

APPROXIMATE GAUSSIAN PROCESS REGRESSION WITH PAIRWISE COMPARISON DATA

Efe Sertkaya
Ilya O. Ryzhov

Robert H. Smith School of Business
University of Maryland
7699 Mowatt Lane
College Park, MD 20742, USA

ABSTRACT

We use approximate Bayesian inference, together with Gaussian process regression, to create a new estimator for an unknown function in a situation where we can only observe pairwise comparisons of function values at different inputs. Preliminary experimental results suggest that, although information is heavily censored in this setting, it may still be possible to learn the local and global minima of the underlying function. We discuss possible sampling criteria, and explore the performance of the “probability of improvement” strategy numerically.

1 INTRODUCTION

We consider the global optimization problem $\min_{x \in \mathcal{X}} f(x)$, where no closed-form expression for f (or its derivative) is available. In the simulation literature, it is common to assume that the function values $f(x)$ can be evaluated at specific choices of the input vector x , either exactly or with noise (Ankenman, Nelson, and Staum 2010). Given a set $\{x_n\}_{n=1}^N$ of such inputs, one may construct a function, often called a “metamodel” or “surrogate model,” that interpolates the observed values $\{f(x_n)\}_{n=1}^N$. The metamodel can be used to predict the function value at any point that has not yet been observed, as well as to estimate the minimizer of the function.

Among interpolation methodologies, perhaps the most widely used is Gaussian process (GP) regression (Rasmussen and Williams 2006), which adopts a Bayesian philosophy and models the simulationist’s uncertainty about f as a random field. This stochastic model is used regardless of whether the observations are deterministic or stochastic. However, the computation of the prediction of the unknown value $f(x)$ turns out to be nearly identical to the calculations used in the methodology of radial basis function interpolation (Ji and Kim 2014), which enables one to choose between a frequentist and Bayesian view (whichever is more convenient for the task at hand) of essentially the same basic model. GP regression is frequently used to model the unknown function in global optimization; see Kleijnen (2014) for an overview. The computer science community refers to this problem class by the name of Bayesian optimization (Garnett 2023).

Although the literature on GP regression is very rich, the vast majority of it focuses on the case where $f(x)$ is observable. In this paper, we consider a variant of the problem where the only available data consists of pairwise comparisons, i.e., we can choose x and x' and observe whether $f(x) > f(x')$, but we do not see either of the function values. To motivate this setting, suppose that \mathcal{X} represents an “innovation space” modeling the design attributes of a new product that has not yet gone into production. The unknown function represents the market response (for example, demand or revenue) to different configurations of the product. Before we finalize the design, we have the opportunity to conduct a limited number of test markets to more accurately evaluate a small number of potential configurations. The customers in the test group may not be able to meaningfully quantify their reactions to a single configuration, but it may be

easier for them to compare two choices and decide which one they like more. In decision analysis, this is known as preference elicitation (Branke et al. 2017). As in standard GP regression, we have to use this censored information to predict the values of untested designs and choose one that appears to be the best.

We explore this problem using the methodology of approximate Bayesian inference, which aims to construct computationally efficient learning models in the presence of censoring. Most standard Bayesian models make use of the property of conjugacy, which means that the prior distribution of the unknown object (before learning) and the posterior distribution (given one or more observations) belong to the same family, e.g., the normal family. GP regression is also an example of a conjugate model: we begin with a Gaussian process prior on f and assume that our observation of $f(x)$ is also normally distributed, which ensures that the posterior distribution of f is another Gaussian process. In the setting of pairwise comparisons, however, the observation is binary-valued and conjugacy no longer holds. The core idea of approximate Bayesian inference is to construct an artificial posterior distribution that belongs to the desired family, while also approximating the actual posterior in some way.

The approximation can be constructed in many ways, including moment-matching (Minka 2001b), variational bounds (Jaakkola and Jordan 2000), and optimization of Kullback-Leibler divergence (Qu et al. 2015). There is an important distinction between methods that first collect n observations (taking $n \rightarrow \infty$ in convergence proofs) and then perform one single approximation, and methods that perform a new approximation after each individual observation. The first class of methods is more amenable to analysis (Wang and Blei 2019), but more computationally intensive. The second class often runs very quickly, but is quite difficult to analyze, though it often exhibits superlative practical performance (one celebrated example being the TrueSkillTM system of Herbrich, Minka, and Graepel 2006 for competitive online games). Chen and Ryzhov (2020) was the first effort to study the asymptotic behavior of the second, sequential class of methods, and established rigorous consistency results for several known algorithms of this type, where the unknown object to be learned was a scalar or vector parameter.

This paper discusses our ongoing work on using approximate Bayesian inference in GP regression with censored observations. Using the method of moment-matching, we provide closed-form, computationally efficient expressions for the posterior mean and covariance functions given a single pairwise comparison (as a first step, we assume that the result of the comparison is noiseless). These results can be used for sequential learning: after collecting one observation, we simply view the approximate Bayesian posterior as the new “prior” and make a new pairwise comparison. Our experimental results suggest that, while the approximate posterior does not accurately estimate f as a whole, it does have the potential to correctly identify both local and global minima. Thus, while the inference technique itself performs estimation rather than optimization, it is most likely to be useful for the purpose of global optimization.

To our knowledge, this is the first sequential mechanism for approximate GP regression. Previous work in machine learning has considered GP regression with censoring, in settings such as ordinal regression (Chu et al. 2005), preference learning (Chu and Ghahramani 2005), and binary classification (Kuss et al. 2005). See also Brochu et al. (2010) for an application to computer animation. These papers use various approximate inference techniques, and some theoretical guarantees have recently been derived by Picheny et al. (2019) and Petit et al. (2022), but all of this work falls into the non-sequential class of methods. On the other hand, in the simulation literature, pairwise comparisons have recently been considered by Xiao et al. (2023), but this paper focuses on a finite-dimensional learning problem closer to ranking and selection than global optimization. Therefore, the direction investigated in our paper has some novelty for both communities, and our preliminary results suggest that it could be of some practical interest.

2 PRELIMINARIES

Section 2.1 reviews the standard GP regression model (without censoring). Section 2.2 gives a high-level explanation of approximate Bayesian inference techniques, with emphasis on the moment-matching method that will be used later in the paper.

2.1 Gaussian Process Regression

Define a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and assume that the domain $\mathcal{X} \subset \mathbb{R}^d$ is compact. Let $(f(x))_{x \in \mathcal{X}}$ be a Gaussian process with mean function $m: \mathcal{X} \rightarrow \mathbb{R}$ and covariance kernel $K: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. By the finite-dimensional characterization of Gaussian processes, we have the following property: for an arbitrary positive integer k , and for arbitrary $x \in \mathcal{X}^k$, the vector $X = [f(x_1), \dots, f(x_k)]^\top$ follows a multivariate Gaussian distribution with mean vector

$$\mu_X = [m(x_1), \dots, m(x_k)]^\top$$

and covariance matrix

$$\Sigma_X = \begin{bmatrix} K(x_1, x_1) & \dots & K(x_1, x_k) \\ \dots & \ddots & \dots \\ K(x_k, x_1) & \dots & K(x_k, x_k) \end{bmatrix}.$$

Suppose that $y \in \mathcal{X}^m$ and $Y = [f(y_1), \dots, f(y_m)]^\top$. Then, Y is similarly multivariate Gaussian with mean vector μ_Y and covariance matrix Σ_Y , and the pair (X, Y) is jointly Gaussian with covariance matrix

$\begin{bmatrix} \Sigma_X & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_Y \end{bmatrix}$, where

$$\Sigma_{XY} = \begin{bmatrix} K(x_1, y_1) & \dots & K(x_1, y_m) \\ \dots & \ddots & \dots \\ K(x_k, y_1) & \dots & K(x_k, y_m) \end{bmatrix}.$$

Then, it follows straightforwardly that the conditional distribution of X given Y is Gaussian with mean

$$\mu_X - \Sigma_{XY} \Sigma_Y^{-1} (Y - \mu_Y) \quad (1)$$

and covariance matrix

$$\Sigma_X - \Sigma_{XY} \Sigma_Y^{-1} \Sigma_{YX}. \quad (2)$$

This property is computationally useful because it reduces the conditional distribution calculations to matrix inversion and multiplication. It forms the basis for (noiseless) Gaussian process regression: given Y , the conditional distribution of the entire process f cannot be written down explicitly, but we can easily compute such a distribution for any finite-dimensional vector X of individual function values. The same basic principle applies when observations are noisy, as we can add a Gaussian white noise vector to Y , maintaining the Gaussianity assumptions required for the computation.

2.2 Approximate Inference Techniques

Suppose now that we do not observe the function values, or any quantity that follows a Gaussian distribution. In general, the conditional density $\mathbb{P}(X \in dx | \mathcal{G})$, for an arbitrary sub-sigma-algebra $\mathcal{G} \subseteq \mathcal{F}$, will not be multivariate Gaussian. The goal of approximate Bayesian inference is to construct an artificial density $\mathbb{Q}(X \in dx)$, for any arbitrary X , that is multivariate Gaussian, and ideally can be efficiently computed from the prior distribution and \mathcal{G} , but also approximates the true non-Gaussian posterior in some way.

One popular approach for constructing such an approximation (which we also adopt in this paper) is the method of moment-matching, first studied by Minka (2001b) and subsequently deployed in settings such as dynamic pricing (Chhabra and Das 2011), competitive gaming (Herbrich, Minka, and Graepel 2006), and ranking and selection (Zhang and Song 2017). This technique chooses \mathbb{Q} to solve the system

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}(X) &= \mathbb{E}_{\mathbb{P}(\cdot | \mathcal{G})}(X), \\ \mathbb{E}_{\mathbb{Q}}(XX^\top) &= \mathbb{E}_{\mathbb{P}(\cdot | \mathcal{G})}(XX^\top). \end{aligned}$$

In other words, as the name suggests, we are simply setting the first and second moments of the artificial distribution to be equal to the corresponding moments of the true posterior. In some situations, it may be necessary to consider higher-order moments, but for our purposes in this paper, the mean and covariance matrix will suffice.

We briefly mention two other approximate inference techniques. The method of assumed density filtering chooses \mathbb{Q} to solve the optimization problem

$$\min_{\mathbb{Q}} \mathbb{E}_{\mathbb{Q}} \left[\log \left(\frac{\mathbb{P}(X | \mathcal{G})}{\mathbb{Q}(X)} \right) \right],$$

whose objective function is the Kullback-Leibler divergence measuring the difference between \mathbb{Q} and the true posterior. This technique was used by, e.g., Qu et al. (2015) in the context of ranking and selection with correlated beliefs. The optimization problem, however, may be more difficult computationally than the solution of the moment-matching equations. In certain situations, as argued by Minka (2001a), the two methods yield the same results.

In problems involving GPs, a fairly popular technique is Laplace approximation; see Chu, Ghahramani, and Williams (2005), Chu and Ghahramani (2005), and Kuss, Rasmussen, and Herbrich (2005). The mean of the artificial density \mathbb{Q} is set to the maximum a posteriori (MAP) estimator, which is straightforward to compute in standard GP regression, but not under censored information. For this reason, we do not explore this technique in the present paper.

In any event, however, once \mathbb{Q} has been selected using one of these techniques, we then treat it as the true density of X , and discard $\mathbb{P}(X \in dx | \mathcal{G})$ entirely. The technique of choice can thus be applied sequentially, treating \mathbb{Q} as the new “prior” in the next approximation. See Chen and Ryzhov (2020) for an overview of sequential approximations in finite-dimensional settings.

3 MOMENT MATCHING FOR GP REGRESSION WITH PAIRWISE COMPARISONS

We now develop a moment-matching approximation for a specific situation in which \mathcal{G} represents the information contained in a pairwise comparison. Let $z_1, z_2 \in \mathcal{X}$ be arbitrary input values, and let $B = 1_{\{f(z_1) > f(z_2)\}}$. Let $Y = [f(z_1), f(z_2)]^\top$, as in Section 2.1. Similarly, define X to be a vector of function values corresponding to an arbitrary finite set of design points, as before. Standard GP regression relies on the fact that the random vectors Y and $X - \Sigma_{XY}\Sigma_Y^{-1}Y$ are independent. However, since B is a function of Y , it follows that B is also independent of $X - \Sigma_{XY}\Sigma_Y^{-1}Y$, which turns out to be useful in the computation.

3.1 Deriving the Approximate Posterior Mean

We first calculate the conditional first moment. For notational convenience, let $A = \Sigma_{XY}\Sigma_Y^{-1}$. Then,

$$\begin{aligned} \mathbb{E}[X | B] &= \mathbb{E}[X - AY + AY | B] \\ &= \mathbb{E}[X - AY] + A\mathbb{E}[Y | B] \\ &= \mu_X + A(\mathbb{E}[Y | B] - \mu_Y), \end{aligned} \tag{3}$$

where (3) uses the independence between B and $X - AY$. Because Y is a bivariate Gaussian vector with known parameters, we can analytically calculate

$$\mathbb{E}[Y | B] = \begin{bmatrix} \frac{\bar{\sigma}_1 \frac{J_1}{\mathbb{P}(B)} + \left(\mu_1 - \frac{\sigma_{12}}{\sigma_2^2} \mu_2 \right) + \frac{\sigma_{12}}{\sigma_2^2} \left(-\bar{\sigma}_2 \frac{J_2}{\mathbb{P}(B)} + \left(\mu_2 - \frac{\sigma_{12}}{\sigma_1^2} \mu_1 \right) \right)}{1 - \frac{\sigma_{12}^2}{\sigma_1^2 \sigma_2^2}} \\ -\bar{\sigma}_2 \frac{J_2}{\mathbb{P}(B)} + \left(\mu_2 - \frac{\sigma_{12}}{\sigma_1^2} \mu_1 \right) + \frac{\sigma_{12}}{\sigma_1^2} \left(\bar{\sigma}_1 \frac{J_1}{\mathbb{P}(B)} + \left(\mu_1 - \frac{\sigma_{12}}{\sigma_2^2} \mu_2 \right) \right) \\ \frac{}{1 - \frac{\sigma_{12}^2}{\sigma_1^2 \sigma_2^2}} \end{bmatrix} \tag{5}$$

where σ_i^2 is the unconditional (prior) variance of Y_i , and σ_{12} is the prior covariance of Y_1 and Y_2 . The other quantities that appear in (5) are given by

$$\begin{aligned}\bar{\sigma}_1 &= \sqrt{\sigma_1^2 - \frac{\sigma_{12}^2}{\sigma_2^2}}, \\ \bar{\sigma}_2 &= \sqrt{\sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}}, \\ J_1 &= \frac{1}{\sqrt{2\pi}} \cdot \sqrt{\frac{\bar{\sigma}_1^2}{\sigma_2^2 + \bar{\sigma}_1^2}} \cdot \exp \left\{ \frac{\left(\frac{\bar{\mu}_1 \sigma_2^2 + \mu_2 \bar{\sigma}_1^2}{\sigma_2^2 + \bar{\sigma}_1^2} \right)^2}{2 \frac{\bar{\sigma}_1^2 \sigma_2^2}{\sigma_2^2 + \bar{\sigma}_1^2}} - \frac{\bar{\mu}_1^2}{2\bar{\sigma}_1^2} - \frac{\mu_2^2}{2\sigma_2^2} \right\}, \\ J_2 &= \frac{1}{\sqrt{2\pi}} \cdot \sqrt{\frac{\bar{\sigma}_2^2}{\sigma_1^2 + \bar{\sigma}_2^2}} \cdot \exp \left\{ \frac{\left(\frac{\bar{\mu}_2 \sigma_1^2 + \mu_1 \bar{\sigma}_2^2}{\sigma_1^2 + \bar{\sigma}_2^2} \right)^2}{2 \frac{\bar{\sigma}_2^2 \sigma_1^2}{\sigma_1^2 + \bar{\sigma}_2^2}} - \frac{\bar{\mu}_2^2}{2\bar{\sigma}_2^2} - \frac{\mu_1^2}{2\sigma_1^2} \right\}, \\ \bar{\mu}_1 &= \mu_1 - \frac{\sigma_{12}}{\sigma_2^2} \mu_2, \\ \bar{\mu}_2 &= \mu_2 - \frac{\sigma_{12}}{\sigma_1^2} \mu_1.\end{aligned}$$

All of these quantities depend only on the prior mean and covariance values for z_1, z_2 , with no dependence on X . In other words, we use the same J_1, J_2 values regardless of which or how many values we wish to predict.

We can interpret (5) as a pair of “hypothesized” values for $f(z_1), f(z_2)$. Comparing (4) with (1) shows that, given a pairwise comparison, the posterior mean of X is updated as if we had actually observed the values $\mathbb{E}[Y | B]$ at the inputs z_1, z_2 . By construction, the first component of (5) is larger than the second because B is being assumed. In the case where we observe $f(z_1) < f(z_2)$, we may simply change the order of z_1 and z_2 in the preceding computation.

Thus, given a prior mean function m_0 and a single observation B , we may use (4) to compute an approximate posterior mean $m_1(x)$ for any $x \in \mathcal{X}$. We may then apply the same technique to m_1 and a new pairwise comparison to obtain a sequence $\{m_n\}$ of posterior mean functions.

3.2 Deriving the Approximate Posterior Covariance

We proceed similarly to Section 3.1. Recalling $A = \Sigma_{XY} \Sigma_Y^{-1}$, we derive the conditional second moment of X as

$$\begin{aligned}\mathbb{E} [XX^\top | B] &= \mathbb{E} [(X - AY + AY)(X - AY + AY)^\top | B] \\ &= \mathbb{E} [(X - AY)(X - AY)^\top + (X - AY)(AY)^\top + AY(X - AY)^\top + AY(AY)^\top | B] \\ &= \mathbb{E} [(X - AY)(X - AY)^\top] + \mathbb{E} [(X - AY)] (\mathbb{E} [AY | B])^\top \\ &\quad + \mathbb{E} [AY | B] \mathbb{E} [(X - AY)]^\top + A \mathbb{E} [YY^\top | B] A^\top \\ &= \Sigma_X - A \Sigma_Y A^\top + \mathbb{E} [(X - AY)] (\mathbb{E} [(X - AY)])^\top + \mathbb{E} [(X - AY)] (\mathbb{E} [AY | B])^\top \\ &\quad + \mathbb{E} [AY | B] \mathbb{E} [(X - AY)]^\top + A \mathbb{E} [YY^\top | B] A^\top.\end{aligned}$$

Some of the terms in this expression disappear when plugged into the computation of the posterior variance. We obtain

$$\begin{aligned} \mathbb{E} \left[(X - \mathbb{E}[X | B]) (X - \mathbb{E}[X | B])^\top | B \right] &= \mathbb{E} \left[XX^\top | B \right] - \mathbb{E}[X | B] (\mathbb{E}[X | B])^\top \\ &= \Sigma_X - A \Sigma_Y A^\top + AVA^\top, \end{aligned} \quad (6)$$

where

$$V = \mathbb{E} \left[YY^\top | B \right] - \mathbb{E}[Y | B] (\mathbb{E}[Y | B])^\top$$

is the conditional covariance matrix of Y given B . The term $\mathbb{E}[Y | B]$ has already been computed in (5). The second moment $\mathbb{E}[YY^\top | B]$ looks somewhat unwieldy, and so we do not state the complete expression here, but it can also be computed. Like (5), these quantities all depend only on the inputs z_1, z_2 for which the pairwise comparison is performed, not on X .

Comparing (6) with (2), we see that the approximation is almost identical except for the extra term AVA^\top . This quantity can be interpreted as additional variability or uncertainty due to the presence of censoring. It is easy to verify that AVA^\top is a positive semidefinite matrix, so the approximate posterior covariance derived in this way yields a valid covariance function.

4 DISCUSSION OF PAIRWISE SAMPLING

Just as in standard GP regression, the approximate posterior update is affected by the choice of the pair (z_1, z_2) . Our approach naturally lends itself to a sequential sampling scheme where these pairs are selected adaptively, based on the outcomes of all previous decisions. Many sequential learning methods have been developed for standard GP regression, including probability of improvement (Kushner 1964), expected improvement (Jones et al. 1998), sequential kriging (Huang et al. 2006) and knowledge gradient (Scott et al. 2010). In this paper, we will work with a simple adaptation of one such scheme, but an important question for future work is the development of new acquisition functions specialized to the setting of pairwise comparisons.

Suppose that we have made n pairwise comparisons, and let m_n and K_n be the approximate mean and covariance functions obtained by computing (4) and (6) sequentially after each comparison. Let \mathbb{P}_n denote the measure induced by the current approximate posterior distribution (note that this measure is the result of n sequential approximations, and therefore is not a conditional probability). For two inputs z_1, z_2 , define the probability of improvement (PI) as

$$g_n(z_1, z_2) = \mathbb{P}_n(f(z_1) > f(z_2)) = \Phi \left(- \frac{m_n(z_2) - m_n(z_1)}{\sqrt{K_n(z_1, z_1) + K_n(z_2, z_2) - 2K_n(z_1, z_2)}} \right), \quad (7)$$

where Φ is the standard normal cdf. In words, this is the probability that z_2 has a better value than z_1 . Note that (7) incorporates uncertainty about both $f(z_1)$ and $f(z_2)$ into the computation, analogous to the “complete expected improvement” criterion of Salemi et al. (2019). We do not use expected improvement in this paper because we are not able to observe function values directly, while (7) deals with precisely the kind of event that we can observe. One could perhaps refer to (7) as “complete probability of improvement” by analogy.

We then choose a new pair $(z_{n,1}, z_{n,2})$ as follows. Let $x_n^* = \arg \min_{x \in \mathcal{X}} m_n(x)$ be the current estimated minimizer. Then, set

$$z_{n,1} = x_n^*, \quad z_{n,2} = \arg \max_{z \in \mathcal{X}} g_n(x_n^*, z). \quad (8)$$

In words, we compare the current-best solution with the solution that appears to be the most likely to improve on the current-best, under the beliefs that we currently have. Similar reasoning could potentially be used to adapt other existing acquisition functions like expected improvement or knowledge gradient,




Figure 1: True contour plots of Rosenbrock (left) and Rastrigin (right) functions.

but, again, we believe that it is more principled to use a criterion that is directly connected to the nature of the information that we can observe.

In this paper, we will explore the performance of (8) numerically, but theoretical guarantees remain an important open question. Calvin et al. (2018) derived strong convergence rate results for probability of improvement in a setting where rectangular partitioning (not GP regression) was used to build a metamodel and there was no censoring. Bull (2011) studied the convergence rates of expected improvement in standard GP regression. However, it is still not fully clear whether such methods can retain their advantages in the present situation.

5 NUMERICAL EXPERIMENTS

In the following, we present a preliminary numerical study to illustrate how the approximate Bayesian algorithm works and to investigate practical issues related to its performance. We use two different test functions in both 1 and 2 dimensions, illustrated in Figure 1. We also consider two types of prior covariance functions, namely the squared exponential kernel $K_0(x, y) = \sigma^2 \exp\{-\frac{1}{2}(\|x - y\|_2)^2 / (2l^2)\}$ and the Matérn $3/2$ kernel $K_0(x, y) = \sigma^2 \exp\{-\frac{\sqrt{3}}{2}\|x - y\|_2 / l\} (1 + \frac{\sqrt{3}}{2}\|x - y\|_2 / l)$. As a simple benchmark, we also implemented a sampling policy that selected pairs uniformly at random on the domain. Table 1 summarizes the settings of the various experiments that were conducted.

Each experiment is started by running 10 uniformly chosen pairs, then proceeding with the corresponding sampling policy. As a measure of performance, we use the current estimated minimum x_n^* and the true minimum $x^* = \arg \min_x f(x)$ to calculate $f(x_n^*) - f(x^*)$, which measures the suboptimality of the current-best solution. For computational convenience, we discretized the domain into a 141×141 grid. It is possible to avoid grid approximation if a symbolic mathematics language is used, however, in that case the computation gets more and more expensive at each iteration.

Table 1: Parameters used in experiments.

E#	f	d	Kernel	(σ, l)	Policy
1	Rastrigin	2	SE	(1000,0.1)	PI
2	Rastrigin	2	Matérn	(1000,0.1)	PI
3	Rastrigin	2	Matérn	(1000,0.1)	Uniform
4	Rastrigin	1	SE	(100,0.15)	Uniform
5	Rosenbrock	2	SE	(1000,0.5)	PI
6	Rosenbrock	2	Matérn	(1000,0.5)	PI
7	Rosenbrock	2	SE	(1000,0.5)	Uniform




Figure 2: Estimated surface (left) and $f(x_n^*) - f(x^*)$ (right) for experiment 1.

The first two experiments, shown in Figures 2 and 3, indicate that the choice of kernel has little impact on our ability to find the minimum of the two-dimensional Rastrigin function. In both cases, the PI policy quickly finds a reasonably good local minimum. The surface as a whole is estimated poorly, because of the presence of censoring, and because one of the points in each pair is always set to the current minimizer, causing the policy to quickly narrow in on a particular region of the domain. In general, there seems to be little reason to expect that we would be able to accurately estimate the entire function; yet, it certainly seems that we can still optimize it.

Figure 4 presents results for the same function with the pairs being uniformly generated. Overall, performance is worse than PI because considerable sampling effort is being expended on relatively unimportant regions. However, we have captured a bit more of the overall shape of the function – not necessarily the function values themselves, but the ordering between them. This can be seen more clearly in experiment 4, where we consider a single dimension. Figure 5 shows that the estimated surface is highly inaccurate in terms of pure statistical error; nonetheless, the approximate Bayesian procedure is quite accurate in identifying local optima and their relationship to each other. In our opinion, this suggests that approximate Bayesian inference may be quite useful for optimization, where the main concern is finding the global minimizer rather than correctly estimating the function.

As we can see in Figures 6-7, experiments 5 and 6 behave similarly to experiments 1 and 2: that is, the PI policy quickly narrows in on a single region and finds a local optimum. The shape of the function on the whole domain is not captured, but a good solution is found more quickly. We note (since this is not apparent from Figure 1) that the Rosenbrock function attains its global minimum at the point (1, 1), so the optimal solution is found exactly in experiment 6. However, in Experiment 5, PI gets stuck in the




Figure 3: Estimated surface (left) and $f(x_n^*) - f(x^*)$ (right) for experiment 2.

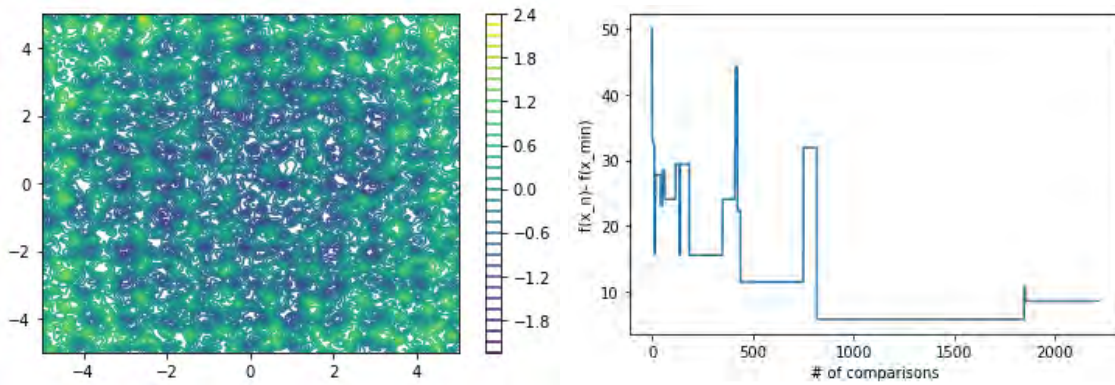


Figure 4: Estimated surface (left) and $f(x_n^*) - f(x^*)$ (right) for experiment 3.




Figure 5: Estimated surface (left, blue curve) compared to true f (left, violet curve) and estimated surface with 95% confidence bands (right) for experiment 4.

long “saddle” of the Rosenbrock function and is not able to improve on it within the time horizon we considered.

In experiment 7, we again consider the Rosenbrock function, but use uniform sampling to generate pairs. Since this function can experience abrupt changes between pairs of points, we can see in Figure 8 that the estimated minimum can actually become worse after certain comparisons. However, the overall shape of the function begins to emerge, similarly to experiments 3 and 4.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we have presented evidence that one can effectively optimize an unknown function using only data from pairwise comparisons, together with a sequential approximate Bayesian estimator. We have also illustrated the exploration-exploitation tradeoff by contrasting probability of improvement with uniform sampling. We find that PI allows for quicker estimation of the minimizer at the cost of poor overall accuracy and the risk of getting stuck at local minima.

In our ongoing work, we are studying ways to construct sampling policies that can be shown to converge to the true minimizer. Although the observations are heavily censored, our numerical results suggest that this is possible – even if we do not learn the function itself, it may be sufficient to correctly learn the ordering between pairs of inputs.

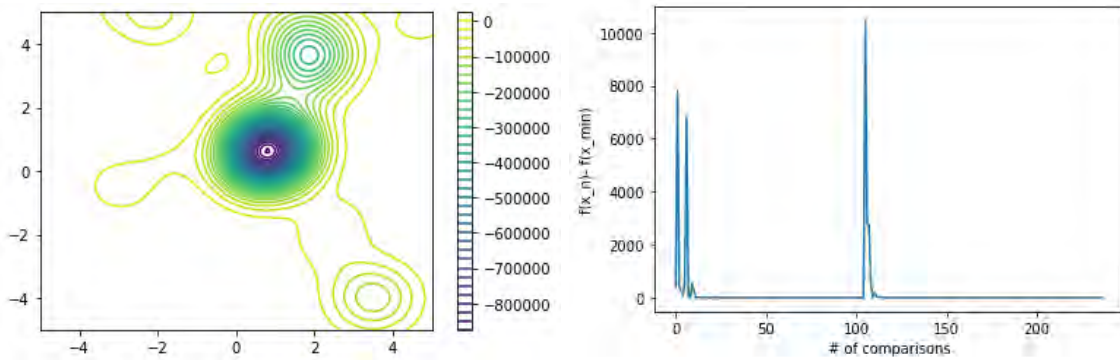


Figure 6: Estimated surface (left) and $f(x_n^*) - f(x^*)$ (right) for experiment 5.

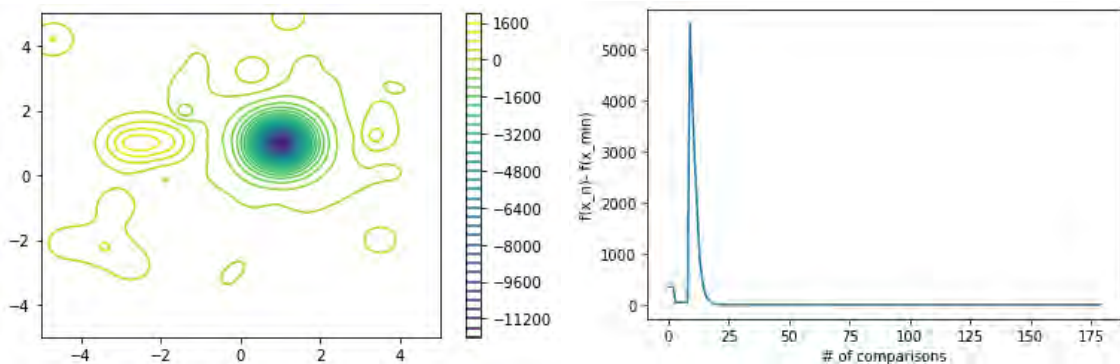


Figure 7: Estimated surface (left) and $f(x_n^*) - f(x^*)$ (right) for experiment 6.

REFERENCES

- Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Metamodeling". *Operations Research* 58(2):371–382.
- Branke, J., S. Corrente, S. Greco, and W. Gutjahr. 2017. "Efficient Pairwise Preference Elicitation Allowing for Indifference". *Computers & Operations Research* 88:175–186.
- Brochu, E., T. Brochu, and N. De Freitas. 2010. "A Bayesian Interactive Optimization Approach to Procedural Animation Design". In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 103–112.
- Bull, A. D. 2011. "Convergence Rates of Efficient Global Optimization Algorithms". *Journal of Machine Learning Research* 12:2879–2904.
- Calvin, J., G. Gimbutienė, W. O. Phillips, and A. Žilinskas. 2018. "On Convergence Rate of a Rectangular Partition Based Global Optimization Algorithm". *Journal of Global Optimization* 71(1):165–191.
- Chen, Y., and I. O. Ryzhov. 2020. "Technical Note – Consistency Analysis of Sequential Learning under Approximate Bayesian Inference". *Operations Research* 68(1):295–307.
- Chhabra, M., and S. Das. 2011. "Learning the Demand Curve in Posted-Price Digital Goods Auctions". In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems*, edited by K. Tumer, P. Yolum, L. Sonenberg, and P. Stone, Volume 1, 63–70.
- Chu, W., and Z. Ghahramani. 2005. "Preference Learning with Gaussian Processes". In *Proceedings of the 22nd International Conference on Machine Learning*, 137–144.
- Chu, W., Z. Ghahramani, and C. K. I. Williams. 2005. "Gaussian Processes for Ordinal Regression". *Journal of Machine Learning Research* 6(7):1019–1041.
- Garnett, R. 2023. *Bayesian Optimization*. Cambridge University Press.
- Herbrich, R., T. P. Minka, and T. Graepel. 2006. "TrueSkill™: A Bayesian Skill Rating System". In *Advances in Neural Information Processing Systems*, edited by B. Schölkopf, J. C. Platt, and T. Hoffman, Volume 19, 569–576.




Figure 8: Estimated surface (left) and $f(x_n^*) - f(x^*)$ (right) for experiment 7.

- Huang, D., T. T. Allen, W. I. Notz, and N. Zeng. 2006. “Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models”. *Journal of Global Optimization* 34(3):441–466.
- Jaakkola, T. S., and M. I. Jordan. 2000. “Bayesian Parameter Estimation via Variational Methods”. *Statistics and Computing* 10:25–37.
- Ji, Y., and S. Kim. 2014. “Regularized Radial Basis Function Models for Stochastic Simulation”. In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 3833–3844. IEEE.
- Jones, D., M. Schonlau, and W. J. Welch. 1998. “Efficient Global Optimization of Expensive Black-Box Functions”. *Journal of Global Optimization* 13(4):455–492.
- Kleijnen, J. P. C. 2014. “Simulation-Optimization via Kriging and Bootstrapping: a Survey”. *Journal of Simulation* 8(4):241–250.
- Kushner, H. J. 1964. “A New Method of Locating the Maximum Point of an Arbitrary Multippeak Curve in the Presence of Noise”. *Journal of Basic Engineering* 86(1):97–106.
- Kuss, M., C. E. Rasmussen, and R. Herbrich. 2005. “Assessing Approximate Inference for Binary Gaussian Process Classification”. *Journal of Machine Learning Research* 6(10):1679–1704.
- Minka, T. P. 2001a. “Expectation Propagation for Approximate Bayesian Inference”. In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, 362–369.
- Minka, T. P. 2001b. *A Family of Algorithms for Approximate Bayesian Inference*. Ph. D. thesis, Massachusetts Institute of Technology.
- Petit, S., J. Bect, and E. Vazquez. 2022. “Relaxed Gaussian Process Interpolation: a Goal-Oriented Approach to Bayesian Optimization”. *arXiv preprint arXiv:2206.03034*.
- Picheny, V., S. Vakili, and A. Artemev. 2019. “Ordinal Bayesian Optimisation”. *arXiv preprint arXiv:1912.02493*.
- Qu, H., I. O. Ryzhov, M. C. Fu, and Z. Ding. 2015. “Sequential Selection with Unknown Correlation Structures”. *Operations Research* 63(4):931–948.
- Rasmussen, C. E., and C. K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. MIT Press.
- Salemi, P. L., E. Song, B. L. Nelson, and J. Staum. 2019. “Gaussian Markov Random Fields for Discrete Optimization via Simulation: Framework and Algorithms”. *Operations Research* 67(1):250–266.
- Scott, W. R., W. B. Powell, and H. P. Simão. 2010. “Calibrating Simulation Models Using the Knowledge Gradient with Continuous Parameters”. In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Huan, and E. Yücesan, 1099–1109.
- Wang, Y., and D. M. Blei. 2019. “Frequentist Consistency of Variational Bayes”. *Journal of the American Statistical Association* 114(527):1147–1161.
- Xiao, H., Y. Zhang, G. Kou, S. Zhang, and J. Branke. 2023. “Ranking and Selection for Pairwise Comparison”. *Naval Research Logistics (to appear)*.
- Zhang, Q., and Y. Song. 2017. “Moment-Matching-Based Conjugacy Approximation for Bayesian Ranking and Selection”. *ACM Transactions on Modeling and Computer Simulation* 27(4):26:1–26:23.

AUTHOR BIOGRAPHIES

EFE SERTKAYA is a Ph.D. student in Operations Management at the Robert H. Smith School of Business, University of Maryland. His research interests include simulation optimization and stochastic optimization. His email address is sertkaya@umd.edu.

ILYA O. RYZHOV is a Professor of Operations Management at the Robert H. Smith School of Business, University of Maryland. His research primarily focuses on simulation optimization and statistical learning, with applications in business analytics, revenue management, and nonprofit/humanitarian operations. He is a coauthor of the book *Optimal Learning* (Wiley, 2012). His work was recognized in WSC's Best Theoretical Paper Award competition on three separate occasions (winner in 2012, finalist in 2009 and 2016), and he received I-SIM's Outstanding Publication Award in 2017. His email address is iryzhov@rhsmith.umd.edu.