
BAYESIAN DECISION THEORY FOR GAUSSIAN PROCESS MODELS: WITH SOME RESULTS ON LOSS FUNCTION EQUIVALENCE AND A DECISION THEORETIC EXTENSION THEOREM

Ryan Warnick

Microsoft Security Research
Redmond, Washington 98052
ryanwarnick@microsoft.com

October 13, 2025

ABSTRACT

We provide a plug-and-drop decision theoretic recipe to recover source terms Q in general order linear ordinary differential equations ($\mathcal{L}y = \sum_{k=0}^n a_k(\circ) y^{(k)} = Q$, $x \in \Omega$) and partial differential equations ($(\mathcal{L}y)(x) = \sum_{|\alpha| \leq k} A_\alpha(x) D^\alpha y(x)$, $x \in \Omega$), assuming a Gaussian process prior on y with derivatives up to order k in mean square. Under squared error loss, the Bayes action is the posterior mean of $\mathcal{L}y$; we give closed-form expressions for both the mean and covariance of Q , handle general linear observations (including exact/noisy boundary conditions), and provide an implementation-ready matrix formulation. The method permits evaluation of infinite-dimensional Bayes actions on function spaces (with associated uncertainty quantification), or, in the finite approximation setting, an uncountably infinite family of Bayes actions corresponding to the counting measure on the domain $x \in \Omega$. An equivalence between a type of 0-1 loss and the mean squared loss on Q is shown under Gaussian process measures, and an extension theorem for infinite-dimensional decision rules on function spaces is formalized. Additionally, an extension to nonlinear operators is also outlined utilizing second-order Frechet expansions, with associated probability bounds.

1 Introduction

The Bayesian decision-theoretic perspective offers a powerful and conceptually unified approach to inference, uncertainty quantification, and optimal decision making in parametric problems; however, extensions to the infinite-dimensional setting have been minimally explored. While Gaussian processes (GPs) have become a central tool for modeling and inference in statistical and machine learning applications, the decision-theoretic dimension of GP modeling concerned with the definition and computation of optimal actions under posterior uncertainty remains comparatively underexplored.

The following work establishes foundational results for Bayesian decision theory for Gaussian process models, proving the equivalence of 0–1 loss and mean-squared loss under Gaussian measures, and formalizing the extension theorem for Bayes actions on Hilbert spaces; an additional result on decision boundaries between two Gaussian processes have is outlined, allowing the design of Bayes-optimal decision boundaries between two distinct Gaussian processes defined on a shared Hilbert space. These results connect classical Bayesian inverse problems [Dashti and Stuart, 2015, Kaipio and Somersalo, 2004, Kaipio et al., 2007] to the emerging field of probabilistic numerics [Cockayne et al., 2019], where algorithms are treated as stochastic estimators and evaluated under decision theoretic loss functions.

In this continuation, we propose an extension the theoretical and practical scope of these proceeding frameworks by addressing three interconnected challenges: (i) The formal analysis of decision theoretic risk in infinite dimensional Hilbert space problems. (ii) The integration of decision-theoretic risk into experimental design and data acquisition. (iii) The extension of risk bounds and uncertainty quantification to nonlinear and operator-valued mappings. Much of

the following paper is expository results outlining known properties of inverse problems in differential equations (both partial and ordinary) using GPs. However, the decision theoretic framing is novel, and the ensuing decision theoretic results outlined in Section 2 permit generalizations of the same approaches to broader classes of decision problems.

The paper then operates as a survey outlining general background and a computational approaches for source function evaluation from Gaussian processes (GPs) believed to be noisy realizations from an ODE or PDE, but also with extensions to nonlinear operators and three principal results extending GP based decision theory as well as decision theory on general Hilbert spaces.

1.1 Motivation

Many scientific and engineering problems involve making optimal decisions under uncertainty over continuous domains. Inverse problems, operator identification, and physical system modeling all depend on recovering unknown functions or operators from noisy data. Gaussian processes provide a nonparametric and flexible prior over such functions, while Bayesian inference supplies the engine for updating beliefs in light of data. Typically these are treated as separate problems, as the camp of practitioners in the scientific computing and inverse problems camp use the Bayesian machinery (and primarily GP-based machinery) only for functional inference, and then the ensuing inversion employs general tools from probabilistic-numeric [Cockayne et al., 2019, Dashti and Stuart, 2015]. However, in the Bayesian community more broadly there is broad effort to optimize decision making in uncertain environments through loss function design and finding optimal bayes actions through minimizing the posterior expected loss [Berger, 2010].

This combined perspective permits a decision theoretic and Bayesian experimental design [Chaloner and Verdinelli, 1995] approach to optimal decision making in Bayesian inverse problems (ordinary/partial differential equation based or otherwise). In the Gaussian process solution to inverse problems, this permits an additional benefit: the results here make no effort to diminutize the longstanding $O(n^3)$ for number of samples $n \in \mathbb{N}$ computational bottleneck of Gaussian process inference, however, if one can get $m \in \mathbb{N} < n$ for which the Bayes-risk under the m is less than under the n , or which has lower variance of the Bayes action solution to minimizing Bayes-risk, then the computational bottleneck can be ever-so-slightly circumvented.

1.2 Contributions

This work makes the following contributions, expanding the theoretical and computational space of Bayesian decision theory for nonparametric models (and particularly GPs):

1. We prove three principal novel results in Bayesian decision theory on Hilbert spaces:

- *Extension theorem for Bayes actions on Hilbert spaces* We prove that under mild regularity conditions on the loss function and underlying Hilbert space upwardly convergent subspaces (e.g. finite approximations to the underlying Hilbert space) converge to the solution to the loss acting on the global Hilbert space.
- *0-1 small ball probability loss vs. squared-error loss Bayes action equivalence for Gaussian Process Measures on Hilbert spaces* An equivalence theorem between the squared error loss and small ball probability loss (a variant of 0-1 loss) is shown for Gaussian process measures under certain conditions. This extends the result outlining equivalence between the mean minimizing the L_2 loss in finite dimensional Gaussians and the median (which is the mean) minimizing the 0-1 loss in finite dimensional Gaussians.
- *Decision boundary for two Gaussian-process posteriors* For two Gaussian processes defined on a shared Hilbert Space, we give an explicit characterization of a Hilbert manifold contained in the said Hilbert space which is locally diffeomorphic to the global Hilbert manifold. E.g., the decision boundary between two GPs locally appears like the shared Hilbert space and Frechet calculus can be performed on it.

2. **Bayesian Risk-Based Experimental Design.** As outlined in section 1.1, we introduce a new formalism for optimal data acquisition based on posterior expected loss rather than posterior variance. This leads to a family of risk driven experimental design procedures that directly target decision functionals (e.g., differential operators, source terms, or physical observables) and can be efficiently computed from the GP posterior mean and covariance. This contribution bridges the gap between Gaussian process regression and Bayesian optimal design.

3. **Frechet-Based Nonlinear Operator Analysis.** We extend the Frechet derivative framework to quantify risk and uncertainty in nonlinear operator mappings applied to Gaussian process posteriors over function. The analysis establishes explicit high probability bounds on decision error as a function of the operator norm of the Frechet derivative and curvature. These results generalize the previous linear Gaussian guarantees outlined in the preceding sections to the nonlinear setting and provide the theoretical underpinning for risk certification in operator learning and physics-informed modeling.

4. **Unified Computational Framework.** We provide a matrix-based formulation for implementing decision-theoretic inference and experimental design, compatible with general linear observation operators and flexible boundary conditions. The proposed algorithms admit efficient implementations via standard Gaussian process conditioning, while preserving theoretical rigor in infinite-dimensional Hilbert spaces.
5. **Applications and Implications.** We discuss applications of the developed theory to Bayesian inverse problems and probabilistic numerics methods, highlighting how decision theoretic risk can serve as a universal design and validation metric in Gaussian process modeling. This framework lays the groundwork for principled risk-aware learning in physics informed neural and kernel-based models.

1.3 Relation to Prior Work and Looking Forward

The framework developed here connects three strands of research. First, it builds on the theory of Bayesian inverse problems Dashti and Stuart [2015], Kaipio and Somersalo [2004], extending classical results on posterior consistency to decision-based optimality. Second, it aligns with probabilistic numerical methods Cockayne et al. [2019], where the objective is to quantify uncertainty in numerical computations. Finally, it contributes to the growing literature on physics-informed and operator-based learning Raissi et al. [2019], offering a theoretical foundation for uncertainty-aware operator inference. The framework similarly extends the rich Bayesian tradition within decision theory [Berger, 2010] to broader classes of decision problems on Hilbert spaces, and particularly, as outlined in the following Section, Theorem 1 serves as an extension theorem for any decision problem on Hilbert spaces for which the loss satisfies certain regularity conditions. This opens the door to other rich classes of nonparametric posteriors in Bayesian analysis.

1.4 Remaining Gap: Computational Bottleneck

Reiterating what was outlined in section 1.1, the most substantial missing component of the framework is that we do not address the substantial challenge of the $O(n^3)$ computational bottleneck of Gaussian process inference. However, a major contribution is that, as noted in Section 1.2, the Bayesian experimental design approach allows optimizing performance for computable n by optimizing the experimental design. This permits the practitioner to optimize for various types of loss functions and assess the overall contribution to reducing uncertainty in Source function estimation of each data point (and, more broadly, general Gaussian process decision problems), and reduce the posterior uncertainty optimally.

As outlined in Section 1.3, there exist many rich classes of Bayesian nonparametric priors and ensuing posterior formulations for which Theorem 1 permits analysis. Some of these include Polya-trees [Lavine, 1992], Gibbs-type species sampling priors [De Blasi et al., 2015], and completely random measures governed by a Levy-intensity [Kingman, 1967]. We hope this explicit foray into what at first pass appears to be the most tractable of all nonparametric priors (GPs) inspires others to pursue similar directions with broader families.

1.5 Paper Organization

2 Results

We outline three principal results and prove them in the appendix. These results support the decision theoretic framing and extend the mean squared loss as equivalent to a larger family of losses and decision problems.

Theorem 1 (Extension theorem for Bayes actions on Hilbert spaces). *($\mathcal{H}, \langle \circ, \circ \rangle$) be a (real, separable) Hilbert space with norm $\| \circ \|$. Define a loss function $\Psi : (a, \mathcal{H}) \rightarrow \mathbb{R}$ for $a \in \mathcal{H}$, and assume Ψ on $\mathcal{H} \times \mathcal{H}$. Assume that there exists a posterior process $\mathcal{Q}|z$ with elements almost surely $q \sim \mathcal{Q}|z \in \mathcal{H}$. Let $\{H_n\}_{n \geq 1}$ be an increasing sequence of finite dimensional subspaces $H_n \uparrow \mathcal{H}$ and let P_n be the corresponding projection operator $P_n : \mathcal{H} \rightarrow H_n$. Suppose the following to inequality hold for some $C, L > 0$ and $p \geq 1$. (Globally Lipschitz and linear growth):*

$$|\Psi(a, h) - \Psi(a, h')| \leq L \|h - h'\| \tag{1}$$

$$|\Psi(a, h)| \leq C(1 + \|a\|^p + \|h\|^p) \tag{2}$$

And $\mathbb{E}_{\mathcal{Q}|z}[\|Q\|^p|z] < \infty$. Then $\mathbb{E}_{\mathcal{Q}}[\|\Psi(a_n, P_n Q)\|] \xrightarrow{n \rightarrow \infty} \mathbb{E}_{\mathcal{Q}}[\Psi(a, Q)]$.

If $\lim_{n \rightarrow \infty} a_n = a \in \mathcal{H}$ exists and is finite, then it is the Bayes action for $a = \operatorname{argmin}_{a \in \mathcal{H}} \mathbb{E}_{\mathcal{Q}}[\Psi(a, Q)|z]$

Proof. Proof available in Appendix section A □

Theorem 2 (0–1 small ball probability loss vs. squared-error loss Bayes action equivalence for Gaussian Process Measures on Hilbert spaces). *Let $(\mathcal{H}, \langle \cdot, \cdot \rangle)$ be a (real, separable) Hilbert space with norm $\|\cdot\|$, and suppose the posterior distribution of a random element $Q \in \mathcal{H}$ given data z is Gaussian:*

$$Q | z \sim N(\mu, \mathcal{C}),$$

with mean $\mu \in \mathcal{H}$ and covariance operator $\mathcal{C} : \mathcal{H} \rightarrow \mathcal{H}$ (self-adjoint, positive, trace-class on the closed linear span of the support). For an action $a \in \mathcal{H}$ consider the two losses

$$L_2(q, a) = \|q - a\|^2, \quad L_{01,\tau}(q, a) = \mathbf{1}\{\|q - a\| > \tau\},$$

where $\tau > 0$ is any fixed radius. Then the Bayes action under both losses equals the posterior mean:

$$a^* = \arg \min_{a \in \mathcal{H}} \mathbb{E}[L_2(Q, a) | z] = \arg \min_{a \in \mathcal{H}} \mathbb{E}[L_{01,\tau}(Q, a) | z] = \mu.$$

Consequently, the Bayes rules under L_2 and under $L_{01,\tau}$ (for any $\tau > 0$) are identical.

Proof. Proof available in Appendix section B □

Theorem 3 (Decision boundary for two Gaussian-process posteriors). *Let \mathcal{H} be a real, separable Hilbert space. Let $Y \sim \mathcal{N}(\mu, \mathcal{C})$ and $Y' \sim \mathcal{N}(\mu', \mathcal{C}')$ be Gaussian random elements in \mathcal{H} with means $\mu, \mu' \in \mathcal{H}$ and covariance operators $\mathcal{C}, \mathcal{C}'$ (self-adjoint, positive, trace-class on the closed span of their supports). For any $a \in \mathcal{H}$, define the small-ball 0–1 risks*

$$R_{01,\tau}^Y(a) := \mathbb{E}[\mathbf{1}\{\|Y - a\| > \tau\}], \quad R_{01,\tau}^{Y'}(a) := \mathbb{E}[\mathbf{1}\{\|Y' - a\| > \tau\}].$$

Define the decision boundary as the zero-level set of the limiting risk difference,

$$\mathcal{B} := \left\{ a \in \mathcal{H} : \lim_{\tau \downarrow 0} (R_{01,\tau}^Y(a) - R_{01,\tau}^{Y'}(a)) = 0 \right\}.$$

Then \mathcal{B} is a Hilbert manifold in \mathcal{H} , and coincides with the zero-level set of the squared-loss risk difference,

$$\mathcal{B} = \left\{ a \in \mathcal{H} : \mathbb{E}\|Y - a\|^2 = \mathbb{E}\|Y' - a\|^2 \right\},$$

and therefore admits the explicit affine description

$$\boxed{\|a - \mu\|^2 - \|a - \mu'\|^2 = \text{tr}(\mathcal{C}') - \text{tr}(\mathcal{C})} \Leftrightarrow \boxed{\frac{\|a - \mu\|^2 - \|a - \mu'\|^2}{\text{tr}(\mathcal{C}') - \text{tr}(\mathcal{C})} = 1}.$$

Which is an affine Hilbert manifold in \mathcal{H} locally diffeomorphic to \mathcal{H} . In particular, when $\mathcal{C} = \mathcal{C}'$, the boundary is the (codimension-one) bisector hyperplane Hilbert manifold in \mathcal{H}

$$\{ a \in \mathcal{H} : \langle a, \mu' - \mu \rangle = \frac{1}{2}(\|\mu'\|^2 - \|\mu\|^2) \},$$

i.e. the locus of points equidistant (in $\|\cdot\|$) from μ and μ' .

Proof. Proof available in Appendix Section C □

3 Linear k'th Order ODEs

3.1 Model and observations

Operator. Let

$$(\mathcal{L}y)(u) = \sum_{k=0}^n a_k(u) y^{(k)}(u), \quad u \in \mathcal{U} \subset \mathbb{R},$$

with known coefficient functions $a_k : \mathcal{U} \rightarrow \mathbb{R}$ and order $n \geq 0$.

Prior. Place a zero-mean GP prior on the latent solution:

$$y \sim \text{GP}(0, K),$$

with kernel K differentiable up to order n in each argument.

Observations (very general). Assume we observe $\mathbf{z} \in \mathbb{R}^m$ through a *linear* observation operator \mathcal{H} with additive Gaussian noise:

$$\mathbf{z} = \mathcal{H}[y] + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, R).$$

This covers point evaluations $y(t_i)$, derivative samples $y^{(j)}(t_i)$, linear functionals (e.g., integrals), and (noiseless) boundary constraints by taking the corresponding rows of R to have vanishing variance (practically, a very small jitter).

3.2 Decision problem and Bayes estimator

We seek a decision rule $\hat{Q} : \mathcal{U} \rightarrow \mathbb{R}$ minimizing posterior expected squared loss

$$\mathcal{R}(\hat{Q}) = \mathbb{E} \left[\int_{\mathcal{U}} (\mathcal{L}y(u) - \hat{Q}(u))^2 du \mid \mathbf{z} \right].$$

By convexity and linearity of expectation, the pointwise minimizer is

$$\hat{Q}(u) = \mathbb{E}[(\mathcal{L}y)(u) \mid \mathbf{z}] = \sum_{k=0}^n a_k(u) \mathbb{E}[y^{(k)}(u) \mid \mathbf{z}].$$

Uncertainty (useful for UQ and design) comes from the posterior covariance

$$\text{Cov}(Q(u), Q(v) \mid \mathbf{z}) = \text{Cov}((\mathcal{L}y)(u), (\mathcal{L}y)(v) \mid \mathbf{z}).$$

3.3 Closed forms via linear-Gaussian conditioning

Let $U_* = \{u_1, \dots, u_s\}$ be test locations where we want Q . Stack the derivatives of y needed at U_* as

$$\mathbf{y}_* = [y^{(0)}(U_*); y^{(1)}(U_*); \dots; y^{(n)}(U_*)] \in \mathbb{R}^{(n+1)s}.$$

Similarly, let Z denote the set of inputs (times/locations and derivative orders) used by the observation operator \mathcal{H} . Because $(\mathbf{y}_*, \mathcal{H}[y])$ are jointly Gaussian, the posterior is

$$\mathbf{y}_* \mid \mathbf{z} \sim \mathcal{N}(\mu_*, \Sigma_*), \quad \mu_* = K_{*Z} (K_{ZZ} + R)^{-1} \mathbf{z}, \quad \Sigma_* = K_{**} - K_{*Z} (K_{ZZ} + R)^{-1} K_{Z*},$$

where the blocks are obtained by differentiating the kernel:

$$\begin{aligned} [K_{**}]_{(k,i),(l,j)} &= \frac{\partial^{k+l}}{\partial u^k \partial v^l} K(u_i, v_j), \\ [K_{*Z}]_{(k,i),r} &= \mathbb{E}[y^{(k)}(u_i) (\mathcal{H}[y])_r] \quad (\text{differentiate } K \text{ under } \mathcal{H}), \\ K_{ZZ} &= \mathbb{E}[\mathcal{H}[y] \mathcal{H}[y]^\top] \quad (\text{apply } \mathcal{H} \text{ to both kernel arguments}). \end{aligned}$$

Linear map from derivatives to Q . Define the *block-row* linear operator $L_* \in \mathbb{R}^{s \times (n+1)s}$ acting pointwise as

$$(L_* \mathbf{y}_*)_i = \sum_{k=0}^n a_k(u_i) y^{(k)}(u_i), \quad i = 1, \dots, s.$$

Then the posterior for $\mathbf{Q}_* := Q(U_*)$ is Gaussian:

$$\mathbf{Q}_* \mid \mathbf{z} \sim \mathcal{N} \left(\underbrace{L_* \mu_*}_{\text{Bayes estimator}}, \underbrace{L_* \Sigma_* L_*^\top}_{\text{posterior covariance}} \right).$$

Equivalently, the *pointwise* Bayes estimator is

$$\hat{Q}(u_i) = \sum_{k=0}^n a_k(u_i) \mu_*^{(k)}(u_i).$$

3.4 Matrix Approach

For many users it is convenient to think entirely in matrices at a grid U_* . Let $D_k \in \mathbb{R}^{s \times s}$ denote the linear map that returns the k^{th} derivative of the GP at U_* when applied to the latent vector $y(U_*)$.¹ Let $A_k = \text{diag}(a_k(U_*))$. Define the assembled operator

$$L = \sum_{k=0}^n A_k D_k \in \mathbb{R}^{s \times s_{\text{stack}}},$$

where D_k selects the k -th derivative block from \mathbf{y}_* (or, equivalently, we realize L_* directly as above). Then

$$\hat{\mathbf{Q}} = L \mu_*, \quad \text{Cov}(\mathbf{Q}) = L \Sigma_* L^\top.$$

This is the implementation once we have μ_*, Σ_* .

3.5 Observations (including boundary conditions included)

Common choices of \mathcal{H} and corresponding kernel manipulations:

- **Noisy point values:** $z_i = y(t_i) + \epsilon_i$. Then $K_{ZZ} = K(t_i, t_j)$, $R = \text{diag}(\sigma_i^2)$, and K_{*Z} uses $\partial_u^k K(u_i, t_j)$.
- **Noisy derivatives:** $z_i = y^{(r_i)}(t_i) + \epsilon_i$. Differentiate kernel in the t argument accordingly.
- **Exact boundary conditions:** $y(b_\ell) = c_\ell$ or $y^{(r)}(b_\ell) = c$. Encode as rows of \mathcal{H} with a tiny noise variance (e.g. 10^{-10}) so that conditioning stays numerically stable.
- **Linear functionals:** $z_i = \int w_i(t) y^{(r_i)}(t) dt + \epsilon_i$. Apply \mathcal{H} to the kernel by pushing the integral/derivative through K .

3.6 Kernels and differentiability

We need K differentiable up to order n in each argument. Examples:

- Squared Exponential (RBF): infinitely differentiable.
- Matérn ν : k -th derivatives exist in mean-square for $k < \nu$. For an n -th order ODE, choose $\nu > n + 1/2$ for comfortable margins.

3.7 Practical Implementation

1. Choose order n and coefficient functions $a_k(u)$.
2. Choose kernel K with sufficient smoothness.
3. Specify linear observations \mathcal{H} and noise R (including boundary conditions).
4. Build kernel blocks by differentiating K as required to form K_{ZZ}, K_{*Z}, K_{**} .
5. Compute μ_*, Σ_* via Gaussian conditioning.
6. Evaluate $\hat{\mathbf{Q}} = L_* \mu_*$ and (optionally) $\text{Cov}(\mathbf{Q}) = L_* \Sigma_* L_*^\top$.

3.8 Example: 3'rd Order Linear ODE

4 Linear k'th Order PDEs

4.1 Setup: multi-index operators on $\Omega \subset \mathbb{R}^d$

Let $\Omega \subset \mathbb{R}^d$ be open with boundary $\partial\Omega$. For a multi-index $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}_0^d$, denote $D^\alpha \equiv \partial^{|\alpha|} / \partial x_1^{\alpha_1} \circ \dots \circ \partial x_d^{\alpha_d}$, with order $|\alpha| = \sum_j \alpha_j$. Consider a (possibly vector-valued) field $y : \Omega \rightarrow \mathbb{R}^p$ and a linear differential operator of order m :

$$(\mathcal{L}y)(x) = \sum_{|\alpha| \leq m} A_\alpha(x) D^\alpha y(x), \quad x \in \Omega,$$

¹In exact GP algebra, derivatives are obtained via differentiated kernels against observations; D_k here is a mnemonic for ‘‘apply the k -th derivative with respect to the first kernel argument and evaluate at U_* ’’. In practice we never build D_k explicitly; we build the corresponding kernel blocks.




Figure 1: The data and the fit Gaussian Process for the 3'rd order ODE template scenario.

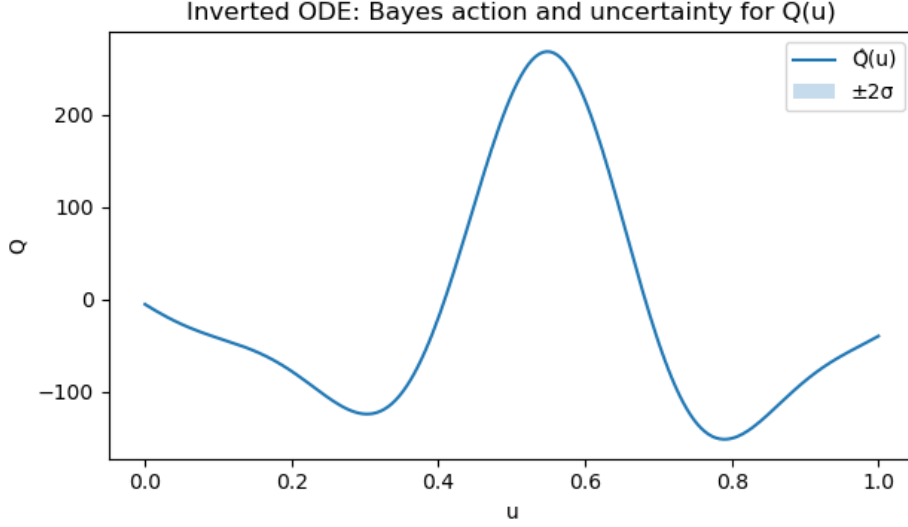


Figure 2: Bayes action and Uncertainty for the 3'rd order linear ODE.

where $A_\alpha(x) \in \mathbb{R}^{q \times p}$ are known coefficient matrices, and $Q : \Omega \rightarrow \mathbb{R}^q$ is the unknown source/forcing term:

$$\mathcal{L}y = Q \quad \text{in } \Omega, \quad \mathcal{B}y = b \quad \text{on } \partial\Omega,$$

with \mathcal{B} a linear boundary operator (e.g. Dirichlet/Neumann/Robin) and known boundary data b .

Prior. Place a (possibly vector-valued) GP prior on y :

$$y \sim \text{GP}(0, K), \quad K : \Omega \times \Omega \rightarrow \mathbb{R}^{p \times p},$$

with mean-square derivatives up to order m in each argument. Popular choices include separable kernels $K(x, x') = k(x, x') \otimes \Sigma$ for $\Sigma \in \mathbb{R}^{p \times p}$, or anisotropic products $k(x, x') = \prod_{j=1}^d k_j(x_j, x'_j)$.

Observations (interior + boundary, general linear). Let \mathcal{H} be any linear operator mapping y to \mathbb{R}^m : pointwise values $y(x_i)$, derivatives $D^\beta y(x_i)$, integrals $\int_\Omega w(x)^\top y(x) dx$, or boundary traces $(\mathcal{T}y)(\xi_i)$ at $\xi_i \in \partial\Omega$ (Dirichlet/Neumann/Robin). Assume

$$\mathbf{z} = \mathcal{H}[y] + \varepsilon, \quad \varepsilon \sim \mathbf{N}(0, R).$$

Noiseless constraints (e.g. exact boundary data) are obtained by taking very small diagonal entries in R for the corresponding rows.

4.2 Decision problem and Bayes action for $Q = \mathcal{L}y$

Under squared loss over Ω ,

$$\mathcal{R}(\hat{Q}) = \mathbb{E} \left[\int_{\Omega} \|\mathcal{L}y(x) - \hat{Q}(x)\|_2^2 dx \mid \mathbf{z} \right],$$

the pointwise Bayes action is

$$\hat{Q}(x) = \mathbb{E}[(\mathcal{L}y)(x) \mid \mathbf{z}] = \sum_{|\alpha| \leq m} A_{\alpha}(x) \mathbb{E}[D^{\alpha}y(x) \mid \mathbf{z}].$$

Uncertainty is given by the posterior covariance

$$\text{Cov}(Q(x), Q(x') \mid \mathbf{z}) = \text{Cov}((\mathcal{L}y)(x), (\mathcal{L}y)(x') \mid \mathbf{z}).$$

4.3 Block-Gaussian conditioning with multi-indices

Let $X_* = \{x_1, \dots, x_s\} \subset \Omega$ be query points for Q . Stack the needed derivatives of y at X_* up to order m :

$$\mathbf{y}_* = [D^{\alpha}y(X_*)]_{|\alpha| \leq m} \in \mathbb{R}^{p(N_{\alpha})s}, \quad N_{\alpha} = \#\{\alpha : |\alpha| \leq m\}.$$

Let Z denote the set of functionals used by \mathcal{H} . By linearity, $(\mathbf{y}_*, \mathcal{H}[y])$ are jointly Gaussian with

$$\mathbf{y}_* \mid \mathbf{z} \sim \mathcal{N}(\mu_*, \Sigma_*), \quad \mu_* = K_{*Z}(K_{ZZ} + R)^{-1}\mathbf{z}, \quad \Sigma_* = K_{**} - K_{*Z}(K_{ZZ} + R)^{-1}K_{Z*}.$$

The blocks are obtained by differentiating the kernel with respect to multi-indices:

$$\begin{aligned} [K_{**}]_{(\alpha, i), (\beta, j)} &= D_x^{\alpha} D_{x'}^{\beta} K(x_i, x_j), \\ [K_{*Z}]_{(\alpha, i), r} &= \mathbb{E}[D^{\alpha}y(x_i) (\mathcal{H}[y])_r] = \mathcal{H}_{x'}[D_x^{\alpha}K(x_i, x')], \\ K_{ZZ} &= \mathbb{E}[\mathcal{H}[y] \mathcal{H}[y]^{\top}] = \mathcal{H}_x \mathcal{H}_{x'}[K(x, x')]. \end{aligned}$$

Here \mathcal{H}_x and $\mathcal{H}_{x'}$ denote applying \mathcal{H} to the first/second kernel argument, respectively; boundary traces apply via standard trace operators (e.g. Dirichlet $y|_{\partial\Omega}$, Neumann $\partial_n y|_{\partial\Omega}$).

Linear map to Q . Define L_* so that, for each x_i ,

$$(L_* \mathbf{y}_*)_i = \sum_{|\alpha| \leq m} A_{\alpha}(x_i) D^{\alpha}y(x_i).$$

Then

$$\mathbf{Q}_* \mid \mathbf{z} \sim \mathcal{N}(L_* \mu_*, L_* \Sigma_* L_*^{\top}), \quad \hat{Q} = L_* \mu_*.$$

4.4 Boundary operators \mathcal{B} as observations

Dirichlet, Neumann, and Robin (or mixed) conditions are all linear:

$$\mathcal{B}y = B_0 y + \sum_{j=1}^d B_j \partial_{x_j} y \quad \text{on } \partial\Omega,$$

with matrix-valued B_j possibly depending on the boundary location. Encode these as rows of \mathcal{H} at boundary points (or boundary integrals) with tiny noise for exactness. Their covariance contributions come from tracing the kernel to $\partial\Omega$, e.g. $\mathbb{E}[y(\xi) y(\xi')^{\top}] = K(\xi, \xi')$ for $\xi, \xi' \in \partial\Omega$, $\mathbb{E}[\partial_n y(\xi) y(\xi')^{\top}] = \partial_{n_{\xi}} K(\xi, \xi')$, etc.

4.5 Computational structure and separable kernels

If K is separable across dimensions, $k(x, x') = \prod_{j=1}^d k_j(x_j, x'_j)$, then many derivative blocks factorize, and on tensor grids we can exploit Kronecker algebra for K_{**}, K_{*Z}, K_{ZZ} . For vector-valued y with $K(x, x') = k(x, x') \otimes \Sigma$, use block-diagonalization via the eigen-decomposition of Σ .

4.6 Two templates

(i) **Poisson** For scalar y and

$$\Delta y(x) = Q(x) \quad \text{in } \Omega,$$

we have $m = 2$ with $A_{e_j+e_j}(x) = -\kappa(x)$ (second derivatives), cross-terms from spatially varying κ via product rule if desired Build L_* using second-order derivatives $D^\alpha y$ for $|\alpha| \leq 2$. Dirichlet or Neumann boundaries become \mathcal{H} -rows on $\partial\Omega$.

(ii) **Advection–Diffusion–Reaction.**

$$-\nabla \circ (D(x)\nabla y) + \mathbf{v}(x) \circ \nabla y + \rho(x)y = Q.$$

Here $m = 2$; include first- and second-order derivative blocks for ∇y and $\nabla\nabla y$, assemble L_* with the appropriate $A_\alpha(x)$ derived from D, \mathbf{v}, ρ .

4.7 Practical Implementation (PDE)

1. Choose m , coefficient matrices $A_\alpha(x)$, and boundary operator \mathcal{B} .
2. Pick K with mean-square derivatives up to m (e.g. RBF; Matérn with $\nu > m + d/2$).
3. Specify \mathcal{H} including boundary rows; set R (tiny variance for exact constraints).
4. Build kernel blocks by applying D_x^α, D_x^β , traces, and integrals to K .
5. Compute μ_*, Σ_* and return $\hat{\mathbf{Q}} = L_*\mu_*, \text{Cov}(\mathbf{Q}_*) = L_*\Sigma_*L_*^\top$.

4.8 Example: Poisson Boundary Conditions

We illustrate the example on the Poisson equation. We sample 250 values uniformly from $[0, 1] \times [0, 1]$ evaluate the following formula on the values:

$$f(x, y) = 5 * \sin(2\pi x) \cos(2\pi y) + 2 * \sin(6\pi xy) \tag{3}$$

and add *iid* Gaussian noise $z = f(x, y) + \mathcal{N}(0, 2)$. An image depicting the data is available in Figure 3.




Figure 3: The data distributed on a 2 dimensional grid which the Gaussian process $y|z$ on \mathbb{R}^2 that the Gaussian Process is fit to.

We then fit a Gaussian process with radial basis function kernel (RBF) given by:

$$k(a, b; \sigma^2, s) = \sigma^2 \exp\left[-\frac{\|a - b\|^2}{2s^2}\right] \quad (4)$$

This is fit to the data using a zero mean GP with RBF kernel and we get a posterior process $y|z \sim N(\mu_*, \sigma_*)$.

Laplacian coefficients are built over a 50×50 uniform grid on $[0, 1] \times [0, 1]$ using finite differences. This giving us an approximation to the Laplacian $\hat{\Delta}$; which is then applied to the original GP to give $\hat{\Delta}(y|z) \sim N(\hat{\Delta}(\mu), \hat{\Delta}\sigma_*\hat{\Delta}^\top)$. The mean of this is the Bayes risk under squared error loss for the source term. A plot illustrating the mean is given in Figure 4. We note that by Theorem 2 this is equivalent to the 0-1 loss minimizer, and by Theorem 1 it is an approximation on the grid projection of the Hilbert space Bayes action for the Laplacian operator (as was equivalent in Section /3.8).




Figure 4: Bayes action of of the resultant Dirichlet loss resulting from taking the mean of the Gaussian process after applying the Laplacian operator.

For uncertainty estimates we take the trace of the discretized operator $\Lambda_*\sigma_*\Lambda_*^\top$ and evaluate the $\mu_* \pm 2\hat{\sigma}$. This is provide on the discretized $[0, 1] \times [0, 1]$ in Figure 5.

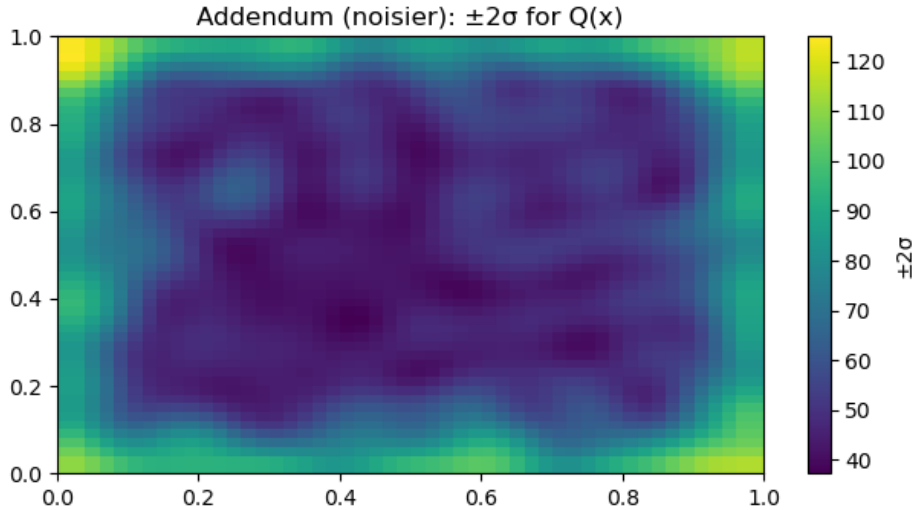


Figure 5: Uncertainty around the Bayes action after application of the Dirichlet operator.

5 Performance guarantees for nonlinear operators via Frechet derivatives

Let \mathcal{Y}, \mathcal{Q} be Hilbert spaces equipped with inner products $\langle \cdot, \cdot \rangle_{\mathcal{Y}}$ and $\langle \cdot, \cdot \rangle_{\mathcal{Q}}$, respectively. Suppose we are interested in a (generally nonlinear) operator $\mathcal{N} : \mathcal{Y} \rightarrow \mathcal{Q}$ and the quantity $Q = \mathcal{N}(y)$. From the Gaussian process (GP) posterior over the latent field y given data z we have

$$y \mid z \sim \mathcal{N}(\mu_{\star}, \Sigma_{\star}), \quad \widehat{Q} := \mathcal{N}(\mu_{\star}).$$

Writing the posterior error as $e := y - \mu_{\star}$, we expand \mathcal{N} around μ_{\star} using Frechet calculus.

5.1 First-order linearization and Bayes risk

If \mathcal{N} is Frechet differentiable at μ_{\star} with derivative $D\mathcal{N}[\mu_{\star}] \in \mathcal{L}(\mathcal{Y}, \mathcal{Q})$, then

$$\mathcal{N}(\mu_{\star} + e) = \mathcal{N}(\mu_{\star}) + D\mathcal{N}[\mu_{\star}](e) + R_2(e), \quad \|R_2(e)\|_{\mathcal{Q}} = o(\|e\|_{\mathcal{Y}}). \quad (5)$$

Under squared loss the first-order approximation of the posterior Bayes risk is

$$\mathbb{E}[\|Q - \widehat{Q}\|_{\mathcal{Q}}^2 \mid z] \approx \mathbb{E}[\|D\mathcal{N}[\mu_{\star}](e)\|_{\mathcal{Q}}^2 \mid z] = \text{tr}(D\mathcal{N}[\mu_{\star}] \Sigma_{\star} D\mathcal{N}[\mu_{\star}]^*). \quad (6)$$

Hence a conservative spectral bound is

$$\mathbb{E}[\|Q - \widehat{Q}\|_{\mathcal{Q}}^2 \mid z] \leq \|D\mathcal{N}[\mu_{\star}]\|_{\text{op}}^2 \text{tr}(\Sigma_{\star}), \quad (7)$$

where $\|\cdot\|_{\text{op}}$ denotes the operator norm and $\text{tr}(\cdot)$ the Hilbert–Schmidt trace.

5.2 High-probability certificate

Let d be the dimension of a finite discretization and define the $(1 - \alpha)$ -credible ellipsoid $\mathcal{E}_{\alpha} = \{e : e^{\top} \Sigma_{\star}^{-1} e \leq \chi_{d,\alpha}^2\}$. Then, with probability at least $1 - \alpha$,

$$\|Q - \widehat{Q}\|_{\mathcal{Q}} \lesssim \|D\mathcal{N}[\mu_{\star}]\|_{\text{op}} \sqrt{\lambda_{\max}(\Sigma_{\star})} \sqrt{\chi_{d,\alpha}^2} + \sup_{e \in \mathcal{E}_{\alpha}} \|R_2(e)\|_{\mathcal{Q}}. \quad (8)$$

5.3 Second-order curvature control

If \mathcal{N} is twice Frechet differentiable in a neighborhood of μ_{\star} and $\|D^2\mathcal{N}[y]\|_{\text{op}} \leq M$ there, the remainder admits the quadratic bound

$$\|R_2(e)\|_{\mathcal{Q}} \leq \frac{1}{2} M \|e\|_{\mathcal{Y}}^2. \quad (9)$$

Taking expectations yields

$$\mathbb{E}[\|Q - \widehat{Q}\|_{\mathcal{Q}}^2 \mid z] \leq \|D\mathcal{N}[\mu_{\star}]\|_{\text{op}}^2 \text{tr}(\Sigma_{\star}) + \frac{1}{2} M^2 \mathbb{E}\|e\|_{\mathcal{Y}}^4 + \text{cross terms}, \quad (10)$$

where for Gaussian e the fourth moment is determined by $\text{tr}(\Sigma_{\star})$ and $\|\Sigma_{\star}\|_F^2$.

5.4 Lipschitz (global) guarantees

If \mathcal{N} is (locally) Lipschitz with constant L near μ_{\star} ,

$$\|Q - \widehat{Q}\|_{\mathcal{Q}} \leq L \|e\|_{\mathcal{Y}}, \quad \mathbb{E}[\|Q - \widehat{Q}\|_{\mathcal{Q}}^2 \mid z] \leq L^2 \text{tr}(\Sigma_{\star}). \quad (11)$$

Typically one can bound L via the derivative as $L \geq \sup_{y \in \mathcal{U}} \|D\mathcal{N}[y]\|_{\text{op}}$ on a neighborhood \mathcal{U} .

5.5 Operator templates

Pointwise nonlinearity $\mathcal{N}(y) = f \circ y$. For smooth $f : \mathbb{R} \rightarrow \mathbb{R}$, $[D\mathcal{N}[\mu_{\star}](h)](x) = f'(\mu_{\star}(x)) h(x)$. Hence $\|D\mathcal{N}[\mu_{\star}]\|_{\text{op}} = \|f'(\mu_{\star})\|_{L^{\infty}}$. If f'' is bounded by M_f , then $M \leq \|f''(\mu_{\star})\|_{L^{\infty}}$ in (9).

Bilinear convective term $\mathcal{N}(y) = B(y, y)$. For a bounded bilinear form B , $D\mathcal{N}[y](h) = B(h, y) + B(y, h)$, so $\|D\mathcal{N}[\mu_{\star}]\|_{\text{op}} \leq 2\|B\|_{\text{op}}\|\mu_{\star}\|$ and $R_2(e) = B(e, e)$ obeys $\|R_2(e)\| \leq \|B\|_{\text{op}}\|e\|^2$.

Coefficient-nonlinear elliptic term $\mathcal{N}(y) = -\nabla \cdot (a(y)\nabla y)$. The derivative is

$$D\mathcal{N}[\mu_\star](h) = -\nabla \cdot (a'(\mu_\star) h \nabla \mu_\star + a(\mu_\star) \nabla h),$$

leading to bounds (up to domain-dependent constants) of the form $\|D\mathcal{N}[\mu_\star]\|_{\text{op}} \lesssim \|a'(\mu_\star)\|_{L^\infty} \|\nabla \mu_\star\|_{L^\infty} + \|a(\mu_\star)\|_{L^\infty}$.

5.6 Reporting and design

For transparency and reproducibility, we recommend reporting:

1. the first-order covariance $\Sigma_Q^{(1)} := D\mathcal{N}[\mu_\star] \Sigma_\star D\mathcal{N}[\mu_\star]^*$ and its trace/spectral radius;
2. an estimate of $\|D\mathcal{N}[\mu_\star]\|_{\text{op}}$ (e.g., power iteration);
3. a curvature constant M (or an empirical ratio $\|R_2(e)\|/\|D\mathcal{N}[\mu_\star](e)\|$ on samples $e \sim \mathcal{N}(0, \Sigma_\star)$);
4. a high-probability radius r_α from (8).

These quantities also drive active data acquisition by targeting reductions in $\text{tr}(\Sigma_Q^{(1)})$ or the spectral radius of $\Sigma_Q^{(1)}$.

6 Acknowledgements

Code available at the github repo ****INSERT LINK HERE**** and was generated with the help of GPT5.

7 Notes for me:

Regularity. Interchange of derivatives/expectation and trace operations is justified under standard GP regularity; ensure K is sufficiently smooth and the trace/normal derivatives are defined (e.g. via Sobolev embeddings).

Learning. Hyperparameters of K and any unknown coefficients in A_α can be learned by marginal likelihood or hierarchical Bayes using K_{ZZ} .

Vector fields. For $p, q > 1$, use matrix-valued kernels or LMC constructions; the formulas above apply verbatim with matrix blocks.

Nonconstant coefficients. All formulas above already allow $a_k = a_k(u)$ varying with u (they only enter via the diagonal A_k at test points).

Model selection. We can learn kernel hyperparameters by maximizing the marginal likelihood of \mathbf{z} (standard GP training); $\hat{\mathbf{Q}}$ then updates accordingly. Possibly extend to quantification of overfitting of data in this context and associated Bayes Risk perturbations.

Design/active sampling. Since we have $\text{Cov}(\mathbf{Q})$, we can design new measurements to reduce uncertainty in Q , e.g. choose new times to minimize $\text{tr}(L_\star \Sigma_\star L_\star^\top)$.

A Proof of Theorem 1

Statement 1. $(\mathcal{H}, \langle \circ, \circ \rangle)$ be a (real, separable) Hilbert space with norm $\|\circ\|$. Define a loss function $\Psi : (a, \mathcal{H}) \rightarrow \mathbb{R}$ for $a \in \mathcal{H}$, and assume Ψ on $\mathcal{H} \times \mathcal{H}$. Assume that there exists a posterior process $\mathcal{Q}|z$ with elements almost surely $q \sim \mathcal{Q}|z \in \mathcal{H}$. Let $\{H_n\}_{n \geq 1}$ be an increasing sequence of finite dimensional subspaces $H_n \uparrow \mathcal{H}$ and let P_n be the corresponding projection operator $P_n : \mathcal{H} \rightarrow H_n$. Suppose the following inequality hold for some $C, L > 0$ and $p \geq 1$. (Globally Lipschitz and linear growth):

$$|\Psi(a, h) - \Psi(a, h')| \leq L \|h - h'\| \tag{12}$$

$$|\Psi(a, h)| \leq C(1 + \|a\|^p + \|h\|^p) \tag{13}$$

And $\mathbb{E}_{\mathcal{Q}|z}[\|Q\|^p|z] < \infty$. Then $\mathbb{E}_{\mathcal{Q}}[|\Psi(a_n, P_n Q)|] \xrightarrow{n \rightarrow \infty} \mathbb{E}_{\mathcal{Q}}[\Psi(a, Q)]$.

If $\lim_{n \rightarrow \infty} a_n = a \in \mathcal{H}$ exists and is finite, then it is the Bayes action for $a = \text{argmin}_{a \in \mathcal{H}} \mathbb{E}_{\mathcal{Q}}[\Psi(a, Q)|z]$

Proof. 1. *Stability in the first coordinate.* By continuity in the first coordinate ($a_n \in H_n$) and the growth bound:

$$|\Psi(a_n, Q) - \Psi(a, Q)| \leq G(Q) \|a_n - a\|_{H_n} \quad (14)$$

for some integrable $G(Q)$ ($G(Q) = C(1 + \|Q\|^p)$). Hence $\mathbb{E}_{Q|z}[\Psi(a_n, Q) - \Psi(a, Q)] \rightarrow 0$. so $\mathbb{E}_{Q|z}[\Psi(a_n, Q)] \rightarrow \mathbb{E}_{Q|z}[\Psi(a, Q)]$

2. *Stability in the second coordinate.* Add and subtract:

$$R_n(a_n) = \mathbb{E}_{Q|z}[\Psi(a_n, Q_n)] = \mathbb{E}_{Q|z}[\Psi(a_n, Q)] + \mathbb{E}_{Q|z}[\Psi(a_n, Q_n)] - \mathbb{E}_{Q|z}[\Psi(a_n, Q)] \quad (15)$$

By assumption $|\Psi(a_n, Q_n) - \Psi(a_n, Q)| \leq L\|Q_n - Q\|$, so:

$$\mathbb{E}_{Q|z}[\Psi(a_n, Q_n) - \Psi(a_n, Q)] \leq L\mathbb{E}_{Q|z}[P_n Q - Q] \rightarrow 0 \quad (16)$$

Coming steps 1 and 2 gives $R_n(a_n) \rightarrow R(a)$.

3. *Optimality of the limit:* For any candidate $b \in \mathcal{H}$ let $b_n = P_n(b) \in H_n$. Since a_n minimizes R_n in H_n we have:

$$R_n(a_n) \leq R_n(b_n) \quad (17)$$

Taking limits as $n \rightarrow \infty$ and using continuity of the risk functional we get:

$$\mathbb{E}_{\mathcal{Q}}[\Psi(a, Q)|z] \leq \mathbb{E}_{\mathcal{Q}}[\Psi(b, Q)|z] \quad \forall b \in \mathcal{H} \quad (18)$$

□

B Proof of Theorem 2

Statement 2 (0-1 vs. squared-error Bayes equivalence in Gaussian Process Measures on Hilbert spaces). *Let $(\mathcal{H}, \langle \cdot, \cdot \rangle)$ be a (real, separable) Hilbert space with norm $\|\cdot\|$, and suppose the posterior distribution of a random element $Q \in \mathcal{H}$ given data z is Gaussian:*

$$Q|z \sim N(\mu, \mathcal{C}),$$

with mean $\mu \in \mathcal{H}$ and covariance operator $\mathcal{C} : \mathcal{H} \rightarrow \mathcal{H}$ (self-adjoint, positive, trace-class on the closed linear span of the support). Suppose that μ is almost surely finite for $\mu(x) < \infty \quad \forall x \in \Omega$. For an action $a \in \mathcal{H}$ consider the two losses.

$$L_2(q, a) = \|q - a\|^2, \quad L_{01,\tau}(q, a) = \mathbf{1}\{\|q - a\| \geq \tau\},$$

where $\tau > 0$ is any fixed radius. Then the Bayes action under both losses equals the posterior mean:

$$a^* = \arg \min_{a \in \mathcal{H}} \mathbb{E}[L_2(Q, a) | z] = \arg \min_{a \in \mathcal{H}} \mathbb{E}[L_{01,\tau}(Q, a) | z] = \mu.$$

Consequently, the Bayes rules under L_2 and under $L_{01,\tau}$ (for any $\tau > 0$) are identical.

Proof. We rely on Corollary 3.5 introduced in the review Alberts and Khoshnevisan [2018] and Theorem 1.

Squared loss. Write $Q = \mu + X$ with $X|z \sim N(0, \mathcal{C})$. Then

$$\mathbb{E}[\|Q - a\|^2 | z] = \mathbb{E}[\|\mu - a + X\|^2] = \mathbb{E}[\langle \mu - a, \mu - a \rangle + \langle X, X \rangle + 2\langle \mu - a, X \rangle]$$

where $\mathbb{E}\|X\|^2 = \text{tr}(\mathcal{C})$ (possibly $+\infty$ on components outside the effective support, but independent of a). The term $\langle \mu - a, \mu - a \rangle$ is minimized at $a = \mu$. $\langle \mu - a, X \rangle \rightarrow \langle \mu - a, \mathbb{E}_X[X] \rangle - \langle \mu - a, 0 \rangle = 0$ is minimized at $a = \mu$ by noting that the operator \mathcal{C} is self adjoint and thus the expectation of this term (noting the X is zero mean) is going to be symmetric about 0. This gives us the desired solution.

0-1 (ball) loss. As in Theorem 1, fix $\{H_n\}_{n \geq 1} \xrightarrow{n \rightarrow \infty} \mathcal{H}$. We note that this gives us a Gaussian measure on H_n for each $n \in \mathbb{N}$. For fixed $\tau > 0$, the Bayes risk is

$$R_{01,\tau}(a) = \mathbb{P}(\|Q_n - a_n\| > \tau | z) = 1 - \mathbb{P}(\|Q_n - a_n\| \leq \tau | z).$$

Which has a solution at the posterior median by Anderson's inequality (Corollary 3.5 in Alberts and Khoshnevisan [2018]). This always exists for sufficiently regular Gaussian distributions (finite mean and positive variance). This gives us that $a_n = P_n(\mu)$. Note that we are in a real, separable Hilbert space, and that the 0 - 1, so we examine the requirements of Theorem 1.

0-1 small-ball loss using ϵ -approximations to 01-loss The direct application of Theorem 1 fails because the indicator loss $L_{01,\tau}(q, a) = \mathbf{1}\{\|q - a\| > \tau\}$ is not Lipschitz in its second argument. To address this, we introduce a Lipschitz *ramp* approximation that satisfies the conditions in Equation 12-13 and then pass to the limit $\epsilon \downarrow 0$.

For $\epsilon > 0$, define the upper and lower ϵ -ramps

$$\ell_\epsilon^+(r) = \begin{cases} 0, & r \leq \tau, \\ \frac{r - \tau}{\epsilon}, & \tau < r < \tau + \epsilon, \\ 1, & r \geq \tau + \epsilon, \end{cases} \quad \ell_\epsilon^-(r) = \begin{cases} 0, & r \leq \tau - \epsilon, \\ \frac{r - (\tau - \epsilon)}{\epsilon}, & \tau - \epsilon < r < \tau, \\ 1, & r \geq \tau. \end{cases} \quad (19)$$

Each ℓ_ϵ^\pm is globally $1/\epsilon$ -Lipschitz and bounded in $[0, 1]$, hence satisfies inequalities 12-13 with $L = 1/\epsilon$ and $C = 1$. Define corresponding losses

$$L_\epsilon^\pm(q, a) := \ell_\epsilon^\pm(\|q - a\|), \quad R_\epsilon^\pm(a) := \mathbb{E}[L_\epsilon^\pm(Q, a) \mid z].$$

By Theorem 1, each R_ϵ^\pm admits a Bayes action $a_\epsilon^\pm = \mu$ since, for fixed ϵ , L_ϵ^\pm is continuous and depends on Q only through the centered Gaussian $X = Q - \mu$; thus Anderson's inequality implies that $a = \mu$ maximizes $\mathbb{P}(\|Q - a\| \leq r)$ for every $r > 0$ in finite-dimensional projections and hence in H .

Finally, note that

$$L_\epsilon^-(q, a) \uparrow L_{01,\tau}(q, a) \downarrow L_\epsilon^+(q, a) \quad (\epsilon \downarrow 0),$$

and by dominated and monotone convergence together with $\mathbb{P}(\|Q - a\| = \tau) = 0$ (for nondegenerate Gaussian Q),

$$\lim_{\epsilon \downarrow 0} R_\epsilon^\pm(a) = \mathbb{E}[L_{01,\tau}(Q, a) \mid z] = R_{01,\tau}(a).$$

Because each R_ϵ^\pm is minimized uniquely at $a = \mu$, the same holds in the limit:

$$a^* = \arg \min_{a \in H} R_{01,\tau}(a) = \mu.$$

Hence the Bayes action under the small-ball loss coincides with the posterior mean, completing the proof. \square

C Proof of Theorem 3

Statement 3 (Decision boundary for two Gaussian-process posteriors). *Let \mathcal{H} be a real, separable Hilbert space. Let $Y \sim \mathcal{N}(\mu, \mathcal{C})$ and $Y' \sim \mathcal{N}(\mu', \mathcal{C}')$ be Gaussian random elements in \mathcal{H} with means $\mu, \mu' \in \mathcal{H}$ and covariance operators $\mathcal{C}, \mathcal{C}'$ (self-adjoint, positive, trace-class on the closed span of their supports). For any $a \in \mathcal{H}$, define the small-ball 0-1 risks*

$$R_{01,\tau}^Y(a) := \mathbb{E}[\mathbf{1}\{\|Y - a\| > \tau\}], \quad R_{01,\tau}^{Y'}(a) := \mathbb{E}[\mathbf{1}\{\|Y' - a\| > \tau\}].$$

Define the decision boundary as the zero-level set of the limiting risk difference,

$$\mathcal{B} := \left\{ a \in \mathcal{H} : \lim_{\tau \downarrow 0} (R_{01,\tau}^Y(a) - R_{01,\tau}^{Y'}(a)) = 0 \right\}.$$

Then \mathcal{B} is a Hilbert manifold in \mathcal{H} , and coincides with the zero-level set of the squared-loss risk difference,

$$\mathcal{B} = \left\{ a \in \mathcal{H} : \mathbb{E}\|Y - a\|^2 = \mathbb{E}\|Y' - a\|^2 \right\},$$

and therefore admits the explicit affine description

$$\|a - \mu\|^2 - \|a - \mu'\|^2 = \text{tr}(\mathcal{C}') - \text{tr}(\mathcal{C}) \iff \frac{\|a - \mu\|^2 - \|a - \mu'\|^2}{\text{tr}(\mathcal{C}') - \text{tr}(\mathcal{C})} = 1.$$

Which is an affine Hilbert manifold in \mathcal{H} locally diffeomorphic to \mathcal{H} . In particular, when $\mathcal{C} = \mathcal{C}'$, the boundary is the (codimension-one) bisector hyperplane Hilbert manifold in \mathcal{H}

$$\{ a \in \mathcal{H} : \langle a, \mu' - \mu \rangle = \frac{1}{2}(\|\mu'\|^2 - \|\mu\|^2) \},$$

i.e. the locus of points equidistant (in $\|\cdot\|$) from μ and μ' .

Proof. The proof proceeds with the following steps:

1. **Squared-loss risks.** For any Gaussian $Z \sim \mathcal{N}(m, \Sigma)$ in \mathcal{H} ,

$$\mathbb{E}\|Z - a\|^2 = \|m - a\|^2 + \text{tr}(\Sigma).$$

Applying this to Y and Y' yields

$$\mathbb{E}\|Y - a\|^2 - \mathbb{E}\|Y' - a\|^2 = (\|a - \mu\|^2 - \|a - \mu'\|^2) + (\text{tr}(\mathcal{C}) - \text{tr}(\mathcal{C}')).$$

Hence the squared-loss boundary satisfies

$$\|a - \mu\|^2 - \|a - \mu'\|^2 = \text{tr}(\mathcal{C}') - \text{tr}(\mathcal{C}).$$

2. **Equivalence of small-ball and squared-loss risks.** By Theorem 2, for a Gaussian process posterior $Q \sim \mathcal{N}(m, \Sigma)$, the Bayes action minimizing either squared loss $\|Q - a\|^2$ or small-ball 0–1 loss $\mathbf{1}\{\|Q - a\| > \tau\}$ (for fixed $\tau > 0$) is the posterior mean $a^* = m$. Thus, for both (Y, μ) and (Y', μ') , the ordering of risks under the small-ball loss agrees with that under squared loss as $\tau \downarrow 0$.

3. **Equality of risks.** Define

$$\Delta_2(a) := \mathbb{E}\|Y - a\|^2 - \mathbb{E}\|Y' - a\|^2, \quad \Delta_{01,\tau}(a) := R_{01,\tau}^Y(a) - R_{01,\tau}^{Y'}(a).$$

By Step 1, Δ_2 is affine in a , and by Step 2, $\text{sign}(\Delta_{01,\tau}(a)) = \text{sign}(\Delta_2(a))$ for small τ . Consequently,

$$\lim_{\tau \downarrow 0} \Delta_{01,\tau}(a) = 0 \iff \Delta_2(a) = 0.$$

Hence \mathcal{B} equals the affine set described above.

4. **Manifold structure.** Unless $\mu = \mu'$ and $\text{tr}(\mathcal{C}) = \text{tr}(\mathcal{C}')$, the boundary is given by one nontrivial affine equation in \mathcal{H} , so it is an affine hyperplane (locally diffeomorphic to \mathcal{H}).

□

D More Simulated Examples

References

- Tom Alberts and Davar Khoshnevisan. Calculus on gauss space: An introduction to gaussian analysis, 2018. URL <https://www.math.utah.edu/~davar/math7880/F18/GaussianAnalysis.pdf>. Version 5.0, Lecture notes.
- James O Berger. *Statistical decision theory and Bayesian analysis*. Springer, 2010.
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10:273–304, 1995.
- Jon Cockayne, Chris J. Oates, T. J. Sullivan, and Mark Girolami. Bayesian probabilistic numerical methods. *SIAM Review*, 61(4):756–789, 2019.
- Masoumeh Dashti and Andrew M. Stuart. The bayesian approach to inverse problems, 2015. URL <https://arxiv.org/abs/1302.6989>.
- Pierpaolo De Blasi, Stefano Favaro, Antonio Lijoi, Ramses H. Mena, Igor Prunster, and Matteo Ruggiero. Are gibbs-type priors the most natural generalization of the dirichlet process? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):212–229, February 2015. ISSN 2160-9292. doi: 10.1109/tpami.2013.217. URL <http://dx.doi.org/10.1109/TPAMI.2013.217>.
- Jari Kaipio and Erkki Somersalo. *Statistical and Computational Inverse Problems*. Springer, 2004.
- Jari Kaipio, Ville Kolehmainen, Erkki Somersalo, and Marko Vauhkonen. Statistical inversion and monte carlo sampling methods in electrical impedance tomography. *Inverse Problems*, 23(6):P1–P39, 2007.
- J F C Kingman. Completely random measures. *Pacific Journal of Mathematics*, 21(1):59–78, 1967. URL <https://msp.org/pjm/1967/21-1/pjm-v21-n1-p06-p.pdf>.
- Michael Lavine. Polya trees and random distributions. *The Annals of Statistics*, 20(3):1222–1235, 1992.
- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear pdes. *Journal of Computational Physics*, 378:686–707, 2019.