Constraint active search for experimental design

Bolong Cheng
SigOpt, an Intel company
San Francisco, CA, USA
harvey@sigopt.com

Gustavo Malkomes
SigOpt, an Intel company
San Francisco, CA, USA
gustavo@sigopt.com

Michael McCourt
SigOpt, an Intel company
San Francisco, CA, USA
mccourt@sigopt.com

Abstract

Many problems in engineering and design require balancing competing objectives under the presence of uncertainty. The standard approach in the literature characterizes the relationship between design decisions and their corresponding outcomes as a Pareto frontier, which is discovered through multiobjective optimization. In this position paper, we suggest that this approach is not ideal for reasoning about practical design decisions. Instead of multiobjective optimization, we propose soliciting desired minimum performance constraints on all objectives to define regions of satisfactory. We present work-in-progress which visualizes the design decisions that consistently satisfy user-defined thresholds in an additive manufacturing problem.

1 Introduction

Design problems are ubiquitous in science and engineering. These design problems are inherently decision making tasks that require balancing complex choices under competing metrics to satisfy real-world constraints. Using numerical simulation to study the impact of design decisions prior to manufacturing has become a common strategy to reduce the number of fabrications required (and the cost to find an effective design). This may be called sim-to-real in some communities such as robotics [30]. Optimization has been a powerful ally for solving design problems [9, 24]. Scientists and engineers often pose design problems as optimization of an objective function $f\colon \mathcal{X} \to \mathbb{R}$ that codifies their preference over choices within the design space $\mathcal{X}$ of possibilities. In most cases, there are multiple competing objectives that need to be investigated [19, 29].

Traditional optimization methods typically rely on structural assumptions about these objective functions; for example, one might require that f is differentiable, linear, convex, deterministic or very cheap to evaluate. In many real-world scenarios, however, f will be nonlinear, noisy and/or with unknown function form. In these scenarios, the Bayesian optimization (BO) [10, 20, 13] (and more recently, multiobjective Bayesian optimization [1, 28, 16]) framework has become a popular approach to address these problems.

Bayesian optimization is especially popular (in contrast to, for example, genetic algorithms [15, 8]) when evaluating very expensive objectives. Despite the popularity of these techniques, they fail to address the crux of design problems: understanding the impact of choices (independent variable) to outcomes (dependent variable). While finding a high performing parameter, or approximate Pareto frontier (in a multiobjective setting), produces a good design choice, the sample efficiency limits any understanding of the impact of varying designs on the resulting performance.

Our motivation comes from materials science, where the materials design process can be accelerated by numerical simulation prior to eventual fabrication through additive manufacturing. As can be seen in Figure 1, a CAD model, on which numerical simulations would be conducted, will encounter some amount of discrepancy during actual fabrication. This will be especially pronounced when

34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada.

approaching the fabrication limits of the additive manufacturing device (in this case, the Nanoscribe Photonic Professional GT 3D printer [22]).

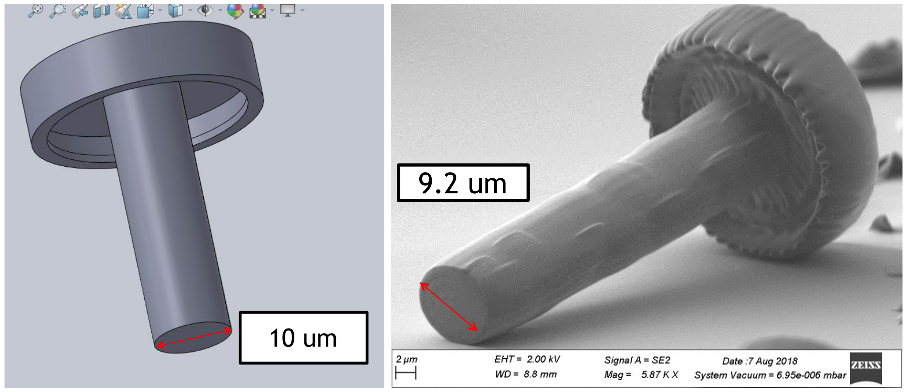

Figure 1: *left*: A CAD model of a proposed pin. *right*: The resulting fabrication of this pin; the extremely small size of this pin causes imprecision on a normally small scale to be more noticeable.

Of course, the purpose of BO algorithms is not to provide understanding on the design/objective relationship; it is to balance exploration and exploitation of the possible design options to find high performing outcomes. In using these, however, we fail to account for a key element in the design process: that our metrics from the numerical simulations are not exactly equal to what will be seen in production.

In this setting, the discrepancy between simulated and real objective values can be significant. Multiobjective Bayesian optimization does not account for this variation in design space, making it unsuitable for many real-world design tasks. We propose an alternative to Bayesian multiobjective optimization based on minimum performance constraints, with the goal of better preparing for the additive manufacturing process. To do so, we suggest that the balance of exploration and exploitation should be adjusted to provide design decisions which are more likely to be reliable in production.

2 Additive manufacturing example

In [14], the authors studied a numerical simulation for helping design an additive manufacturing strategy for minimizing the reflection of light at multiple incidence angles. A search was conducted for different nanostructures which balance a desire for minimizing normal reflection with reflection at an oblique angle:

$$\min_{\mathbf{x}\in\mathcal{X}} R_{\text{normal}}(\mathbf{x}),\ \min_{\mathbf{x}\in\mathcal{X}} R_{\text{oblique}}(\mathbf{x}).$$

The results of the multiobjective optimizations were approximate Pareto frontiers containing the efficient designs, shown in Figure 3.

The simulation, conducted using the Lumerical software [21], has a level of inaccuracy associated with the numerical methods. In addition, there is a level of imprecision associated with the actual manufacturing process which means that the desired design parameters are only realized to limited precision during the manufacturing process. The scanning electron microscope image in Figure 2 shows the fabrication result of one of the efficient designs.

For [14], the authors studied the designs of the Pareto efficient nanostructures and showed that physical properties predicted by theory were confirmed to be efficient in the numerical simulation. But the physical limitations of the manufacturing process complicate the use of these results (as seen in Figure 2).

- The "true" Pareto frontier for these two objectives is a 1D curve in parameter space (3 or 4 dimensions for these fabrication problems).
  - The true Pareto frontier cannot be learned, both because the numerical simulation has limited accuracy and because the Pareto frontier is an uncountable set of points. The multiobjective optimization process only estimates it.

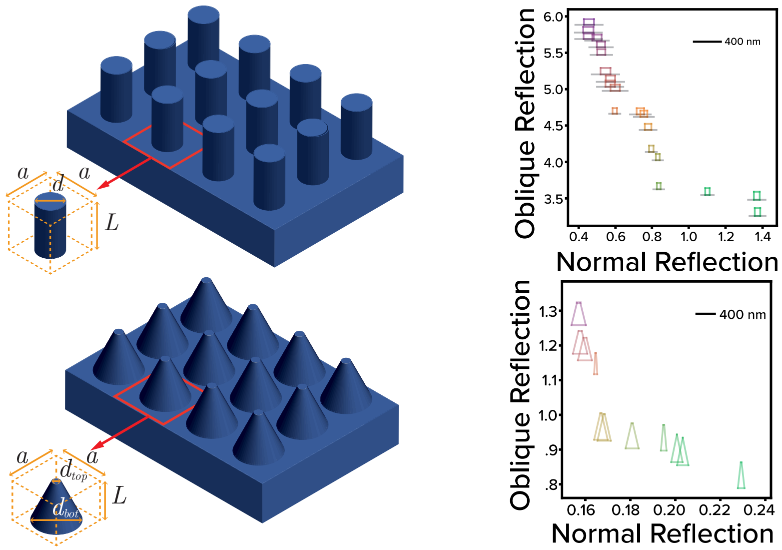

Figure 2: Schematic drawings (*left*) of the nanowire (*top*) and nanocone (*bottom*) design spaces and associated Pareto frontiers (*right*), as estimated through numerical simulation (condensed presentation from [14]).

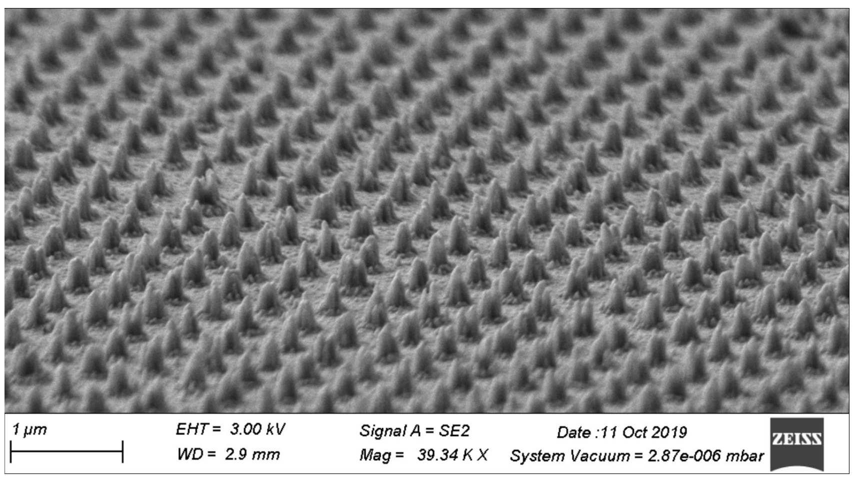

Figure 3: SEM image of nanocones built on top of a piece of bare glass using additive manufacturing. The scale of the manufacturing is approaching the limits of the fabrication device, which prevents the structures from meeting the exact specifications used in the simulation.

- During manufacturing, the outcome will not exactly match the desired design.
  - Therefore, even if we knew the true Pareto frontier, the manufactured result would not be Pareto efficient.

3 Alternative to the Pareto frontier

The problem with the Pareto frontier is that deviation in parameter space generally yields uncertain results in objective space. During a multiobjective optimization, effort is focused entirely on trying to expand the Pareto frontier (in objective space) without much consideration of how eventual inaccuracies in the design implementation will affect the performance.

This problem is amplified in a multiobjective Bayesian optimization setting, where the expensive cost in the simulation (or, more generally, function evaluation) demands sample efficiency in the search for the Pareto frontier. When a genetic algorithm is used for multiobjective optimization, there may be a very large number of designs simulated, leaving many results to help choose viable outcomes for fabrication. For a multiobjective Bayesian optimization, many fewer designs are simulated, and much less is known about how imprecision in manufacturing will affect the finished product.

We propose constraint search as an alternative to the multiobjective optimization for finding effective designs. In this formulation, we solicit constraints τ_j for each metric of interest f_j: these constraints serve as minimum performance thresholds that the eventual fabrication must meet to be acceptable. Given these constraints, our goal for this reflection minimization problem is to find the regions in parameter space which satisfy these constraints:

$$\text{Find} : \{\mathbf{x} : R_{\text{normal}}(\mathbf{x}) \leq \tau_{\text{normal}},\ R_{\text{oblique}}(\mathbf{x}) \leq \tau_{\text{oblique}}\}.$$

This is written with upper performance thresholds (as a minimization problem), but, in general, each objective could have either lower or upper bounds without loss of generality.

To demonstrate the potential of this alternate consideration, we reconsider the fabrication problems in [14]; to avoid the cost of the simulation, we built surrogate models from the 500 results used to create the nanowire and nanocone results in that article. With these surrogates, we estimated the normal and oblique reflection values from 10000 uniformly random design configurations. The reflection values are shown in Figures 4 and 5 (left), and the respective design space on the right.

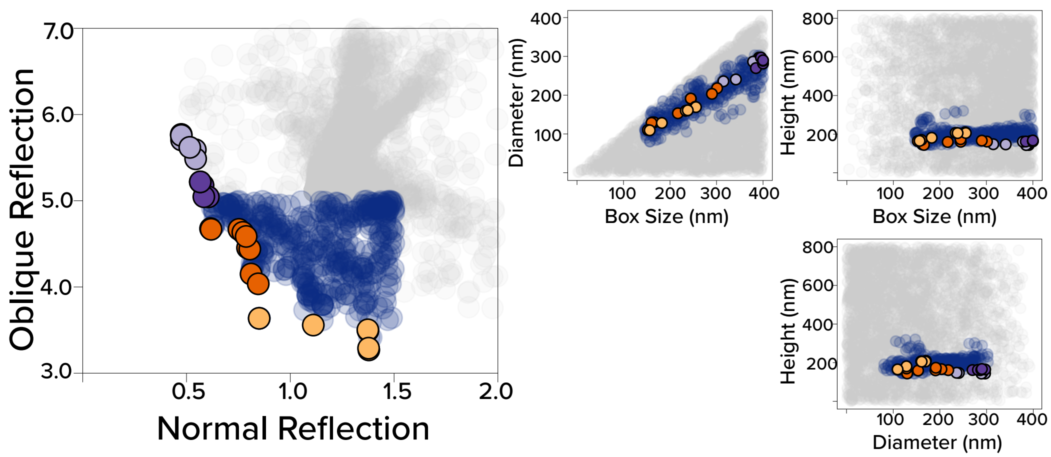

Figure 4: Nanowire results from uniform sampling of the surrogate with $\tau_{\text{normal}} = 1.5$ and $\tau_{\text{oblique}} = 5$. The initial 500 points were also included to better estimate the Pareto frontier. *left*: Plots of objective values. *right*: Design space of three parameters (box size, diameter and height). Pareto frontier points are highlighted in colors. Feasible (blue) and infeasible (light gray) points also shown. The domains are triangular where physical limitations prevent certain configurations

In these examples, the Pareto frontier (colored points) corresponds to a scattered set of points throughout the parameter space. It is unclear that a particular choice of normal reflection vs oblique reflection (for example the yellow points) would lead to a particular design configuration — notice how the box size varies considerable. In contrast, the parameter region associated with the constraint satisfied results produces a more dense set of results. We find this much more actionable when dealing with uncertainty in the final fabrication process – the region gives a better sense of which designs could be used to achieve the desired performance thresholds when the fabrication process imparts some error. It also naturally captures the imprecision of the objective functions in the transition from simulation to reality.

4 Constraint active search, related work and open challenges

We prefer this constraint search formulation, in lieu of the standard multiobjective formulation, for the following reasons.

- The explore-exploit balance which powers an efficient constraint search would find more actionable information for the actual manufacturing (rather than simply trying to increase

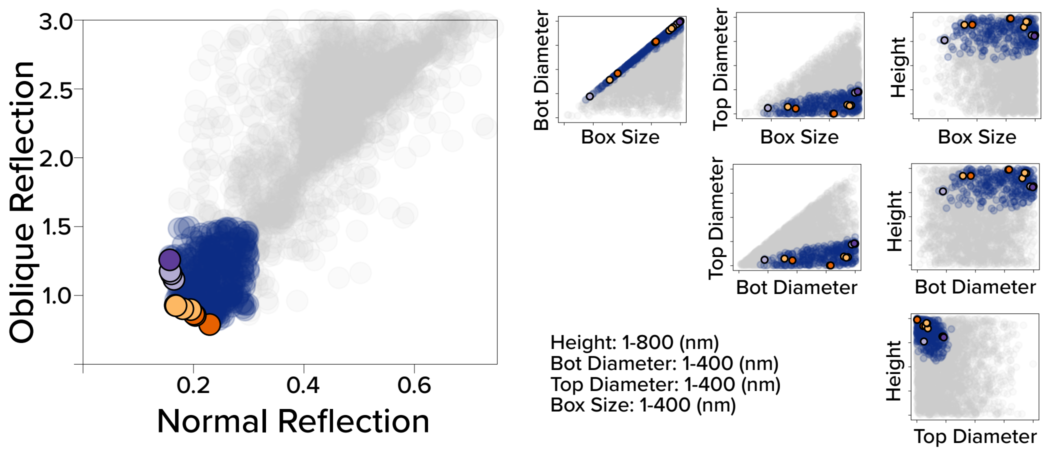

Figure 5: Nanocone results from similar setup of Figure 4. $\tau_{\text{normal}} = 0.5$ and $\tau_{\text{oblique}} = 1.5$. *left*: Plots of objective values. *right*: Design space of four parameters (box size, bottom diameter, top diameter, and height). Pareto frontier points are highlighted in colors. Feasible (blue) and infeasible (light gray) points also shown.

the hypervolume of the feasible region in metric space, the goal of any multiobjective optimization algorithm).

- The use of constraints on objectives as performance thresholds naturally fits many industrial conditions. These thresholds are also naturally interpretable, even by people not involved in the design process.
- No explicitly defined noise distribution or knowledge of a design-objective relationship must be known.
- Conceivably, an efficient constraint would run independently of the number of objectives, unlike a multiobjective optimization which becomes more difficult as the number of objectives grows.
  - The logic we are using is also applicable in a setting with a single objective.

The demonstration in Section 3 required 10000 design evaluations, which is impractical when using the actual simulation (rather than the surrogate). To move this from a constraint search to a constraint active search, where we intelligently choose each design to test so as to be sample efficient, we look to the broader framework of decision theory [23, 7]. A typical approach in decision theory is to quantify the quality of each design by a utility function; the design with highest utility is the next to be tested. Designing such a utility function, which appropriately values our desire to effectively explore the constraint-satisfying region of parameter space, remains an open problem.

To start to address this problem, we note that our proposed constraint search formulation shares elements which appear in strategies across numerous fields. The most immediate connection is to the multiobjective optimization field, as we described earlier. The difference between these is that multiobjective optimization seeks to maximize the hypervolume in metric space, whereas constraint search seeks to maximize the hypervolume of constraint-satisfying results in parameter space.

The topic of robust optimization deals with preventing undesirable outcomes. Deterministic robust optimization avoids undesirable outcomes through a stability radius in design space [27, 5, 6], which is generalized in probabilistic robust optimization through an uncertainty set in design space [4, 2, 3, 25]. Instead of preventing undesirable outcomes, constraint search seeks to efficiently map out the space of desirable ones, implicitly defined through thresholds on the objective function. This makes constraint search a fundamentally different problem, more suited to revealing the structure and relationships of parameter space to objective values.

Finally, our formulation is closely related to the work the field of active search [12, 11, 17]. In active search, the goal is to select as many relevant, rare elements from a finite design space as quickly as possible, e.g., chemical configurations in drug discovery, [26, 18]. This matches our goal here (which is why we use the term constraint search), where the combination of constraints on all the metrics have changed the problem into a search for positive outcomes. Unlike active search, constraint search will likely take place in a continuously parametrized design space. Active search algorithms typically require one positive outcome for initialization, which would be potentially extremely expensive to produce in our setting. Furthermore, our goal is to explore the constraint-satisfying parameter space, not simply exploit known satisfying results to maximize the number of positive outcomes found.

Each of these fields provides ideas and methodologies which we think will help constraint active search become a viable strategy for helping simulation improve efficiency in additive manufacturing, as well as other applications. Significant work must still be done to effectively visualize the results of such as search; those visualizations are fundamental to choosing designs which will undergo fabrication, but are increasingly complicated as the design space grows in complexity and contains both continuous and discrete parameters.

Acknowledgments and Disclosure of Funding

Special thanks to Sajad Haghanifar for providing the scanning electron microscope images seen in this article. Also, special thanks to Eric Lee for discussions during the writing of this article.

References

[1] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 7825–7835, 2019.

[2] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton University Press, 2009.

[3] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.

[4] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.

[5] Hans-Georg Beyer and Bernhard Sendhoff. Robust optimization–a comprehensive survey. *Computer methods in applied mechanics and engineering*, 196(33-34):3190–3218, 2007.

[6] Ilija Bogunovic, Jonathan Scarlett, Stefanie Jegelka, and Volkan Cevher. Adversarially robust optimization with gaussian processes. In *Advances in neural information processing systems*, pages 5760–5770, 2018.

[7] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statist. Sci.*, 10(3):273–304, 08 1995.

[8] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

[9] Alexander Forrester, Andras Sobester, and Andy Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, 2008.

[10] Peter I. Frazier. Bayesian optimization. In Esma Gel and Lewis Ntaimo, editors, *Recent Advances in Optimization and Modeling of Contemporary Problems*, pages 255–278. INFORMS, 2018.

[11] Roman Garnett, Thomas Gärtner, Martin Vogt, and Jürgen Bajorath. Introducing the 'active search' method for iterative virtual screening. *Journal of Computer-Aided Molecular Design*, 29(4):305–314, April 2015.

[12] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian optimal active search and surveying. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*, ICML'12, page 843–850, Madison, WI, USA, 2012. Omnipress.

[13] Aldair E. Gongora, Bowen Xu, Wyatt Perry, Chika Okoye, Patrick Riley, Kristofer G. Reyes, Elise F. Morgan, and Keith A. Brown. A bayesian experimental autonomous researcher for mechanical design. *Science Advances*, 6(15), 2020.

[14] Sajad Haghanifar, Michael McCourt, Bolong Cheng, Jeffrey Wuenschell, Paul Ohodnicki, and Paul W Leu. Discovering high-performance broadband and broad angle antireflection surfaces by machine learning. *Optica*, 7(7):784–789, 2020.

[15] Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). *Evolutionary computation*, 11(1):1–18, 2003.

[16] Florian Häse, Loïc M. Roch, and Alán Aspuru-Guzik. Chimera: enabling hierarchy based multi-objective optimization for self-driving laboratories. *Chem. Sci.*, 9:7642–7655, 2018.

[17] Shali Jiang, Gustavo Malkomes, Matthew Abbott, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic batch active search. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 1099–1109. Curran Associates, Inc., 2018.

[18] Shali Jiang, Gustavo Malkomes, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic active search with applications in drug and materials discovery. In *NeurIPS 2018 Workshop on Machine Learning for Molecules and Materials*, 11 2018.

[19] Slawomir Koziel, Adrian Bekasiewicz, Ivo Couckuyt, and Tom Dhaene. Efficient multi-objective simulation-driven antenna design using co-kriging. *IEEE Transactions on Antennas and Propagation*, 62(11):5900–5905, 2014.

[20] Benjamin Letham, Brian Karrer, Guilherme Ottoni, Eytan Bakshy, et al. Constrained bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2):495–519, 2019.

[21] Lumerical device multiphysics simulation suite, FDTD solver, 2019. Lumerical Inc.

[22] Nanoscribe photonoic professional GT2, 2020. Nanoscribe GmbH.

[23] David JC MacKay and David JC Mac Kay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.

[24] Sean Molesky, Zin Lin, Alexander Y Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12(11):659–670, 2018.

[25] José Nogueira, Ruben Martinez-Cantin, Alexandre Bernardino, and Lorenzo Jamone. Unscented bayesian optimization for safe robot grasping. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1967–1972. IEEE, 2016.

[26] Dino Oglic, Steven A. Oatley, Simon J. F. Macdonald, Thomas Mcinally, Roman Garnett, Jonathan D. Hirst, and Thomas Gärtner. Active search for computer-aided drug design. *Molecular Informatics*, 37(1-2):1700130, 2018.

[27] Luc Pronzato and Éric Walter. Robust experiment design via maximin optimization. *Mathematical Biosciences*, 89(2):161–176, 1988.

[28] Koji Shimoyama, Shinkyu Jeong, and Shigeru Obayashi. Kriging-surrogate-based optimization considering expected hypervolume improvement in non-constrained many-objective test problems. In *2013 IEEE Congress on Evolutionary Computation*, pages 658–665. IEEE, 2013.

[29] Prashant Singh, Ivo Couckuyt, Khairy Elsayed, Dirk Deschrijver, and Tom Dhaene. Shape optimization of a cyclone separator using multi-objective surrogate-based optimization. *Applied Mathematical Modelling*, 40(5-6):4248–4259, 2016.

[30] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30. IEEE, 2017.