Kernel Methods for Bayesian Elliptic Inverse Problems on Manifolds

John Harlim [1] Daniel Sanz-Alonso [2] Ruiyi Yang [2]

jharlim@psu.edu {sanzalonso,yry}@uchicago.edu

Introduction

We study the inverse problem of determining the diffusion coefficient of a second-order elliptic PDE on a closed manifold from noisy measurements of the solution. Inspired by manifold learning techniques, we approximate the elliptic differential operator with a kernel-based integral operator that can be discretized via Monte-Carlo without reference to the Riemannian metric. We adopt a Bayesian perspective to the inverse problem, and establish an upper-bound on the total variation distance between the true posterior and the approximate posterior. Supporting numerical results show the effectiveness of the proposed methodology.

Bayesian Formulation of the Problem

We consider the elliptic equation

$$\mathcal{L}^\kappa u := -\mathrm{div}(\kappa \nabla u) = f, \qquad x \in \mathcal{M}, \tag{1}$$

where κ is a function on $\mathcal{M}$. We are interested in the inverse problem of determining the diffusion coefficient function κ from noisy measurements of u of the form

$$y = \mathcal{D}(u) + \eta,$$

where the observation map $\mathcal{D} : L^2 \to \mathbb{R}^J$ will be assumed to be known and $\eta \sim \mathcal{N}(0, \Gamma)$ for given positive definite $\Gamma \in \mathbb{R}^{J\times J}$.

The Bayesian formulation consists of three ingredients: the prior, the forward map, and the posterior.

Prior: Writing $\kappa = e^\theta$, we put a prior π over θ that is supported on some Banach space $\mathcal{B}$.

Forward map: Under certain regularity assumptions on κ and f, equation (1) has a unique weak solution in the space L_0^2 of mean-zero square integrable functions on $\mathcal{M}$. This allows us to define a forward map $\mathcal{F} : \theta \in \mathcal{B} \mapsto u \in L_0^2$, i.e., the map that associates each coefficient θ with its unique solution.

Posterior: Provided that the map $\mathcal{D} \circ \mathcal{F} : \mathcal{B} \to \mathbb{R}^J$ is measurable, the posterior π^y can be written as a change of measure with respect to the prior

$$\frac{d\pi^y}{d\pi}(\theta) \propto \exp\left(-\frac{1}{2}|y - \mathcal{D}\circ\mathcal{F}(\theta)|_\Gamma^2\right),$$

with $|\cdot|_\Gamma^2 := \langle \cdot, \Gamma^{-1}\cdot\rangle$. Therefore one can use the posterior mean for estimating the truth and the posterior credible intervals for uncertainty quantification.

Approximating the Forward Map

The forward map involves the solution operator of the PDE and is usually not known analytically. Various numerical methods such as finite element method have shown success in approximating the solution. However, most of the methods rely on knowing certain representations of the manifold $\mathcal{M}$, which is somehow restrictive. Therefore, we propose an approximate solution that can be discretized without full knowledge of $\mathcal{M}$. In particular, we will be interested in the case where only an unstructured sample of points on $\mathcal{M}$ is available.

We exploit ideas from diffusion maps proposed by [1] to approximate the Laplace-Beltrami operator $\Delta_\mathcal{M} := -\mathrm{div}(\nabla\cdot)$. Let

$$G_\varepsilon u(x) := \frac{1}{\sqrt{4\pi\varepsilon}^{\frac{m}{2}}}\int_\mathcal{M} \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) u(\tilde{x}) dV(\tilde{x}),$$

We have

$$G_\varepsilon u(x) = u(x) + \varepsilon\big(\omega u(x) - \Delta_\mathcal{M} u(x)\big) + O(\varepsilon^2), \qquad x \in \mathcal{M}. \tag{2}$$

We then approximate $\mathcal{L}^\kappa = -\mathrm{div}(\kappa\nabla\cdot)$ by the following relation

$$-\mathrm{div}(\kappa\nabla u) = \sqrt{\kappa}\left[\Delta_\mathcal{M}(u\sqrt{\kappa}) - u\Delta_\mathcal{M}\sqrt{\kappa}\right] = \frac{\sqrt{\kappa}}{\varepsilon}\left[uG_\varepsilon\sqrt{\kappa} - G_\varepsilon(u\sqrt{\kappa})\right] + O(\varepsilon^2) := \mathcal{L}_\varepsilon^\kappa u + O(\varepsilon^2).$$

Writing out explicitly,

$$\mathcal{L}_\varepsilon^\kappa u(x) = \frac{1}{\sqrt{4\pi}\varepsilon^{\frac{m}{2}+1}}\int_\mathcal{M}\exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right)\sqrt{\kappa(x)\kappa(\tilde{x})}[u(x) - u(\tilde{x})]dV(\tilde{x}).$$

It turns out the approximate equation

$$\mathcal{L}_\varepsilon^\kappa u = f$$

also has a unique weak solution in L_0^2 so that we can define an approximate forward map $\mathcal{F}_\varepsilon : \mathcal{B} \mapsto L_0^2$. The approximate posterior is then

$$\frac{d\pi_\varepsilon^y}{d\pi}(\theta) \propto \exp\left(-\frac{1}{2}|y - \mathcal{D}\circ\mathcal{F}_\epsilon(\theta)|_\Gamma^2\right).$$

Choice of Prior

We will consider Gaussian priors of the form

$$\pi = \mathcal{N}(0, \mathcal{C}_{\tau,s}), \quad \mathcal{C}_{\tau,s} = (\tau I + \Delta_\mathcal{M})^{-s},$$

where $\Delta_\mathcal{M}$ is the Laplace-Beltrami operator on $\mathcal{M}$. Random samples of π admit a Karhunen-Loève expansion

$$v = \sum_{i=1}^\infty (\tau + \lambda_i)^{-s/2}\xi_i\varphi_i, \qquad \xi_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$$

where (λ_i, φ_i)'s are eigenpairs of $\Delta_\mathcal{M}$. Such priors are related to Gaussian processes with Matérn covariance functions in the Euclidean setting [3]. The parameter τ controls the essential frequencies and hence the inverse length scale of the samples paths. The parameter s controls the decay of the coefficients and hence regularity of the sample paths. By Sobolev embedding, a sufficiently large s will guarantee that the samples paths of π belong to C^4, which will be needed in the following theoretical results.

Theoretical Results

Forward map approximation: Suppose that $f \in \mathcal{C}^{3,\alpha}$ and $\kappa \in \mathcal{C}^4$, with κ bounded below by $\kappa_{\min} > 0$. Let u solve $\mathcal{L}^\kappa u = f$, and u_ε solve $\mathcal{L}_\varepsilon^\kappa u_\varepsilon = f$ weakly in L_0^2. Then for $\frac{1}{4} < \beta < \frac{1}{2}$ and ε small enough depending on β,

$$\|u - u_\varepsilon\|_{L^2} \le CA(\kappa)\|f\|_{H^3}\varepsilon^{4\beta-1},$$

where C is a constant depending only on $\mathcal{M}$ and $A(\kappa)$ is a function of $\kappa_{min}, \|\kappa\|_{\mathcal{C}^3}, \|\kappa\|_{\mathcal{C}^4}$.

Posterior approximation: Let π be a Gaussian measure on $\mathcal{C}^4$, and suppose that $f \in \mathcal{C}^{3,\alpha}$ for $0 < \alpha < 1$. Then for any $\frac{1}{4} < \beta < \frac{1}{2}$ and ε sufficiently small depending on β,

$$d_{\mathrm{TV}}(\pi^y, \pi_\varepsilon^y) \le C\varepsilon^{4\beta-1},$$

where C is constant depending only on $\mathcal{M}$.

Discretization

We now demonstrate how to discretize the approximate operator $\mathcal{L}_\varepsilon^\kappa$ for computation purposes. Suppose we are only given an unstructured point cloud $\{x_i\}_{i=1}^n$ distributed according to some unknown density q. Then $\mathcal{L}_\varepsilon^\kappa u$ evaluated at the point cloud is approximated by

$$\mathcal{L}_\varepsilon^\kappa u(x_i) \approx \frac{1}{\sqrt{4\pi}n\varepsilon^{\frac{m}{2}+1}}\sum_{j=1}^n \exp\left(-\frac{|x_i - x_j|^2}{4\epsilon}\right)\sqrt{\kappa(x_i)\kappa(x_j)}q_\varepsilon(x_j)^{-1}[u(x_i) - u(x_j)] := L_{\varepsilon,n}^\kappa u(x_i), \tag{3}$$

where q_ε is an estimate of q. Using the relation (2) we can approximate q up to first order in ε by $G_\varepsilon q$, i.e. we set

$$q_\varepsilon(x_j) := \frac{1}{\sqrt{4\pi}n\varepsilon^{\frac{m}{2}}}\sum_{k=1}^n \exp\left(-\frac{|x_j - x_k|^2}{4\varepsilon}\right).$$

The discretization of equation (1) becomes

$$L_{\varepsilon,n}^\kappa u_n = f_n,$$

where f_n is the n-dimensional vector with entries $f(x_i)$. One can see from (3) that $L_{\varepsilon,n}^\kappa$ is positive semidefinite and self-adjoint under the weighted inner product $\langle u, v\rangle_q := \frac{1}{n}\sum_{j=1}^n u(x_i)v(x_i)q_\varepsilon(x_i)^{-1}$. Hence $L_{\varepsilon,n}^\kappa$ admits a nonnegative spectrum $\{\lambda_i\}_{i=1}^n$ with $\lambda_1 = 0$ and an orthonormal basis of eigenfunctions $\{v_i\}_{i=1}^n$ wih respect to $\langle\cdot,\cdot\rangle_q$, with $v_1 \equiv 1$. We then set the solution to be

$$u_n = \sum_{i=2}^n \frac{f_n^i}{\lambda_i}v_i,$$

where the $f_n^i = \langle f_n, v_i\rangle_q$. The prior π can be discretized similarly and we will use the pCN algorithm for posterior sampling.

Numerical Results

One Dimensional Problem on an Unknown Ellipse

We consider $\mathcal{M}$ as the ellipse parametrized by $\iota(\omega) = (\cos\omega, 3\sin\omega)^T, \omega \in [0, 2\pi]$. Set the truth as

$$\kappa^\dagger = 2 + \cos\omega, \quad u^\dagger = \cos\omega$$

and compute f analytically. $\mathcal{M}$ is then given as a sample of 400 points. Below are the reconstructions given f and J noise observations of the form $y_i = u^\dagger(x_i) + \mathcal{N}(0, \sigma^2)$.

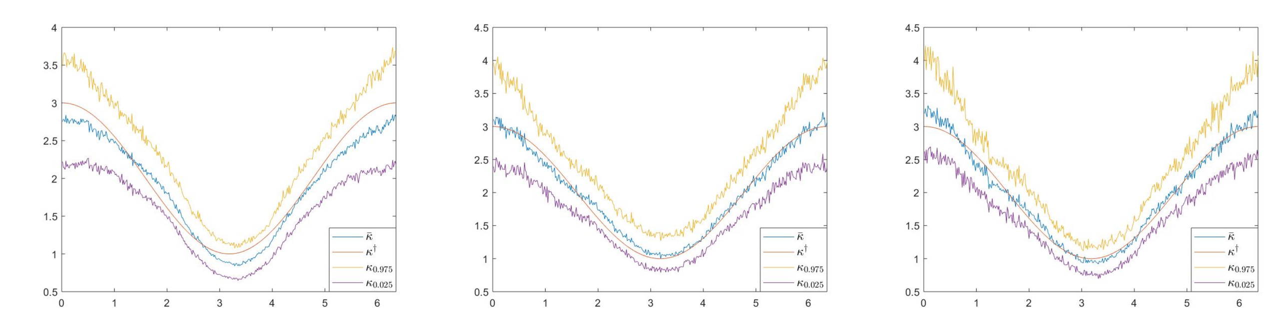

$J = 100$. $J = 200$. $J = 400$.

Figure: Posterior means and 95% credible intervals for $\sigma = 0.1$ and different J's. Here $\kappa_{0.025}$ and $\kappa_{0.975}$ represent the 2.5% and 97.5% posterior quantiles respectively.

Two Dimensional Problem on an Unknown Artificial Surface

We consider $\mathcal{M}$ as a cow-shaped surface, which consists of 2930 data points. To avoid inverse crime, we generate our truth using 2930 points and then set up our inverse problem on a subset of 1000 points. Let Δ_{2930} be a graph Laplacian constructed with the full 2930 points. $\kappa^\dagger$ is generated from $\mathcal{N}(0, \tau I + \Delta_{2930})^{-s}$ for $\tau = 0.7, s = 6$ and $u^\dagger$ is the rescaled second eigenvector of Δ_{2930}. Below are the reconstructions.

Posterior mean $\bar{\kappa}$. Truth $\kappa^\dagger$. Error $|\bar{\kappa} - \kappa^\dagger|$. Standard deviation.

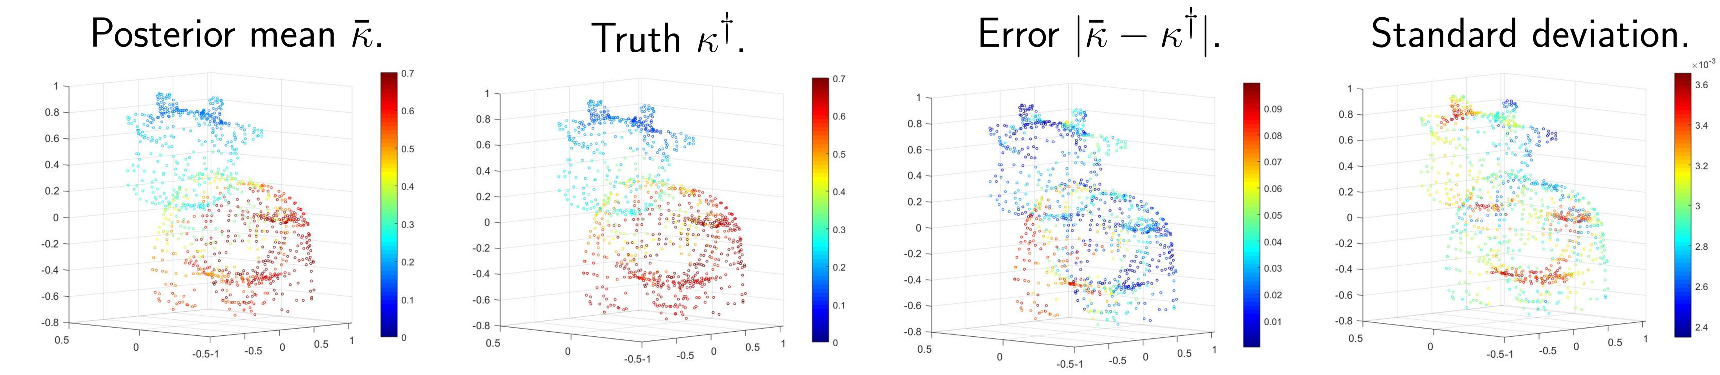

References

[1] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.

[2] J. Harlim, D. Sanz-Alonso, and R. Yang. Kernel methods for bayesian elliptic inverse problems on manifolds. *SIAM/ASA Journal on Uncertainty Quantification*, 8(4):1414–1445, 2020.

[3] D. Sanz-Alonso and R. Yang. The spde approach to mat\'ern fields: Graph representations. *arXiv preprint arXiv:2004.08000*, 2020.