

# UC Santa Barbara

## UC Santa Barbara Previously Published Works

### Title

Probabilistic bisection with spatial metamodels

### Permalink

<https://escholarship.org/uc/item/2t53w5pz>

### Journal

European Journal of Operational Research, 286(2)

### ISSN

0377-2217

### Authors

Rodriguez, Sergio  
Ludkovski, Michael

### Publication Date

2020-10-01

### DOI

10.1016/j.ejor.2020.03.049

Peer reviewed

# Probabilistic Bisection with Spatial Metamodels

Sergio Rodriguez and Michael Ludkovski

Department of Statistics and Applied Probability,  
University of California, Santa Barbara, Santa Barbara, CA, USA 93106  
{srodriguez,ludkovski}@pstat.ucsb.edu

## Abstract

Probabilistic Bisection Algorithm performs root finding based on knowledge acquired from noisy oracle responses. We consider the generalized PBA setting (G-PBA) where the statistical distribution of the oracle is *unknown* and location-dependent, so that model inference and Bayesian knowledge updating must be performed simultaneously. To this end, we propose to leverage the spatial structure of a typical oracle by constructing a statistical surrogate for the underlying logistic regression step. We investigate several non-parametric surrogates, including Binomial Gaussian Processes (B-GP), Polynomial, Kernel, and Spline Logistic Regression. In parallel, we develop sampling policies that adaptively balance learning the oracle distribution and learning the root. One of our proposals mimics active learning with B-GPs and provides a novel look-ahead predictive variance formula. The resulting gains of our Spatial PBA algorithm relative to earlier G-PBA models are illustrated with synthetic examples and a challenging stochastic root finding problem from Bermudan option pricing.

Keywords: Stochastic Root-Finding, Simulation metamodeling, Uncertainty Quantification.

## 1 Introduction

The Probabilistic Bisection Algorithm (PBA) is a numeric estimation procedure for learning an unknown parameter  $x^*$  (defined on a bounded search space, without loss of generality  $[0, 1]$ ) based on the information provided by noisy responses observed independently at sampling/querying sites  $x_{1:n} := (x_1, \dots, x_n)$ . In the context of the Stochastic Root Finding Problem (SRFP) (Pasupathy and Kim, 2011; Waeber et al., 2011), the PBA can be used to learn the root,  $x^* := h^{-1}(0)$ , of a noisily observed real-valued function  $h : [0, 1] \rightarrow \mathbb{R}$ . Specifically we consider an oracle of the form

$$Y(x_n) := \text{sign } Z(x_n), \quad (1)$$

where the structural form of the (random) responses  $Z(x_n)$  in (1) is given by

$$Z(x_n) := h(x_n) + \epsilon(x_n). \quad (2)$$

The noise component  $\epsilon(x_n)$  in (2) is assumed to be a *symmetric* heteroscedastic (i.e., input-dependent) random term with mean  $\mathbb{E}[\epsilon(x_n)] = 0$  and variance  $\text{Var}(\epsilon(x_n)) := \sigma^2(x)$ , with independent realizations across different oracle calls.

The PBA leverages the classical *bisection search* strategy in a noise-free setting: repeatedly halve the search region and then select a subinterval in which a root must lie for further processing. The stochastic PBA accounts for the noise in the oracle responses by considering  $x^*$  as the realization of an absolutely continuous random variable  $X^* \sim g_0$  with prior density  $g_0$  supported on  $[0, 1]$ . The PBA then works with the sign of the noisy function evaluations (1), which provide information as to whether  $x^*$  lies to the left or to the right of a given  $x_n$ , in order to subsequently update a posterior density for  $X^*$ ,

$$g_n(X^*) := p(X^* | Y_{1:n}, x_{1:n}). \quad (3)$$

Thus,  $g_n$  is the pdf of the root location  $X^*$  conditional on the history  $Y_{1:n} := (Y_1(x_1), \dots, Y_n(x_n))$  of oracle responses, the sampling locations  $x_{1:n}$  and the prior  $g_0$ . The posterior (3) then serves for the twin purposes of guiding the election of the next sampling location  $x_{n+1}$  at which to query (1), as well as to provide a point estimator  $\hat{x}_n$  for  $X^*$  (e.g., the posterior median or mean of  $g_n(\cdot)$ ).

Due to the noise term  $\epsilon(x_n)$  in the simulation outputs  $Z(x_n)$  in (1), the responses  $Y(x_n) = \text{sign } Z(x_n)$  translate into potentially inaccurate oracle directions. To account for such “mistakes”, the PBA considers the probability of *correct sign*,

$$p(x_n) := \mathbb{P}(Y(x_n) = \text{sign}\{x_n - x^*\}), \quad (4)$$

henceforth referred to as oracle specificity or *accuracy*, which is then used to update knowledge about  $X^*$  by re-weighting the current  $g_n$  proportionally to  $p(x_n)$ . Waeber et al. (2013) provided an explicit recursive updating formula under the restrictive condition that the oracle accuracy is a known constant  $p(x) = p^* > 1/2$  for all  $x \in (0, 1)$ . This assumption of spatial oracle stationarity would tend to be met in applications where the transition between regions in  $h$  is abrupt. As an example, if a city’s water supply were contaminated with a dangerous chemical we would want to localize the extent of contamination as quickly as possible, and if the chemical did not dissolve well in water but instead tended to stay concentrated, we would face a situation with such abrupt transition between contaminated and uncontaminated water (Powell and Ryzhov, 2012).

However, in the more general and practical case, including the SRFP in (2),  $p(x)$  is unknown and location-dependent and hence must be itself *estimated*. The Generalized PBA (G-PBA) that we developed in Rodriguez and Ludkovski (2017) extends the classical PBA by using the observed data to construct a point estimate,  $\hat{p}(x)$ , for  $p(x)$ , as well as to learn the root location  $X^*$  in parallel. The proposed estimators  $\hat{p}(x_{n+1})$  under the aforementioned G-PBA paradigm were constructed *locally* at  $x_{n+1}$  (i.e., without using information from previous locations  $x_{1:n}$ ). As such, they were robust to arbitrary specification of  $p(\cdot)$  and made minimal assumptions about the oracle.

**Surrogate modeling.** In this article we construct a *spatial* G-PBA by modeling the entire oracle accuracy  $x \mapsto p(x)$  using a *surrogate*.

The surrogate relies on two main premises: (i) Due to symmetrical noise distribution of the functional responses (2), the oracle accuracy (4) can be re-formulated as  $p(x) = \max\{\theta(x), 1 - \theta(x)\}$ , where

$$\theta(x) := \mathbb{E} [1_{\{Z(x) > 0\}}] \quad (5)$$

is the probability of observing a *positive* oracle response. Thus, inference on  $p(x)$  can be performed by inferring  $\theta(x)$  first and then plugging in a spatial-based estimate  $\hat{\theta}_n(x)$  into  $p(x) = \max\{\hat{\theta}_n(x), 1 - \hat{\theta}_n(x)\}$ , and (ii) the smoothness of the map  $x \mapsto \theta(x)$ , implies that  $p(x)$  and  $p(x')$  should be similar when  $x$  and  $x'$  are deemed close to each other.

The spatial structure is natural in the root-finding context and provides two key benefits. On the one hand, it improves estimation of a given  $p(x_n)$  through leveraging the knowledge acquired at previous sampling locations  $x_{1:n-1}$ . On the other hand, it enables better sampling strategies by furnishing a prediction  $\hat{p}(x)$  at arbitrary, unsampled sites  $x$ . In contrast, in G-PBA,  $\hat{p}(x_{n+1})$  was only available *a posteriori* after sampling at  $x_{n+1}$ .

The resulting Spatial G-PBA strategy blends the root-centric framework of PBA and the function-centric paradigm of response surface modeling (RSM). Indeed, a further alternative for solving the SRFP would be to learn the entire  $\theta(\cdot)$  and then take  $\hat{x} = \hat{\theta}^{-1}(0.5)$  since  $h(x^*) = 0 \Leftrightarrow \theta(x^*) = 0.5$ . Thus, stochastic root-finding can be recast as a (localized) learning task, namely contour-finding for  $\theta(\cdot)$  at the level 0.5. Strategies similar to Bayesian optimization (Jones et al., 1998; Chevalier et al., 2014) can then be employed to efficiently target this objective during sequential design. Nevertheless, several challenges are encountered with such an approach that are circumvented in PBA. First, a major feature of PBA is full uncertainty quantification: the algorithm provides not only the point estimate  $\hat{x}$  but also the entire posterior distribution  $f_n$  of  $X^*$  conditional on the data. Typical RSM models return only point estimates (or pointwise credible intervals) of  $\theta(x)$ ; the latter are difficult to “invert” into uncertainty about  $\theta^{-1}(0.5)$  (Azzimonti et al., 2016). Second, existing experimental design approaches for contour-finding are developed only for simple models (e.g. with zero or constant observation noise), and their performance in a complex stochastic setting like ours is poorly understood. In contrast, the PBA framework explicitly targets the goal of reducing uncertainty about  $X^*$ . PBA moreover exploits the structural knowledge of a *unique* root to speed up estimation, an option that is not available in contour-finding. Third, contour-finding usually assumes continuous response, and nontrivial modifications (essentially “logistic” contour-finding) are necessary to handle binary  $Y_n \in \{-1, 1\}$ . In contrast, PBA is intrinsically designed for binomial responses.

Given the above discussion, we construct a hybrid algorithm that borrows the best of both worlds. We exploit the smoothness of  $h$  that implies spatial dependence in  $\theta(\cdot)$  and hence accelerates learning the oracle. At the same time, we employ the paradigm of PBA to construct the knowledge state  $f_n$  (a pseudo-posterior of  $X^*$ ) that is the primary driver of sampling decisions and uncertainty quantification. For the RSM component, we rely on two key concepts. First, we investigate non-parametric architectures that have the flexibility to consistently learn the entire response  $x \mapsto \theta(x)$  and to handle non-uniform simulation designs. The latter point is key as we wish to organically refine the surrogate in regions where more

inputs are placed (namely close to the root), but at the same time give a good global fit. To handle the binary responses (1) we employ *logistic regression* which represents the probability of observing a positive response  $\theta(x) = \mathbb{E}[1_{Z_n(x)>0}]$  via a *latent* process  $\varphi(x) := \text{logit}(\theta(x))$ . Other link functions can also be used but as we show in the sections below, the canonical Bernoulli logit link is best suited for our needs. For capturing the spatial surrogate  $\varphi$ , we consider Gaussian Process (GP) models, as well as spline, kernel and polynomial logistic regressions. Second, we apply *batched sampling* that significantly lowers the computational overhead of surrogate construction and improves the learning of  $\theta(\cdot)$ . Replicated experimental designs allow to blend the local inference of  $\theta(x_n)$  with the global fitting of the surrogate. They also offer a new aspect of sequential design, namely adaptive replication, linking to the active learning literature in Bayesian optimization.

**Summary of Contributions and Related Literature.** Our contributions can be traced along two directions. First, the developed G-PBA algorithm extends existing probabilistic bisection schemes in Jedynek et al. (2012); Waeber et al. (2011, 2013); Waeber (2013); Frazier et al. (2016), in particular making them much more efficient even when the oracle distribution is a priori unknown. Thus, we contribute to the stochastic root-finding toolkit. Second, our work has independent interest in terms of applications of binomial GP (B-GP) surrogates. To this end, we provide an original result for the look-ahead variance formula of a binomial GP, that to our knowledge is not available in existing literature. This provides a new application of B-GPs in the context of active learning, linking to related work in Kapoor et al. (2007); Tesch et al. (2013); Wang et al. (2016).

In the extensive numerical section, we demonstrate that by introducing a spatial surrogate we are able to improve PBA’s accuracy in the root estimation. By using three different synthetic examples, we show that absolute residuals decay faster using spatial surrogates than their corresponding (local) G-PBA methods. Additionally, the posterior uncertainty of the root estimate, as measured by the width of the posterior credible interval (CI), is reduced and, most importantly, the probability coverage (i.e., the proportion of macro-runs where the CI contains the actual root value) drastically increases with respect to their local counterparts.

The rest of the paper is organized as follows. In Section 2 we describe the model methodology used to provide a spatial estimate for  $p(\cdot)$ . Section 2.2 then describes an adaptive batching/replication scheme in order to determine the number of replicates  $a_{n+1}$  given an estimated surrogate model. Section 3 develops the surrogate-based sampling schemes for the SRFP. In Sections 4 and 5 we illustrate the developed Spatial G-PBA with several synthetic examples and a challenging real-world application coming from an Optimal Stopping problem.

## 2 Spatial Modeling of the Oracle

PBA works in the sequential setting, adaptively picking query sites  $x_{n+1}$  given information from previous queries. The latter is summarized via a knowledge state  $f_n$  which captures the Bayesian formulation of the SRFP, translating the task of learning the root  $X^*$  into quantifying the corresponding posterior uncertainty. At each iteration  $n = 1, \dots$ , the oracle is

called  $a_n \geq 1$  times at a fixed sampling location  $x_n$  with the responses  $(Z_1(x_n), \dots, Z_{a_n}(x_n))$  aggregated via the total number of positive signs observed at  $x_n$ :

$$B_n(x_n) := \sum_{j=1}^{a_n} 1_{\{Z_j(x_n) > 0\}}. \quad (6)$$

The overall information set by round  $n$  is therefore  $\mathcal{D}_n := (B_{1:n}, a_{1:n}, x_{1:n})$ . We shall distinguish between the macro counter  $n$  that keeps track of PBA iterations, i.e. the number of distinct sites  $x_{1:n}$ , and the wall clock  $T_n := \sum_{i=1}^n a_i$  that counts total number of function evaluation and hence the overall computational expense. Occasionally, we abuse the notation, switching between writing  $f_n$  and  $f_{T_n}$ .

Given the current knowledge state  $f_n$  and a total simulation budget of  $T > 0$  wall-clock iterations, the fundamental G-PBA loop is:

```

Initialize  $T_0$  and  $\hat{p}_0$ ;
while  $T_n < T$  do
    Choose  $x_{n+1}$  based on  $f_n$  and  $\hat{p}_n(\cdot)$ ;
    Call the oracle  $a_{n+1}$  times at  $x_{n+1}$  and record  $B_{n+1}$  as in (6);
    Use  $(a_{n+1}, B_{n+1})$  to update  $f_{n+1} \leftarrow \Psi(f_n, B_{n+1}; \hat{p}(x_{n+1}), a_{n+1})$  and re-fit  $\hat{p}_{n+1}(\cdot)$ ;
    Update wall-clock time:  $T_{n+1} \leftarrow T_n + a_{n+1}$  and increment  $n \leftarrow n + 1$ ;
end
return Knowledge state  $f_n \simeq g_n$  and estimator for the root location  $\hat{x}_n$ .

```

**Algorithm 1:** G-PBA iterations

In the general case of unknown and varying oracle specificity, the key ingredients of Algorithm 1 are:

(GPBA-I) statistical learning sub-routine for  $\hat{p}(\cdot)$ .

(GPBA-II) the mechanism to update knowledge states  $\Psi : f_n \rightarrow f_{n+1}$ ;

(GPBA-III) the sampling rule  $\eta$  for selecting  $x_{n+1} = \eta(f_n; \hat{p}_n)$  given  $f_n$  and  $\hat{p}_n$ .

**Learning sub-routine for  $\hat{p}(\cdot)$ .** For estimating the oracle specificity, G-PBA relies on the aggregated number  $B_n$  of positive signs observed at  $x_n$  across  $a_n$  oracle queries. Replicates decouple the problems of learning  $X^*$  and of learning  $p(\cdot)$ ; they also boost the signal-to-noise ratio which allows *faster* convergence at the macro-level. The original G-PBA did this locally, returning an estimate  $\hat{p}(x_n)$  depending solely on  $(x_n, a_n, B_n)$ . In this paper we extend (GPBA-I) by introducing a surrogate model  $x \mapsto \varphi(x)$  on (5) which is built upon the *history* of binomial responses  $B_{1:n} := (B_1(x_1), \dots, B_n(x_n))$  observed at all queried  $x_{1:n}$ . In particular, we have that  $B_n(x) \sim \text{Bin}(a_n, \theta(x))$  is a binomial random variable which is statistically *sufficient* and unbiased for  $\theta(x)$ . To learn  $\theta(\cdot)$  we therefore regress  $B_{1:n}$  against the locations  $x_{1:n}$ , linking each  $x_i$  to  $\theta(x_i)$  via the canonical Bernoulli link function:

$$\log \left( \frac{\theta(x_i)}{1 - \theta(x_i)} \right) = \varphi(x_i), \quad i = 1, \dots, n. \quad (7)$$

We consider two families for  $\varphi(\cdot)$ : (A) Gaussian random field approach (Williams and Rasmussen, 2006) that takes  $\varphi$  as a latent Gaussian process (GP) and outputs the posterior distribution  $p(\varphi_*|\mathcal{D}_n)$ ; (B) a linear additive model that assumes that  $\varphi$  is an element of a linear space  $\mathcal{H}$  spanned by a collection of basis functions, i.e.,  $\varphi(x) = \sum_{j=1}^p \beta_j \phi_j(x)$ , with the coefficients  $\beta := (\beta_1, \dots, \beta_p)$  fitted, for example, by penalized MLE.

Given the fitted surrogate  $\hat{\varphi}_n$ , the estimate for  $p(\cdot)$  is a plug-in estimate of the form:

$$\hat{p}_n(x) := \max\{\hat{\theta}_n(x), 1 - \hat{\theta}_n(x)\}; \text{ where } \hat{\theta}_n(x) \equiv \Theta(\hat{\varphi}_n(x)) := [1 + e^{-\hat{\varphi}_n(x)}]^{-1}. \quad (8)$$

### Updating knowledge states.

The knowledge state  $f_n$  is intended to capture all available information about  $X^*$  given  $\mathcal{D}_n$ . Since the true Bayesian posterior  $g_n$  is not attainable due to unknown  $p(\cdot)$ , we notationally distinguish between the *approximate* knowledge state  $f_n$  and the true  $g_n$  (3). For assimilating information, we mimic the exact Bayesian updating from Waeber et al. (2011) and use the batched knowledge state transition introduced in Rodriguez and Ludkovski (2017). Thus we take  $f_{n+1} = \Psi(f_n, x_{n+1}, B_{n+1}; \hat{p}_{n+1}, a_{n+1})$  with

$$\Psi(f_n, x_{n+1}, B_{n+1}; p, a)(u) \propto \begin{cases} [p(x_{n+1})^{B_{n+1}}(1 - p(x_{n+1}))^{a - B_{n+1}}] f_n(u), & x_{n+1} < u \\ [(1 - p(x_{n+1}))^{B_{n+1}} p(x_{n+1})^{a - B_{n+1}}] f_n(u), & x_{n+1} \geq u. \end{cases} \quad (9)$$

Note that we replace the unknown  $p(x_{n+1})$  with the surrogate-based  $\hat{p}_{n+1}(x_{n+1})$ . Over multiple rounds, this implies that  $f_n$  depends on the historical estimates  $\hat{\varphi}_{1:n}(x_{1:n})$  introducing a complex path-dependency between the latest knowledge state and the past surrogates of  $p(\cdot)$ .

**Sampling strategies.** A sampling policy  $\eta$  is a rule which maps knowledge states to actions, namely sampling decisions. The sampling decision to be made at step  $(n + 1)$  concerns the new query site  $x_{n+1}$  and the respective number of replicates  $a_{n+1}$ . We consider two complementary ideas: (i) first select  $a_{n+1}$  and then  $x_{n+1}$ ; (ii) choose  $x_{n+1}$  and then determine the respective  $a_{n+1}$ .

Approach (i) utilizes fixed replication amount  $a \geq 1$  and selects the new  $x_{n+1}$  using an information-theoretic criterion. In analogy to the Information Directed Sampling (IDS) policy used in the G-PBA context (Rodriguez and Ludkovski, 2017), we consider a criterion based on the *batched* expected Kullback-Leibler (KL) divergence  $\mathbb{E}[D(f_{n+1}; f_n)]$  between  $f_n$  and the updated knowledge state  $f_{n+1} = \Psi(f_n, x, B_{n+1}; \hat{p}_{n+1}, a)$ ,

$$\mathcal{I}(x, f_n; \hat{p}_n(x), a) := \mathbb{E}_{\hat{p}}^{B(x)} \left[ \int_0^1 \log_2 \left( \frac{f_n(u|B(x), a)}{f_n(u)} \right) f_n(u) du \right]; \quad (10)$$

where the expectation is taken with respect to the random variable  $B(x) \sim \text{Bin}(a, \hat{p}_n(x))$  and  $\hat{p}_n(\cdot)$  is recovered using (8). Given the acquisition function (10), the next sampling location is its greedy maximizer

$$x_{n+1}^{\text{IDS}} := \arg \sup_{x \in (0,1)} \mathcal{I}(x, f_n; \hat{p}_n(x), a).$$

The IDS rule was shown to be optimal for the base case of known and constant  $p(x)$  and  $a = 1$  (Jedynak et al., 2012). In that case it is known to correspond to selecting  $x_{n+1}$  which

maximizes the conditional mutual information between the oracle response  $Y_{n+1}(x_{n+1})$  and  $X^*$  given  $f_n$ . Approach (ii), dubbed Randomized Quantile Sampling (RQS), selects locations using the knowledge state  $f_n$  as a proposal density, i.e.,  $x_{n+1}^{\text{RQS}} \sim f_n(\cdot)$  and then adaptively picks  $a_{n+1}$ . The RQS strategy resembles Thompson sampling (Russo and Van Roy, 2016) and was shown to be competitive with IDS (and frequently slightly better) in the earlier G-PBA context. Conditional on  $x_{n+1}$ ,  $a_{n+1}$  is then picked to control the surrogate accuracy at  $x_{n+1}$  in order to ensure the right amount of learning.

**Estimating the root  $X^*$ .** The final ingredient is the rule  $\hat{x}_n$  to construct a point estimate of the root  $x^*$  based on  $f_n$ . In analogy to the classical PBA setting (Waeber et al., 2013), we utilize the posterior median which we find is generally more robust than say the mean, as  $f_n$  is often skewed or multi-modal,

$$\hat{x}_n := \text{median}(f_n). \quad (11)$$

## 2.1 Binomial Gaussian Process Regression

GPs can conveniently be used to specify prior distributions for Bayesian inference in the regression context and are widely adopted for sequential design tasks. In G-PBA they facilitate managing the sample budget for calling (2) by quantifying the predictive uncertainty at the next sampling site  $x_{n+1}$  in terms of the number of replicates  $a_{n+1}$  (Kamiński, 2015; Binois et al., 2018b). The related Binomial Gaussian processes (B-GPs) (also known as GP classification and originally introduced in Williams and Barber (1998)) arise naturally in the context of *latent variable* regression for (5). In this case, the  $\varphi$  is seen as a realization of a random process whose finite dimensional distribution follows a Multivariate Normal (MVN) distribution and whose spatial dependency is described by a (stationary) covariance function.

While for plain regression with Gaussian noise inference can be done in closed form, since for a given election of covariance kernel the posterior corresponds also to a GP (Williams and Rasmussen, 2006), this is no longer the case for B-GPs. Indeed, since the binomial data likelihood is not conjugate to the Gaussian prior, exact inference is analytically intractable and therefore approximations to the predictive posterior must be conducted. One route summarized in Nickisch and Rasmussen (2008) is based on approximating the non-Gaussian posterior with a tractable Gaussian distribution. Some of the most common instances of such schemes are the Laplace Approximation (LA) (Williams and Barber, 1998) and Expectation Propagation (EP) (Minka, 2001).

Let us assume that the surrogate  $\varphi$  in (7) is drawn from a GP prior,  $\varphi \sim \text{GP}(0, \kappa_{\vartheta}(\cdot, \cdot))$ , characterized by a *covariance kernel* function  $\kappa_{\vartheta}(\cdot, \cdot)$  and parameterized by a vector of *hyperparameters*  $\vartheta \equiv (\tau^2, l)$ . One of the most commonly used kernels is the Matérn-5/2 family,

$$\kappa_{\vartheta}(x_i, x_j) := \tau^2 \left[ 1 + \sqrt{5}r/l + 5r^2/(3l^2) \right] e^{-\sqrt{5}r/l} \quad r := |x_i - x_j|; \quad (12)$$

where  $\tau^2 \geq 0$  is the intrinsic GP variance, and  $l > 0$  is the length-scale, which governs how fast the correlation decreases as the distance  $|x_i - x_j|$  between inputs increases.



**Binomial GPs as latent variable models.** For fixed hyper-parameter  $\vartheta$ , the joint distribution of the vector  $\varphi_{1:n} := (\varphi_1(x_1), \dots, \varphi_n(x_n))$  is a MVN

$$\varphi_{1:n} \sim \mathbf{N}(\mathbf{0}, \mathbf{K}_n), \quad (13)$$

where  $\mathbb{E}[\varphi_{1:n}|x_{1:n}] = \mathbf{0}$  is the mean vector and  $\mathbf{K}_n \equiv Cov(\varphi_{1:n}|x_{1:n})$  is the covariance matrix with entries  $\kappa_{\vartheta}(x_i, x_j) = Cov(\varphi_i, \varphi_j|x_{i,j})$ . Inference of  $\theta(\cdot)$  in (5) is conducted in two stages. First, we compute the posterior distribution of the vector  $\varphi_{1:n}$  given the training data  $\mathcal{D}_n := (B_{1:n}, a_{1:n})$ , consisting of the history of binomial  $B_i$  responses and number of queries  $a_i$  at each location  $x_i$ ,

$$p(\varphi_{1:n}|\mathcal{D}_n) \propto p(B_{1:n}|\varphi_{1:n}, a_{1:n})p(\varphi_{1:n}); \quad (14)$$

which is proportional to the binomial data likelihood  $p(B_{1:n}|\varphi_{1:n}, a_{1:n})$  times the MVN prior  $p(\varphi_{1:n})$  given by (13). Second, the posterior predictive distribution  $\varphi_* \equiv \varphi_*(x)$  at a location  $x \in (0, 1)$  is

$$p(\varphi_*|\mathcal{D}_n) := \int p(\varphi_*, \tilde{\varphi}_{1:n}|\mathcal{D}_n, x)d\tilde{\varphi}_{1:n}, \quad (15)$$

which is calculated by marginalizing the distribution of  $\varphi_*$  over the joint posterior distribution of  $(\varphi_{1:n}, \varphi_*)$  given by (14). Finally, the predicted  $\hat{\theta}^{GP}(x)$  is produced by averaging the inverse link function with respect to (15); i.e.,  $\hat{\theta}_n^{GP}(x) := \int (1 + e^{-\varphi_*})^{-1} \cdot p(\varphi_*|\mathcal{D}_n)d\varphi_*$ .

*Remark 1.* Following the classical inference paradigm for binomial regression we assume that  $\theta(\varphi(x_i))$  is related to the random variable  $\varphi(x_i)$  via the canonical *logistic* link function (7). Although other link functions can be entertained (such as the probit link), we use the logistic one since this link is used to obtain closed-form expressions for adaptive replication (see Lemma 2.1 in Section 2.2).

The main challenge in computing the joint posterior (14) is that the MVN prior over  $\varphi_{1:n}$  does not correspond to a conjugate prior for the binomial likelihood, so either analytic approximations of integrals or solutions based on MCMC sampling are required. A commonly used method is to approximate the non-Gaussian posterior  $p(\varphi_{1:n}|\mathcal{D}_n)$  with a Gaussian one via Laplace Approximation.

**Laplace Approximation.** The Laplace method is constructed from the second order Taylor expansion of the *score function*,  $\mathcal{L}(\varphi_{1:n}) := \log p(\varphi_{1:n}|\mathcal{D}_n)$ , around its mode:

$$\hat{\varphi}_n = \arg \max_{\varphi_n} p(\varphi_n|\mathcal{D}_n).$$

In Appendix A we show that this method yields a MVN approximation:

$$p(\cdot|\mathcal{D}_n) \simeq q(\cdot|\mathcal{D}_n, \hat{\varphi}_n) = \mathbf{N}(\cdot; \hat{\varphi}_n, (\mathbf{K}_n^{-1} + \hat{\mathbf{W}}_n)^{-1}), \quad (16)$$

where

$$\hat{\varphi}_n := (\hat{\varphi}_{1:n}, \dots, \hat{\varphi}_{n:n}) \quad (17)$$

is found numerically via Newton-Raphson iterations using the training data  $\mathcal{D}_n$ , and  $\hat{\mathbf{W}}_n$  is the  $n \times n$  Fisher Information matrix of the binomial (negative) log-likelihood  $l(\varphi_{1:n}) := \log p(B_{1:n}|a_{1:n}, \varphi_{1:n})$ . Importantly, if the canonical link is used, then the  $i$ -th entry of  $\hat{\mathbf{W}}_n$  corresponds to the variance of the binomial response  $B_i$  at  $x_i$ :

*Lemma 2.1.* Under the Bernoulli link function (7), the Hessian  $\hat{\mathbf{W}}_n(\varphi_{1:n}) = -\Delta l(\varphi_{1:n})$  is diagonal:

$$w_{ij} = \begin{cases} a_i \Theta(\varphi_{i;n})(1 - \Theta(\varphi_{i;n})), & i = j, \\ 0 & i \neq j, \quad \text{for } i, j = 1, \dots, n. \end{cases} \quad (18)$$

Hence, we have that  $\hat{\mathbf{W}}_n = \text{diag}(\hat{w}_{1;n}, \dots, \hat{w}_{n;n})$ ; where  $\hat{w}_{i;n} := a_i \Theta(\hat{\varphi}_{i;n})(1 - \Theta(\hat{\varphi}_{i;n}))$  are evaluated at the posterior mode (17). Having found the joint (16), the (approximated) predictive posterior density  $\varphi_* \sim \mathbf{N}(m_n(x), s_n^2(x))$  is also Gaussian with mean  $m_n(x) \equiv m_n(x; \hat{\varphi}_n)$  and posterior variance  $s_n^2(x) \equiv s_n^2(x; \hat{\varphi}_n)$ :

$$m_n(x) := \mathbf{K}_n^T \mathbf{K}_n^{-1} \hat{\varphi}_n; \quad (19a)$$

$$s_n^2(x) := \boldsymbol{\kappa}_n^T (\mathbf{K}_n + \hat{\mathbf{W}}_n^{-1})^{-1} \boldsymbol{\kappa}_n, \quad (19b)$$

where  $\boldsymbol{\kappa}_n := (\kappa(x, x_1), \dots, \kappa(x, x_n))^T$  is the  $n \times 1$  vector of covariances between  $\varphi_*$  and  $\varphi_{1:n}$ . The resulting point estimate for  $\theta(x)$  given  $\mathcal{D}_n$  is thus

$$\hat{\theta}_n^{GP}(x) := \int_{\mathbb{R}} (1 + e^{-\varphi_*})^{-1} \mathbf{N}(\varphi_*; m_n(x), s_n^2(x)) d\varphi_*, \quad x \in (0, 1). \quad (20)$$

Numerically,  $\hat{\theta}_n^{GP}(x)$  is obtained by approximating the integral in (20) via a quadrature method. In particular we use `integrate()` which is part of the core distribution of **R** and relies on the Gauss-Kronrod quadrature method (R Core Team, 2016).

**Hyper-parameter estimation.** The above model specification is valid for fixed hyper-parameters  $\vartheta$ . To optimize the latter, we consider a maximum a posteriori estimate (MAP),  $\hat{\vartheta} := \arg \max_{\vartheta} \{\log q(\mathcal{D}_n | \vartheta) + \log q_0(\vartheta)\}$  based on a prior  $q_0(\cdot)$ . In order to obtain  $\hat{\vartheta}$  we use the package `GPstuff` (Vanhatalo et al., 2013), which uses interleaved numerical optimization: at iteration  $m$  given  $\hat{\vartheta}^{(m)}$ , evaluate the covariance matrix  $\mathbf{K}_n(\hat{\vartheta}^{(m)}) = (\kappa_{\hat{\vartheta}^{(m)}}(x_i, x_j))_{i,j=1}^n$  and so estimate the mode  $\hat{\varphi}_n^{(m)}$ ; then fix  $\hat{\varphi}_n^{(m)}$  and find  $\hat{\vartheta}^{(m+1)} = \arg \max_{\vartheta} \log q(\mathcal{D}_n | \vartheta, \hat{\varphi}_n^{(m)}) + \log q_0(\vartheta)$ , where  $q(\mathcal{D}_n | \vartheta, \hat{\varphi}_n^{(m)})$  is the data marginal log-likelihood,

$$\log q(\mathcal{D}_n | \vartheta, \hat{\varphi}_n) = -\frac{1}{2} \hat{\varphi}_n^T \mathbf{K}_n(\vartheta)^{-1} \hat{\varphi}_n + \log p(B_{1:n} | a_{1:n}, \hat{\varphi}_n) - \frac{1}{2} \log \{ |\mathbf{K}_n(\vartheta)| \cdot |\mathbf{K}_n(\vartheta)^{-1} + \hat{\mathbf{W}}_n(\hat{\varphi}_n)| \},$$

which is available in closed-form, see Algorithm 5.1 in Williams and Rasmussen (2006).

## 2.2 Adaptive Batching using the Posterior GP Variance

The posterior variance  $s_n(\cdot)$  of the surrogate quantifies the quality of learning the latent GP. It can be used to guide sampling decisions via the associated information gain regarding  $\varphi(\cdot)$ . This is achieved by considering the look-ahead  $s_{n+1}(\cdot)$  conditional on sampling at  $x_{n+1}$ . For plain GPs,  $s_{n+1}$  is independent of the future response and hence can be evaluated exactly. Unfortunately, for binomial GPs the look-ahead predictive variance *does* depend on the future  $B_{n+1}(x_{n+1})$ . Specifically, Equation (21) expresses the fact that  $s_{n+1}^2(x_{n+1})$  depends on the entire  $\hat{\varphi}_{n+1}$  (computed based on  $\mathcal{D}_{n+1}$ ).

*Theorem 2.2.* The look-ahead variance  $s_{n+1}^2(x_{n+1})$  at a new location  $x_{n+1}$  under the Laplace approximation (19) is given by

$$s_{n+1}^2(x_{n+1}) = \left( \frac{1}{s_n^2(x_{n+1}; \hat{\varphi}_{1:n,n+1})} + \frac{1}{a_{n+1} \cdot \Theta(\hat{\varphi}_{n+1,n+1})(1 - \Theta(\hat{\varphi}_{n+1,n+1}))} \right)^{-1} \quad (21)$$

$$\simeq \left( \frac{1}{s_n^2(x_{n+1})} + \frac{1}{a_{n+1} \hat{\theta}_n^{GP}(x_{n+1})(1 - \hat{\theta}_n^{GP}(x_{n+1}))} \right)^{-1}, \quad (22)$$

where  $s_n^2(x_{n+1})$  is the iteration- $n$  posterior variance from (19b) and  $\hat{\theta}_n^{GP}(x_{n+1})$  is from (20).

The approximation in (22) aims to remove the dependence of (21) on  $B_{n+1}$  by using only information available at iteration  $n$ . To do so, we approximate the denominator of the first term in (21) via  $s_n^2(x_{n+1}; \hat{\varphi}_n) \simeq s_n^2(x_{n+1}; \hat{\varphi}_{1:n,n+1})$ , that is, using the estimated posterior mode at time  $n$ . Similarly, the future local binomial variance in the second term of (21) is approximated by its iteration- $n$  counterpart  $a_{n+1} \hat{\theta}_n^{GP}(x_{n+1})(1 - \hat{\theta}_n^{GP}(x_{n+1}))$ ; see the full proof in the Appendix.

The look-ahead variance forms the basis of numerous *expected improvement* (EI) design heuristics that quantify the gain from sampling at  $x_{n+1}$ , see e.g. Jones et al. (1998); Chevalier et al. (2014). Below we adapt these concepts to the setting of binomial GPs by quantifying the approximate reduction in posterior variance of  $\varphi(x_{n+1})$  due to sampling  $a_{n+1}$  replicates at  $x_{n+1}$  and hence allowing optimization of  $a_{n+1}$  conditional on  $x_{n+1}$ . Related batched EI criteria have recently appeared in Kamiński (2015); see also Binois et al. (2018a).

The idea of adaptive replication is to aim for driving the iteration- $n+1$  variance  $s_{n+1}^2(x_{n+1}) \leq \nu_n$  below a threshold  $\nu_n$ . Using the variance decomposition formula in the RHS of (22) and solving for  $a_{n+1}^\nu$  we have that:

$$a_{n+1}^\nu \geq \frac{1}{\hat{\theta}_n(x_{n+1})(1 - \hat{\theta}_n(x_{n+1}))} \cdot \left( \frac{1}{\nu_n} - \frac{1}{s_n^2(x_{n+1})} \right).$$

We therefore consider the following adaptive replication scheme:

$$\hat{a}_{n+1}^\nu := a_0^\nu \cdot \mathbb{1}_{\{s_n^2(x_{n+1}) < \nu_n\}} + \frac{1}{\hat{\theta}_n(x_{n+1})(1 - \hat{\theta}_n(x_{n+1}))} \left( \frac{1}{\nu_n} - \frac{1}{s_n^2(x_{n+1})} \right) \cdot \mathbb{1}_{\{s_n^2(x_{n+1}) \geq \nu_n\}}. \quad (23)$$

*Remark 2.* We focus on the predictive uncertainty in the latent process  $\varphi$  as a measure to determine  $a_{n+1}$  —as opposed to the predictive variance of the random variable  $\theta(\varphi(x_{n+1}))$ . Focusing on the uncertainty of the latent GP is a common strategy in sequential design (especially when the data likelihood is Gaussian), see for example Ankenman et al. (2010); Chen and Zhou (2017). Another common measure for constructing sequential designs is the posterior predictive entropy (Kapoor et al., 2007) which is the preferred uncertainty measure in the active learning framework.

### 2.3 MLE-Based Binomial Regression

An alternative approach to B-GPs is to fit a linear surrogate of the form  $\varphi(x) := \beta^T \phi(x)$  for a given set of *basis functions*. Thus we seek the best fit in the function space  $\mathcal{H} = \text{span}(\phi_j :$

$j = 1, \dots, p$ ). The coefficients  $\beta \in \mathbb{R}^p$  can be found by optimizing the *penalized* binomial log-likelihood criterion

$$\min_{\beta} \sum_{i=1}^n \left\{ B_i \sum_{j=1}^p \beta_j \phi_j(x_i) + a_i \log \left( 1 + \exp \left( \sum_{j=1}^p \beta_j \phi_j(x_i) \right) \right) \right\} + \frac{1}{2} \lambda \mathcal{J} \left( \sum_{j=1}^p \beta_j \phi_j \right), \quad (24)$$

where  $\mathcal{J}(\varphi)$  is a penalty functional. The above specification includes the classical *logistic regression model* when the basis elements in  $\mathcal{H}$  are monomials and  $\lambda = 0$ , which we also implemented with AIC-based selection of the degree of the polynomial.

**Kernel Logistic Regression (KLR).** Another choice is the family of positive definite kernel functions  $\phi_j(\cdot) := \kappa_{l_j}(\cdot; \xi_j)$ , where each basis element  $\kappa_{l_j}(\cdot, \xi_j)$  is indexed by a location parameter  $\xi_j$  and a scale parameter  $l_j$ . The corresponding space of functions  $\mathcal{H}$  is a Reproducing Kernel Hilbert Space with penalty functional  $\mathcal{J}(\varphi) = \|\varphi\|_2^2 = \beta^T \Phi \beta$ , where  $\Phi_{ij} = \phi_j(x_i)$ . A popular choice is the *Gaussian radial kernel*:

$$\kappa_l(x; \xi) := \exp \left( -\frac{|x - \xi|^2}{l^2} \right). \quad (25)$$

KLR behaves similarly to Support Vector Machines: data inputs are mapped to a space spanned by positive definite kernel functions, and the loss function being optimized are also similar (Zhu and Hastie, 2005). For our purposes, it is natural to use  $\xi_i = x_i$ , i.e. a separate kernel function for each query location.

**Spline Logistic Regression (SLR).** A further commonly used functional space  $\mathcal{H}$  is the B-spline basis where the  $\phi_j$ 's are piecewise continuous functions defined in terms of a set of *knots*. Namely, an order- $P$  spline with knots  $(\xi_j)_{j=1}^p$  is a piecewise-polynomial of order  $P$ , and has continuous derivatives up to order  $P - 2$ . The B-spline family takes  $P = 4$  and can be represented in terms of  $p > 0$  basis functions:  $\phi_1(x) = 1$ ,  $\phi_2(x) = x$ , and for  $j = 2, \dots, p$ ,  $\phi_{j+1}(x) = d_j(x) - d_{j-1}(x)$ , where

$$d_j(x) := \frac{(x - \xi_j)_+^3 - (x - \xi_p)_+^3}{\xi_p - \xi_j}, \quad j = 2, \dots, p. \quad (26)$$

The basis coefficients  $\beta$  are fitted by penalizing the curvature of  $\varphi(\cdot)$  using  $\mathcal{J}(\varphi) = \|\varphi''\|_2^2$ .

### 3 Sampling Policies

To make sampling decisions based on the surrogate  $\varphi$  and the information about  $X^*$  contained in  $f_n$  we investigate three types of policies.

**Batched Information-Directed Sampling.** Our first approach utilizes fixed replication  $a \geq 1$  and selects the next  $x_{n+1}$  to maximize the estimated batched expected KL divergence between the knowledge state at  $T_n$  and  $T_n + a$  as in (10), available in closed-form according to Theorem 3.1.

*Theorem 3.1.* Let  $x \in (0, 1)$  and  $f_n$  be the current knowledge state with cumulative distribution function (CDF)  $F_n(\cdot)$ . The expected KL divergence,  $\mathbb{E}[D(f_{n+1}; f_n)]$  between  $f_{n+1}$  and  $f_n$  from  $a$  queries at  $x$  is given by

$$\mathbb{E}[D(f_{n+1}; f_n)] = \mathbb{E} \left[ \log_2 \left( \frac{(1-p(x))^B p(x)^{a-B}}{c_n(x, B)} \right) \right] F_n(x) \quad (27a)$$

$$+ \mathbb{E} \left[ \log_2 \left( \frac{p(x)^{\tilde{B}} (1-p(x))^{a-\tilde{B}}}{c_n(x, \tilde{B})} \right) \right] (1 - F_n(x)). \quad (27b)$$

where the expected values (27a) and (27b) are taken with respect to  $B \sim \text{Bin}(a, 1-\theta(x))$  and  $\tilde{B} \sim \text{Bin}(a, \theta(x))$ , respectively.

*Proof.* By definition, the KL divergence between  $f_n$  and  $f_{n+1}$  is:

$$D(f_{n+1}; f_n) = \int_0^1 \log_2 \left( \frac{f_{n+1}(u)}{f_n(u)} \right) f_n(u) du.$$

Since  $\mathbb{P}_p(B = j | a, x, u) := \text{Bin}(j; a, 1 - \theta(x)) 1_{\{u \leq x\}} + \text{Bin}(j; a, \theta(x)) 1_{\{u > x\}}$  and for  $0 < u < x^*$  we have that  $f_{n+1}(u) := [(1-p(x))^B p(x)^{a-B}] f_n(u) / c_n(x, B)$ , cf. (9), so taking expectations end up with

$$\begin{aligned} \mathbb{E}[D(f_{n+1}; f_n)] &= \int_0^x \mathbb{E} \left[ \log_2 \left( \frac{(1-p(x))^B p(x)^{a-B}}{c_n(x, B)} \right) \right] f_n(u) du \\ &\quad + \int_x^1 \mathbb{E} \left[ \log_2 \left( \frac{p(x)^{\tilde{B}} (1-p(x))^{a-\tilde{B}}}{c_n(x, \tilde{B})} \right) \right] f_n(u) du \end{aligned}$$

which simplifies to (27). Above

$$c_n(x, B) := \left[ (1-p(x))^{B(x)} p(x)^{a-B(x)} \right] F_n(x) + \left[ p(x)^{B(x)} (1-p(x))^{a-B(x)} \right] (1 - F_n(x))$$

is the normalizing constant of the updating (9).  $\square$

We now re-use KL divergence to define an acquisition function  $\mathcal{I}(x, f_n; p(x), a) := \mathbb{E}_p[D(f_{n+1}; f_n)]$  as in (27) (emphasizing the dependence on the oracle accuracy  $p$ ) and which is to be maximized over  $x$ . To illustrate the relationship between the knowledge state  $f_n$  and the batched information-criterion  $x \mapsto \mathcal{I}(x, f_n; p(\cdot), a)$ , Figure 1 shows a realization of the Spatial IDS/PBA algorithm for a fixed batch size  $a = 100$  for  $n = 0, \dots, 9$  starting with  $g_0 \equiv \text{Unif}(0, 1)$  prior on  $X^*$ . The underlying response is (31) consisting of a decreasing linear function  $h_1$  with root at  $x^* = 1/3$ , and for now we assume access to the true oracle accuracy  $p(x) = \Phi(5|x - 1/3|)$ , so that the knowledge state is the exact Bayesian posterior  $g_n$ . We notice that sampling at  $x_{n+1}^{\text{IDS}}$  concentrates  $g_n$  rapidly around the root  $x^*$ . Furthermore,  $\mathcal{I}$  typically has two local maxima, along with a global minimum at  $x^*$  (sampling at the root is completely uninformative due to  $p(x^*) = 0.5$ ). The right panel of Figure 1 plots  $x \mapsto \mathcal{I}(x, g_9; p(x), a_{10})$  across different replication values  $a_{10} \in \{1, 10, 5, 100, 250\}$ . It can be seen that as the batch size  $a$  is increased, information gain increases, but the maximizer  $\arg \max \mathcal{I}(x, g_9; p(x), a)$  (vertical dotted lines) does not change significantly. This is partly because the posterior  $g_9$  is already concentrated.

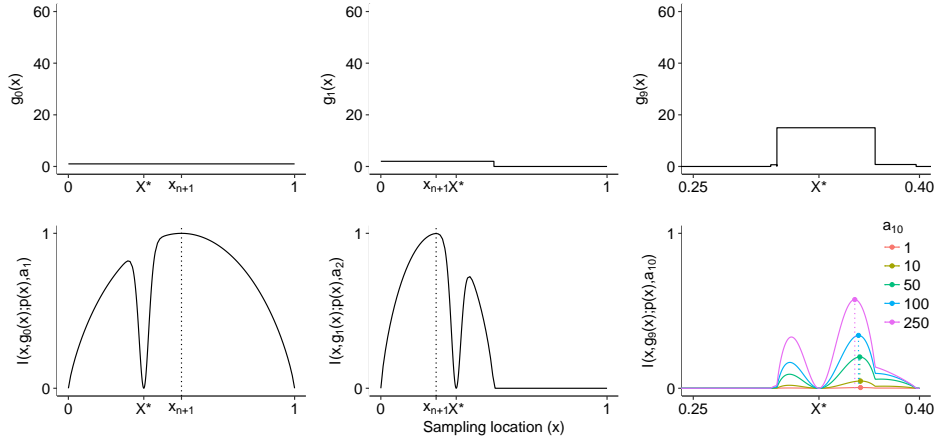


Figure 1: Data acquisition procedure using the batched information criterion  $\mathcal{I}(\cdot, g_n; p(\cdot), a_{n+1})$  starting with a Uniform prior  $g_0$  for the linear test function (31) and  $x^* = 1/3$ . The first row shows the true Bayesian posterior  $g_n$  for  $n \in \{0, 1, 9\}$ . The second row depicts the information gain function along with its maximizer  $x_{n+1}$  (vertical dotted lines). The right-bottom plot shows the information criterion  $\mathcal{I}$  for several replication sizes  $a_{10} \in \{1, 10, 50, 100, 250\}$  with the corresponding maximizers  $x_{10}^{\text{sIDS}}$  of the information criterion  $\mathcal{I}$  (vertical dotted lines) given the knowledge state  $g_9$  obtained by updating  $g_0$  using fixed replication amounts  $a_{1:9} = 100$ .

Crucially, maximizing (27) requires knowledge of the entire  $x \mapsto p(x)$ . This was one of the main challenges in the original G-PBA, where IDS was applied ad hoc *after* estimating  $p(\tilde{x}_i)$  at a set of  $M \geq 2$  candidate locations  $\tilde{x}_{1:M}$ . However, under our spatial modeling setting one can plug-in the surrogate  $\hat{p}_n(x)$  and compute the maximizer of the resulting  $\mathcal{I}$  conditional on sampling  $a_{n+1} \geq 1$  times at any  $x \in (0, 1)$ . Thus,  $x_{n+1}$  is chosen greedily as the maximizer of  $\mathcal{I}(\cdot, f_n, \hat{p}_n(\cdot), a_{n+1})$ , that is,

$$x_{n+1}^{\text{sIDS}} := \arg \max_{x \in (0,1)} \mathcal{I}(x, f_n; \hat{p}_n(x), a_{n+1}). \quad (28)$$

A numeric optimization procedure is needed to find  $x_{n+1}^{\text{sIDS}}$ . In our experiments below we utilize the R package `NLOpt` (Johnson, 2014). In particular, we use the `DIRECT` (DIviding `RECT`angles) algorithm (Jones et al., 1993) that implements gradient-free deterministic-search global optimization.

The next two schemes switch the order, first picking  $x_{n+1}$  and then  $a_{n+1}$ .

**Adaptive One-Step IDS policy.** Note that (28) requires specifying the replication amount  $a_{n+1}$ . To implement the adaptive replication scheme (23) within the IDS approach, we use an ad hoc heuristic which first maximizes  $\mathcal{I}$  using  $a = 1$  to get  $x_{n+1}$  and then selects the replication amount  $a_{n+1}^\nu$ . Let

$$x_{n+1}^{\text{Ada-sIDS}} := \arg \max_{x \in (0,1)} \mathcal{I}(x; f_n, \hat{p}_n(x), 1). \quad (29)$$

Conditional on  $x_{n+1}$ ,  $a_{n+1}$  is then picked to control the surrogate uncertainty at  $x_{n+1}$  according to (23). Observe that Ada-IDS is only feasible with a B-GP surrogate furnishing the predictive variance  $s_n(x_{n+1})$ .

**Randomized Quantiles Sampling.** The RQS strategy randomizes the next sampling location according to

$$x_{n+1}^{\text{RQS}} := F_n^{-1}(U_{n+1}), \quad \text{where } U_{n+1} \sim \text{Unif}(0, 1). \quad (30)$$

The RQS policy can be interpreted as sampling based on the posterior distribution of  $X^*$ . This tends to sample close to the mean of  $f_n$  but will also occasionally explore the latter’s tails, capturing the trade-off between exploitation and exploration. An attractive feature of RQS is that it relies solely on  $f_n$  so the surrogate  $\theta_n$  is only used for updating  $f_n$  in (9).

### 3.1 The Spatial Generalized Probabilistic Bisection Algorithm

Summarizing the above developments, Algorithm 2 specifies the ingredients for blending surrogate modeling with probabilistic bisection. Two remarks are in order. First, the initialization step is non-sequential: we begin by employing  $N_0 \times a_0 = T_0 \ll T$  oracle evaluations to build  $\hat{\varphi}_{N_0}$ , picking equidistant (i.e. space-filling) sites  $x_{1:N_0}$  in  $(0, 1)$  and  $a_0 \geq 1$  replications per site. The corresponding  $f_{T_0}$  is constructed via (9). Second, the surrogate re-fitting step in Algorithm 2 is user-controlled, since re-fitting can be expensive. In principle, re-fitting could be stopped entirely once  $n$  is large enough, keeping the overhead cost of predicting  $\theta_n(x)$  fixed, rather than increasing in  $n$ . We also note that the chosen surrogates are non-sequential, i.e. re-estimating  $\hat{\varphi}_n$  is done from scratch, rather than via an updating formula (like is done for  $f_n$ ).

## 4 Numeric Examples

We proceed to empirically assess the performance of Algorithm 2. To do so, we mix-and-match the three components that the user must pick: the sampling policy  $\eta$ , surrogate model for  $\hat{\theta}$  and the batch size  $a$  (fixed or adaptive). To analyze the algorithm sensitivity to  $(\eta, \hat{p}, a)$ , we consider multiple metrics regarding the quality of the root estimates, namely absolute residuals, credible interval length, and corresponding coverage. Furthermore, we benchmark against schemes that are allowed to use the true posterior  $g_n$  and  $p(\cdot)$ , quantifying the impact of learning the oracle. Our numeric examples are based on three test functions which capture different aspects and difficulties typically encountered in SRFPs, such as heteroscedasticity or zero curvature at the root location.

### 4.1 Experimental Setup

In analogy to Waeber (2013); Rodriguez and Ludkovski (2017), we utilize the following three test functions  $h_i(x)$  defined for  $x \in (0, 1)$ :

**PBA parameters:** Prior  $f_0$ ;  $T_0$  and  $a_0 \geq 1$ . Set  $N_0 := T_0/a_0$ ;  
**Surrogate initialization:** Regress  $B_{1:N_0}$  on locations  $x_{1:N_0}$  to obtain the surrogate model  $\hat{\theta}_{N_0}$ ;  
Update knowledge state starting from  $f_0$  to  $f_{T_0}$  given  $\hat{\theta}_{N_0}$ ,  $B_{1:N_0}$  and  $x_{1:N_0}$ ;  
 $n \leftarrow N_0$ ,  $T_n \leftarrow T_0$ ,  $\mathcal{D}_n \leftarrow (B_{1:N_0}, a_{1:N_0})$ ;  
**while**  $T_n < T$  **do**  
    Using  $f_n$  generate next sampling location  $x_{n+1}$  and batch size  $a_{n+1}$ ;  
    Query oracle  $a_{n+1}$  times at  $x_{n+1}$  to observe  $B_{n+1}(x_{n+1})$ ;  
    **if** (*OPTIONAL*) **then**  
        Re-fit surrogate for  $\hat{\theta}_{n+1}$  based on  $\mathcal{D}_{n+1} = (\mathcal{D}_n, B_{n+1}, a_{n+1})$ ;  
    **else**  
         $\hat{\theta}_{n+1} \leftarrow \hat{\theta}_n$ ;  
    Update knowledge state at  $x_{n+1}$   $f_{n+1} \leftarrow \Psi(f_n, x_{n+1}, B_{n+1}; \hat{p}_{n+1}, a_{n+1})$  using  
         $\hat{p}_{n+1} = \max\{\hat{\theta}_{n+1}(x_{n+1}), 1 - \hat{\theta}_{n+1}(x_{n+1})\}$ ;  
    Update  $T_n \leftarrow T_n + a_{n+1}$  and  $n \leftarrow n + 1$ ;  
**end**  
**return** Knowledge state  $f_N$  and estimator for the root location  $\hat{x}_N = \text{median}(f_N)$ ;

**Algorithm 2:** Spatial Generalized-PBA.

$$\text{linear} \quad h_1(x) = X^* - x, \quad \sigma_1(x) = 0.2; \quad (31)$$

$$\text{exponential} \quad h_2(x) = e^{2(X^* - x)} - 1, \quad \sigma_2(x) = 0.2 \cdot 1_{\{x < X^*\}} + 1 \cdot 1_{\{x > X^*\}}; \quad (32)$$

$$\text{cubic} \quad h_3(x) = (X^* - x)^3, \quad \sigma_3(x) = 0.025. \quad (33)$$

In all cases the stochastic simulator (2) consists of a Normally distributed  $\epsilon(x) \sim \mathbf{N}(0, \sigma_i^2(x))$  random noise and the root location  $X^* \sim \text{Unif}(0, 1)$  is drawn from a Uniform distribution on  $(0, 1)$ . We thus have that the ground-truth oracle  $\theta(x)$  is given by  $\theta_i(x) := \Phi(-h_i(x)/\sigma_i(x))$  for  $i = 1, 2, 3$ .

Figure 2 displays the test functions  $h_i(x)$  with  $x^* = 1/3$  (first row), the maps  $x \mapsto \theta_i(x)$  (second row), and the corresponding  $\text{logit}(\theta_i(x))$  (third row) used for constructing  $\hat{\varphi}$ . The base example we investigate is the linear function (31) whose slope is constant and significantly different from zero in locations close to the root  $X^*$  and therefore leads to a simpler SRFP. In contrast, the curvature of (32) together with the non-constant  $\sigma_2(x)$  create a skew in the oracle and the posterior  $f_n$ . Finally, example (33) represents a difficult root-finding setting due to  $h_3'(X^*) = 0$ , which implies that  $p(x) \simeq 1/2$  in the vicinity of  $X^*$ .

**Performance Evaluation Metrics.** For a given configuration  $(\eta, \hat{p}, a)$  we use the following four performance metrics of the resulting Spatial G-PBA that all rely on  $f_n$ :

1. *Absolute residuals:* to determine the accuracy of the estimator  $\hat{x}_n := \text{median}(f_n)$  we consider the  $L_1$ -residuals,

$$r(f_n) := |\hat{x}_n - x^*|; \quad (34)$$






Figure 2: Synthetic test functions (31), (32) and (33) for Section 4.

2. *Credible intervals*: we evaluate the degree of uncertainty associated to the unknown root location  $X^*$  through the length of a symmetric  $(1 - \alpha)\%$  credible interval (CI) between the  $\alpha/2$  and  $(1 - \alpha/2)$  percentiles of  $f_n$ :

$$l_{1-\alpha}(f_n) := F_n^{-1}(1 - \alpha/2) - F_n^{-1}(\alpha/2); \quad (35)$$

3. *Coverage*: to measure the accuracy of the above CI we evaluate

$$c_{1-\alpha}(f_n) := \Pr \{x^* \in [F_n^{-1}(\alpha/2), F_n^{-1}(1 - \alpha/2)]\}, \quad (36)$$

where the averaging in  $\Pr\{\cdot\}$  is across Monte-Carlo (MC) runs of the algorithm to capture the sampling distribution. If  $c_{1-\alpha}(f_n) \ll (1 - \alpha)$  the coverage test indicates that  $f_n$  prematurely collapses or equivalently overstates its confidence about  $X^*$ . Small CI length  $l_{1-\alpha}$  relative to residuals  $r$  will lead to low coverage  $c$ . For both  $c$  and  $l$  we use  $\alpha = 0.05$ .

4. *KL divergence*: given the chosen querying sites  $x_{1:n}$ , we compare  $f_n$  to the true posterior  $g_n$  (which is available for our three synthetic examples but not for the case-study in Section 5) using the KL divergence,  $D(f_n; g_n)$ . Since both  $f_n$  and  $g_n$  are updated at the same set of knots (sorted in increasing order)  $\tilde{x}_{1:n}$ , we may write  $g_n(x) := \sum_{j=1}^n g(\tilde{x}_{j-1})1_{x \in [\tilde{x}_{j-1}, \tilde{x}_j]}$  and  $f_n(x) := \sum_{j=1}^n f(\tilde{x}_{j-1})1_{x \in [\tilde{x}_{j-1}, \tilde{x}_j]}$ , with  $\tilde{x}_0 := 0$  and  $\tilde{x}_n := 1$ . We then obtain

$$D(f_n; g_n) := \sum_{j=1}^n \log \left( \frac{f(\tilde{x}_j)}{g(\tilde{x}_j)} \right) f(\tilde{x}_j)(\tilde{x}_j - \tilde{x}_{j-1}). \quad (37)$$

We make the usual convention that  $\log(f(x)/g(x))f(x) = 0$  if  $f(x) = 0$  (including when  $g(x) = 0$ ); as well as  $\log(f(x)/g(x))f(x) = +\infty$  if  $g(x) = 0$  and  $f(x) > 0$  (Cover and

Thomas, 2006). Practically, to estimate the average KL divergence we consider only finite values.

Performance metrics 1-4 are averaged using a total of  $MC = 100$  Monte Carlo macro-iterations. To make all schemes comparable, we fix  $X_{(i)}^* \sim \text{Unif}(0, 1)$  and each combination of  $(\eta, \hat{p}, a)$  is applied using the same root value  $X_{(i)}^*$  during the  $i$ -th MC iteration,  $i = 1, \dots, MC$ .

**Surrogates for  $p(\cdot)$ :**

- *B-GP*. For the binomial GP (B-GP) we use the 5/2-Matérn covariance kernel (12). The hyper-parameters  $\vartheta = (\tau^2, l)$  are estimated via a Bayesian MAP estimation procedure, placing a square root uniform prior (i.e.,  $q_0(\sqrt{\tau^2}) \propto 1$ ) on  $\tau^2$  and a Student- $t$  prior on the length scale parameter  $l$  (both default priors for binomial GPs in `GPstuff`). Although parameter estimation can be expensive, the B-GP is *re-fitted* and *updated* every  $T_n = a_n$  simulation outputs; that is, the hyper-parameters  $\hat{\vartheta}$  are re-fitted and the posterior mode  $\hat{\varphi}_n$  is re-computed every time a new pair of sampling location/binomial response is observed, such that the surrogate is able to assimilate acquired information.
- *KLR*. Kernel Logistic Regression (KLR) is implemented with the Gaussian kernel basis function (25) using a *fixed* length scale parameter  $l \equiv 1$  and centering  $\phi_j$  at each sampling location  $\xi_j \equiv x_j$ ,  $j = 1, \dots, n$  (implying that we use as many kernel functions as sampling points to learn  $\varphi$ ). Since we would like to induce a surrogate model  $\hat{\varphi}$  that closely resembles the local estimators  $\hat{p}(\cdot)$ , we use a (small) *fixed* value  $\lambda = 0.01$  as the penalty parameter for optimizing (24). Numerically, we implement KLR as stated in Algorithm 1 of Zhu and Hastie (2005).
- *SLR*: We consider a *smoothing* spline logistic regression (SLR) model where the penalty coefficient  $\hat{\lambda}$  (aka smoothing parameter) is estimated via Generalized Cross-Validation (Friedman et al., 2001) jointly with the spline basis coefficients. In this case, the spline knots  $\xi_j$  are placed at percentiles of the sampling locations  $x_{1:n}$ . Thus, as the mass of  $f_n$  concentrates around  $x^*$  (and hence sampling locations  $x_{1:n}$  concentrate around the root), more knots  $\xi_j$ 's are also placed near  $X^*$ , making the surrogate more localized in regions where the variability of the binomial responses  $B_n$  is maximal.
- *LR*. Polynomial logistic regression with  $\varphi(x) = \beta_0 + \sum_{j=1}^5 \beta_j x^j$ , a quintic polynomial and zero penalty  $\lambda = 0$  (to enforce surrogate flexibility). Both the SLR and LR surrogates are implemented using the `gam()` routine from the `mgcv` package in R (Wood, 2001).

**Sampling Policies  $\eta$ :**

- Spatial Information-Directed Sampling (sIDS) (28);
- Spatial Randomized Quantile Sampling (sRQS) (30);
- One-step sIDS (29) combined with the adaptive replication scheme  $a_{n+1}^\nu$ .

For the initialization stage in Algorithm 2 we use  $N_0$  equally spaced  $x_{1:N_0}$  to learn  $\varphi_{N_0}(\cdot)$  non-sequentially. In our experiments all surrogates are initialized using  $T_0 := 5000$  (i.e., 25% of total sampling budget) oracle evaluations with  $a_0 \in \{100, 250\}$  which results in  $N_0 := T_0/a_0 \in \{50, 20\}$  initial training locations.

**Adaptive Replication**  $a'_{n+1}$ . The scheme (23) has two parameters: the *minimum replication amount*  $a'_0$  and the variance *thresholding sequence*  $(\nu_n)_{n \geq 1}$ . In our experiments, we use  $a'_0 := 1$  in order to favor exploration in regions where the spatial surrogate  $\varphi$  already learned  $p(\cdot)$  sufficiently well, as quantified in terms of the predictive posterior GP variance (19b). For  $\nu_n$  we use the following two variants (see Algorithm 2):

$$\nu_n^{(100)} := 0.1/n \quad \text{when } a_0 = 100 \quad \text{and} \quad \nu_n^{(250)} := 0.05/n \quad \text{when } a_0 = 250. \quad (38)$$

This choice is linked to the fact that since the initialization stage budget  $T_0$  is fixed, larger  $a_0$  makes  $N_0$  smaller and hence leads to larger  $s_n$ , and so we take the thresholds  $\nu_n$  larger as well. To avoid excessive batching which could occasionally arise in our implementation we bound  $a'_{n+1}$  (specifically by 1000 in all experiments). This allows to manage the overall sampling budget in order to enforce exploration.

Figure 3 depicts the *realized* replication amounts  $n \mapsto a'_{n+1}(x_{n+1})$  using the one-step sIDS policy (29) applied to our running example (31) (during initialization,  $n \leq N_0 := T_0/a_0$ ,  $a_n \equiv a_0$  is fixed). We observe that  $a'_{n+1}$  generally slowly decreases as  $n$  rises, although the local behavior can be quite “spiky”: sometimes a large batch is required to bring  $s_n^2(x_{n+1})$  below  $\nu_n$ , see top panels of Figure 3. One reason is that as  $n$  increases, the sampling concentrates around  $x^*$ . Since this region quickly becomes well-explored, we usually obtain quite low  $s_n^2(x_{n+1})$  making  $a_{n+1}$  low as well. To give a sense of the macro-time behavior, by  $T = 2 \cdot 10^4$  the median number of sampling locations is  $N_T = 291$  and  $N_T = 116$  (for mean replication amounts of 70 and 170 respectively) for the thresholding sequences  $\nu_n^{(100)}$  and  $\nu_n^{(250)}$  in (38).

## 4.2 Illustrating Spatial G-PBA

Figure 4 compares the fitted surrogate models using a fixed dataset  $\mathcal{D}_N^\eta$  (in order to remove the effect of the design and target surrogate accuracy) generated using two different sampling policies  $\eta$ : sIDS (first row) and sRQS (second row) implemented using the true posterior  $g_n$  and the linear test function (31). For all models a fixed batch size  $a \equiv 100$  and  $N = 200$  total training locations is used. Figure 4 depicts three fundamental features of spatial G-PBA: (i) the sIDS strategy achieves lower posterior  $X^*$ -uncertainty relative to the sRQS policy, as seen in the narrower confidence bands depicted in the right panel; (ii) the design of the sIDS strategy brackets the root, gradually squeezing the posterior  $f_n$  towards  $X^*$ ; and (iii) the spatial surrogates succeed in learning the true  $\theta_1(x) := \Phi(-\frac{1/3-x}{0.20})$  especially around the root, cf. the left panels of the Figure. As a result, root estimation is significantly improved and leads to reliable posterior CIs on the right panels of Figure 4.




Figure 3: *First row*: estimated predictive variance  $s_n^2(x_{n+1})$  and the thresholding sequences  $\nu_n^{(100)} = 0.1/n$  (first column) and  $\nu_n^{(250)} = 0.05/n$  (second column), when the initial batch size for initializing the B-GP is  $a_0 = 100, 250$ , respectively. *Second row*: adaptive replication amount  $a_{n+1}^v(x_{n+1})$  ( $y$ -axis) in macro-time  $n$  ( $x$ -axis) selecting  $x_{n+1}$  using the one-step IDS criterion (29).

### 4.3 Results

Table 1 shows the results for the linear test function (31). To allow a direct comparison to the non-spatial G-PBA, the last few rows present the performance of the best *local* G-PBA schemes as identified in Rodriguez and Ludkovski (2017):

- the *empirical majority proportion*,  $\bar{p}(x_n) := \max\{B_n/a_n, 1 - B_n/a_n\}$ ; and
- the *posterior mode* given  $\bar{p}(x_n)$ ,  $\hat{p}_{L_0}(x) := \arg \max_{p \in (1/2, 1)} \pi(p|\bar{p}(x))$ ; where  $\pi(\cdot|\bar{p}(x))$  is the posterior density of  $p$  seen as a random variable with prior  $\pi_0 \equiv \text{Unif}(1/2, 1)$ .

The local estimators  $\bar{p}$  and  $\hat{p}_{L_0}$  are then applied within two non-spatial G-PBA policies:

- Deterministic-IDS (Det-IDS) which chooses  $x_{n+1}$  by maximizing  $\mathcal{I}(\tilde{x}_{n,i}; f_n, \hat{p}(\tilde{x}_{n,i}), a_0)$  among the two candidates  $\tilde{x}_{n,i} \in \{F_n^{-1}(0.25), F_n^{-1}(0.75)\}$  (i.e., the 25-th and 75-th quantiles of  $f_n$ );
- Local RQS which selects  $x_{n+1}$  according to (30).

Table 1 demonstrates that using surrogate modeling substantially improves root estimation relative to the original G-PBA. Indeed, we obtain significantly lower residuals (roughly half as

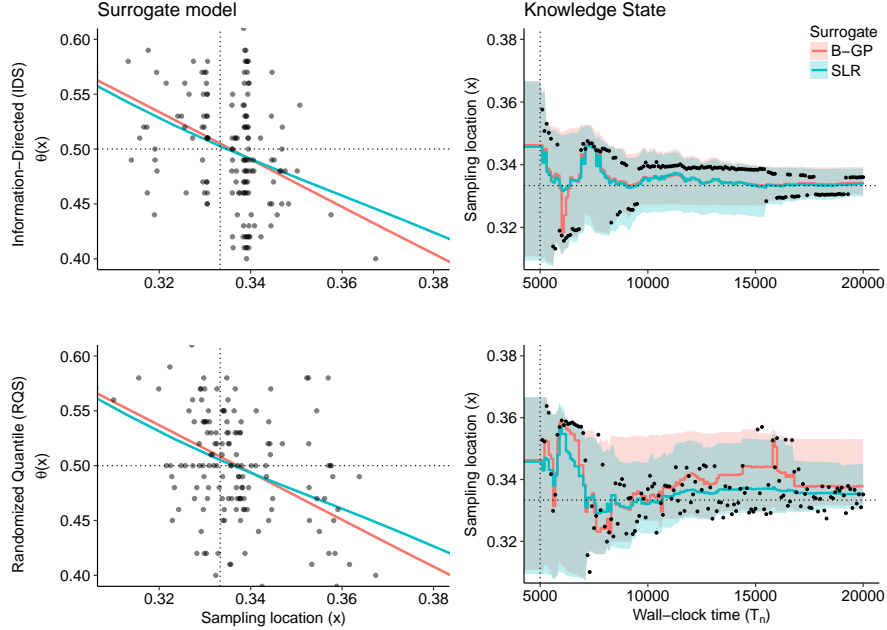


Figure 4: Spatial G-PBA with the linear test function (31). *Left*: B-GP, and SLR surrogates trained on a fixed dataset obtained using the sIDS policy (first row) and the sRQS policy (second row) at  $T = 2 \cdot 10^4$  and batch size  $a = 100$  (so that  $N = 200$ ). The  $x$ -axis is zoomed to the neighborhood of  $x^* = 1/3$ , so does not show the full  $\mathcal{D}_N$ . *Right*: posterior inter-quantile range (shaded regions) across spatial surrogates (colors) and sampling policies (rows) as a function of  $T_n$ . We also show the corresponding root estimates  $\hat{x}_n = \text{median}(f_n)$  (lines). The true  $\theta(\cdot)$ , as well as the estimated median using  $g_n$  (i.e., knowledge state with the true  $\theta(\cdot)$ ) is shown with a dashed line, respectively.

big), and narrower CI across while maintaining a high probability coverage. The latter shows that the knowledge state  $f_T$  is correctly converging to the true root value  $X^*$ . In particular, polynomial logistic regression (LR) offers a good choice as it minimizes the average residuals and length of CI, as well as matches the nominal coverage  $\hat{c}(f_T) \approx 0.95$ , confirming that  $f_T$  is close to the true posterior  $g_T$ .

Importantly, we can see that spatial modeling leads to a nearly two orders of magnitude reduction in the average KL divergence between  $f_T$  and  $g_T$ , primarily due to the lower bias in the estimation of  $p(\cdot)$  relative to the two local estimators considered. Thus, spatial G-PBA successfully resolves the problem of  $f_T$  experiencing premature collapse which was a major concern in G-PBA where  $c_\alpha$  was frequently unacceptably low. We note that  $\hat{D}(f_T; g_T)$  is consistently low across all surrogate models  $\varphi$ , indicating that the goodness-of-fit for  $\theta(\cdot)$  is not overly sensitive to the choice of the surrogate type.

In terms of the sampling policies, sIDS outperforms sRQS since the respective average residuals and CI length are lower while preserving a high coverage probability. For the replication regime  $a_n$ , we note a preference for  $a = 250$  (i.e., a total of  $N = 80$  design sites) which

Table 1: Performance Monte-Carlo metrics for the test function  $h_1$  at  $T = 20,000$ .

Policy $\eta$	$\hat{p}$	$\hat{r}(f_T)$ ( $10^{-2}$ )		$\hat{l}_{0.95}(f_T)$ ( $10^{-2}$ )		$\hat{c}_{0.95}(f_T)$		$\hat{D}(f_T; g_T)$	
		$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$
sIDS	B-GP	0.2241	0.1874	0.8931	0.9215	0.88	0.98	0.62	0.39
	KLR	0.2106	0.2037	0.8998	0.9496	0.95	0.96	0.57	0.38
	SLR	0.1864	0.1954	0.8669	0.8810	0.87	0.89	0.64	0.59
	LR	0.1956	0.1709	0.8852	0.8708	0.94	0.98	0.56	0.38
sRQS	B-GP	0.2230	0.1985	1.2683	1.3497	0.95	0.99	0.61	0.48
	KLR	0.2152	0.1734	1.2052	1.3843	0.99	0.99	0.51	0.41
	SLR	0.1935	0.2181	1.2027	1.2302	1.00	0.96	0.57	0.60
	LR	0.1840	0.2012	1.2543	1.3174	0.96	0.97	0.56	0.50
Ada-sIDS	B-GP	0.2016	0.2060	0.9730	1.0051	0.97	0.96	0.34	0.33
Ada-sRQS		0.3025	0.2398	1.4612	1.5013	0.99	1.00	0.16	0.22
Det-IDS	$\bar{p}$	0.3692	0.2996	0.0196	0.0773	0.01	0.05	26.58	7.85
	$\hat{p}_{\mathcal{L}_0}$	0.4377	0.3576	0.0769	0.2068	0.03	0.13	21.78	6.40
RQS	$\bar{p}$	0.4422	0.2528	0.0000	0.0038	0.00	0.01	31.74	19.59
	$\hat{p}_{\mathcal{L}_0}$	0.4099	0.2735	0.0384	0.0202	0.01	0.03	32.60	17.63

tends to yield better learning rates about  $p(\cdot)$  (and therefore about  $X^*$ ) compared to  $a = 100$ , as measured by the average KL divergence. Adaptive batching generally under-performs, especially sRQS that frequently uses excessive batch sizes (see Fig. 3) far from the root, and hence does not exploit sufficiently. At the same time, adaptive batching achieves the lowest KL divergence. It remains an open question how to best select the thresholding sequence.

**Empirical results for the exponential and cubic test functions.** Tables 2 and 3 show the performance metrics for the test functions  $h_2$  and  $h_3$ , respectively. Results are largely similar. As for  $h_1$  we observe a large improvement in performance relative to non-spatial G-pBA, especially in terms of the coverage probability  $\hat{c}$ , which was improved from  $\hat{c}_{0.95} \approx 0$  (meaning the algorithm fails completely in providing a CI for  $x^*$ ) to the actual nominal CI coverage value, see the right-most columns of Table 2 and Table 3. In terms of sampling policies, we note that sIDS again outperforms sRQS in terms of average absolute residuals and length of CI for both  $h_2$  and  $h_3$ . Furthermore, polynomial logistic regression (LR) continues to be the best surrogate choice combined with fixed batch of  $a_n \equiv 100$  (i.e., using a total of  $N = 200$  design points) implying a preference for *exploration* in these harder problems. We note that B-GP performs worse, especially for  $h_2$ , possibly due to the non-smoothness of  $\theta_2(\cdot)$  at the root (cf. Figure 2). Because B-GP assumes a smooth response surface it fails to properly capture such “cusp” that calls for a spatially non-stationary covariance structure.

#### 4.4 Evaluating the Quality of the Design

To focus on the sampling aspect of spatial G-PBA, we examine more closely the designs  $x_{1:n}^{(a,\eta)}$  obtained from implementing the sampling policy  $\eta$  and batch size  $a$ . For this analysis we return to  $h_1$  in (31) and use B-GP as the representative surrogate with a fixed batch size

Table 2: Performance Monte-Carlo metrics for the test function  $h_2$  at  $T = 20,000$ .

Policy $\eta$	$\hat{p}$	$\hat{r}(f_T)$ ( $10^{-2}$ )		$\hat{l}_{0.95}(f_T)$ ( $10^{-2}$ )		$\hat{c}_{0.95}(f_T)$		$\hat{D}(f_T; g_T)$	
		$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$
sIDS	B-GP	0.5330	0.4439	1.1043	1.1666	0.45	0.53	1.61	1.41
	KLR	0.4537	0.4098	0.7974	0.8751	0.46	0.44	3.23	2.79
	SLR	0.4352	0.4076	1.0995	1.2555	0.67	0.80	1.21	0.86
	LR	0.3814	0.4128	1.0641	1.1795	0.60	0.56	2.08	1.61
sRQS	B-GP	0.4817	0.5162	1.4630	1.6117	0.70	0.73	1.35	1.42
	KLR	0.4602	0.5440	1.0787	1.3580	0.57	0.60	2.68	1.91
	SLR	0.3956	0.4250	1.6651	1.7434	0.82	0.83	0.93	0.86
	LR	0.4653	0.5143	1.7161	1.5902	0.79	0.67	1.40	1.34
Ada-sIDS	B-GP	0.5095	0.4883	1.3638	1.2129	0.49	0.52	1.62	1.83
Ada-sRQS		0.5586	0.5088	1.6736	1.7562	0.74	0.77	1.00	0.97
Det-IDS	$\bar{p}$	0.6848	0.4418	0.0211	0.1158	0.02	0.07	27.40	8.85
	$\hat{p}_{\mathcal{L}_0}$	0.6570	0.5756	0.0090	0.3639	0.01	0.18	24.10	6.99
RQS	$\bar{p}$	0.7075	0.4846	0.0428	0.0649	0.02	0.07	27.88	12.36
	$\hat{p}_{\mathcal{L}_0}$	0.8442	0.4686	0.0477	0.0527	0.02	0.06	24.43	10.71

of  $a = 100$ . To judge the quality of  $x_{1:n}^{(a,\eta)}$  for the SRF, we compute the resulting exact posterior  $g_n^{(a,\eta)}(\cdot|x_{1:n}^{(a,\eta)})$  and evaluate the resulting absolute residual  $|\text{median}(g_n^{(a,\eta)}) - x^*|$  and corresponding length of  $(1-\alpha)\%$ -CI. A design that is better quantifying uncertainty about  $X^*$  should have lower residuals and lower CI. We then benchmark the resulting metrics against the following *baseline* schemes which utilize the true  $p(x)$  (and therefore the actual posterior density  $g_n$ ):

$$x_{n+1} := \arg \max_{x \in (0,1)} \mathcal{I}(x, g_n; p(x), a) \quad (\text{IDS})$$

$$x_{n+1} := G_n^{-1}(U_{n+1}), \quad U_{n+1} \sim \text{Unif}(0, 1); \quad (\text{RQS})$$

$$x_{n+1} \sim \text{Unif}(0, 1). \quad (\text{Unif})$$

The sampling strategy (IDS) is optimal in the sense of maximizing the expected KL distance between  $g_n$  and  $g_{n+a}$ , and hence we use it as an upper bound on performance; (Unif) is a *passive policy* used as a lower bound. To make the baseline policies comparable with the spatial G-PBA strategies, we implement the former with batched sampling using the transition function (9) and  $a = 100$ . We also match the initialization step, employing  $N_0 = T_0/a_0$  equidistant locations  $x_{1:N_0}$  (with  $T_0 = 5000$ ) to construct  $g_{T_0}$ , from which (IDS), (RQS) and (Unif) are implemented.

Figure 5 visualizes the results. We observe that sIDS is the sampling policy which best approximates the true IDS, and that all G-PBA strategies significantly outperform the (Unif) baseline strategy. Interestingly, both randomized and information-directed policies appear to have similar asymptotic performance in terms of average residuals and CI length.

Table 3: Performance Monte-Carlo metrics for the test function  $h_3$  at  $T = 20,000$ .

Policy $\eta$	$\hat{p}$	$\hat{r}(f_T)$ ( $10^{-2}$ )		$\hat{l}_{0.95}(f_T)$ ( $10^{-2}$ )		$\hat{c}_{0.95}(f_T)$		$\hat{D}(f_T; g_T)$	
		$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$	$a_0=100$	$a_0=250$
sIDS	B-GP	4.3661	4.2959	8.2188	9.7472	0.57	0.64	1.87	1.57
	KLR	4.3403	4.5771	13.0221	12.0456	0.76	0.76	1.32	1.36
	SLR	4.4470	4.6160	7.3444	7.7896	0.49	0.46	2.35	2.23
	LR	3.7645	3.6936	10.6028	10.5738	0.71	0.70	1.52	1.39
sRQS	B-GP	4.1913	4.0209	10.7298	10.8774	0.67	0.67	1.70	1.45
	KLR	3.9131	3.7121	14.2680	14.0897	0.81	0.84	1.17	0.98
	SLR	4.0451	4.1825	10.3663	10.2469	0.69	0.68	1.79	2.03
	LR	3.6513	4.1276	12.9502	11.5623	0.80	0.66	1.27	1.31
Ada-sIDS	B-GP	4.1540	4.2334	11.1918	11.3152	0.68	0.67	1.11	1.11
Ada-sRQS		4.1874	4.0915	11.6808	12.9052	0.67	0.76	1.39	1.05
Det-IDS	$\bar{p}$	5.3257	4.8835	0.0187	0.4446	0.00	0.03	33.34	11.93
	$\hat{p}_{\mathcal{L}_0}$	5.7587	5.3403	0.0001	0.3862	0.00	0.01	27.94	10.01
RQS	$\bar{p}$	5.1556	4.7262	0.0000	0.2978	0.00	0.01	37.77	15.16
	$\hat{p}_{\mathcal{L}_0}$	5.3406	4.7325	0.0001	0.7267	0.00	0.02	31.77	12.68

## 5 Case-Study: Root-Finding for Optimal Stopping

In this section, we apply the spatial G-PBA Algorithm 2 to solve the root-finding sub-routine for pricing a Bermudan Put option (Ludkovski, 2018). Valuing a Bermudan option with maturity  $\tilde{T}$  is equivalent to solving an optimal stopping problem

$$V(t, x) := \sup_{\tilde{T} \geq \tau \geq t, \tau \in \mathcal{S}} \mathbb{E}[H(\tau, X_\tau) | X_t = x],$$

where  $(X_t)$  is a stochastic process and  $H(t, x)$  is the *reward function*. Assuming the classical discretized Black-Scholes model with time step  $\Delta t$  we have that  $(X_t)$  is a log-normal random walk and  $H(t, x) := e^{-rt}(K^{Put} - x)_+$ , where  $K^{Put}$  is the strike price and  $r > 0$  is the interest rate. In this setting, it is well-known that there is a unique *exercise boundary*  $x^*(t) \leq K^{Put}$ , and one should exercise at the first  $t$  when  $X_t$  drops below this boundary. Finding the exercise boundaries  $\{x^*(t)\}$  reduces to solving a sequence of SRFPs, that is, pricing the Bermudan Put is equivalent to finding the solutions  $x^*(t)$  of the equation  $h(t, x) = 0$  for  $t = \tilde{T} - \Delta t, \tilde{T} - 2\Delta t, \dots, 0$ , where  $h(t, x) := V(t, x) - H(t, x)$  is the *timing value*.

The Longstaff and Schwartz (2001) method (LSM) recursively builds noisy simulators for  $h(t, x)$  by generating forward paths  $x_{t:\tilde{T}}$  of the state process  $(X_t)$  and computing corresponding path-wise stopping times  $\tau \equiv \tau(t + \Delta t, x_{t:\tilde{T}}) = \min\{s > t : x_s \leq x^*(s)\} \wedge \tilde{T}$ . Namely, the pathwise difference  $Z_t^{LSM}(x_t) := H(\tau, x_\tau) - H(t, x_t)$  between future and immediate reward over the given trajectory  $x_{t:\tilde{T}}$  satisfies  $\mathbb{E}[Z_t(x_t)] = h(x_t, t)$ , matching the structure of the oracle (2). The random component  $\epsilon(t, x)$  arises intrinsically from the randomness in the  $X$ -trajectory.

We implement the spatial G-PBA for the Bermudan Put oracle using  $K^{Put} = 40, r = 0.06, \sigma = 0.25, \tilde{T} = 1$  and  $\Delta t = 0.04$ , restricting the root-finding to the “in-the-money” domain






Figure 5: Comparison of spatial sampling policies with respect to baseline policies using the true posteriors  $g_T$ . *Left panel*: average absolute residuals,  $\text{median}(g_T)$ , against wall-clock time  $T$ ; *Right panel*: average length of 95% CI  $l_{0.95}(g_T)$  against  $T$ . Replication amounts fixed at  $a_n = 250 \forall n$ .

$x \in (25, 40)$ . Following the discussion in Section 5 in Rodriguez and Ludkovski (2017), to conform to the assumption of a symmetric noise distribution, we employ pre-averaging that considers the sign of an average of  $R > 1$  oracle evaluations:

$$\bar{Y}_R^{LSM}(x) := \text{sign } \bar{Z}_R(x), \quad \bar{Z}_R(x) := \frac{1}{R} \sum_{r=1}^R Z_r^{LSM}(x). \quad (39)$$

The role of pre-averaging is to alleviate statistical anomalies of  $\epsilon(\cdot)$  via the Central Limit Theorem; the resulting oracle accuracy for this case-study is:  $p_R^{LSM}(x) := \mathbb{P}(\bar{Y}_R^{LSM}(x) = \text{sign}\{x^* - x\})$ . Below we continue to record the clock-time  $T_n$  based on underlying oracle evaluations (rather than the pre-averaged ones).

Due to the non-standard noise distribution and very low signal-to-noise ratio, this is a difficult root-finding problem; in particular since we keep the simulation budget to  $T = 20,000$ . We implement the simulator (39) with  $R = 25$ , which is roughly the minimal level of pre-averaging required to alleviate the skew of  $Z^{LSM}$  (Rodriguez and Ludkovski, 2017), and a batch size of  $a_n \equiv 500$ . Thus, the *effective number of replicates* at each location is  $\tilde{a} := a/R = 20$ . All surrogates are initialized with  $\hat{\varphi}_{N_0}$  based on  $T_0 = 0.25 \times T = 5000$  function evaluations using  $N_0 := T_0/a = 10$  design sites over the interval  $(25, 40)$ . The adaptive replication scheme (23) is implemented with  $a_{n+1}^\nu = \min\{\hat{a}_{n+1}^\nu, a/R\}$ , where  $a_0^\nu := 1$  and  $\nu_n^{Put} := 0.5/n$ , so that the maximum number of oracle evaluations is at most  $a = 500$  per querying location.

## 5.1 Results

Table 4 shows the average residuals, length of CI, and coverage probability of the spatial G-PBA schemes compared against the baseline root location  $\hat{x}^*(t) \simeq 35.1249$  (here  $t = 0.6$ )

Table 4: Spatial G-PBA Performance in the Bermudan Put SRFP with simulation budget  $T = 20,000$ . All metrics are averages across  $MC = 100$  macro-replications of the algorithms.

$\eta$	$\hat{p}$	$\hat{r}(f_T)$	$\hat{l}_{0.95}(f_T)$	$\hat{c}_{0.95}(f_T)$ (in %)
sIDS	B-GP	0.3210	0.8903	69.00
	KLR	0.3598	0.6351	53.33
	SLR	0.3158	0.9878	77.00
	LR	0.2753	1.0687	88.50
sRQS	B-GP	0.2988	1.4064	86.00
	KLR	0.3121	0.8209	62.00
	SLR	0.3180	1.2005	74.50
	LR	0.2913	1.4039	90.50
Ada-sIDS	B-GP	0.2225	0.6944	80.00
Ada-sRQS		0.3011	1.1121	76.47

found in Rodriguez and Ludkovski (2017). This time, adaptive replication with the one-step sIDS policy (29) is the best-performing scheme. One reason could be that it allows for more sampling locations (median number of sampling locations was  $\text{median}(N_T) = 55$ , as apposed to 40 for the fixed  $a_n$  schemes). Among the rest, sIDS policy coupled with the polynomial logistic regression model (LR) also performs very well, consistent with our findings in Section 4. Relative to the non-spatial PBA in Rodriguez and Ludkovski (2017) two important improvements are noted: (i) much better coverage probabilities, indicating the gains in learning  $p(\cdot)$  and hence maintaining a reliable knowledge state; (ii) residuals below 0.25 while they used to be about 0.35.

## 6 Conclusion

We have developed a family of numerical schemes that extend generalized probabilistic bisection (Rodriguez and Ludkovski, 2017) by modeling the unknown oracle accuracy  $p(\cdot)$  through a spatial surrogate based on non-parametric binomial regression. The spatial structure yields two key benefits: (I) given the surrogate, the IDS criterion  $\mathcal{I}$  can be predicted for any  $x$ , allowing direct optimization of next querying site selection like in standard PBA; (II) employing a GP surrogate quantifies the predictive uncertainty of additional samples and hence allows for adaptive batching schemes. Adaptive replication allows to automatically fine-tune exploration by reducing replication amounts in regions where  $p(\cdot)$  is already learned well. Our numeric experiments confirm the advantages of Spatial G-PBA relative to the original proposals in Rodriguez and Ludkovski (2017) with the new algorithm inducing more accurate root estimates and better quantifying the posterior uncertainty about  $x^*$ .

Looking ahead, one motivation for considering PBA in the context of SRFP is its Bayesian flavor that allows in particular to apply informative priors  $f_0$  as a way to warm-start the root

search. This offers one way to lift PBA, which is intrinsically limited to a one-dimensional setting, to higher dimensions. The analogue of SRF in two-dimensions is noisy (zero-)contour-finding, which can be viewed as a collection of root-finding problems in the first coordinate  $x_1$ , indexed by the second coordinate  $x_2$ . Assuming the zero-contour is smooth, one may then try to solve for a few  $x^*(x_2)$  and then “connect the dots” through interpolation (or a further surrogate model). Such searches can be made efficient with G-PBA by using  $f_N(\cdot; x_2)$  as a basis for an informative prior  $f_0(\cdot; x_2')$  at a new  $x_2'$ . We leave such investigations to future research.

## Acknowledgments

Rodriguez is partially supported by the National Science and Technology Council of Mexico (CONACYT) and University of California Institute for Mexico and the United States (UCMEXUS) under grant CONACYT-216011. Ludkovski is partially supported by NSF DMS-1521743. We are also grateful to the UCSB Center for Scientific Computing from the CNSI and MRL: an NSF MRSEC (DMR-1720256).

# Appendices

## Appendix A Binomial GPs and Laplace Approximation

**Binomial log-likelihood Gradient and Hessian.** We use the Bernoulli link function  $\Theta(\varphi) = (1 + e^{-\varphi})^{-1}$  which implies that conditional on  $\varphi_i = \varphi(x_i)$ , the number of positive responses  $B_i := \sum_{j=1}^{a_i} 1_{\{Z_j > 0\}}$  follows a binomial distribution  $B_i \sim \text{Bin}(a_i, \Theta(\varphi_i))$  with log-likelihood function (in the latent  $\varphi_{1:n}$ ):

$$l(\varphi_{1:n}) \equiv \log p(B_{1:n} | \varphi_{1:n}, a_{1:n}) = \sum_{j=1}^n \left\{ \log \binom{a_i}{B_i} + B_i \log \theta(\varphi_i) + (a_i - B_i) \log [1 - \theta(\varphi_i)] \right\}.$$

Since  $\Theta'(\varphi) := \Theta(\varphi)[1 - \Theta(\varphi)]$ , the corresponding gradient vector  $\mathbf{u}_n(\varphi_{1:n}) := \nabla l(\varphi_{1:n})$  is given by

$$\begin{aligned} \partial_{\varphi_i} l(\varphi_{1:n}) = u_i &= B_i \frac{\Theta'(\varphi_i)}{\Theta(\varphi_i)} - (a_i - B_i) \frac{\Theta'(\varphi_i)}{1 - \Theta(\varphi_i)} = B_i [1 - \Theta(\varphi_i)] + (B_i - a_i) \Theta(\varphi_i) \\ &= B_i - a_i \theta(\varphi_i), \quad i = 1, \dots, n, \end{aligned} \tag{40}$$

which is a function of  $\varphi_i$  only. Differentiating  $\mathbf{u}_n$  again therefore yields the diagonal  $n \times n$  Hessian matrix  $\mathbf{W}_n(\varphi_{1:n}) = -\Delta \log p(B_{1:n} | \varphi_{1:n}, a_{1:n})$  as specified in (18).

**Normal Approximation to the Joint Posterior Distribution.** By Bayes’ rule the posterior  $p(\varphi_{1:n} | \mathcal{D}_n)$  is proportional to the Binomial likelihood  $p(B_{1:n} | \varphi_{1:n}, a_{1:n}, x_{1:n})$  times the zero-mean GP prior  $p(\varphi_{1:n} | x_{1:n})$ . Taking the log of the unnormalized joint posterior we

obtain

$$\begin{aligned}\mathcal{L}(\varphi_{1:n}) &\propto \log p(B_{1:n}|\varphi_{1:n}, a_{1:n}, x_{1:n}) + \log p(\varphi_{1:n}|x_{1:n}) \\ &:= \log p(B_{1:n}|\varphi_{1:n}, a_{1:n}, x_{1:n}) - \frac{1}{2}\varphi_{1:n}^T \mathbf{K}_n^{-1} \varphi_{1:n} - \frac{1}{2} \log |\mathbf{K}_n| - \frac{n}{2} \log 2\pi.\end{aligned}\quad (41)$$

Denote by  $\hat{\varphi}_n := \arg \max_{\varphi_n} \mathcal{L}(\varphi_n) = \arg \max_{\varphi_n} p(\varphi_n|\mathcal{D}_n)$ . Expanding  $\mathcal{L}(\cdot)$  around  $\hat{\varphi}_n$  gives  $\mathcal{L}(\varphi_n) = \mathcal{L}(\hat{\varphi}_n) + \frac{1}{2}(\varphi_n - \hat{\varphi}_n)^T [\Delta \mathcal{L}(\hat{\varphi}_n)](\varphi_n - \hat{\varphi}_n) + \dots$ ; where the linear term in the expansion is zero because the log-posterior density has zero derivative at its mode. As discussed in Gelman et al. (2014), the remainder terms of higher order fade in importance relative to the quadratic term when  $\varphi_n$  is close to  $\hat{\varphi}_n$  and the sample size  $n$  is large. Taking first and second partial derivatives of  $\mathcal{L}(\varphi_{1:n})$  with respect to  $\varphi_{1:n}$  and combining with (40)-(18) we obtain:

$$\nabla \mathcal{L}(\varphi_{1:n}) = \mathbf{u}_n(\varphi_{1:n}) - \mathbf{K}_n^{-1} \varphi_{1:n}, \quad (42)$$

$$\Delta \mathcal{L}(\varphi_{1:n}) = -\mathbf{W}_n(\varphi_{1:n}) - \mathbf{K}_n^{-1}; \quad (43)$$

At the mode of  $\mathcal{L}(\varphi_{1:n})$  we have

$$\nabla \mathcal{L}(\hat{\varphi}_n) = \mathbf{0} \quad \Rightarrow \quad \hat{\varphi}_n = \mathbf{K}_n \mathbf{u}_n(\hat{\varphi}_n) \quad (44)$$

as a self-consistent nonlinear equation determining  $\hat{\varphi}_n$ . In order to solve (44), an iterative procedure based on classical Newton-Raphson search is employed.

Next, the Hessian of the score function  $\Delta \mathcal{L}(\varphi_n)$  is interpreted as the inverse covariance matrix, leading to the Gaussian approximation  $q(\cdot|\mathcal{D}_n)$  to the true posterior  $p(\cdot|\mathcal{D}_n)$

$$q(\cdot|\mathcal{D}_n) := \mathbf{N}(\cdot; \hat{\varphi}_n, \Sigma_n) \quad \text{where} \quad \Sigma_n \equiv (\mathbf{K}_n^{-1} + \mathbf{W}_n(\hat{\varphi}_n))^{-1}. \quad (45)$$

**Predictive distribution.** The approximated predictive pdf  $\varphi_*(x)$  at a test location  $x \in (0, 1)$  is Gaussian  $\varphi_*(x) \sim \mathbf{N}(m_n(x; \hat{\varphi}_n), s_n^2(x; \hat{\varphi}_n))$  with the mean  $m_n(x; \hat{\varphi}_n)$  given by:

$$\begin{aligned}m_n(x; \hat{\varphi}_n) &:= \int \mathbb{E}[\varphi(x)|\tilde{\varphi}_{1:n}] p(\tilde{\varphi}_{1:n}|\mathcal{D}_n) d\tilde{\varphi}_{1:n} \\ &= \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \int \tilde{\varphi}_{1:n} p(\tilde{\varphi}_{1:n}|\mathcal{D}_n) d\tilde{\varphi}_{1:n} \\ &= \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \mathbb{E}[\varphi_{1:n}|\mathcal{D}_n] \simeq \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \hat{\varphi}_n;\end{aligned}\quad (46)$$

where  $\boldsymbol{\kappa}_n^T \equiv (\kappa(x_1, x), \dots, \kappa(x_n, x))$ , matching (19a). Likewise, the approximated predictive variance,  $s_n(x; \hat{\varphi}_n) \equiv \text{Var}(\varphi(x)|\mathcal{D}_n, \hat{\varphi}_n, x)$ , is given by (cf. (19b)):

$$\begin{aligned}s_n(x; \hat{\varphi}_n) &:= \mathbb{E}[\text{Var}(\varphi(x)|\varphi_{1:n}, x_{1:n})|\mathcal{D}_n] + \text{Var}(\mathbb{E}[\varphi(x)|\varphi_{1:n}, x_{1:n}, x]|\mathcal{D}_n) \\ &= \mathbb{E}[\tau^2 - \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \boldsymbol{\kappa}_n|\mathcal{D}_n] + \text{Var}(\boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \varphi_{1:n}|\mathcal{D}_n) \\ &= \tau^2 - \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \boldsymbol{\kappa}_n + \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \text{Var}(\varphi_{1:n}|\mathcal{D}_n) \mathbf{K}_n^{-1} \boldsymbol{\kappa}_n \\ &\simeq \tau^2 - \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} \boldsymbol{\kappa}_n + \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} (\mathbf{K}_n^{-1} + \mathbf{W}_n(\hat{\varphi}_n))^{-1} \mathbf{K}_n^{-1} \boldsymbol{\kappa}_n \\ &= \tau^2 - \boldsymbol{\kappa}_n^T (\mathbf{K}_n + \mathbf{W}_n(\hat{\varphi}_n))^{-1} \boldsymbol{\kappa}_n,\end{aligned}$$

where the last line is true via the matrix inversion lemma applied to  $(\mathbf{K}_n + \mathbf{W}_n(\hat{\varphi}_n))^{-1}$ :

$$\begin{aligned}\boldsymbol{\kappa}_n^T (\mathbf{K}_n + \mathbf{W}_n(\hat{\varphi}_n))^{-1} \boldsymbol{\kappa}_n &= \boldsymbol{\kappa}_n^T \{ \mathbf{K}_n^{-1} - \mathbf{K}_n^{-1} (\mathbf{K}_n^{-1} + \mathbf{W}_n(\hat{\varphi}_n))^{-1} \mathbf{K}_n^{-1} \} \boldsymbol{\kappa}_n \\ &= \boldsymbol{\kappa}_n^T \mathbf{K}_n^{-1} - \boldsymbol{\kappa}_n^T(x) \mathbf{K}_n^{-1} (\mathbf{K}_n^{-1} + \mathbf{W}_n(\hat{\varphi}_n))^{-1} \mathbf{K}_n^{-1} \boldsymbol{\kappa}_n.\end{aligned}$$

## Appendix B Predictive Variance Decomposition for Binomial GPs under Laplace Approximation (Theorem 2.2)

*Proof.* Set  $\hat{\varphi}_{n+1} \equiv (\hat{\varphi}_{1;n+1}, \dots, \hat{\varphi}_{n+1;n+1})$  to be the  $(n+1)$ -dimensional estimated mode based on training data  $\mathcal{D}_{n+1}$  obtained at locations  $x_{1:n+1}$ ; and let

$$\hat{\mathbf{W}}_{n+1;n+1} := \text{diag}\{\hat{w}_{1;n+1}, \dots, \hat{w}_{n+1;n+1}\}, \quad \hat{w}_i \equiv w_i(\hat{\varphi}_{i;n+1})$$

be the Hessian matrix (18) evaluated at  $\hat{\varphi}_{n+1}$ . Then, we have that the  $(n+1) \times (n+1)$  covariance matrix  $\Sigma_{n+1} \equiv (\mathbf{K}_{n+1} + \hat{\mathbf{W}}_{n+1})^{-1}$  of the joint approximated posterior (45) can be partitioned as:

$$\Sigma_{n+1} = \begin{pmatrix} \Sigma_{1:n;n+1} & \boldsymbol{\kappa}_n^* \\ (\boldsymbol{\kappa}_n^*)^T & \tau^2 + \hat{w}_{n+1;n+1}^{-1} \end{pmatrix}, \quad (47)$$

where  $\boldsymbol{\kappa}_n^* := (\kappa(x_1, x_{n+1}), \dots, \kappa(x_n, x_{n+1}))^T$  is a  $n \times 1$  column vector of covariances of  $\varphi_{1:n}$  against  $\varphi_{n+1}$ , and  $\tau^2 = \kappa(x_{n+1}, x_{n+1})$  is scalar. Applying the Matrix Inversion Theorem Henderson and Searle (1981), the inverse of (47) is:

$$\begin{aligned} \Sigma_{n+1}^{-1} &= \begin{pmatrix} \Sigma_{1:n;n+1}^{-1} + (\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*) (\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*)^T a^{-1} & -\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^* a^{-1} \\ -(\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*)^T a^{-1} & a^{-1} \end{pmatrix} \\ &= \begin{pmatrix} \Sigma_{n;n+1}^{-1} & 0 \\ 0 & 0 \end{pmatrix} + a^{-1} \begin{pmatrix} (\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*) (\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*)^T & -\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^* \\ -(\Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*)^T & 1 \end{pmatrix}, \end{aligned}$$

where the scalar is  $a := (\tau^2 + \hat{w}_{n+1;n+1}^{-1}) - (\boldsymbol{\kappa}_n^*)^T \Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*$   
 $= \hat{w}_{n+1;n+1}^{-1} + (\tau^2 - (\boldsymbol{\kappa}_n^*)^T \Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*) = \hat{w}_{n+1;n+1}^{-1} + s_n^2(x_{n+1}; \hat{\varphi}_{1:n,n+1})$ .

Substituting the expression for  $\Sigma_{n+1}^{-1}$  obtained above in the predictive variance formula (19b), we have that the posterior predictive variance given the dataset  $\mathcal{D}_{n+1}$  is:

$$\begin{aligned} s_{n+1}^2(x_{n+1}; \hat{\varphi}_{n+1}) &:= \tau^2 - (\boldsymbol{\kappa}_n^*)^T \Sigma_{n+1}^{-1} \boldsymbol{\kappa}_n^* \\ &= \tau^2 - \mathbf{u}^T \left\{ \begin{pmatrix} \Sigma_{1:n;n+1}^{-1} & 0 \\ 0 & 0 \end{pmatrix} + a^{-1} \begin{pmatrix} \mathbf{v}\mathbf{v}^T & -\mathbf{v} \\ -\mathbf{v}^T & 1 \end{pmatrix} \right\} \mathbf{u} \\ &= \tau^2 - \mathbf{u}^T \begin{pmatrix} \Sigma_{n;n+1}^{-1} & 0 \\ 0 & 0 \end{pmatrix} \mathbf{u} - a^{-1} \mathbf{u}^T \begin{pmatrix} \mathbf{v}\mathbf{v}^T & -\mathbf{v} \\ -\mathbf{v}^T & 1 \end{pmatrix} \mathbf{u} \\ &= \left[ \tau^2 - (\boldsymbol{\kappa}_n^*)^T \Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^* \right] - a^{-1} [b^2 - b\tau^2 - \tau^2 b + \tau^4], \quad b_{1 \times 1} \equiv (\boldsymbol{\kappa}_n^*)^T \mathbf{v} \\ &= s_n^2(x_{n+1}; \hat{\varphi}_n) - a^{-1} (\tau^2 - b)^2; \end{aligned}$$

where we set  $\mathbf{v}_{n \times 1} \equiv \Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*$  and let  $\mathbf{u}_{(n+1) \times 1} \equiv (\boldsymbol{\kappa}_n^* \ \tau^2)^T$  be the concatenation of the vector  $\boldsymbol{\kappa}_n^*$  and the scalar  $\tau^2$ . Simplifying, we finally get:

$$\begin{aligned} s_{n+1}^2(x_{n+1}; \hat{\varphi}_{n+1}) &= s_n^2(x_{n+1}; \hat{\varphi}_{1:n,n+1}) - \frac{1}{a} (\tau^2 - (\boldsymbol{\kappa}_n^*)^T \Sigma_{1:n;n+1}^{-1} \boldsymbol{\kappa}_n^*)^2 \\ &= s_n^2(x_{n+1}; \hat{\varphi}_{1:n,n+1}) - \frac{(s_n^2(x_{n+1}; \hat{\varphi}_{1:n,n+1}))^2}{\hat{w}_{n+1}^{-1} + s_n^2(x_{n+1}; \hat{\varphi}_{n;n+1})} \end{aligned}$$

$$\begin{aligned}
&= \frac{s_n^2(x_{n+1}; \hat{\varphi}_{1:n;n+1}) \hat{w}_{n+1;n+1}^{-1}}{\hat{w}_{n+1;n+1}^{-1} + s_n^2(x_{n+1}; \hat{\varphi}_{1:n;n+1})} \\
&= \left( \frac{1}{s_n^2(x_{n+1}; \hat{\varphi}_{1:n;n+1})} + \frac{1}{\hat{w}_{n+1;n+1}^{-1}} \right)^{-1}.
\end{aligned}$$

Finally, we notice that  $\hat{w}_{n+1;n+1} := a_{n+1} \Theta(\hat{\varphi}_{n+1;n+1}) (1 - \Theta(\hat{\varphi}_{n+1;n+1}))$  which leads to (21).  $\square$

## References

- B. Ankenman, B. L. Nelson, and J. Staum. Stochastic kriging for simulation metamodeling. *Operations research*, 58(2):371–382, 2010.
- D. Azzimonti, J. Bect, C. Chevalier, and D. Ginsbourger. Quantifying uncertainties on excursion sets under a Gaussian random field prior. *SIAM/ASA J. Uncertainty Quantification*, 4(1):850–874, 2016.
- M. Binois, R. B. Gramacy, and M. Ludkovski. Practical heteroskedastic Gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics*, 0(ja):1–41, 2018a.
- M. Binois, J. Huang, R. B. Gramacy, and M. Ludkovski. Replication or exploration? Sequential design for stochastic simulation experiments. *Technometrics*, 0(ja):1–43, 2018b.
- X. Chen and Q. Zhou. Sequential design strategies for mean response surface metamodeling via stochastic kriging with adaptive exploration and exploitation. *European Journal of Operational Research*, 262(2):575–585, 2017.
- C. Chevalier, J. Bect, D. Ginsbourger, E. Vazquez, V. Picheny, and Y. Richet. Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set. *Technometrics*, 56(4):455–465, 2014.
- T. M. Cover and J. A. Thomas. *Elements of Information theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006. ISBN 0471241954.
- P. I. Frazier, S. G. Henderson, and R. Waeber. Probabilistic bisection converges almost as quickly as stochastic approximation. Technical report, arXiv preprint arXiv:1612.03964, 2016.
- J. Friedman, T. Hastie, and R. Tibshirani. *The Elements of Statistical Learning*, volume 1. Springer series in Statistics New York, 2001.
- A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian data analysis*, volume 2. Taylor & Francis, 2014.

- H. V. Henderson and S. R. Searle. On deriving the inverse of a sum of matrices. *SIAM Review*, 23(1):53–60, 1981.
- B. Jedynak, P. I. Frazier, R. Sznitman, et al. Twenty questions with noise: Bayes optimal policies for entropy loss. *Journal of Applied Probability*, 49(1):114–136, 2012.
- S. G. Johnson. The NLOpt nonlinear-optimization package, 2014.
- D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993.
- B. Kamiński. A method for the updating of stochastic kriging metamodels. *European Journal of Operational Research*, 247(3):859–866, 2015.
- A. Kapoor, K. Grauman, R. Urtasun, and T. Darrell. Active learning with Gaussian processes for object categorization. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007.
- F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: a simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147, 2001.
- M. Ludkovski. Kriging metamodels and experimental design for Bermudan option pricing. *Journal of Computational Finance*, to Appear, 2018.
- T. P. Minka. Expectation propagation for approximate Bayesian inference. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 362–369. Morgan Kaufmann Publishers Inc., 2001.
- H. Nickisch and C. E. Rasmussen. Approximations for binary Gaussian process classification. *Journal of Machine Learning Research*, 9(Oct):2035–2078, 2008.
- R. Pasupathy and S. Kim. The Stochastic Root-Finding Problem: Overview, Solutions, and Open Questions. *ACM Trans. Model. Comput. Simul.*, 21(3):19:1–19:23, Feb. 2011. ISSN 1049-3301.
- W. B. Powell and I. O. Ryzhov. *Optimal learning*, volume 841. John Wiley & Sons, 2012.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016. URL <https://www.R-project.org/>.
- S. Rodriguez and M. Ludkovski. Generalized Probabilistic Bisection for Stochastic Root-Finding. *arXiv preprint arXiv:1711.00843*, 2017.
- D. Russo and B. Van Roy. An information-theoretic analysis of Thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016.

- M. Tesch, J. Schneider, and H. Choset. Expensive function optimization with stochastic binary outcomes. In *International Conference on Machine Learning*, pages 1283–1291, 2013.
- J. Vanhatalo, J. Riihimäki, J. Hartikainen, P. Jylänki, V. Tolvanen, and A. Vehtari. GPstuff: Bayesian modeling with Gaussian processes. *Journal of Machine Learning Research*, 14 (Apr):1175–1179, 2013.
- R. Waeber. *Probabilistic Bisection Search for Stochastic Root-Finding*. PhD thesis, Cornell University, 2013.
- R. Waeber, P. I. Frazier, and S. G. Henderson. A Bayesian approach to stochastic root finding. In *Proceedings of the 2011 Winter Simulation Conference (WSC)*, pages 4033–4045. IEEE, 2011.
- R. Waeber, P. I. Frazier, and S. G. Henderson. Bisection search with noisy responses. *SIAM Journal on Control and Optimization*, 51(3):2261–2279, 2013.
- Z. Wang, B. Zhou, and S. Jegelka. Optimization as estimation with Gaussian processes in bandit settings. In *Artificial Intelligence and Statistics*, pages 1022–1031, 2016.
- C. K. Williams and D. Barber. Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1342–1351, 1998.
- C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. the MIT Press, 2006.
- S. N. Wood. mgcv: GAMs and generalized ridge regression for R. *R News*, 1(2):20–25, 2001.
- J. Zhu and T. Hastie. Kernel logistic regression and the import vector machine. *Journal of Computational and Graphical Statistics*, 14(1):185–205, 2005.