

Adaptive Experimental Design to Accelerate Scientific Discovery and Engineering Design

Aryan Deshwal

University of Minnesota

adeshwal@umn.edu

Abstract

Artificial Intelligence (AI) and Machine Learning hold immense potential to accelerate scientific discovery and engineering design. A fundamental challenge in these domains involves efficiently exploring a large space of designs or hypotheses using expensive experiments in a resource-efficient manner. This paper surveys novel adaptive experimental design methods to address this broad challenge. Specifically, we discuss new probabilistic modeling and decision-making techniques that are applicable in small data settings. These approaches have shown substantial improvements in sample-efficiency, particularly for black-box optimization over high-dimensional combinatorial spaces (e.g., sequences and graphs) and a variety of goals ranging from multiobjective to multi-fidelity optimization. This paper outlines key methods and their real-world sustainability applications in areas such as nanoporous materials discovery, hardware design, surfactant design, and additive manufacturing.

systems etc.) to minimize the number of expensive evaluations needed to achieve our design goals? We refer to such problems as experimental design problems in this paper.

The key objective in experimental design is to achieve a specific design goal or learn about an unknown quantity of interest in a resource-efficient manner [Deshwal, 2024; Mutn̄y, 2024]. Different design goals lead to distinct instantiations of experimental design. For instance, active learning is a widely studied problem setting where the goal is to iteratively select data examples to be labeled, so as to learn a predictor that achieves highest accuracy using as few labeled examples as possible. Similarly, in bandit problems, the goal is to maximize the cumulative reward over time by iteratively selecting an arm from a set of candidates. Another widely applicable problem setting is expensive *black-box optimization* where the goal is to optimize an objective function while minimizing the number of function evaluations. For example, searching the space of materials to maximize a desired property while minimizing the total resource-cost of physical lab experiments for their evaluation [Terayama *et al.*, 2021].

1 Introduction

Accelerating scientific discovery and engineering design can help us tackle many sustainability challenges facing our society, ranging from climate change to health, water, food, and manufacturing. Indeed, we are at the cusp of an AI-driven revolution in science and engineering [Jumper *et al.*, 2021; Hopfield and Hinton, 2024; Karpatne *et al.*, 2025]. For example, AI-driven rapid discovery of nanoporous materials has the potential to solve some of society’s biggest challenges, from absorbing carbon dioxide from air to storing hydrogen gas for fuel [Ding *et al.*, 2019]. Similarly, AI-driven drug/vaccine design could help find significantly better therapeutics for major human ailments [Schneider *et al.*, 2020]. The design of high-performance and low-power hardware by overcoming Moore’s law will enable sustainable computing from edge to cloud [Huang *et al.*, 2021]. However, there is one recurring technical challenge that arise across many such science and engineering applications: How can we design sequences of experiments that efficiently explore large combinatorial design spaces (molecules, materials, manycore

In this paper, we first survey a series of methods focusing primarily on under-studied challenges in the black-box optimization setting that arise in science and engineering applications. First, many scientific domains require handling structured design spaces which can be discrete or hybrid (a mixture of discrete and continuous variables). Second, there is a need to handle a variety of optimization goals, including finding a single best design (single-objective optimization), approximating the optimal Pareto front that trades off multiple competing objectives (multi-objective optimization), and leveraging different fidelities of experiments (multi-fidelity optimization). We describe novel adaptive experimental design algorithms developed within the framework of Bayesian optimization (BO) to address these challenges. Subsequently, we discuss multiple applications of these methods to real-world science and engineering problems in collaboration with domain-scientists. As we will discuss later, the underlying principles behind these approaches extend to related settings and to multiple experimental design goals beyond optimization, including discovering feasible candidates and parameter estimation [Hammond *et al.*, 2026].

2 Novel Bayesian Optimization Methods

Bayesian optimization is an effective framework for tackling the challenge of expensive black-box optimization [Shahriari *et al.*, 2015]. There are two key components of this framework: 1) Surrogate modeling from past experimental data to make predictions and quantify their uncertainty; and 2) Decision-making policies (acquisition functions) to select one or more candidate experiments depending on the goal. We discuss advances in both components next.

2.1 Probabilistic Surrogate Models over Combinatorial Structures

Design spaces in scientific domains (e.g., molecules, proteins) are naturally represented as combinatorial objects such as attributed graphs and high-dimensional categorical sequences, which pose hard challenges in probabilistic modeling, especially in the small supervised data regime. We addressed this challenge by developing new Gaussian process (GP) based probabilistic models over different kinds of combinatorial structures [Deshwal *et al.*, 2023; Deshwal *et al.*, 2021b; Deshwal *et al.*, 2022; Deshwal *et al.*, 2021c]. We discuss a subset of these approaches below.

In BODI [Deshwal *et al.*, 2023], we developed a novel dictionary embeddings based GP model for handling the challenge of high-dimensional binary/categorical spaces. The key idea is to select a number of discrete structures from the input space (the dictionary) and use them to define an ordinal embedding for high-dimensional combinatorial structures. This allows us to use existing Gaussian process models for continuous spaces. The paper provides two dictionary construction approaches including a principled approach based on binary wavelets to construct dictionaries for binary spaces, and a randomized construction method that generalizes to categorical spaces. Additionally, it provides a theoretical justification in terms of improved regret bounds to support the effectiveness of the dictionary-based embeddings. The efficacy of this approach is demonstrated on diverse real-world benchmarks from domains including communications engineering, maximum satisfiability, pest contamination control, and Auto ML.

HyBO [Deshwal *et al.*, 2021b] addresses the problem of black-box optimization over hybrid structures (mixture of discrete and continuous input variables). Hybrid spaces generalize continuous and discrete spaces and commonly arise in applications such as microbiome design and AutoML. The key challenge is to accurately model the complex interactions between discrete and continuous variables. We develop a GP surrogate model for hybrid spaces by designing a novel diffusion kernel via an additive GP formulation, and prove a universal approximation property for this kernel.

Recently, there is growing interest in leveraging deep generative models (DGM) for solving black-box optimization problems over combinatorial spaces. One such approach is to run the BO procedure over the continuous latent space of generative models such as variational autoencoders. However, the surrogate model defined over the latent space only uses the information learned by the DGM, which may not have the desired inductive bias to approximate the target black-box function in small data settings. In order to address this

challenge, we developed the LADDER approach [Deshwal and Doppa, 2021] that combines continuous space kernels over the latent space of a deep generative model (DGM) with expert-designed structured kernels to improve surrogate models. This approach has the advantage that it can leverage any advances in both DGMs and structured kernels in a plug-and-play manner. Recent work used LADDER to develop new methods for prompt optimization [Chen *et al.*, 2024].

2.2 Information-Theoretic Decision-Making Policies

Given a surrogate model, the second key component of Bayesian optimization is a decision-making policy (acquisition function) that selects the next experiment(s) to perform. In this section, we discuss a line of work that develops information-theoretic policies based on a common principle: select the experiment that maximizes the information gain per unit resource cost about the unknown quantity of interest, where the quantity of interest depends on the target goal (e.g., the optimal input for single-objective optimization, or the optimal Pareto front for multi-objective optimization) [Belakaria *et al.*, 2019; Belakaria *et al.*, 2020].

Many applications involve reasoning about multiple objectives (e.g., power and performance in hardware design) and multi-fidelity experiments that evaluate each objective via cheaper approximations. The goal is to approximate the optimal Pareto front by minimizing the resource cost of experiments. We developed a novel output space entropy (OSE) [Belakaria *et al.*, 2021] based decision-making framework, along with effective instantiations of the framework of this principle for the multi-objective [Belakaria *et al.*, 2019], multi-fidelity settings [Belakaria *et al.*, 2020]. The OSE framework was also extended to handle the challenge of black-box constraints [Belakaria *et al.*, 2021].

3 Applications in Science and Engineering

In this section, we discuss the application of adaptive experimental design methods to real-world problems with high societal impact across science and engineering.

Nanoporous materials discovery: In collaboration with chemistry experts, we applied BO to demonstrate significant successes for nanoporous materials discovery. For example, in [Deshwal *et al.*, 2021d], BO found the optimal covalent organic framework (COF) with the highest simulated methane deliverable capacity after evaluating 120 COFs (about 0.3% of the overall space of 70K materials). In follow-up work on multi-fidelity optimization, it found the optimal COF with the largest adsorptive selectivity for separating the noble gases xenon (Xe) over krypton (Kr) at room temperature using only 6.3% and 1.5% of COFs evaluated with low-fidelity and high-fidelity experiments, respectively [Gantzler *et al.*, 2023].

Electronic design automation (EDA): With the rising needs of emerging applications from edge to cloud, a key challenge in EDA is to make the development of application-specific hardware architectures easy, inexpensive, and as seamless as developing software. In collaboration with EDA experts, adaptive experimental design methods have been used successfully to design 3D heterogeneous manycore systems

for optimizing power, performance, and thermal trade-offs [Deshwal *et al.*, 2019], to create runtime resource management policies [Deshwal *et al.*, 2021a; Narang *et al.*, 2023] and for hardware and software co-design [Joardar *et al.*, 2021].

Surfactant design : In a recent work, we applied adaptive experimental design to characterize surfactants in a sample-efficient manner. Surfactants are workhorse molecules used in products such as detergents, paints, and coatings. Because most existing surfactants are petrochemical-based, environmentally friendly and biodegradable alternatives are increasingly needed. A key property of a surfactant is its critical micelle concentration (CMC), a latent parameter that is traditionally inferred from a large number of concentration-dependent lab measurements. We formulated CMC determination as active estimation of this latent parameter using an information-theoretic policy. This approach determined the CMC with roughly half as many physical wet-lab experiments as standard practice [Hammond *et al.*, 2026].

Additive manufacturing: Additive manufacturing has an exciting potential to realize new solutions for many applications including healthcare, aerospace, and civil infrastructure. In [Chen *et al.*, 2025], BO is leveraged for configuring 3D printing of multi-organ models. This approach achieved better shape precision and porosity of 3D-printed organs compared to existing approaches. Recently, we applied adaptive experimental design methods for feasibility discovery towards additive manufacturing of GRCop copper alloys, which are important for thermal and aerospace applications. The approach discovered, for the first time, feasible printing configurations for this alloy under 1000 Watts, with physical experiments in the loop [Fadhel *et al.*, 2026]. Beyond these applications, adaptive experimental design methods are also actively used in a large variety of other settings including robotics, human-computer interface design, algorithm configuration, and AutoML [Garnett, 2023].

Collaboration and Community: A key lesson we gained from our work is that close collaboration between domain scientists/engineers and AI experts is critical for achieving real-world impact. Engaging closely with domain experts serves two purposes: it allows incorporating valuable domain knowledge as an inductive bias for learning robust models, and it surfaces novel, well-motivated problems that can drive AI research [Rolnick *et al.*, 2024]. A sustained large community effort is essential to achieve this goal of bringing the two research communities together [AI2ASE Workshops, 2026].

4 Future Opportunities

There are many promising directions for developing new adaptive experimental design methods inspired by science and engineering applications. For instance, one promising direction exploits the *side-information* that experiments provide beyond the objective value. As an example, gradient information is commonly available from density functional theory based energy calculations for materials and crystal structures [Ament *et al.*, 2021]. Related problems are also studied within the broader framework of *grey-box optimization* [Astudillo and Frazier, 2021]. Another interesting research direction which has received some attention recently is the

intersection of adaptive experimentation and causal modeling [Aglietti *et al.*, 2020; Duong *et al.*, 2025].

A second frontier arises from foundation models [Bommasani *et al.*, 2021], where many training and inference problems are computationally or economically expensive and thus well suited for adaptive experimental design. These include prompt or exemplar optimization [Hu *et al.*, 2024; Luo and Deshwal, 2025] data mixture selection [Luo and Deshwal, 2025], active fine-tuning [Hübotter *et al.*, 2025] and efficient model evaluation [Berrada *et al.*, 2026].

5 Conclusion

In this paper, we present author’s and his collaborators’ recent work on developing new methods for adaptive experimental design and applying them to multiple high-impact problems in science and engineering applications. We emphasize the importance of inter-disciplinary collaboration and end by discussing opportunities for future work in this research area.

References

- [Aglietti *et al.*, 2020] Virginia Aglietti, Xiaoyu Lu, Andrei Paleyes, and Javier González. Causal bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 3155–3164. PMLR, 2020.
- [AI2ASE Workshops, 2026] AI2ASE Workshops. Annual aaai workshop on AI to accelerate science and engineering, 2026.
- [Ament *et al.*, 2021] Sebastian Ament, Maximilian Amsler, Duncan R Sutherland, Ming-Chiang Chang, Dan Guevarra, Aine B Connolly, John M Gregoire, Michael O Thompson, Carla P Gomes, and R Bruce Van Dover. Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams. *Science Advances*, 7(51):eabg4930, 2021.
- [Astudillo and Frazier, 2021] Raul Astudillo and Peter I Frazier. Thinking inside the box: A tutorial on grey-box Bayesian optimization. In *2021 Winter Simulation Conference (WSC)*, pages 1–15. IEEE, 2021.
- [Belakaria *et al.*, 2019] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective Bayesian optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- [Belakaria *et al.*, 2020] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Multi-fidelity multi-objective Bayesian optimization: An output space entropy search approach. In *Proceedings of the AAAI Conference on artificial intelligence*, volume 34, pages 10035–10043, 2020.
- [Belakaria *et al.*, 2021] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Output space entropy search framework for multi-objective Bayesian optimization. *Journal of Artificial Intelligence Research*, 72:667–715, 2021.
- [Berrada *et al.*, 2026] Gabrielle Berrada, Jannik Kossen, Freddie Bickford Smith, Muhammed Razzak, Yarin Gal, and Thomas Rainforth. Scaling up active testing to large

- language models. *Advances in Neural Information Processing Systems*, 38:173670–173697, 2026.
- [Bommasani *et al.*, 2021] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [Chen *et al.*, 2024] Lichang Chen, Jiuhai Chen, Tom Goldstein, Heng Huang, and Tianyi Zhou. InstructZero: Efficient instruction optimization for black-box large language models. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 6503–6518. PMLR, 21–27 Jul 2024.
- [Chen *et al.*, 2025] Eric S Chen, Alaleh Ahmadianshalchi, Sonja S Sparks, Chuchu Chen, Aryan Deshwal, Janardhan R Doppa, and Kaiyan Qiu. Machine learning enabled design and optimization for 3d-printing of high-fidelity presurgical organ models. *Advanced Materials Technologies*, 10(1):2400037, 2025.
- [Deshwal and Doppa, 2021] Aryan Deshwal and Jana Doppa. Combining latent space and structured kernels for Bayesian optimization over combinatorial spaces. *Advances in Neural Information Processing Systems*, 34:8185–8200, 2021.
- [Deshwal *et al.*, 2019] Aryan Deshwal, Nitthilan Kanappan Jayakodi, Biresh Kumar Joardar, Janardhan Rao Doppa, and Partha Pratim Pande. MOOS: A multi-objective design space exploration and optimization framework for NoC enabled manycore systems. *ACM Transactions on Embedded Computing Systems (TECS)*, 18(5s):1–23, 2019.
- [Deshwal *et al.*, 2021a] Aryan Deshwal, Syrine Belakaria, Ganapati Bhat, Janardhan Rao Doppa, and Partha Pratim Pande. Learning pareto-frontier resource management policies for heterogeneous SoCs: An information-theoretic approach. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pages 607–612. IEEE, 2021.
- [Deshwal *et al.*, 2021b] Aryan Deshwal, Syrine Belakaria, and Janardhan Rao Doppa. Bayesian optimization over hybrid spaces. In *International Conference on Machine Learning*, pages 2632–2643. PMLR, 2021.
- [Deshwal *et al.*, 2021c] Aryan Deshwal, Syrine Belakaria, and Janardhan Rao Doppa. Mercer features for efficient combinatorial Bayesian optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7210–7218, 2021.
- [Deshwal *et al.*, 2021d] Aryan Deshwal, Cory M Simon, and Janardhan Rao Doppa. Bayesian optimization of nanoporous materials. *Molecular Systems Design & Engineering*, 6(12):1066–1086, 2021.
- [Deshwal *et al.*, 2022] Aryan Deshwal, Syrine Belakaria, Janardhan Rao Doppa, and Dae Hyun Kim. Bayesian optimization over permutation spaces. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6515–6523, 2022.
- [Deshwal *et al.*, 2023] Aryan Deshwal, Sebastian Ament, Maximilian Balandat, Eytan Bakshy, Janardhan Rao Doppa, and David Eriksson. Bayesian optimization over high-dimensional combinatorial spaces via dictionary-based embeddings. In *International Conference on Artificial Intelligence and Statistics*, pages 7021–7039. PMLR, 2023.
- [Deshwal, 2024] Aryan Deshwal. *Adaptive Experimental Design for Optimizing Combinatorial Structures*. Washington State University, 2024.
- [Ding *et al.*, 2019] Meili Ding, Robinson W Flaig, Hai-Long Jiang, and Omar M Yaghi. Carbon capture and conversion using metal–organic frameworks and mof-based materials. *Chemical Society Reviews*, 48(10):2783–2828, 2019.
- [Duong *et al.*, 2025] Bao Duong, Sunil Gupta, and Thin Nguyen. Causal discovery via bayesian optimization. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [Fadhel *et al.*, 2026] Azza Fadhel, Nathaniel W Zuckschwerdt, Aryan Deshwal, Susmita Bose, Amit Bandyopadhyay, and Jana Doppa. Discovery of feasible 3d printing configurations for metal alloys via AI-driven adaptive experimental design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026.
- [Gantzler *et al.*, 2023] Nickolas Gantzler, Aryan Deshwal, Janardhan Rao Doppa, and Cory Simon. Multi-fidelity Bayesian Optimization of Covalent Organic Frameworks for Xenon/Krypton Separations. *Digital Discovery*, 2023.
- [Garnett, 2023] Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.
- [Hammond *et al.*, 2026] Makaila Hammond, Qia Ke, Aryan Deshwal, Cory M Simon, and Alison Bain. Adaptive, Bayesian experimental design to efficiently determine the critical micelle concentration of a surfactant. *Langmuir*, 42(15):10325–10336, 2026.
- [Hopfield and Hinton, 2024] John J Hopfield and Geoffrey E Hinton. The nobel prize in physics 2024. *Stockholm, Sweden: The Royal Swedish Academy of Sciences*, 2024.
- [Hu *et al.*, 2024] Wenyang Hu, Yao Shu, Zongmin Yu, Zhaoxuan Wu, Xiaoqiang Lin, Zhongxiang Dai, See-Kiong Ng, and Bryan Kian Hsiang Low. Localized zeroth-order prompt optimization. *Advances in Neural Information Processing Systems*, 37:86309–86345, 2024.
- [Huang *et al.*, 2021] Guyue Huang, Jingbo Hu, Yifan He, Jialong Liu, Mingyuan Ma, Zhaoyang Shen, Juejian Wu, Yuanfan Xu, Hengrui Zhang, Kai Zhong, et al. Machine learning for electronic design automation: A survey. *ACM Transactions on Design Automation of Electronic Systems*, 26(5):1–46, 2021.
- [Hübötter *et al.*, 2025] Jonas Hübötter, Sascha Bongni, Ido Hakimi, and Andreas Krause. Efficiently learning at test-time: Active fine-tuning of llms. In *International Con-*

ference on Learning Representations, volume 2025, pages 74978–75035, 2025.

- [Joardar *et al.*, 2021] Biresh Kumar Joardar, Aryan Deshwal, Janardhan Rao Doppa, Partha Pratim Pande, and Krishnendu Chakrabarty. High-throughput training of deep CNNs on ReRAM-based heterogeneous architectures via optimized normalization layers. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(5):1537–1549, 2021.
- [Jumper *et al.*, 2021] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.
- [Karpatne *et al.*, 2025] Anuj Karpatne, Aryan Deshwal, Xiaowei Jia, Wei Ding, Michael Steinbach, Aidong Zhang, and Vipin Kumar. AI-enabled scientific revolution in the age of generative ai: second nsf workshop report. *npj Artificial Intelligence*, 1(1):18, 2025.
- [Luo and Deshwal, 2025] Gaoxiang Luo and Aryan Deshwal. Com-bom: Bayesian exemplar search for efficiently exploring the accuracy-calibration pareto frontier. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20350–20363, 2025.
- [Mutnỳ, 2024] Mojmír Mutnỳ. *Modern Adaptive Experiment Design: Machine Learning Perspective*. PhD thesis, ETH Zurich, 2024.
- [Narang *et al.*, 2023] Gaurav Narang, Aryan Deshwal, Raid Ayoub, Michael Kishinevsky, Janardhan Rao Doppa, and Partha Pratim Pande. Dynamic power management in large manycore systems: A learning-to-search framework. *ACM Transactions on Design Automation of Electronic Systems*, 28(5):1–21, 2023.
- [Rolnick *et al.*, 2024] David Rolnick, Alan Aspuru-Guzik, Sara Beery, Bistra Dilkina, Priya L Donti, Marzyeh Ghassemi, Hannah Kerner, Claire Monteleoni, Esther Rolf, Milind Tambe, et al. Position: Application-driven innovation in machine learning. In *Forty-first International Conference on Machine Learning*, 2024.
- [Schneider *et al.*, 2020] Petra Schneider, W Patrick Walters, Alleyn T Plowright, Norman Sieroka, Jennifer Listgarten, Robert A Goodnow Jr, Jasmin Fisher, Johanna M Jansen, José S Duca, Thomas S Rush, et al. Rethinking drug design in the artificial intelligence era. *Nature Reviews Drug Discovery*, 19(5):353–364, 2020.
- [Shahriari *et al.*, 2015] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- [Terayama *et al.*, 2021] Kei Terayama, Masato Sumita, Ryo Tamura, and Koji Tsuda. Black-box optimization for automated discovery. *Accounts of Chemical Research*, 54(6):1334–1346, 2021.