Preferential Heteroscedastic Bayesian Optimization with Informative Noise Priors

Marshal Sinaga MARSHAL.SINAGA@AALTO.FI
Department of Computer Science, Aalto University

Julien Martinelli JULIEN.MARTINELLI@AALTO.FI
Department of Computer Science, Aalto University

Samuel Kaski SAMUEL.KASKI@AALTO.FI
Department of Computer Science, Aalto University
Department of Computer Science, University of Manchester

Abstract

Preferential Bayesian optimization (PBO) is a sample-efficient framework for optimizing a black-box function by utilizing human preferences between two candidate solutions as a proxy. Conventional PBO relies on homoscedastic noise to model human preference structure. However, such noise fails to accurately capture the varying levels of human aleatoric uncertainty among different pairs of candidates. For instance, a chemist with solid expertise in glucose-related molecules may easily compare two compounds and struggle for alcohol-related molecules. Furthermore, PBO ignores this uncertainty when searching for a new candidate, consequently underestimating the risk associated with human uncertainty. To address this, we propose heteroscedastic noise models to learn human preference structure. Moreover, we integrate the preference structure with the acquisition functions that account for aleatoric uncertainty. The noise models assign noise based on the distance of a specific input to a predefined set of reliable inputs known as *anchors*. We empirically evaluate the proposed approach on a range of synthetic black-box functions, demonstrating a consistent improvement over homoscedastic PBO.

Keywords: Preferential Bayesian Optimization, Heteroscedastic noise

1. Introduction

Preferential Bayesian Optimization (PBO) utilizes the Gaussian process (GP) to learn human preference structure (Brochu et al., 2010). Typically, existing works model human preferences together with a notion of *epistemic* uncertainty, coming from the fact that we only have a finite number of observations to learn the latent function describing the human's utility (Chu and Ghahramani, 2005). Here, we stress that a key component to model a user's incomplete knowledge is the input-dependent *aleatoric* uncertainty, as the relevance of an expert is not uniform across the design space. Consider a scenario where a chemist proficient in glucose-related compounds can effortlessly contrast two substances but encounter challenges when dealing with alcohol-related molecules. In such a case, assigning a uniform level of uncertainty to both types of molecules is not reasonable.

Another challenge is that PBO approaches are typically *risk-neutral*, only seeking to optimize the expected preference structure value (Gonzalez et al., 2017). However, the varying

human uncertainty results in situations where the preference structure yields two solutions that deliver similar expected function values, yet one may be noisier. Consequently, the optimization problem shifts towards a *risk-averse* paradigm, necessitating a balance between maximizing the expected preference structure value and minimizing uncertainty (Makarova et al., 2021). For example, a chemist might opt for slightly inferior quality molecules but, at the same time, lead to a smaller variance.

This research aims to develop heteroscedastic PBO models by leveraging the information provided by humans. Our contributions are summarized as follows:

- Development of a heteroscedastic PBO framework employing a kernel density estimator-based noise model to account for human aleatoric uncertainty.

- Extensive empirical validation showcasing the effectiveness and superiority of our proposed models through experiments conducted on various synthetic black-box functions, consistently outperforming homoscedastic PBO approaches.

2. Method

For a function $f : \mathcal{X} \to \mathbb{R}$ defined over a subset $\mathcal{X} \subset \mathbb{R}^d$, we aim at finding $\mathbf{x}^* = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{argmax}}\, f(\mathbf{x})$. We assume the presence of heteroscedastic noise, such that for all $\mathbf{x} \in \mathcal{X}$,

$$g(\mathbf{x}) = f(\mathbf{x}) + \varepsilon(\mathbf{x}), \tag{1}$$

where ε represents the input-dependent noise, assumed to be independent of the latent function. In contrast to vanilla Bayesian Optimization (BO), in which the black-box function can be queried for a specific input, resulting in a scalar output, we assume that f can only be queried by means of preferences $\mathbf{x} \succ \hat{\mathbf{x}}$. The latter is a binary random variable taking value 1 when $\mathbf{x}$ is preferred to $\hat{\mathbf{x}}$ and 0 otherwise. A preference $\mathbf{x} \succ \hat{\mathbf{x}}$ is then such that

$$\mathbf{x} \succ \hat{\mathbf{x}} \iff g(\mathbf{x}) > g(\hat{\mathbf{x}}) \iff f(\mathbf{x}) + \varepsilon(\mathbf{x}) > f(\hat{\mathbf{x}}) + \varepsilon(\hat{\mathbf{x}}),$$

We assume access to a dataset of m comparisons $\mathcal{D} = \{\mathbf{x}_k \succ \hat{\mathbf{x}}_k\}_{k=1}^m$, thus involving $2m$ inputs. Denoting $v \triangleq f(\hat{\mathbf{x}}) + \varepsilon(\hat{\mathbf{x}}) - f(\mathbf{x}) - \varepsilon(\mathbf{x})$, we have $\mathcal{D} = \{v_k < 0\}_{k=1}^m$, which we write $\mathbf{v}_m$ for conciseness. Denote the winner (resp. loser) of duel $\mathbf{x}_k \succ \hat{\mathbf{x}}_k$ as $\mathbf{x}_{k,w}$ (resp. $\mathbf{x}_{k,l}$), and consider $\mathbf{X} = (\mathbf{x}_{1,*}, \ldots, \mathbf{x}_{n,*}, \mathbf{x}_{1,w}, \ldots, \mathbf{x}_{m,w}, \mathbf{x}_{1,l}, \ldots, \mathbf{x}_{m,l})$. We begin by placing a zero-mean Gaussian process prior over the latent function f with kernel k. For a $(n+2m)$-dimensional vector of latent function values $\mathbf{f} = \{f(\mathbf{x})\}_{\mathbf{x}\in\mathbf{X}}$, this yields

$$p(\mathbf{f}) = |2\pi\mathbf{K}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\mathbf{f}^\top\mathbf{K}^{-1}\mathbf{f}\right), \tag{2}$$

where $\mathbf{K} = \{k(\mathbf{x}, \mathbf{x}')\}_{\mathbf{x},\mathbf{x}' \in \mathbf{X}} \in \mathbb{R}^{(n+2m)\times(n+2m)}$ is the kernel matrix. The joint distribution for $(\mathbf{f}_*, \mathbf{v}_m)$ can be written as follows (Takeno et al., 2023):

$$\begin{aligned}
&\begin{bmatrix} \mathbf{f}_* \\ \mathbf{v}_m \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}) \\
&\quad \boldsymbol{\Sigma} \triangleq \mathbf{A}(\mathbf{K} + \mathbf{B})\mathbf{A}^\top \in \mathbb{R}^{(n+m)\times(n+m)} \\
&\quad \mathbf{A} \triangleq \begin{bmatrix} \mathbf{I}_n & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -\mathbf{I}_m & \mathbf{I}_m \end{bmatrix} \in \mathbb{R}^{(n+m)\times(n+2m)} \\
&\quad \mathbf{B} \triangleq \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_{\text{noise}} \end{bmatrix} \in \mathbb{R}^{(n+2m)\times(n+2m)}
\end{aligned} \tag{3}$$

Modeling the noise of our process boils down to specifying $\mathbf{V}_{\text{noise}}$, the lower right component of the matrix $\mathbf{B}$. In the homoscedastic setting, $\mathbf{V}_{\text{noise}}$ is a diagonal matrix with constant entries $\mathbf{V}_{\text{noise}} = \sigma^2_{\text{noise}} I_{2m}$. In the heteroscedastic setting, we may consider different values for the diagonal entries of $\mathbf{V}_{\text{noise}}$, and even take into account correlations among inputs at the noise level. In the next two subsections, we develop two models.

2.1 Anchors-based input-dependent noise

As a starting point for our noise process ε, we draw inspiration from the framework described by Eduardo and Gutmann (2023) and introduce a set of initial designs $\mathbf{x_0} = \{\mathbf{x}_0^{(q)}\}_{q=1}^{Q}$ provided by the expert. This set is also referred to as the *anchor* set and contains samples over which the user has high confidence over their possible value, meaning that the noise should have low variance. Comparisons involving samples close to the anchors should, therefore, be easy for the expert. This intuition leads to consider a mapping ϕ,

$$\phi(\mathbf{x}) = \frac{1}{Q} \sum_{q=1}^{Q} \frac{1}{h^d} \; k_h(\|\mathbf{x} - \mathbf{x}_0^{(q)}\|_2^2), \tag{4}$$

which associates a high value to a sample $\mathbf{x}$ the closer it lies from the anchors $\mathbf{x_0}$. This notion of closeness is captured by a ℓ_2 distance and a kernel k with a lengthscale h. In this work, k is chosen to be a squared exponential kernel.

As ϕ assigns higher values to points in proximity to the anchors $\mathbf{x_0}$, which are believed to characterize regions with lower noise, we define our noise model as follows:

$$\varepsilon(\mathbf{x}_k) \sim \mathcal{N}(0, \exp(-\sigma_p^2 \, \phi(\mathbf{x}_k))). \tag{5}$$

with $\sigma_p^2 > 0$ is a scaling factor. The negative squared exponentiation effectively diminishes the noise variance for higher values of ϕ in an exponentially decreasing manner. Hence, ϕ can be envisioned as a hyperprior governing the noise. Finally, in reference to Equation 3, this model implies that $\mathbf{V}_{\text{noise}} = \text{diag}(\mathbb{V}[\varepsilon(\mathbf{x}_k)]_{1\leq k\leq 2m})$.

2.2 Gaussian Process-based noise modeling

The previous model described by Equation 5 does not model noise correlation between inputs. To tackle this issue, we extend it using a GP formulation as follows

$$\begin{aligned} &\tilde{y}|\mathbf{f}, \boldsymbol{\varepsilon}, \tilde{\mathbf{x}} \sim \mathcal{N}(f(\tilde{\mathbf{x}}), \boldsymbol{\varepsilon}(\tilde{\mathbf{x}})) \\ &\boldsymbol{\varepsilon} \sim \mathcal{GP}(\mathbf{0}, \mathbf{L}) \\ &L(\mathbf{x}, \tilde{\mathbf{x}}) = \sigma_0^2(\mathbf{x}, \tilde{\mathbf{x}}) \mathrm{C}_s(\mathbf{x}, \tilde{\mathbf{x}}) \\ &\sigma_0^2(\mathbf{x}, \tilde{\mathbf{x}}) = \sigma_p^2 \sqrt{\phi(\mathbf{x})\phi(\tilde{\mathbf{x}})} \end{aligned} \tag{6}$$

C_s is a stationary covariance function and $\sigma_p^2 > 0$ a scaling factor. Notably, when considering two points $\mathbf{x}_i$ and $\mathbf{x}_j$ that are close to anchors—where both $\phi(\mathbf{x}_i)$ and $\phi(\mathbf{x}_j)$ are large—this prior suggests a high correlation between $\varepsilon(\mathbf{x}_i)$ and $\varepsilon(\mathbf{x}_j)$. For this model, the derivation of the joint distribution $(\mathbf{f}_*, \mathbf{v}_m)$ slightly differs but can again be written in closed form.

Regarding inference, we resort to the Most Likely Heteroscedastic Gaussian Process (MLHGP) scheme introduced by Kersting et al. (2007) and adapt it to the Preferential setting, yielding MLHPGP:

1. If first iteration $\rightarrow$ Fit preferential GP $\mathbf{f}$ given $\mathcal{D} = \{\mathbf{x}_k \succ \hat{\mathbf{x}}_k\}_{k=1}^m$ with $\varepsilon(\mathbf{x}_k), \varepsilon(\hat{\mathbf{x}}_k) \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ using HB.
   else $\rightarrow$ Fit $\mathbf{f}$ given $\mathcal{D} = \{\mathbf{x}_k \succ \hat{\mathbf{x}}_k\}_{k=1}^m$ with $\varepsilon(\mathbf{x}_k) \sim \mathcal{N}(0, \exp(\mu_{\mathbf{g}}^*(\mathbf{x}_k)), \varepsilon(\hat{\mathbf{x}}_k) \sim \mathcal{N}(0, \exp(\mu_{\mathbf{g}}^*(\hat{\mathbf{x}}_\mathbf{k}))$ using HB.
   This step yields the posterior $\mathbf{f}^* \sim \mathcal{N}(\mu_{\mathbf{f}}^*, \Sigma_{\mathbf{f}}^*)$.

2. Fit $\mathbf{g}$ given $\mathcal{D} = \{\mathbf{x}_k \succ \hat{\mathbf{x}}_k\}_{k=1}^m$ and $\mathbf{f}^* \sim \mathcal{N}(\mu_{\mathbf{f}}^*, \Sigma_{\mathbf{f}}^*)$ using HB. This step yields $\mathbf{g} \sim \mathcal{N}(\mu_{\mathbf{g}}^*, \Sigma_{\mathbf{g}}^*)$

3. Repeat until convergence

The GP f learned over $\mathcal{D}$ is used to subtract the effect of the latent function on the observations, creating pseudo-observations of the noise, which are then fitted by another GP g. The details of MLHPGP can be found in the appendix.

2.3 Inference and Acquisition strategy

At acquisition time, the idea is to leverage the introduced noise models to factor out the noise at candidate acquisition time. Several acquisition functions have been designed for heteroscedastic BO, namely, Aleatoric Noise-Penalized Expected Improvement (ANPEI) and Heteroscedastic Augmented Expected Improvement (HAEI) (Griffiths et al., 2021).

$$\mathrm{HAEI}(\mathbf{x}) = \mathbb{E}[(f(\mathbf{x}^*) - f(\mathbf{x}))_+] \left(1 - \frac{\gamma \sigma_\varepsilon(\mathbf{x})}{\sqrt{\mathbb{V}[f(\mathbf{x})] + \gamma^2 \sigma_\varepsilon^2(\mathbf{x})}}\right) \tag{7}$$

$$\mathrm{ANPEI}(\mathbf{x}) = \beta \mathbb{E}[(f(\mathbf{x}^*) - f(\mathbf{x}))_+] - (1 - \beta)\sigma_\varepsilon(\mathbf{x}) \tag{8}$$

where γ denotes a positive penalty term and $\beta \in [0, 1]$ is a constant.

These acquisition functions can be integrated into the preferential BO setting in a straightforward manner. One only requires small changes in the Hallucination Believer

Algorithm 1 Hallucination Believer for HPBO

Input: Initial dataset $\mathcal{D} = \{\mathbf{x}_k \succ \hat{\mathbf{x}}_k\}_{k=1}^m$
for $t = 1, \dots$ **do**
  $\mathbf{x}_{t,w}^{(1)} \leftarrow \mathbf{x}_{t-1,w}^{(1)}$
  Draw $\tilde{\mathbf{v}}_{t-1}$ from the posterior $p(\mathbf{v}_{t-1}|\mathbf{v}_{t-1} < \mathbf{0})$
  Sequentially estimate f and ε (Section 2.1 and Section 2.2)
  $\mathbf{x}_t^{(2)} \leftarrow \text{argmax}_{\mathbf{x}\in\mathcal{X}}\, \alpha(\mathbf{x})$ based on GPs $f|\tilde{\mathbf{v}}_{t-1}$ and ε
  Set $\mathbf{x}_{t,w}$ and $\mathbf{x}_{t,l}$ as the winner and loser of the duel between $\mathbf{x}_t^{(1)}$ and $\mathbf{x}_t^{(2)}$, respectively.
  $\mathcal{D}_t \leftarrow \mathcal{D}_{t-1} \cup (\mathbf{x}_{t,w} \succ \mathbf{x}_{t,l})$
end for

strategy proposed by Takeno et al. (2023). The latter builds on the observation that the posterior predictive distribution over latent function values $\mathbf{f}_*$ at a set of newly observed points $\mathbf{x}_*$ given observed duels $\mathbf{v}_m < 0$ and $\mathbf{v}_m$ can be computed in closed-form (Takeno et al., 2023, Proposition 3.1)

$$
\begin{aligned}
p(\mathbf{f}_*|\mathbf{v}_m < 0, \mathbf{v}_m) &= \mathcal{N}(\boldsymbol{\mu}_{*|\mathbf{v}}, \boldsymbol{\Sigma}_{*|\mathbf{v}}) \\
\boldsymbol{\mu}_{*|\mathbf{v}} &= \boldsymbol{\Sigma}_{*,\mathbf{v}}\boldsymbol{\Sigma}_{\mathbf{v},\mathbf{v}}^{-1}\mathbf{v}_{m-1} \\
\boldsymbol{\Sigma}_{*|\mathbf{v}} &= \boldsymbol{\Sigma}_{*,*} + \mathbf{V}_{\text{noise}*} - \boldsymbol{\Sigma}_{*,\mathbf{v}}\boldsymbol{\Sigma}_{\mathbf{v},\mathbf{v}}^{-1}\boldsymbol{\Sigma}_{*,\mathbf{v}}^{\top}.
\end{aligned}
\tag{9}
$$

The procedure and highlighted changes are presented in Algorithm 1. Instead of using a classical acquisition function to obtain the second member of the duel, we use a heteroscedastic acquisition function like ANPEI or HAEI. The latter involves the GPs for the latent function f and the noise ε, estimated with the MLHPGP method (Section 2.2).

Furthermore, in a preferential setting, the incumbent η required by the acquisition functions (Equations 7 and 8) is not available. Thus at iteration t, we use $\eta \sim p(f(\mathbf{x}_{t,w}^{(1)})|\mathbf{x}_{t,w}^{(1)})$, where $\mathbf{x}_{t,w}^{(1)}$ is the winner of the previous duel performed at iteration $t-1$.

Considering the noise in the preferential acquisition function, the tradeoff is to select informative and easy duels for humans, as they are believed to involve samples with low noise. Intuitively, the comparison between two points may be informative but merely impossible to carry for a human due to incomplete knowledge of a human.

3. Experiments

We evaluate our method on several 2-dimensional synthetic functions: Branin, Beale, and Styblinski-Tang. A linear Gaussian noise is added to the function observations following the work of (Griffiths et al., 2021). We perform $N = 3$ repetitions per baseline with different random seeds. As kernel, we use the squared exponential kernel to exhibit stationary kernel. We select $Q = 5$ anchors as initial points. Results are presented in Figure 1. To a lesser extent, the Branin and the Beale function, the ANPEI and HAEI acquisition functions equipped with our simple model perform on par with the homoscedastic approach EI. Next, the GP-based noise models yield poor results compared to the other baselines. This suggests computational issues, for instance, at inference time with Laplace approximation, as this approach is known to be suboptimal for preferential GPs (Takeno et al., 2023).

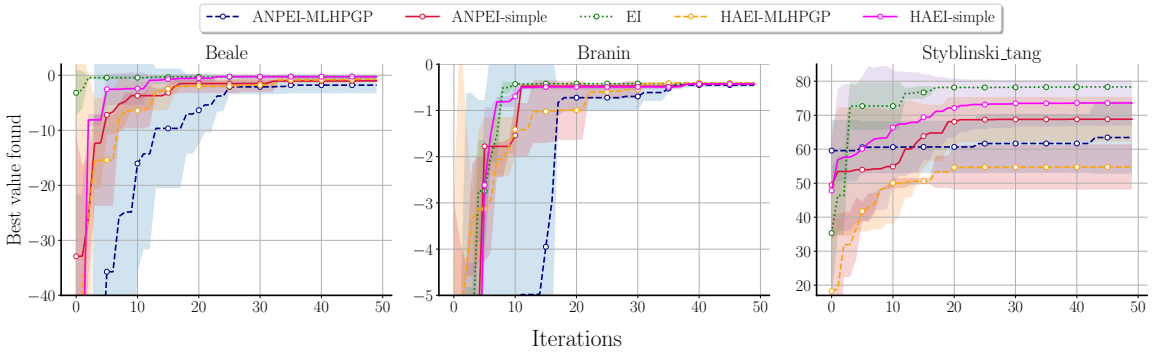

Figure 1: Best value found over BO trial for three synthetic test functions. Mean and standard deviation computed across $N = 3$ repetitions. Plain curves are associated with the simple model (Section 2.1), dashed curves with the GP model (Section 2.2, while the dotted baseline does not consider heteroscedastic noise.

4. Conclusion

We presented HPBO, a method to integrate heteroscedastic noise in Preferential BO with heteroscedastic noise. Borrowing ideas from (Eduardo and Gutmann, 2023), we based our noise models on *anchors*, a set of inputs for which the user has high confidence. In several synthetic experiments, preliminary results are not conclusive and highlight the need for further investigations from the modeling, experimental, and theoretical point of view. For instance, the inference was carried out using Laplace approximation, which is known to be suboptimal for preferential GPs (Takeno et al., 2023). Next, different synthetically injected noise models could be considered, together with real-world examples.

This being said, we want to draw attention to the fact that the combination between preferential BO and heteroscedastic noise is sensible, specifically when considering human subjects as labelers and when the dimensionality of the input space increases. For humans, performing comparisons bears an inherent noise: typically, inputs that are similar are expected to be harder to compare than inputs that are far apart. However, the notion of similarity, or distance, is not straightforward here. For instance, two D-dimensional inputs may be close in ℓ_2-norm, but share equal values on several input dimensions, thus making them potentially easier to compare than two inputs that are far in ℓ_2-norm but with all dimensions taking different values. This argument does not extend to structured data like images, graphs or molecules, which suggests the use of kernels, for that matter. We have relied on a set of anchors to encompass this notion, together with the mapping ϕ, involving a kernel k, which could accommodate structured inputs. Encompassing these different aspects in our framework represents promising avenues for future work.

Appendix - Model Selection for anchors-based input-dependent noise

For model selection, we minimize the negative log marginal likelihood approximated with Laplace approximation. The approximation requires obtaining $\mathbf{f}_{\text{MAP}} = \arg\min_{\mathbf{f}} -\log p(\mathbf{f}|\mathbf{v}_m) \approx \arg\min_{\mathbf{f}} S(\mathbf{f})$ where we define $S(\mathbf{f})$ as

$$S(\mathbf{f}) = -\log \Phi(\mathbf{z}) + \frac{1}{2}\mathbf{f}^\top K^{-1}\mathbf{f}$$

with $-\log \Phi(\mathbf{z}) = -\sum_{k=1}^{m} \log \Phi(\mathbf{z}_k)$. We derive the likelihood $\Phi(\mathbf{z}_k)$ as follows:

$$p(\mathbf{x}_k \succ \hat{\mathbf{x}}_k | f(\mathbf{x}_k), f(\hat{\mathbf{x}}_k), \varepsilon(\mathbf{x}_k), \varepsilon(\hat{\mathbf{x}}_k)) = \Phi\left(\mathbf{z}_k \triangleq \frac{f(\mathbf{x}_k) - f(\hat{\mathbf{x}}_k)}{\sqrt{\exp(-\sigma_p^2\phi(\mathbf{x}_k)) + \exp(-\sigma_p^2\phi(\hat{\mathbf{x}}_k))}}\right)$$

where $\varepsilon(\mathbf{x}_k) \sim \mathcal{N}(0, \exp(-\sigma_p^2\phi(\mathbf{x}_k)))$ and $\varepsilon(\hat{\mathbf{x}}_k) \sim \mathcal{N}(0, \exp(-\sigma_p^2\phi(\hat{\mathbf{x}}_k)))$. Taking the first and the second derivates, respectively gives us

$$\begin{aligned}
\frac{\partial S(\mathbf{f})}{\partial \mathbf{f}} &= \begin{bmatrix} -\frac{\phi(\mathbf{z})}{\sqrt{\exp(-\sigma_p^2\phi(\mathbf{x})) + \exp(-\sigma_p^2\phi(\hat{\mathbf{x}}))}\Phi(\mathbf{z})} \\ \frac{\phi(\mathbf{z})}{\sqrt{\exp(-\sigma_p^2\phi(\mathbf{x})) + \exp(-\sigma_p^2\phi(\hat{\mathbf{x}}))}\Phi(\mathbf{z})} \end{bmatrix} + K^{-1}\mathbf{f} \\
\frac{\partial^2 S(\mathbf{f})}{\partial^2 \mathbf{f}} &= \Lambda + K^{-1} \\
\Lambda &= \begin{bmatrix} \mathbf{z}'\text{diag}(\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} + \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) & \mathbf{z}'\text{diag}(-\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} - \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) \\ \mathbf{z}'\text{diag}(-\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} - \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) & \mathbf{z}'\text{diag}(\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} + \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) \end{bmatrix} \\
\mathbf{z}' &= \frac{1}{\exp(-\sigma_p^2\phi(\mathbf{x})) + \exp(-\sigma_p^2\phi(\hat{\mathbf{x}}))}
\end{aligned}$$

Given $\frac{\partial S(\mathbf{f})}{\partial \mathbf{f}}$ and $\frac{\partial^2 S(\mathbf{f})}{\partial^2 \mathbf{f}}$, we construct the evidence as

$$p(\mathcal{D}|\theta) \approx \exp(-S(\mathbf{f}_{\text{MAP}}))|I + K\Lambda_{\text{MAP}}|^{-1/2}$$

with $\Lambda_{\text{MAP}} = \Lambda|_{\mathbf{f}_{\text{MAP}}}$. We then write the model selection problem as

$$\theta^* = \arg\min_{\theta} -\log p(\mathcal{D}|\theta)$$

with $\theta = \{\lambda, \lambda_{\text{KDE}}\}$ denotes the GP hyperparameters. We utilize L-BFGS-B to obtain the solution.

Appendix - Preferential MLHGP

We rely on the Laplace approximation to fit $\mathbf{g}$ and perform model selection on $\mathbf{f}, \mathbf{g}$. The following subsections provide the details of model selection $\mathbf{f}$ and $\mathbf{g}$.

Model selection of f

While the posterior inference depends on the hallucination believer method, we perform model selection by minimizing the negative log marginal likelihood approximated by Laplace approximation. The approximation requires obtaining $\mathbf{f}_{\text{MAP}} = \arg\min_{\mathbf{f}} -\log p(\mathbf{f}|\mathcal{D}) \approx \arg\max_{\mathbf{f}} S(\mathbf{f})$ where we define $S(\mathbf{f})$ as

$$S(\mathbf{f}) = -\log \Phi(\mathbf{z}) + \frac{1}{2}\mathbf{f}^\top K^{-1}\mathbf{f}$$

with $-\log \Phi(\mathbf{z}) = -\sum_{k=1}^{m} \log \Phi(\mathbf{z}_k)$. We derive the likelihood $\Phi(\mathbf{z}_k)$ as follows:

$$p(\mathbf{x}_k \succ \hat{\mathbf{x}}_k | f(\mathbf{x}_k), f(\hat{\mathbf{x}}_k), \varepsilon(\mathbf{x}_k), \varepsilon(\hat{\mathbf{x}}_k)) = \Phi\left(\mathbf{z}_k \triangleq \frac{f(\mathbf{x}_k) - f(\hat{\mathbf{x}}_k)}{\sqrt{\exp(\mu^*_{\mathbf{g}}(\mathbf{x}_k)) + \exp(\mu^*_{\mathbf{g}}(\hat{\mathbf{x}}_k))}}\right)$$

with $\varepsilon(\mathbf{x}_k) \sim \mathcal{N}(0, \exp(\mu^*_{\mathbf{g}}(\mathbf{x}_k))$ and $\varepsilon(\hat{\mathbf{x}}_k) \sim \mathcal{N}(0, \exp(\mu^*_{\mathbf{g}}(\mathbf{x}_k))$. Taking the first and the second derivates, respectively, gives us

$$\begin{aligned}
\frac{\partial S(\mathbf{f})}{\partial \mathbf{f}} &= \begin{bmatrix} -\frac{\phi(\mathbf{z})}{\sqrt{\exp(\mu^*_{\mathbf{g}}(\mathbf{x})) + \exp(\mu^*_{\mathbf{g}}(\hat{\mathbf{x}}))}\Phi(\mathbf{z})} \\ \frac{\phi(\mathbf{z})}{\sqrt{\exp(\mu^*_{\mathbf{g}}(\mathbf{x})) + \exp(\mu^*_{\mathbf{g}}(\hat{\mathbf{x}}))}\Phi(\mathbf{z})} \end{bmatrix} + K^{-1}\mathbf{f} \\
\frac{\partial^2 S(\mathbf{f})}{\partial^2 \mathbf{f}} &= \Lambda + K^{-1} \\
\Lambda &= \frac{1}{\exp(\mu^*_{\mathbf{g}}(\mathbf{x})) + \exp(\mu^*_{\mathbf{g}}(\hat{\mathbf{x}}))} \begin{bmatrix} \text{diag}(\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} + \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) & \text{diag}(-\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} - \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) \\ \text{diag}(-\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} - \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) & \text{diag}(\frac{\phi(\mathbf{z})^2}{\Phi(\mathbf{z})^2} + \frac{\phi(\mathbf{z})}{\Phi(\mathbf{z})}\mathbf{z}) \end{bmatrix}
\end{aligned}$$

The model selection problem is the same as in the simple noise model.

Model Selection g

The model selection minimizes the negative log marginal likelihood approximated with Laplace approximation. The approximation requires obtaining $\mathbf{g}_{\text{MAP}} = \arg\min_{\mathbf{g}} -\log p(\mathbf{g}|\mathcal{D}) \approx \arg\min_{\mathbf{g}} S(\mathbf{g})$ where we define $S(\mathbf{g})$ as

$$S(\mathbf{g}) = -\log \Phi(\mathbf{z}) + \frac{1}{2}\mathbf{g}^\top \Sigma^{-1}\mathbf{g}$$

with $-\log \Phi(\mathbf{z}) = -\sum_{k=1}^{m} \log \Phi(\mathbf{z}_k)$. We derive the likelihood $\Phi(\mathbf{z}_k)$ as follows:

$$\begin{aligned}
p(\mathbf{x}_k \succ \hat{\mathbf{x}}_k | f(\mathbf{x}_k), f(\hat{\mathbf{x}}_k), \varepsilon(\mathbf{x}_k), \varepsilon(\hat{\mathbf{x}}_k)) &= p(f(\mathbf{x}_k) + \varepsilon(\mathbf{x}_k) > f(\hat{\mathbf{x}}_k) + \varepsilon(\hat{\mathbf{x}}_k)) \\
&= p(f(\hat{\mathbf{x}}_k) - f(\mathbf{x}_k) < \varepsilon(\mathbf{x}_k) - \varepsilon(\hat{\mathbf{x}}_k)) \\
&= \Phi\left(\mathbf{z}_k \triangleq \frac{\delta_{\mathbf{x}_k}\sqrt{\exp(g(\mathbf{x}_k))} - \delta_{\hat{\mathbf{x}}_k}\sqrt{\exp(g(\hat{\mathbf{x}}_k))} + \mu^*_{f\mathbf{x}_k} - \mu^*_{f\hat{\mathbf{x}}_k}}{\sqrt{\Sigma^*_{f\mathbf{x}_k,\mathbf{x}_k} + \Sigma^*_{f\hat{\mathbf{x}}_k,\hat{\mathbf{x}}_k} - 2\Sigma^*_{f\mathbf{x}_k,\hat{\mathbf{x}}_k}}}\right)
\end{aligned}$$

with $\delta_{\mathbf{x}_k} \sim \mathcal{N}(0,1)$ and $\delta_{\hat{\mathbf{x}}_k} \sim \mathcal{N}(0,1)$. In GP-based noise model, we have $\varepsilon(\mathbf{x}_k) \sim \mathcal{N}(0, \exp(g(\mathbf{x}_k)))$ and $\varepsilon(\hat{\mathbf{x}}_k) \sim \mathcal{N}(0, \exp(g(\mathbf{x}_k)))$. Since we aim to obtain $\mathbf{g}_{\text{MAP}}$, we perform the parameterization trick on $\varepsilon(\mathbf{x}_k)$ and $\varepsilon(\hat{\mathbf{x}}_k)$ (third row). Taking the first and the second derivatives of $S(\mathbf{g})$ respectively gives us

$$
\begin{aligned}
\frac{\partial S(\mathbf{g})}{\partial \mathbf{g}} &= \begin{bmatrix} -\frac{0.5\,\phi(\mathbf{z})(\delta_{\mathbf{x}}\sqrt{\exp(g(\mathbf{x}))} - \delta_{\hat{\mathbf{x}}}\sqrt{\exp(g(\hat{\mathbf{x}}))})}{\Phi(\mathbf{z})\sqrt{\Sigma^*_{f\mathbf{x},\mathbf{x}} + \Sigma^*_{f\hat{\mathbf{x}},\hat{\mathbf{x}}} - 2\Sigma^*_{f\mathbf{x},\hat{\mathbf{x}}}}} \\ \frac{0.5\,\phi(\mathbf{z})(\delta_{\mathbf{x}}\sqrt{\exp(g(\mathbf{x}))} - \delta_{\hat{\mathbf{x}}}\sqrt{\exp(g(\hat{\mathbf{x}}))})}{\Phi(\mathbf{z})\sqrt{\Sigma^*_{f\mathbf{x},\mathbf{x}} + \Sigma^*_{f\hat{\mathbf{x}},\hat{\mathbf{x}}} - 2\Sigma^*_{f\mathbf{x},\hat{\mathbf{x}}}}} \end{bmatrix} + L^{-1}\mathbf{g} \\
\frac{\partial^2 S(\mathbf{g})}{\partial^2 \mathbf{g}} &= \Lambda + L^{-1} \\
\Lambda &= \begin{bmatrix} H & -H \\ -H & H \end{bmatrix}
\end{aligned}
$$

Let $\hat{\mathbf{z}} \triangleq 0.5\,\delta_{\mathbf{x}}\sqrt{\exp(g(\mathbf{x}))} - 0.5\,\delta_{\hat{\mathbf{x}}}\sqrt{\exp(g(\hat{\mathbf{x}}))}$, we define H as

$$
H = \frac{-1}{\sqrt{\Sigma^*_{f\mathbf{x},\mathbf{x}} + \Sigma^*_{f\hat{\mathbf{x}},\hat{\mathbf{x}}} - 2\Sigma^*_{f\mathbf{x},\hat{\mathbf{x}}}}} \frac{[\phi(\mathbf{z})\hat{\mathbf{z}}]'\Phi(\mathbf{z}) - \phi(\mathbf{z})\hat{\mathbf{z}}[\Phi(\mathbf{z})]'}{\Phi^2(\mathbf{z})}
$$

with

$$
\begin{aligned}
[\phi(\mathbf{z})\hat{\mathbf{z}}]' &= -\frac{\hat{\mathbf{z}}^2\phi(\mathbf{z})\mathbf{z}}{\sqrt{\Sigma^*_{f\mathbf{x},\mathbf{x}} + \Sigma^*_{f\hat{\mathbf{x}},\hat{\mathbf{x}}} - 2\Sigma^*_{f\mathbf{x},\hat{\mathbf{x}}}}} + \frac{1}{2}\phi(\mathbf{z})\hat{\mathbf{z}} \\
&= \left(-\frac{\hat{\mathbf{z}}^2\mathbf{z}}{\sqrt{\Sigma^*_{f\mathbf{x},\mathbf{x}} + \Sigma^*_{f\hat{\mathbf{x}},\hat{\mathbf{x}}} - 2\Sigma^*_{f\mathbf{x},\hat{\mathbf{x}}}}} + \frac{1}{2}\hat{\mathbf{z}} \right) \phi(\mathbf{z}) \\
\Phi'(\mathbf{z}) &= \frac{\phi(\mathbf{z})\hat{\mathbf{z}}}{\sqrt{\Sigma^*_{f\mathbf{x},\mathbf{x}} + \Sigma^*_{f\hat{\mathbf{x}},\hat{\mathbf{x}}} - 2\Sigma^*_{f\mathbf{x},\hat{\mathbf{x}}}}}
\end{aligned}
$$

The model selection problem is the same as in the simple model. Finally, we define the predictive distribution as

$$
\begin{aligned}
\mu^*_{\mathbf{g}} &= L_{*\mathbf{x}} L_{\mathbf{xx}} \mathbf{g}_{\text{MAP}} \\
\Sigma^*_{\mathbf{g}} &= L_{**} - L_{*\mathbf{x}}(L_{\mathbf{xx}} + \Lambda^{-1}_{\text{MAP}})^{-1} L_{*\mathbf{x}}
\end{aligned}
$$

References

Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

Wei Chu and Zoubin Ghahramani. Preference learning with gaussian processes. In *Proceedings of the 22nd international conference on Machine learning*, 2005.

Afonso Eduardo and Michael U. Gutmann. Bayesian optimization with informative covariance. *Transactions on Machine Learning Research*, 2023.

Javier Gonzalez, Zhenwen Dai, Andreas Damianou, and Neil D. Lawrence. Preferential bayesian optimization, 2017.

Ryan-Rhys Griffiths, Alexander Aldrick, Miguel Garcia-Ortegon, Vidhi Lalchand, and Alpha Lee. Achieving robustness to aleatoric uncertainty with heteroscedastic bayesian optimisation. *Machine Learning: Science and Technology*, 2021.

Kristian Kersting, Christian Plagemann, Patrick Pfaff, and Wolfram Burgard. Most-likely heteroscedastic gaussian process regression. In *Proceedings of the 24th International Conference on Machine Learning*, 2007.

Anastasia Makarova, Ilnura Usmanova, Ilija Bogunovic, and Andreas Krause. Risk-averse heteroscedastic bayesian optimization. *Advances in Neural Information Processing Systems*, 34:17235–17245, 2021.

Shion Takeno, Masahiro Nomura, and Masayuki Karasuyama. Towards practical preferential Bayesian optimization with skew Gaussian processes. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.