

UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

Evaluating Gaussian process metamodels and sequential designs for noisy level set estimation

Permalink

<https://escholarship.org/uc/item/1t5601rr>

Journal

STATISTICS AND COMPUTING, 31(4)

ISSN

0960-3174

Authors

Lyu, Xiong
Binois, Mickael
Ludkovski, Michael

Publication Date

2021

DOI

10.1007/s11222-021-10014-w

Peer reviewed

Evaluating Gaussian Process Metamodels and Sequential Designs for Noisy Level Set Estimation

Xiong Lyu*

Mickaël Binois[†]Michael Ludkovski[‡]

March 3, 2020

Abstract

We consider the problem of learning the level set for which a noisy black-box function exceeds a given threshold. To efficiently reconstruct the level set, we investigate Gaussian process (GP) metamodels. Our focus is on strongly stochastic samplers, in particular with heavy-tailed simulation noise and low signal-to-noise ratio. To guard against noise misspecification, we assess the performance of three variants: (i) GPs with Student- t observations; (ii) Student- t processes (TPs); and (iii) classification GPs modeling the sign of the response. In conjunction with these metamodels, we analyze several acquisition functions for guiding the sequential experimental designs, extending existing stepwise uncertainty reduction criteria to the stochastic contour-finding context. This also motivates our development of (approximate) updating formulas to efficiently compute such acquisition functions. Our schemes are benchmarked by using a variety of synthetic experiments in 1–6 dimensions. We also consider an application of level set estimation for determining the optimal exercise policy of Bermudan options in finance.

1 Introduction

1.1 Statement of Problem

Metamodeling has become widespread for approximating expensive black-box functions that arise in applications ranging from engineering to environmental science and finance (Santner et al., 2013). Rather than aiming to capture the precise shape of the function over the entire region, in this article we are interested in estimating the *level set* where the function exceeds some particular threshold. Such problems are common in cases where we need to quantify the reliability of a system or its performance relative to a benchmark. It also arises intrinsically in control frameworks where one wishes to rank the pay-off from several available actions (Hu and Ludkovski, 2017).

We consider a setup where the latent $f : D \rightarrow \mathbb{R}$ is a continuous function over a d -dimensional input space $D \subseteq \mathbb{R}^d$. The level-set estimation problem consists in classifying every input $x \in D = S \cup N$

*Department of Statistics and Applied Probability, University of California UCSB, Santa Barbara, CA 93106-3110, USA lyu@pstat.ucsb.edu

[†]Argonne National Laboratory, Mathematics and Computer Science Division, 9700 South Cass Avenue, Lemont, IL 60439, USA mickaël.binois@inria.fr

[‡]Department of Statistics and Applied Probability, University of California UCSB, Santa Barbara, CA 93106-3110, USA ludkovski@pstat.ucsb.edu

according to

$$S = \{x \in D : f(x) \geq 0\}, \quad N = \{x \in D : f(x) < 0\}. \quad (1.1)$$

Without loss of generality the threshold is taken to be zero, so that the level set estimation is equivalent to learning the sign of the response function f . For later use we also define the corresponding zero-contour of f , namely the partition boundary $\partial S = \partial N = \{x \in D : f(x) = 0\}$.

For any $x \in D$, we have access to a simulator $Y(x)$ that generates noisy samples of $f(x)$:

$$Y(x) = f(x) + \epsilon(x), \quad (1.2)$$

where $\epsilon(x)$ are realizations of independent, mean zero random variables with variance $\tau^2(x)$.

To assess a level-set estimation algorithm, we compare the resulting estimate \hat{S} with the true S in terms of their symmetric difference. Let μ be a probability measure on the Borel σ -algebra $\mathcal{B}(D)$ (e.g., $\mu = \text{Leb}_D$). Then our loss function is

$$L(S, \hat{S}) = \mu(S \Delta \hat{S}), \quad (1.3)$$

where $S_1 \Delta S_2 := (S_1 \cap S_2^C) \cup (S_1^C \cap S_2)$. Frequently, the inference is carried out by first producing an estimate \hat{f} of the response function; in that case we take $\hat{S} = \{x \in D : \hat{f}(x) \geq 0\}$ and rewrite the loss as

$$L(f, \hat{f}) = \int_{x \in D} \mathbb{I}(\text{sign } \hat{f}(x) \neq \text{sign } f(x)) \mu(dx), \quad (1.4)$$

where $\mathbb{I}(\cdot)$ is the indicator function.

1.2 Motivation

As a concrete example of level set estimation, consider the problem of evaluating the probability of failure, determined via the limit state S of a performance function $f(\cdot)$ (Picheny and Ginsbourger, 2013). The system is safe when $f(x) \leq h$, and fails otherwise. In the context where the performance function can be evaluated via deterministic experiments, the estimation of the safe zone (more precisely its volume $\mu(S)$) was carried out in Bect et al. (2012) and Mukhopadhyay et al. (2005) employing a Gaussian Process approach with a sequential design. A related example dealing with the probability of failure in a nuclear fissile chain reaction appeared in Chevalier et al. (2014a).

Another application, which motivated this present investigation, comes from simulation-based algorithms for valuation of Bermudan options (Gramacy and Ludkovski, 2015; Ludkovski, 2018). This problem consists of maximizing the expected reward $h(\tau, X_\tau)$ over all stopping times $\tau \in \{0, \Delta t, 2\Delta t, \dots, T\}$ bounded by the specified horizon T :

$$V(t, x) = \sup_{\tau \geq t, \tau \in \mathcal{S}} \mathbb{E}[h(\tau, X_\tau) | X_t = x], \quad (1.5)$$

where (X_t) is the underlying asset price at time t , typically satisfying a stochastic differential equation and Δt is the frequency of exercising. The approach in the so-called Regression Monte Carlo methods (Longstaff and Schwartz, 2001; Tsitsiklis and Van Roy, 2001) is to convert the decision of whether to exercise the option $\tau(t, x) = t$ or continue $\tau(t, x) > t$ when $X_t = x$ at intermediate step t , into comparing the immediate reward $h(t, x)$ vis à vis the reward-to-go $C(t, x)$. In turn this is equivalent to determining the zero level set (known as the continuation region) $S_t = \{x \in D : f(x; t) \geq 0\}$ of the timing value

$f(x; t) := C(t, x) - h(t, x)$. The stopping problem (1.5) is now solved recursively by backward induction over $t = T - \Delta t, T - 2\Delta t, \dots$, which allows noisy samples of $f(x; t)$ to be generated by simulating a trajectory $X_{t:T}^x$ emanating from x and evaluating the respective *pathwise* reward-to-go. Probabilistically, this means that we are interested in (1.2) where f corresponds to a *conditional expectation* related to a path-dependent functional of the Markov process X ; the loss function (1.3) arises naturally as a metric regarding the quality of the estimated stopping rule in terms of the underlying distribution $\mu(\cdot; t)$ of X_t . We refer to Ludkovski (2018) for a summary of existing state of the art and the connection to employing a GP metamodel for learning the timing value $T(\cdot; t)$.

1.3 Design of Experiments for Contour Finding

Reconstructing S via a metamodel can be divided into two steps: the construction of the response model and the development of methods for efficiently selecting the simulation inputs $x_{1:N}$, known as design of experiments (DoE). Since the level set is intrinsically defined in terms of the unknown f , an *adaptive* DoE approach is needed that selects x_n 's sequentially.

For the response modeling aspect, GP regression, or kriging, has emerged as the most popular nonparametric approach for both deterministic and stochastic black-box functions (Bect et al., 2012; Gramacy and Lee, 2009; Picheny et al., 2013a; Jalali et al., 2017). GPs have also been widely used for the level-set estimation problem; see Bryan and Schneider (2008); Gotovos et al. (2013); Hu and Ludkovski (2017); Picheny et al. (2010) and Ranjan et al. (2008). In a nutshell, at step n the GP paradigm constructs a metamodel $\hat{f}^{(n)}$ that is then used to guide the selection of x_{n+1} and also to construct the estimate $\hat{S}^{(n)}$. To this end, GPs are well suited for sequential design by offering a rich uncertainty quantification aspect that can be (analytically) exploited to construct information-theoretic DoE heuristics. The standard framework is to develop an acquisition function $\mathcal{I}_n(x)$ that quantifies the value of information from taking a new sample at input x conditional on an existing dataset $(x_{1:n}, y_{1:n})$ and then to myopically maximize \mathcal{I}_n :

$$x_{n+1} = \arg \max_{x \in D} \mathcal{I}_n(x). \quad (1.6)$$

Early level-set sampling criteria were proposed by Bichon et al. (2008), Picheny et al. (2010), and Ranjan et al. (2008) based on modifications to the Expected Improvement criterion (Jones et al., 1998) for response function optimization. A criterion more targeted to reduce the uncertainty about S itself was first developed by Bect et al. (2012) using the concept of stepwise uncertainty reduction (SUR). Specifically, the SUR strategy aims to myopically maximize the global learning rate about S ; see also Chevalier et al. (2014a) for related computational details. Recently, further criteria using tools from random set theory were developed in Chevalier et al. (2013); Azzimonti et al. (2016). Specifically, those works use the notions of Vorob'ev expectation and Vorob'ev deviation to choose inputs that minimize the posterior expected distance in measure between the level set S and its estimate \hat{S} . This approach is computationally expensive however, and requires conditional simulations of the posterior Gaussian field. Other works dealing with more conservative estimates are Bolin and Lindgren (2015); Azzimonti et al. (2019). Clear analysis comparing all these choices in the stochastic setting is currently lacking.

1.4 Summary of Contributions

Most of the cited papers consider only the deterministic setting without any simulation noise. The main goal of this article is to present a comprehensive assessment of GP-based surrogates for stochastic contour-finding. In that sense, our analysis complements the work of Picheny et al. (2013b) and Jalali et al.

(2017), who benchmarked GP metamodels for Bayesian optimization (BO) where the objective is to evaluate $\max_x f(x)$.

While simple versions (with constant or prespecified Gaussian noise) are easily handled, the literature on GP surrogates for complex stochastic simulators remains incomplete. Recently, several works focused on heteroskedastic simulation variance; see the Stochastic Kriging approach of Ankenman et al. (2010) and the earlier works by two of the authors (Binois et al., 2018, 2019). In the present article we instead target the non-Gaussian aspects, in particular the likely heavy-tailed property. This issue is fundamental to any realistic stochastic simulator where there is no justification for assuming Gaussian-distributed $\epsilon(x)$ (as opposed to the physical experimental setup where ϵ represents observation noise and is expected to be Gaussian thanks to the central limit theorem). This motivates us to study *alternative GP-based metamodels* for learning \hat{S} that are more robust to non-Gaussian ϵ in (1.2). In parallel, we investigate which of the contour-finding heuristics outlined above perform best in such setups.

To stay within the sequential design paradigm, we continue to work with a GP-based setup but investigate several modifications that are relevant for learning \hat{S} .

- To relax the Gaussian noise assumption, we investigate t -observation GPs (Williams and Rasmussen, 2006; Jylänki et al., 2011); the use of the Student- t likelihood nests both the heavy-tailed and Gaussian cases.
- As another non-Gaussian specification we consider Student- t processes (TPs) (Shah et al., 2014; Wang et al., 2017), as one replacement of GPs, that are also resistant to observation outliers.
- To target the classification-like objective underlying (1.3), we consider the use of classification GPs that model the sign of the response $Y(x)$ via a probit logistic model driven by a latent GP $Z(\cdot)$: $\mathbb{P}(Y(x) > 0|x) = \text{probit}(Z(x))$. Deployment of the logistic regression is expected to “wash out” non-Gaussian features in $\epsilon(x)$ beyond its effect on the sign of the observations.
- In a different vein, to exploit a structure commonly encountered in applications where the level set S is *connected*, we study the performance of *monotone* GP regression/classification metamodels (Riihimäki and Vehtari, 2010) that force f (or Z) to be monotone in the specified coordinates.

Our analysis is driven by the primal effect of noise on contour-finding algorithms. This effect was already documented in related studies, such as that of Jalali et al. (2017) who observed the strong impact of $\epsilon(\cdot)$ on performance of BO. Consequently, specialized metamodeling frameworks and acquisition functions are needed that can best handle the stochasticity for the given loss specification. Thus, the combination of the above tools with the GP framework aims to strike the best balance in carrying out uncertainty quantification and constructing a robust surrogate that is not too swayed by the simulation noise structure. In the context of GPs, this means accurate inference of the mean response and sampling noise that in turn drive the posterior mean \hat{f} and the posterior GP variance $s(x)^2$. Both of the latter ingredients are needed to blend the exploitation objective to locally learn the contour ∂S and to explore less-sampled regions. These issues drive our choices of the metamodels and also factor in developing the respective acquisition functions $\mathcal{I}_n(x)$; see cf. Section 3. On the latter front we consider four choices (MCU, cSUR, tMSE, ICU), including heuristics that depend only on the posterior standard deviation $s^{(n)}(\cdot)$, as well as those that anticipate information gain from sampling at x_{n+1} via the look-ahead standard deviation $s^{(n+1)}(\cdot)$. Because in the non-Gaussian GPs $s^{(n+1)}$ depends on $Y(x_{n+1})$, we develop tractable approximations $\hat{s}^{(n+1)}$ for that purpose, see Propositions 4.2-4.3-4.4.

To recap, our contributions can be traced along five directions. First, we investigate two ways to handle heavy-tailed simulation noise via a GP with t -observations and via TP. As far as we are aware, this is the first application of either tool in sequential design and contour-finding contexts. Second, we present an original use of monotonic GP metamodels for level set estimation. This idea is related to a gray-box approach that aims to exploit known structural properties of f (or S) so as to improve on the agnostic black-box strategies. Third, we analyze the performance of classification GP metamodels for contour-finding. This context offers an interesting and novel comparison between regression and classification approaches benchmarked against a shared loss function. Fourth, we develop and implement approximate *look-ahead* formulas for all our metamodels that are used for the evaluation of acquisition functions. To our knowledge, this is the first presentation of such formulas for non-Gaussian GPs, as well as TPs. Fifth, beyond the metamodels themselves, we also provide a detailed comparison among the proposed acquisition functions, identifying the best-performing combinations of $\mathcal{I}(\cdot)$ and metamodel \hat{f} and documenting the complex interplay between design geometry and surrogate architecture.

The rest of the article is organized as follows. Section 2 describes the metamodels we employ. Section 3 develops the sequential designs for the level-set estimation problem, and Section 4 discusses the look-ahead variance formulas for non-Gaussian GPs. Section 5 compares the models using synthetic data where ground truth is known. Two case studies from derivative pricing are investigated in Section 6. In Section 7 we summarize our conclusions.

2 Statistical Model

2.1 Gaussian Process Regression with Gaussian Noise

We begin by discussing regression frameworks for contour finding that target learning the latent $f(\cdot)$ based on the loss (1.4). The Gaussian process paradigm treats f as a random function whose posterior distribution is determined from its prior and the collected samples $\mathcal{A}_n \equiv \{(x_i, y_i), 1 \leq i \leq n\}$. We view $f(\cdot) \sim GP(m(\cdot), K(\cdot, \cdot))$, a priori, as a realization of a Gaussian process completely specified by its mean function $m(x) := \mathbb{E}[f(x)]$ and covariance function $K(x, x') := \mathbb{E}[(f(x) - m(x))(f(x') - m(x'))]$.

In the classical case (Williams and Rasmussen, 2006), the noise distribution is homoskedastic Gaussian $\epsilon(x) \sim \mathcal{N}(0, \tau^2)$, and the prior mean is zero, $m(x) = 0$. Given observations $\mathbf{y}_{1:n} = [y_1, \dots, y_n]^T$ at inputs $\mathbf{x}_{1:n} = [x_1, \dots, x_n]^T$, the conditional distribution $f|\mathcal{A}_n$ is then another Gaussian process, with posterior marginal mean $\hat{f}_{\text{Gsn}}^{(n)}(x_*)$ and covariance $v_{\text{Gsn}}^{(n)}(x_*, x'_*)$ given by (throughout we use subscripts to indicate the metamodel type, e.g., *Gsn* for Gaussian noise)

$$\hat{f}_{\text{Gsn}}^{(n)}(x_*) = k(x_*)[\mathbf{K} + \tau^2\mathbf{I}]^{-1}\mathbf{y}_{1:n}, \quad (2.1)$$

$$v_{\text{Gsn}}^{(n)}(x_*, x'_*) = K(x_*, x'_*) - k(x_*)[\mathbf{K} + \tau^2\mathbf{I}]^{-1}k(x'_*)^T, \quad (2.2)$$

with the $1 \times n$ vector $k(x_*)$ and $n \times n$ matrix \mathbf{K} defined by $k(x_*) := K(x_*, \mathbf{x}_{1:n}) = [K(x_*, x_1), \dots, K(x_*, x_n)]$, and $\mathbf{K}_{i,j} := K(x_i, x_j)$.

The posterior mean $\hat{f}_{\text{Gsn}}^{(n)}(x_*)$ is treated as a point estimate of $f(x_*)$ and the posterior standard deviation $s_{\text{Gsn}}^{(n)}(x_*)^2 = v_{\text{Gsn}}^{(n)}(x_*, x_*)$ as the uncertainty of this estimate. We use \mathbf{f} to denote the random posterior vector $f(\mathbf{x}_{1:n})|\mathcal{A}_n$.

Model Fitting: In this article, we model the covariance between the values of f at two inputs x and x'

with the squared exponential (SE) function:

$$K_{\text{se}}(x, x') := \sigma_{\text{se}}^2 \exp \left(- \sum_{i=1}^d \frac{(x^i - x'^i)^2}{2\theta_i^2} \right), \quad (2.3)$$

defined in terms of the hyperparameters $\boldsymbol{\vartheta} = \{\sigma_{\text{se}}, \theta_1, \dots, \theta_d, \tau\}$ known as the process variance and length-scales, respectively. Simulation variance τ is also treated as unknown and part of $\boldsymbol{\vartheta}$. Several common ways exist for estimating $\boldsymbol{\vartheta}$. Within a Bayesian approach we integrate against the prior $p(\boldsymbol{\vartheta})$ using

$$p(\mathbf{f}|\mathbf{y}_{1:n}, \mathbf{x}_{1:n}, \boldsymbol{\vartheta}) = \frac{p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \mathbf{f})p(\mathbf{f}|\boldsymbol{\vartheta})}{p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \boldsymbol{\vartheta})}, \quad (2.4)$$

where $p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \mathbf{f})$ is the likelihood and $p(\mathbf{f}|\boldsymbol{\vartheta})$ is the latent function prior. Notice that following the Gaussian noise assumption, the likelihood $p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \mathbf{f})$ is Gaussian. With a Gaussian prior $p(\mathbf{f}|\boldsymbol{\vartheta})$, the posterior $p(\mathbf{f}|\mathbf{y}_{1:n}, \mathbf{x}_{1:n}, \boldsymbol{\vartheta})$ is tractable and also follows a Gaussian distribution. The normalizing constant in the denominator $p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \boldsymbol{\vartheta})$ is independent of the latent function and is called the marginal likelihood, given by

$$p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \boldsymbol{\vartheta}) = \int p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \mathbf{f})p(\mathbf{f}|\boldsymbol{\vartheta})d\mathbf{f}. \quad (2.5)$$

One may similarly express the posterior over the hyperparameters $\boldsymbol{\vartheta}$, where $p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \boldsymbol{\vartheta})$ plays the role of the likelihood. To avoid expensive MCMC integration, we use the Maximum Likelihood (ML) estimate $\hat{\boldsymbol{\vartheta}}$ which maximizes the likelihood (2.5). Given the estimated hyperparameters $\hat{\boldsymbol{\vartheta}}$, we take the posterior of \mathbf{f} as $p(\mathbf{f}|\mathbf{y}_{1:n}, \mathbf{x}_{1:n}, \hat{\boldsymbol{\vartheta}})$.

2.2 Gaussian Process Regression with Student t -Noise

Taking the noise term $\epsilon(x)$ as Gaussian is widely used since the marginal likelihood is then analytically tractable. In a stochastic simulation setting however, the exact distribution of the outputs relative to their mean is unknown and often is clearly non-Gaussian. A more robust choice is to assume that $\epsilon(x)$ has a Student- t distribution (Jylänki et al., 2011). In particular, this may work better when the noise is heavy-tailed by making inference more resistant to outliers (O’Hagan, 1979). In the resulting t -GP formulation $\epsilon(x)$ is assumed to be t -distributed with variance τ^2 and $\nu > 2$ degrees of freedom (the latter is treated as another hyperparameter). The marginal likelihood of observing $\mathbf{y}_{1:n}$ can be written as

$$p_{t\text{GP}}(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \mathbf{f}) = \prod_{i=1}^n \frac{\Gamma((\nu+1)/2)}{\Gamma(\nu/2)\sqrt{\nu\pi}\sigma_n} \left(1 + \frac{(y_i - f_i)^2}{\nu\sigma_n^2} \right)^{-(\nu+1)/2}, \quad (2.6)$$

where $\Gamma(\cdot)$ is the incomplete Gamma function. The likelihood $p_{t\text{GP}}(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}, \mathbf{f})$ in (2.4) is no longer Gaussian, and integrating (2.6) against the Gaussian prior $p(\mathbf{f}|\boldsymbol{\vartheta})$ is intractable; we therefore use the Laplace approximation (LP) method (Williams and Barber, 1998) to calculate the posterior. A second-order Taylor expansion of $\log p_{t\text{GP}}(\mathbf{f}|\mathbf{x}_{1:n}, \mathbf{y}_{1:n})$ around its mode, $\tilde{\mathbf{f}}_{t\text{GP}}^{(n)} := \arg \max_{\mathbf{f}} p_{t\text{GP}}(\mathbf{f}|\mathbf{x}_{1:n}, \mathbf{y}_{1:n})$, gives a Gaussian approximation

$$p_{t\text{GP}}(\mathbf{f}|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}) \approx q_{t\text{GP}}(\mathbf{f}|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}) = \mathcal{N} \left(\tilde{\mathbf{f}}_{t\text{GP}}^{(n)}, \boldsymbol{\Sigma}_{t\text{GP}}^{-1} \right), \quad (2.7)$$

where Σ_{tGP}^{-1} is the Hessian of the negative conditional log posterior density at $\tilde{\mathbf{f}}_{tGP}^{(n)}$:

$$\Sigma_{tGP} = -\nabla^2 \log p_{tGP}(\mathbf{f}|\mathbf{x}_{1:n}, \mathbf{y}_{1:n})|_{\mathbf{f}=\tilde{\mathbf{f}}_{tGP}^{(n)}} = \mathbf{K}^{-1} + \mathbf{W}_{tGP}, \quad (2.8)$$

and $\mathbf{W}_{tGP} = -\nabla^2 \log p_{tGP}(\mathbf{y}_{1:n}|\mathbf{f}, \mathbf{x}_{1:n})|_{\mathbf{f}=\tilde{\mathbf{f}}_{tGP}^{(n)}}$ is diagonal, since the likelihood factorizes over observations.

Using (2.7), the approximate posterior distribution is also Gaussian $f(x_*)|\mathcal{A}_n \sim \mathcal{N}(\hat{f}_{tGP}^{(n)}(x_*), s_{tGP}^2(x_*))$, defined by its mean $\hat{f}_{tGP}^{(n)}(x_*)$ and covariance $v_{tGP}^{(n)}(x_*, x'_*)$:

$$\hat{f}_{tGP}^{(n)}(x_*) = k(x_*)\mathbf{K}^{-1}\tilde{\mathbf{f}}_{tGP}^{(n)}, \quad (2.9)$$

$$v_{tGP}^{(n)}(x_*, x'_*) = K(x_*, x'_*) - k(x_*)[\mathbf{K} + \mathbf{W}_{tGP}^{-1}]^{-1}k(x'_*). \quad (2.10)$$

Note the similarity to (2.1)–(2.2): with Student- t likelihood the mode $\tilde{\mathbf{f}}_{tGP}^{(n)}$ plays the role of $\mathbf{y}_{1:n}$ and \mathbf{W}_{tGP}^{-1} replaces the noise matrix $\tau^2\mathbf{I}$. Critically, the latter implies that the posterior variance is a function of both designs $\mathbf{x}_{1:n}$ and observations $\mathbf{y}_{1:n}$.

2.3 Gaussian Process Classification

Our target in (1.1) is to learn where the mean response is positive, which is equivalent to classifying each $x \in D$ as belonging either to S or to N . Assuming that $\epsilon(x)$ is symmetric, $\{x \in S\} = \{f(x) \geq 0\} = \{\mathbb{P}(Y(x) > 0) > 0.5\}$. This motivates us to consider the alternative of directly modeling the response sign (rather than overall magnitude) via a classification GP model (CI-GP) (Williams and Barber, 1998; Williams and Rasmussen, 2006). The idea is to model the probability of a positive observation $Y(x)$ by using a probit logistic regression: $\mathbb{P}(Y(x) > 0|x) = \Phi(Z(x))$, with $\Phi(\cdot)$ the standard normal cdf. The latent classifier function is taken as the GP $Z \sim GP(0, K(\cdot, \cdot))$. After learning Z we then set $\hat{S} = \{x \in D : \hat{Z}(x) > 0\}$.

To compute the posterior distribution of Z conditional on \mathcal{A}_n , we use the fact that for an observation (x_i, y_i) and conditional on $z_i = Z(x_i)$ the likelihood of $y_i > 0$ is $\Phi(z_i)1_{\{y_i \geq 0\}} + (1 - \Phi(z_i))1_{\{y_i < 0\}}$. To simplify notation we use $\check{Y}(x) = \text{sign } Y(x) \in \{-1, 1\}$ to represent the signed responses driving CI-GP, leading to $p_{Cl}(\check{\mathbf{y}}_{1:n}|\mathbf{z}, \mathbf{x}_{1:n}) = \prod_{i=1}^n \Phi(\check{y}_i z_i)$. The posterior of the latent $\mathbf{z} = Z(\mathbf{x}_{1:n})$ is therefore

$$p_{Cl}(\mathbf{z}|\mathbf{x}_{1:n}, \check{\mathbf{y}}_{1:n}) = \frac{p(\mathbf{z}|\mathbf{x}_{1:n}) \prod_{i=1}^n \Phi(\check{y}_i z_i)}{p(\check{\mathbf{y}}_{1:n}|\mathbf{x}_{1:n})}. \quad (2.11)$$

Similar to t -GP, we use a Laplace approximation for the non-Gaussian $p_{Cl}(\mathbf{z}|\mathbf{x}_{1:n}, \check{\mathbf{y}}_{1:n})$ in Eq. (2.11) (details to be found in Appendix B). The posterior mean for $Z(\cdot)$ at x_* is then expressed by using the GP predictive mean equation (2.1) and LP approximation (B.1):

$$\hat{z}^{(n)}(x_*) = k(x_*)\mathbf{K}^{-1}\tilde{\mathbf{z}}^{(n)}, \quad (2.12)$$

$$v_{Cl}^{(n)}(x_*, x'_*) = K(x_*, x'_*) - k(x_*)[\mathbf{K} + \mathbf{V}^{-1}]^{-1}k(x'_*)^T. \quad (2.13)$$

We again see the same algebraic structure, with $\tilde{\mathbf{z}}^{(n)}$ a stand-in for $\mathbf{y}_{1:n}$ in (2.1) and \mathbf{V}^{-1} a stand-in for $\tau^2\mathbf{I}$ in (2.2). Also note that we may formally link the Z of the CI-GP metamodel to the GP f used previously via the posterior probability that $x \in S$:

$$\begin{aligned} \mathbb{P}(f(x) \geq 0|\mathcal{A}_n) &= \mathbb{P}(Y(x) > 0|\mathcal{A}_n) = \int_{\mathbb{R}} \Phi(z)p_{Z(x)}(z|\mathcal{A}_n)dz \\ &= \int \Phi(z)\phi\left(\frac{z - \hat{z}^{(n)}(x)}{s_{Cl}^{(n)}(x)}\right)dz = \Phi\left(\frac{\hat{z}^{(n)}(x)}{\sqrt{1 + s_{Cl}^{(n)}(x)^2}}\right). \end{aligned} \quad (2.14)$$

2.4 Student- t Process Regression with Student- t Noise

Instead of just adding Student- t likelihood to the observations, Shah et al. (2014) proposed t -processes (TPs) as an alternative to GPs, deriving closed-form expressions for the marginal likelihood and posterior distribution of the t -process by imposing an inverse Wishart process prior over the covariance matrix of a GP model. They found the t -process to be more robust to model misspecification and to be particularly promising for BO. Moreover, Shah et al. (2014) showed that TPs retain most of the appealing properties of GPs, including analytical expressions, with increased flexibility.

As noticed for example in Williams and Rasmussen (2006), dealing with noisy observations is less straightforward with TPs, since the sum of two independent Student- t distributions has no closed form. Still, this drawback can be circumvented by incorporating the noise directly in the kernel. The corresponding data-generating mechanism is taken to be multivariate- t $\mathbf{y}_{1:n} \sim \mathcal{T}(\nu, m(\mathbf{x}_{1:n}), \mathbf{K} + \tau^2 \mathbf{I})$, where the degrees of freedom are $\nu \in (2, \infty)$. The posterior predictive distribution is then $f(x_*) | \mathcal{A}_n \sim \mathcal{T}(\nu + n, \hat{f}_{\text{TP}}^{(n)}(x_*), v_{\text{TP}}^{(n)}(x_*, x_*))$, where (Shah et al., 2014)

$$\hat{f}_{\text{TP}}^{(n)}(x_*) = k(x_*)[\mathbf{K} + \tau^2 \mathbf{I}]^{-1} \mathbf{y}_{1:n}, \quad (2.15)$$

$$v_{\text{TP}}^{(n)}(x_*, x'_*) = \frac{\nu + \beta^{(n)} - 2}{\nu + n - 2} \{K(x_*, x'_*) - k(x_*)[\mathbf{K} + \tau^2 \mathbf{I}]^{-1} k(x'_*)^T\}, \quad (2.16)$$

with

$$\beta^{(n)} := \mathbf{y}_{1:n}^\top [\mathbf{K} + \tau^2 \mathbf{I}]^{-1} \mathbf{y}_{1:n}.$$

Comparing with the regular GPs, we have the same posterior mean $\hat{f}_{\text{TP}}^{(n)}(x_*) = \hat{f}_{\text{Gsn}}^{(n)}(x_*)$, but the posterior covariance now depends on observations $\mathbf{y}_{1:n}$ and is inflated: $v_{\text{TP}}^{(n)}(x_*, x'_*) = \frac{\nu + \beta^{(n)} - 2}{\nu + n - 2} v_{\text{Gsn}}^{(n)}(x_*, x'_*)$. Moreover, the latent function f and the noise are uncorrelated but not independent. As noticed in Shah et al. (2014), assuming the same hyperparameters, as n goes to infinity, the above predictive distribution becomes Gaussian.

Inference of TPs can be performed similarly as for a GP, for instance based on the marginal likelihood:

$$p_{\text{TP}}(\mathbf{y}_{1:n} | \mathbf{x}_{1:n}, \boldsymbol{\vartheta}) = \frac{\Gamma(\frac{\nu+n}{2})}{((\nu-2)\pi)^{\frac{n}{2}} \Gamma(\frac{\nu}{2})} |\mathbf{K}|^{-1/2} \left(1 + \frac{\mathbf{y}_{1:n}^\top \mathbf{K}^{-1} \mathbf{y}_{1:n}}{\nu-2}\right)^{-\frac{\nu+n}{2}}. \quad (2.17)$$

One issue is estimation of ν , which plays a central role in the TP predictions. We find that restricting ν to be small is important in order to avoid degenerating to the plain Gaussian GP setup.

2.5 Metamodel Performance for Level Set Inference

To evaluate the performance of different metamodels, we consider several metrics. The first statistic is the error rate \mathcal{ER} based on the loss function L defined in Eq. (1.3), measuring the distance between the level set S and its estimate \hat{S} :

$$\mathcal{ER} := \mu(S \Delta \hat{S}) = \int_{x \in D} \mathbb{I}[\text{sign } f(x) \neq \text{sign } \hat{f}(x)] \mu(dx). \quad (2.18)$$

For CI-GP, we replace $f(x)$ with $z(x)$ in the above, namely, use $\mu(S \Delta \hat{S}) = \mu\{x \in D : \hat{z}(x) < 0 < z(x) \cup \hat{z}(x) > 0 > z(x)\}$. A related statistic is the bias \mathcal{B} , which is based on the *signed* (μ -weighted) difference between S and \hat{S} :

$$\mathcal{B} = \mu(S \setminus \hat{S}) - \mu(\hat{S} \setminus S) = \int_{x \in D} \left\{ \mathbb{I}[\hat{f}(x) < 0 < f(x)] - \mathbb{I}[\hat{f}(x) > 0 > f(x)] \right\} \mu(dx). \quad (2.19)$$

The error rate \mathcal{ER} and bias \mathcal{B} evaluate the accuracy of the point estimate \hat{S} when the ground truth is known. In a realistic case study when the latter is unavailable, we replace \mathcal{R} by its empirical counterpart, based on quantifying the uncertainty in \hat{S} through the associated uncertainty of \hat{f} . Following Azzimonti et al. (2016), we define the empirical error \mathcal{E} as the expected distance in measure between the random set $S|\mathcal{A}$ and its estimate \hat{S} :

$$\mathcal{E} := \mathbb{E} \left[\mu(S \Delta \hat{S}) | \mathcal{A} \right] = \int_{x \in D} \bar{E}(x) \mu(dx), \quad (2.20)$$

with $\bar{E}(x)$ calculated by using (2.1) and (2.2):

$$\begin{aligned} \bar{E}(x) &:= \mathbb{E} \left[\mathbb{I}[\text{sign } f(x) \neq \text{sign } \hat{f}(x)] | \mathcal{A} \right] \\ &= \int_{\mathbb{R}} \mathbb{I}[\text{sign } f(x) \neq \text{sign } \hat{f}(x)] p(f(x) | \mathcal{A}) df(x) = \Phi \left(\frac{-|\hat{f}(x)|}{s(x)} \right). \end{aligned}$$

The local empirical error $\bar{E}(x)$ is the posterior probability of wrongly classifying x conditional on the training dataset \mathcal{A} . It is intrinsically tied to the point estimate $\hat{f}(x)$ and the associated posterior variance $s(x)^2$ through the Gaussian uncertainty quantification. For the TPs, the predictive distribution is Student- t , so that the Gaussian cdf Φ is replaced with the respective survival function.

Uncertainty Quantification: To quantify the overall uncertainty about S (rather than local uncertainty about $f(x)$), a natural criterion is the *volume* of the credible band $CI_{\partial S}$ that captures inputs x whose sign classification remains ambiguous given \mathcal{A} . A simple definition at a credibility level α (e.g., $\alpha = 0.05$) would be

$$CI_{\partial S}^{(n)} = \left\{ x \in D : \left(\hat{f}^{(n)}(x) + z_{1-\frac{\alpha}{2}} s^{(n)}(x) \right) \left(\hat{f}^{(n)}(x) - z_{1-\frac{\alpha}{2}} s^{(n)}(x) \right) < 0 \right\}, \quad (2.21)$$

where $z_{1-\frac{\alpha}{2}}$ is the appropriate Gaussian/Student- t α -quantile. Thus (2.21) evaluates the region where the sign of f is nonconstant over the posterior α -CI of f . Heuristically however, $CI_{\partial S} \simeq \{x \in D : \bar{E}(x) > \alpha\}$ is effectively equivalent to empirical error $\bar{E}(x)$ exceeding α , so that the volume of $CI_{\partial S}$ is roughly proportional to the integrated empirical error \mathcal{E} .

In a more sophisticated approach based on random set theory, Chevalier et al. (2013) used the Vorob'ev deviation to define the uncertainty measure $V_\alpha(\hat{S})$:

$$\begin{aligned} V_\alpha(\hat{S}) &:= \mathbb{E} \left[\mu(\hat{S}^\alpha \Delta S) | \mathcal{A} \right] \\ &= \int_{x \in \hat{S}^\alpha} \mathbb{P}(x \notin S | \mathcal{A}) \mu(dx) + \int_{x \in (\hat{S}^\alpha)^c} \mathbb{P}(x \in S | \mathcal{A}) \mu(dx) \\ &= \int_{x \in \hat{S}^\alpha} (1 - p_V(x)) \mu(dx) + \int_{x \in (\hat{S}^\alpha)^c} p_V(x) \mu(dx), \end{aligned} \quad (2.22)$$

where

$$\hat{S}^\alpha := \left\{ x \in D : \hat{f}(x) - z_{1-\frac{\alpha}{2}} s(x) \geq 0 \right\}$$

and

$$p_V(x) = \mathbb{P}(x \in S | \mathcal{A}) = \Phi \left(\frac{\hat{f}(x)}{s(x)} \right).$$

An α satisfying the unbiasedness condition

$$\int_{x \in D} p_V(x) \mu(dx) = \mathbb{E}[\mu(S) | \mathcal{A}] = \mu(\hat{S}^\alpha)$$

is referred to as the *Vorob'ev threshold* and can be determined through dichotomy (Chevalier et al., 2013). If the Vorob'ev threshold is picked to be zero, then the Vorob'ev deviation is reduced to the empirical error \mathcal{E} . Because of the computational overhead of working with (2.22), we restrict attention to the credible bands defined through \hat{S}^α , which correspond to local uncertainty about f (or Z) as in (2.21).

3 Sequential Design

We estimate the level set S in a sequential design setting that assumes that f is expensive to evaluate, for example because of the complexity of the underlying stochastic simulator. Therefore efficient selection of the inputs $\mathbf{x}_{1:n}$ is important. In sequential design, at each step the next sampling location x_{n+1} is selected given all previous measurements. The Bayesian approach to sequential design is based on greedily optimizing an acquisition function as in (1.6). These strategies got popularized thanks to the success of the expected improvement (EI) criterion and the associated efficient global optimization (EGO) algorithm (Jones et al., 1998). The basic loop for sequential design is as following:

- Initialize $\mathcal{A}_{n_0} = \{(x_i, y_i), 1 \leq i \leq n_0\}$.
- Loop for $n = n_0+1, \dots, N$.
 - Choose the next input $x_{n+1} = \arg \max_{x \in \mathcal{M}} \mathcal{I}_n(x)$, and sample $y_{n+1} = Y(x_{n+1})$.
 - Augment $\mathcal{A}_{n+1} = \mathcal{A}_n \cup \{(x_{n+1}, y_{n+1})\}$.
 - Update $\hat{S}^{(n+1)}$ with \mathcal{A}_{n+1} .

We now propose several metrics for the acquisition function $\mathcal{I}_n(x)$ in Eq. (1.6). The key plan is to target regions close to the boundary $\partial \hat{S}$. A second strategy is to use the look-ahead posterior standard deviation $s^{(n+1)}$ conditional on sampling at x , in order to assess the corresponding *information gain*. This links the constructed design to the metamodel for f , since different surrogate architectures quantify uncertainty differently.

The first metric, dubbed Maximum Contour Uncertainty (MCU), stems from the Upper Confidence Bound (UCB) strategies proposed by Srinivas et al. (2012) for Bayesian optimization. The idea of UCB is to express the exploitation-exploration trade-off through the posterior mean $\hat{f}(x)$ and standard deviation $s(x)$. Following the spirit of UCB, MCU blends the minimization of $|\hat{f}^{(n)}(x)|$ (exploitation) with maximization of the posterior uncertainty $s^{(n)}(x)$ (exploration):

$$\mathcal{I}_n^{\text{MCU}}(x) := -|\hat{f}^{(n)}(x)| + \gamma^{(n)} s^{(n)}(x), \quad (3.1)$$

where $\gamma^{(n)}$ is a step-dependent sequence of weights. Thus, MCU targets inputs with high uncertainty (large $s^{(n)}(x)$) and close to the boundary $\partial \hat{S}$ (small $|\hat{f}^{(n)}(x)|$). Small $\gamma^{(n)}$ leads to aggressive sampling concentrated along the estimated $\partial \hat{S}$; large $\gamma^{(n)}$ leads to space-filling sampling that effectively minimizes the integrated mean-squared error. Thus, the choice of γ 's is critical for the performance; in particular $\gamma^{(n)}$ should be increasing to avoid being trapped in local minima of $|\hat{f}^{(n)}(x)|$. In the original application to BO (Srinivas et al., 2012) it is proved that with $\gamma^{(n)} = (2 \log(\frac{|D| \pi^2 n^2}{6\delta}))^{1/2}$, the regret (a metric measuring the distance between estimated optima and the truth in BO) of the estimate is guaranteed to converge. Further recipes for (3.1) for level set estimation were proposed in Gotovos et al. (2013) and Bogunovic et al. (2016); both papers mention that the above recommendation is too conservative and tends to over-explore. A constant choice of $\gamma^{(n)} = 1.96$ corresponds to the Straddle scheme in Bryan et al. (2006) and leads to $\mathcal{I}_n(x) \geq 0 \Leftrightarrow x \in (95\%$

CI band of ∂S). Similarly, Gotovos et al. (2013) employed $\gamma^{(n)} = 3$ and Bogunovic et al. (2016) suggested to use $\gamma^{(n)} = \sqrt{\log(|D|n^2)}$. Based on our experiments (see Appendix A), we find that a constant value of $\gamma^{(n)}$ may be problematic and recommend to adapt $\gamma^{(n)}$ to the relative ratio between $f(x)$ (for steeper response surfaces γ should be larger) and $s(x)$ (γ needs to rise as posterior uncertainty decreases). One recipe is to use $\gamma^{(n)} = \text{IQR}(\hat{f}^{(n)}) \setminus 3 \text{Ave}(s^{(n)})$ ($\text{Ave}(s^{(n)})$ denotes the average of standard deviation and IQR the inter-quantile range) which keeps both terms in (3.1) approximately comparable as n changes.

Remark 3.1. The local empirical error $\bar{E}(x)$ as defined in Eq. (??) could be directly used as an acquisition function, i.e.,

$$\mathcal{I}_n^{\text{MEE}}(x) \equiv \bar{E}(x) = \Phi\left(-\frac{|\hat{f}^{(n)}(x)|}{s^{(n)}(x)}\right). \quad (3.2)$$

This Maximal Empirical Error (MEE) acquisition function measures the local probability of misclassification and is similar to the sequential criteria in Bect et al. (2012); Echard et al. (2010); Ranjan et al. (2008); Bichon et al. (2008), all based on the idea of sampling at x where the event $\{f(x) \geq 0\} | \mathcal{A}_n$ is most uncertain. However, (3.2) is not suitable for our purposes since it is maximized across the entire $\partial \hat{S}$ (namely $\mathcal{I}_n^{\text{MEE}}(x) = 0.5$ for any x where $\hat{f}^{(n)}(x) = 0$), so does not possess a unique maximizer as soon as $\partial \hat{S}$ is non-trivial. One potential solution could be to maximize (3.2) over a finite candidate set, which however requires significant fine-tuning.

Our second strategy focuses on quickly *reducing* \bar{E} by comparing the current $\bar{E}(x)$ given \mathcal{A}_n and the expected $\bar{E}(x)$ conditional on the one-step-ahead sample, $\mathcal{A}_n \cup \{x_{n+1}, y_{n+1}\}$. This is achieved by integrating out the effect of $Y(x_{n+1})$ on $\bar{E}(x_{n+1})$:

$$\begin{aligned} \mathcal{I}_n^{\text{cSUR}}(x) &= \mathcal{I}_n^{\text{MEE}}(x) - \mathbb{E}_{Y(x)} [\mathcal{I}_{n+1}^{\text{MEE}}(x)] \\ &= \Phi\left(-\frac{|\hat{f}^{(n)}(x)|}{s^{(n)}(x)}\right) - \mathbb{E}_{Y(x)} \left[\Phi\left(-\frac{|\hat{f}^{(n+1)}(x)|}{s^{(n+1)}(x)}\right) \right]. \end{aligned} \quad (3.3)$$

The name cSUR is because (3.3) is directly related to the SUR strategy (Bect et al., 2012), modified to target contour-finding. Crucially, $\mathcal{I}_n^{\text{cSUR}}$ ties the selection of x_{n+1} to the look-ahead mean $\hat{f}^{(n+1)}(x_{n+1})$ and look-ahead standard deviation $s^{(n+1)}(x_{n+1})$ that appear on the right-hand side of (3.3). To compute the integral over $Y(x)$, we replace $\hat{f}^{(n+1)}(x)$ with its average $\hat{f}^{(n)}(x) = \mathbb{E}_n[f(x)] = \mathbb{E}_n[\mathbb{E}_{n+1}[f(x)]] = \mathbb{E}_n[\hat{f}^{(n+1)}(x)]$. Similarly, we plug in the approximate one-step-ahead standard deviation $\hat{s}^{(n+1)}$ discussed in Section 4 (especially Equations (4.4), (C.4), and (C.6)) for $s^{(n+1)}(x)$:

$$\hat{\mathcal{I}}_n^{\text{cSUR}}(x) = \Phi\left(-\frac{|\hat{f}^{(n)}(x)|}{s^{(n)}(x)}\right) - \Phi\left(-\frac{|\hat{f}^{(n)}(x)|}{\hat{s}^{(n+1)}(x)|_{x_{n+1}=x}}\right). \quad (3.4)$$

Note that if x is such that $\hat{f}^{(n)}(x) = 0$ then both terms above are 1/2 and $\mathcal{I}_n^{\text{cSUR}}(x) = 0$. Thus, the cSUR criterion will not place samples *directly* on $\partial \hat{S}$, but will aim to bracket the zero-contour.

In (3.4) cSUR only measures the *local* improvement in $\bar{E}(x_{n+1})$ at the sampling location x_{n+1} and consequently might be overly aggressive in targeting $\partial \hat{S}$. This motivates us to target the *global* reduction in the uncertainty of \hat{S} , so as to take into account the spatial structure of D . The resulting Integrated Contour Uncertainty (ICU) is linked to the already defined empirical error \mathcal{E} from Section 2.5:

$$\begin{aligned} \mathcal{I}_n^{\text{ICU}}(x) &:= \mathcal{E}^{(n)} - \mathbb{E}_{Y(x)} [\mathcal{E}^{(n+1)} | x_{n+1} = x] \\ &= \mathcal{E}^{(n)} - \mathbb{E}_{Y(x)} \left[\int_{u \in D} \Phi\left(\frac{-|\hat{f}^{(n+1)}(u)|}{s^{(n+1)}(u)|_{x_{n+1}=x}}\right) \mu(du) \right]. \end{aligned} \quad (3.5)$$

We apply the same approximation as for cSUR to simplify the expectation over $Y(x)$ and replace the integral over D with a sum over a finite subset \mathcal{D} of size M :

$$\hat{\mathcal{I}}_n^{\text{ICU}}(x) = - \sum_{x_m \in \mathcal{D}} \Phi \left(\frac{-|\hat{f}^{(n)}(x_m)|}{\hat{s}^{(n+1)}(x_m)|_{x_{n+1}=x}} \right) \mu(x_m). \quad (3.6)$$

Then $\mathcal{I}^{\text{ICU}}(x)$ can be viewed as measuring the overall information gain about S from sampling at x . The motivation behind ICU is to myopically minimize the expected one-step-ahead empirical error \mathcal{E} , which would correspond to 1-step Bayes-optimal design.

As a last alternative, we utilize the targeted mean squared error (tMSE) criterion, a localized form of targeted IMSE criterion in Picheny et al. (2010):

$$\mathcal{I}_n^{\text{tMSE}}(x) := s^{(n)}(x)^2 \cdot W_n^{\text{tMSE}}(x), \quad (3.7)$$

where

$$W_n^{\text{tMSE}}(x) := \frac{1}{\sqrt{2\pi}s^{(n)}(x)} \exp \left(-\frac{\hat{f}_n(x)^2}{2s^{(n)}(x)^2} \right). \quad (3.8)$$

The tMSE criterion upweighs regions close to the zero contour through the weight function $W_n^{\text{tMSE}}(x)$ which measures the distance of x to $\partial\hat{S}^{(n)}$ using the Gaussian posterior density $\mathcal{N}(\hat{f}^{(n)}, s^{(n)}(x)^2)$. Like MCU, tMSE is based only on the posterior at step n and does not integrate over future $Y(x)$'s.

Remark 3.2. In Picheny et al. (2010) an additional parameter σ_ϵ was added to the definition of $W_n^{\text{tMSE}}(x)$ by replacing $s^{(n)}(x)$ everywhere with $\sqrt{s^{(n)}(x)^2 + \sigma_\epsilon^2}$. Larger σ_ϵ yields more space-filling as $W_n^{\text{tMSE}}(x)$ becomes flatter. Since Picheny et al. (2010) dealt with deterministic experiments, σ_ϵ was necessary to ensure that $W_n^{\text{tMSE}}(x)$ is well defined at existing $x_{1:n}$ and the recommendation was for σ_ϵ to be 5% of the range of f . In our case $s^{(n)}(x)$ is intrinsically bounded away from zero and (3.8) works well as is. Additional experiments (available upon request) indicate that the performance of (3.7) is not sensitive to σ_ϵ , so to minimize the number of tuning parameters we stick to $\sigma_\epsilon = 0$ in (3.8).

In the TP case, for MCU, cSUR, and ICU, we replace the standard normal cdf $\Phi(\cdot)$ appearing in the formulas by its Student- t counterpart (with the estimated degrees of freedom ν_n). For tMSE, to maintain tractability, we keep the same expression (3.8) for the weights W^{tMSE} .

Illustration

For instructive purposes, we consider a one-dimensional case where we use the Gaussian observation GP to learn the sign of the quadratic $f(x) = x^2 - 0.75^2$ on $D = [0, 1]$, where $S = [0, 0.75]$ and with the unique zero contour at $\partial S = 0.75$. The initial design $\mathbf{x}_{1:10}$ consists of $n = 10$ inputs drawn according to Latin hypercube sampling (LHS). The observations are $Y(x) = f(x) + \epsilon$, where $\epsilon \sim t_3(0, 0.1^2)$. In the top plot in Figure 1, we plot the true $f(\cdot)$, the posterior mean $\hat{f}^{(100)}(\cdot)$, and associated 95%-CI. We also show the credible band for $\partial\hat{S}$; in the respective bottom panel, we plot the acquisition functions $\mathcal{I}_n^{\text{MCU}}(\cdot)$, $\mathcal{I}_n^{\text{cSUR}}(\cdot)$, $\mathcal{I}_n^{\text{ICU}}(\cdot)$ and $\mathcal{I}_n^{\text{tMSE}}(\cdot)$ as defined in Equations (3.1), (3.4), (3.6), and (3.7).

Comparing the acquisition functions of the four criteria, we find that, besides ICU, all of the others have maxima within the shaded credible interval of the boundary $CI_{\partial S}$. In practice, we care only about the maximizer of the acquisition function, rather than its full shape, since the former drives the selection of the next sample x_{n+1} . The x_{n+1} 's selected by MCU and tMSE criteria are close. For the cSUR criterion,




Figure 1: Comparison of acquisition functions. *Upper panel*: true function $f = (x + 0.75)(x - 0.75)$ (black solid line), the posterior mean $\hat{f}(\cdot)$ (dashed line) and 95% CI_f (shaded area) based on observed samples $(\mathbf{x}_{1:100}, \mathbf{y}_{1:100})$ (blue dots). Along the x -axis we also show the credible interval of the partition boundary $CI_{\partial S}$ (grey solid line) relative the true zero level set $S = [0, 0.75]$ (red triangle). *Lower panel*: acquisition functions $\mathcal{I}_n(\cdot)$ for MCU, cSUR, ICU, and tMSE criteria, with vertical lines marking the respective maxima $\arg \max_x \mathcal{I}_n(x)$.

because $\mathcal{I}_n^{\text{cSUR}}(x) = 0$ at $\partial \hat{S}$, there are two local maxima with a “valley” between them. The interval between the two local maxima is roughly the confidence interval $CI_{\partial S}$ for the boundary (2.21). Both MCU and tMSE select a location very close to the boundary $\hat{f}^{(n)}(x_{n+1}) \simeq 0$. We note that MCU has a flatter acquisition function, i.e., tMSE is more aggressive. In contrast, the ICU and cSUR criteria are more “global”; in particular, ICU is the flattest among all the criteria.

After using the various acquisition functions to select x_{n+1} at $n = 11, \dots, 100$, we show in Figure 2 the resulting designs $\mathbf{x}_{1:n}$ and the final estimate $\hat{f}^{(100)}$ with a Gaussian observation GP metamodel. As desired, all methods target the true zero-contour at $\partial S = 0.75$. As a result, the posterior variance $s^{(n)}(x)^2$ is much lower in this neighborhood; in contrast, especially for tMSE and MCU, few samples are taken far from $x = 0.75$, and the posterior uncertainty there remains high. The true zero contour is within the estimated posterior CI for all the criteria. However, the CIs for MCU and tMSE are much wider than those for the others.

The bottom row in Figure 2 shows the sampled location x_n as a function of step n . We observe that MCU and tMSE heavily concentrate their search around the zero contour, leading to few samples (and consequently relatively large empirical errors $\mathcal{E}^{(n)}$) in other areas, although the overall error rate \mathcal{R} is comparable. The ICU and cSUR criteria exhibit an “edge” effect; that is, besides the desired zero contour $x = 0.75$, multiple samples are taken close to the edges of the input space at $x = 0$ and $x = 1$. This occurs due to the relatively large posterior variance $s^2(\cdot)$ in those regions (which arises intrinsically with any spatial-based metamodel) that in turn strongly influences $\mathcal{I}^{\text{cSUR}}$ in (3.4) and \mathcal{I}^{ICU} in (3.6). Inputs sampled by the cSUR criterion bracket the contour ∂S from both directions, matching the two-hill-and-a-valley shape of $\mathcal{I}^{\text{cSUR}}$ in Figure 1. We note that the two sampling “curves” get closer as n grows, indicating a gradual convergence of the estimated zero contour $\partial \hat{S}^{(n)}$, akin to a shrinking credible interval of $\hat{S}^{(n)}$. The ICU criterion generates a much more diffuse design: it engages in more exploration and is less dependent on the current levels of the empirical error \mathcal{E} . This eventually creates a flatter profile for $\bar{E}(x)$.




Figure 2: *Top row*: Fitted metamodel $\hat{f}^{(100)}$ (dashed red line) and its 95%-CI (shaded region) versus the true $f = (x + 0.75)(x - 0.75)$ (solid black), for each of the four design strategies. The estimated 95% CI for the zero-contour ∂S is marked on the x -axis with a grey interval; red triangle indicates the true zero-contour $\partial S = 0.75$. *Bottom row*: sampled inputs x_n (on the x -axis to match the top row) as a function of step $n = 1, \dots, 100$ (on the y -axis, moving from top to bottom) for MCU, tMSE, cSUR, and ICU criteria. The rug plots at the bottom visualize the overall distribution of $\mathbf{x}_{1:n}$ at $n = 100$. The first ten inputs are selected using a (fixed-across schemes) LHS design on $D = [0, 1]$.

The preceding discussion considered a single metamodel choice for f . Other metamodels will generate different design features; in particular, sensitivity to $\epsilon(x)$ will lead to a different mix of exploration (x_n 's far from the zero-contour) and exploitation even for the same choice of a \mathcal{I}_n criterion. Figures 6 and 7, as well as Table 3, emphasize our message that one must jointly investigate the *combinations* of $\mathcal{I}(\cdot)$ and \hat{f} when benchmarking the ultimate performance of the algorithm.

4 Look-Ahead Variance

The cSUR and ICU acquisition functions \mathcal{I}_n require estimates of the look-ahead standard deviation $s^{(n+1)}(x_*)$ conditional on sampling at $x_{n+1} = x$. A related computation is also important for efficient updating of the GP/TP metamodels during sequential design, assimilating the observation (x_{n+1}, y_{n+1}) into \mathcal{A}_n . As is well known, usage of GP necessitates inverting the covariance matrix \mathbf{K}^{-1} which presents a computational bottleneck as n grows. Updating hinges on computing $[\mathbf{K}^{(n+1)}]^{-1}$ via applying the Woodbury identities to the current $[\mathbf{K}^{(n)}]^{-1}$.

A major advantage of the classical GP paradigm is that the posterior variance $s^{(n)}(x)^2$ is a function only of the design $\mathbf{x}_{1:n}$; that is, it is independent of the observations $\mathbf{y}_{1:n}$. This allows an exact analytic expression for $s^{(n+1)}(x)|_{x_{n+1}=x}$ in terms of x_{n+1} . Recall that for an existing design $\mathbf{x}_{1:n}$, after adding a new (x_{n+1}, y_{n+1}) , the mean and variance at location x_* are updated via (Chevalier et al., 2014b)

$$\hat{f}_{\text{Gsn}}^{(n+1)}(x_*) = \hat{f}_{\text{Gsn}}^{(n)}(x_*) + \lambda^{(n)}(x_*, x_{n+1})(y_{n+1} - \hat{f}_{\text{Gsn}}^{(n)}(x_{n+1})), \quad (4.1)$$

$$s_{\text{Gsn}}^{(n+1)}(x_*)^2 = s_{\text{Gsn}}^{(n)}(x_*)^2 - \lambda^{(n)}(x_*, x_{n+1})^2(\tau^2 + s_{\text{Gsn}}^{(n)}(x_{n+1})^2), \quad (4.2)$$

where $\lambda^{(n)}(x_*, x_{n+1})$ is a weight function that measures the influence of the new sample at x_{n+1} on x_* conditioned on the existing inputs $\mathbf{x}_{1:n}$.

Lemma 4.1 (Woodbury formula). *Assume \mathbf{b} is a $n \times 1$ vector, \mathbf{A} is a $n \times n$ matrix, and d and c are nonzero scalars; then we have*

$$\begin{bmatrix} \mathbf{b}^T & d \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{b}^T & c \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{b} \\ d \end{bmatrix} = \mathbf{b}^T \mathbf{A}^{-1} \mathbf{b} - \frac{1}{c - \mathbf{b}^T \mathbf{A}^{-1} \mathbf{b}} (d - \mathbf{b}^T \mathbf{A}^{-1} \mathbf{b})^2. \quad (4.3)$$

Using Lemma 4.1, we obtain the one-step-ahead variance at x_* :

Proposition 4.1. *For any x_* ,*

$$\begin{aligned} \lambda^{(n)}(x_*, x_{n+1}) &= \frac{v_{\text{Gsn}}^{(n)}(x_*, x_{n+1})}{\tau^2 + s_{\text{Gsn}}^{(n)}(x_{n+1})^2} \\ \Rightarrow s_{\text{Gsn}}^{(n+1)}(x_*)^2 &= s_{\text{Gsn}}^{(n)}(x_*)^2 - \frac{v_{\text{Gsn}}^{(n)}(x_*, x_{n+1})^2}{\tau^2 + s_{\text{Gsn}}^{(n)}(x_{n+1})^2}. \end{aligned} \quad (4.4)$$

In particular, after sampling at x_{n+1} the local updated posterior variance is proportional to the current $s_{\text{Gsn}}^{(n)}(x_{n+1})^2$ with a proportionality factor (Hu and Ludkovski, 2017):

$$\frac{s_{\text{Gsn}}^{(n+1)}(x_{n+1})^2}{s_{\text{Gsn}}^{(n)}(x_{n+1})^2} = \frac{\tau^2}{\tau^2 + s_{\text{Gsn}}^{(n)}(x_{n+1})^2}. \quad (4.5)$$

The above lemma is our basis for calculating the acquisition function for the cSUR criterion (3.4) that requires only (4.5) and the ICU criterion (3.6). As we see below, because (4.4) holds only in the Gaussian prior/Gaussian likelihood setting, further approximations are required to apply (4.1)–(4.5) for the alternative metamodels. Such look-ahead variance expressions are of independent interest, applicable beyond the context of level set estimation.

A limitation of using a non-Gaussian observation or classification likelihood is that, unlike for Gaussian observation GP, there are no exact variance look-ahead formulas for the resulting t -GP, CI-GP and TP metamodels. There are two main reasons for this. First, both the posterior mean $\hat{f}^{(n+1)}(x_*)$ in (2.9) and (2.12) and the posterior variance $s^{(n+1)}(x_*)^2$ in (2.10) and (2.13) for t -GP and CI-GP depend on the posterior mode $\tilde{\mathbf{f}}_{t\text{GP}}^{(n+1)}$ or $\tilde{\mathbf{z}}_{CI}^{(n+1)}$, which changes every step. Therefore, they cannot be accessed in advance. Furthermore, for t -GP and CI-GP $s^{(n+1)}(x_*)$ depends on the next-step Hessian \mathbf{W} (namely on $w_{n+1}^{(n+1)}$), and for TP $s^{(n+1)}(x_*)$ depends on $\beta^{(n+1)}$. Both of these again depend on y_{n+1} . To overcome this challenge, we develop an approximation $\hat{s}^{(n+1)}(\cdot)$ for each metamodel. Our strategy is to replace each inaccessible term with its expected value from the point of view of step n . For example, we calculate the expectation of $\tilde{\mathbf{f}}_{t\text{GP}}^{(n+1)}$, $\tilde{\mathbf{z}}_{CI}^{(n+1)}$ and $\beta^{(n+1)}$ with respect to \mathcal{A}_n . Propositions 4.2–4.3–4.4 provide the resulting look-ahead formulas for t -GP, CI-GP and TP respectively, with derivation details in Appendix C.

Proposition 4.2. *For any x_* , the formula for the look-ahead variance for t -GP is*

$$\hat{s}_{t\text{GP}}^{(n+1)}(x_*)^2 := s_{t\text{GP}}^{(n)}(x_*)^2 - \frac{v_{t\text{GP}}^{(n)}(x_*, x_{n+1})^2}{(\tau^2 \frac{\nu+1}{\nu-1}) + s_{t\text{GP}}^{(n)}(x_{n+1})^2}. \quad (4.6)$$

Proposition 4.3. Let $\check{v}_{n+1} = v_{n+1}^+ p_+ + v_{n+1}^- p_-$, where

$$v_{n+1}^+ = \frac{\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2}{\Phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2} + \frac{\hat{z}_{\text{Cl}}^{(n)}(x_{n+1})\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}{\Phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}, \quad (4.7)$$

$$p_+ = \Phi\left(\frac{\hat{z}^{(n)}(x_{n+1})}{\sqrt{1 + s_C^{(n)}(x_{n+1})^2}}\right), \quad (4.8)$$

and

$$v_{n+1}^- = \frac{\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2}{\Phi(-\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2} - \frac{\hat{z}_{\text{Cl}}^{(n)}(x_{n+1})\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}{\Phi(-\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}, \quad (4.9)$$

$$p_- = 1 - p_+. \quad (4.10)$$

For any x_* , the formula for the look-ahead variance for CI-GP is

$$\hat{s}_{\text{Cl}}^{(n+1)}(x_*)^2 := s_{\text{Cl}}^{(n)}(x_*)^2 - \frac{v_{\text{Cl}}^{(n)}(x_*, x_{n+1})^2}{(\check{v}_{n+1})^{-1} + s_{\text{Cl}}^{(n)}(x_{n+1})^2}. \quad (4.11)$$

Proposition 4.4. For any x_* , the formula for the look-ahead variance for TP is

$$s_{\text{TP}}^{(n+1)}(x_*)^2 = \frac{\nu + \check{\beta}^{(n+1)} - 2}{\nu + n - 1} s_{\text{Gsn}}^{(n+1)}(x_*)^2, \quad (4.12)$$

where $\check{\beta}^{(n+1)} = \beta^{(n)} + \frac{\nu}{\nu-2}$.

We note that in our experiments we only use the above to evaluate \mathcal{I}_n , and directly re-estimate $\tilde{f}^{(n+1)}$ at each step of the sequential design.

5 Synthetic Experiments

5.1 Benchmark Construction

As synthetic experiments, we consider three benchmark problems in dimension $d = 1, 2$, and 6. For the latter two we employ the widely used *Branin-Hoo* 2-D and *Hartman* 6-D functions; see, for example, Picheny et al. (2013b). The original functions have been rescaled to map their sample space D onto $[0, 1]^d$; see Table 1.

The latent functions are chosen to cover a variety of problem properties. The quadratic f in 1-D is strictly monotonically increasing, yielding a single boundary ∂S . The original Branin-Hoo function (Picheny et al., 2013b) is modified so that f is increasing in x^1 and the zero-level set has a non-trivial shape in x^2 . The *Hartman* is a multimodal function with a complex zero contour. The parameters in the original *Hartman* function described in Picheny et al. (2013b) are adjusted to reduce the "bumps" in the zero contour and make the problem more appropriate for the sign classification task.

A large number of factors can influence the performance of metamodels and designs. In line with the stochastic simulation perspective, we concentrate on the impact of the simulation noise and consider four observation setups. These cover a variety of noise distributions and signal-to-noise ratio, measured through the proportion of standard deviation σ_τ to the range R_f of the response. The first two settings use *Student-t* distributed noise, with (i) low σ_τ and (ii) high σ_τ . The third setting uses (iii) Gaussian mixture noise to

Table 1: Response surfaces $x \mapsto f(x)$ for synthetic experiments.

Quadratic (1-D)	$f(x) = (x + 0.75)(x - 0.75)$ with $x \in [0, 1]$	
Branin-Hoo (2-D)	$f(x) = \frac{1}{178} \left[\left(\bar{x}^1 - \frac{5.1(\bar{x}^2)^2}{4\pi^2} + \frac{5\bar{x}^2}{\pi} - 20 \right)^2 + \left(10 - \frac{10}{8\pi} \right) \cos(\bar{x}^1) - 181.47 \right]$ with: $\bar{x}^1 = 15x^1, \bar{x}^2 = 15x^2 - 5, x^1, x^2 \in [0, 1]$	
Hartman6 (6-D)	$f(x) = \frac{-1}{0.1} \left[\sum_{i=1}^4 C_i \exp \left(- \sum_{j=1}^6 a_{ji} (x^j - p_{ji})^2 \right) - 0.1 \right]$ with: $\mathbf{C} = [0.2, 0.22, 0.28, 0.3]$	
	$\mathbf{a} = \begin{bmatrix} 8.00 & 0.50 & 3.00 & 10.00 \\ 3.00 & 8.00 & 3.50 & 6.00 \\ 10.00 & 10.00 & 1.70 & 0.50 \\ 3.50 & 1.00 & 8.00 & 8.00 \\ 1.70 & 6.00 & 10.00 & 1.00 \\ 6.00 & 9.00 & 6.00 & 9.00 \end{bmatrix},$	$\mathbf{P} = \frac{1}{10^4} \begin{bmatrix} 1312 & 2329 & 2348 & 4047 \\ 1696 & 4135 & 1451 & 8828 \\ 5569 & 8307 & 3522 & 8732 \\ 124 & 3736 & 2883 & 5743 \\ 8283 & 1004 & 3047 & 1091 \\ 5886 & 9991 & 6650 & 381 \end{bmatrix}$

further test misspecification of ϵ . The fourth setting considers the challenging case of (iv) a heteroscedastic Student- t noise with state-dependent degrees of freedom. In total we have $3 \times 4 \times 4 \times 6$ experiments (indexed by their dimensionality, noise setting, design heuristic, and metamodel type).

Besides the noise distribution, we fix all other metamodeling aspects. All schemes are initialized with $n_0 = 10d$ inputs drawn from an LHS design on $[0, 1]^d$ and use the SE kernel (2.3) for the covariance matrix \mathbf{K} . To analyze for the variability due to the initial design and the noise realizations, we perform 100 macroruns of each design/acquisition function combination. For each run, the same initial inputs are used across all GP metamodels and designs, but otherwise the initial $\mathbf{x}_{1:n_0}$ vary across runs.

 Table 2: Stochastic simulation setup for synthetic experiments. ($R_f \equiv \max_x f(x) - \min_x f(x) = 1$)

Initial design	Latin hypercube sampling of size $n_0 = 10d$
Total budget n	$d = 1, n = 100$; $d = 2, n = 150$; $d = 6, n = 1000$
Test set size $M = \mathcal{D} $	$d = 1, M = 1000$; $d = 2, M = 500$; $d = 6, M = 1000$
Noise setting for $\epsilon(x)$	(i) t /small : $t_3(0, (0.1R_f)^2)$ (ii) t /large : $t_3(0, (0.5R_f)^2)$ (iii) Gsn /mix: 50/50 mix of $\mathcal{N}(0, (0.5R_f)^2)$ and $\mathcal{N}(0, R_f^2)$ (iv) t /hetero : $t_{6-4x^1}(0, (0.4(4x^1 + 1))^2)$

Optimization of the Improvement Metric: We employed the MCU, ICU, tMSE and cSUR criteria to maximize the improvement metric \mathcal{I} and select the next input x_{n+1} . This maximization task is nontrivial in higher dimensions because \mathcal{I} is frequently multimodal and can be flat around its local maxima. We use a genetic optimization approach as implemented in the `ga` function in MATLAB, with tolerance of 10^{-3} and 200 generations. This is a global, gradient-free optimizer that uses an evolutionary algorithm to explore the input space D .

Evaluation of Performance Metrics: Recall that evaluating the quality of $\partial \hat{S}$ is based on \mathcal{R} and \mathcal{E} from (2.18) and (2.20) that require integration over D . In practice, these are computed based on a weighted sum over a finite \mathcal{D} , $\hat{\mathcal{E}} := \sum_{m=1}^M \Phi \left(\frac{-|f(x_m)|}{s(x_m)} \right) \mu(x_m)$ for a space-filling sequence $\mathcal{D} \equiv x_{1:M} \in D$ of test points. In 1-D experiments \mathcal{D} was an equispaced grid of size $M = 1000$. In higher dimensions, to avoid the use of

a lot of test points that are required to ensure an accurate approximation, we adaptively pick \mathcal{D} that targets the critical region close to the zero contour. To do so, we replace the integral with a weighted sum:

$$\mathcal{R} \simeq \frac{p_c}{M_1} \sum_{x_{1:M_1} \in D_1} \mathbb{I}(\text{sign } f(x_m) \neq \text{sign } \hat{f}(x_m)) + \frac{(1-p_c)}{M_2} \sum_{x_{1:M_2} \in D_2} \mathbb{I}(\text{sign } f(x_m) \neq \text{sign } \hat{f}(x_m)), \quad (5.1)$$

where $M = M_1 + M_2$ and the test locations $x_{1:M_1}$ and $x_{1:M_2}$ are subsampled from a large space-filling (scrambled Sobol) sequence on D . The weight p_c determines the relative volume of D_1 and $D_2 = D \setminus D_1$, where on $D_1 = \{x : f(x) \simeq 0\}$ we are close to the zero contour. In the experiments below we use $M_1 = 0.8M$, $M_2 = 0.2M$, and $p_c = 0.4$, so that the density of test points close to ∂S is double relative to those far from the zero contour. We employ the same strategy for speeding the evaluation of the empirical error \mathcal{E} .

Surrogate Inference: Values of hyperparameters ϑ are crucial for good performance of GP metamodels. We estimate ϑ using maximum likelihood. Except for TP, all models are fitted with the open source package `GPstuff` (Vanhatalo et al., 2013) in MATLAB. TPs are fitted with the `hetGP` (Binois et al., 2018) package in R. Auxiliary tests did not reveal any significant effects from using other available tools for plain GPs and t -GP, such as `GPML` (Rasmussen and Nickisch, 2010).

In principle, the hyperparameters ϑ change at every step of the sequential design, in other words, whenever \mathcal{A}_n is augmented with (x_{n+1}, y_{n+1}) . To save time however, we do not update ϑ at each step. Instead, we first estimate the hyperparameters ϑ based on the initial design \mathcal{A}_{n_0} and then freeze them, updating their values only every few steps. Specifically, ϑ is re-estimated at steps $n_0 + 1, n_0 + 2, n_0 + 4, n_0 + 8, n_0 + 16, \dots$ (as the sample size becomes large, the inference of hyperparameters becomes more stable).

The lengthscales θ_i are the most significant for surrogate goodness of fit. A too-small lengthscales will make the estimated \hat{f} look “wiggly” and might lead to overfitting, while θ_i too large will fail to capture an informative shape of the true f and hence S . Since our input domain is always $[0, 1]^d$, we restrict $\theta_i \in [0.3, 2] \forall i$ to be on the order of the length of the sample space D .

Computational Overhead: All the considered metamodels are computationally more demanding than the baseline Gaussian GP. For t -GP and CI-GP, additional cost arises due to the Laplace approximation. TP necessitates estimation of the parameter ν and also the computation of β in (2.16). In the experiments considered, the respective computation times were roughly double to triple relative to the Gaussian GP. In terms of sequential design, MCU, tMSE, and cSUR have approximately equal overhead; ICU is significantly more expensive because it requires evaluating the sum in (3.6). Note that all heuristics include two expensive steps: optimization for x_{n+1} and computation of $\hat{f}^{(n)}$ and $s^{(n)}$ (and/or $\hat{s}^{(n+1)}$).

Overall timing of the schemes is complicated because of the combined effects of n (design budget), M (size of test set), and the use of different software (some schemes run in R and others in Matlab). Most important, the ultimate computation time is driven by the simulation cost of generating $Y(x)$ -samples, which is trivial in the synthetic experiments but assumed to be large in the motivating context.

5.2 Comparison of GP Metamodels

Figure 3 shows the boxplots of the error rate \mathcal{ER} of $\hat{S}^{(N)}$ at the final design ($N = 100$ in 1-D; $N = 150$ in 2-D; $N = 1000$ in 6-D). The plots are sorted by noise settings and design strategies, facilitating comparison between the discussed metamodels. In Table 3, we list the best metamodel and design combination in each case. Several high-level observations can be made. First, we observe the limitations of the baseline Gaussian GP metamodel, which cannot tolerate too much model misspecification. As the noise structure gets more complex, the classical GP surrogate begins to show increasing strain; in the last t /*hetero* setup, it is both

unstable (widely varying performance across runs) and inaccurate, with error rates upward of 30% on “bad” runs. In addition, according to results shown in Table 3, across all of the twelve cases, besides 1d example with $t/small$ noise, the Gaussian GP never performs as the best model. This result is not surprising but confirms that the noise distribution is key for the contour-finding task and illustrates the nonrobustness of the Gaussian observation model, due to which outliers strongly influence the inference.

Second, we document that the simple adjustment of using Student- t observations significantly mitigates the above issue. t -GP performs consistently and significantly better than Gaussian GP in essentially all settings. This result is true even when both models are misspecified (the Gsn/mix and $t/hetero$ cases). The performance of t -GP was still better (though not statistically significantly so) when we tested it in the setting of homoscedastic Gaussian noise (not shown in the plots). The latter fact is not surprising— t -GP adaptively learns the degrees-of-freedom parameter ν and hence can “detect” Gaussian noise by setting ν to be large. Conversely, in heavy-tailed noise cases, the use of t samples will effectively ignore outliers (O’Hagan, 1979) and thus produce more accurate predictions than working with a Gaussian observation assumption. We find that t -GP can handle complex noise structures and offers a good choice for all-around performance, making it a good default selection for applications. It brings smaller error rate \mathcal{ER} , more stable hyperparameter estimation, less contour bias, and tighter contour CI. Moreover t -GP is significantly better than all the other GPs in eight of the twelve setups, indicating that t -GP is essentially the best out of all GP metamodels in most cases.

Third, we also inspect the performance of the TP metamodel. As shown in Table 3, TP is the best in two cases out of the twelve, both of which are with the $t/small$ noise. We note that TP works worst in $t/hetero$ cases, having both large error rate \mathcal{ER} and empirical error \mathcal{E} . Therefore, TP does not work well in cases with low signal-to-noise ratio or greatly misspecified noise. This may be related to the parameterization of TPs, with noise handled in the kernel, which seems less robust to misspecification. Also, since TPs revert to GPs as n increases, the advantage of flexibility offered by the modeling decreases as iterations goes and thus may not last enough for low signal-to-noise ratios, which require more samples. It is apparent for instance in Figure 5, where the error at step 0 is lower than for its counterparts.

Fourth, CI-GP is also better than Gaussian GP in some cases with tMSE and MCU designs (except for the 6-D cases, where the error rate \mathcal{ER} of MCU is not significantly different from that of ICU, although mean of ICU is slightly smaller). There is significant improvement for models with low signal-to-noise ratio; the only exception is for the low-noise setup where CI-GP underperforms classical GP. This matches the intuition that employing classification “flattens” the signal by removing outliers. By considering only the sign of the response, the classification model disregards its magnitude, simplifying the noise structure at the cost of some information loss. The net effect is helpful when the noise is mis-specified or too strong so as to interfere with learning the mean response. It is detrimental if the above gain is outweighed by the information loss, as apparently happens in the 6-D experiments. Of note, CI-GP with MCU design has the smallest error rate among all models in one ($t/hetero$ in 1-D) out of 12 cases shown in Table 3. We also observe, however, that the stability of CI-GP is highly dependent on the design: some designs create large across-run variations in performance. We hypothesize that this is due to a more complex procedure for learning the hyperparameters of CI-GP; therefore, designs that are not aggressive enough to explore the zero contour region (such as cSUR) face difficulties in estimating ϑ . As a result, relative to t -GP, CI-GP has larger sampling variances.




Figure 3: Boxplots of final error rate $\mathcal{ER}^{(n)}$ from (2.18) across designs (rows) and noise setups (columns). Colors correspond to different GP metamodells. Note that x -axis limits are different across columns. Top row is for the 1-D experiment and design size $n = 100$; middle row: 2-D Branin-Hoo function with $n = 150$; bottom row: 6-D Hartman6 function with $n = 1000$.

5.3 Empirical Errors and Uncertainty Quantification

Figure 4 shows the empirical errors \mathcal{E} that are supposed to proxy the true error rates \mathcal{ER} . Overall, we find that MCU tends to produce the largest \mathcal{E} , and ICU the smallest. These results are consistent with their design construction and local behavior: MCU heavily concentrates around $\partial\hat{S}$, which leads to little information collected about other regions, especially around the boundaries of sample space D and hence relatively large $\bar{E}(x)$ there, inflating \mathcal{E} . Conversely, the objective function of ICU is precisely the myopic minimization of \mathcal{E}_{n+1} . The other two designs are intermediate versions in terms of minimizing \mathcal{E} . The tMSE heuristic tends to target the zero contour plus the edges of D , while cSUR tends to broadly target a “credible band” around $\partial\hat{S}$. Both approaches are better at reducing \mathcal{E} compared with MCU but are not directly aimed at this. This logic is less consistent for the classification models, where tMSE often yields the lowest \mathcal{E} . This result echoes Section 5.2, namely, that classification GPs tend to perform better with MCU and tMSE designs in lower dimensional cases. TPs tend to have a greater empirical error \mathcal{E} when the noise is misspecified or in higher dimensional experiments, consistent with the conclusions obtained regarding the error rate \mathcal{ER} .

As a further visualization, Figure 5 shows the median error rate \mathcal{ER} (2.18) and empirical error \mathcal{E} in Eq. (2.20) as a function of step n in the 2-D *Gsn/mix* experiments. This illustrates the learning rates of different schemes as data is collected and offers a further comparison between the true \mathcal{ER} and the self-reported \mathcal{E} of the same scheme. We observe that some metamodells underperform for very low n , even if they eventually “catch up” after sufficiently large simulation budget. This is especially pronounced for the

Table 3: Mean (w/standard deviation) error rate \mathcal{R} and corresponding best-performing sequential design heuristic for the 1-D, 2-D, and 6-D synthetic case studies. Results are based on 100 macroreplications of each scheme.

Model		$t/small$		$t/large$		Gsn/mix		$t/hetero$
<i>1-D Quadratic</i>								
GP	$tMSE$	0.73% (0.60%)	tMSE	3.24% (2.79%)	MCU	3.87% (3.17%)	cSUR	15.68% (12.15%)
t -GP	tMSE	0.80% (0.93%)	$tMSE$	3.15% (1.83%)	cSUR	3.28% (3.74%)	cSUR	12.50% (9.05%)
TP	MCU	0.97% (0.84%)	MCU	5.93% (5.60%)	tMSE	5.09% (4.40%)	ICU	16.44% (10.14%)
CI-GP	tMSE	0.87% (0.64%)	tMSE	3.39% (4.16%)	MCU	4.99% (3.77%)	MCU	8.83% (7.35%)
<i>2-D Branin-Hoo</i>								
GP	MCU	1.78% (0.57%)	cSUR	4.75% (1.95%)	ICU	4.92% (1.86%)	MCU	10.36 % (3.94%)
t -GP	MCU	1.70% (0.29%)	$tMSE$	3.95% (1.47%)	ICU	4.10% (2.07%)	$tMSE$	9.00% (8.66%)
TP	$tMSE$	1.27% (0.41%)	MCU	4.79% (1.84%)	ICU	5.19% (1.68%)	MCU	12.75 % (9.02%)
CI-GP	MCU	1.56% (0.51%)	MCU	4.27% (1.59%)	MCU	5.71% (1.85%)	tMSE	13.23% (7.74%)
<i>6-D Hartman6</i>								
GP	ICU	3.81% (0.34%)	ICU	5.33% (0.54%)	ICU	5.19% (0.70%)	MCU	11.67% (2.89%)
t -GP	ICU	3.75% (0.40%)	ICU	3.98% (0.47%)	ICU	4.86% (0.67%)	ICU	8.25% (1.60%)
TP	ICU	1.25% (0.20%)	MCU	5.66% (1.98%)	MCU	4.88% (0.88%)	MCU	10.69% (2.34%)
CI-GP	MCU	7.99% (4.69%)	ICU	7.20% (0.66%)	ICU	8.31% (2.44%)	ICU	11.11% (2.20%)

classification CI-GP metamodel, which yields very high $\mathcal{E}\mathcal{R}^{(n)}$ (which is also much higher than the self-reported \mathcal{E}) for n small. We also note that CI-GP appears to enjoy faster reduction in $\mathcal{E}\mathcal{R}^{(n)}$ compared with the baseline Gaussian GP, which we conjecture is due to better resistance against Y -outliers that distract plain GP’s inference of S . Comparing the two rows of the figure, we note that discrepancies between $\mathcal{E}\mathcal{R}$ and \mathcal{E} correlate with degraded performance, namely, the metamodel being unable to properly learn the response surface is reflected in poor uncertainty quantification. Moreover, the results suggest that the wedge in performance of different design criteria tends to persist; for example MCU and ICU frequently have not only the highest/lowest $\mathcal{E}^{(n)}$ but also the slowest/fastest rate of *reduction* in $\mathcal{E}^{(n)}$ as n grows. Consistent with results in Section 5.2, CI-GP with ICU criterion yields both greater error rate $\mathcal{E}\mathcal{R}$ and empirical error \mathcal{E} in 2-D experiments.

5.4 Designs for Contour Finding

A key goal of our study is qualitative insights about experimental designs most appropriate for noisy level set estimation. Through identifying the best-performing heuristics we get an inkling regarding the structure of near-optimal designs for (1.1). In this section we illustrate the latter within a 2-D setup that can be conveniently visualized. Taking the $t/large$ experiment as an example, in Figure 6 we plot the fitted zero contour $\partial\hat{S}$ at $N = 150$ together with the chosen inputs $\mathbf{x}_{1:150}$ across the six metamodels and the four \mathcal{I} heuristics. As expected, most of the designs are around the contour ∂S , which is the intuitive approach to minimize the error $\mathcal{E}\mathcal{R}$. Nevertheless, we observe significant differences in designs produced by different \mathcal{I} ’s. The MCU criterion places most of the samples close to the estimated zero contour $\partial\hat{S}$, reflecting its aggressive exploitation nature. For tMSE, the samples tend to cluster at several subregions of $\partial\hat{S}$ and on the edges of D . For cSUR, $\mathbf{x}_{1:n}$ cover a band along $\partial\hat{S}$, resembling the shape of the MCU design but more

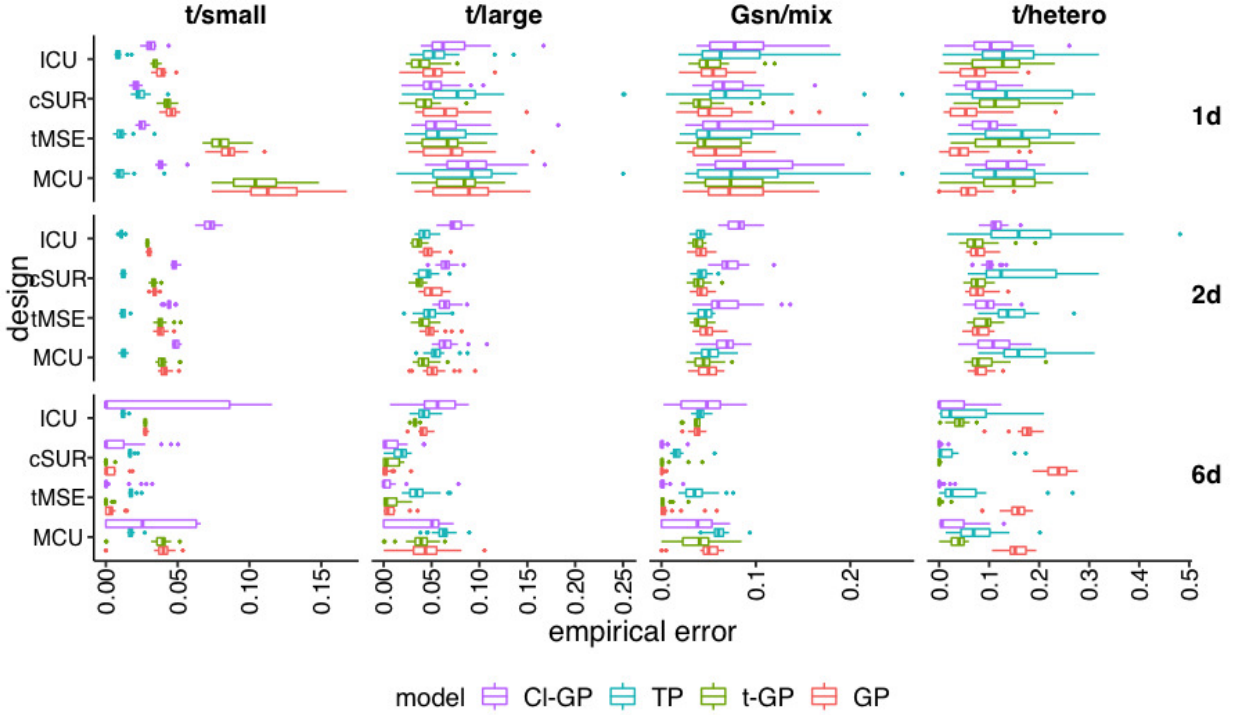


Figure 4: Empirical error $\mathcal{E}^{(n)}$ in Eq. (2.19) for GP, t -GP, TP, CI-GP, and MCI-GP metamodels (colors), using MCU, tMSE, cSUR and ICU-based designs (sub rows) with $n = 100$ in 1-D, $n = 150$ in 2-D, and $n = 1000$ in the 6-D experiments (rows).

dispersed. For ICU the design is much more exploratory, covering a large swath of D . All these findings echo the 1-D setting in Figure 2.

One feature we observe is a so-called edge effect, that is, designs that focus on the edges of the input space. This effect arises due to the intrinsically high posterior uncertainty $s(x)$ for x around ∂D . It features strongly in tMSE and cSUR (which have about 45% of the inputs along the edge) and to some extent in ICU (about 30% of inputs in this example). In contrast, MCU strongly discounts any region that is far from $\partial \hat{S}$. In the given 2-D experiment, we obtain some inputs directly on the boundary $\partial D = \{x^1 \in \{0, 1\} \cup \{x^2 \in \{0, 1\}\}$, that is, the constraint $x \in D$ is binding and the maximizer of $\mathcal{I}_n(\cdot)$ lies at its upper/lower bound. A related phenomenon is the concentration of inputs in the top/left and bottom/right corners of D , which are associated with the highest uncertainty about the level set due to the confluence of the zero contour passing there and reduced spatial information from being on the edge of D .

Another noteworthy feature is *replication* of some inputs, that is, repeated selection of the same x sites. This does not occur for MCU, but happens for ICU, tMSE and cSUR that frequently (across algorithm runs) sample repeatedly at the vertices of D (indicated by the size of the corresponding marker in Figure 6). The replication is typically mild (we observe 145+ unique designs among a total of 150 x_n 's). This finding echoes Binois et al. (2019) the importance for the metamodel to distinguish between signal and noise, which is a key distinction with the noise-free setting $\epsilon(x) \equiv 0$.

Given the above discussion and the relative overhead of the different heuristics, we conclude that in lower dimensional problems, there is little benefit to using the more sophisticated ICU criterion, while for higher dimensional problems, ICU criterion is significantly better than the others. Beyond that, tMSE appears to be an adequate and cheaper choice. However, as the input space becomes more complicated, we

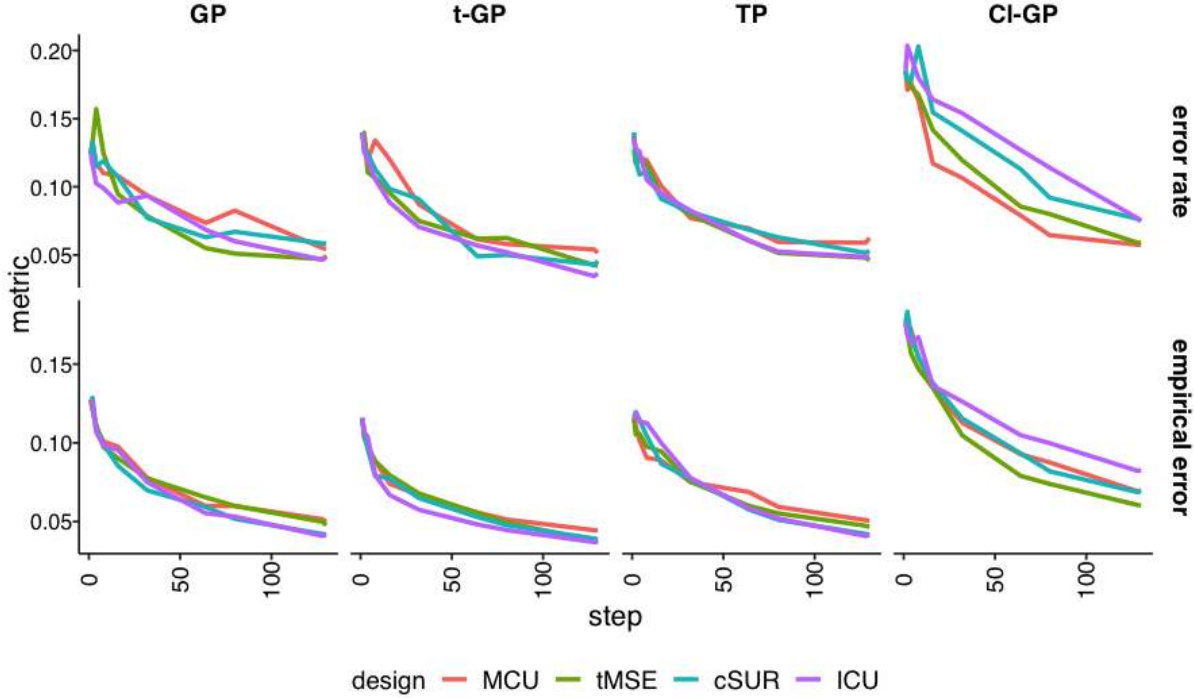


Figure 5: Error rate $\mathcal{R}^{(n)}$ (2.18) and surrogate-based uncertainty measure $\mathcal{E}^{(n)}$ (2.20) as a function of step n in the 2-D *Gsn/mix* setting. We compare six metamodels (columns) and four DoE’s (colors). The y -axis limits differ across rows. We plot median results across 20 macroreplications of each scheme.

need more exploration over the input space and the explorative criteria like ICU start to shine.

The performance of designs differs when combined with different GP metamodels. Table 3 shows that there is no one overall “best” design for all metamodels across all cases. However, it does suggest some design/metamodel “combos” that work better than others, especially in the 1-D and 2-D experiments. The classification GPs seem to prefer more aggressive designs, such as MCU, while the regression GPs prefer more exploratory designs, such as ICU. In higher dimensions, ICU usually wins across all metamodels in accuracy; see the results of 6-D experiments in Table 3.

6 Application to Optimal Stopping Problems in Finance

In our next case study we consider contour finding for determining the optimal exercise policy of a Bermudan financial derivative, cf. Section 1.2. The underlying simulator is based on a d -dimensional geometric Brownian motion (\mathbf{X}_t) that represents prices of d assets and follows the log-normal dynamics

$$\mathbf{X}_{t+\Delta t} = \mathbf{X}_t \exp \left(\left(r - \frac{1}{2} \sigma^2 \right) \Delta t + \Sigma \Delta \mathbf{W}_t \right), \quad \Delta \mathbf{W}_t \sim \mathcal{N}(0, \Delta t \mathbf{I}), \quad (6.1)$$

where \mathbf{I} is the $d \times d$ identity matrix. Let $h(t, x)$ be the option payoff from exercising when $\mathbf{X}_t = x \in \mathbb{R}^d$. Exercising is allowed every Δt time units, up to the option maturity T , so that we wish to determine the collection $\{S_t : t \in \{\Delta t, 2\Delta t, \dots, T - \Delta t\}\}$, which are the zero level sets of the timing function $x \mapsto T(t, x)$. During the backward dynamic programming, we iterate over $t = T, T - \Delta t, \dots, 0$, and the




Figure 6: Estimates of the zero contour $\partial\hat{S}$ for the 2-D Branim-Hoo example with t/large noise setting. We show $\partial\hat{S}^{(n)}$ (red solid line) at step $n = 150$, with its 95% credible band (red dashed lines), the true zero contour ∂S (black solid line) and the sampled inputs $\mathbf{x}_{1:n}$ (replicates indicated with larger symbols). We compare across the six metamodels (rows) and four DoE heuristics (columns).

simulator of $T(t, x)$ returns the difference between the pathwise payoff along a trajectory of $(\mathbf{X}_{t:T})$ that is based on the forward exercise strategy summarized by the forward-looking $\{\hat{S}_s, s > t\}$, and $h(t, x)$.

As discussed in Ludkovski (2018), this setting implies a skewed, non-Gaussian, heteroskedastic distribution of the simulation noise and is a challenging stochastic contour-finding problem. Note that in order to reflect the underlying distribution of \mathbf{X}_t at time t (conditional on the given initial value $\mathbf{X}_0 = x_0$) the weighting measure $\mu(x) = p_{X_t}(x|x_0)$ is used. Thus, $\mu(\cdot)$ is log-normal based on (6.1) and is multiplied by the respective \mathcal{I}_n criteria when selecting $x_{n+1} = \arg \max_{x \in D} \mathcal{I}_n(x)\mu(x)$. In line with the problem context, we no longer directly measure the accuracy of learning $\{S_t\}$ but instead focus on the ultimate output of RMC, which is the estimated option value in (1.5). The latter must itself be numerically evaluated via an out-of-sample Monte Carlo simulation that averages realized payoffs along a large database of M paths $x_{0:T}^{1:M}$:

$$\hat{V}(0, x_0) = \frac{1}{M} \sum_{m=1}^M h(\tau^m, x_{\tau^m}^{(m)}), \quad \tau^m = \inf\{t : x_t^{(m)} \in \hat{S}_t\}. \quad (6.2)$$

Since our goal is to find the *best* exercise value, higher \hat{V} 's indicate a better approximation of $\{S_t\}$.

To allow a direct comparison, we set parameters matching the test cases in Ludkovski (2018), considering a 2-D and 3-D example. In both cases the volatility matrix $\Sigma = \sigma \mathbf{I}$ in (6.1) is diagonal with constant terms; that is, the coordinates $\mathbf{X}_{1:n}^1, \dots, \mathbf{X}_{1:n}^d$ are independently and identically distributed. As a first example, we consider a 2-D basket Put option with parameters $r = 0.06, \sigma = 0.2, \Delta t = 0.04, \mathcal{K} = 40, T = 1$.

The payoff is $h(t, x) = e^{-rt}(\mathcal{K} - \frac{x^1+x^2}{2})_+$ with $K = 40$. Here it is known that stopping becomes optimal once both asset prices x^1 and x^2 become sufficiently low, so the level set S_t is toward the bottom-left of D ; see Fig 7. In contrast, stopping is definitely suboptimal when $h(t, x) = 0 \Leftrightarrow (x^1 + x^2)/2 > \mathcal{K}$. Consequently, the input sample space is taken to be $D = [25, 55] \times [25, 55] \cap \{x^1 + x^2 \leq 80\}$.

In this first case study, the timing value $h(t, x)$ is known to be *monotonically* increasing in the asset price x . To incorporate this constraint, we augment the four main metamodels (GP, t-GP, CI-GP and TP) with two monotonic versions, M-GP and MCI-GP. By constraining the fitted \hat{f} to be monotone, we incorporate structural knowledge about the ground truth, which in turn reduces posterior uncertainty and thus might produce more accurate estimates of S . Monotonicity of the metamodel for f is also one sufficient way to guarantee that the outputted level set \hat{S} is a *connected* subset of D .

Our monotone GPs are based on Riihimäki and Vehtari (2010). In general, any infinite-dimensional Gaussian process is intrinsically non monotone, since the multivariate Gaussian distribution is always supported on the entire \mathbb{R}^d , rather than an orthant. Nevertheless, local monotonicity in \hat{f} can be enforced by considering the gradient ∇f of f which is also a Gaussian process. Specifically, Riihimäki and Vehtari (2010) proposed to adaptively add virtual observations for ∇f ; we employ the resulting implementation in the public `GPstuff` library (Vanhatalo et al., 2013) to build our own version dubbed M-GP. We employ the same strategy to restrict the coordinates z^j of the latent logistic GP Z to be increasing (decreasing) across D . Implementation details are included in Appendix D.

As a second example, we consider a 3-D max-Call $x \in \mathbb{R}^3$ with payoff $h(t, x) = e^{-rt}(\max(x^1, x^2, x^3) - \mathcal{K})_+$. The parameters are $r = 0.05, \delta = 0.1, \sigma = 0.2, X_0 = (90, 90, 90), \mathcal{K} = 100, T = 3$ and $\Delta t = 1/3$. Since stopping is ruled out when $h(t, x) = 0 \Leftrightarrow \max(x^1, x^2, x^3) < \mathcal{K}$, the sample space is taken to be $D = [50, 150]^3 \cup \{\max(x^1, x^2, x^3) > \mathcal{K}\}$. In this case, stopping is optimal if *one* of the coordinates x^i is significantly higher than the other two, so S_t consists of three disconnected components. In this problem there is no monotonicity, so we employ only the GP, t-GP, CI-GP, and TP metamodels.

Because of the iterative construction of the simulator, the signal-to-noise ratio gets low for small t 's. The variance $\tau^2(x)$ is also highly state-dependent, tending to be smaller for sites further from the zero-contour. To alleviate this misspecification and reduce metamodel overhead, we employ *batched* designs (Ludkovski, 2018; Ankenman et al., 2010), reusing $x \in D$ for r replications to collect observations $y^{(1)}(x), \dots, y^{(r)}(x)$ from the corresponding simulator $Y(x)$. Then, we treat the mean of the r observations,

$$\bar{y}(x) = \frac{1}{r} \sum_{i=1}^r y^{(i)}(x), \quad (6.3)$$

as the response for input x and use $(x, \bar{y}(x))$ as a single design entry. The statistical properties of \bar{y} are improved compared with the raw observations y : it is more consistent with the Gaussian assumption thanks to the Central Limit Theorem (CLT), and its noise variance $\bar{\tau}^2(x) = \tau^2(x)/r$ is much smaller. Since the expense of sequential design of GP metamodels comes mainly from choosing the new input at each step, the reduction in budget $n = N/r$ by a factor of r significantly speeds their fitting and updating, with n for the number of unique inputs.

For the 2-D Put case study, we then test a total of three budget settings: (i) $r = 3, n = 80$ (low budget of $N = 240$ simulations); (ii) $r = 15, n = 80$ (high budget $N = 800$ with moderate replication); (iii) $r = 48, n = 25$ (high $N = 800$ with high replication). Comparing (ii) and (iii) shows the competing effects of having non-Gaussian noise (for lower r) and small design size (low n). The initial design size $n_0 = 10$. In this example, taking $n \gg 80$ gives only marginally better performance but significantly

raises the computation time and hence is ruled out as impractical. Three setups are investigated for the 3-D example: $r = 3, n = 100$ (low-budget of $N = 300$), $r = 20, n = 100$ (moderate-budget of $N = 2000$) and $r = 20, n = 200$ (high budget $N = 4000$), both with $n_0 = 30$. In all examples, the results are based on 25 runs of each scheme and are evaluated through the resulting expected reward $\hat{V}(0, x_0)$ (6.2) on a fixed out-of-sample testing set of $M = 160,000$ paths of $\mathbf{X}_{0:T}$.

Table 4: Performance of different designs and models on the 2-D Bermudan Put option in Section 6. Results are the mean (standard deviation) payoff of 25 runs of experiments evaluating on the same out-of-sample testing set of $M = 160000$ $\mathbf{X}_{0:T}$ -paths at each run.

	LHS	MCU	tMSE	cSUR	ICU
R = 3, n* = 80					
GP	1.211(0.120)	1.425(0.008)	1.427(0.007)	1.431(0.009)	1.431(0.007)
t-GP	1.125(0.113)	1.409(0.013)	1.417(0.008)	1.409(0.010)	1.406(0.013)
TP	1.179 (0.133)	1.408 (0.022)	1.414 (0.008)	1.378 (0.044)	1.316 (0.037)
M-GP	1.403(0.014)	1.438(0.007)	1.440(0.006)	1.442(0.009)	1.433(0.005)
CI-GP	1.111(0.121)	1.395(0.015)	1.402 (0.013)	1.393(0.013)	1.391(0.013)
MCI-GP	1.407(0.008)	1.429(0.010)	1.429(0.013)	1.431(0.007)	1.396(0.019)
R = 15, n* = 80					
GP	1.425 (0.017)	1.448 (0.003)	1.450 (0.002)	1.450 (0.003)	1.449 (0.003)
t-GP	1.406 (0.033)	1.445 (0.003)	1.447 (0.002)	1.444 (0.005)	1.446 (0.004)
TP	1.414 (0.023)	1.443 (0.003)	1.443 (0.004)	1.441 (0.004)	1.430 (0.006)
M-GP	1.407 (0.008)	1.449 (0.003)	1.451 (0.002)	1.454 (0.002)	1.451 (0.003)
CI-GP	1.353 (0.050)	1.441 (0.004)	1.440 (0.003)	1.435 (0.004)	1.436 (0.005)
MCI-GP	1.416 (0.010)	1.448 (0.004)	1.449 (0.003)	1.443 (0.003)	1.418 (0.008)
R = 48, n* = 25					
GP	1.341 (0.068)	1.450 (0.003)	1.449 (0.003)	1.443 (0.004)	1.448 (0.003)
t-GP	1.336 (0.126)	1.449 (0.003)	1.452 (0.003)	1.442 (0.004)	1.449 (0.003)
TP	1.367 (0.063)	1.433 (0.006)	1.430 (0.011)	1.421 (0.039)	1.423 (0.023)
M-GP	1.415 (0.007)	1.446 (0.002)	1.444 (0.002)	1.445 (0.004)	1.442 (0.004)
CI-GP	1.110 (0.144)	1.430 (0.010)	1.434 (0.005)	1.409 (0.008)	1.388 (0.016)
MCI-GP	1.423 (0.015)	1.446 (0.004)	1.448 (0.003)	1.413 (0.024)	1.414 (0.024)

6.1 Results

Tables 4 and 5 compare the different designs and metamodels. To assess the sequential design gains, we also report the results from using a baseline nonadaptive LHS design on D . At low budget, we observe the dramatic gains of using adaptive designs for level set estimation, which allow us to obtain the same performance with an order-of-magnitude smaller simulation budget. The tMSE and cSUR criteria work best for the 2-D Put, while ICU is the best for the 3-D max-Call, indicating that the exploratory designs start to win out in more complex settings with higher d .

Regarding the metamodels, in the low-budget setups, the monotonic GP metamodel works best for

the 2-D Put and t -GP for the 3-D max-Call. For the higher budget, which also coincides with higher $r \in \{10, 50\}$, the metamodel performance is similar, with t -GP slightly better than the other GP variants. In particular, once the SNR is high, classical Gaussian GP is effectively as good as any alternative. In both examples, TP and classification metamodels do not work well, possibly because of being more sensitive to the heteroscedastic aspect. We note that TP as well as the classification metamodels suffer from instability, so that lower $\hat{V}(0, x_0)$ is matched with a high sampling standard deviation. Another observation is that CI-GP and MCI-GP perform badly with exploratory heuristic like ICU, especially with high budget.

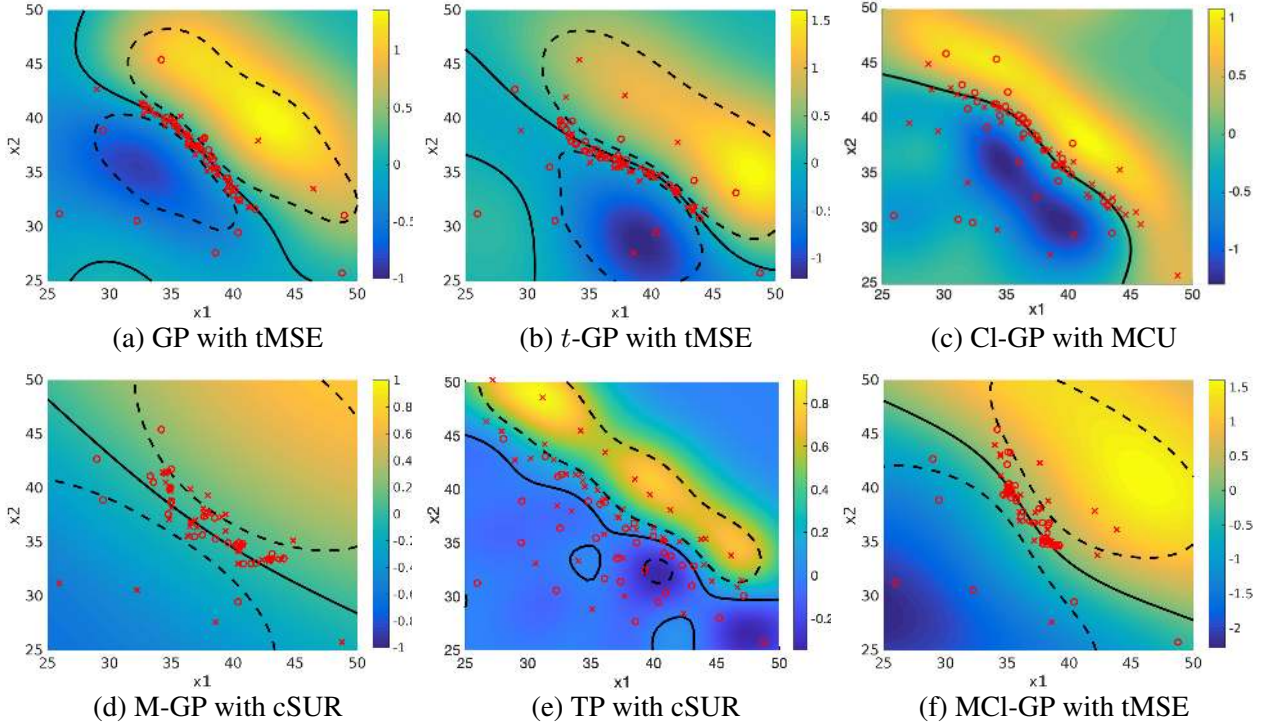


Figure 7: The estimated exercise boundary $\partial \hat{S}$ (solid line with 95% CI as dashed lines) at $t = 0.6$ for 2-D Bermudan Put from Section 6. Shading, which varies panel to panel, indicates the point estimate for the latent $\hat{f}(x)$ or $\hat{z}(x)$. We also show the design $(x_{1:n}, y_{1:n})$ with positive y_n 's marked by \times and negative y_n 's by \circ . All schemes used $R = 15, n^* = 80$.

Figure 7 shows the estimated exercise boundary $\partial \hat{S}_t$ with its 95% CI at $t = 0.4$ for the 2-D Put, for each of the five metamodels, each with the design yielding the highest payoff. We observe that all the best-performing designs look similar, placing about a dozen x_n 's (some of which are from the initial design $x_{1:n_0}$) throughout D and the rest tightly along the zero contour. The results suggest that the criteria are largely interchangeable and that simpler \mathcal{I}_n heuristics are able to reproduce the features of the more sophisticated or expensive ICU. The heuristics *do* differ in their uncertainty quantification; t -GP and GP generate tightest CI bands, while those of classification GPs and TP are too wide, indicating lack of confidence in the estimate. Of note, the regression GP metamodels (GP, t -GP and M-GP) also generate the lowest sampling variance for $\hat{V}(0, x_0)$.

Based on these results, our take-aways are threefold. First, similar to Ludkovski (2018) we document significant gains from sequential design. Second, we find that while using ICU is helpful for more complicated settings with higher dimension d and larger budget, tMSE is the recommended DoE heuristic for lower dimensional cases, achieving excellent results with minimal overhead (in particular without requir-

ing look-ahead variance). Third, we find that for applications with thousands of simulations, the Gaussian observation model is sufficient, since the underlying design needs to be replicated $r \gg 1$ in order to avoid excessively large \mathbf{K} -matrices. Therefore, there is little need for more sophisticated metamodels, although useful gains can be realized from enforcing the monotonic structure, if available.

Table 5: Performance of different designs and models on the 3-D Bermudan max-Call in Section 6. Results are the mean (w/standard deviation) payoff of 25 macroreplications evaluating on the same out-of-sample testing set of $M = 160000$ $\mathbf{X}_{0:T}$ -paths at each run.

	LHS	MCU	tMSE	cSUR	ICU
$\mathbf{R} = 3, \mathbf{n}^* = 100$					
GP	10.036 (0.331)	10.725 (0.095)	10.773 (0.071)	10.711 (0.086)	10.753 (0.072)
<i>t</i> -GP	9.894 (0.447)	10.736 (0.088)	10.747 (0.087)	10.720 (0.104)	10.782 (0.076)
TP	9.169 (0.354)	10.101 (0.218)	9.872 (0.102)	8.867 (0.357)	10.482 (0.156)
CI-GP	9.552 (0.567)	10.566 (0.084)	10.657 (0.097)	10.586 (0.099)	10.604 (0.119)
$\mathbf{R} = 20, \mathbf{n}^* = 100$					
GP	10.924 (0.076)	11.078 (0.029)	11.072 (0.028)	11.055 (0.032)	11.101 (0.023)
<i>t</i> -GP	10.923 (0.071)	11.061 (0.039)	11.055 (0.027)	11.044 (0.029)	11.100 (0.027)
TP	10.385 (0.178)	10.815 (0.039)	10.745 (0.045)	10.620 (0.087)	10.507 (0.087)
CI-GP	10.761 (0.112)	11.026 (0.032)	10.991 (0.037)	10.901 (0.049)	10.937 (0.041)
$\mathbf{R} = 20, \mathbf{n}^* = 200$					
GP	11.105(0.036)	11.147(0.021)	11.119(0.022)	11.131(0.018)	11.178(0.020)
<i>t</i> -GP	11.090(0.034)	11.141(0.019)	11.126(0.020)	11.115(0.027)	11.175(0.021)
TP	10.585 (0.118)	10.896 (0.030)	10.811 (0.035)	10.764 (0.041)	10.638 (0.038)
CI-GP	10.995(0.059)	11.109(0.025)	11.056(0.040)	10.985(0.027)	11.010(0.029)

7 Conclusion

We have carried a comprehensive comparison of five metamodels and four design heuristics on 18 case studies (4×3 synthetic, plus six real-world). In sum, the considered alternatives to standard Gaussian-observation GP do perform somewhat better. In particular, *t*-GP directly nests plain GP and hence essentially always matches or exceeds the performance of the latter. We also observe gains from using CI-GP when SNR is low and from monotonic surrogates when the underlying response is monotone. That being said, final recommendation regarding the associated benefit depends on computational considerations, as the respective overhead becomes larger (and exact updating of the metamodel no longer possible).

In terms of design, we advocate the benefits of tMSE in low dimensional simulations, which generates high-performing experimental designs without requiring expensive acquisition function (or even look-ahead variance). The tMSE criterion does sometimes suffer from the tendency to put many designs at the edge of the input space but otherwise tends to match the performance of more complex and computationally intensive \mathcal{I}_n 's. For complex simulations, ICU is probably still the best choice (although in that case, random-set-based heuristics should also be considered). Especially in higher dimensions with misspecified noise, ICU

is the best choice among all designs. We also stress that the user ought to thoughtfully pick the *combination* of sequential design and metamodel, since cross-dependencies are involved (e.g., classification metamodels generally do not work well with the ICU criterion in lower dimension).

Acknowledgements

XL and ML are partially supported by NSF DMS-1521743 and DMS-1821240. The work of MB is partially supported by NSF DMS-1521702 and the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under Contract No. DE-AC02-06CH11357.

References

- Ankenman B, Nelson BL, Staum J (2010) Stochastic kriging for simulation metamodeling. *Operations Research* 58(2):371–382
- Azzimonti D, Bect J, Chevalier C, Ginsbourger D (2016) Quantifying uncertainties on excursion sets under a Gaussian random field prior. *SIAM/ASA Journal on Uncertainty Quantification* 4(1):850–874
- Azzimonti D, Ginsbourger D, Chevalier C, Bect J, Richet Y (2019) Adaptive design of experiments for conservative estimation of excursion sets. *Technometrics* pp 1–14
- Bect J, Ginsbourger D, Li L, Picheny V, Vazquez E (2012) Sequential design of computer experiments for the estimation of a probability of failure. *Statistics and Computing* 22(3):773–793
- Bichon BJ, Eldred MS, Swiler LP, Mahadevan S, McFarland JM (2008) Efficient global reliability analysis for nonlinear implicit performance functions. *AIAA Journal* 46(10):2459–2468
- Binois M, Gramacy RB, Ludkovski M (2018) Practical heteroskedastic Gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics* 27(4):808–821
- Binois M, Huang J, Gramacy RB, Ludkovski M (2019) Replication or exploration? Sequential design for stochastic simulation experiments. *Technometrics* 61(1):7–23
- Bogunovic I, Scarlett J, Krause A, Cevher V (2016) Truncated variance reduction: A unified approach to Bayesian optimization and level-set estimation. In: *Advances in neural information processing systems*, pp 1507–1515
- Bolin D, Lindgren F (2015) Excursion and contour uncertainty regions for latent Gaussian models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 77(1):85–106
- Bryan B, Schneider J (2008) Actively learning level-sets of composite functions. In: *Proceedings of the 25th International Conference on Machine Learning, ACM*, pp 80–87
- Bryan B, Nichol RC, Genovese CR, Schneider J, Miller CJ, Wasserman L (2006) Active learning for identifying function threshold boundaries. In: *Advances in neural information processing systems*, pp 163–170
- Chevalier C, Ginsbourger D, Bect J, Molchanov I (2013) Estimating and quantifying uncertainties on level sets using the Vorob’ev expectation and deviation with Gaussian process models. In: *mODa 10—Advances in Model-Oriented Design and Analysis*, Springer, pp 35–43

- Chevalier C, Bect J, Ginsbourger D, Vazquez E, Picheny V, Richet Y (2014a) Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set. *Technometrics* 56(4):455–465
- Chevalier C, Ginsbourger D, Emery X (2014b) Corrected kriging update formulae for batch-sequential data assimilation. In: *Mathematics of Planet Earth*, Springer, pp 119–122
- Echard B, Gayton N, Lemaire M (2010) Kriging based Monte Carlo simulation to compute the probability of failure efficiently: AK-MCS method. *6emes Journées Nationales de Fiabilité*, 24–26 mars, Toulouse, France
- Gotovos A, Casati N, Hitz G, Krause A (2013) Active learning for level set estimation. In: *Twenty-Third International Joint Conference on Artificial Intelligence*
- Gramacy RB, Lee HK (2009) Adaptive design and analysis of supercomputer experiments. *Technometrics* 51(2):130–145
- Gramacy RB, Ludkovski M (2015) Sequential design for optimal stopping problems. *SIAM Journal on Financial Mathematics* 6(1):748–775
- Hu R, Ludkovski M (2017) Sequential design for ranking response surfaces. *SIAM/ASA Journal on Uncertainty Quantification* 5(1):212–239
- Jalali H, Van Nieuwenhuysse I, Picheny V (2017) Comparison of Kriging-based methods for simulation optimization with heterogeneous noise. *European Journal of Operational Research* 261(1):279–301
- Jones DR, Schonlau M, Welch WJ (1998) Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* 13(4):455–492
- Jylänki P, Vanhatalo J, Vehtari A (2011) Robust Gaussian process regression with a Student-t likelihood. *Journal of Machine Learning Research* 12(Nov):3227–3257
- Longstaff FA, Schwartz ES (2001) Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies* 14(1):113–147
- Ludkovski M (2018) Kriging metamodels for Bermudan option pricing. *Journal of Computational Finance* to appear, arXiv preprint arXiv:1509.02179
- Minka TP (2001) Expectation propagation for approximate Bayesian inference. In: *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers Inc., pp 362–369
- Mukhopadhyay S, Mahmoodi H, Roy K (2005) Modeling of failure probability and statistical design of SRAM array for yield enhancement in nanoscaled CMOS. *IEEE Transactions on Computer-aided Design of Integrated Circuits and Systems* 24(12):1859–1880
- O’Hagan A (1979) On outlier rejection phenomena in Bayes inference. *Journal of the Royal Statistical Society Series B (Methodological)* pp 358–367
- Picheny V, Ginsbourger D (2013) A nonstationary space-time Gaussian process model for partially converged simulations. *SIAM/ASA Journal on Uncertainty Quantification* 1(1):57–78

- Picheny V, Ginsbourger D, Roustant O, Haftka RT, Kim NH (2010) Adaptive designs of experiments for accurate approximation of a target region. *Journal of Mechanical Design* 132(7):071008
- Picheny V, Ginsbourger D, Richet Y, Caplin G (2013a) Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics* 55(1):2–13
- Picheny V, Wagner T, Ginsbourger D (2013b) A benchmark of kriging-based infill criteria for noisy optimization. *Structural and Multidisciplinary Optimization* 48(3):607–626
- Ranjan P, Bingham D, Michailidis G (2008) Sequential experiment design for contour estimation from complex computer codes. *Technometrics* 50(4):527–541
- Rasmussen CE, Nickisch H (2010) Gaussian processes for machine learning (gpml) toolbox. *Journal of Machine Learning Research* 11(Nov):3011–3015
- Riihimäki J, Vehtari A (2010) Gaussian processes with monotonicity information. In: *AISTATS*, vol 9, pp 645–652
- Santner TJ, Williams BJ, Notz WI (2013) *The Design and Analysis of Computer Experiments*. Springer Science & Business Media
- Shah A, Wilson A, Ghahramani Z (2014) Student-t processes as alternatives to Gaussian processes. In: *Artificial Intelligence and Statistics*, pp 877–885
- Srinivas N, Krause A, Kakade SM, Seeger MW (2012) Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory* 58(5):3250–3265
- Tsitsiklis JN, Van Roy B (2001) Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* 12(4):694–703
- Vanhatalo J, Riihimäki J, Hartikainen J, Jylänki P, Tolvanen V, Vehtari A (2013) Bayesian modeling with Gaussian processes using the GPstuff toolbox. *The Journal of Machine Learning Research* 14(1):1175–1179
- Wang Z, Shi JQ, Lee Y (2017) Extended t-process regression models. *Journal of Statistical Planning and Inference* 189:38–60
- Williams CK, Barber D (1998) Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(12):1342–1351
- Williams CK, Rasmussen CE (2006) *Gaussian Processes for Machine Learning*. MIT Press

A Choice of $\gamma^{(n)}$ for MCU

Following up on the discussion in Section 3, we investigate the role of $\gamma^{(n)}$ in the performance of MCU. Table 6 shows the error rate \mathcal{R} for GP and t -GP metamodels with MCU acquisition function in the 2D synthetic experiments across three constant values of $\gamma^{(n)}$ (constant $\gamma^{(n)}$ was also employed in Gotovos et al. (2013) and Bryan et al. (2006)). We observe that generally the impact of $\gamma^{(n)}$ is secondary (with Gaussian GP being more sensitive), and moreover there is no single choice that works the best across all

cases. As illustrated in Figure 8, large $\gamma^{(n)}$ favors space-filling, while small $\gamma^{(n)}$ favors exploitation in regions close to the boundary ∂S .

Generally, smaller $\gamma^{(n)}$'s work better in cases with less noise (e.g. $\gamma^{(n)} = 10$ is worst in the t /small noise scenario). This validates our recommendation that $\gamma^{(n)}$ should be adaptive to the signal-to-noise ratio of $\hat{f}^{(n)}(x)$ and $s^{(n)}(x)$. Clearly $s^{(n)}(x)$ depends strongly on the original noise specification which is another reason why a fixed ‘‘universal’’ $\gamma^{(n)}$ is inappropriate (unlike in deterministic experiments, there is no simple way to normalize the variance of ϵ). Note that since $s^{(n)}$ decreases in n , the signal-to-noise ratio increases over time, which is consistent with the theoretical results that $\gamma^{(n)}$ should increase with n .

Model	$\gamma^{(n)} = 0.5$	$\gamma^{(n)} = 1.96$	$\gamma^{(n)} = 10$
t/small			
GP	1.87% (0.36%)	1.82% (0.51%)	2.09% (0.54%)
t -GP	1.80% (0.52%)	1.73% (0.22%)	1.84% (0.42%)
t/large			
GP	5.20% (2.33%)	5.59% (2.22%)	4.94% (1.79%)
t -GP	3.80% (1.25%)	4.24% (2.12%)	4.01% (1.43%)
Gsn/mix			
GP	5.10% (2.36%)	5.53% (1.79%)	6.01% (3.08%)
t -GP	4.63% (1.74%)	3.92% (1.26%)	4.39% (1.40%)
t/hetero			
GP	11.23% (5.08%)	10.52% (7.05%)	13.63% (6.32%)
t -GP	7.34% (3.96%)	10.58% (8.25%)	7.77% (3.55%)

Table 6: Mean (w/standard deviation) error rate \mathcal{R} for MCU in 2D synthetic experiments. Results are based on 20 macro-replications of each scheme.

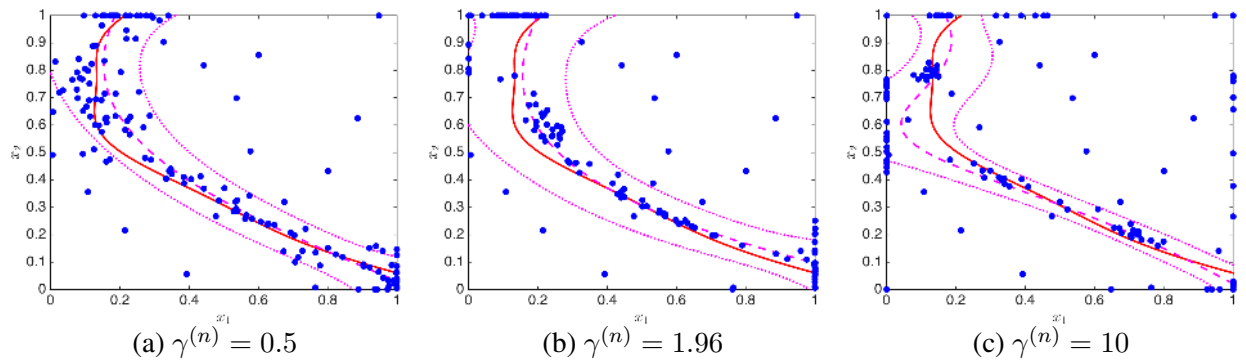


Figure 8: The estimated boundary $\partial \hat{S}$ (dashed line with 95% CI as dotted lines). Blue dots are samples selected by MCU with $\gamma^{(n)} = 0.5$, $\gamma^{(n)} = 1.96$, and $\gamma^{(n)} = 10$.

B Gaussian Process Classification

Similar to the t -GP, we use a Laplace approximation for the non-Gaussian $p_{\text{Cl}}(\mathbf{z}|\mathbf{x}_{1:n}, \check{\mathbf{y}}_{1:n})$ in Eq. (2.11),

$$p_{\text{Cl}}(\mathbf{z}|\mathbf{x}_{1:n}, \check{\mathbf{y}}_{1:n}) \approx q_{\text{Cl}}(\mathbf{z}|\mathbf{x}_{1:n}, \check{\mathbf{y}}_{1:n}) = \mathcal{N}(\tilde{\mathbf{z}}^{(n)}, \Sigma_{\text{Cl}}^{-1}), \quad (\text{B.1})$$

where we again use the mode $\tilde{\mathbf{z}}^{(n)} := \arg \max_{\mathbf{z}} p(\mathbf{z}|\mathcal{A}_n)$ and Σ_{Cl} is the Hessian of the negative log posterior at $\tilde{\mathbf{z}}^{(n)}$:

$$\Sigma_{\text{Cl}} = -\nabla^2 \log p_{\text{Cl}}(\mathbf{z}|\mathcal{A}_n)|_{\mathbf{z}=\tilde{\mathbf{z}}^{(n)}} = \mathbf{K}^{-1} + \mathbf{V}, \quad (\text{B.2})$$

and $\mathbf{V} = -\nabla^2 \log p(\check{\mathbf{y}}_{1:n}|\mathbf{z})|_{\mathbf{z}=\tilde{\mathbf{z}}^{(n)}}$ is diagonal with elements

$$v_i = \mathbf{V}_{ii} = -\frac{\partial^2}{\partial z_i^2} \log p(\check{y}_i|z_i)|_{z_i=\tilde{z}_i^{(n)}} = \frac{\phi(\tilde{z}_i^{(n)})^2}{\Phi(\check{y}_i \tilde{z}_i^{(n)})^2} + \frac{\check{y}_i \tilde{z}_i^{(n)} \phi(\tilde{z}_i^{(n)})}{\Phi(\check{y}_i \tilde{z}_i^{(n)})}, \quad (\text{B.3})$$

for $i = 1, \dots, n$, $\phi(\cdot)$ denoting the density of the standard normal distribution.

C Computation Details for Look-Ahead Variance

t -GP: To approximate $\tilde{\mathbf{f}}_{t\text{GP}}^{(n+1)}$ in t -GP, we recall that the posterior mode and the posterior mean coincide:

$$\hat{\mathbf{f}}_{t\text{GP}}^{(n)}(\mathbf{x}_{1:n}) = \mathbf{K}\mathbf{K}^{-1}\tilde{\mathbf{f}}_{t\text{GP}}^{(n)} = \tilde{\mathbf{f}}_{t\text{GP}}^{(n)}. \quad (\text{C.1})$$

Hence we can compute the expected value of $\tilde{\mathbf{f}}_{t\text{GP}}^{(n+1)}$ using the tower property:

$$\begin{aligned} \mathbb{E}[\tilde{\mathbf{f}}_{t\text{GP}}^{(n+1)}|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}] &= \mathbb{E}[\hat{\mathbf{f}}_{t\text{GP}}^{(n+1)}(\mathbf{x}_{1:n+1})|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}] \\ &= \mathbb{E}[\mathbb{E}[f(\mathbf{x}_{1:n+1})|\mathbf{x}_{1:n+1}, \mathbf{y}_{1:n+1}]|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}] \\ &= \mathbb{E}[f(\mathbf{x}_{1:n+1})|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}] = [\hat{\mathbf{f}}_{t\text{GP}}^{(n)}(\mathbf{x}_{1:n}), \hat{f}_{t\text{GP}}^{(n)}(x_{n+1})] = [\tilde{\mathbf{f}}_{t\text{GP}}^{(n)}, \hat{f}_{t\text{GP}}^{(n)}(x_{n+1})], \end{aligned} \quad (\text{C.2})$$

where the last equality follows from the BLUP property of GP estimates. Therefore, we approximate the $(n+1)$ -dimensional vector $\tilde{\mathbf{f}}_{t\text{GP}}^{(n+1)}$ with $\tilde{\mathbf{f}}_{t\text{GP}}^{(n+1)} = [\tilde{\mathbf{f}}_{t\text{GP}}^{(n)}, \hat{f}_{t\text{GP}}^{(n)}(x_{n+1})]$, where the first component is n -dimensional and the second component is a scalar. In turn, this step allows us to update the matrices $\mathbf{W}_{t\text{GP}}^{(n)}$ and $\mathbf{K}^{(n)}$ assuming a new input x_{n+1} is added. Specifically, the new entry in $\mathbf{W}_{t\text{GP}}^{(n+1)}$ is

$$\begin{aligned} w_{n+1}^{(n+1)} &= (\nu + 1) \frac{\nu\tau^2 - (y_{n+1} - \hat{f}_{t\text{GP}}^{(n+1)}(x_{n+1}))^2}{\left((y_{n+1} - \tilde{f}_{t\text{GP}}^{(n+1)}(x_{n+1}))^2 + \nu\tau^2\right)^2} \\ &\simeq (\nu + 1) \frac{\nu\tau^2 - (y_{n+1} - \hat{f}_{t\text{GP}}^{(n)}(x_{n+1}))^2}{\left((y_{n+1} - \hat{f}_{t\text{GP}}^{(n)}(x_{n+1}))^2 + \nu\tau^2\right)^2}. \end{aligned} \quad (\text{C.3})$$

Matching terms with the Gaussian observation GP, the updated variance $s_{t\text{GP}}^{(n+1)}(x_{n+1})^2$ is then approximately proportional to the current variance:

$$\frac{s_{t\text{GP}}^{(n+1)}(x_{n+1})^2}{s_{t\text{GP}}^{(n)}(x_{n+1})^2} \simeq \frac{(w_{n+1}^{(n+1)})^{-1}}{(w_{n+1}^{(n+1)})^{-1} + s_{t\text{GP}}^{(n)}(x_{n+1})^2}. \quad (\text{C.4})$$

To make this implementable at step n , we need to remove the inaccessible y_{n+1} term in both the numerator and denominator of (C.3). In principle, we could attempt to (numerically) integrate the predictive distribution $Y(x_{n+1}) \sim t_\nu(\hat{f}_t^{(n)}(x_{n+1}), \tau^2)$ against $f_t^{(n)}(x_{n+1}) \sim \mathcal{N}(\hat{f}_t^{(n)}(x_{n+1}), s_t^{(n)}(x_{n+1})^2)$; for simplicity we instead replace $(y_{n+1} - \hat{f}_{t\text{GP}}^{(n)}(x_{n+1}))^2$ with its expectation: $\mathbb{E}[(y_{n+1} - \hat{f}_{t\text{GP}}^{(n)}(x_{n+1}))^2] = \text{Var}[y_{n+1}] = \tau^2$ and therefore obtain the approximation $w_{n+1}^{(n+1)} \simeq (\nu + 1) \frac{(\nu-1)\tau^2}{(\nu+1)^2\tau^4} = \frac{\nu-1}{(\nu+1)\tau^2}$. This leads to the final look-ahead variance formula (cf. (4.4)):

$$\hat{s}_{t\text{GP}}^{(n+1)}(x_*; x_{n+1})^2 := s_{t\text{GP}}^{(n)}(x_*)^2 - \frac{v_{t\text{GP}}^{(n)}(x_*, x_{n+1})^2}{(\tau^2 \frac{\nu+1}{\nu-1}) + s_{t\text{GP}}^{(n)}(x_{n+1})^2}. \quad (\text{C.5})$$

CI-GP: Similar to the t -GP, the look-ahead variance for the classification GP is intractable since $s_{\text{Cl}}^{(n+1)}$ is based on the mode $\tilde{\mathbf{z}}_{\text{Cl}}^{(n+1)}$ of the posterior $p_{\text{Cl}}(\mathbf{z}|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}, x_{n+1}, y_{n+1})$. Similar to (C.2) we use the approximation $\tilde{\mathbf{z}}_{\text{Cl}}^{(n+1)} \simeq \check{\mathbf{z}}_{\text{Cl}}^{(n+1)} := [\check{\mathbf{z}}_{\text{Cl}}^{(n)}, \hat{z}_{\text{Cl}}^{(n)}(x_{n+1})]$. In that case we obtain an expression similar to (C.4), with $w_{n+1}^{(n+1)}$ replaced by $v_{n+1}^{(n+1)}$ from Eq. (B.3):

$$\frac{s_{\text{Cl}}^{(n+1)}(x_{n+1})^2}{s_{\text{Cl}}^{(n)}(x_{n+1})^2} \simeq \frac{(v_{n+1}^{(n+1)})^{-1}}{(v_{n+1}^{(n+1)})^{-1} + s_{\text{Cl}}^{(n)}(x_{n+1})^2}. \quad (\text{C.6})$$

The Hessian element $v_{n+1}^{(n+1)}$ is given by

$$v_{n+1}^{(n+1)} = \frac{\phi(\tilde{z}_{n+1}^{(n+1)})^2}{\Phi(\check{y}_{n+1}\tilde{z}_{n+1}^{(n+1)})^2} + \frac{\check{y}_{n+1}\tilde{z}_{n+1}^{(n+1)}\phi(\tilde{z}_{n+1}^{(n+1)})}{\Phi(\check{y}_{n+1}\tilde{z}_{n+1}^{(n+1)})},$$

which depends on the next-step signed response \check{y}_{n+1} . To develop an approximation in terms of step- n values, we once more replace $\tilde{z}_{n+1}^{(n+1)}$ with the current mean $\hat{z}_{\text{Cl}}^{(n)}(x_{n+1})$. Moreover, the next response \check{y}_{n+1} will take only two values, so $v_{n+1}^{(n+1)}$ will take on just two values v_{n+1}^\pm . Hence, we can compute the ‘‘expected value’’

$$\check{v}_{n+1} := v_{n+1}^+ p_+ + v_{n+1}^- p_-, \quad (\text{C.7})$$

$$\text{where } v_{n+1}^+ = \frac{\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2}{\Phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2} + \frac{\hat{z}_{\text{Cl}}^{(n)}(x_{n+1})\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}{\Phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}, \quad (\text{C.8})$$

$$\text{and } v_{n+1}^- = \frac{\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2}{\Phi(-\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))^2} - \frac{\hat{z}_{\text{Cl}}^{(n)}(x_{n+1})\phi(\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}{\Phi(-\hat{z}_{\text{Cl}}^{(n)}(x_{n+1}))}, \quad (\text{C.9})$$

with $p_+ := \mathbb{P}(Y(x_{n+1}) > 0 | \mathcal{A}_n) = \int_{\mathbb{R}} \Phi(z) p_{Z(x_{n+1})}(z | \mathcal{A}_n) dz = \Phi\left(\frac{\hat{z}^{(n)}(x_{n+1})}{\sqrt{1+s_C^{(n)}(x_{n+1})^2}}\right)$, and $p_- = 1 - p_+$.

The final formula for the look-ahead variance becomes

$$\hat{s}_{\text{Cl}}^{(n+1)}(x_{n+1})^2 := s_{\text{Cl}}^{(n)}(x_{n+1})^2 \cdot \frac{(\check{v}_{n+1})^{-1}}{(\check{v}_{n+1})^{-1} + s_{\text{Cl}}^{(n)}(x_{n+1})^2}.$$

TP: In terms of update formulas, TPs are in between GPs and t -GPs, with closed-form expressions available but depending on y_{n+1} . Specifically, the effect of adding a new observation (x_{n+1}, y_{n+1}) can be highlighted in closed form, since $f(x_*) | \mathbf{y}_{1:n}, y_{n+1} \sim \mathcal{T}(\nu + n + 1, \hat{f}_{\text{TP}}^{(n+1)}(x_*), s_{\text{TP}}^{(n+1)}(x_*))$, where

$$\hat{f}_{\text{TP}}^{(n+1)}(x_*) = \hat{f}_{\text{Gsn}}^{(n+1)}(x_*) \quad (\text{C.10})$$

$$s_{\text{TP}}^{(n+1)}(x_*)^2 = \frac{\nu + \beta^{(n+1)} - 2}{\nu + n - 1} s_{\text{Gsn}}^{(n+1)}(x_*)^2. \quad (\text{C.11})$$

The effect of y_{n+1} is inside $\beta^{(n+1)} = \mathbf{y}_{1:n+1}^\top [\mathbf{K}^{(n+1)}]^{-1} \mathbf{y}_{1:n+1} = \beta^{(n)} + s_{\text{Gsn}}^{(n)}(x_{n+1})^{-1} (h^{(n)}(x_{n+1})^2 + 2y_{n+1}h^{(n)}(x_{n+1}) + y_{n+1}^2)$ using the partition inverse equation, with $h^{(n)}(x) := -\mathbf{y}_{1:n}^\top [\mathbf{K}^{(n)}]^{-1} k(x) = -\hat{f}_{\text{Gsn}}(x)$. Since y_{n+1} is unknown beforehand, we use a plugin value $\check{\beta}^{(n+1)}$ for $\beta^{(n+1)}$, relying again on the tower property:

$$\begin{aligned} \check{\beta}^{(n+1)} &= \mathbb{E}[\beta^{(n+1)} | \mathbf{x}_{1:n}, \mathbf{y}_{1:n}] \\ &= \mathbb{E} \left[\mathbb{E}[\beta^{(n+1)} | \mathbf{x}_{1:n+1}, \mathbf{y}_{1:n+1}] | \mathbf{x}_{1:n}, \mathbf{y}_{1:n} \right] \\ &= \beta^{(n)} + s_{\text{Gsn}}^{(n)}(x_{n+1})^{-2} \left(h^{(n)}(x_{n+1})^2 + 2\hat{f}_{\text{Gsn}}^{(n)}(x_{n+1})h^{(n)}(x_{n+1}) + \hat{f}_{\text{Gsn}}^{(n)}(x_{n+1})^2 + \frac{\nu}{\nu-2} s_{\text{Gsn}}^{(n)}(x_{n+1})^2 \right) \\ &= \beta^{(n)} + \frac{\nu}{\nu-2}. \end{aligned} \tag{C.12}$$

D Gaussian Process with Monotonicity Constraint

Recall that since differentiation is a linear operator, the derivative of a GP f is another GP. Using \mathbf{f}' as a shorthand notation for the gradient ∇f at locations $\mathbf{x}_{1:n}$, we have

$$\mathbb{E}[\partial_{x_j} f(x_*) | \mathcal{A}] = \frac{\partial \mathbb{E}[f(x_*) | \mathcal{A}]}{\partial x_*^j} = \frac{\partial \hat{f}(x_*)}{\partial x_*^j}; \tag{D.1}$$

$$\text{Cov}(\partial_{x_j} f(x_*), f(x'_*) | \mathcal{A}) = K_{\mathbf{f}', \mathbf{f}}(x_*, x'_*) = \frac{\partial}{\partial x_*^j} K(x_*, x'_*) \tag{D.2}$$

$$\text{and } \text{Cov}(\partial_{x_j} f(x_*), \partial_{x'_j} f(x'_*) | \mathcal{A}) = K_{\mathbf{f}', \mathbf{f}'}(x_*, x'_*) = \frac{\partial^2}{\partial x_*^j \partial (x'_*)^{j'}} K(x_*, x'_*). \tag{D.3}$$

In addition to the data set $(\mathbf{x}_{1:n}, \mathbf{y}_{1:n})$, we now introduce virtual observations $(\mathbf{x}_v, \mathbf{y}_v)$ with the dummy responses $y_{v,i} \in \{-1, 1\} \times \{1, \dots, d\}$ set according to whether f is required to be decreasing ($y_{v,i} = (-1, j)$) or increasing ($y_{v,i} = (+1, j)$) with respect to the j th input dimension at input $x_{v,i}$. The key “trick” is to use a probit likelihood $p(y_{v,i} = (+1, j) | \mathbf{x}_{1:n}, \mathbf{x}_v) = \Phi(\frac{1}{\eta} \partial_{x_j} \hat{f}(x_{v,i}))$, where the small parameter η controls the strictness of the monotonicity constraint (Riihimäki and Vehtari, 2010). The probit function approaches the Heaviside step function when $\eta \rightarrow 0$ and forces the fitted $\partial_{x_j} \hat{f}(x_{v,i})$ (computed via (D.1)) to match during likelihood maximization the predetermined sign of $y_{v,i}$. An adaptive method to sequentially add the virtual inputs \mathbf{x}_v is suggested in Riihimäki and Vehtari (2010). Note that if there are multiple monotonic dimensions, then the same $x_{v,i}$ might be reused multiple times to satisfy the constraints on $\partial_{x_j} \hat{f}$ across different j -coordinates, leading to a replicated design. We also remark that monotonic metamodels are more expensive to run, since they require the use of virtual observations that increase the effective sample size to $(\mathbf{x}_{1:n}, \mathbf{x}_v)$ and hence require inversion of larger \mathbf{K} -matrices.

The joint prior for \mathbf{f} and its gradient \mathbf{f}' is given by

$$p_{\text{Mon}} \left(\begin{bmatrix} \mathbf{f} \\ \mathbf{f}' \end{bmatrix} | \mathbf{x}_{1:n}, \mathbf{x}_v \right) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{\text{joint}}), \tag{D.4}$$

$$\text{where } \mathbf{K}_{\text{joint}} = \begin{bmatrix} K_{\mathbf{f}, \mathbf{f}}(\mathbf{x}_{1:n}, \mathbf{x}_{1:n}) & K_{\mathbf{f}, \mathbf{f}'}(\mathbf{x}_{1:n}, \mathbf{x}_v) \\ K_{\mathbf{f}', \mathbf{f}}(\mathbf{x}_v, \mathbf{x}_{1:n}) & K_{\mathbf{f}', \mathbf{f}'}(\mathbf{x}_v, \mathbf{x}_v) \end{bmatrix}.$$

Using Bayes rule, the joint posterior is then

$$p_{\text{Mon}}(\mathbf{f}, \mathbf{f}' | \mathbf{x}_{1:n}, \mathbf{y}_{1:n}, \mathbf{x}_v, \mathbf{y}_v) = \frac{p_{\text{Mon}}(\mathbf{f}, \mathbf{f}' | \mathbf{x}_{1:n}, \mathbf{x}_v) p(\mathbf{y}_{1:n} | \mathbf{f}) \prod_i \Phi \left(y_{v,i} \partial_{x_j} f(x_{v,i}) \frac{1}{\eta} \right)}{p(\mathbf{y}_{1:n}, \mathbf{y}_v | \mathbf{x}_{1:n}, \mathbf{x}_v)}. \tag{D.5}$$

Like for the classification GP to handle the non-Gaussian terms $p(y_{v,i} | \partial_{x_j} f(x_{v,i}))$ we approximate them with a local Gaussian likelihood

$$p(\mathbf{y}_v | \mathbf{f}') \approx q(\mathbf{y}_v | \mathbf{f}') = \mathcal{N}(\tilde{\boldsymbol{\mu}}_{\text{Mon}}, \tilde{\boldsymbol{\Sigma}}_{\text{Mon}}). \quad (\text{D.6})$$

We use the Expectation Propagation (EP) algorithm Minka (2001) to determine the vector of local means $\tilde{\boldsymbol{\mu}}_{\text{Mon}}^i$, and the diagonal EP variance matrix $\tilde{\boldsymbol{\Sigma}}_{\text{Mon}}$, with local variances $(\tilde{\sigma}_{\text{Mon}}^i)^2$. Details about the computation can be found in Riihimäki and Vehtari (2010). The approximate posterior to (D.5) is a product of Gaussian distributions and is simplified to

$$p_{\text{Mon}}(\mathbf{f}, \mathbf{f}' | \mathbf{x}_{1:n}, \mathbf{x}_v, \mathbf{y}_{1:n}, \mathbf{y}_v) \approx q_{\text{Mon}}(\mathbf{f}, \mathbf{f}' | \mathbf{x}_{1:n}, \mathbf{x}_v, \mathbf{y}_{1:n}, \mathbf{y}_v) = \mathcal{N}(\boldsymbol{\mu}_{\text{joint}}, \boldsymbol{\Sigma}_{\text{joint}}).$$

The covariance matrix is $\boldsymbol{\Sigma}_{\text{joint}}^{-1} = \mathbf{K}_{\text{joint}}^{-1} + \tilde{\boldsymbol{\Sigma}}_{\text{joint}}^{-1}$, with $\tilde{\boldsymbol{\Sigma}}_{\text{joint}} = \begin{bmatrix} \sigma^2 \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \tilde{\boldsymbol{\Sigma}}_{\text{Mon}} \end{bmatrix}$, and the posterior mean is

$$\boldsymbol{\mu}_{\text{joint}} = \boldsymbol{\Sigma}_{\text{joint}} \tilde{\boldsymbol{\Sigma}}_{\text{joint}}^{-1} \tilde{\boldsymbol{\mu}}_{\text{joint}}, \text{ with } \tilde{\boldsymbol{\mu}}_{\text{joint}} = \begin{bmatrix} \mathbf{y}_{1:n} \\ \tilde{\boldsymbol{\mu}}_{\text{Mon}} \end{bmatrix}.$$

The posterior mean $\hat{f}_{\text{Mon}}(x_*)$ and posterior covariance $v_{\text{Mon}}(x_*, x'_*)$ for the M-GP metamodel are

$$\hat{f}_{\text{Mon}}^{(n)}(x_*) = [k(x_*), K_{\mathbf{f}, \mathbf{f}'}(x_*, \mathbf{x}_v)] \mathbf{K}_{\text{joint}}^{-1} \boldsymbol{\mu}_{\text{joint}}, \quad (\text{D.7})$$

$$v_{\text{Mon}}^{(n)}(x_*, x'_*) = K_{\mathbf{f}, \mathbf{f}}(x_*, x'_*) - [k(x_*), K_{\mathbf{f}, \mathbf{f}'}(x_*, \mathbf{x}_v)] \times [\mathbf{K}_{\text{joint}} + \tilde{\boldsymbol{\Sigma}}_{\text{joint}}]^{-1} \begin{bmatrix} k(x_*) \\ K_{\mathbf{f}', \mathbf{f}}(\mathbf{x}_v, x_*) \end{bmatrix}, \quad (\text{D.8})$$

analogously to the standard GP prediction equations (2.1) and (2.2).

In M-GP, replacing f with z and again applying the EP algorithm, we reach similar expressions for posterior mean/variance as in (D.7) and (D.8).

Similar to the t -GP and CI-GP, look-ahead variance is intractable for the monotonic GP, since the EP mean $\tilde{\boldsymbol{\mu}}_{\text{Mon}}$ and variance $\tilde{\boldsymbol{\Sigma}}_{\text{Mon}}$ are changing as the designs are augmented. Rewriting (D.8), we obtain

$$\tilde{v}_{\text{Mon}}^{(n)}(x_*, x'_*) = K_{\mathbf{f}, \mathbf{f}}(x_*, x'_*) - [K_{\mathbf{f}, \mathbf{f}'}^{(n)}(x_*, \mathbf{x}_v), k(x_*)] [\tilde{\mathbf{K}}_{\text{joint}}^{(n)} + \tilde{\boldsymbol{\Sigma}}_{\text{joint}}^{(n)}]^{-1} \begin{bmatrix} K_{\mathbf{f}', \mathbf{f}}^{(n)}(\mathbf{x}_v, x_*) \\ k(x_*) \end{bmatrix}, \quad (\text{D.9})$$

where $\tilde{\mathbf{K}}_{\text{joint}}^{(n)} = \begin{bmatrix} K_{\mathbf{f}', \mathbf{f}'}^{(n)}(\mathbf{x}_v, \mathbf{x}_v) & K_{\mathbf{f}', \mathbf{f}}^{(n)}(\mathbf{x}_v, \mathbf{x}_{1:n}) \\ K_{\mathbf{f}, \mathbf{f}'}^{(n)}(\mathbf{x}_{1:n}, \mathbf{x}_v) & K_{\mathbf{f}, \mathbf{f}}^{(n)}(\mathbf{x}_{1:n}, \mathbf{x}_{1:n}) \end{bmatrix}$ and $\tilde{\boldsymbol{\Sigma}}_{\text{joint}}^{(n)} = \begin{bmatrix} \tilde{\boldsymbol{\Sigma}}_{\text{Mon}}^{(n)} & \mathbf{0} \\ \mathbf{0} & \sigma^2 \mathbf{I}_{n \times n} \end{bmatrix}$.

$K_{\mathbf{f}', \mathbf{f}'}^{(n)}(\mathbf{x}_v, \mathbf{x}_v)$ is the step- n covariance matrix for the gradient of virtual observations, and $\tilde{\boldsymbol{\Sigma}}_{\text{Mon}}^{(n)}$ is the approximate covariance matrix for $p_{\text{Mon}}(\mathbf{y}_v | \mathbf{f}')$. When calculating the one-step-ahead variance for monotonic GP, we freeze the virtual observations and their gradient, which in consequence freezes the $K_{\mathbf{f}', \mathbf{f}'}^{(n)}(\mathbf{x}_v, \mathbf{x}_v)$, $K_{\mathbf{f}', \mathbf{f}}^{(n)}(\mathbf{x}_v, \mathbf{x}_{1:n})$, $K_{\mathbf{f}, \mathbf{f}'}^{(n)}(\mathbf{x}_{1:n}, \mathbf{x}_v)$, and $\tilde{\boldsymbol{\Sigma}}_{\text{Mon}}^{(n)}$ matrices. Therefore, the virtual observations are treated as fixed inputs. Then, as a new observation is added, only the last row and column of the covariance matrix are updated, while the other parts remain unchanged. This approach transforms computing the look-ahead standard deviation $s_{\text{Mon}}^{(n+1)}$ into the classical Gaussian observation GP as in (4.5). Therefore, following exactly the same procedures discussed in Section 4, similar to equations (4.4) and (4.5), we obtain the local updated variance $s_{\text{Mon}}^{(n+1)}(x_{n+1})^2$ at x_{n+1} , and the step-ahead variance $s_{\text{Mon}}^{(n+1)}(x_*)^2$ at any input x_* :

$$\frac{s_{\text{Mon}}^{(n+1)}(x_{n+1})^2}{s_{\text{Mon}}^{(n)}(x_{n+1})^2} = \frac{\tau^2}{\tau^2 + s_{\text{Mon}}^{(n)}(x_{n+1})^2}, \quad (\text{D.10})$$

$$s_{\text{Mon}}^{(n+1)}(x_*)^2 = s_{\text{Mon}}^{(n)}(x_*)^2 - \frac{v_{\text{Mon}}^{(n)}(x_*, x_{n+1})^2}{\tau^2 + s_{\text{Mon}}^{(n)}(x_{n+1})^2}. \quad (\text{D.11})$$