

Lessons Learned from Applying Bayesian Optimization to Hyperparameter Tuning in Ad Selection

Joobin Gharibshah*, Mahmuda Rahman

eBay, Inc.

Abstract

Bayesian Optimization (BO) has become a key technique for hyperparameter tuning in ad selection within search monetization. Despite its efficiency, BO faces significant challenges in real-world deployments, including scalability issues in high-dimensional spaces, slow convergence, high computational costs, and difficulties in adapting to dynamic environments. This paper explores these challenges through the lens of industrial application of the AFA (Auto-tuning Filters for Ads) framework and proposes practical solutions, including reinforcement learning techniques, hybrid optimization models, and constraint-aware Bayesian approaches.

Keywords

Bayesian Optimization, Ad Selection, Hyperparameter Tuning, Search Monetization

1. Introduction

Optimizing ad selection in search monetization is essential for balancing platform revenue with user experience. Hyperparameter tuning plays a critical role in this process, and Bayesian Optimization (BO) has emerged as a powerful technique due to its sample efficiency and ability to navigate complex, high-dimensional search spaces. In our previous work, we introduced AFA (Auto-tuning Filters for Ads)[1] a BO-based framework that automates the tuning of ad eligibility filters, reducing manual overhead and accelerating experimentation.

While BO offers clear advantages over traditional methods such as grid or random search, its application in large-scale, real-world ad systems reveals a number of practical challenges. During the industrial deployment of AFA, we encountered several limitations, including scalability in high-dimensional parameter spaces, slow convergence, cold-start behavior, high computational costs, sensitivity to noisy evaluations, difficulty in managing multi-objective trade-offs, and the integration of domain-specific business constraints.

These limitations do not undermine the value of BO, but rather highlight areas where enhancements are needed to make it robust, adaptable, and efficient in production environments. In this paper, we share the lessons learned from deploying AFA at scale and explore strategies including reinforcement learning, hybrid optimization models, constraint-aware Bayesian methods, and noise-resilient techniques to overcome these challenges and extend the practical applicability of BO in ad selection systems.

In the AFA framework[1], we applied BO to automate the tuning of ad filter thresholds and key parameters that govern ad eligibility and directly impact auction volume, user experience, and revenue. AFA iteratively proposed candidate threshold values, evaluated their impact through live A/B testing, and updated a surrogate model to guide further exploration. This approach significantly reduced the manual burden of parameter tuning and improved experimentation velocity. However, as we scaled AFA across diverse traffic segments and filter types, several limitations in standard BO surfaced, motivating the deeper analysis and improvements discussed in this paper.

ECOM'25: SIGIR Workshop on eCommerce, Jul 17, 2025, Padua, Italy

*Corresponding author.

✉ jgharibshah@ebay.com (J. Gharibshah); mahrahman@ebay.com (M. Rahman)

🆔 0000-0002-0877-7063 (J. Gharibshah); 0000-0001-7116-9338 (M. Rahman)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. Limitations of Bayesian Optimization

While Bayesian Optimization (BO) offers strong theoretical foundations and performs well in controlled settings, its application to large-scale ad selection systems presents several real-world challenges. Through the deployment of AFA, we encountered practical limitations that impacted both the performance and applicability of BO in production. These challenges arose from system constraints, data noise, market dynamics, and the complexity of balancing business objectives.

In this section, we describe the key limitations we observed and the lessons we learned, organized across three categories: algorithmic, operational, and statistical challenges.

2.1. Challenges in Adapting to Non-Stationary Environments

Bayesian Optimization assumes a stationary objective function, but in real-world ad selection, this assumption breaks down as user behavior, auction dynamics, and advertiser strategies shift frequently. During the deployment of AFA, we found that configurations optimized under one set of conditions often became suboptimal within days due to external factors such as seasonality, policy changes, or new campaign launches.

While retraining the model more frequently helped, BO's reliance on historical data made it prone to stale recommendations and lacked mechanisms to weigh recency or detect drift. To address this, we introduced an always-on A/B testing strategy that continuously collected fresh feedback from a rotating set of configurations. This design discussed in Section 3.1 acted as a lightweight, pseudo-reinforcement learning loop, enabling more responsive adaptation to market dynamics.

2.2. Slow Convergence and Cold Start Problem

Another practical limitation of Bayesian Optimization is its reliance on initial function evaluations to build an accurate surrogate model. This leads to a cold start phase, during which BO's predictions may be highly uncertain or misinformed. In real-time ad systems, where rapid decision-making and short feedback cycles are crucial, this slow start can hinder the optimizer's effectiveness.

In our deployment of AFA, this challenge was especially visible during the launch of new models or ad filters, where BO lacked prior information about system response. To mitigate this, we seeded the optimizer with a batch of initial observations collected using expert-informed configurations. These initial points were selected to cover a diverse range of the search space and reflect practical edge cases from past deployments.

This approach improved early surrogate accuracy and reduced the number of non-informative suggestions. While it did not eliminate the cold start entirely, it helped BO make meaningful progress within the limited number of iterations allowed in production experimentation cycles.

2.3. High Computational Costs in Real-Time Applications

Ad selection systems operate in dynamic, latency-sensitive environments where decisions must be made within milliseconds. However, traditional Bayesian Optimization is not designed for such real-time constraints. Each iteration involves updating a surrogate model and solving an acquisition function, both of which introduce non-trivial computational overhead—particularly when using Gaussian Processes or optimizing over high-dimensional spaces.

However, the most significant cost is often not computational in the traditional sense: it is the operational and business cost of running live A/B tests to evaluate new configurations. Each evaluation requires deploying a candidate parameter into production traffic, allocating user impressions, and collecting statistically meaningful metrics. These tests are costly in terms of time, risk, and infrastructure, and became the dominant bottleneck in our AFA deployment.

This dual burden—computational and operational—limits the frequency and responsiveness of BO in production. It motivated our exploration of more parallel, simulation-informed, and fidelity-aware optimization methods discussed in Section 3.4.

2.4. Challenges with Multi-Objective Optimization Constraints

Ad selection involves optimizing multiple, often competing goals such as maximizing revenue, maintaining user engagement, and satisfying advertiser fairness requirements. Standard BO, designed for single-objective optimization, becomes limiting when there are inherent trade-offs.

In AFA, we used a scalarization strategy to combine multiple objectives into a single metric. While this reduced complexity and testing cost, it imposed a fixed weighting scheme that did not generalize well across traffic segments. It also obscured the trade-off surface, limiting insight into how decisions affected different objectives.

The scalarization strategy we applied was defined by:

$$f(t) = \bar{t} = \arg \max_{t \in \text{Thresholds}} (\delta_s(t) - \delta_p^*(t))$$

where $\delta_s(t)$ denotes the scaled lift in the engagement or user experience objective. $\delta_p^*(t)$ denotes the penalized lift in platform performance (e.g., revenue). $\bar{t}$ is the threshold that maximizes the trade-off between these competing goals.

This limitation revealed the need for more flexible techniques like Multi-Objective Bayesian Optimization (MOBO), which can surface the Pareto frontier, but often at higher computational and testing costs. We will discuss the details in section 3.4

2.5. Impact of Noise on Optimization Reliability

Ad system evaluations are noisy by nature. Even with identical configurations, A/B test results can vary due to traffic fluctuations, auction randomness, or overlapping experiments. BO's assumption of consistent function evaluations is violated, leading to instability.

In AFA, we observed slower convergence and occasional model misdirection caused by high-variance points. To mitigate this, in AFA, we applied manual techniques such as result averaging and outlier filtering, but these added latency and operational complexity.

This motivated our exploration of noise-aware modeling and adaptive evaluation strategies—covered in Section 3.5.

2.6. Incorporating Business Constraints

Unlike in traditional machine learning tasks, every configuration evaluated in ad selection has real business consequences. Poor filter thresholds can lead to revenue loss, advertiser dissatisfaction, or user churn, impacts that are not always directly captured in the optimization objective, but must be monitored and respected to ensure business viability.

In AFA, we had to operate within guardrails such as minimum quality guarantees, advertiser fairness, compliance policies, and system latency. However, Bayesian Optimization offers no native mechanism to encode hard or soft constraints into its decision process, forcing us to rely on external workarounds.

To enforce these constraints, we embedded business logic into the candidate generation step and deprioritized configurations with unacceptable outcomes. While this approach worked in practice, the rules were often static and brittle, requiring manual updates and limiting BO's flexibility.

These challenges highlight the gap between theoretical BO frameworks and production realities. Constraint-aware Bayesian Optimization methods explored in Section 3.6 offer a more robust and scalable alternative.

2.7. Scalability in High-Dimensional Search Spaces

Bayesian Optimization is known to struggle in high-dimensional spaces due to the computational cost of fitting surrogate models and optimizing acquisition functions. In AFA, we found that single-parameter tuning converged in around three iterations. With two parameters, convergence stretched to ten or more iterations. The scalability ceiling was clear, as shown in Table 1.

Table 1

Observed convergence behavior in AFA as dimensionality increased. As more parameters were tuned, convergence slowed significantly.

# Parameters	Avg. Iterations	Notes
1	~3	Fast convergence
2	~10	Partial convergence observed
3+	>15	Not attempted in production deployment

Although we did not attempt to solve this limitation in the AFA deployment, the issue highlights the need for scalable surrogate models or dimensionality reduction techniques when applying BO to larger configuration spaces. We discuss potential solutions in Section 3.7.

3. Enhancement for Bayesian Optimization

Bayesian Optimization (BO) offers significant potential for improving ad selection systems, but its successful deployment in production environments requires several adaptations. Each of the limitations discussed in Section 2 reveals a critical friction point that we encountered during the development and rollout of AFA. In this section, we propose corresponding solutions that map directly to those challenges.

Although some solutions may span multiple limitations, such as parallelization improving both latency and scalability, or always-on testing enhancing both adaptability and noise robustness, we maintain a one-to-one correspondence in structure for clarity.

3.1. Enhancing Adaptation to Non-Stationary Environments with Reinforcement Learning

Since BO assumes a static objective function, it struggles to adapt to rapidly changing market conditions in ad selection. A more flexible approach is to integrate reinforcement learning (RL), which continuously learns from real-time interactions and updates optimization strategies accordingly [2, 3].

Policy-based RL techniques, such as Proximal Policy Optimization (PPO), are particularly well suited for ad selection. These methods allow stable policy updates while preventing abrupt changes that could negatively impact performance [4]. RL frameworks have also shown success in real-time bidding and online personalization settings, where optimization must evolve continuously based on new signals [5]. Recent work has also proposed combining Bayesian Optimization with local policy search to actively guide exploration and reduce sample variance in RL [6].

In AFA, rather than fully replacing BO with RL, we implemented a hybrid strategy that mimicked RL’s feedback loop while retaining the sample-efficiency of BO. This involved deploying an *always-on A/B testing* framework where a rotating pool of candidate configurations was continuously evaluated. The surrogate model was incrementally updated with fresh data, allowing the optimizer to adjust to environmental changes.

This pseudo-RL setup bridged the gap between traditional Bayesian Optimization and full reinforcement learning by preserving BO’s sample efficiency while introducing the adaptability required for dynamic ad marketplaces. Although not a complete RL policy learner, the hybrid approach provided a practical compromise—allowing continuous adaptation through incremental updates without incurring the complexity of full RL deployment. This structure laid the groundwork for future exploration of BO-RL hybrids in constrained experimentation settings.

Recent advances reinforce this direction. For example, Müller et al. [6] combine BO with local policy search to actively select informative samples and improve sample efficiency in dynamic environments. Other work has explored time-aware BO using Gaussian processes with temporal kernels [7], as well as bandit-based methods designed for non-stationary objectives [8], further highlighting the potential of adaptive and hybrid approaches for real-world optimization.




Figure 1: The always on AB testing pipeline to tackle multiple limitations such as non-stationary environment

3.2. Accelerating Convergence and Addressing Cold Start with Informed Initialization

To address the slow convergence and cold start limitations outlined in Section 2.2, we adopted several techniques to better initialize the optimization process and improve sample efficiency in early iterations.

First, we incorporated expert-informed initialization by seeding the optimizer with configurations that were known to be safe or representative, based on domain knowledge. This significantly improved early surrogate quality and reduced the number of non-informative or poor-performing suggestions.

Second, we emphasized diversity in initial sampling. Instead of relying solely on random exploration, we sampled a wide range of thresholds from the feasible space—ensuring the surrogate had a good prior across the input space.

These steps helped mitigate the cold start problem and accelerated convergence. While BO still required several iterations to refine the search, our initialization strategy reduced the time needed to reach competitive configurations. This approach resulted in a 25% reduction in the number of experiments and data points required, significantly lowering the overall cost of A/B testing.

Similar strategies have been explored in prior work using meta-learning or transfer learning to warm start BO in new domains [9, 10]. Other work emphasizes the importance of diverse initial sampling to build well-calibrated surrogate models early in the process [11].

3.3. Reducing Real-Time Computational and Evaluation Costs via Parallel and Multi-Fidelity BO

The iterative nature of BO contributes to significant computational overhead, making it challenging for real-time ad selection where decisions must be made within a short period of time. One effective solution is parallel Bayesian Optimization, which enables multiple configurations to be evaluated simultaneously.

This can be achieved through batch BO methods, such as Thompson Sampling-based Batch BO, which allows multiple candidate solutions to be explored in parallel, reducing the sequential dependency of traditional BO [12].

Another promising approach is multi-fidelity Bayesian Optimization, where approximate, lower-cost evaluations are used alongside high-fidelity ones to accelerate convergence. In ad selection, simulated auction environments or historical campaign data can serve as low-fidelity proxies [13].

In our deployment of AFA, we experimented with the multi-fidelity approach by incorporating lower-cost approximations (e.g., simulated environments) alongside live traffic evaluations. However, we found that this introduced noise and bias into the surrogate model, occasionally leading to misleading optimization guidance. As a result, this technique proved less effective in our setting, where fidelity

mismatches between offline and online evaluations were difficult to calibrate.

3.4. Addressing Multi-Objective Optimization with MOBO and Pareto Strategies

Ad selection involves optimizing multiple, often competing objectives, such as maximizing revenue while maintaining user engagement and ensuring fair ad distribution. Traditional BO is inherently designed for single-objective optimization, which limits its ability to handle complex trade-offs.

In AFA, we addressed this limitation using a scalarization strategy, where objectives were combined using predefined weights into a single optimization target. While this simplified the optimization and reduced the number of required A/B tests, it locked in trade-off preferences and failed to reveal the broader Pareto landscape. It reduced the number of AB test needed by 50%.

To improve flexibility, we can explore Multi-Objective Bayesian Optimization (MOBO) approaches, particularly those that rely on random scalarizations or directly construct Pareto frontiers [14]. These methods allow systems to uncover a set of non-dominated solutions and enable downstream decision-makers to select from among desirable trade-offs. However, we should expect longer convergence time and higher computation cost relatively.




Figure 2: Scalarization vs. Pareto Front in Multi-Objective Optimization. The plot shows a convex Pareto front representing non-dominated solutions across two competing objectives. A scalarized solution (black star) selects a single point based on predefined weights, while Pareto-based methods allow for flexible trade-offs.

3.5. Improving Robustness to Noisy Evaluations in Bayesian Optimization

A critical challenge in production optimization arises when repeated A/B tests for the same configuration yield different outcomes due to stochastic variability. This inconsistency can mislead BO's surrogate model, slowing down convergence and promoting suboptimal decisions.

We addressed this issue in AFA by experimenting with several strategies:

Input Noise Modeling, where the surrogate explicitly accounts for observed variance. This can be achieved using models that estimate both the mean and variance of outcomes [15].

Bootstrapped Ensembles, where multiple surrogates are trained on resampled data to capture uncertainty in predictions. This ensemble-based method improves robustness by averaging across noisy estimates [16].

Adaptive Re-evaluation, where high-variance configurations are automatically re-tested to stabilize the model's belief. This technique helps avoid overreacting to outliers in early evaluations [17].

These strategies improved BO's resilience to noisy data and reduced the likelihood of overfitting to outlier outcomes observed in short-term experiments. We also found that integrating these techniques with always-on A/B testing (Figure 1) further improved the stability and adaptivity of the optimization process in dynamic environments. This combination led to an 11% improvement in performance metrics and a 50% reduction in the number of required experiment and evaluation, significantly lowering experimentation costs.

3.6. Embedding Business Constraints into the Optimization Process

Hyperparameter tuning in ad selection often requires real-world A/B testing, live traffic allocation, and continuous evaluation based on revenue impact. Unlike many machine learning tasks that allow for offline training, ad selection systems must operate within latency, fairness, and policy constraints.

Bayesian Optimization does not natively support such hard or soft constraint, which may involve dynamic thresholds or non-differentiable boundaries. To address this, we embedded empirical safety rules and platform-specific heuristics into the candidate generation and filtering process during our AFA deployment. To embed business constraints into the optimization process, AFA applied a heuristic penalization strategy within the objective function:

$$f(t) = \delta_s(t) - \lambda \cdot \delta_p(t)$$

where $\delta_s(t)$ is the scale lift, $\delta_p(t)$ is the performance penalty, and λ balances the trade-off. The penalty term ensured configurations met quality thresholds while maximizing reach, with λ calibrated to reflect acceptable business risk.

Recent advances in constraint-aware Bayesian Optimization provide more principled approaches. For example, acquisition functions can be extended to penalize infeasible regions [18], or surrogate models can be trained to learn feasibility constraints directly [19]. These techniques are promising for high-stakes optimization under real-world business restrictions [20].

3.7. Scaling BO for High-Dimensional Parameter Spaces with Efficient Surrogates

Bayesian Optimization faces well-known challenges in high-dimensional parameter spaces, where Gaussian Process surrogates become computationally expensive and less accurate. In our deployment, we found that performance dropped significantly as the number of tuned parameters increased beyond two.

To address this, we explored more scalable surrogate models such as Tree-structured Parzen Estimators (TPE) [21] and Bayesian Neural Networks (BNNs) [20], both of which handle higher dimensionality more efficiently than GPs.

We also investigated dimensionality reduction techniques, including low-rank kernel approximations and random embedding methods, which project the problem into a lower-dimensional latent space while preserving relevant structure [22].

Recent work has also proposed batch MOBO strategies that emphasize diversity across the Pareto front, improving exploration and sample efficiency in complex, multi-objective settings [23].

These approaches significantly reduced computational cost and improved convergence behavior in settings with five or more parameters.

3.8. Integrated Deployment of AFA: Real-World Gains in Efficiency and Performance

Taken together, the strategies outlined in this section were implemented as part of the AFA (Auto-tuning Filters for Ads) framework, which integrated multiple enhancements to overcome the limitations of standard Bayesian Optimization in real-world ad systems. By combining techniques such as expert-informed initialization, always-on A/B testing, noise-aware modeling, constraint filtering, and hybrid BO-RL mechanisms, AFA significantly improved both tuning speed and operational efficiency. In our production deployment, AFA reduced the average number of A/B test iterations required to reach a satisfactory configuration from over 10 to fewer than 5, while maintaining comparable or improved performance metrics by 11% points. Additionally, the average time from initial experiment launch to production deployment decreased by approximately 65%, with measurable gains in CTR stability and impression lift across multiple campaigns. These results demonstrate that the systematic application of enhanced BO strategies can yield robust, scalable, and interpretable optimization pipelines suitable for dynamic, high-stakes environments like ad selection.

4. Background and Related Work

Bayesian Optimization (BO) has become a standard approach for optimizing expensive black-box functions, with applications ranging from hyperparameter tuning in machine learning [24, 20], to scientific experimentation [25, 26], and engineering design [13]. In the context of advertising, BO has shown promise in automating critical system decisions such as ad filtering and ranking.

Dynamic and Non-Stationary Optimization Real-world systems, including ad selection, operate in dynamic environments where the objective function evolves over time. Standard BO assumes a stationary objective, making it ill-suited for such cases. Recent efforts have introduced time-aware BO and hybrid Reinforcement Learning (RL) strategies that update policies based on recent feedback. Proximal Policy Optimization (PPO) [4], model-agnostic meta-learning (MAML) [9], and more recently, BO-guided local policy search frameworks [6] have demonstrated promise in adapting to non-stationarity while maintaining sample efficiency.

Handling Noisy and Uncertain Evaluations A/B testing in ad systems introduces considerable noise due to traffic variability, auction randomness, and overlapping experiments. To handle such uncertainty, methods like input noise modeling [15], ensemble-based surrogates [16], and adaptive re-evaluation strategies [17] have been explored. These methods improve the robustness of BO under stochastic feedback.

Constraint-Aware and Multi-Objective Optimization Advertising platforms often face multiple competing objectives such as maximizing revenue while ensuring user satisfaction. Scalarization techniques remain the most widely used solution to convert multi-objective problems into single-objective form, though they may obscure trade-off dynamics. Multi-objective Bayesian Optimization (MOBO) techniques like Pareto front modeling [14] enable more flexible exploration of trade-offs. Furthermore, constraint-aware BO methods [19, 18] extend acquisition functions or model feasibility to respect domain-specific business constraints during optimization.

Scalability and Surrogate Modeling in High-Dimensional Spaces Traditional BO methods struggle with scalability as dimensionality increases. Gaussian Processes (GPs), while popular, become computationally expensive and less reliable in high-dimensional settings. To address this, scalable surrogate models such as Tree-structured Parzen Estimators (TPE) [21] and Bayesian Neural Networks (BNNs) [20] have been proposed. Dimensionality reduction methods and random embedding techniques [22] also provide a way to operate efficiently in large parameter spaces. Additionally, batch and parallelized BO methods such as Thompson Sampling-based strategies [12] and recent Pareto-diverse batch MOBO methods [23] improve computational efficiency and diversity of exploration.

Bayesian Optimization in Advertising Systems In the context of ad selection, AFA (Auto-tuning Filters for Ads) [1] demonstrated the feasibility of using BO to automate filter threshold tuning in search monetization. While this approach reduced manual experimentation overhead and improved filter deployment efficiency, it also highlighted key limitations of BO in production environments—particularly with respect to system dynamics, scalability, and real-time evaluation constraints. These insights form the foundation of the lessons and adaptations discussed in this paper.

5. Conclusion

Bayesian Optimization (BO) offers a sample-efficient, principled approach to hyperparameter tuning in ad selection. From our deployment of AFA, we identified key challenges—such as cold starts, computational cost, noise, non-stationarity, and business constraints—and proposed practical adaptations including hybrid BO-RL methods, parallelism, scalarization, and constraint-aware modeling. While BO has limitations, many represent opportunities for innovation. Future work should advance hybrid, context-aware, and scalable optimization to make BO more robust and production-ready in complex ad systems.

References

- [1] J. Gharibshah, M. Rahman, A. Bagherjeiran, Auto-tuning filters for ads: Scalable bayesian optimization for ad selection in search, in: Proceedings of the KDD, 2024.
- [2] H. S. Jomaa, J. Grabocka, L. Schmidt-Thieme, Hyp-rl : Hyperparameter optimization by reinforcement learning, 2019. URL: <https://arxiv.org/abs/1906.11527>. arXiv:1906.11527.
- [3] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, A general reinforcement learning algorithm that masters chess, shogi, and go through self-play, Science 362 (2018).
- [4] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, arXiv preprint arXiv:1707.06347 (2017).
- [5] G. Theodorou, P. Thomas, M. Ghavamzadeh, A reinforcement learning framework for real-time bidding advertising, in: Proceedings of the 21th ACM SIGKDD, 2015.
- [6] S. Müller, A. von Rohr, S. Trimpe, Local policy search with bayesian optimization, 2021. URL: <https://arxiv.org/abs/2106.11899>. arXiv:2106.11899.
- [7] I. Bogunovic, J. Scarlett, V. Cevher, Truncated variance reduction: A unified approach to bayesian optimization and level-set estimation, in: NeurIPS, 2016.
- [8] T. Nguyen, A. Rivera, J. Gonzalez, Non-stationary bayesian optimization with hypernetwork surrogates, ICML (2022).
- [9] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in: International Conference on Machine Learning (ICML), 2017.
- [10] M. Feurer, J. Springenberg, F. Hutter, Initializing bayesian hyperparameter optimization via meta-learning, in: AAAI, 2015.
- [11] F. Hutter, H. H. Hoos, K. Leyton-Brown, Sequential model-based optimization for general algorithm configuration, in: LION, 2011.
- [12] K. Kandasamy, G. Dasarthy, J. Schneider, B. Póczos, Parallelised bayesian optimisation via thompson sampling, in: AISTATS, 2018.
- [13] R. Sen, P. Choudhury, S. Das, V. Ravi, Multi-fidelity bayesian optimization with continuous approximations, arXiv preprint arXiv:1803.05490 (2018).
- [14] B. Paria, K. Kandasamy, B. Póczos, A flexible multi-objective bayesian optimization approach using random scalarizations, in: International Conference on Machine Learning (ICML), 2019.
- [15] J. M. Hernández-Lobato, M. W. Hoffman, Z. Ghahramani, Predictive entropy search for bayesian optimization with unknown constraints, in: ICML, 2014.
- [16] B. Lakshminarayanan, A. Pritzel, C. Blundell, Simple and scalable predictive uncertainty estimation using deep ensembles, in: Advances in Neural Information Processing Systems, 2017.
- [17] B. Letham, B. Karrer, G. Ottoni, E. Bakshy, Constrained bayesian optimization with noisy experiments, Bayesian Analysis 14 (2019).
- [18] J. Gardner, M. Kusner, Zhixiang, K. Weinberger, J. Cunningham, Bayesian optimization with inequality constraints, in: E. P. Xing, T. Jebara (Eds.), Proceedings of the 31st International Conference on Machine Learning, volume 32 of *Proceedings of Machine Learning Research*, PMLR, Beijing, China, 2014.
- [19] M. A. Gelbart, J. Snoek, R. P. Adams, Bayesian optimization with unknown constraints, in: Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, 2014.
- [20] J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, G. E. Dahl, R. P. Adams, Scalable bayesian optimization using deep neural networks, in: International Conference on Machine Learning (ICML), 2015.
- [21] S. Watanabe, Tree-structured parzen estimator: Understanding its algorithm components and their roles for better empirical performance, 2023. URL: <https://arxiv.org/abs/2304.11127>. arXiv:2304.11127.
- [22] Z. Wang, S. Jegelka, Batched high-dimensional bayesian optimization via structural kernel learning, in: Proceedings of the 34th International Conference on Machine Learning, 2017.
- [23] S. Belakaria, A. Deshwal, S. Srinivasan, R. Garnett, S. Shukla, Pareto front-diverse batch multi-

objective bayesian optimization, in: International Conference on Machine Learning (ICML), 2023.

- [24] J. Bergstra, R. Bardenet, Y. Bengio, B. Kégl, Algorithms for hyper-parameter optimization, in: Advances in Neural Information Processing Systems, 2011.
- [25] D. Golovin, B. Solnik, S. Moitra, G. Kochanski, J. Karro, D. Sculley, Google vizier: A service for black-box optimization, in: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), ACM, 2017.
- [26] I. Dey, C. Rea, Z. Peterson, K. Olofsson, J. Parker, S. Radhakrishnan, A. Svyatkovskiy, A. R. Field, O. Meneghini, Bayesian optimization accelerates and improves tokamak experiment design, Nature 614 (2023).