errors of covariance functions. Given $\mathbf{z}, \mathbf{z}', \mathbf{x}^* \in \mathcal{X}$, for any $\alpha \in (0, \infty)$ and $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, $\max\{|\mathbf{x} - \mathbf{z}|, |\mathbf{x}' - \mathbf{z}'|\} < \alpha \Rightarrow |k(\mathbf{x}, \mathbf{x}') - k(\mathbf{z}, \mathbf{z}')| < \zeta_1(\alpha, \mathbf{z}, \mathbf{z}')$ and $|k(\mathbf{x}, \mathbf{x}^*) - k(\mathbf{z}, \mathbf{x}^*)| < \zeta_2(\alpha, \mathbf{z}, \mathbf{x}^*)$. See appendix B.1.

likelihood is $O(n^3)$. For new inputs $\mathbf{X}_* \equiv (\mathbf{x}_i^*)_{i=1}^{n_*}$ with $\mathbf{x}_i^* \in \mathcal{X}$, the posterior distribution of $\mathbf{f}_* \equiv (\mathbf{f}(\mathbf{x}_i^*))_{i=1}^{n_*}$ is

$$p(\mathbf{f}_* | \mathbf{y}) \equiv \exp(-\mathcal{L}) \int_{\mathbb{R}^n} p(\mathbf{f}_* | \mathbf{f}) p(\mathbf{y} | \mathbf{f}) p(\mathbf{f}) d\mathbf{f}, \quad (2)$$

where $\mathbf{K}_{\mathbf{f}_*}$ is the Gram matrix of $(\mathbf{X}, \mathbf{X}_*)$, \mathbf{K}_{**} is that of $(\mathbf{X}_*, \mathbf{X}_*)$, and $p(\mathbf{f}_* | \mathbf{f}) \equiv \mathcal{N}(\mathbf{f}_*; \mathbf{K}_{*f} \mathbf{K}_{\mathbf{ff}}^{-1} \mathbf{f}, \mathbf{K}_{**} - \mathbf{K}_{*f} \mathbf{K}_{\mathbf{ff}}^{-1} \mathbf{K}_{f*})$.

4 Input Approximation

To clarify the discussion in this study, we define summarized data and the error $\beta \in (0, \infty)$ of a covariance function. Assume that data points $(\mathbf{x}_i, y_i)_{i=1}^n$ are associated with assignments $\omega \equiv (\omega_i)_{i=1}^n$ with $\omega_i \in \{1, \dots, m\}$. Let summarized data be represented as a set of tuples $(z_j, \bar{y}_j, n_j)_{j=1}^m$, where each $z_j \in \mathcal{X}$ is a representative feature, each $\bar{y}_j \in \mathbb{R}$ is a summary statistic that depends only on the outputs y_i with $\omega_i = j$, and each $n_j > 0$ is the number of indices such that $\omega_i = j$. Let β satisfy the following conditions:

- $|k(\mathbf{x}_i, \mathbf{x}_j) - k(z_{\omega_i}, z_{\omega_j})| < \beta$ for all $1 \leq i \leq j \leq n$.
- $|k(\mathbf{x}_i, \mathbf{x}_j^*) - k(z_{\omega_i}, \mathbf{x}_j^*)| < \beta$ for all $1 \leq i \leq n$ and $1 \leq j \leq n_*$.

Even in situations where only summarized data is available, the range of complete data inputs is often accessible. This range allows us to evaluate the approximation errors of covariance functions when the inputs are representative features, as described in table 1. For Gaussian process regression, we can flexibly design many types of covariance functions (e.g., covariance functions equivalent to an infinitely wide deep network (Lee et al. 2018) or relying on non-Euclidean metric (Feragen, Lauze, and Hauberg 2014)). Lemma 4.1 guarantees that a small range improves the approximation accuracy for any covariance function.

Lemma 4.1. *Given $\mathbf{z}, \mathbf{z}', \mathbf{x}^* \in \mathcal{X}$, there exists $\alpha \in (0, \infty)$ such that $\max\{|\mathbf{x} - \mathbf{z}|, |\mathbf{x}' - \mathbf{z}'|\} < \alpha \Rightarrow \max\{|k(\mathbf{x}, \mathbf{x}') - k(\mathbf{z}, \mathbf{z}')|, |k(\mathbf{x}, \mathbf{x}^*) - k(\mathbf{z}, \mathbf{x}^*)|\} < \beta$, for any $\beta \in (0, \infty)$ and $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$.*

Proof. See appendix B.2. \square

Here we approximate the marginal likelihood and posterior distribution using representative features, and evaluate the approximation errors. Initially, we consider the case where the inputs are given by $(z_{\omega_i})_{i=1}^n$ and the parameters corresponding to the centroids $\mathbf{Z} \equiv (z_i)_{i=1}^m$ are represented by $\mathbf{u} \equiv (\mathbf{f}(z_i))_{i=1}^m$. In this case, the prior distribution and the likelihood function are $p(\mathbf{u}) \equiv \mathcal{N}(\mathbf{u}; \mathbf{0}, \mathbf{K}_{\mathbf{uu}})$ and $p(\mathbf{y} | \mathbf{W}_{\mathbf{fu}} \mathbf{u}) \equiv p(\mathbf{y} | \mathbf{f})|_{\mathbf{f}=\mathbf{W}_{\mathbf{fu}} \mathbf{u}}$, respectively, where

$\mathbf{K}_{\mathbf{uu}}$ is the Gram matrix of (\mathbf{Z}, \mathbf{Z}) , and $\mathbf{W}_{\mathbf{fu}}$ is $n \times m$ matrix with $[\mathbf{W}_{\mathbf{fu}}]_{ij} = 1$ if $\omega_i = j$; $[\mathbf{W}_{\mathbf{fu}}]_{ij} = 0$ otherwise. Using this prior distribution and likelihood function, the log marginal likelihood is

$$\mathcal{E} \equiv \log \int_{\mathbb{R}^m} p(\mathbf{y} | \mathbf{W}_{\mathbf{fu}} \mathbf{u}) p(\mathbf{u}) d\mathbf{u}. \quad (3)$$

The posterior distribution of \mathbf{f}_* is $p(\mathbf{f}_* | \mathbf{y}, \omega)$ defined as

$$\exp(-\mathcal{E}) \int_{\mathbb{R}^m} p(\mathbf{f}_* | \mathbf{W}_{\mathbf{fu}} \mathbf{u}) p(\mathbf{y} | \mathbf{W}_{\mathbf{fu}} \mathbf{u}) p(\mathbf{u}) d\mathbf{u}, \quad (4)$$

where $p(\mathbf{f}_* | \mathbf{W}_{\mathbf{fu}} \mathbf{u}) \equiv p(\mathbf{f}_* | \mathbf{f})|_{\mathbf{f}=\mathbf{W}_{\mathbf{fu}} \mathbf{u}}$. We evaluate their errors in comparison to Gaussian process regression with complete data. Let $p(\mathbf{f} | \mathbf{u})$ be defined as

$$\mathcal{N}(\mathbf{f}; \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{u}, \mathbf{K}_{\mathbf{ff}} - \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{K}_{\mathbf{uf}}), \quad (5)$$

where $\mathbf{K}_{\mathbf{fu}}$ is the Gram matrix of (\mathbf{X}, \mathbf{Z}) . As in eq. (1) and eq. (2), the marginal likelihood and posterior distribution of the original model contain $p(\mathbf{f}) = \int_{\mathbb{R}^m} p(\mathbf{f} | \mathbf{u}) p(\mathbf{u}) d\mathbf{u}$. Considering that eq. (3) and eq. (4) correspond to the case where $\mathbf{f} = \mathbf{W}_{\mathbf{fu}} \mathbf{u}$ holds, we proceed to analyze the integral of $p(\mathbf{f} | \mathbf{u})$. To facilitate this analysis, we introduce the following lemma.

Lemma 4.2. *Let γ denote the maximum absolute value of the elements in $\mathbf{W}_{\mathbf{fu}} - \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1}$. Then we have*

$$\mathbf{K}_{\mathbf{ff}} - \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{K}_{\mathbf{uf}} = O(\beta + m\beta\gamma). \quad (6)$$

Proof. See appendix B.3. \square

The error γ tends to become small when β is sufficiently small. In such cases, from lemma 4.1 and lemma 4.2, the integral of $p(\mathbf{f} | \mathbf{u})$ over the region within a certain distance from the point $\mathbf{f} = \mathbf{W}_{\mathbf{fu}} \mathbf{u}$ approaches 1 as the range of inputs becomes small. The following lemmas evaluate this dynamic. For simplicity, we assume $\gamma \neq 0$.

Lemma 4.3. *For $\delta_1 \in [0, \infty)$ and $\delta_2 \in [0, \delta_1]$, define*

$$\epsilon(\delta_1, \delta_2) \equiv F\left(m, \frac{\kappa(\delta_1 - \delta_2)^2}{\lambda_1} m\right) + F\left(n, \frac{\delta_2^2}{\lambda_2} n\right), \quad (7)$$

where $\kappa \equiv \inf_{\mathbf{u} \in \mathbb{R}^m; |(\mathbf{W}_{\mathbf{fu}} - \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1}) \mathbf{u}|=1} \frac{n}{m} |\mathbf{u}|^2$, λ_1 is the maximum eigenvalue of $\mathbf{K}_{\mathbf{uu}}$, and λ_2 is that of $\mathbf{K}_{\mathbf{ff}} - \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{K}_{\mathbf{uf}}$. Then we have

$$\int_{\mathbb{R}^n} \int_{\mathbb{R}^n \setminus R(\mathbf{u}, \delta_1)} p(\mathbf{f} | \mathbf{u}) p(\mathbf{u}) d\mathbf{f} d\mathbf{u} \leq \epsilon(\delta_1, \delta_2), \quad (8)$$

where $R(\mathbf{u}, \delta_1) \equiv \{\mathbf{f} \in \mathbb{R}^n \mid |\mathbf{f} - \mathbf{W}_{\mathbf{fu}} \mathbf{u}| \leq \sqrt{n} \delta_1\}$.

Proof. See fig. 2 and appendix B.4. \square

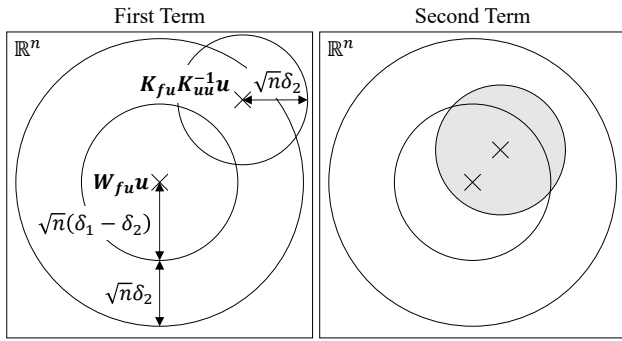


Figure 2: Sketch of the proof of lemma 4.3. The first and second terms in eq. (7) correspond respectively to the left and right sides. Evaluating the integral in eq. (8) requires analyzing $p(\mathbf{f} | \mathbf{u})$ over $\mathbb{R}^n \setminus R(\mathbf{u}, \delta_1)$, posing a significant challenge. Consequently, to assess the upper bound, we consider the hypersphere centered at the mean vector in $p(\mathbf{f} | \mathbf{u})$. The space of \mathbf{u} is divided based on whether $R(\mathbf{u}, \delta_1)$ encompasses the hypersphere. Gray indicates the space that is not integrated with respect to \mathbf{f} .

Lemma 4.4. *The following holds:*

$$\kappa^{-1} = O(m\gamma^2), \quad \lambda_2 = O(\beta + m\beta\gamma). \quad (9)$$

Proof. See appendix B.5. \square

Lemma 4.5. *Suppose $\xi \geq \xi_0$, where ξ_0 is the larger value satisfying $2\sqrt{\pi\xi_0} \exp(-\xi_0) = 1$. Then $F(m, \xi m)$ is monotonically decreasing with respect to m .*

Proof. See appendix B.6. \square

The percentiles of the chi-square distributions appearing in eq. (7) depend on κ and λ_2 . Lemma 4.4 shows that the integral in eq. (8) approaches zero as β and γ decrease. The evaluation of the integral behaves oppositely with respect to m in lemma 4.4 and lemma 4.5. While the integral decreases as δ_1 increases, \mathbf{f} within $R(\mathbf{u}, \delta_1)$ move away from $\mathbf{W}_{fu}\mathbf{u}$. Considering this behavior, we derive the error bounds for the marginal likelihood and posterior distribution.

Theorem 4.6. *Let η be defined as*

$$\eta \equiv \inf_{\xi \in [\xi_0, \infty)} (\sqrt{\xi\lambda_1\kappa^{-1}} + \sqrt{\xi\lambda_2} + F(m, \xi m)). \quad (10)$$

Given the hyperparameters and complete data, we have

$$\mathcal{L} - \mathcal{E} = O(\eta), \quad (11)$$

$$\|\mathbb{E}_{p(\mathbf{f}_* | \mathbf{y})}[\mathbf{f}_*] - \mathbb{E}_{p(\mathbf{f}_* | \mathbf{y}, \boldsymbol{\omega})}[\mathbf{f}_*]\| = O(\eta), \quad (12)$$

where $\|\cdot\|$ denote the norm of a vector.

Proof. See appendix B.7. \square

Theorem 4.6 demonstrates that the performance of input approximation improves as η decreases. While the first and second terms in η increase as ξ increases, the last term decreases. Lemma 4.4 suggests that smaller values of β and γ , relative to the square root of m , help prevent the first and




Figure 3: Behavior of η in a toy model. Let θ represent the length scale. Light blue: $\theta = 0.1$. Light green: $\theta = 1$. Pink: $\theta = 10$. The left and right figures correspond to the covariance functions $\exp(-\frac{1}{\theta}|x - x'|)$ and $\exp(-\frac{1}{2\theta^2}|x - x'|^2)$, respectively. In each figure, the vertical and horizontal axes represent η and m , respectively. The $n = 1000$ inputs and m representative features are equally spaced within $[0, 2\pi]$. The assignment ω_i of a data point x_i is the index j of the closest centroid z_j . Appendix B.8 explains the process of obtaining κ and η .

second terms from increasing. From lemma 4.5, the last term decreases as the summarized data becomes finer. Figure 3 shows that η decreases with finer summarized data. Furthermore, the length scale of the covariance functions significantly affects η . As the length scale increases, η decreases due to smaller β and γ . Note that this result does not directly prove the effect of the length scale on input approximation, as the proportional constants in eq. (11) and eq. (12) contain the hyperparameters.

5 Sample Quasi-Likelihood

Since the likelihood function $p(\mathbf{y} | \mathbf{W}_{fu}\mathbf{u})$ requires the outputs of complete data, the log marginal likelihood \mathcal{E} and the posterior distribution $p(\mathbf{f}_* | \mathbf{y}, \boldsymbol{\omega})$ still cannot be computed using only summarized data. Consequently, we replace the likelihood function with a function that excludes complete data outputs and incorporates summary statistics $\bar{\mathbf{y}} \equiv (\bar{y}_i)_{i=1}^m$. As one such function, we propose the concept of sample quasi-likelihood characterized by a variance function $v : \mathbb{R} \rightarrow (0, \infty)$ as follows:

Definition 5.1. A sample quasi-likelihood function $\bar{Q} : \mathcal{Y}^m \times \mathbb{R}^m \rightarrow \mathbb{R}$ is a function such that

$$\frac{\partial \bar{Q}(\bar{\mathbf{y}}, \mathbf{u})}{\partial \mathbf{u}} = \mathbf{V}_{\mathbf{uu}}^{-1}(\bar{\mathbf{y}} - \mathbf{u}), \quad (13)$$

where $\mathbf{V}_{\mathbf{uu}} \equiv \text{diag}(n_1^{-1}v(\bar{y}_1), \dots, n_m^{-1}v(\bar{y}_m))$.

Suppose that the prior distribution is $p(\mathbf{u})$, and the likelihood function is $\mathcal{N}(\bar{\mathbf{y}}; \mathbf{u}, \mathbf{V}_{\mathbf{uu}})$. Then the log marginal likelihood is

$$\begin{aligned} \mathcal{Q} \equiv & -\frac{m}{2} \log(2\pi) - \frac{1}{2} \log |\mathbf{K}_{\mathbf{uu}} + \mathbf{V}_{\mathbf{uu}}| \\ & - \frac{1}{2} \bar{\mathbf{y}}^\top (\mathbf{K}_{\mathbf{uu}} + \mathbf{V}_{\mathbf{uu}})^{-1} \bar{\mathbf{y}}. \end{aligned} \quad (14)$$

The posterior distribution of \mathbf{f}_* is $\mathcal{N}(\mathbf{f}_*; \boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q)$, where $\boldsymbol{\mu}_q \equiv \mathbf{K}_{*\mathbf{u}}(\mathbf{K}_{\mathbf{uu}} + \mathbf{V}_{\mathbf{uu}})^{-1} \bar{\mathbf{y}}$, $\boldsymbol{\Sigma}_q \equiv \mathbf{K}_{**} - \mathbf{K}_{*\mathbf{u}}(\mathbf{K}_{\mathbf{uu}} +$

$\mathbf{V}_{\mathbf{u}\mathbf{u}}^{-1}\mathbf{K}_{\mathbf{u}\ast}$, and $\mathbf{K}_{\mathbf{u}\ast}$ is the Gram matrix of $(\mathbf{Z}, \mathbf{X}_\ast)$. See appendix C.1. These marginal likelihood and posterior distribution are computed using only summarized data. The implementation is straightforward, as it is equivalent to Gaussian process regression with a Gaussian likelihood. The computational complexity is $O(m^3)$.

Despite these advantages, we still have not discussed the difference between likelihood and sample quasi-likelihood. Here we demonstrate the asymptotic behavior of the sample quasi-likelihood. The assumptions outlined below enable us to apply the Laplace approximation using summary statistics. Note that the asymptotic behavior in Bayesian statistics can be found in (Watanabe 2018).

Assumption 5.2. $-\nabla\nabla\log p(\mathbf{y} | \mathbf{W}_{\mathbf{f}\mathbf{u}}\mathbf{u})|_{\mathbf{u}=\bar{\mathbf{y}}} = \mathbf{V}_{\mathbf{u}\mathbf{u}}^{-1}$.

Assumption 5.3. $p(\mathbf{y} | \mathbf{W}_{\mathbf{f}\mathbf{u}}\mathbf{u})$ is unimodal with respect to \mathbf{u} , having its mode at $\bar{\mathbf{y}}$.

The Laplace approximation is conventionally used to handle non-Gaussian likelihoods by approximating the posterior with a Gaussian centered at the maximum a posteriori estimate. Unlike this, our approach replaces the likelihood with a Gaussian centered at the maximum likelihood estimate.

Theorem 5.4. Suppose that assumption 5.2 and assumption 5.3 hold. Then we have

$$\begin{aligned} \mathcal{E} - \mathcal{Q} &= \log p(\mathbf{y} | \mathbf{W}_{\mathbf{f}\mathbf{u}}\mathbf{u})|_{\mathbf{u}=\bar{\mathbf{y}}} + \frac{m}{2} \log(2\pi) \\ &\quad + \frac{1}{2} \log|\mathbf{V}_{\mathbf{u}\mathbf{u}}| + o_p(m), \end{aligned} \quad (15)$$

where o_p denote convergence in probability. Additionally, the posterior distribution $p(\mathbf{f}_\ast | \mathbf{y}, \boldsymbol{\omega})$ asymptotically converges to $\mathcal{N}(\mathbf{f}_\ast; \boldsymbol{\mu}_p, \boldsymbol{\Sigma}_p)$, where

$$\begin{aligned} \boldsymbol{\mu}_p &\equiv \mathbf{K}_{\ast\mathbf{f}}\mathbf{K}_{\mathbf{f}\mathbf{f}}^{-1}\mathbf{W}_{\mathbf{f}\mathbf{u}}(\mathbf{V}_{\mathbf{u}\mathbf{u}}^{-1} + \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1})^{-1}\mathbf{V}_{\mathbf{u}\mathbf{u}}^{-1}\bar{\mathbf{y}}, \\ \boldsymbol{\Sigma}_p &\equiv \mathbf{K}_{\ast\ast} - \mathbf{K}_{\ast\mathbf{f}}\mathbf{K}_{\mathbf{f}\mathbf{f}}^{-1}\mathbf{K}_{\mathbf{f}\ast} \\ &\quad + \mathbf{K}_{\ast\mathbf{f}}\mathbf{K}_{\mathbf{f}\mathbf{f}}^{-1}\mathbf{W}_{\mathbf{f}\mathbf{u}}(\mathbf{V}_{\mathbf{u}\mathbf{u}}^{-1} + \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1})^{-1}\mathbf{W}_{\mathbf{u}\mathbf{f}}\mathbf{K}_{\mathbf{f}\mathbf{f}}^{-1}\mathbf{K}_{\mathbf{f}\ast}. \end{aligned}$$

Proof. See appendix C.2. \square

Theorem 5.5. Suppose that $(\mathbf{V}_{\mathbf{u}\mathbf{u}}^{-1} + \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1})^{-1}$, $\mathbf{K}_{\mathbf{u}\ast}^{-1}$, $\mathbf{K}_{\mathbf{u}\ast}$, $\mathbf{K}_{\mathbf{f}\mathbf{f}}^{-1}$, $\mathbf{W}_{\mathbf{f}\mathbf{u}}$, and $\mathbf{V}_{\mathbf{u}\mathbf{u}}^{-1}\bar{\mathbf{y}}$ become $O(\beta)$ when multiplied by a matrix of $O(\beta)$. Then we have

$$\boldsymbol{\mu}_p - \boldsymbol{\mu}_q = O(\beta), \quad \boldsymbol{\Sigma}_p - \boldsymbol{\Sigma}_q = O(\beta + m\beta^2). \quad (16)$$

Proof. See appendix C.3. \square

Theorem 5.4 describes the asymptotic behavior of the marginal likelihood and posterior distribution. This approximation performs well when n is sufficiently larger than m . Since the right side of eq. (15) does not depend on the hyperparameters of a covariance function, we can employ \mathcal{Q} to optimize \mathcal{E} . Regarding the posterior distribution, we cannot compute $\mathcal{N}(\mathbf{f}_\ast; \boldsymbol{\mu}_p, \boldsymbol{\Sigma}_p)$ given only summarized data since it contains $\mathbf{K}_{\mathbf{f}\mathbf{f}}$ and $\mathbf{K}_{\ast\mathbf{f}}$. Theorem 5.5 allows us to avoid the computation of them. From eq. (16), the sample quasi-likelihood can be used for a smaller range of inputs. Figure 4 illustrates that the approximation errors of the marginal likelihood and posterior distribution are affected by the relative granularity of summarized data with respect to the length

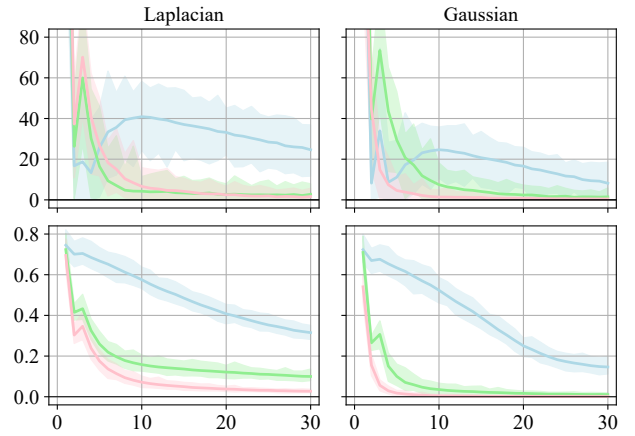


Figure 4: Behavior of marginal likelihood and posterior distribution in a toy model. The vertical axis in the upper figure represents the absolute difference between $\mathcal{Q} + \log p(\mathbf{y} | \mathbf{W}_{\mathbf{f}\mathbf{u}}\mathbf{u})|_{\mathbf{u}=\bar{\mathbf{y}}} + \frac{m}{2} \log(2\pi) + \frac{1}{2} \log|\mathbf{V}_{\mathbf{u}\mathbf{u}}|$ and \mathcal{L} . The vertical axis in the lower figure represents the root mean squared error (RMSE) between $\mathbf{K}_{\ast\mathbf{f}}(\mathbf{K}_{\mathbf{f}\mathbf{f}} + \text{diag}(1, \dots, 1))^{-1}\mathbf{y}$ and $\boldsymbol{\mu}_q$. Each solid line represents the average over 100 trials, with the shaded region in each color showing the range between the maximum and minimum values. The $n_\ast = 1000$ new inputs were uniformly generated within $[0, 2\pi]$. Each output y_i was generated from $\mathcal{N}(\sin(\mathbf{x}_i), 1)$. The summary statistic is the sample mean. The likelihood function is $\prod_{i=1}^n \mathcal{N}(y_i; f(\mathbf{x}_i), 1)$. Other displays and conditions are identical to those in fig. 3.

scale of the covariance functions. Their behavior largely aligns with the dynamics of η in fig. 3. While the input domain of this toy model is one-dimensional, this behavior is expected to persist in higher-dimensional data. However, it is important to note that the range of inputs expands as the input dimension increases.

The discussion so far supports the use of the sample quasi-likelihood for learning from summarized data. Addressing non-Gaussian likelihoods with summarized data presents a significant challenge, just as it does in Gaussian process regression with complete data. Our approach tackles this challenge by specifying an appropriate variance function, ensuring that assumption 5.2 is satisfied for an implicit likelihood function. Furthermore, the sample quasi-likelihood allows for the use of various summary statistics that satisfy assumption 5.3. This assumption requires that the summary statistic corresponds to the maximum likelihood estimate of the implicit likelihood function. For example, the sample median serves as the maximum likelihood estimate of the location parameter in the Laplace distribution. Our approximation performs poorly when the likelihood function associated with the variance function and summary statistic is flat. The concept of sample quasi-likelihood is analogous to the quasi-likelihood proposed by (Wedderburn 1974, 1976). While our approach also benefits from specifying a variance function instead of a likelihood function, its motivation

Output	Covariance	Grid Size					
		1.6×1.6	0.8×0.8	0.4×0.4	0.2×0.2	0.1×0.1	0.05×0.05
MedInc	Laplacian	0.57 ± 0.04	0.60 ± 0.05	0.59 ± 0.05	0.54 ± 0.05	0.44 ± 0.05	0.32 ± 0.05
MedInc	Gaussian	0.56 ± 0.09	0.59 ± 0.11	0.57 ± 0.12	0.50 ± 0.14	0.39 ± 0.12	0.28 ± 0.10
HouseAge	Laplacian	0.63 ± 0.03	0.57 ± 0.03	0.52 ± 0.04	0.41 ± 0.04	0.30 ± 0.04	0.19 ± 0.03
HouseAge	Gaussian	0.62 ± 0.04	0.57 ± 0.04	0.54 ± 0.05	0.40 ± 0.06	0.28 ± 0.07	0.17 ± 0.06
AveRooms	Laplacian	0.62 ± 0.12	0.59 ± 0.11	0.54 ± 0.11	0.49 ± 0.09	0.40 ± 0.10	0.29 ± 0.08
AveRooms	Gaussian	0.62 ± 0.16	0.63 ± 0.16	0.61 ± 0.18	0.59 ± 0.23	0.48 ± 0.26	0.32 ± 0.24
AveBedrms	Laplacian	0.62 ± 0.15	0.58 ± 0.15	0.51 ± 0.12	0.42 ± 0.10	0.32 ± 0.13	0.22 ± 0.10
AveBedrms	Gaussian	0.77 ± 0.23	0.77 ± 0.24	0.75 ± 0.22	0.71 ± 0.27	0.59 ± 0.30	0.44 ± 0.33
Population	Laplacian	0.26 ± 0.10	0.26 ± 0.10	0.25 ± 0.11	0.25 ± 0.12	0.22 ± 0.11	0.19 ± 0.09
Population	Gaussian	0.16 ± 0.12	0.16 ± 0.12	0.15 ± 0.13	0.14 ± 0.13	0.13 ± 0.13	0.11 ± 0.12
AveOccup	Laplacian	0.47 ± 0.17	0.47 ± 0.18	0.46 ± 0.18	0.45 ± 0.17	0.41 ± 0.16	0.30 ± 0.12
AveOccup	Gaussian	0.46 ± 0.33	0.46 ± 0.33	0.45 ± 0.34	0.45 ± 0.33	0.43 ± 0.32	0.36 ± 0.32
MedValue	Laplacian	0.79 ± 0.03	0.78 ± 0.04	0.74 ± 0.05	0.63 ± 0.05	0.47 ± 0.05	0.32 ± 0.04
MedValue	Gaussian	0.85 ± 0.08	0.84 ± 0.11	0.75 ± 0.14	0.58 ± 0.20	0.41 ± 0.18	0.27 ± 0.13

Table 2: Approximation performance. To compare our approximation with the original regression, we present the mean and standard deviation of the RMSE between the mean vectors of their posterior distributions over 100 trials. The RMSE was normalized by dividing it by the standard deviation of the training data outputs. Let the likelihood be Gaussian and let $n = 1000$.

and definition differ slightly. The formulation of the sample quasi-likelihood provides a computational advantage by expressing the marginal likelihood and posterior distribution in closed form. The comparison between both approaches is presented in appendix C.4.

6 Spatial Modeling

We investigate the usage of our method in spatial modeling tasks, using the California housing dataset ¹. Our approach leverages the input approximation via representative features and utilizes the sample quasi-likelihood. We evaluate its performance by comparing it with Gaussian process regression given complete data. To compute the exact marginal likelihood and posterior distribution, we used a Gaussian likelihood in this comparison. We also perform predictive testing with the Gaussian and Poisson likelihoods.

Summarized Data. The dataset was randomly shuffled and then split to obtain the n training data points and $n_* = 20640 - n$ test data points. The inputs were the pairs of latitude and longitude included among the attributes. The seven remaining attributes were used as outputs, respectively. The domain $[32.54, 41.95] \times [-124.35, -114.31]$ in latitude and longitude was divided into grids. The data points within each grid were summarized, with the centers of the grids serving as the representative features, and the summary statistics being the sample means. Assume that the sample variances are also provided and that the distribution of input data points for learning is unavailable.

Functions. The likelihood function is presented in appendix D.1. Let $g(\mathbf{x})$ denote $f(\mathbf{x})$ for the Gaussian, and $\exp(\mathbf{x})$ for the Poisson. We replace the summary statistics

$\bar{\mathbf{y}}$ with $(g^{-1}(\bar{y}_i))_{i=1}^m$. This replacement makes the likelihood functions consistent with assumption 5.3. Similarly, we specified the variance functions based on the likelihoods, ensuring assumption 5.2 is satisfied. The mean function was $\tau(\cdot) = g^{-1}(\frac{1}{n} \sum_{i=1}^n y_i)$. The covariance functions were designed by multiplying a constant kernel with the functions used in fig. 3 and adding white noise.

Hyperparameters. The hyperparameters of the covariance functions were optimized using the L-BFGS-B algorithm, starting from an initial value of 1. This optimization was implemented with version 1.7.3 of the SciPy software ², with the parameters set to their default values. We used a 64-bit Windows machine with the Intel Core i9-9900K @ 3.60 GHz and 64 GB of RAM. The code was implemented in the python programming language 3.7.3 version. As a hyperparameter, we also optimized the variance $\sigma \in (0, \infty)$ of the Gaussian likelihood. In this case, the right side of eq. (15) contains the hyperparameter. Therefore, we employed \mathcal{E} instead of \mathcal{Q} for the Gaussian likelihood, which can be computed using the sample mean and sample variances.

Baselines. For the prediction test baselines, we adopted the sparse variational inference method (Hensman, Fusi, and Lawrence 2013; Hensman, Matthews, and Ghahramani 2015) implemented in version 1.12.0 of the GPy software ³. The parameters in GPy were set to their default values. The initial locations of $\lfloor \sqrt{m^3 n^{-1}} \rfloor$ inducing points were randomly allocated from m representative locations. Two patterns using variational inference were attempted. One involved the standard learning and inference with complete data. In the other, the inputs were replaced with represen-

¹<https://lib.stat.cmu.edu/datasets/>

²<https://scipy.org/>

³<https://sheffieldml.github.io/GPy/>

Output	Likelihood	Complete VI		Summarized VI		Our Approach	
		RMSE	Time	RMSE	Time	RMSE	Time
MedInc	Gaussian	0.881 ± 0.014	41 ± 4	1.015 ± 0.015	30 ± 14	0.970 ± 0.004	6 ± 1
MedInc	Poisson	0.887 ± 0.024	59 ± 14	1.014 ± 0.016	59 ± 12	0.977 ± 0.008	6 ± 2
HouseAge	Gaussian	0.880 ± 0.049	39 ± 11	0.970 ± 0.024	29 ± 17	1.000 ± 0.000	4 ± 0
HouseAge	Poisson	0.861 ± 0.049	55 ± 18	0.966 ± 0.021	51 ± 20	1.032 ± 0.060	7 ± 2
AveRooms	Gaussian	0.988 ± 0.031	15 ± 12	0.995 ± 0.014	12 ± 12	0.975 ± 0.051	10 ± 3
AveRooms	Poisson	0.990 ± 0.031	19 ± 19	0.995 ± 0.015	19 ± 20	0.949 ± 0.012	8 ± 1
AveBedrms	Gaussian	0.989 ± 0.029	15 ± 13	0.996 ± 0.012	14 ± 12	0.964 ± 0.041	12 ± 3
AveBedrms	Poisson	0.927 ± 0.021	62 ± 9	0.941 ± 0.016	61 ± 11	0.962 ± 0.044	10 ± 3
Population	Gaussian	1.242 ± 0.224	5 ± 3	1.041 ± 0.085	18 ± 15	1.000 ± 0.000	1 ± 0
Population	Poisson	0.990 ± 0.009	45 ± 26	1.000 ± 0.005	37 ± 24	1.004 ± 0.013	6 ± 1
AveOccup	Gaussian	1.002 ± 0.008	12 ± 11	1.002 ± 0.006	10 ± 10	1.002 ± 0.006	6 ± 2
AveOccup	Poisson	1.021 ± 0.141	59 ± 23	1.015 ± 0.063	43 ± 15	1.001 ± 0.002	6 ± 2
MedValue	Gaussian	0.699 ± 0.009	41 ± 2	1.008 ± 0.015	40 ± 2	0.904 ± 0.010	6 ± 1
MedValue	Poisson	0.746 ± 0.033	59 ± 15	0.973 ± 0.038	58 ± 14	0.909 ± 0.011	5 ± 0

Table 3: Prediction performance. The mean and standard deviation of the RMSE and execution time [s] over 100 trials are shown. The RMSE between predictors and test outputs was normalized by dividing it by the standard deviation of the test data outputs. The prediction was obtained by applying the mapping g to the mean vector of the posterior distribution of \mathbf{f}_* . The execution time was measured from the start of learning to the end of prediction. The covariance function was the Gaussian. The grid size was 0.4×0.4 . Let $n = 10000$. The displays are bolded when the mean loss of the summarized VI and our method differs by greater than 0.01. Similarly, bolding is applied when the difference in the mean execution time exceeds 1 second. The results of the 0.8×0.8 grid size can be found in appendix D.2.

tative locations (i.e., $(z_{\omega_i})_{i=1}^n$). The complete outputs were used in both cases, contrasting with our setting where only summary statistics are available. We refer to these as the complete VI and summarized VI, respectively. The other settings were identical to those in our approach.

Evaluation. Table 2 shows that the granularity of summarized data strongly impact the approximation performance of our approach. This result is consistent with the theoretical findings and the simple simulation using the toy model. The superiority or inferiority of the two covariance functions depended on the output attributes, even though the same inputs were used. The comparison between the complete VI and summarized VI in table 3 suggests a negative impact of the coarse inputs. Our approach offered advantages in both loss and execution time compared to the summarized VI. For the AveRooms, specifying the variance function corresponding to the Poisson distribution was beneficial.

7 Conclusion

This study focuses on learning and inference when only summarized data is available. The introduction of the sample quasi-likelihood facilitates them, including in cases where the likelihood is non-Gaussian. Our approach is straightforward to implement and incurs low computational costs. We have demonstrated the approximation errors of the marginal likelihood and posterior distribution, both theoretically and experimentally. This analysis highlights when our approach works well. In particular, the approximation performance depends on the granularity relative to the length scale of co-

variance functions. Experiments on spatial modeling show the impact of using summarized data instead of complete data. The feasibility of our approach was supported by the series of discussions in this study.

Limitations. The input approximation does not perform well when the granularity of summarized data is rough. Since our approach relies on the Laplace approximation applied to each summarized group, the prediction performance for non-Gaussian observations deteriorates as the number of data points decreases. Further investigation is needed for extending our approach to cases where statistics other than the sample mean are available. The proportional constants in eq. (11) and eq. (12) contain the hyperparameters of a covariance function, as shown in lemma B.3 and lemma B.4. The theoretical behavior of our approximation concerning the hyperparameters requires further analysis.

Future Works. Spatial modeling is valuable for decision-making in many fields, but a significant concern is protecting privacy. While our discussion in this study is primarily from the perspective of a data user, the design of datasets to maintain confidentiality, such as through differential privacy (Dwork et al. 2006), is also important. Our approach potentially contributes to encouraging discussions on the granularity of summarized data, considering the trade-offs between usability and confidentiality. Additionally, optimizing representative locations and data point assignments to improve approximation performance is an appealing direction for future work.

References

- Beaumont, M. A. 2019. Approximate Bayesian Computation. *The Journal of Annual Review of Statistics and Its Application*, 6: 379–403.
- Beranger, B.; Lin, H.; and Sisson, S. 2022. New models for symbolic data analysis. *The Journal of Advances in Data Analysis and Classification*, 17: 659–699.
- Besag, J. 1975. Spatial interaction and the statistical analysis of lattice systems. *The Journal of the Royal Statistical Society: Series D*, 24: 179–195.
- Cressie, N. A. C. 1993. *Statistics for Spatial Data*. John Wiley and Sons, Inc.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the 3th Theory of Cryptography*, 265–284.
- Feragen, A.; Lauze, F.; and Hauberg, S. 2014. Geodesic Exponential Kernels: When Curvature and Linearity Conflict. In *Proceedings of the 28th Conference on Computer Vision and Pattern Recognition*, 3032–3042.
- Hensman, J.; Fusi, N.; and Lawrence, N. D. 2013. Gaussian Processes for Big data. In *Proceedings of the 29th International Conference on Uncertainty in Artificial Intelligence*, 282–290.
- Hensman, J.; Matthews, A.; and Ghahramani, Z. 2015. Scalable Variational Gaussian Process Classification. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 351–360.
- Law, H. C. L.; Sejdinovic, D.; Cameron, E.; Lucas, T. C.; Flaxman, S.; Battle, K.; and Fukumizu, K. 2018. Variational Learning on Aggregate Outputs with Gaussian Processes. In *Proceedings of the 32th International Conference on Neural Information Processing Systems*, 6084–6094.
- Lee, J.; Bahri, Y.; Novak, R.; Schoenholz, S. S.; Pennington, J.; and Sohl-Dickstein, J. 2018. Deep Neural Networks as Gaussian Processes. In *Proceedings of the 6th International Conference on Learning Representations*.
- Minka, T. P. 2001. *A Family of Algorithms for Approximate Bayesian Inference*. Ph.D. thesis, Massachusetts Institute of Technology.
- Price, L. F.; Drovandi, C. C.; Lee, A.; and Nott, D. J. 2018. Bayesian synthetic likelihood. *The Journal of Computational and Graphical Statistics*, 27(1): 1–11.
- Rasmussen, C. E.; and Williams, C. K. I. 2006. *Gaussian Processes for Machine Learning*. MIT Press.
- Sheth, R.; Wang, Y.; and Khardon, R. 2015. Sparse Variational Inference for Generalized Gaussian Process Models. In *Proceedings of the 32th International Conference on Machine Learning*, volume 37, 1302–1311.
- Tanaka, Y.; Tanaka, T.; Iwata, T.; Kurashima, T.; Okawa, M.; Akagi, Y.; and Toda, H. 2019. Spatially Aggregated Gaussian Processes with Multivariate Areal Outputs. In *Proceedings of the 33th International Conference on Neural Information Processing Systems*, 3005–3015.
- Varin, C.; Reid, N.; and Firth, D. 2011. An Overview of Composite Likelihood Methods. *The Journal of Statistica Sinica*, 21: 5–42.
- Watanabe, S. 2018. *Mathematical Theory of Bayesian Statistics*. Chapman and Hall.
- Wedderburn, R. W. M. 1974. Quasi-likelihood functions, generalized linear models, and the Gauss—Newton method. *Journal of Biometrika*, 61(3): 439–447.
- Wedderburn, R. W. M. 1976. On the Existence and Uniqueness of the Maximum Likelihood Estimates for Certain Generalized Linear Models. *The Journal of Biometrika*, 63(1): 27–32.
- Williams, C. K. I.; and Barber, D. 1998. Bayesian Classification With Gaussian Processes. *Journal of IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12): 1342–1351.
- Wood, S. N. 2010. Statistical inference for noisy nonlinear ecological dynamic systems. *The Journal of Nature*, 466: 1102–1104.