



LAWRENCE  
LIVERMORE  
NATIONAL  
LABORATORY

LLNL-CONF-867786

# Robust Multi-fidelity Bayesian Optimization with Deep Kernel and Partition

F. Zhang, T. A. Desautels, Y. Chen

August 2, 2024

The 39th Annual AAAI Conference on Artificial Intelligence  
Philadelphia, PA, United States  
February 25, 2025 through March 4, 2025

## **Disclaimer**

---

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

---

# Robust Multi-fidelity Bayesian Optimization with Deep Kernel and Partition

---

**Fengxue Zhang**

University of Chicago

zhangfx@uchicago.edu

**Thomas A. Desautels**

Lawrence Livermore

National Laboratory

desautels2@llnl.gov

**Yuxin Chen**

University of Chicago

chenyuxin@uchicago.edu

## Abstract

Multi-fidelity Bayesian optimization (MFBO) is a powerful approach that utilizes low-fidelity, cost-effective sources to expedite the exploration and exploitation of a high-fidelity objective function. Existing MFBO methods with theoretical foundations either lack justification for performance improvements over single-fidelity optimization or rely on strong assumptions about the relationships between fidelity sources to construct surrogate models and direct queries to low-fidelity sources. To mitigate the dependency on cross-fidelity assumptions while maintaining the advantages of low-fidelity queries, we introduce a random sampling and partition-based MFBO framework with deep kernel learning. This framework is robust to cross-fidelity model misspecification and explicitly illustrates the benefits of low-fidelity queries. Our results demonstrate that the proposed algorithm effectively manages complex cross-fidelity relationships and efficiently optimizes the target fidelity function.

## 1 Introduction

Multi-fidelity Bayesian optimization (MFBO) (Dai et al., 2019; Wu et al., 2020; Takeno et al., 2020) is increasingly prevalent in the adaptive design of scientific experiments (Buterez et al., 2023), automated hyperparameter optimization (Eggensperger et al., 2021; Pfisterer et al., 2022), and policy optimization in control problems (Letham and Bakshy, 2019; Wu et al., 2020).

Previous work has often relied on various assumptions about the relationship between different fidelities to analyze efficiency theoretically (Song et al., 2019; Kandasamy et al., 2016, 2017). Similar to transfer learning for Bayesian optimization (BO), a more practical challenge involves handling significant misalignment between fidelities while maintaining cost efficiency. Recent approaches to robust transfer learning for BO (Appice et al., 2015; Probst et al., 2019; Perrone et al., 2019; Reif et al., 2012; Pfisterer et al., 2021; Feurer et al., 2018) and robust single-fidelity BO against model misspecification (Bogunovic and Krause, 2021; Liu et al., 2023) have addressed this issue yet typically do not consider the sample efficiency on the lower fidelities. Some research suggests mitigating the problem by avoiding evaluations or learning from unreliable low-fidelity sources (Mikkola et al., 2023; Foumani et al., 2023), but they do not explicitly deal with errors incurred from the unreliable model learning of the multi-fidelity structure in the model design and acquisition.

Leveraging recent advancements in efficient kernel learning, uncertainty quantification, and error bounds for learning algorithms (Xu and Raginsky, 2017; Robinson et al., 2020; Wang et al., 2021), we propose a general-purpose framework that uses sampling-based cost-aware acquisition. This framework captures complex and potentially misaligned multi-fidelity evaluations while explicitly addressing model misspecification on the fly with robust data acquisition and deep kernel learning.

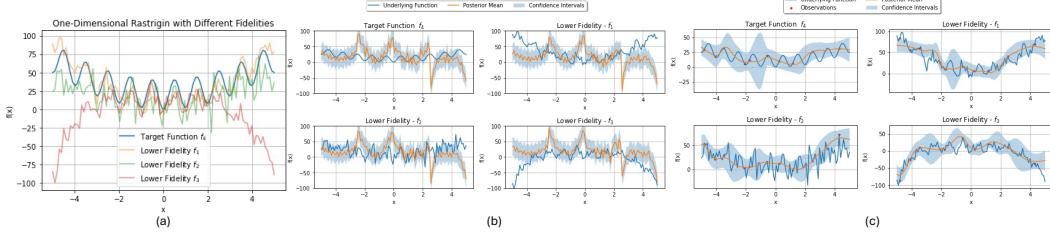


Figure 1: We illustrate the problem of learning and optimization on both target fidelity and misaligned low fidelities. (a) demonstrates the one-dimensional Rastrigin function (Pohlheim, 2007) and the manually constructed low fidelities. (b) demonstrates the posterior of learning the multi-fidelity functions with conventional multi-task GP (Swersky et al., 2013) previously applied in MFBO (Letham and Bakshy, 2019) when feeding 2000 training points densely distributed in the search space. (c) shows the posterior of the proposed model when feeding 10 points from each fidelity.

## 2 Preliminaries

We begin by introducing useful notation, mostly following previous work by Song et al. (2019), and formally stating the problem studied in this paper.

### 2.1 Multi-fidelity Optimization of Unknown Objective

Consider the problem of maximizing an unknown payoff function  $f_M : \mathbf{X} \rightarrow \mathbb{R}$ . We can probe this function by directly querying it at some point  $\mathbf{x} \in \mathbf{X}$ , consequently obtaining a noise-free observation  $y_{\langle \mathbf{x}, M \rangle} = f_M(\mathbf{x})$ . In addition to  $f_M$ , we have access to oracle calls for unknown auxiliary functions  $f_1, \dots, f_{M-1} : \mathbf{X} \rightarrow \mathbb{R}$ . Similarly, querying any  $f_\ell$  at  $\mathbf{x}$  yields a noise-free observation  $y_t = f_{\ell_t}(\mathbf{x}_t)$ . Each auxiliary function  $f_\ell$  can be viewed as a lower-fidelity version of  $f_M$  when  $\ell < M$ . Specifically, we model the unknown target fidelity functions with corresponding Gaussian process (GP):  $f_M \sim \text{GP}(\mu_M(\mathbf{x}), k_M(\mathbf{x}, \mathbf{x}'))$ , where  $\mu_M$  and  $k_M$  denote the prior mean and covariance. Let  $\langle \mathbf{x}, \ell \rangle$  denote the action of querying  $f_\ell$  at  $\mathbf{x}$ . Each action  $\langle \mathbf{x}, \ell \rangle$  incurs a cost of  $\lambda_\ell$  and yields a reward:

$$r(\langle \mathbf{x}, \ell \rangle) = \begin{cases} f_M(\mathbf{x}) & \text{if } \ell = M \\ r_{\min} & \text{otherwise} \end{cases}$$

That is, performing  $\langle \mathbf{x}, M \rangle$  at the target fidelity level achieves a reward  $f_M(\mathbf{x})$ . The collective historical observations after  $T$  iteration is denoted by  $\mathcal{D}_T \triangleq \{\langle \mathbf{x}_t, \ell_t \rangle, y_t\}_{t=1 \dots T}$ . We also define the collective historical observations up to certain fidelity  $\ell$  as  $\mathcal{D}_{\ell, T} \triangleq \{\langle \mathbf{x}_t, \ell_t \rangle, y_t\}_{1 \leq t \leq T, \ell_t \leq \ell}$ . When given a fixed budget, we need to guarantee the cumulative cost does not exceed the budget, i.e.,  $\sum_{t=1}^T \lambda_{\ell_t} \leq \Lambda$ . Lower fidelity actions  $\langle \mathbf{x}, \ell \rangle$  for  $\ell < M$  yield the minimal immediate reward  $r_{\min}$  but can provide valuable information about  $f_M$ , potentially leading to better decisions later. Without loss of generality, we assume  $\max_{\mathbf{x}} f_M(\mathbf{x}) \geq 0$  and  $r_{\min} \equiv 0$ . Note that we define the reward only as incurred based on the target fidelity, and the query on low fidelity does not incur a reward but only helps with the learning. Hence, in the context of multi-fidelity Bayesian optimization, the **simple regret (SR)** is defined as:  $\mathbf{R}(\hat{\mathbf{x}}) = f_M(\mathbf{x}^*) - f_M(\hat{\mathbf{x}})$ , where  $\hat{\mathbf{x}} := \arg \max_{\mathbf{x}: \langle \mathbf{x}, M \rangle, y \in \mathcal{D}_T} f_M(\mathbf{x})$  is a point selected to be evaluated at the target fidelity, and  $\mathbf{x}^*$  is the global maximizer of the function  $f_M$ . Our objective is to find the candidate that minimizes the simple regret after **exhausting a given budget**  $\Lambda$ .

### 2.2 Expected Excess Risk

In addition to the conventional analysis that assumes the prior is properly specified, we explicitly deal with the model misspecification regarding the difference between the posterior mean and true underlying function. In the context of statistical learning theory, the **convergence rate** of the expected excessive risk with respect to the training dataset  $\mathcal{D}_{\ell, T}$  at fidelity  $\ell \in [M]^+ = [1 \dots M]$  after  $T$  iterations is defined as  $\text{Rate}_{MF}(\ell, T) \triangleq \mathbb{E} \left[ \mathcal{L}(f_\ell, \tilde{f}_\ell) | \mathcal{D}_{\ell, T} \right] - \mathcal{L}(f_\ell, \tilde{f}_\ell^*) = O(T^\alpha)$ , Here,  $-1 < \alpha < 0$  is a constant that characterizes the rate of convergence.  $\tilde{f}_\ell$  is the hypothesis of fidelity  $\ell$  produced by the learning algorithm when trained on  $\mathcal{D}_{\ell, T}$ .  $\mathcal{L}(f, \tilde{f})$  is the loss function evaluating the hypothesis  $\tilde{f}$  with respect to the true objective  $f$ , and  $\tilde{f}_\ell^*$  is the hypothesis on fidelity  $\ell$  that minimizes the expected loss.

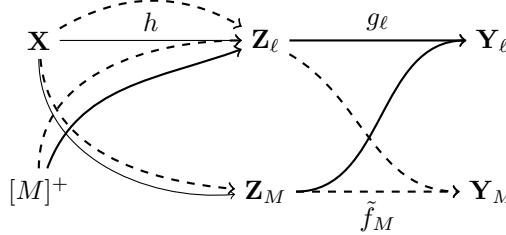


Figure 2: Schema for multi-fidelity learning implemented with deep kernel. The dotted lines denote the flow of target fidelity (strong data), and the solid lines the flow of low fidelities(weak data). Here, we denote the output space for certain fidelity  $\ell \in [M]^+$  as  $\mathbf{Y}_\ell$ . Specifically, the target fidelity output space is denoted as  $\mathbf{Y}_M$ .

### 3 Method

In this section, we discuss the model design and the analysis-inspired data acquisition procedure of the proposed Robust Multi-Fidelity Bayesian Optimization with Deep Kernel Learning and Partition (RMFBO-DP). The pseudo code and detailed design choices are deferred to Appendix B.

#### 3.1 Model

We employ hierarchical deep kernel learning (Wilson et al., 2016) regularized with the spectral norm (Van Amersfoort et al., 2021) with mean absolute error (MAE) as the loss function to deal with overfitting caused by shortage of training data and the specific choice of training loss. We assume that the multiple fidelities  $\{f_\ell\}_{\ell \in [M]}$  are mutually dependent through some fixed, possibly unknown joint probability distribution. Therefore, we seek to approximate the joint underlying function  $f : \mathbf{X} \times [M]^+ \rightarrow \mathbb{R}$ , where both the position  $\mathbf{x} \in \mathbf{X}$  and fidelity  $\ell \in [M]^+$  are inputs, and learn the approximation  $\tilde{f}$  through joint learning of  $h$  and  $g_{\ell \in [M]^+}$ . Here, a *single* latent space mapping  $h : \mathbf{X} \times [M]^+ \rightarrow \mathbf{Z}$  convert the input space  $\mathbf{X}$  to the latent space  $\mathbf{Z}$  which consists of the fidelity independent part  $\mathbf{Z}_M$  and fidelity dependent part  $\mathbf{Z}_\ell$ . On top of the latent space, we construct a set of objective mappings  $g_1 \dots g_M$  for each fidelity. Namely  $\forall \ell \in [M]^+, f_\ell \triangleq f(\cdot, \ell)$  and  $f(\cdot, \ell)$  is approximated by  $\tilde{f}_\ell \triangleq g_\ell(h(\cdot, \ell))$ . We illustrate the model structure in figure 2. The model is trained by first traversing all low fidelities for weak learning, then jointly optimizing the deep kernel.

#### 3.2 Data Acquisition

We extend random exploration to the multi-fidelity regime. We use both the expected generalization error and SR bounds to guide the cost-efficient acquisition. When the error bound contributes more to global regret, we randomly explore until it is more cost-efficient to conduct target fidelity acquisition. When the SR contributes more to the general regret, we conduct target fidelity acquisition on certain partitions of the search space. We leverage the GP posterior calibrated with the excess risk on observed points to exclude from acquisition the partitions of the search space that, with high probability, do not contain the global optimum. To do so, we rely on both the upper confidence bound  $UCB_{f_M, t}(\mathbf{x}) \triangleq \mu_{f_M, t-1}(\mathbf{x}) + \beta_{f_M, t}^{1/2} \sigma_{f_M, t-1}(\mathbf{x})$  and lower confidence bound  $LCB_{f_M, t}(\mathbf{x}) \triangleq \mu_{f_M, t-1}(\mathbf{x}) - \beta_t^{1/2} \sigma_{f_M, t-1}(\mathbf{x})$ , where  $\sigma_{t-1}(\mathbf{x}) = k_{t-1}(\mathbf{x}, \mathbf{x})^{1/2}$  and  $\beta_t$  acts as a scaling factor corresponding to certain confidence. Formally, the acting search space at iteration  $t$  is  $\hat{\mathbf{X}}_t = \{\mathbf{x} \in \mathbf{X} \mid UCB_{f_M, t}(\mathbf{x}) > \max_{\mathbf{x}' \in \mathbf{X}} LCB_{f_M, t}(\mathbf{x}') - \text{Rate}_{MF}(M, t)\}$ .

Regarding the theoretical justification, previous works offer an upper bound for cumulative regret when applying random exploration on the partitions of interest (Salgia et al., 2024) defined as  $\{x \in \mathbf{X} \mid UCB_{f_M, t}(x) > \max_{x' \in \mathbf{X}} LCB_{f_M, t}(x')\}$ , as the target fidelity acquisition function when we ignore the contribution of low-fidelity evaluation to the target fidelity learning. We extend the cumulative regret bound into the following form of SR bound.

**Informal Theorem 1** *Under proper assumptions and choices of parameters, when ignoring the contribution of performance by learning on the low fidelities, we have with probability at least*

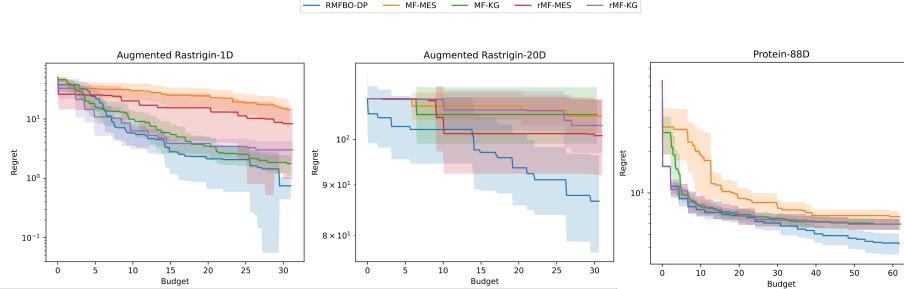


Figure 3: We illustrate the performance of RMFBO-DP compared against MF-MES, MF-KG, rMF-MES, and rMF-KG on both synthetic and real-world datasets. The results are collected from 10 independent trials. The y-axis denotes the simple regret, and the x-axis denotes the consumed budget. The shades area shows the 95% confidence interval. We trim the shared initial single data point. Detailed discussions are deferred to Appendix C.

$1 - \delta, SR(t) = \tilde{\mathcal{O}} \left( \sqrt{\frac{\gamma_{T_M(t)}}{T_M(t)}} \log \frac{T_M(t)}{\delta} \right)$ . Here  $\tilde{\mathcal{O}}$  means up to the logarithmic factor, and  $T_\ell(t) \triangleq |\{\langle \mathbf{x}_{t'}, \ell_{t'} \rangle, y_{t'}\}_{1 \leq t' \leq t, \ell_{t'} = \ell}|$  denotes evaluations at fidelity  $\ell$  among  $t$  evaluations.

We exploit recent advancements in expected excess risk in meta-learning (Robinson et al., 2020) to extend the previous SR results to multiple weak learning sources. We state the informal version of the theoretical results here while deferring the discussion of assumptions, proof, and other details to Appendix A. First, we decompose the ultimate SR into separate components for the conventional regret and the generalization error.

**Informal Theorem 2** *The misspecification-aware simple regret ( $SR_{MA}$ ) of the proposed algorithm can be decomposed into the standard simple regret (SR) and the rate term as  $SR_{MA}(t) \leq SR(t) + Rate_{MF}(M, t)$ .*

In the following, we generalize the meta-learning expected excess risk (Robinson et al., 2020; Xu and Raginsky, 2017) to multi-fidelity learning.

**Informal Theorem 3** *When the lowest single fidelity bears the convergence rate  $Rate_{MF}(1, T) = O(\sqrt{\frac{1}{T_1}})$ , the excess risk bound  $Rate_{MF}(\ell, t) \leq$*

$$O \left( Rate_{MF}(\ell - 1, t) + \frac{\sqrt{\left( \log_{T_\ell(t)} Rate_{MF}(\ell - 1, t) + 1 \right) \log T_\ell(t)}}{T_\ell(t)} \right)$$

This allows us to differentiate the multiple fidelities' contribution to the target fidelity learning and regret minimization. Specifically, a cost-aware multi-fidelity acquisition could be made by minimizing the  $SR_{MA}(T)$  such that the total cost incurred by querying different fidelities does not exceed  $\Lambda$ .

### 3.3 Evaluation

We evaluated the proposed algorithm RMFBO-DP against four baselines on both synthetic datasets corresponding to figure 1 and real-world protein design dataset. The results shown in figure 3 demonstrate the robustness of RMFBO-DP in various tasks and efficiency in tasks of different dimensionalities. We defer detailed description to Appendix C.

## 4 Conclusion

In this paper, we introduced a novel multi-fidelity Bayesian optimization approach focusing on robustness and efficiency. Our method explicitly addresses the misspecification issues in multi-fidelity deep model learning by incorporating budget-sensitive low-fidelity sampling and constraining acquisitions to a subset of the global search space for target fidelity optimization. By tackling the challenges of low-fidelity misalignment and efficient target fidelity optimization in a principled, cost-effective manner, we demonstrated that our approach significantly improves robustness and performance over existing methods, as confirmed by our theoretical and empirical results.

## Acknowledgements

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-CONF-867786. The GUIDE program is executed by the Joint Program Executive Office for Chemical, Biological, Radiological and Nuclear Defense (JPEO-CBRND) Joint Project Lead for CBRND Enabling Biotechnologies (JPL CBRND EB) on behalf of the Department of Defense's Chemical and Biological Defense Program. This effort was in collaboration with the Defense Health Agency (DHA) COVID funding initiative. The views expressed in this paper reflect the views of the authors and do not necessarily reflect the position of the Department of the Army, Department of Defense, nor the United States Government. References to non-federal entities do not constitute nor imply Department of Defense or Army endorsement of any company or organization. This work was completed in part with resources provided by the University of Chicago's Research Computing Center.

## References

Annalisa Appice, Pedro Pereira Rodrigues, Vítor Santos Costa, Carlos Soares, João Gama, and Alípio Jorge. Machine learning and knowledge discovery in databases. In *Proceedings of the European Conference, ECML PKDD, Part II, Porto, Portugal*.–2015. Springer, 2015. [1](#)

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. *Advances in neural information processing systems*, 33:21524–21538, 2020. [9](#)

Kyle A Barlow, Shane O Conchuir, Samuel Thompson, Pooja Suresh, James E Lucas, Markus Heinonen, and Tanja Kortemme. Flex ddg: Rosetta ensemble-based estimation of changes in protein–protein binding affinity upon mutation. *The Journal of Physical Chemistry B*, 122(21): 5389–5399, 2018. [10](#)

Ilija Bogunovic and Andreas Krause. Misspecified gaussian process bandit optimization. *Advances in neural information processing systems*, 34:3004–3015, 2021. [1](#)

Oliver Buß, Jens Rudat, and Katrin Ochsenreither. Foldx as protein engineering tool: better than random based approaches? *Computational and structural biotechnology journal*, 16:25–33, 2018. [10](#)

David Buterez, Jon Paul Janet, Steven J Kiddle, and Pietro Liò. Mf-pcba: Multifidelity high-throughput screening benchmarks for drug discovery and machine learning. *Journal of Chemical Information and Modeling*, 63(9):2667–2678, 2023. [1](#)

Zhongxiang Dai, Haibin Yu, Bryan Kian Hsiang Low, and Patrick Jaillet. Bayesian optimization meets bayesian optimal stopping. In *International conference on machine learning*, pages 1496–1506. PMLR, 2019. [1](#)

Rhiju Das and David Baker. Macromolecular modeling with rosetta. *Annu. Rev. Biochem.*, 77: 363–382, 2008. [10](#)

Thomas Desautels, Adam Zemla, Edmond Lau, Magdalena Franco, and Daniel Faissol. Rapid in silico design of antibodies targeting sars-cov-2 using machine learning and supercomputing. *BioRxiv*, 2020. [10](#)

Thomas A Desautels, Kathryn T Arrildt, Adam T Zemla, Edmond Y Lau, Fangqiang Zhu, Dante Ricci, Stephanie Cronin, Seth Zost, Elad Binshtein, Suzanne M Scheaffer, et al. Computationally restoring the potency of a clinical antibody against sars-cov-2 omicron subvariants. *bioRxiv*, pages 2022–10, 2022. [10](#)

Katharina Eggensperger, Philipp Müller, Neeratoyoy Mallik, Matthias Feurer, René Sass, Aaron Klein, Noor Awad, Marius Lindauer, and Frank Hutter. Hpobench: A collection of reproducible multi-fidelity benchmark problems for hpo. *arXiv preprint arXiv:2109.06716*, 2021. [1](#)

Matthias Feurer, Benjamin Letham, Frank Hutter, and Eytan Bakshy. Practical transfer learning for bayesian optimization. *arXiv preprint arXiv:1802.02219*, 2018. [1](#)

Zahra Zanjani Foumani, Mehdi Shishehbor, Amin Yousefpour, and Ramin Bostanabad. Multi-fidelity cost-aware bayesian optimization. *Computer Methods in Applied Mechanics and Engineering*, 407:115937, 2023. [1](#)

Jacob Gardner, Geoff Pleiss, Kilian Q Weinberger, David Bindel, and Andrew G Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. *Advances in neural information processing systems*, 31, 2018. [9](#)

Kirthevasan Kandasamy, Gautam Dasarathy, Junier B Oliva, Jeff Schneider, and Barnabás Póczos. Gaussian process bandit optimisation with multi-fidelity evaluations. *Advances in neural information processing systems*, 29, 2016. [1](#)

Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabás Póczos. Multi-fidelity bayesian optimisation with continuous approximations. In *International conference on machine learning*, pages 1799–1808. PMLR, 2017. [1](#)

Benjamin Letham and Eytan Bakshy. Bayesian optimization for policy search via online-offline experimentation. *Journal of Machine Learning Research*, 20(145):1–30, 2019. [1](#), [2](#)

Chong Liu, Ming Yin, and Yu-Xiang Wang. No-regret linear bandits beyond realizability. In *Uncertainty in Artificial Intelligence*, pages 1294–1303. PMLR, 2023. [1](#)

Petrus Mikkola, Julien Martinelli, Louis Filstroff, and Samuel Kaski. Multi-fidelity bayesian optimization with unreliable information sources. In *International Conference on Artificial Intelligence and Statistics*, pages 7425–7454. PMLR, 2023. [1](#), [9](#)

Valerio Perrone, Huibin Shen, Matthias W Seeger, Cedric Archambeau, and Rodolphe Jenatton. Learning search spaces for bayesian optimization: Another view of hyperparameter transfer learning. *Advances in neural information processing systems*, 32, 2019. [1](#)

Florian Pfisterer, Jan N van Rijn, Philipp Probst, Andreas C Müller, and Bernd Bischl. Learning multiple defaults for machine learning algorithms. In *Proceedings of the genetic and evolutionary computation conference companion*, pages 241–242, 2021. [1](#)

Florian Pfisterer, Lennart Schneider, Julia Moosbauer, Martin Binder, and Bernd Bischl. Yahpo gym—an efficient multi-objective multi-fidelity benchmark for hyperparameter optimization. In *International Conference on Automated Machine Learning*, pages 3–1. PMLR, 2022. [1](#)

Hartmut Pohlheim. Examples of objective functions. *Retrieved*, 4(10):2012, 2007. [2](#), [10](#)

Philipp Probst, Anne-Laure Boulesteix, and Bernd Bischl. Tunability: Importance of hyperparameters of machine learning algorithms. *Journal of Machine Learning Research*, 20(53):1–32, 2019. [1](#)

Matthias Reif, Faisal Shafait, and Andreas Dengel. Meta-learning for evolutionary parameter optimization of classifiers. *Machine learning*, 87:357–380, 2012. [1](#)

Joshua Robinson, Stefanie Jegelka, and Suvrit Sra. Strength from weakness: Fast learning using weak supervision. In *International Conference on Machine Learning*, pages 8127–8136. PMLR, 2020. [1](#), [4](#), [7](#), [8](#)

Sudeep Salgia, Sattar Vakili, and Qing Zhao. Random exploration in bayesian optimization: Order-optimal regret and computational efficiency. In *Forty-first International Conference on Machine Learning*, 2024. [3](#), [7](#), [8](#)

Yesol Sapozhnikov, Jagdish Suresh Patel, F Marty Ytreberg, and Craig R Miller. Statistical modeling to quantify the uncertainty of foldx-predicted protein folding and binding stability. *BMC bioinformatics*, 24(1):426, 2023. [10](#)

Joost Schymkowitz, Jesper Borg, Francois Stricher, Robby Nys, Frederic Rousseau, and Luis Serrano. The foldx web server: an online force field. *Nucleic acids research*, 33(suppl\_2):W382–W388, 2005. [10](#)

Jialin Song, Yuxin Chen, and Yisong Yue. A general framework for multi-fidelity bayesian optimization with gaussian processes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3158–3167. PMLR, 2019. [1](#), [2](#)

Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task bayesian optimization. *Advances in neural information processing systems*, 26, 2013. [2](#)

Shion Takeno, Hitoshi Fukuoka, Yuhki Tsukada, Toshiyuki Koyama, Motoki Shiga, Ichiro Takeuchi, and Masayuki Karasuyama. Multi-fidelity bayesian optimization with max-value entropy search and its parallelization. In *International Conference on Machine Learning*, pages 9334–9345. PMLR, 2020. [1](#), [9](#)

Joost Van Amersfoort, Lewis Smith, Andrew Jesson, Oscar Key, and Yarin Gal. On feature collapse and deep kernel learning for single forward pass uncertainty. *arXiv preprint arXiv:2102.11409*, 2021. 3

Zifeng Wang, Shao-Lun Huang, Ercan E Kuruoglu, Jimeng Sun, Xi Chen, and Yefeng Zheng. Pac-bayes information bottleneck. *arXiv preprint arXiv:2109.14509*, 2021. 1

Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P. Xing. Deep kernel learning. volume 51 of *Proceedings of Machine Learning Research*, pages 370–378, Cadiz, Spain, 09–11 May 2016. PMLR. 3

Jian Wu, Saul Toscano-Palmerin, Peter I Frazier, and Andrew Gordon Wilson. Practical multi-fidelity bayesian optimization for hyperparameter tuning. In *Uncertainty in Artificial Intelligence*, pages 788–798. PMLR, 2020. 1, 9

Aolin Xu and Maxim Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. *Advances in neural information processing systems*, 30, 2017. 1, 4

## A Theorems and Proofs

We begin by introducing the necessary notion.

**Definition 1 (RKHS Salgia et al. (2024))** Consider a positive definite kernel  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ . A Hilbert space  $\mathcal{H}_k$  of functions on  $\mathcal{X}$  equipped with an inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$  is called a Reproducing Kernel Hilbert Space (RKHS) with reproducing kernel  $k$  if the following conditions are satisfied: (i)  $\forall x \in \mathcal{X}, k(\cdot, x) \in \mathcal{H}_k$ ; (ii)  $\forall x \in \mathcal{X}, \forall f \in \mathcal{H}_k, f(x) = \langle f, k(\cdot, x) \rangle_{\mathcal{H}_k}$ . For simplicity, we use  $\psi_x$  to denote  $k(\cdot, x)$ . The inner product induces the RKHS norm,  $\|f\|_{\mathcal{H}_k}^2 = \langle f, f \rangle_{\mathcal{H}_k}$ .

The following discusses the necessary assumptions, formal theorems, and corresponding proofs. We state the typical assumption for BO performance analysis first.

**Assumption 1** Throughout the optimization,  $\tilde{f}_M$  bears an upper bound  $B$  on the RKHS norm (Salgia et al., 2024) corresponding to the learned deep kernels  $k_t$ . Namely  $\|\tilde{f}_M\|_{\mathcal{H}_{k_t}} \leq B$ .

**Assumption 2** (Generalized Assumption 4.1 of Salgia et al. (2024)) For all  $n \in \mathbb{N}$ , there exists a discretization  $\mathcal{D}_n$  of  $\mathcal{X}$  such that for all  $f \in \mathcal{H}_k$ ,

$$|f(\mathbf{x}) - f([\mathbf{x}]_{\mathcal{D}_n})| \leq \|f\|_{\mathcal{H}_k} / n \quad \text{and} \quad |\mathcal{D}_n| = \text{poly}(n)^3,$$

where  $[\mathbf{x}]_{\mathcal{D}_n} = \arg \min_{y \in \mathcal{D}_n} \|\mathbf{x} - y\|_2$ , is the point in  $\mathcal{D}_n$  that is closest to  $\mathbf{x}$ .

**Assumption 3** (Assumption 4.2 of Salgia et al. (2024)) Let  $\mathcal{L}_\eta = \{x \in \mathcal{X} \mid f(x) \geq \eta\}$  denote the level set of  $f$  for  $\eta \in [-B, B]$ . We assume that for all  $\eta \in [-B, B]$ ,  $\mathcal{L}_\eta$  is a disjoint union of at most  $M_f < \infty$  components, each of which is closed and connected. Moreover, for each such component, there exists a bi-Lipschitzian map<sup>1</sup> between each such component and  $\mathcal{X}$  with normalized Lipschitz constant pair  $L_f, L'_f < \infty$ .

We then state the assumption for analysis of the statistical learning.

**Assumption 4** The hypothesis space for each fidelity  $\mathcal{H}_\ell$  contains the underlying functions  $f_\ell$  for  $\forall \ell \in [M]^+$ .

**Assumption 5** We assume the  $\forall \ell \in \{2, \dots, M\}$ ,  $\tilde{f}_\ell$  is  $L$ -Lipschitz relative to the function space  $\mathcal{H}_h$ , meaning  $\forall \mathbf{x} \in \mathbf{X}, \forall y \in \mathbf{Y}_\ell, \forall h, h' \in \mathcal{H}_h$ , we have  $|\mathcal{L}(y, g_\ell(h(\mathbf{x}))) - \mathcal{L}(y, g_\ell(h'(\mathbf{x})))| \leq L \mathcal{L}(g_{\ell-1}(h'(\mathbf{x})), g_{\ell-1}(h(\mathbf{x})))$ .

This generalizes the assumption of Theorem 10 from Robinson et al. (2020).

In the following section, we restate the key insights for this work.

<sup>1</sup>A map  $f : X \rightarrow Y$  is bi-Lipschitzian if there exist constants  $c_1, c_2 > 0$  such that  $c_1\|x - y\| \leq \|f(x) - f(y)\| \leq c_2\|x - y\|$  for all  $x, y \in X$ .

**Theorem 1** When Assumptions 1, 2, and 3 hold, constraining the random exploration on the target fidelity on  $\{x \in \mathbf{X} \mid \text{UCB}_{f_M, t}(x) > \max_{x' \in \mathbf{X}} \text{LCB}_{f_M, t}(x')\}$ , and choosing  $\beta = B^2$ , we have the following bound with probability at least  $1 - \delta$ ,

$$SR(t) = \tilde{\mathcal{O}}\left(\sqrt{\frac{\gamma_{T_M}(t)}{T_M(t)}} \log(T_M(t)/\delta)\right) \quad (1)$$

*Proof:* A direct extension of Theorem 4.3 and Theorem 4.5 of [Salgia et al. \(2024\)](#) by expanding the noise-free bound in its Theorem 4.3 to the noisy scenario bound in Theorem 4.5 is that when only considering single fidelity optimization applying random exploration in the region of interest defined above, with probability at least  $1 - \delta$ , the cumulative regret bears the upper bound  $\tilde{\mathcal{O}}(\sqrt{\gamma_{T_M}(t)} T_M(t) \log(T_M(t)/\delta))$ . Adding  $T_M(t)$  to the denominator converts it to a high probability upper bound of SR.  $\square$

**Theorem 2** Under the assumptions of Theorem 1 except for constraining random exploration in  $\hat{\mathbf{X}}_t$  as defined in equation 3. The misspecification-aware Bayesian simple regret ( $SR_{MA}$ ) of the proposed algorithm can be decomposed into the standard Bayesian simple regret (SR) and the rate term as follows:

$$SR_{MA}(t) \leq SR(t) + \text{Rate}_{MF}(M, t).$$

*Proof:* The result is a simple extension as we extend the region of interest to  $\hat{\mathbf{X}}_t$  as defined in equation 3. The additional term  $\text{Rate}_{MF}(M, t)$  count for the introduced excess risk.  $\square$

**Theorem 3** (Generalized Theorem 10 of [Robinson et al. \(2020\)](#)) With the aforementioned assumptions 1-5 hold, and the lowest single fidelity bears the convergence rate  $\text{Rate}_{MF}(1, T) = O(\sqrt{\frac{1}{T_1}})$ . The excessive risk bears the bound

$$\text{Rate}_{MF}(\ell, t) \leq O\left(\text{Rate}_{MF}(\ell - 1, t) + \frac{\sqrt{\left(\log_{T_\ell(t)} \text{Rate}_{MF}(\ell - 1, t) + 1\right) \log T_\ell(t)}}{T_\ell(t)}\right) \quad (2)$$

*Proof:* With aforementioned assumptions hold, for  $\forall 1 < \ell \leq M$ , the learning on fidelity  $\ell - 1$  and  $\ell$  meets the assumption of Theorem 10 of [Robinson et al. \(2020\)](#), then the bound above could be direct results of recursively applying the theorem 10 for  $\ell = 2 \dots M$ .  $\square$

## B Additional Algorithm Details

In the following, we offer the additional details of implementing the proposed algorithm RMFBO-DP.

**Reliable search space exclusion** We leverage the GP posterior calibrated with the excess risk on observed points to exclude from acquisition the regions that, with high probability, do not contain the global optimum. To do so, we rely on both the upper confidence bound  $\text{UCB}_{f_M, t}(x) \triangleq \mu_{f_M, t-1}(x) + \beta_t^{1/2} \sigma_{f_M, t-1}(x)$  and lower confidence bound  $\text{LCB}_{f_M, t}(x) \triangleq \mu_{f_M, t-1}(x) - \beta_t^{1/2} \sigma_{f_M, t-1}(x)$ , where  $\beta_t$  is the scaling factor corresponding to certain confidence. Formally, the acting search space at iteration  $t$  is  $\hat{\mathbf{X}}_t$  defined as

$$\left\{x \in \mathbf{X} \mid \text{UCB}_{f_M, t}(x) > \max_{x' \in \mathbf{X}} \text{LCB}'_{f_M, t}(x')\right\}. \quad (3)$$

Here  $\text{LCB}'_{f_M, t}(x') \triangleq \text{LCB}_{f_M, t}(x') - \text{Rate}_{MF}(M, t)$  generalize the lower confidence bound to the incorporate the expected generalization error.

---

**Algorithm 1** Robust Multi-Fidelity Bayesian Optimization with Deep Kernel Learning and Partition (RMFBO-DP)

---

```

1: Input: Search space  $\mathbf{X}$ ,  $t_1$  initial observation  $\mathcal{D}_{t_1}$ , total budget  $\Lambda$ ;
2: Initialize timestamp as  $t \leftarrow t_1$ .
3: Initialize remaining budget  $\Lambda \leftarrow \Lambda - \sum_{t'=1}^t \lambda_{\ell_{t'}}$ 
4: while  $\Lambda > 0$  do
5:   Update the model  $\tilde{f}$ .
6:   Identify ROIs  $\hat{\mathbf{X}}_t$  according to equation 3.
7:   Maximize cost-sensitive  $\text{SR}_{\text{MA}}(t+1) - \text{SR}_{\text{MA}}(t)$  reduction.

$$\Delta_{\text{SR},t} \leftarrow \frac{\text{SR}(t+1) - \text{SR}(t)}{\lambda_M}$$


$$\Delta_{\text{Rate}_{MF},t} \leftarrow \max_{\ell \in [M]^+} \frac{\text{Rate}_{MF}(\ell, t+1) - \text{Rate}_{MF}(\ell, t)}{\lambda_\ell}$$

8:   if  $\Delta_{\text{SR},t} \geq \Delta_{\text{Rate}_{MF},t}$  then
9:     Draw random sample  $\mathbf{x}_t$  from  $\hat{\mathbf{X}}_t$  on fidelity  $\ell_t$ 

$$\ell_t \leftarrow M$$


$$\mathbf{x}_t \leftarrow \max_{\mathbf{x} \in \hat{\mathbf{X}}} \alpha_{f_M}(\mathbf{x})$$

10:  else
11:    Sample candidate on on fidelity  $\ell_t$ 

$$\ell_t \leftarrow \arg \max_{\ell \in [M]^+} \frac{\text{Rate}_{MF}(\ell, t+1) - \text{Rate}_{MF}(\ell, t)}{\lambda_\ell}$$


$$\mathbf{x}_t \leftarrow \text{sample from uniform distribution on } \mathbf{X}.$$

12:  end if
13:  Update the observation  $\mathcal{D}_{t+1} \leftarrow \mathcal{D}_t \cup \{(\mathbf{x}_t, \ell'_t), y_t\}$ 
14:  Update the remaining budget  $\Lambda \leftarrow \Lambda - \lambda_\ell$ 
15:  Update the timestamp  $t \leftarrow t + 1$ 
16: end while

```

---

**Estimation of SR rates** Due to the difficulty of analyzing exact SR, we rely on the following simple approximation. For two consecutive evaluations of the target fidelity, if we observe improvement in the best reward, we leverage the improvement to regress the SR. Namely, for  $\forall 1 \leq t_1 < t_2 \leq T$ , if  $\Delta_{f_M} \triangleq y_{\langle \mathbf{x}_{t_2}, \ell_t = M \rangle} - y_{\langle \mathbf{x}_{t_1}, \ell_t = M \rangle} > 0$ , we update the approximation for  $\text{SR}(t_2)$  by solving  $\Delta_{f_M} = \text{SR}(t_2) - \text{SR}(t_1)$ .

**Estimation of expected excess risk** Similar to the above approximation of SR, we approximate the excess risk reduction by regressing to the observed fitting error improvement.  $\text{Rate}_{MF}(M, t)$ . for  $\forall 1 \leq t_1 < t_2 \leq T, \ell \in [M]^+$ , we resort to 5-fold cross-validation on  $\mathcal{D}_{\ell, t_1}$  and  $\mathcal{D}_{\ell, t_2}$  to estimate the model fitting improvement  $\Delta_{\mathcal{L}(f_\ell, \tilde{f}_\ell), t} = \text{Rate}_{MF}(\ell, t_2) - \text{Rate}_{MF}(\ell, t_1)$ . Solving the equation allows us to approximate  $\text{Rate}_{MF}(\ell, t_2)$ .

**Constraining acquisition** We rely on random discretization to constrain the acquisition within  $\hat{\mathbf{X}}_t$ , which rejects the candidates outside  $\hat{\mathbf{X}}_t$ . Note popular BO frameworks typically allow optimizing the acquisition function subject to constraints, e.g., Botorch (Balandat et al., 2020).

## C Experiments

We compare the proposed RMFBO-DP against four baselines. The first one is the entropy-based method denoted as MF-MES proposed by Takeno et al. (2020); the second one is denoted as MF-KG, which is the cost-efficient knowledge gradient method proposed by Wu et al. (2020). The third and fourth algorithms are corresponding variants when applying the robust MFBO framework proposed by Mikkola et al. (2023), denoted as rMF-MES and rMF-KG correspondingly. We rely on BoTorch (Balandat et al., 2020) and gpytorch (Gardner et al., 2018) to implement RMFBO-DP and the baselines.

## C.1 Dataset

We evaluate algorithm performance on synthetic datasets and a real-world multi-fidelity protein design task.

**Rastrigin dataset** As illustrated in figure 1, we construct the four-fidelity version of Rastrigin function (Pohlheim, 2007) on 1D search space. We further extend the construction to 20D search space. Here, the first two lower fidelities generally share the same trend as the target-fidelity underlying function, while the lowest fidelity that incurs the cheapest evaluation cost disagrees with the target fidelity function except for limited central area and largely diverges in the border areas.

**Multi-fidelity protein design** We use a protein engineering dataset describing a set of antigen/antibody binding calculations. These calculations, executed using supercomputing resources, estimate the change in binding free energy at the interface between each of the 71769 modified antibodies and the SARS-CoV-2 spike protein, as compared to the single reference antibody from which they are derived. Estimations of binding free energy ( $\Delta\Delta G$ ) are calculated using protein-structure-based Rosetta Flex simulation software (Das and Baker, 2008; Barlow et al., 2018) and FoldX (Schymkowitz et al., 2005; Sapozhnikov et al., 2023; Buß et al., 2018). We treat Rosetta’s outcomes as the objective of the target fidelity. These calculations took several CPU hours each and were produced during an antibody design process (Desautels et al., 2020, 2022).

## C.2 Evaluation

We evaluated the proposed algorithm Robust Multi-fidelity Bayesian Optimization (RMFBO-DP) against four baselines on both synthetic datasets corresponding to figure 1 and four real-world tasks. We’ve shown that the proposed algorithm outperforms the baselines in terms of simple regret on Rastrigin-1D, Rastrigin 20d, and Protein-88D.

## C.3 Ablation Study

We conduct an ablation study to investigate the impact of the proposed algorithm components. We compare the performance of the additional variants of the baselines when applying the same deep kernel learning yet without random sampling on low-fidelities. As is shown in table 1, the performance of the variants is not consistently improved upon the corresponding baselines and lags behind RMFBO-DP. This observation suggests that data acquisition is crucial to the performance improvement of the proposed RMFBO-DP.

Method	Rastrigin-1D	Rastrigin-20D
<b>RMFBO-DP</b>	<b><math>0.75 \pm 0.30</math></b>	<b><math>86.64 \pm 9.80</math></b>
MF-MES	$15.00 \pm 6.83$	$105.66 \pm 4.25$
MF-KG	$1.84 \pm 0.61$	$106.03 \pm 6.96$
MF-MES-DK	$2.83 \pm 2.57$	$106.62 \pm 5.72$
MF-KG-DK	$7.04 \pm 4.60$	$110.00 \pm 0.00$
rMF-MES	$8.33 \pm 7.36$	$100.93 \pm 8.87$
rMF-KG	$3.02 \pm 1.18$	$103.37 \pm 6.41$
rMF-MES-DK	$14.88 \pm 7.65$	$103.20 \pm 6.24$
rMF-KG-DK	$3.37 \pm 1.23$	$103.52 \pm 5.32$
Random	$5.05 \pm 2.43$	$104.06 \pm 6.15$

Table 1: Ultimate simple regrets for Rastrigin-1D and Rastrigin-20D tasks. The best-performing algorithm for each task is highlighted in bold. The results are collected from at least ten independent trials. We mark the variants using a deep kernel with “-DK”.