
Max-value Entropy Search for Multi-Objective Bayesian Optimization with Constraints

Syrine Belakaria, Aryan Deshwal, Janardhan Rao Doppa

School of EECS, Washington State University

{syrine.belakaria, aryan.deshwal, jana.doppa}@wsu.edu

Abstract

We consider the problem of constrained multi-objective blackbox optimization using expensive function evaluations, where the goal is to approximate the true Pareto set of solutions satisfying a set of constraints while minimizing the number of function evaluations. For example, in aviation power system design applications, we need to find the designs that trade-off total energy and the mass while satisfying specific thresholds for motor temperature and voltage of cells. This optimization requires performing expensive computational simulations to evaluate designs. In this paper, we propose a new approach referred to as *Max-value Entropy Search for Multi-objective Optimization with Constraints (MESMOC)* to solve this problem. MESMOC employs an output-space entropy based acquisition function to efficiently select the sequence of inputs for evaluation to uncover high-quality pareto-set solutions while satisfying constraints. We apply MESMOC to two real-world engineering design applications to demonstrate its effectiveness.

1 Introduction

Many engineering and scientific applications involve making design choices to optimize multiple objectives. Some examples include tuning the knobs of a compiler to optimize performance and efficiency of a set of software programs; and designing new materials to optimize strength, elasticity, and durability. There are three common challenges in solving this kind of multi-objective optimization (MO) problems: **1)** The objective functions are unknown and we need to perform expensive experiments to evaluate each candidate design choice. For example, performing computational simulations and physical lab experiments for compiler optimization and material design applications respectively. **2)** The objectives are conflicting in nature and all of them cannot be optimized simultaneously. **3)** The problem involves several black-box constraints that need to be satisfied. Therefore, we need to find the *Pareto optimal* set of solutions satisfying the constraints. A solution is called Pareto optimal if it cannot be improved in any of the objectives without compromising some other objective. The overall goal is to approximate the true Pareto set satisfying the constraints while minimizing the number of function evaluations.

Bayesian Optimization (BO) ([20]) is an effective framework to solve blackbox optimization problems with expensive function evaluations. The key idea behind BO is to build a cheap surrogate model (e.g., Gaussian Process ([22]) using the real experimental evaluations; and employ it to intelligently select the sequence of function evaluations using an acquisition function, e.g., expected improvement (EI). There is a large body of literature on single-objective BO algorithms ([20, 1, 9, 10]) and their applications including hyper-parameter tuning of machine learning methods ([21, 17]). However, there is relatively less work on the more challenging problem of BO for multiple objectives ([16, 11, 14, 2, 3]) and very limited prior work to address constrained MO problems ([13, 12]). PSMOC ([13]) is the current state-of-the-art method in this problem setting. PSMOC is an information-theoretic approach that relies on the principle of input space entropy search. However, it is computationally

expensive to optimize the acquisition function behind PESMOC. A series of approximations are performed to improve the efficiency potentially at the expense of accuracy.

In this paper, we propose a new and principled approach referred to as **Max-value Entropy Search for Multi-objective Optimization with Constraints** (MESMOC). MESMOC employs an *output space entropy* based acquisition function to select the candidate inputs for evaluation. The key idea is to evaluate the input that maximizes the information gain about the optimal Pareto front in each iteration while satisfying the constraints. Output space entropy search has many advantages over algorithms based on input space entropy search ([2]). MESMOC is an extension of the MESMO algorithm [2] based on output space information gain, which was shown to be efficient and robust, to the challenging constrained MO setting.

2 MESMOC for Multi-Objective Optimization with Constraints

In this section, we explain the technical details of our proposed algorithm. We describe mathematically the MESMOC acquisition function and provide an algorithmic approach to efficiently compute it.

Multi-Objective Optimization Problem with Constraints. Without loss of generality, our goal is to minimize real-valued objective functions $f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_K(\mathbf{x})$, with $K \geq 2$, while satisfying L black-box constraints of the form $C_1(x) \geq 0, C_2(x) \geq 0, \dots, C_L(x) \geq 0$ over continuous space $\mathfrak{X} \subseteq \mathbb{R}^d$. Each evaluation of an input $\mathbf{x} \in \mathfrak{X}$ produces a vector of objective values and constraint values $\mathbf{y} = (y_{f_1}, y_{f_2}, \dots, y_{f_K}, y_{c_1}, \dots, y_{c_L})$ where $y_{f_j} = f_j(x)$ for all $j \in \{1, 2, \dots, K\}$ and $y_{c_i} = C_i(x)$ for all $i \in \{1, 2, \dots, L\}$. We say that a valid point \mathbf{x} (satisfies all constraints) *Pareto-dominates* another point \mathbf{x}' if $f_j(\mathbf{x}) \leq f_j(\mathbf{x}') \forall j$ and there exists some $j \in \{1, 2, \dots, K\}$ such that $f_j(\mathbf{x}) < f_j(\mathbf{x}')$. The optimal solution of MOO problem with constraints is a set of points $\mathcal{X}^* \subset \mathfrak{X}$ such that no point $\mathbf{x}' \in \mathfrak{X} \setminus \mathcal{X}^*$ Pareto-dominates a point $\mathbf{x} \in \mathcal{X}^*$ and all points in \mathcal{X}^* satisfy the problem constraints. The solution set \mathcal{X}^* is called the optimal *Pareto set* and the corresponding set of function values \mathcal{Y}^* is called the optimal *Pareto front*. Our goal is to approximate \mathcal{X}^* by minimizing the number of function evaluations.

Surrogate models. We model the objective functions and blackbox constraints by independent Gaussian processes models $\mathcal{M}_{f_1}, \mathcal{M}_{f_2}, \dots, \mathcal{M}_{f_K}$ and $\mathcal{M}_{c_1}, \mathcal{M}_{c_2}, \dots, \mathcal{M}_{c_L}$ with zero mean and i.i.d. observation noise. Let $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{t-1}$ be the training data from past $t-1$ function evaluations, where $\mathbf{x}_i \in \mathfrak{X}$ is an input and $\mathbf{y}_i = \{y_{f_1}^i, \dots, y_{f_K}^i, y_{c_1}^i, \dots, y_{c_L}^i\}$ is the output vector resulting from evaluating the functions and constraints at \mathbf{x}_i . We learn surrogate models from \mathcal{D} .

Output space entropy based acquisition function. Input space entropy based methods such as PESMO [14] select the next candidate input by maximizing the information gain about the Pareto set \mathcal{X}^* . Information gain is defined as the expected reduction in entropy of the posterior distribution over the optimal Pareto set. This formulation has several limitations: 1) relies on a very expensive and high-dimensional distribution 2) requires a series of approximations [14] which can be potentially sub-optimal; and 3) performance is strongly dependent on the number of Monte-Carlo samples.

To overcome the above challenges of computing input space entropy based acquisition function, [2] proposed to maximize the information gain about the optimal **Pareto front** \mathcal{Y}^* . However, MESMO did not address the challenge of constrained Pareto front. We propose an extension of MESMO's acquisition function to maximize the information gain between the next candidate input for evaluation \mathbf{x} and constrained Pareto front \mathcal{Y}^* given as:

$$\alpha(\mathbf{x}) = I(\{\mathbf{x}, \mathbf{y}\}, \mathcal{Y}^* | D) = H(\mathbf{y} | D, \mathbf{x}) - \mathbb{E}_{\mathcal{Y}^*}[H(\mathbf{y} | D, \mathbf{x}, \mathcal{Y}^*)] \quad (1)$$

Due to space constraint, we provide the full derivation of the acquisition function in the appendix. The final form of our acquisition function is shown below:

$$\alpha(\mathbf{x}) \simeq \frac{1}{S} \sum_{s=1}^S \left[\sum_{j=1}^K \frac{\gamma_s^{f_j}(\mathbf{x}) \phi(\gamma_s^{f_j}(\mathbf{x}))}{2\Phi(\gamma_s^{f_j}(\mathbf{x}))} - \ln \Phi(\gamma_s^{f_j}(\mathbf{x})) + \sum_{i=1}^L \frac{\gamma_s^{c_i}(\mathbf{x}) \phi(\gamma_s^{c_i}(\mathbf{x}))}{2\Phi(\gamma_s^{c_i}(\mathbf{x}))} - \ln \Phi(\gamma_s^{c_i}(\mathbf{x})) \right] \quad (2)$$

where $\gamma_s^{c_i}(x) = \frac{y_s^{c_i^*} - \mu_{c_i}(\mathbf{x})}{\sigma_{c_i}(\mathbf{x})}$, $\gamma_s^{f_j}(x) = \frac{y_s^{f_j^*} - \mu_{f_j}(\mathbf{x})}{\sigma_{f_j}(\mathbf{x})}$, $y_s^{c_i^*}$ and $y_s^{f_j^*}$ are the maximum values of constraint \tilde{c}_i and function \tilde{f}_j reached after the cheap multi-objective optimization over sampled functions and constraints. ϕ and Φ are the p.d.f and c.d.f of a standard normal distribution respectively.

A complete description of the MESMOC algorithm is given in Algorithm 1. The blue colored steps correspond to computation of our output space entropy based acquisition function via sampling.

Algorithm 1 MESMOC Algorithm

Input: input space \mathcal{X} ; K blackbox functions $f_1(x), f_2(x), \dots, f_K(x)$; L blackbox constraints $C_1(x), C_2(x), \dots, C_L(x)$; and maximum no. of iterations T_{max}

- 1: Initialize Gaussian process models $\mathcal{M}_{f_1}, \mathcal{M}_{f_2}, \dots, \mathcal{M}_{f_K}$ and $\mathcal{M}_{c_1}, \mathcal{M}_{c_2}, \dots, \mathcal{M}_{c_L}$ by evaluating at N_0 initial points
- 2: **for** each iteration $t = N_0 + 1$ to T_{max} **do**
- 3: Select $\mathbf{x}_t \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha_t(\mathbf{x})$
 s.t $(\mu_{c_1} \geq 0, \dots, \mu_{c_L} \geq 0)$
- 4: $\alpha_t(\cdot)$ is computed as:
- 5: **for** each sample $s \in 1, \dots, S$:
- 6: Sample $\tilde{f}_j \sim \mathcal{M}_{f_j}, \quad \forall j \in \{1, \dots, K\}$
- 7: Sample $\tilde{C}_i \sim \mathcal{M}_{c_i}, \quad \forall i \in \{1, \dots, L\}$
- 8: // Solve *cheap* MOO over $(\tilde{f}_1, \dots, \tilde{f}_K)$ constrained by $(\tilde{C}_1, \dots, \tilde{C}_L)$
- 9: $\mathcal{Y}_s^* \leftarrow \arg \max_{x \in \mathcal{X}} (f_1, \dots, f_K)$
 s.t $(\tilde{C}_1 \geq 0, \dots, \tilde{C}_L \geq 0)$
- 10: Compute $\alpha_t(\cdot)$ based on the S samples of \mathcal{Y}_s^* as given in Equation 2
- 11: Evaluate \mathbf{x}_t ; $\mathbf{y}_t \leftarrow (f_1(\mathbf{x}_t), \dots, f_K(\mathbf{x}_t), C_1(\mathbf{x}_t), \dots, C_L(\mathbf{x}_t))$
- 12: Aggregate data: $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\mathbf{x}_t, \mathbf{y}_t)\}$
- 13: Update models $\mathcal{M}_{f_1}, \mathcal{M}_{f_2}, \dots, \mathcal{M}_{f_K}$ and $\mathcal{M}_{c_1}, \mathcal{M}_{c_2}, \dots, \mathcal{M}_{c_L}$
- 14: $t \leftarrow t + 1$
- 15: **end for**
- 16: **return** Pareto front of $f_1(x), f_2(x), \dots, f_K(x)$ based on \mathcal{D}

3 Experiments and Results

In this section, we describe our experimental evaluation of MESMOC on two real-world engineering applications, namely, electrified aviation power system design and analog circuit design.

3.1 Experimental Setup

We compare MESMOC with PESMOC ([13]). Due to lack of BO approaches for constrained MO, we compare to known genetic algorithms (NSGA-II and MOEAD). However, they require a large number of function evaluations to converge which is not practical for optimization of expensive functions. We use a GP based statistical model with squared exponential kernel. The hyperparameters are estimated after every 5 function evaluations. We initialize the GP models for all functions by sampling initial points at random. The code is available at (github.com/belakaria/MESMOC)

Electrified aviation power system design. We consider optimizing the design of electrified aviation power system of unmanned aerial vehicle (UAV) via a time-based static simulation. The UAV system architecture consists of a central Li-ion battery pack, hex-bridge DC-AC inverters, PMSM motors, and necessary wiring ([4]). Each candidate input consists of a set of 5 ($d=5$) variable design parameters such as the battery pack configuration (battery cells in series, battery cells in parallel) and motor size (number of motors, motor stator winding length, motor stator winding turns). We minimize two objective functions: mass and total energy. This problem has 5 black-box constraints:

- C_0 : Maximum final depth of discharge $\leq 75\%$
- C_1 : Minimum cell voltage $\geq 3V$
- C_2 : Maximum motor temperature $\leq 125^\circ C$
- C_3 : Maximum inverter temperature $\leq 120^\circ C$
- C_5 : Maximum modulation index ≤ 1.3

For a design to be valid, the simulated UAV must be capable of completing the specified mission without violating any of the constraints. The overall design space has a total of 250,000 possible candidate designs. Out of the entire design space, only 9% of the designs are valid (i.e., satisfy all the constraints), which makes it a very challenging task and only five points are in the Pareto front.

Analog circuit optimization domain. We consider optimizing the design of a multi-output switched-capacitor voltage regulator via Cadence circuit simulator that imitates the real hardware [5]. This circuit relies on a dynamic frequency switching clock. Each candidate circuit design is defined by 33 input variables ($d=33$). The first 24 variables are the width, length, and unit of the eight capacitors of the circuit $W_i, L_i, M_i \forall i \in 1 \dots 8$. The remaining input variables are four output voltage references $V_{ref_i} \forall i \in 1 \dots 4$ and four resistances $R_i \forall i \in 1 \dots 4$ and a switching frequency f . We optimize nine objectives: maximize efficiency Eff , maximize four output voltages $V_{o_1} \dots V_{o_4}$, and minimize four output ripples $OR_1 \dots OR_4$. Our problem has a total of nine constraints. OR_{lb} and OR_{ub} are the predefined lower-bound and upper-bound of OR_i respectively. $C_{p_{total}}$ is the total capacitance.

$$C_0 : C_{p_{total}} \simeq 20nF \text{ with } C_{p_{total}} = \sum_{i=1}^8 (1.955W_iL_i + 0.54(W_i + L_i))M_i$$

$$C_1 \text{ to } C_4 : V_{o_i} \geq V_{ref_i} \forall i \in 1 \dots 4$$

$$C_5 \text{ to } C_8 : OR_{lb} \leq OR_i \leq OR_{ub} \forall i \in 1 \dots 4$$

$$C_9 : Eff \leq 100\%$$

3.2 Results and Discussion

We evaluate the performance of our algorithm and the baselines using the Pareto hypervolume (PHV) metric. PHV is a commonly employed metric to measure the quality of a given Pareto front [23]. Figure 1 shows that MESMOC outperforms existing baselines. It recovers a better Pareto front with a significant gain in the number of function evaluations. Both of these experiments are motivated by real-world engineering applications where further analysis of the designs in the Pareto front is crucial.

Electrified aviation power system design. In this setting, the input space is discrete with 250,000 combinations of designs. Out of the entire design space, only 9% of designs passed all the constraints and only five points are in the Pareto front. From a domain expert perspective, satisfying all the constraints is critical. Hence, the results reported for the hypervolume include only points that satisfy all the constraints. Despite the hardness of the problem, 50% of the designs selected by MESMOC satisfy all the constraints while for PESMOC, MOEAD, and NSGA-II, this was 1.5%, 9.5%, and 7.5% respectively. MESMOC was not able to recover all the five points of the Pareto front. However, it was able to reach closely the true Pareto front and recover better designs than the baselines.

Analog circuit design optimization. In this setting, the input space is continuous, consequently there is an infinite number of candidate designs. From a domain expert perspective, satisfying all the constraints is not critical and is impossible to achieve. The main goal is to satisfy most of the constraints (and getting close to satisfying the threshold for violated constraints) while reaching the best possible objective values. Therefore, the results reported for the hypervolume include all the evaluated points. In this experiment, the efficiency of circuit is the most important objective function. The table in Figure 2 in the appendix shows the optimized circuit parameters from different algorithms. All algorithms can generate design parameters for the circuit that meets the voltage




Figure 1: Results of different constrained multi-objective algorithms including MESMOC. The hypervolume metric is shown as a function of the number function evaluations.

reference requirements. The optimized circuit using MESMOC can achieve the highest conversion efficiency of 88.81% (12.61% improvement when compared with PESMOC and 17.86% improvement when compared with NSGA-II) with similar output ripples. The circuit with optimized parameters can generate the target output voltages within the range of 0.52V to 0.76V (1/3x ratio) and 0.99V to 1.17V (2/3x ratio) under the loads varying from 14 Ohms to 1697 Ohms.

Broader Impact

Our work contributes to general algorithmic approaches for solving optimization problems with multiple black-box objectives. The biggest impact of our algorithm will be on the scientific and engineering community, where these optimization algorithms can be employed to accelerate scientific discovery and minimize the resource cost to perform experiments. Most of the current applications of Bayesian optimization techniques are in hyper-parameter tuning of machine learning models and accelerating scientific discoveries in various fields such as chemistry, materials design, and physics.

References

- [1] Ricardo Baptista and Matthias Poloczek. Bayesian optimization of combinatorial structures. In *Proceedings of the 35th International Conference on Machine Learning*, pages 462–471, 2018.
- [2] Syrine Belakaria, Aryan Deshwal, and Janardhan Doppa. Max-value entropy search for multi-objective Bayesian optimization. In *Proceedings of International Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [3] Syrine Belakaria, Aryan Deshwal, Nitthilan Kannappan Jayakodi, and Janardhan Rao Doppa. Uncertainty-aware search framework for multi-objective Bayesian optimization. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 10044–10052, 2020.
- [4] Syrine Belakaria, Derek Jackson, Yue Cao, Janardhan Rao Doppa, and Xiaonan Lu. Machine learning enabled fast multi-objective optimization for electrified aviation power system design. In *IEEE Energy Conversion Congress and Exposition (ECCE)*, 2020.
- [5] Syrine Belakaria, Zhiyuan Zhou, Aryan Deshwal, Janardhan Rao Doppa, Partha Pande, and Deuk Heo. Design of multi-output switched-capacitor voltage regulator via machine learning. In *Proceedings of the Twenty-Third IEEE/ACM Design Automation and Test in Europe Conference (DATE)*, 2020.
- [6] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley and Sons, 2012.
- [7] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, T Meyarivan, and A Fast. Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [8] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.
- [9] Aryan Deshwal, Syrine Belakaria, and Janardhan Rao Doppa. Scalable combinatorial Bayesian optimization with tractable statistical models. *CoRR*, abs/2008.08177, 2020.
- [10] Aryan Deshwal, Syrine Belakaria, Janardhan Rao Doppa, and Alan Fern. Optimizing discrete spaces via expensive evaluations: A learning to search framework. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 3773–3780, 2020.
- [11] Michael Emmerich and Jan-willem Klinkenberg. The computation of the expected improvement in dominated hypervolume of pareto front approximations. *Technical Report, Leiden University*, 34, 2008.
- [12] Paul Feliot, Julien Bect, and Emmanuel Vazquez. A bayesian approach to constrained single-and multi-objective optimization. *Journal of Global Optimization*, 67(1-2):97–133, 2017.
- [13] Eduardo C Garrido-Merchán and Daniel Hernández-Lobato. Predictive entropy search for multi-objective bayesian optimization with constraints. *Neurocomputing*, 361:50–68, 2019.
- [14] Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. Predictive entropy search for multi-objective bayesian optimization. In *Proceedings of International Conference on Machine Learning (ICML)*, pages 1492–1501, 2016.

- [15] José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems*, pages 918–926, 2014.
- [16] Joshua Knowles. Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, 2006.
- [17] Lars Kotthoff, Chris Thornton, Holger H Hoos, Frank Hutter, and Kevin Leyton-Brown. Auto-weka 2.0: Automatic model selection and hyperparameter optimization in weka. *Journal of Machine Learning Research (JMLR)*, 18(1):826–830, 2017.
- [18] Joseph Victor Michalowicz, Jonathan M Nichols, and Frank Bucholtz. *Handbook of differential entropy*. Chapman and Hall/CRC, 2013.
- [19] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184, 2008.
- [20] 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, 2016.
- [21] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2951–2959, 2012.
- [22] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT Press, 2006.
- [23] Eckart Zitzler. *Evolutionary algorithms for multiobjective optimization: Methods and applications*, volume 63. Ithaca: Shaker, 1999.

4 Full Derivation

Input space entropy based methods such as PESMO [14] selects the next candidate input \mathbf{x}_t (for ease of notation, we drop the subscript in below discussion) by maximizing the information gain about the optimal Pareto set \mathcal{X}^* . The acquisition function based on input space entropy is given as follows:

$$\alpha(\mathbf{x}) = I(\{\mathbf{x}, \mathbf{y}\}, \mathcal{X}^* \mid D) \tag{3}$$

$$= H(\mathcal{X}^* \mid D) - \mathbb{E}_{\mathbf{y}}[H(\mathcal{X}^* \mid D \cup \{\mathbf{x}, \mathbf{y}\})] \tag{4}$$

$$= H(\mathbf{y} \mid D, \mathbf{x}) - \mathbb{E}_{\mathcal{X}^*}[H(\mathbf{y} \mid D, \mathbf{x}, \mathcal{X}^*)] \tag{5}$$

Information gain is defined as the expected reduction in entropy $H(\cdot)$ of the posterior distribution $P(\mathcal{X}^* \mid D)$ over the optimal Pareto set \mathcal{X}^* as given in Equations 4 and 5 (resulting from symmetric property of information gain). This mathematical formulation relies on a very expensive and high-dimensional ($m \cdot d$ dimensions) distribution $P(\mathcal{X}^* \mid D)$, where m is the size of the optimal Pareto set \mathcal{X}^* . Furthermore, optimizing the second term in r.h.s poses significant challenges: a) requires a series of approximations [14] which can be potentially sub-optimal; and b) optimization, even after approximations, is expensive c) performance is strongly dependent on the number of Monte-Carlo samples.

To overcome the above challenges of computing input space entropy based acquisition function, [2] proposed to maximize the information gain about the optimal **Pareto front** \mathcal{Y}^* . However, MESMO did not address the challenge of constrained Pareto front. We propose an extension of MESMO’s acquisition function to maximize the information gain between the next candidate input for evaluation \mathbf{x} and constrained Pareto front \mathcal{Y}^* given as:

$$\alpha(\mathbf{x}) = I(\{\mathbf{x}, \mathbf{y}\}, \mathcal{Y}^* \mid D) \tag{6}$$

$$= H(\mathcal{Y}^* \mid D) - \mathbb{E}_{\mathbf{y}}[H(\mathcal{Y}^* \mid D \cup \{\mathbf{x}, \mathbf{y}\})] \tag{7}$$

$$= H(\mathbf{y} \mid D, \mathbf{x}) - \mathbb{E}_{\mathcal{Y}^*}[H(\mathbf{y} \mid D, \mathbf{x}, \mathcal{Y}^*)] \tag{8}$$

In this case, the output vector \mathbf{y} is $K + L$ dimensional: $\mathbf{y} = (y_{f_1}, y_{f_2}, \dots, y_{f_K}, y_{c_1}, \dots, y_{c_L})$ where $y_{f_j} = f_j(x)$ for all $j \in \{1, 2, \dots, K\}$ and $y_{c_i} = C_i(x)$ for all $i \in \{1, 2, \dots, L\}$. Consequently, the first term in the r.h.s of equation 8, entropy of a factorizable $(K + L)$ -dimensional Gaussian distribution $P(\mathbf{y} | D, \mathbf{x})$, can be computed in closed form as shown below:

$$H(\mathbf{y} | D, \mathbf{x}) = \frac{(K + C)(1 + \ln(2\pi))}{2} + \sum_{j=1}^K \ln(\sigma_{f_j}(\mathbf{x})) + \sum_{i=1}^L \ln(\sigma_{c_i}(\mathbf{x})) \quad (9)$$

where $\sigma_{f_j}^2(\mathbf{x})$ and $\sigma_{c_i}^2(\mathbf{x})$ are the predictive variances of j^{th} function and i^{th} constraint GPs respectively at input \mathbf{x} . The second term in the r.h.s of equation 8 is an expectation over the Pareto front \mathcal{Y}^* . We can approximately compute this term via Monte-Carlo sampling as shown below:

$$\mathbb{E}_{\mathcal{Y}^*}[H(\mathbf{y} | D, \mathbf{x}, \mathcal{Y}^*)] \simeq \frac{1}{S} \sum_{s=1}^S [H(\mathbf{y} | D, \mathbf{x}, \mathcal{Y}_s^*)] \quad (10)$$

where S is the number of samples and \mathcal{Y}_s^* denote a sample Pareto front. The main advantages of our acquisition function are: computational efficiency and robustness to the number of samples ([2]).

There are two key algorithmic steps to compute Equation 10: 1) How to compute Pareto front samples \mathcal{Y}_s^* ?; and 2) How to compute the entropy with respect to a given Pareto front sample \mathcal{Y}_s^* ? We provide solutions for these two questions below.

1) Computing Pareto front samples via cheap multi-objective optimization. To compute a Pareto front sample \mathcal{Y}_s^* , we first sample functions and constraints from the posterior GP models via random fourier features ([15, 19]) and then solve a cheap multi-objective optimization over the K sampled functions and L sampled constraints.

Cheap MO solver. We sample \tilde{f}_i from GP model \mathcal{M}_{f_j} for each of the K functions and \tilde{C}_i from GP model \mathcal{M}_{c_i} for each of the L constraints. A *cheap* constrained multi-objective optimization problem over the K sampled functions $\tilde{f}_1, \tilde{f}_2, \dots, \tilde{f}_k$ and the L sampled constraints $\tilde{C}_1, \tilde{C}_2, \dots, \tilde{C}_L$ is solved to compute the sample Pareto front \mathcal{Y}_s^* . This cheap multi-objective optimization also allows us to capture the interactions between different objectives while satisfying the constraints. We employ the popular constrained NSGA-II algorithm ([7, 8]) to solve the constrained MO problem with cheap objective functions noting that any other algorithm can be used to similar effect.

2) Entropy computation with a sample Pareto front. Let $\mathcal{Y}_s^* = \{\mathbf{z}^1, \dots, \mathbf{z}^m\}$ be the sample Pareto front, where m is the size of the Pareto front and each \mathbf{z}^i is a $(K + L)$ -vector evaluated at the K sampled functions and L sampled constraints $\mathbf{z}^i = \{z_{f_1}^i, \dots, z_{f_K}^i, z_{c_1}^i, \dots, z_{c_L}^i\}$. The following inequality holds for each component y_j of the $(K + L)$ -vector $\mathbf{y} = \{y_{f_1}, \dots, y_{f_K}, y_{c_1}, \dots, y_{c_L}\}$ in the entropy term $H(\mathbf{y} | D, \mathbf{x}, \mathcal{Y}_s^*)$:

$$y_j \leq \max\{z_j^1, \dots, z_j^m\} \quad \forall j \in \{f_1, \dots, f_K, c_1, \dots, c_L\} \quad (11)$$

The inequality essentially says that the j^{th} component of \mathbf{y} (i.e., y_j) is upper-bounded by a value obtained by taking the maximum of j^{th} components of all m $(K + L)$ -vectors in the Pareto front \mathcal{Y}_s^* . This inequality had been proven by a contradiction in ([2]) for $j \in \{f_1, \dots, f_K\}$. We assume the same for $j \in \{c_1, \dots, c_L\}$.

By combining the inequality 11 and the fact that each function is modeled as an independent GP, we can model each component y_j as a truncated Gaussian distribution since the distribution of y_j needs to satisfy $y_j \leq \max\{z_j^1, \dots, z_j^m\}$. Furthermore, a common property of entropy measure allows us to decompose the entropy of a set of independent variables into a sum over entropies of individual variables [6]:

$$H(\mathbf{y} | D, \mathbf{x}, \mathcal{Y}_s^*) \simeq \sum_{j=1}^K H(y_{f_j} | D, \mathbf{x}, \max\{z_{f_j}^1, \dots, z_{f_j}^m\}) + \sum_{i=1}^L H(y_{c_i} | D, \mathbf{x}, \max\{z_{c_i}^1, \dots, z_{c_i}^m\}) \quad (12)$$

The r.h.s is a summation over entropies of $(K + L)$ -variables $\mathbf{y} = \{y_{f_1}, \dots, y_{f_K}, y_{c_1}, \dots, y_{c_L}\}$. The differential entropy for each y_j is the entropy of a truncated Gaussian distribution ([18]) and given by the following equations:

$$H(y_{f_j}|D, \mathbf{x}, y_s^{f_j^*}) \simeq \left[\frac{(1 + \ln(2\pi))}{2} + \ln(\sigma_{f_j}(\mathbf{x})) + \ln \Phi(\gamma_s^{f_j}(\mathbf{x})) - \frac{\gamma_s^{f_j}(\mathbf{x})\phi(\gamma_s^{f_j}(\mathbf{x}))}{2\Phi(\gamma_s^{f_j}(\mathbf{x}))} \right] \quad (13)$$

$$H(y_{c_i}|D, \mathbf{x}, y_s^{c_i^*}) \simeq \left[\frac{(1 + \ln(2\pi))}{2} + \ln(\sigma_{c_i}(\mathbf{x})) + \ln \Phi(\gamma_s^{c_i}(\mathbf{x})) - \frac{\gamma_s^{c_i}(\mathbf{x})\phi(\gamma_s^{c_i}(\mathbf{x}))}{2\Phi(\gamma_s^{c_i}(\mathbf{x}))} \right] \quad (14)$$

Consequently we have:

$$H(\mathbf{y} | D, \mathbf{x}, \mathcal{Y}_s^*) \simeq \sum_{j=1}^K \left[\frac{(1 + \ln(2\pi))}{2} + \ln(\sigma_{f_j}(\mathbf{x})) + \ln \Phi(\gamma_s^{f_j}(\mathbf{x})) - \frac{\gamma_s^{f_j}(\mathbf{x})\phi(\gamma_s^{f_j}(\mathbf{x}))}{2\Phi(\gamma_s^{f_j}(\mathbf{x}))} \right] + \sum_{i=1}^L \left[\frac{(1 + \ln(2\pi))}{2} + \ln(\sigma_{c_i}(\mathbf{x})) + \ln \Phi(\gamma_s^{c_i}(\mathbf{x})) - \frac{\gamma_s^{c_i}(\mathbf{x})\phi(\gamma_s^{c_i}(\mathbf{x}))}{2\Phi(\gamma_s^{c_i}(\mathbf{x}))} \right] \quad (15)$$

where $\gamma_s^{c_i}(x) = \frac{y_s^{c_i^*} - \mu_{c_i}(\mathbf{x})}{\sigma_{c_i}(\mathbf{x})}$, $\gamma_s^{f_j}(x) = \frac{y_s^{f_j^*} - \mu_{f_j}(\mathbf{x})}{\sigma_{f_j}(\mathbf{x})}$, $y_s^{c_i^*}$ and $y_s^{f_j^*}$ are the maximum values of constraint \tilde{c}_i and function \tilde{f}_j reached after the cheap multi-objective optimization over sampled functions and constraints: $y_s^{c_i^*} = \max\{z_{c_i}^1, \dots, z_{c_i}^m\}$, $y_s^{f_j^*} = \max\{z_{f_j}^1, \dots, z_{f_j}^m\}$, ϕ and Φ are the p.d.f and c.d.f of a standard normal distribution respectively. By combining equations 9 and 15 with Equation 8, we get the final form of our acquisition function as shown below:

$$\alpha(\mathbf{x}) \simeq \frac{1}{S} \sum_{s=1}^S \left[\sum_{j=1}^K \frac{\gamma_s^{f_j}(\mathbf{x})\phi(\gamma_s^{f_j}(\mathbf{x}))}{2\Phi(\gamma_s^{f_j}(\mathbf{x}))} - \ln \Phi(\gamma_s^{f_j}(\mathbf{x})) + \sum_{i=1}^L \frac{\gamma_s^{c_i}(\mathbf{x})\phi(\gamma_s^{c_i}(\mathbf{x}))}{2\Phi(\gamma_s^{c_i}(\mathbf{x}))} - \ln \Phi(\gamma_s^{c_i}(\mathbf{x})) \right] \quad (16)$$

5 Additional results

SPECS	NSGA-II		PESMOC		MESMOC	
$V_{ref1}(V)$	0.6	0.5	0.52	0.53	0.63	0.52
$V_{ref2}(V)$	0.55	0.62	0.55	0.61	0.51	0.53
$V_{ref3}(V)$	1.06	1.06	1.07	1.12	1.05	1.13
$V_{ref4}(V)$	1.07	1.09	1.09	1.06	1.05	1.06
$V_{o1}(mV)$	699.6	713.1	677.10	760.60	678.40	551.62
$V_{o2}(mV)$	700.4	712.2	690.70	725.70	520.61	632.80
$V_{o3}(V)$	1.10	1.06	1.08	1.15	1.12	1.16
$V_{o4}(V)$	1.09	1.09	1.08	0.99	1.14	1.08
Eff (%)	73.26	71.85	76.20	74.82	88.81	88.53

Figure 2: Comparison table of optimized circuit parameters obtained with different algorithms (designs are selected from the Pareto set prioritized by efficiency)