

ACCEPTED MANUSCRIPT • OPEN ACCESS

Achieving Robustness to Aleatoric Uncertainty with Heteroscedastic Bayesian Optimisation

To cite this article before publication: Ryan-Rhys Griffiths *et al* 2021 *Mach. Learn.: Sci. Technol.* in press <https://doi.org/10.1088/2632-2153/ac298c>

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is © 2021 The Author(s). Published by IOP Publishing Ltd.

As the Version of Record of this article is going to be / has been published on a gold open access basis under a CC BY 3.0 licence, this Accepted Manuscript is available for reuse under a CC BY 3.0 licence immediately.

Everyone is permitted to use all or part of the original content in this article, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by/3.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions may be required. All third party content is fully copyright protected and is not published on a gold open access basis under a CC BY licence, unless that is specifically stated in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

Achieving Robustness to Aleatoric Uncertainty with Heteroscedastic Bayesian Optimisation

Ryan-Rhys Griffiths¹, Alexander A Aldrick¹, Miguel Garcia-Ortegon^{2,3}, Vidhi Lalchand² and Alpha A. Lee¹

¹ Department of Physics, University of Cambridge

² Department of Engineering, University of Cambridge

³ Department of Mathematics University of Cambridge

E-mail: rrg27@cam.ac.uk

Abstract. Bayesian optimisation is a sample-efficient search methodology that holds great promise for accelerating drug and materials discovery programs. A frequently-overlooked modelling consideration in Bayesian optimisation strategies however, is the representation of heteroscedastic aleatoric uncertainty. In many practical applications it is desirable to identify inputs with low aleatoric noise, an example of which might be a material composition which displays robust properties in response to a noisy fabrication process. In this paper, we propose a heteroscedastic Bayesian optimisation scheme capable of representing and minimising aleatoric noise across the input space. Our scheme employs a heteroscedastic Gaussian process (GP) surrogate model in conjunction with two straightforward adaptations of existing acquisition functions. First, we extend the augmented expected improvement (AEI) heuristic to the heteroscedastic setting and second, we introduce the aleatoric noise-penalised expected improvement (ANPEI) heuristic. Both methodologies are capable of penalising aleatoric noise in the suggestions. In particular, the ANPEI acquisition yields improved performance relative to homoscedastic Bayesian optimisation and random sampling on toy problems as well as on two real-world scientific datasets. Code is available at: <https://github.com/Ryan-Rhys/Heteroscedastic-BO>

Keywords: Bayesian Optimisation, Gaussian Processes, Heteroscedasticity

1. Introduction

Bayesian optimisation is proving to be a highly effective search methodology in areas such as drug discovery [1, 2, 3], materials discovery [4, 5, 6], chemical reaction optimisation [7, 8, 9], robotics [10], sensor placement [11], tissue engineering [12] and genetics [13]. Heteroscedastic aleatoric noise however, is rarely accounted for in these settings despite being an important consideration for real-world applications. Aleatoric uncertainty refers to uncertainty inherent in the observations (measurement noise) [14]. In contrast, epistemic uncertainty corresponds to model uncertainty and may be explained away given sufficient data. Heteroscedastic aleatoric noise refers to aleatoric noise which varies across the input domain and is a prevalent feature of many scientific datasets; perhaps suprisingly not only experimental datasets, but also datasets where properties are predicted computationally. One such source of heteroscedasticity in the computational case might be situations in which the accuracy of first-principles calculations deteriorate as a function of the chemical complexity of the molecule being studied [15].




Figure 1. (a) The density histogram of computational errors (kcal/mol) for the FreeSolv hydration energy dataset ([16]). The computational errors in the hydration free energy arise from systematic errors in the force field used in alchemical free energy calculations based on classical molecular dynamics (MD) simulations. (b) A similar density histogram for the experimental errors where the source of uncertainty stems from the instrumentation used to obtain the measurement. The histograms are overlaid with kernel density estimates.

In Figure 1 we illustrate real-world sources of heteroscedasticity using the FreeSolv dataset of [16]. The consequences of misrepresenting heteroscedastic noise as being homoscedastic, i.e. constant across the input domain, are illustrated using a second dataset [17] in Figure 2. The homoscedastic model can underestimate noise in certain regions of the input space which in turn could induce a Bayesian optimisation scheme to suggest values possessing large aleatoric noise. In an application such as high-throughput virtual screening [18] the cost of misrepresenting noise during the screening process could

Robust Bayesian Optimisation




Figure 2. Comparison of homoscedastic and heteroscedastic GP fits to the soil phosphorus fraction dataset [17].

lead to a substantial loss of time in material fabrication [19]. In this paper we present a heteroscedastic Bayesian optimisation algorithm capable of both representing and minimising aleatoric noise in its suggestions. Our contributions are:

- (1) The introduction of a novel combination of surrogate model and acquisition function designed to minimise heteroscedastic aleatoric uncertainty.
- (2) A demonstration of our scheme’s ability to outperform naive schemes based on homoscedastic Bayesian optimisation and random sampling on toy problems as well as two real-world scientific datasets.
- (3) The provision of an open-source implementation.

The paper is structured as follows: section 2 introduces related work on heteroscedastic Bayesian optimisation. Section 3 provides background on Bayesian optimisation and homoscedastic GP surrogate models. Section 4 provides background on the heteroscedastic GP surrogate model used in this work and introduces the novel HAEI and ANPEI acquisitions functions. Section 5 considers experiments on synthetic and scientific datasets possessing heteroscedastic noise where the goal is to be robust to, i.e. minimise, aleatoric noise in the suggestions. Section 6 presents an ablation study on noiseless tasks as well as tasks with homoscedastic and heteroscedastic noise in order to determine whether there is a detrimental effect to using a heteroscedastic surrogate when the noise properties of the problem are a priori unknown. Section 7 concludes with some limitations of the approach presented as well as fruitful sources for future work.

2. Related Work

The most similar work to our own is that of [20] where experiments are reported on a heteroscedastic Branin-Hoo toy function using the variational heteroscedastic Gaussian

Robust Bayesian Optimisation

process (GP) approach of [21]. This work defines and optimises a robustness index, making a compelling case for penalisation of aleatoric noise in real-world Bayesian optimisation problems. A modification to expected improvement (EI), expected risk improvement is introduced in [22] and is applied to problems in robotics where robustness to aleatoric noise is desirable. In this framework however, the relative weights of performance and robustness cannot be tuned [20]. [23, 24] implement heteroscedastic Bayesian optimisation but do not introduce an acquisition function that penalises aleatoric noise. [25, 26] consider the related problem of safe Bayesian optimisation through implementing constraints in parameter space. In this instance, the goal of the algorithm is to enforce a performance threshold for each evaluation of the black-box function. Recently, the winners of the 2020 NeurIPS Black-Box Optimisation Competition applied non-linear output transformations in their solution to tackle heteroscedasticity. The authors however are not interested in explicitly penalising aleatoric noise in this case. In terms of acquisition functions, [27, 28] propose principled approaches to handling aleatoric noise in the homoscedastic setting that could be extended to the heteroscedastic setting. Our primary focus in this work however, is to highlight that heteroscedasticity in the surrogate model is beneficial and so an examination of a subset of acquisition functions is sufficient for this purpose. We take the opportunity here to note earlier unpublished workshop versions of this paper which consider the same problem [29, 30].

3. Background

3.1. Bayesian Optimisation

Bayesian optimisation [31, 32, 33] solves the global optimisation problem defined as

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$$

where $\mathbf{x}^*$ is the global optimiser of a black-box function $f : \mathcal{X} \rightarrow \mathcal{Y}$. $\mathcal{X}$ is the design space and is typically a compact subset of $\mathbb{R}^d$. What makes this optimisation problem practically relevant in applications are the following properties:

- (i) Black-Box Objective: We do not have the analytic form of f . We can however evaluate f pointwise anywhere in the design space $\mathcal{X}$.
- (ii) Expensive Evaluations: Choosing an input location $\mathbf{x}$ and evaluating $f(\mathbf{x})$ takes a very long time.
- (iii) Noise: The evaluation of a given $\mathbf{x}$ is a noisy process. In addition, this noise may vary across $\mathcal{X}$, making the underlying process heteroscedastic.

We have a dataset $\mathcal{D} = \{(\mathbf{x}_i, t_i)\}_{i=1}^n$ consisting of observations of the black-box function f and fit a probabilistic surrogate model to these datapoints. We then leverage the predictive mean as well as the uncertainty estimates of the surrogate model to guide the acquisition of the next data point $\mathbf{x}_{n+1}$ according to a heuristic known as an

Robust Bayesian Optimisation

5

acquisition function. In Bayesian optimisation, exact GPs are the most popular choice of surrogate model because of their ability to represent posterior uncertainty without resorting to approximate Bayesian inference.

3.2. Gaussian Processes

In the terminology of stochastic processes we may formally define a GP as follows:

Definition 1. A GP [34] is a collection of random variables, any finite number of which have a joint Gaussian distribution.

GPs can be used to set a prior over functions in Bayesian modelling applications. In this setting, the random variables consist of function values $f(\mathbf{x})$ at different locations $\mathbf{x}$ within the design space. The GP is characterised by a mean function

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$$

and a covariance function

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))].$$

The process is written as follows

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')).$$

In our experiments, the prior mean function will be set to the empirical mean of the data. The covariance function or kernel computes the pairwise covariance between two random variables (function values). The covariance between a pair of output values $f(\mathbf{x})$ and $f(\mathbf{x}')$ is a function of an input pair $\mathbf{x}$ and $\mathbf{x}'$. As such, the kernel encodes smoothness assumptions about the latent function being modelled. The most widely-utilised kernel is the squared exponential (SE) kernel

$$k_{\text{SQE}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right) \quad (1)$$

where σ_f^2 is the signal amplitude hyperparameter (vertical lengthscale) and ℓ is the (horizontal) lengthscale hyperparameter. Although Equation 1 is written with a single lengthscale shared across dimensions, for multidimensional input spaces we optimise a lengthscale per dimension. For consistency, we use the squared exponential kernel in all experiments reported in the main paper. In Appendix C we compare the performance of different kernels on a set of synthetic optimisation functions. For a more detailed introduction to GPs the reader is referred to [34].

4. Heteroscedastic Bayesian Optimisation

We wish to perform Bayesian optimisation whilst minimising input-dependent aleatoric noise. In order to represent input-dependent aleatoric noise, a heteroscedastic surrogate model is required.

4.1. The Most Likely Heteroscedastic Gaussian Process

We adopt the most likely heteroscedastic Gaussian process (MLHGP) approach of [35], and for consistency, we use the same notation as the source work in our presentation. We have a dataset $\mathcal{D} = \{(\mathbf{x}_i, t_i)\}_{i=1}^n$ in which the target values t_i have been generated according to $t_i = f(\mathbf{x}_i) + \epsilon_i$. We assume independent Gaussian noise terms $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$ with variances given by $\sigma_i^2 = r(\mathbf{x}_i)$. In the heteroscedastic setting r is typically a non-constant function over the input domain $\mathbf{x}$. In order to perform Bayesian optimisation, we wish to model the predictive distribution $P(\mathbf{t}^* | \mathbf{x}_1^*, \dots, \mathbf{x}_q^*)$ at the query points $\mathbf{x}_1^*, \dots, \mathbf{x}_q^*$. Placing a GP prior on f and taking $r(\mathbf{x})$ as the assumed noise function, the predictive distribution is multivariate Gaussian $\mathcal{N}(\mu^*, \Sigma^*)$ with mean

$$\mu^* = E[\mathbf{t}^*] = K^*(K + R)^{-1}\mathbf{t}$$

and covariance matrix

$$\Sigma^* = \text{var}[\mathbf{t}^*] = K^{**} + R^* - K^*(K + R)^{-1}K^{*T},$$

where $K \in \mathbb{R}^{n \times n}$, $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, $K^* \in \mathbb{R}^{q \times n}$, $K_{ij}^* = k(\mathbf{x}_i^*, \mathbf{x}_j)$, $K^{**} \in \mathbb{R}^{q \times q}$, $K_{ij}^{**} = k(\mathbf{x}_i^*, \mathbf{x}_j^*)$, $\mathbf{t} = (t_1, t_2, \dots, t_n)^T$, $R = \text{diag}(\mathbf{r})$ with $\mathbf{r} = (r(\mathbf{x}_1), r(\mathbf{x}_2), \dots, r(\mathbf{x}_n))^T$, and $R^* = \text{diag}(\mathbf{r}^*)$ with $\mathbf{r}^* = (r(\mathbf{x}_1^*), r(\mathbf{x}_2^*), \dots, r(\mathbf{x}_q^*))^T$.

The MLHGP algorithm [35] executes the following steps:

- (i) Estimate a homoscedastic GP, G_1 on the dataset $\mathcal{D} = \{(\mathbf{x}_i, t_i)\}_{i=1}^n$
- (ii) Given G_1 , we estimate the empirical noise levels for the training data using $z_i = \log(\text{var}[t_i, G_1(\mathbf{x}_i, \mathcal{D})])$ where $\text{var}[t_i, G_1(\mathbf{x}_i, \mathcal{D})] \approx \frac{1}{s} \sum_j^s 0.5 (t_i - t_i^j)^2$ with t_i^j a sample from the predictive distribution induced by the GP at $\mathbf{x}_i$, forming a new dataset $\mathcal{D}' = \{(\mathbf{x}_i, z_i)\}_{i=1}^n$
- (iii) Estimate a second GP, G_2 on $\mathcal{D}'$.
- (iv) Estimate a combined GP, G_3 on $\mathcal{D}$ using G_2 to predict the logarithmic noise levels r_i .
- (v) If not converged, set G_3 to G_1 and repeat.

In essence, the defining characteristic of the MLHGP approach is that G_1 learns the latent function and G_2 learns the noise function.

Robust Bayesian Optimisation

4.2. Bayesian Optimisation with Aleatoric Noise Penalisation

Our heteroscedastic Bayesian optimisation problem may be framed as

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} h(\mathbf{x}),$$

where the black-box objective h , to be minimised has the form

$$h(\mathbf{x}) = \alpha f(\mathbf{x}) + (1 - \alpha)g(\mathbf{x}).$$

where $f(\mathbf{x})$ is the black-box function of the principal objective i.e. the objective corresponding to classical Bayesian optimisation where noise is not optimised, and $g(\mathbf{x})$ is the latent heteroscedastic noise function which governs the magnitude of the noise at a given input location $\mathbf{x}$. α is a parameter chosen, for the purposes of evaluation, by a domain expert that trades off the weight of the principal objective relative to the noise objective. It is worth noting that α is a parameter that affects only the evaluation of an algorithm and not the execution. The evaluation criteria however, will dictate the optimal hyperparameters of the acquisition function.

4.3. Heteroscedastic Acquisition Functions

We investigate extensions of the expected improvement [33] acquisition criterion, the form of which may be written in terms of the targets t and the incumbent best objective function value, η , found so far as

$$\text{EI}(\mathbf{x}) = \mathbb{E}[(\eta - t)_+] = \int_{-\infty}^{\infty} (\eta - t)_+ p(t | \mathbf{x}) dt$$

where $p(t | \mathbf{x})$ is the posterior predictive marginal density of the objective function evaluated at $\mathbf{x}$. $(\eta - t)_+ \equiv \max(0, \eta - t)$ is the improvement over the incumbent best objective function value η . Evaluations of the objective are noisy in all of the problems we consider and so we use expected improvement with plug-in [36], the plug-in value being the GP predictive mean [37].

We propose two extensions to the expected improvement criterion. The first is an extension of the augmented expected improvement criterion

$$\text{AEI}(\mathbf{x}) = \mathbb{E}[(\eta - t)_+] \left(1 - \frac{\sigma_n}{\sqrt{\text{var}[t] + \sigma_n^2}} \right),$$

of [38] where σ_n is the fixed aleatoric noise level. AEI is introduced as a heuristic for the optimisation of noisy functions. EI is recovered in the case that $\sigma_n^2 = 0$ and in the case that $\sigma_n^2 > 0$ AEI operates as a rescaling of the EI acquisition function, penalising test locations where the GP predictive variance is small relative to the fixed noise level σ_n^2 .

Robust Bayesian Optimisation

We extend AEI to the heteroscedastic setting by exchanging the fixed aleatoric noise level with the input-dependent one:

$$\text{HAEI}(\mathbf{x}) = \mathbb{E}[(\eta - t)_+] \left(1 - \frac{\gamma \sqrt{r(\mathbf{x})}}{\sqrt{\text{var}[t] + \gamma^2 r(\mathbf{x})}} \right), \quad (2)$$

where $r(\mathbf{x})$ is the predicted aleatoric uncertainty at input $\mathbf{x}$ under the MLHGP and $\text{var}[t]$ is the predictive variance of the MLHGP at input $\mathbf{x}$. γ in this instance is defined to be a positive penalty parameter for regions with high aleatoric noise.

Proposition 1 (Limit of Large Epistemic Uncertainty). *The HAEI acquisition function reduces to EI when the ratio of epistemic uncertainty to aleatoric uncertainty is much greater than γ^2 .*

Proof. Let $k = \frac{\text{var}[t]}{r(\mathbf{x})}$ denote the ratio of epistemic to aleatoric uncertainty at an arbitrary input location $\mathbf{x}$. Dividing the numerator and the denominator of the second term in the second factor of Equation 2 by $\sqrt{r(\mathbf{x})}$ yields

$$\text{HAEI}(\mathbf{x}) = \text{EI}(\mathbf{x}) \left(1 - \frac{\gamma}{\sqrt{k + \gamma^2}} \right). \quad (3)$$

Taking the limit analytically as k tends to infinity and assuming finite γ

$$\lim_{k \rightarrow \infty} \text{EI}(\mathbf{x}) \left(1 - \frac{\gamma}{\sqrt{k + \gamma^2}} \right) = \text{EI}(\mathbf{x}),$$

recovers the expected improvement acquisition. □

Proposition 2 (Limit of Large Aleatoric Uncertainty). *The HAEI acquisition function goes to zero as the ratio of epistemic uncertainty to aleatoric uncertainty goes to zero.*

Proof. Taking the limit as k tends to zero in Equation 3 yields

$$\lim_{k \rightarrow 0} \text{EI}(\mathbf{x}) \left(1 - \frac{\gamma}{\sqrt{k + \gamma^2}} \right) = 0. \quad \square$$

Remark. *In the limit of large aleatoric uncertainty there is an approximation that is linear in k for the HAEI scaling factor.*

Letting $S(k) = 1 - \frac{\gamma}{\sqrt{k + \gamma^2}}$ such that $\text{HAEI} = \text{EI}(\mathbf{x})S(k)$, consider the MacLaurin expansion of $S(k)$,

$$S(k) = S(0) + S'(0)k + \frac{S''(0)}{2!}k^2 + \frac{S'''(0)}{3!}k^3 + \dots,$$




Figure 3. The HAEI scaling factor $S(k)$, now written as a function of k for different values of γ . When k , the ratio of epistemic to aleatoric uncertainty is small, the scaling factor goes to zero to reflect the penalty for regions of high aleatoric uncertainty. γ controls the decay rate of this penalty. Also shown is the linear approximation to the scaling factor for $\gamma = 10$.

Dropping terms of $O(k^2)$ and higher we obtain

$$S(k) \approx \frac{k}{2\gamma^2}.$$

This approximation may be used when k is small relative to γ and could provide guidance in setting the γ parameter if prior knowledge about k and the desired trade-off between the principal and noise objectives is available. In [Figure 3](#) we provide insight into the effect that different values of γ will have on the scaling factor $S(k)$.

In addition to HAEI, we propose a simple modification to EI that explicitly penalises regions of the input space with large aleatoric noise. We call this acquisition function aleatoric noise-penalised expected improvement (ANPEI) and denote it

$$\text{ANPEI} = \beta \text{EI}(\mathbf{x}) - (1 - \beta) \sqrt{r(\mathbf{x})}, \quad (4)$$

where β is a scalarisation constant. In the multiobjective optimisation setting a particular value of β will correspond to a point on the Pareto frontier. We showcase the advantages of both HAEI and ANPEI acquisition functions in conjunction with the MLHGP surrogate model in [section 5](#).

Robust Bayesian Optimisation




Figure 4. Illustrative Toy Problem. The latent function in a) is corrupted with heteroscedastic Gaussian noise according to the function in b) where $g(x)$ is a constant multiplier of a sample from a standard Gaussian. The combined objective is given in c) and is obtained by subtracting the noise function from the latent function.




Figure 5. Noisy samples $y_i = f(x_i) + g(x_i)\epsilon$ from the heteroscedastic sin wave function.

5. Experiments on Robustness to Aleatoric Uncertainty

5.1. Implementation

Experiments were run using a custom NumPy [39] implementation of GP regression and MLHGP regression. All code to reproduce the experiments is available at <https://github.com/Ryan-Rhys/Heteroscedastic-BO>. The squared exponential kernel was chosen as the covariance function for both the homoscedastic GP as well as G_1 and G_2 of the MLHGP. Across all datasets, the lengthscales, ℓ , of the homoscedastic GP were initialised to 1.0 for each input dimension. The signal amplitude σ_f^2 was initialised to a value of 1.0. The lengthscale, ℓ , of G_2 of the MLHGP [35] was initialised to 1.0, the initial noise level of G_2 was set to 1.0. The EM-like procedure required to train the MLHGP was run for 10 iterations and the sample size required to construct the variance estimator producing the auxiliary dataset was 100. All standard error confidence bands are computed using 50 independent random seed initialisations. Hyperparameter values, including the noise level of the homoscedastic GP, were obtained by optimising the

Robust Bayesian Optimisation

11

marginal likelihood using the scipy implementation of the L-BFGS-B optimiser [40], taking the best of 20 random restarts. The objective function is

$$h(x) = \alpha f(x) - (1 - \alpha)g(x)$$

for the one-dimensional sin wave experiment which is a maximisation problem and as such has a subtractive penalty for regions of large noise. For the remaining experiments, which are minimisation problems, the objective is

$$h(\mathbf{x}) = \alpha f(\mathbf{x}) + (1 - \alpha)g(\mathbf{x}) \quad (5)$$

The sin wave and Branin-Hoo tasks are initialised with 25 and 100 data points respectively drawn uniformly at random within the bounds of the design space. The soil and FreeSolv experiments are initialised with 36 and 129 data points respectively drawn uniformly at random from the datasets. α is set to 0.5 for all experiments while β is set to 0.5, $\frac{1}{11}$, 0.5 and 0.5 for the sin, Branin-Hoo, soil and FreeSolv experiments. γ is set to 1, 500, 1 and 1 for the sin, Branin-Hoo, soil and FreeSolv experiments. We run 5 acquisition functions in all experiments: random sampling, homoscedastic EI, AEI, HAEI and ANPEI. Homoscedastic EI is included as a baseline to demonstrate the difference consideration of aleatoric noise yields in the optimisation of the objective. AEI is included to demonstrate the difference consideration of heteroscedastic aleatoric noise yields and random sampling is included as a baseline as it is known to be highly competitive with Bayesian optimisation in noisy settings.

5.2. Heteroscedastic Sin Wave Function

The objective function has the form

$$h(x) = f(x) - g(x)$$

where $f(x) = \sin(x) + 0.2(x) + 3$ and $g(x) = 0.5(x)$. In this instance α from subsection 5.1 has a setting of 0.5 but we omit it explicitly as the objectives have equal weight. Over the course of the experiment samples

$$y_i = f(x_i) + g(x_i)\epsilon, \quad \epsilon \sim \mathcal{N}(0, 1)$$

are observed. The problem setup is depicted in Figure 4 and Figure 5. The Bayesian optimisation problem is constructed such that the first maximum in Figure 4(a) is to be preferred as samples from this region of the input space will have low aleatoric noise. The black-box objective in Figure 4(c) illustrates this trade-off. In Figure 6 we compare the performance of all surrogate model/acquisition function combinations. We observe the low aleatoric noise-seeking behaviour of HAEI and ANPEI on $g(x)$ as well as their ability to optimise the composite objective $h(x)$.

Robust Bayesian Optimisation




Figure 6. Comparison of heteroscedastic and homoscedastic Bayesian optimisation on the sin wave problem. (a) shows the optimisation of $h(x) = f(x) - g(x)$ (higher is better) whereas (b) shows the values $g(x)$ obtained over the course of the optimisation of $h(x)$. This latter plot demonstrates the propensity of ANPEI to seek low aleatoric noise solutions.

5.3. Heteroscedastic Branin-Hoo Function

In the second experiment we consider the objective

$$h(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x})$$

with an additive penalty because the task is a minimisation problem and an α setting of 0.5 for equal-weight objectives.

$$f(\mathbf{x}) = \frac{1}{51.95} \left[\left(\bar{x}_2 - \frac{5.1\bar{x}_1^2}{4\pi^2} + \frac{5\bar{x}_1}{\pi} - 6 \right)^2 + \left(10 - \frac{10}{8\pi} \right) \cos(\bar{x}_1) - 44.81 \right] \quad (6)$$

with $\bar{x}_1 = 15x_1 - 5$, $\bar{x}_2 = 15x_2$ and $\mathbf{x} = (x_1, x_2)$ is the standardised Branin-Hoo function introduced in [36]. The noise function $g(\mathbf{x})$ is in this instance

$$g(\mathbf{x}) = 15 - 8x_1 + 8x_2^2. \quad (7)$$

Samples are again generated according to

$$y_i = f(\mathbf{x}_i) + g(\mathbf{x}_i)\epsilon, \quad \epsilon \sim \mathcal{N}(0, 1)$$

Robust Bayesian Optimisation

13




Figure 7. Branin-Hoo Optimisation Problem. The latent function in a) is corrupted by heteroscedastic Gaussian noise function according to the function in b) The combined objective function is given in c) and is obtained by summing the functions in a) and b). The sum is required to penalise regions of large aleatoric noise because the objective is being minimised.




Figure 8. Comparison of heteroscedastic and homoscedastic Bayesian optimisation on the Branin-Hoo problem. (a) shows the optimisation of $h(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x})$ (lower is better) whereas (b) shows the values $g(\mathbf{x})$ obtained over the course of the optimisation of $h(\mathbf{x})$.

The problem setup is shown in Figure 7 and the performance of all surrogate model/acquisition function pairs is depicted in Figure 8. The gulf in performance between the heteroscedastic and homoscedastic surrogate models is more pronounced in this case because the noise function is more severe relative to the sin wave problem.

5.4. Soil Phosphorus Fraction Optimisation

In this experiment we consider the optimisation of the phosphorus fraction of soil. Soil phosphorus is an essential nutrient for plant growth and is widely used as a fertiliser in agriculture. While the amount of arable land worldwide is declining, global population is expanding concomitantly with food demand. As such, understanding the availability of plant nutrients that increase crop yield is a topic worthy of attention. To this end, [17] have curated a dataset on soil phosphorus, relating phosphorus content to variables such as soil particle size, total nitrogen, organic carbon and bulk density. We choose to study the relationship between bulk soil density and the phosphorus fraction, the goal being to minimise the phosphorus content of soil subject to heteroscedastic noise. In lieu of performing a formal test for heteroscedasticity, we provide evidence that there is heteroscedasticity in the dataset by comparing the fits of a homoscedastic GP and the MLHGP in Figure 2 and provide a predictive performance comparison based on negative log predictive density values in Appendix A.

In this problem, we do not have access to a continuous-valued black-box function or a ground truth noise function. As such, the surrogate models were initialised with a subset of the data and the query locations selected by Bayesian optimisation were mapped to the closest datapoints in the heldout data. The following kernel smoothing procedure was used to generate pseudo ground-truth noise values:

- (1) Fit a homoscedastic GP to the full dataset.
- (2) At each point x_i , compute the corresponding squared error $s_i^2 = (y_i - \mu(x_i))^2$.
- (3) Estimate variances by computing a moving average of the squared errors, where the relative weight of each s_i^2 was assigned with a Gaussian kernel.

The performances of heteroscedastic and homoscedastic Bayesian optimisation are compared in Figure 9. Given that regions of low phosphorus fraction coincide with regions of small aleatoric noise, we apply an α value of $\frac{1}{6}$ to the composite objective $h(x)$ to admit a finer granularity for distinguishing between degrees of low aleatoric noise in the solutions.

5.5. Molecular Hydration Free Energy Optimisation

We perform a retrospective virtual screening experiment with the aim of identifying molecules with favourable hydration free energy, a property important in determining the binding affinity of a drug candidate. Experiments were performed with an initialisation of 129 out of the 642 molecules in the FreeSolv dataset [41, 16] over 10 iterations of data collection. Unlike the soil phosphorus fraction dataset, ground truth measurement error (aleatoric noise $g(\mathbf{x})$) values are available for the FreeSolv dataset. The remaining 513 molecules were reserved as a heldout set where at each iteration of data collection one of the heldout molecules was selected. Chemical fragments computed using RDKit [42] were used as the molecular representation based on the fact that these global features, unlike local Morgan fingerprints, act as good predictors of the hydration free energy.

Robust Bayesian Optimisation

15





Figure 9. Comparison of heteroscedastic and homoscedastic Bayesian optimisation on the soil phosphorus fraction optimisation problem. (a) shows the optimisation of $h(x) = f(x) + g(x)$ (lower is better) where x is the dry bulk density of the soil. (b) shows the values $g(x)$ obtained over the course of the optimisation of $h(x)$.

The fragment features were projected down to 14 components using principal component analysis, retaining more than 90% of the variance on average across random trials. The results are shown in Figure 10. Compared to previous experiments, the noise is smaller in this instance relative to the magnitude of the hydration free energy (Signal-to-noise ratio of approximately 10) and as such the heteroscedastic modelling problem is more difficult, leading to only very marginal gains in obtaining low noise solutions. While ANPEI obtains the lowest objective function value over the Bayesian optimisation trace, the results are unlikely to be statistically significant according to the standard error bands.

5.6. Heteroscedastic Acquisition Function Hyperparameters

The β hyperparameter of ANPEI in Equation 4 and the γ hyperparameter of HAEI in Equation 3 are designed to modulate the avoidance of aleatoric noise in the acquisitions. In Figure 11 we offer some intuition as to the effect of various settings of β and γ by examining the heteroscedastic Branin-Hoo function introduced in subsection 5.3. The results demonstrate that the performance of the algorithms is strongly dependent on the setting of the β hyperparameter for ANPEI whereas γ is less influential on the performance of HAEI. It is worth noting in Figure 11(b) that if too large a value of γ is chosen the principal objective $f(\mathbf{x})$ may be compromised through overly aggressive avoidance of aleatoric noise. In practice choosing the value of β in line with the value of the evaluation criterion parameter α in Equation 5 is likely to be a sensible approach i.e.

Robust Bayesian Optimisation



(a) Best Objective Value Found so Far

(b) Lowest Aleatoric Noise Found so Far

Figure 10. Comparison of heteroscedastic and homoscedastic Bayesian optimisation on the FreeSolv hydration free energy optimisation problem. (a) shows the optimisation of $h(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x})$ (lower is better) where $\mathbf{x}$ is the fragment set of molecular descriptors, $f(\mathbf{x})$ is the hydration free energy and $g(\mathbf{x})$ is the aleatoric noise. (b) shows the values $g(\mathbf{x})$ obtained over the course of the optimisation of $h(\mathbf{x})$.

if the noise objective is more important relative to the principal objective by a factor of 10 then the value of β should be $\frac{1}{11}$.

5.7. Robustness Experiments Summary

The experiments of this section provide strong evidence that modelling heteroscedasticity in Bayesian optimisation is a useful approach for problems in which there is a strong degree of aleatoric noise present. The ANPEI acquisition tends to outperform HAEI on the majority of the tasks where there is a small degree of aleatoric noise whilst the acquisitions are more evenly matched when the extent of the aleatoric noise is high. The outstanding questions for these methods however, is how well they perform on tasks where heteroscedastic noise is not present. Such a situation may easily arise for real-world problems where the noise properties of the tasks are a priori unknown and as such, it is important to ascertain whether there is a deleterious effect on performance in noiseless and homoscedastic noise settings.

6. Ablation Study on Noiseless, Homoscedastic Noise and Heteroscedastic Noise Tasks

In this section we perform an ablation study where components of the ablation constitute different noise properties. We examine the noiseless case as a base task before adding first a homoscedastic noise component and second, a heteroscedastic noise component.




Figure 11. Performance of ANPEI and HAEI plotted for different values of the β and γ hyperparameters respectively. *Smaller* values of β encourage avoidance of regions of high aleatoric noise whilst *larger* values of γ encourage avoidance of regions of high aleatoric noise.

Additionally, we examine the effect of the size of the initialisation grid on performance in the heteroscedastic noise tasks.

6.1. Ablation

The ablation study makes use of three synthetic optimisation functions: The Branin-Hoo function, the Hosaki function and the Goldstein-Price function. The form of the Branin-Hoo function is the same standardised Branin-Hoo function introduced in Equation 6 with heteroscedastic noise function given in Equation 7. The Hosaki function, defined on the domain $x_1, x_2 \in [0, 5]$, is

$$\text{Hosaki}(x_1, x_2) = \left(1 - 8x_1 + 7x_1^2 - \frac{7}{3}x_1^3 + \frac{1}{4}x_1^4\right)x_2^2 \exp(-x_2).$$

To facilitate the GP fit, the Hosaki function is subsequently standardised by its mean (0.817) and standard deviation (0.573). The noise function is

$$g_{\text{Hosaki}}(x_1, x_2) = 50 \cdot \frac{1}{(x_1 - 3.5)^2 + 2.5} \cdot \frac{1}{(x_2 - 2)^2 + 2.5}. \quad (8)$$

The logarithmic Goldstein-Price function [36] is

Robust Bayesian Optimisation




Figure 12. (a) The latent Hosaki Function $f(\mathbf{x})$ together with (b) its heteroscedastic noise function $g(\mathbf{x})$ and (c) the objective function $f(\mathbf{x}) + g(\mathbf{x})$.

$$G-P(x_1, x_2) = \frac{1}{2.427} \left[\log \left([1 + (\bar{x}_1 + \bar{x}_2 + 1)^2 (19 - 14\bar{x}_1 + 3\bar{x}_1^2 - 14\bar{x}_2 + 6\bar{x}_1\bar{x}_2 + 3\bar{x}_2^2)] \right. \right. \\ \left. \left. [30 + (2\bar{x}_1 - 3\bar{x}_2)^2 (18 - 32\bar{x}_1 + 12\bar{x}_1^2 + 48\bar{x}_2 - 36\bar{x}_1\bar{x}_2 + 27\bar{x}_2^2)] \right) - 8.693 \right] \quad (9)$$

$$(10)$$

where $\bar{x}_1 = 4x_1 - 2$ and $\bar{x}_2 = 4x_2 - 2$. The Goldstein-Price noise function is

$$g_{G-P}(x_1, x_2) = \frac{3}{2} \cdot \frac{1}{(x_1 - 0.5)^2 + 0.2} \cdot \frac{1}{(x_2 - 0.3)^2 + 0.3}. \quad (11)$$

For clarity, only the results of the Hosaki function are presented in the main paper with the Branin-Hoo and Goldstein-Price results presented in [Appendix B](#). The Hosaki function is visualised in [Figure 12](#). The value of β for ANPEI is set to 0.5 and the value of γ is set to 500 for all Hosaki function experiments.

6.1.1. Noiseless Case In this case, the synthetic functions do not possess any observation noise and the optimisation function corresponds to the situation in [Figure 12\(a\)](#). 9 points sampled uniformly at random are used for initialisation and the results are displayed in [Figure 13](#). As expected, all Bayesian optimisation methods outperform random search in the noiseless case. In this example it is unclear as to whether heteroscedastic Bayesian optimisation methods are detrimental as HAEI performs best whereas ANPEI performs worst.

6.1.2. Homoscedastic Noise Case In this case the functions are subject to homoscedastic noise of the form 25ϵ where epsilon is noise sampled from a standard Gaussian $\mathcal{N}(0, 1)$. The GP surrogates are again initialised with 9 points. The results are displayed in




Figure 13. Hosaki function noiseless case. All Bayesian optimisation methods outperform random search. HAEI performs best and ANPEI performs worst.

Figure 14. The Bayesian optimisation methods perform worse in the homoscedastic noise case relative to the noiseless case although the rank order of the methods mirrors that of the noiseless case.

6.1.3. Heteroscedastic Noise In the heteroscedastic noise case the Hosaki function is subject to the noise function given in Equation 8 and visualised in Figure 12. 144 points were used to initialise the GP surrogates. The results are shown in Figure 15. In this instance, given that the extent of heteroscedastic noise is very strong (relative to the homoscedastic noise case), random search is highly competitive with the Bayesian optimisation methods. ANPEI however, is the best-performing algorithm. The large number of initialisation points chosen for this experiment reflects one limitation of the heteroscedastic surrogate approach; for the MLHGP to effectively learn a decomposition of the function into signal and noise components it needs access to more samples. As such, this merits an investigation into the effect of the number of samples on the performance of the heteroscedastic acquisitions.

6.2. Effect of Initialisation Set Size

The effect of the size of the initialisation set on the heteroscedastic Branin-Hoo task is investigated in Figure 16. The value of β used for ANPEI is $\frac{1}{11}$ and the value of γ used for HAEI is 500. The performance of the heteroscedastic acquisitions ANPEI and




Figure 14. Hosaki function homoscedastic noise case. All Bayesian optimisation methods outperform random search with HAEI the best and ANPEI the worst.




Figure 15. Comparison of heteroscedastic and homoscedastic Bayesian optimisation on the heteroscedastic 2D Hosaki function. (a) shows the optimisation of $h(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x})$ (lower is better) where $g(\mathbf{x})$ is the aleatoric noise. (b) shows the values $g(\mathbf{x})$ obtained over the course of the optimisation of $h(\mathbf{x})$.

Robust Bayesian Optimisation

21




Figure 16. The effect of the initialisation set size on the heteroscedastic Branin-Hoo function. The performance of heteroscedastic acquisitions ANPEI and HAEI increases as they are given access to more samples. An excess of samples do not help the homoscedastic Bayesian optimisation methods as they are unable to model the heteroscedastic noise component.

HAEI is observed to improve as the size of the initialisation set increases. In contrast, the homoscedastic methods EI and AEI do not improve on obtaining access to more samples as they are unable to model the heteroscedastic noise component of the task.

6.3. Conclusions from Ablation Experiments

Synthesising the results from the additional ablation experiments in [Appendix B](#) some trends may be observed:

- (i) All Bayesian optimisation methods outperform random search in the noiseless case and homoscedastic noise cases on aggregate across the three synthetic functions.
- (ii) On aggregate there is no significant difference between Bayesian optimisation methods in the noiseless or homoscedastic noise cases (HAEI marginally outperforms ANPEI on 2/3 noiseless tasks and 2/3 homoscedastic noise tasks).
- (iii) The heteroscedastic acquisitions ANPEI and HAEI perform competitively on the noiseless and homoscedastic noise tasks most likely because the MLHGP is capable of effecting nonstationary behaviour by "fantasising" heteroscedastic noise. As such, the MLHGP surrogate may be achieving enhanced flexibility relative to the homoscedastic GP in this setting.
- (iv) The heteroscedastic acquisitions tend to outperform other Bayesian optimisation approaches on the heteroscedastic noise tasks although crucially this depends on the size of the initialisation set. In order to detect heteroscedastic noise the MLHGP surrogate needs access to more samples relative to the noiseless and homoscedastic cases.
- (v) ANPEI outperforms HAEI.

In summary, the experiments would appear to show that there is no significant downside to employing a heteroscedastic surrogate and acquisition function on noiseless tasks or tasks with homoscedastic noise save for the increased training time for the model.

7. Conclusions

We have presented an approach for performing Bayesian optimisation with the explicit goal of minimising aleatoric noise in the suggestions. We posit that such an approach can prove useful for the natural sciences in the search for molecules and materials that are robust to experimental measurement noise. The synthetic function ablation study highlights no particular downside to the use of the MLHGP in conjunction with ANPEI or HAEI in cases where the noise structure of the problem is a priori unknown i.e the black-box optimisation problem is either noiseless or homoscedastic. Nonetheless, we anticipate that this type of approach may be particularly relevant for the experimental natural sciences where noiseless objectives or those with homoscedastic noise are highly uncommon. In terms of concrete recommendations on when to apply the algorithm, we foresee the best performance in situations where the user has access to a moderately-sized initialisation set in order to provide the MLHGP with enough samples to distinguish heteroscedastic noise from intrinsic function variability. There are a number of possible extensions to the current approach which may facilitate its application to high-dimensional datasets and act as fruitful sources for future work:

- (1) **Surrogate Model:** One disadvantage of the MLHGP model is the lack of convergence guarantees for the EM-like procedure required for fitting. Various other forms of heteroscedastic GP exist [43, 44, 45, 46, 47, 48, 49] and have demonstrated success in modelling applications [50, 51, 52, 53]. Of particular interest for real-world problems are scalable heteroscedastic GPs [54, 55] which could circumvent the computationally-intensive bottleneck of fitting multiple exact GPs as a subroutine of the MLHGP Bayesian optimisation procedure.
- (2) **Advances in Surrogate Model Machinery:** Advances in areas such as efficient sampling of GPs [56] are liable to yield improvements to sampled-based acquisition functions such as Thompson sampling [57] while fully Bayesian approaches to hyperparameter estimation for sparse GPs [58] are liable to yield improvements in model fitting procedures.
- (3) **Scalable Bayesian Optimisation:** Scalable Bayesian optimisation can also be enabled via dimensionality reduction techniques [59, 60, 61]. Such approaches, when combined with efficient libraries [62, 63] could facilitate heteroscedastic Bayesian optimisation in high-dimensional settings.
- (4) **Acquisition Function Optimisation:** Recent developments in acquisition function optimisation including Monte Carlo reformulations [64, 65], compositional

optimisers [66, 65] and tight relaxations [67] of common acquisition functions have the potential to yield gains in empirical performance.

- (5) **Data Transformation:** Input-warping [68] and output transformations [69] have recently shown success towards addressing heteroscedastic datasets.
- (6) **Approaches for Molecular Bayesian Optimisation:** In relation to molecules, the use of tailored GP kernels such as Tanimoto kernels [70, 71] and more expressive dimensionality reduction techniques [72] could lead to performance gains and enhanced scalability respectively.
- (7) **Exploration in the Noise Objective:** Incorporating exploration in the noise objective in the multi-objective setting as in [22].

Lastly, a further use-case of the machinery developed in this paper is obtained by turning the noise minimisation problem into a noise maximisation problem. As an example, in materials discovery, we may derive benefit from being antifragile [73] towards (i.e. derive benefit from) high aleatoric noise. In an application such as the search for performant perovskite solar cells, we are faced with an extremely large compositional space, with millions of potential candidates possessing high aleatoric noise for identical reproductions [74]. In this instance we may want to guide search towards a candidate possessing a high photoluminescence quantum efficiency with high aleatoric noise. If the cost of repeating material syntheses is small relative to the cost of the search, the large aleatoric noise will present opportunities to synthesise materials possessing efficiencies far in excess of their mean values.

8. Acknowledgements

The authors would like to thank James T. Wilson for discussion about the experimental setup and applications domains of heteroscedastic Bayesian optimisation, Henry Moss for useful discussions regarding Bayesian optimisation in the noisy setting and Luke Corcoran for discussions relating to the limiting behaviour of the HAEI acquisition. We would additionally like to thank the anonymous reviewers who were instrumental in improving the quality of the empirical analysis as well as the clarity of the manuscript.

REFERENCES

24

References

- [1] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018.
- [2] Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chemical Science*, 11(2):577–586, 2020.
- [3] Samuel Hoffman, Vijil Chenthamarakshan, Kahini Wadhawan, Pin-Yu Chen, and Payel Das. Optimizing molecules using efficient queries from property evaluations. *arXiv preprint arXiv:2011.01921*, 2020.
- [4] Florian Hase, Loic M Roch, and Alan Aspuru-Guzik. Gryffin: An algorithm for Bayesian optimization for categorical variables informed by physical intuition with applications to chemistry. *arXiv preprint arXiv:2003.12127*, 2020.
- [5] Florian Hase, Matteo Aldeghi, Riley Hickman, Loic Roch, Elena Liles, Melodie Christensen, Jason Hein, and Alán Aspuru-Guzik. Olympus: a benchmarking framework for noisy optimization and experiment planning. *Machine Learning: Science and Technology*, 2021.
- [6] Kei Terayama, Masato Sumita, Ryo Tamura, Daniel T Payne, Mandeep K Chahal, Shinsuke Ishihara, and Koji Tsuda. Pushing property limits in materials discovery via boundless objective-free exploration. *Chemical Science*, 2020.
- [7] Kobi Felton, Daniel Wigh, and Alexei Lapkin. Multi-task Bayesian optimization of chemical reactions. *ChemRxiv*, 2020.
- [8] Kobi Felton, Jan Rittig, and Alexei Lapkin. Summit: Benchmarking machine learning methods for reaction optimisation. *ChemRxiv*, 2020.
- [9] Chonghuan Zhang, Yehia Amar, Liwei Cao, and Alexei A. Lapkin. Solvent selection for Mitsunobu reaction driven by an active learning surrogate model. *Organic Process Research & Development*, 24(12):2864–2873, 2020. doi: 10.1021/acs.oprd.0c00376.
- [10] Roberto Calandra, André Seyfarth, Jan Peters, and Marc Peter Deisenroth. Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1-2):5–23, 2016.
- [11] James Grant, Alexis Boukouvalas, Ryan-Rhys Griffiths, David Leslie, Sattar Vakili, and Enrique Munoz De Cote. Adaptive sensor placement for continuous spaces. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2385–2393, 2019.
- [12] Simon Olofsson, Mohammad Mehrian, Roberto Calandra, Liesbet Geris, Marc Peter Deisenroth, and Ruth Misener. Bayesian multiobjective optimisation with mixed analytical and black-box functions: Application to tissue engineering. *IEEE Transactions on Biomedical Engineering*, 66(3):727–739, 2018.

REFERENCES

25

- [13] Henry Moss, David Leslie, Daniel Beck, Javier Gonzalez, and Paul Rayson. Boss: Bayesian optimization over string spaces. *Advances in Neural Information Processing Systems*, 33, 2020.
- [14] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584, 2017.
- [15] Ryan-Rhys Griffiths, Philippe Schwaller, and Alpha A. Lee. Dataset bias in the natural sciences: A case study in chemical reaction prediction and synthesis design. *ChemRxiv*, 2018.
- [16] Guilherme Duarte Ramos Matos, Daisy Y Kyu, Hannes H Loeffler, John D Chodera, Michael R Shirts, and David L Mobley. Approaches for calculating solvation free energies and enthalpies demonstrated with an update of the freesolv database. *Journal of Chemical & Engineering Data*, 62(5):1559–1569, 2017.
- [17] Enqing Hou, Xiang Tan, Marijke Heenan, and Dazhi Wen. A global dataset of plant available and unavailable phosphorus in natural soils derived by Hedley method. *Scientific Data*, 5(1):180166, 2018.
- [18] Edward O Pyzer-Knapp, Changwon Suh, Rafael Gómez-Bombarelli, Jorge Aguilera-Iparraguirre, and Alán Aspuru-Guzik. What is high-throughput virtual screening? a perspective from organic materials discovery. *Annual Review of Materials Research*, 45:195–216, 2015.
- [19] José Miguel Hernández-Lobato, James Requeima, Edward O Pyzer-Knapp, and Alán Aspuru-Guzik. Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. In *International Conference on Machine Learning*, pages 1470–1479, 2017.
- [20] Roberto Calandra. *Bayesian Modeling for Optimization and Control in Robotics*. PhD thesis, Technische Universität Darmstadt, 2017.
- [21] Miguel Lázaro-Gredilla and Michalis K Titsias. Variational heteroscedastic Gaussian process regression. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 841–848. Omnipress, 2011.
- [22] Scott R Kuindersma, Roderic A Grupen, and Andrew G Barto. Variable risk control via stochastic optimization. *The International Journal of Robotics Research*, 32(7): 806–825, 2013.
- [23] John-Alexander M Assael, Ziyu Wang, Bobak Shahriari, and Nando de Freitas. Heteroscedastic treed Bayesian optimisation. *arXiv preprint arXiv:1410.7172*, 2014.
- [24] Ryo Ariizumi, Matthew Tesch, Howie Choset, and Fumitoshi Matsuno. Expensive multiobjective optimization for robotics with consideration of heteroscedastic noise. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2230–2235, 2014.
- [25] Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. Safe exploration

REFERENCES

26

- for optimization with Gaussian processes. In *International Conference on Machine Learning*, pages 997–1005, 2015.
- [26] Felix Berkenkamp, Andreas Krause, and Angela P Schoellig. Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *arXiv preprint arXiv:1602.04450*, 2016.
- [27] Peter Frazier, Warren Powell, and Savas Dayanik. The knowledge-gradient policy for correlated normal beliefs. *INFORMS Journal on Computing*, 21(4):599–613, 2009.
- [28] Benjamin Letham, Brian Karrer, Guilherme Ottoni, Eytan Bakshy, et al. Constrained Bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2):495–519, 2019.
- [29] Ryan-Rhys Griffiths, Aldrick Alexander A Garcia-Ortegon, Miguel, and Alpha A Lee. Heteroscedastic Bayesian optimisation in scientific discovery. *2019 NeurIPS Workshop on Machine Learning and the Physical Sciences*, .
- [30] Ryan-Rhys Griffiths, Miguel Garcia-Ortegon, Alexander A Aldrick, and Alpha A Lee. Achieving robustness to aleatoric uncertainty with heteroscedastic Bayesian optimisation. *2019 NeurIPS Workshop on Safety and Robustness in Decision-Making*, .
- [31] H. J. Kushner. A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *Journal of Basic Engineering*, 86(1):97+, 1964.
- [32] Tiesis V. Mockus J. and Žilinskas. The application of Bayesian methods for seeking the extremum. In I. Dixon and G. Szego, editors, *Toward Global Optimization*. Elsevier, Amsterdam, 1978.
- [33] Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- [34] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [35] 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*, pages 393–400, 2007.
- [36] Victor Picheny, Tobias Wagner, and David Ginsbourger. A benchmark of kriging-based infill criteria for noisy optimization. *Structural and Multidisciplinary Optimization*, 48(3):607–626, 2013.
- [37] Emmanuel Vazquez, Julien Villemonteix, Maryan Sidorkiewicz, and Eric Walter. Global optimization based on noisy evaluations: An empirical study of two statistical approaches. In *Journal of Physics: Conference Series*, volume 135, page 012100. IOP Publishing, 2008.

REFERENCES

27

- [38] Deng Huang, Theodore T Allen, William I Notz, and Ning Zeng. Global optimization of stochastic black-box systems via sequential kriging meta-models. *Journal of Global Optimization*, 34(3):441–466, 2006.
- [39] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [40] Ciyu Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)*, 23(4):550–560, 1997.
- [41] David L Mobley and J Peter Guthrie. Freesolv: a database of experimental and calculated hydration free energies, with input files. *Journal of Computer-Aided Molecular Design*, 28(7):711–720, 2014.
- [42] Greg Landrum. Rdkit: Open-source cheminformatics. URL <http://www.rdkit.org>.
- [43] Quoc V Le, Alex J Smola, and Stéphane Canu. Heteroscedastic Gaussian process regression. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 489–496, 2005.
- [44] Mickael Binois, Robert B Gramacy, and Mike Ludkovski. Practical heteroscedastic Gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics*, 27(4):808–821, 2018.
- [45] Ibrahim Almosallam. *Heteroscedastic Gaussian processes for uncertain and incomplete data*. PhD thesis, University of Oxford, 2017.
- [46] Luis Muñoz-González, Miguel Lázaro-Gredilla, and Aníbal R Figueiras-Vidal. Heteroscedastic Gaussian process regression using expectation propagation. In *2011 IEEE International Workshop on Machine Learning for Signal Processing*, pages 1–6. IEEE, 2011.
- [47] Zilong Wang and Marianthi Ierapetritou. A novel surrogate-based optimization method for black-box simulation with heteroscedastic noise. *Industrial & Engineering Chemistry Research*, 56(38):10720–10732, 2017.
- [48] Chunyi Wang and Radford M Neal. Gaussian process regression with heteroscedastic or non-Gaussian residuals. *arXiv preprint arXiv:1212.6246*, 2012.
- [49] Qin-Hu Zhang and Yi-Qing Ni. Improved most likely heteroscedastic Gaussian process regression via Bayesian residual moment estimator. *IEEE Transactions on Signal Processing*, 68:3450–3460, 2020.

REFERENCES

28

- [50] Filipe Rodrigues and Francisco C Pereira. Heteroscedastic Gaussian processes for uncertainty modeling in large-scale crowdsourced traffic data. *Transportation Research Part C: Emerging Technologies*, 95:636–651, 2018.
- [51] Lucie Tabor, James-A Goulet, Jean-Philippe Charron, and Clelia Desmettre. Probabilistic modeling of heteroscedastic laboratory experiments using Gaussian process regression. *Journal of Engineering Mechanics*, 144(6):04018038, 2018.
- [52] TJ Rogers, P Gardner, N Dervilis, K Worden, AE Maguire, E Papatheou, and EJ Cross. Probabilistic modelling of wind turbine power curves with application of heteroscedastic Gaussian process regression. *Renewable Energy*, 148:1124–1136, 2020.
- [53] Qi-Ang Wang and Yi-Qing Ni. Measurement and forecasting of high-speed rail track slab deformation under uncertain shm data using variational heteroscedastic Gaussian process. *Sensors*, 19(15):3311, 2019.
- [54] Wenjing Wang and Xi Chen. Distributed variational inference-based heteroscedastic Gaussian process metamodeling. In *2019 Winter Simulation Conference (WSC)*, pages 380–391. IEEE, 2019.
- [55] Haitao Liu, Yew-Soon Ong, and Jianfei Cai. Large-scale heteroscedastic regression via Gaussian process. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [56] James T. Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *International Conference on Machine Learning*, 2020.
- [57] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- [58] Vidhi Lalchand and Carl Edward Rasmussen. Approximate inference for fully Bayesian Gaussian process regression. *arXiv preprint arXiv:1912.13440*, 2019.
- [59] Riccardo Moriconi, Marc Peter Deisenroth, and KS Sesh Kumar. High-dimensional Bayesian optimization using low-dimensional feature spaces. *Machine Learning*, 109(9):1925–1943, 2020.
- [60] Antonio Candelieri and Riccardo Perego. Dimensionality reduction methods to scale Bayesian optimization up. *Numerical Computations: Theory and Algorithms NUMTA 2019*, page 167, 2019.
- [61] Antoine Grosnit, Rasul Tutunov, Alexandre Max Maraval, Ryan-Rhys Griffiths, Alexander I Cowen-Rivers, Lin Yang, Lin Zhu, Wenlong Lyu, Zhitang Chen, Jun Wang, et al. High-dimensional Bayesian optimisation with variational autoencoders and deep metric learning. *arXiv preprint arXiv:2106.03609*, 2021.
- [62] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 33, pages 21524–21538, 2020.

REFERENCES

29

- [63] Kirthivasan Kandasamy, Karun Raju Vysyaraju, Willie Neiswanger, Biswajit Paria, Christopher R. Collins, Jeff Schneider, Barnabas Poczos, and Eric P. Xing. Tuning hyperparameters without grad students: Scalable and robust Bayesian optimisation with dragonfly. *Journal of Machine Learning Research*, 21(81):1–27, 2020.
- [64] James Wilson, Frank Hutter, and Marc Deisenroth. Maximizing acquisition functions for Bayesian optimization. *Advances in Neural Information Processing Systems*, 31: 9884–9895, 2018.
- [65] Antoine Grosnit, Alexander I Cowen-Rivers, Rasul Tutunov, Ryan-Rhys Griffiths, Jun Wang, and Haitham Bou-Ammar. Are we forgetting about compositional optimisers in Bayesian optimisation? *arXiv preprint arXiv:2012.08240*, 2020.
- [66] Rasul Tutunov, Minne Li, Alexander I Cowen-Rivers, Jun Wang, and Haitham Bou-Ammar. Compositional adam: An adaptive compositional solver. *arXiv preprint arXiv:2002.03755*, 2020.
- [67] Artur M Schweidtmann, Dominik Bongartz, Daniel Grothe, Tim Kerkenhoff, Xiaopeng Lin, Jaromil Najman, and Alexander Mitsos. Global optimization of Gaussian processes. *arXiv preprint arXiv:2005.10902*, 2020.
- [68] Johannes Wiebe, Inês Cecílio, Jonathan Dunlop, and Ruth Misener. A robust approach to warped Gaussian process-constrained optimization. *arXiv preprint arXiv:2006.08222*, 2020.
- [69] Alexander I Cowen-Rivers, Wenlong Lyu, Rasul Tutunov, Zhi Wang, Antoine Grosnit, Ryan Rhys Griffiths, Hao Jianye, Jun Wang, and Haitham Bou Ammar. An empirical study of assumptions in Bayesian optimisation. *arXiv preprint arXiv:2012.03826*, 2020.
- [70] Henry B Moss and Ryan-Rhys Griffiths. Gaussian process molecule property prediction with FlowMO. *arXiv preprint arXiv:2010.01118*, 2020.
- [71] Aditya R Thawani, Ryan-Rhys Griffiths, Arian Jamasb, Anthony Bourached, Penelope Jones, William McCorkindale, Alexander A Aldrick, and Alpha A Lee. The photoswitch dataset: A molecular machine learning benchmark for the advancement of synthetic chemistry. *arXiv preprint arXiv:2008.03226*, 2020.
- [72] Bingqing Cheng, Ryan-Rhys Griffiths, Simon Wengert, Christian Kunkel, Tamas Stenczel, Bonan Zhu, Volker L Deringer, Noam Bernstein, Johannes T Margraf, Karsten Reuter, et al. Mapping materials and molecules. *Accounts of Chemical Research*, 53(9):1981–1991, 2020.
- [73] Nassim Nicholas Taleb. *Antifragile: Things That Gain from Disorder*. Random House, New York, 1st ed edition, 2012. ISBN 978-1-4000-6782-4.
- [74] Yuanyuan Zhou and Yixin Zhao. Chemical stability and instability of inorganic halide perovskites. *Energy & Environmental Science*, 12(5):1495–1511, May 2019. ISSN 1754-5706. doi: 10.1039/C8EE03559H.

Appendix A. Heteroscedasticity of the Soil Phosphorus Fraction Dataset

[Table A1](#) is used to demonstrate the efficacy of modelling the soil phosphorus fraction dataset using a heteroscedastic GP. The heteroscedastic GP outperforms the homoscedastic GP on prediction based on the metric of negative log predictive density (NLPD)

$$\text{NLPD} = \frac{1}{n} \sum_{i=1}^n -\log p(t_i | \mathbf{x}_i)$$

which penalises both over and under-confident predictions.

Table A1. Comparison of NLPD values on the soil phosphorus fraction dataset. Standard errors are reported for 10 independent train/test splits. Lower scores are better.

Soil Phosphorus Fraction Dataset	GP	Het GP
NLPD	1.35 ± 1.33	1.00 ± 0.95

Appendix B. Additional Ablation Experiments

In this section we present the ablation results on noiseless, homoscedastic and heteroscedastic noise tasks in line with [section 6](#) of the main paper.

Appendix B.1. Goldstein-Price Function

The form of the Goldstein-Price function is given in [Equation 9](#) with noise function in [Equation 11](#). The function is visualised in [Figure B1](#). 9 data points are used for initialisation in the noiseless and homoscedastic noise cases whereas 100 data points are used for initialisation in the heteroscedastic noise case. β is set to 0.5 for the noiseless and homoscedastic noise tasks and $\frac{1}{11}$ for the heteroscedastic noise task. γ is set to 500 for all experiments.

Appendix B.1.1. Noiseless Case The results of the noiseless case for Goldstein-Price are given in [Figure B2](#). All Bayesian optimisation methods outperform random search with ANPEI best and HAEI second best.

Appendix B.1.2. Homoscedastic Noise Case The results of the homoscedastic noise case for Goldstein-Price are shown in [Figure B3](#). In this instance HAEI performs best.

Appendix B.1.3. Heteroscedastic Noise The results of the heteroscedastic noise case for Goldstein-Price are shown in [Figure B4](#). ANPEI performs best whilst HAEI performs worse than random search.




Figure B1. (a) The latent Goldstein-Price Function $f(\mathbf{x})$ together with (b) its heteroscedastic noise function $g(\mathbf{x})$ and (c) the objective function $f(\mathbf{x}) + g(\mathbf{x})$.




Figure B2. Goldstein-Price function noiseless case. All Bayesian optimisation methods outperform random search. ANPEI performs best and HAEI is runner-up.

Appendix B.2. Branin-Hoo Function

The form of the Branin-Hoo function is given in Equation 6 with noise function in Equation 7. The function is visualised in Figure B5, a figure from the main paper repeated here for clarity. 9 data points are used for initialisation in the noiseless and homoscedastic noise cases whereas 100 data points are used for initialisation in the heteroscedastic noise case. β is set to 0.5 and γ is set to 500 for all experiments.

Appendix B.2.1. Noiseless Case The results of the noiseless case for the Branin-Hoo





Figure B3. Goldstein-Price function homoscedastic noise case. HAEI performs best.



(a) Best Objective Value Found so Far

(b) Lowest Aleatoric Noise Found so Far

Figure B4. Comparison of heteroscedastic and homoscedastic Bayesian optimisation on the heteroscedastic 2D Goldstein-Price function. (a) shows the optimisation of $h(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x})$ (lower is better) where $g(\mathbf{x})$ is the aleatoric noise. (b) shows the values $g(\mathbf{x})$ obtained over the course of the optimisation of $h(\mathbf{x})$.

function are given in Figure B6. HAEI performs best in this case whereas ANPEI




Figure B5. Heteroscedastic Branin Function.




Figure B6. Branin-Hoo function noiseless case. HAEI performs best. ANPEI performs worst.

performs worst.

Appendix B.2.2. Homoscedastic Noise Case The results of the homoscedastic noise case for the Branin-Hoo function are given in Figure B7. All Bayesian optimisation methods outperform random search yet perform comparably against each other.

Appendix B.2.3. Heteroscedastic Noise The results of the heteroscedastic noise case for the Branin-Hoo function are shown in Figure B8. ANPEI performs best whilst HAEI performs worse than random search.




Figure B7. Branin-Hoo function homoscedastic noise case. All Bayesian optimisation methods outperform random search.




Figure B8. Comparison of heteroscedastic and homoscedastic Bayesian optimisation on the heteroscedastic 2D Branin function. (a) shows the optimisation of $h(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x})$ (lower is better) where $g(\mathbf{x})$ is the aleatoric noise. (b) shows the values $g(\mathbf{x})$ obtained over the course of the optimisation of $h(\mathbf{x})$.

Appendix C. Performance Impact of the Kernel Choice

In this section we analyse the impact that the choice of GP kernel has on Bayesian optimisation performance. We select three kernels for this purpose: the squared exponential kernel

$$k_{\text{SQE}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right)$$

used for all experiments in the main paper, the exponential kernel

$$k_{\text{exp}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|}{\ell}\right),$$

a special instance of the Matérn kernel for values of $\nu = \frac{1}{2}$ [34] as well as the Matérn 5/2 kernel

$$k_{\text{Matérn}(5/2)}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \left(1 + \frac{\sqrt{5}\|\mathbf{x} - \mathbf{x}'\|}{\ell} + \frac{5\|\mathbf{x} - \mathbf{x}'\|^2}{3\ell^2}\right) \cdot \exp\left(\frac{-\sqrt{5}\|\mathbf{x} - \mathbf{x}'\|}{\ell}\right)$$

which is one of the most popular kernels for large scale empirical studies [64, 65]. It should be noted that while the equations are written assuming a single scalar lengthscale, in practice for the experiments in greater than 1D, each lengthscale is optimised per dimension under the marginal likelihood. For all experiments we choose the same kernel for both GPs of the MLHGP model i.e. the GP modelling the objective as well as the GP modelling the noise. 100 points are used for initialisation in the Branin-Hoo and Goldstein-Price functions and 144 points are used for the Hosaki function. β is set to 0.5 for the Branin-Hoo and Hosaki functions and $\frac{1}{11}$ for the Goldstein-Price function. γ is set to 500 for all experiments. The results are shown in Figure C1, Figure C2 and Figure C3 for the Branin-Hoo function, Goldstein-Price function and Hosaki functions respectively. There is no significant difference in performance using each kernel save for the Branin-Hoo function where ANPEI underperforms using the somewhat rougher exponential kernel.

Appendix

36




Figure C1. Branin-Hoo function kernel comparison.




Figure C2. Goldstein-Price function kernel comparison.




Figure C3. Hosaki Function kernel comparison.