
If Optimizing for General Parameters in Chemistry is Useful, Why is it Hardly Done?

Stefan P. Schmid^{*,a,b} Ella Miray Rajaonson^{c,d} Cher Tian Ser^{c,d} Mohammad Haddadnia^{e,f}
Shi Xuan Leong^{c,g} Alán Aspuru-Guzik^{c,d,h,i,j,k,l} Agustinus Kristiadi^d Kjell Jorner^{a,b}
Felix Strieth-Kalthoff^m

^aDepartment of Chemistry and Applied Biosciences, ETH Zurich

^bNCCR Catalysis

^cDepartment of Chemistry, University of Toronto

^dVector Institute

^eDepartment of Biological Chemistry & Molecular Pharmacology, Harvard University

^fDana-Farber Cancer Institute

^gSchool of Chemistry, Chemical Engineering and Biotechnology, Nanyang Technological University

^hDepartment of Computer Science, University of Toronto

ⁱDepartment of Chemical Engineering and Applied Chemistry, University of Toronto

^jDepartment of Materials Science and Engineering, University of Toronto

^kAcceleration Consortium

^lCanadian Institute for Advanced Research (CIFAR)

^mSchool of Mathematics and Natural Sciences, University of Wuppertal

Abstract

General parameters are highly desirable in the natural sciences — e.g., reaction conditions that enable high yields across a range of related transformations. This has a significant practical impact since those general parameters can be transferred to related tasks without the need for laborious and time-intensive re-optimization. While Bayesian optimization (BO) is widely applied to find optimal parameter sets for specific tasks, it has remained underused in experiment planning towards such general optima. In this work, we consider the the real-world problem of condition optimization for chemical reactions to study whether performing generality-oriented BO can accelerate the identification of general optima, and whether these optima also translate to unseen examples. This is achieved through a careful formulation of the problem as an optimization over curried functions, as well as systematic benchmarking of generality-oriented strategies for optimization tasks on real-world experimental data. We find that the optimization for general reaction conditions are determined by the sampling of substrates, with random selection outperforming more data-driven strategies.

1 Introduction

Identifying parameters that deliver satisfactory performance on a wide set of tasks, which we refer to as *general parameters*, is crucial for numerous real-world challenges. Examples are the identification of sensor settings that allow accurate measurements in different environments [27], or the design of footwear that provides good performance for a wide range of people on different undergrounds [51]. A particularly prominent example comes from the domain of chemical synthesis, where finding reaction

*Correspondence to: stefan.schmid@chem.ethz.ch, aspuru@utoronto.ca, akristiadi@vectorinstitute.ai, kjell.jorner@chem.ethz.ch, strieth-kalthoff@uni-wuppertal.de.




Figure 1: *Left*: While conditions can be optimized to maximize the reaction outcome for only one substrate (red), generality-optimized conditions provide a satisfactory reaction outcome for multiple substrates. *Right*: Optimization loop for generality-oriented optimization under partial monitoring.

conditions under which different starting materials can be reliably converted into the corresponding products, remains a critical challenge [9, 50, 52, 54, 58, 65]. Such general reaction conditions are of particular interest, *e.g.*, in the pharmaceutical industry, where thousands of reactions of each type are carried out regularly, and optimizing each reaction is unfeasible [65].

In general, the ability to rapidly identify optimal reaction conditions has posed a long-standing challenge to synthetic chemists, even for a single starting material [62]. At the same time, this is crucial for achieving an efficient and sustainable chemical industry [5, 18, 57]. In recent years, traditional human-centric optimization workflows have been complemented by data-driven methods such as Bayesian Optimization (BO) in order to increase optimization efficiency [11, 12, 25, 60]. Coupled with robotic reaction execution, autonomous reaction optimization has emerged as a key application of self-driving labs in the chemical domain to date [64]. However, in the vast majority of cases, reaction conditions are specifically optimized for single substrates, neglecting any generality considerations during optimization (Figure 1, left-hand side) [13, 62].

This can be attributed to the fact that directly observing the generality of selected conditions is associated with largely increased experimental costs, as experimental evaluations on multiple substrates are required. Attempts to reduce the required number of experiments inevitably increase the complexity of the decision-making process, thus, in the absence of appropriate decision-making algorithms, hindering its widespread usage in real-world scenarios. Here, generality-oriented optimization turns into a *partial monitoring* scenario, in which each condition can only be evaluated on a subset of all possible substrates. As a consequence, any iterative experiment planning algorithm needs to recommend both the condition and the substrate for the next experimental evaluation (Figure 1, right-hand side). Experimentally measuring the outcome of the recommended experiment corresponds to a partial observation of the generality objective, which needs to be taken into account when recommending the next experiment.

Although, in the past two years, isolated studies have targeted the identification of general reaction conditions through a variation of BO [6], multi-armed bandit optimization [66], or genetic algorithms [17], generality-oriented optimizations are still not commonly performed in real-world experiments (see Section 2.2.3). This likely arises from the fact that the applicability and limitations of these algorithms are yet to be understood, which is crucial for their effective integration into real-world laboratory workflows [64].

For these reasons, we herein perform a systematic study into generality-oriented optimization. Therefore, we formulate generality-oriented optimization as an optimization problem over curried functions, and perform systematic benchmarks on various real-world chemical reaction optimization tasks. Specifically for the latter, we (i) confirm the expectation that optimization over multiple substrates leads to more general optima, and (ii) demonstrate that the finding these optima crucially depends on effective strategies for navigating the space of all possible substrates.

In summary, our contributions are four-fold:

- Principled formulation of generality-oriented optimization as an optimization problem over a curried function.
- Expansion and adaptation of established reaction optimization benchmark tasks, improving their utility as benchmarks for generality-oriented BO.
- Evaluation of different optimization algorithms for identifying general optima.
- CURRYBO as an open-source extension to BOTORCH [7] for generality-oriented optimization problems [1].

2 Foundations of generality-oriented Bayesian optimization

To formalize the generality-oriented optimization problem, we provide a principled outline by considering it as an extension of established global optimization approaches over curried functions. For clarity, we also discuss its distinction to different variations of global optimization, including multiobjective, multifidelity, and mixed-variable optimization.

2.1 Global optimization

Global black-box optimization is concerned with finding the optimum of an unknown objective function $f(\mathbf{x})$:

$$\hat{\mathbf{x}} = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmax}} f(\mathbf{x}) \quad (1)$$

Suppose $f(\mathbf{x})$ is a function that (a) is not analytically tractable, (b) is very expensive to evaluate, and (c) can only be evaluated without obtaining gradient information. In this scenario, BO has emerged as a ubiquitous approach for finding the global optimum $\hat{\mathbf{x}} \in \mathcal{X}$ in a sample-efficient manner [19]. The working principle of BO involves a probabilistic surrogate model $g(\mathbf{x})$ to approximate $f(\mathbf{x})$, which can be used to compute a predictive posterior distribution over g under all previous observations $\mathcal{D} = \{(\mathbf{x}_i, f(\mathbf{x}_i))\}_{i=1}^k$. The most prominent choice for $p(g(\mathbf{x}) | \mathcal{D})$ are Gaussian processes [GPs; 53], with various types of Bayesian neural networks becoming increasingly popular in the past decade [29, 36, 37, 44]. Based on the predictive posterior, an acquisition function α over the input space \mathcal{X} is used to decide at which $\mathbf{x}_{\text{next}} \in \mathcal{X}$ the objective function should be evaluated next. Key to the success of BO is the implicit exploitation–exploration tradeoff in α , which makes use of the posterior distribution $p(g(\mathbf{x}) | \mathcal{D})$ [47]. Common choices of α are Upper Confidence Bound [UCB; 3, 33, 34], Expected Improvement [32], Knowledge Gradient [15, 16, 26] or Thompson Sampling [TS; 63]. The hereby selected \mathbf{x}_{next} is evaluated experimentally, resulting in $f(\mathbf{x}_{\text{next}})$, and the described procedure is repeated until a satisfactory outcome is observed, or the experimentation budget is exhausted.

2.2 Global optimization for generality

2.2.1 Problem formulation

Extending the global optimization framework, we consider a black-box function $f : \mathcal{X} \times \mathcal{W} \rightarrow \mathbb{R}$ in joint space $\mathcal{X} \times \mathcal{W}$, where $\mathbf{x} \in \mathcal{X}$ can be continuous, discrete or mixed-variable. $\mathcal{W} = \{\mathbf{w}_i\}_{i=1}^n$ is a discrete parameter space of size n (see Figure 2). Each evaluation of f is expensive and does not provide any gradient information. In the example of reaction condition optimization, \mathbf{x} corresponds to conditions from the reaction condition space \mathcal{X} (e.g. the reaction temperature), and $\mathbf{w} \in \mathcal{W}$ are the substrates (starting materials of a reaction) that are considered for generality-oriented optimization. Let curry be a currying operator on the second argument, i.e., $\operatorname{curry}(f) : \mathcal{W} \rightarrow (\mathcal{X} \rightarrow \mathbb{R})$, that transforms a function taking two arguments into a family of functions, each taking one argument. Then, for some $\mathbf{w} \in \mathcal{W}$, evaluating $\operatorname{curry}(f)(\mathbf{w})$ yields a new function $f(\cdot; \mathbf{w}) : \mathcal{X} \rightarrow \mathbb{R}$, where $f(\mathbf{x}; \mathbf{w}) = f(\mathbf{x}, \mathbf{w})$. Importantly, these $f(\cdot; \mathbf{w}) : \mathcal{X} \rightarrow \mathbb{R}$ correspond to functions that can be evaluated experimentally (i.e. the reaction of a specific substrate as a function of conditions), even though evaluations are expensive. This allows us to describe all observable functions through an n -sized set $\mathcal{F} = \{f(\cdot; \mathbf{w}_i) : \mathcal{X} \rightarrow \mathbb{R}\}_{i=1}^n$. In the context of reaction condition optimization, \mathcal{F} consists of all functions that describe the reaction outcome for each substrate.

Evaluation of a specific $f(\mathbf{x}_{\text{obs}}; \mathbf{w}_{\text{obs}})$ then corresponds to measuring the reaction outcome of a substrate (described by \mathbf{w}_{obs}) under specific reaction conditions \mathbf{x}_{obs} .




Figure 2: Conceptual overview of the generality-oriented optimization problem. *Left*: The function values across the joint space $\mathcal{X} \times \mathcal{W}$. *Right*: Mean aggregation applied to the function family $f(\mathbf{x}; \mathbf{w})$, that is obtained via currying of the joint space $\mathcal{X} \times \mathcal{W}$. The quantity $\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W})$ constitutes the partially observable objective function, of which $\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \in \mathcal{X}} \phi(\mathbf{x})$ is the optimum that should be identified.

In generality-oriented optimization, the goal is to identify the optimum $\hat{\mathbf{x}} \in \mathcal{X}$ that is generally optimal across \mathcal{W} , meaning $\hat{\mathbf{x}}$ maximizes a user-defined generality metric over all $\mathbf{w} \in \mathcal{W}$ (see Figure 2 for illustration). We refer to this generality metric as the *aggregation function* ϕ :

$$\hat{\mathbf{x}} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \phi(\mathbf{x}) := \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W}) \quad (2)$$

In the reaction optimization example, this corresponds to conditions (e.g. reaction temperature) that give e.g. the highest average yield over all considered substrates. In this scenario, the choice of ϕ is the mean $\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W}) = 1/|\mathcal{W}| \sum_{\mathbf{w} \in \mathcal{W}} f(\mathbf{x}; \mathbf{w})$. An alternative choice of ϕ could be the number of function values $\{f(\mathbf{x}; \mathbf{w}_i)\}_{i=1}^n$ above a user-defined threshold [9]. Further practically relevant aggregation functions are described in Appendix A.1.1.

While equation 2 appears like a standard global optimization problem over \mathcal{X} , evaluating $\phi(\mathbf{x})$ itself is intractable due to the aggregation over \mathcal{W} . Indeed, to evaluate $\phi(\mathbf{x})$ on a single \mathbf{x} , one must perform n -many expensive function evaluations to first obtain $\{f(\mathbf{x}; \mathbf{w}_i)\}_{i=1}^n$. Due to this impracticality, ideally, the number of such function evaluations is minimized. Thus, this setting differs from the conventional global optimization problem, due to its *partial observation* nature: One can only compute $\phi(\mathbf{x})$ via a subset of observations $\{f(\mathbf{x}; \mathbf{w}_j)\}_{j=1}^m$ where $m < n$.

To maximize sample efficiency, an optimizer should always recommend a new pair $(\mathbf{x}_{\text{next}}, \mathbf{w}_{\text{next}})$ to evaluate next — in other words: $\phi(\mathbf{x}_{\text{next}})$ is only observed partially via a single evaluation of f , i.e., $m = 1$. Treating this in the conventional framework of BO, we can build a probabilistic surrogate model $g(\mathbf{x}_i; \mathbf{w}_i)$ from all k available observations $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{w}_i, f(\mathbf{x}_i; \mathbf{w}_i))\}_{i=1}^k$. From the posterior distribution over g , we estimate the posterior distribution over ϕ via Monte-Carlo integration [7], see Appendix A.1.2 for further details.

Unlike the conventional BO case, we now need a specific acquisition policy A to decide which $\mathbf{x} \in \mathcal{X}$ and $\mathbf{w} \in \mathcal{W}$ the aggregated objective function $\phi(\mathbf{x})$ should be partially evaluated. Note that A plays an important role since it must respect the partial observability constraint. That is, it must also propose a *single* \mathbf{w} at each BO step such that the general (over *all* \mathbf{w}_i 's) optimum $\hat{\mathbf{x}}$ is obtained in as few steps as possible. We discuss specific choices of A in Table 1; more details in Appendix A.1.3.

Given the pair $(\mathbf{x}_{k+1}, \mathbf{w}_{k+1})$, the aggregated objective $\phi(\mathbf{x}_{k+1})$ is partially observed, \mathcal{D} is updated, and the discussed steps are repeated until the experimentation budget is exhausted. Eventually, owing to the partial monitoring scenario [42, 43, 55], the final optimum after a budget of k experiments, $\hat{\mathbf{x}}_k$, is returned greedily, i.e. as the $\mathbf{x} \in \mathcal{X}$ that maximizes the mean of the predictive posterior of ϕ . A summary of this is provided in Algorithm 1.

2.2.2 Distinction from existing variants of the BO formalism

Despite seeming similarities with *multiobjective*, *multifidelity*, and *mixed-variable* optimization, the generality-oriented approach describes a distinctly different scenario, necessitating distinct algorithmic solutions.

Algorithm 1 Generality-oriented Bayesian optimization

Input:

Set of observable functions $\mathcal{F} = \{f(\cdot; \mathbf{w}_i) : \mathcal{X} \rightarrow \mathbb{R}\}_{i=1}^n$
Initial dataset $\mathcal{D}_k = \{\mathbf{x}_j, \mathbf{w}_j, f(\mathbf{x}_j; \mathbf{w}_j)\}_{j=1}^k$
Aggregation function $\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W})$
Surrogate model $g(\mathbf{x}, \mathbf{w})$ and acquisition policy A
Budget K

- 1: **while** $k \leq K$ **do**
- 2: Compute posterior distribution $p(g_k(\mathbf{x}, \mathbf{w}) \mid \mathcal{D}_k)$
- 3: Acquire $\mathbf{x}_{k+1}, \mathbf{w}_{k+1} = A(p(g_k(\mathbf{x}, \mathbf{w}) \mid \mathcal{D}_k), \phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W}))$
- 4: Observe $f(\mathbf{x}_{k+1}; \mathbf{w}_{k+1})$
- 5: Update $\mathcal{D}_{k+1} = \mathcal{D}_k \cup \{(\mathbf{x}_{k+1}, \mathbf{w}_{k+1}, f(\mathbf{x}_{k+1}; \mathbf{w}_{k+1}))\}$
- 6: $k = k + 1$
- 7: **end while**
- 8: **return** $\hat{\mathbf{x}} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \mathbb{E} \left[p(\phi(\mathbf{x}) \mid \mathcal{D}_K) \mid \mathbf{x} \right]$

- In contrast to *multiobjective* optimization, here, we consider a single optimization objective, i.e. $\phi(\mathbf{x})$. However, this objective can only be partially observed. Whereas the overall optimization problem aims to identify $\hat{\mathbf{x}} \in \mathcal{X}$, finding the next recommended observation requires a joint optimization over \mathcal{X} and \mathcal{W} .
- In contrast to *multifidelity* BO, the functions parameterized by $\mathbf{w} \in \mathcal{W}$ do not correspond to the same objective with different fidelities. Rather, they are independent functions which all contribute equally to the objective function $\phi(\mathbf{x})$.
- Unlike *mixed-variable* BO [14], the goal of generality-oriented BO is not to find (\mathbf{x}, \mathbf{w}) that maximizes the objective in the *joint* space. Rather, the goal is to find the set optimum $\hat{\mathbf{x}}$ that maximizes $\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W})$ over $f(\mathbf{x}; \mathbf{w})$. In the case of ϕ being a sum, this bears resemblance to maximizing the *marginal* over \mathbf{x} (see Figure 2). Moreover, \mathcal{X} itself can be continuous, discrete or mixed, thus, $\mathcal{X} \times \mathcal{W}$ can be a fully-discrete space.

2.2.3 Related works

Given the enormous importance of general reaction conditions for accelerating molecular discovery, the concept of "reaction generality" has been discussed on multiple occasions in the past years [9, 50, 52, 54, 65], and has been included in computational molecular design studies [17].

The first example of actual generality-oriented optimization has been reported by Angello et al. [6], who describe a modification of BO, sequentially acquiring \mathbf{x}_{next} (maximized probability of improvement over all $\mathbf{x} \in \mathcal{X}$) and \mathbf{w}_{next} (minimized posterior variance over all $\mathbf{w} \in \mathcal{W}$ given \mathbf{x}_{next}). The authors demonstrate its applicability in automated experiments on Suzuki–Miyaura cross couplings. Their algorithm is evaluated herein as the CURRYBO 1LA-VAR strategy.

Following an alternative strategy, Wang et al. [66] recently formulated generality-oriented optimization as a multi-armed bandit problem, where each arm corresponds to a possible reaction condition. It should be noted that, in this scenario, uniform sampling of all arms is required for initializing the algorithm. While the approach has mainly been tested for optimization campaigns with few possible reaction conditions, this requirement renders its application impractical for a high number of discrete conditions or even continuous variables. However, such cases are commonly faced in real-world optimization campaigns and benchmarked herein.

Despite these advances, their applicability and limitations have remained unclear. Thus, after the problem formulation provided in Section 2.2.1, our work provides a comparison over different generality-oriented strategies. Our analysis focuses on published approaches, as well as additional BO-based strategies, owing to their demonstrated real-world applicability in reaction condition optimization [12, 64].

Due to the partial monitoring nature of generality-oriented optimization, we want to highlight work that has been conducted on the partial monitoring case in bandit optimization scenarios [42, 43, 55]. However, to the best of our knowledge, work in this field has mostly dealt with an information-theoretic approach towards optimally scaling algorithms. We refer the readers to select publications [35, 39–41]. However, systematic evaluations in scenarios that require sample-efficient optimization, particularly addressing the early stages of optimization campaigns, has not been targeted.

3 Methods

3.1 Experimental benchmark problems

We consider two real-world chemical reaction benchmark problems stemming from high-throughput experimentation [HTE; 10, 68]. Each problem evaluates the optimization of a chemically relevant reaction outcome (such as enantioselectivity $\Delta\Delta G^\ddagger$, or starting material conversion), and contains an experimental dataset of substrates, conditions and measured outcomes. For each benchmark, generality is evaluated over substrates or substrate pairs (16 - 25 discrete options), and is optimized over one or two categorical reaction conditions (43 - 96 discrete conditions). Since it has been well-demonstrated that these problems can be effectively modeled by regression approaches [4, 56, 68], we trained a random forest regressor on each dataset, which was used as the ground-truth emulator for all experiments [30]. Extensive analysis of the benchmark problems can be found in Appendix A.2. Moreover, to overcome well-known biases towards high-outcome reactions, we augmented the existing search spaces via a chemically sensible expansion workflow (see Appendix A.2.2). The STONED algorithm [49] was used to generate a pool of new in-domain substrates which possess the required functional groups for the reaction, and exhibit structural similarity (fingerprint distance) above a user-defined threshold. From this pool, substrates with low reaction outcomes, as predicted by the ground-truth emulator, were preferentially sampled for expanding the original search space, ensuring enrichment with low-outcome reactions.

To analyze the utility of generality-oriented BO for each benchmark problem, we first use exhaustive grid search over all potential reaction conditions to identify the most general reaction conditions across a varying number of "train" substrates. The generality of these conditions was subsequently evaluated on a held-out test set of substrates to derive a test-set independent generality score, that scales the generality of selected conditions with the minimal and maximal obtainable generality across the test set substrates. With this generality score, we investigate whether generality-oriented optimization across more substrates indeed yields more transferable reaction conditions. Further details on the generality score and the data-analysis are provided in Appendix A.3.

3.2 Optimization algorithms

Using the benchmark problems outlined above, we perform systematic evaluations of multiple methods for the identification of general optima. Table 1 summarizes the strategies for recommending the next data point $(\mathbf{x}_{\text{next}}, \mathbf{w}_{\text{next}})$ (cf. Section 2.2) that are discussed in the main text. Each strategy is evaluated under two different definitions of generality, namely the *mean* and the *number-above-threshold* aggregation (threshold aggregation) functions described in Section 2.2.1 (see Appendix A.1.1 for further details).

In all BO experiments, we used a GP regressor, as provided in *BoTorch*, with the Tanimoto kernel as implemented in *Gauche* [24]. Molecules were represented using Morgan Fingerprints [46] with 1024 bits and a radius of 2. Fingerprints were generated using RDKit [2]. For each experiment, we provide statistics over 30 independent runs, each performed over different substrates and initial conditions, but with fixed random seeds to ensure cross-method comparability. Further baseline experiments are discussed in Appendix A.4.

4 Results and discussion

We commenced our analysis by investigating the overall utility of generality-oriented optimization to yield general and transferable optima. In our initial studies, we evaluated generality as the average outcome over all $\mathbf{w} \in \mathcal{W}$ (*mean* aggregation function). Exhaustive grid search, which enables us to identify the ground truth optima for each set of train substrates evaluated during optimization, reveals

Table 1: Benchmarked acquisition policies

Optimization strategy	Description
CURRYBO 2LA-RAND	Sequential two-step lookahead acquisition policy, using Thompson Sampling acquisition function to select \mathbf{x}_{next} and a Random selection of \mathbf{w}_{next} with a greedy final recommendation
CURRYBO 1LA-VAR	Sequential one-step lookahead acquisition policy, using an UCB acquisition function with $\beta = 0.5$ to select \mathbf{x}_{next} and a posterior variance acquisition function to select \mathbf{w}_{next} with a greedy final recommendation. This strategy resembles the algorithm described by Angello et al. [6]
CURRYBO 1LA-RAND	Sequential one-step lookahead acquisition policy, using an UCB acquisition function with $\beta = 0.5$ to select \mathbf{x}_{next} and a random acquisition of \mathbf{w}_{next} with a greedy final recommendation.
BANDIT	Multi-armed bandit algorithm as outlined and implemented by Wang et al. [66]
RANDOM GREEDY	Sequential acquisition policy, using a random acquisition function to select \mathbf{x}_{next} and a random acquisition function to select \mathbf{w}_{next} with a greedy final recommendation
RANDOM	Random acquisition policy of \mathbf{x}_{next} and \mathbf{w}_{next} and final random recommendation

that with an increasing number of train substrates, the transferability of the found optima to a held-out set of test substrate increases (Figure 3, left). While this finding is arguably unsurprising, and merely confirms a common assumption in the field [65], it indicates possible caveats concerning the use of these problems as benchmarks for generality-oriented optimization: Even with larger numbers of train substrates, the ground-truth optima for the train sets did not lead to optimal outcomes on the corresponding test sets, as evidenced by test set generality scores significantly below 1.0.

These observations indicate the existence of multiple scattered local optima. We attribute this to an overall high density of high-outcome results across the search space (see Appendix A.2) – a well-known bias in chemical reaction databases [8, 61]. This is, however, not necessarily reflective of the "real-world" scenario encountered in reaction condition optimization. At this stage, it should be noted that, while widely used as such, the problems have not been designed as benchmarks for reaction condition optimization. These issues motivated the search space expansion, as discussed in Section 3.1. On the augmented benchmark tasks, we find that transferability of the identified optima




Figure 3: *Top*: Normalized test-set generality score as determined by exhaustive grid search for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems. Average and standard error are taken from thirty different train/test substrates splits. *Bottom*: Normalized test-set generality score as determined by exhaustive grid search for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems. Average and standard error are taken from thirty different train/test substrates splits.




Figure 4: Generalization-oriented optimization trajectories of different algorithms, averaged over all benchmark problems on the original (left) and augmented (right) search spaces. Optimization algorithms are described in Table 1.

to a held-out test set is significantly improved (Figure 3, right), and the consideration of multiple substrates during optimization robustly leads to well-generalizing optima. Notably, these observations are not limited to the definition of generality as the average over all $\mathbf{w} \in \mathcal{W}$, but remain valid for further aggregation functions (see Appendix A.6.1). These findings underline that – especially in "needle in a haystack scenarios" – generality-oriented optimization is indeed necessary for finding transferable optima. Most importantly, these scenarios are most commonly found in real-world reaction optimization, where for most reactions, the majority of possible conditions do not lead to observable product quantities. This re-emphasizes the need for benchmark problems that reflect experimental reality.

With these findings at hand, we investigated how to select the $\mathcal{W}_{\text{train}} \subset \mathcal{W}$ to identify maximally transferable optima. Surprisingly, we find that molecular-distance-based selection, aiming at a broad coverage of \mathcal{W} [20, 21, 28, 59], did not outperform random sampling approaches. We provide an extended discussion of this observation in Appendix A.6.1.

Eventually, having established the fundamental utility of generality-oriented optimization, we set out to perform a systematic evaluation into how to identify those optima using iterative optimization under partial monitoring of the optimization objective. We evaluate previously reported strategies [6, 66], as well as additional BO-based approaches (Table 1). As a summary, Figure 4 shows the optimization trajectories of these different algorithms, quantified as a normalized, problem-independent optimization metric ($\text{GAP} = (y_i - y_0)/(y^* - y_0)$, where y_i 's are maximum observed values and y^* is the true global optimum) [31], and averaged across both benchmark problems. Individual optimization trajectories for each benchmark problem are shown in Figure 5 and in Appendix A.6.

In spite of performance differences across the benchmark tasks, we were surprised to find overall good performance of random data point acquisition combined with a model-based greedy prediction of the final global optimum (RANDOM GREEDY). As this strategy is seldomly employed as a baseline in optimization, it is interesting to observe optimization behavior which is almost on par with iterative data-driven techniques — particularly in the context of fixed-budget optimization. In fact, in many scenarios, it outperformed data-driven algorithms for iteratively picking and evaluating \mathbf{x}_{next} and \mathbf{w}_{next} (see Figure 4 and Figure 5). The optimization behavior of the BANDIT algorithm is significantly affected by the necessity to uniformly sample all arms (i.e. all $\mathbf{x} \in \mathcal{X}$) at the outset of a campaign, which disfavors it in scenarios with large \mathcal{X} , such as the Pd-catalyzed carbon-heteroatom coupling benchmark (see Figure 5). In contrast, for small to medium-sized \mathcal{X} (e.g. the N,S-acetal formation benchmark), competitive optimization behavior is observed after sampling all $\mathbf{x} \in \mathcal{X}$ (see Figure 5).

In contrast, the BO-based algorithm (CURRYBO 1LA-VAR) recommends improved general conditions already after few experimental evaluations. However, we observe that sequentially acquiring \mathbf{x}_{next} by optimizing a conventional acquisition function, and \mathbf{w}_{next} based on posterior variance (i.e. uncertainty)[6], leads to comparatively slow optimization, with little indication of reliably converging towards the global optimum (Figure 4). To identify the origin of this behavior, we subsequently




Figure 5: Generality-oriented optimization trajectories of different algorithms, shown for the augmented Pd-catalyzed carbon-heteroatom coupling benchmark (left) and the augmented N,S-Acetal formation benchmark (right). Optimization algorithms are described in Table 1.




Figure 6: Generality-oriented optimization trajectories of different algorithms, averaged over all benchmark problems on the original (left) and augmented (right) search spaces. Optimization algorithms are described in Table 1.

evaluated variations of this algorithm, using different strategies for acquiring w_{next} . To our surprise, randomly selecting w_{next} (similar to the BANDIT algorithm) showed significantly more sample-efficient optimization than uncertainty-based selection, and consistently converged towards the global optimum [CURRYBO 1LA-RAND, Figure 6; 38, 48]. Importantly, this is observed across all tasks, emphasizing the importance of performing broad and systematic benchmarks in different real-world settings. Eventually, we evaluated the possibility to improve the strategy for acquiring x_{next} . As opposed to conventional BO, where the goal is to pick the next data point by maximizing the chance that this next point is the optimum (one-step lookahead), the partial monitoring scenario would conceptually require a different approach. Here, x_{next} should be picked to maximize the chance of recommending the optimum in the subsequent greedy pick. To account for this, we implemented and benchmarked a BO strategy building on two-step lookahead acquisition functions [CURRYBO 2LA-RAND; 22, 23, 67]. Combined with random selection of w_{next} , however, only marginal improvements in optimization performance were observed (Figure 6). At the same time, the use of different acquisition functions due to implementation constraints (see Appendix A.1.3) prevents a direct, rigorous comparison. Yet, these findings clearly suggest that the efficacy of generality-oriented optimization is largely governed by sampling w_{next} , while a data-driven acquisition of x_{next} is required for efficient optimization.

5 Conclusion

We explored the extension of global optimization frameworks for the identification of general optima, exemplified by the real-life problem of chemical reaction condition optimization. Systematic analysis

of common reaction optimization benchmarks supports the hypothesis that optimization over multiple related tasks can indeed lead to more general optima, particularly in such scenarios where the density of high-outcome experiments across the search space is low. We provide augmented versions of these benchmarks to reflect these real-life considerations. To identify general optima through iterative optimization, we find that, due to the partial observation nature of the problem, efficient exploration of the \mathcal{W} -space is crucial for the success of generality-oriented optimization, and that in our experiments, random sampling outperforms all other strategies for this task. Whilst our findings present an important step towards applying generality-oriented optimization in chemical laboratories, our study also re-emphasizes the need for further benchmark problems that are reflective of the real-world scenarios [45]. Such studies, as well as more systematic evaluations of chemical reaction representations and substrate selection policies, will be beneficial for a more principled usage of generality-oriented optimization. It should be added that within our experiments, we assume the existence of general reaction conditions that are transferable to new substrates. While we show that, within the chemical space explored in the benchmarks, these general reaction conditions exist, such an assumption may not hold as the search space further expands. In such cases, identifying multiple domains with domain-general conditions may present a more practical strategy. In addition, we want to emphasize that the identified general optima may differ significantly from the specific global optimum for each substrate. In practice, however, this is usually sufficient: In molecular discovery (e.g. drug discovery), the primary objective is to establish conditions that yield a non-zero outcome, thus enabling the evaluation of a new, unknown substrate's properties. In process optimization, general conditions can serve as a valuable starting point for further, substrate-specific optimization. Therefore, building on our results, we anticipate that generality-oriented optimization will be increasingly adopted in experimental chemistry and beyond, contributing to the development of more applicable and sustainable reactions.

References

- [1] Anonymized Repository - Anonymous GitHub. URL <https://anonymous.4open.science/r/general-bayesopt>.
- [2] RDKit: Open-source cheminformatics, 2023. URL <http://www.rdkit.org>.
- [3] Rajeev Agrawal. Sample mean based index policies by $O(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27(4):1054–1078, December 1995. ISSN 0001-8678, 1475-6064. doi: 10.2307/1427934. URL <https://www.cambridge.org/core/journals/advances-in-applied-probability/article/sample-mean-based-index-policies-by-olog-n-regret-for-the-multiarmed-bandit-problem/F79B49DC58E1070F6DFBE6F5D6DFD6FE>.
- [4] Derek T. Ahneman, Jesús G. Estrada, Shishi Lin, Spencer D. Dreher, and Abigail G. Doyle. Predicting reaction performance in C–N cross-coupling using machine learning. *Science*, 360(6385):186–190, April 2018. doi: 10.1126/science.aar5169. URL <https://www.science.org/doi/10.1126/science.aar5169>. Publisher: American Association for the Advancement of Science.
- [5] Paul Anastas, John Warner, Paul Anastas, and John Warner. *Green Chemistry: Theory and Practice*. Oxford University Press, Oxford, New York, March 2000. ISBN 978-0-19-850698-0.
- [6] Nicholas H. Angello, Vandana Rathore, Wiktor Beker, Agnieszka Wołos, Edward R. Jira, Rafał Roszak, Tony C. Wu, Charles M. Schroeder, Alán Aspuru-Guzik, Bartosz A. Grzybowski, and Martin D. Burke. Closed-loop optimization of general reaction conditions for heteroaryl Suzuki–Miyaura coupling. *Science*, 378(6618):399–405, October 2022. doi: 10.1126/science.adc8743. URL <https://www.science.org/doi/10.1126/science.adc8743>. Publisher: American Association for the Advancement of Science.
- [7] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems*, volume 33, pages 21524–21538. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/f5b1b89d98b7286673128a5fb112cb9a-Abstract.html>.

- [8] Wiktor Beker, Rafał Roszak, Agnieszka Wołos, Nicholas H. Angello, Vandana Rathore, Martin D. Burke, and Bartosz A. Grzybowski. Machine Learning May Sometimes Simply Capture Literature Popularity Trends: A Case Study of Heterocyclic Suzuki–Miyaura Coupling. *Journal of the American Chemical Society*, 144(11):4819–4827, March 2022. ISSN 0002-7863. doi: 10.1021/jacs.1c12005. URL <https://doi.org/10.1021/jacs.1c12005>. Publisher: American Chemical Society.
- [9] Isaiah O. Betinol, Junshan Lai, Saumya Thakur, and Jolene P. Reid. A Data-Driven Workflow for Assigning and Predicting Generality in Asymmetric Catalysis. *Journal of the American Chemical Society*, 145(23):12870–12883, June 2023. ISSN 0002-7863. doi: 10.1021/jacs.3c03989. URL <https://doi.org/10.1021/jacs.3c03989>. Publisher: American Chemical Society.
- [10] Alexander Buitrago Santanilla, Erik L. Regalado, Tony Pereira, Michael Shevlin, Kevin Bate-man, Louis-Charles Campeau, Jonathan Schneeweis, Simon Berritt, Zhi-Cai Shi, Philippe Nan-termet, Yong Liu, Roy Helmy, Christopher J. Welch, Petr Vachal, Ian W. Davies, Tim Cernak, and Spencer D. Dreher. Nanomole-scale high-throughput chemistry for the synthesis of complex molecules. *Science*, 347(6217):49–53, January 2015. doi: 10.1126/science.1259203. URL <https://www.science.org/doi/full/10.1126/science.1259203>. Publisher: American Association for the Advancement of Science.
- [11] Lung-Yi Chen and Yi-Pei Li. Machine Learning-Guided Strategies for Reaction Condition Design and Optimization, September 2024. URL <https://chemrxiv.org/engage/chemrxiv/article-details/66d71904cec5d6c142016bca>.
- [12] Adam D. Clayton, Jamie A. Manson, Connor J. Taylor, Thomas W. Chamberlain, Brian A. Taylor, Graeme Clemens, and Richard A. Bourne. Algorithms for the self-optimisation of chemical reactions. *Reaction Chemistry & Engineering*, 4(9):1545–1554, August 2019. ISSN 2058-9883. doi: 10.1039/C9RE00209J. URL <https://pubs.rsc.org/en/content/articlelanding/2019/re/c9re00209j>. Publisher: The Royal Society of Chemistry.
- [13] Ian W. Davies. The digitization of organic synthesis. *Nature*, 570(7760):175–181, June 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-1288-y. URL <https://www.nature.com/articles/s41586-019-1288-y>. Publisher: Nature Publishing Group.
- [14] Erik Daxberger, Anastasia Makarova, Matteo Turchetta, and Andreas Krause. Mixed-Variable Bayesian Optimization. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 2633–2639, July 2020. doi: 10.24963/ijcai.2020/365. URL <http://arxiv.org/abs/1907.01329>. arXiv:1907.01329 [cs, stat].
- [15] Peter Frazier, Warren Powell, and Savas Dayanik. The Knowledge-Gradient Policy for Correlated Normal Beliefs. *INFORMS Journal on Computing*, 21(4):599–613, November 2009. ISSN 1091-9856, 1526-5528. doi: 10.1287/ijoc.1080.0314. URL <https://pubsonline.informs.org/doi/10.1287/ijoc.1080.0314>.
- [16] Peter I. Frazier, Warren B. Powell, and Savas Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47(5):2410–2439, 2008. ISSN 0363-0129. doi: 10.1137/070693424. URL <https://collaborate.princeton.edu/en/publications/a-knowledge-gradient-policy-for-sequential-information-collection>. Publisher: Society for Industrial and Applied Mathematics Publications.
- [17] Simone Gallarati, Puck van Gerwen, Ruben Laplaza, Lucien Brey, Alexander Makaveev, and Clemence Corminboeuf. A genetic optimization strategy with generality in asymmetric organocatalysis as a primary target. *Chemical Science*, 15(10):3640–3660, 2024. doi: 10.1039/D3SC06208B. URL <https://pubs.rsc.org/en/content/articlelanding/2024/sc/d3sc06208b>. Publisher: Royal Society of Chemistry.
- [18] Krishna N. Ganesh, Deqing Zhang, Scott J. Miller, Kai Rossen, Paul J. Chirik, Marisa C. Kozłowski, Julie B. Zimmerman, Bryan W. Brooks, Phillip E. Savage, David T. Allen, and Adelina M. Voutchkova-Kostal. Green Chemistry: A Framework for a Sustainable Future. *Organic Process Research & Development*, 25(7):1455–1459, July 2021. ISSN 1083-6160.

doi: 10.1021/acs.oprd.1c00216. URL <https://doi.org/10.1021/acs.oprd.1c00216>. Publisher: American Chemical Society.

- [19] Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2023.
- [20] Tobias Gensch, Gabriel dos Passos Gomes, Pascal Friederich, Ellyn Peters, Théophile Gaudin, Robert Pollice, Kjell Jorner, AkshatKumar Nigam, Michael Lindner-D'Addario, Matthew S. Sigman, and Alán Aspuru-Guzik. A Comprehensive Discovery Platform for Organophosphorus Ligands for Catalysis. *Journal of the American Chemical Society*, 144(3):1205–1217, January 2022. ISSN 0002-7863. doi: 10.1021/jacs.1c09718. URL <https://doi.org/10.1021/jacs.1c09718>. Publisher: American Chemical Society.
- [21] Tobias Gensch, Sleight R. Smith, Thomas J. Colacot, Yam N. Timsina, Guolin Xu, Ben W. Glasspoole, and Matthew S. Sigman. Design and Application of a Screening Set for Monophosphine Ligands in Cross-Coupling. *ACS Catalysis*, 12(13):7773–7780, July 2022. doi: 10.1021/acscatal.2c01970. URL <https://doi.org/10.1021/acscatal.2c01970>. Publisher: American Chemical Society.
- [22] David Ginsbourger and Rodolphe Le Riche. Towards Gaussian Process-based Optimization with Finite Time Horizon. In Alessandra Giovagnoli, Anthony C. Atkinson, Bernard Torsney, and Caterina May, editors, *mODa 9 – Advances in Model-Oriented Design and Analysis*, pages 89–96, Heidelberg, 2010. Physica-Verlag HD. ISBN 978-3-7908-2410-0. doi: 10.1007/978-3-7908-2410-0_12.
- [23] Javier Gonzalez, Michael Osborne, and Neil Lawrence. GLASSES: Relieving The Myopia Of Bayesian Optimisation. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 790–799. PMLR, May 2016. URL <https://proceedings.mlr.press/v51/gonzalez16b.html>. ISSN: 1938-7228.
- [24] Ryan-Rhys Griffiths, Leo Klärner, Henry B. Moss, Aditya Ravuri, Sang Truong, Samuel Stanton, Gary Tom, Bojana Rankovic, Yuanqi Du, Arian Jamasb, Aryan Deshwal, Julius Schwartz, Austin Tripp, Gregory Kell, Simon Frieder, Anthony Bourached, Alex Chan, Jacob Moss, Chengzhi Guo, Johannes Durholt, Saudamini Chaurasia, Felix Strieth-Kalthoff, Alpha A. Lee, Bingqing Cheng, Alán Aspuru-Guzik, Philippe Schwaller, and Jian Tang. GAUCHE: A Library for Gaussian Processes in Chemistry, February 2023. URL <http://arxiv.org/abs/2212.04450>. arXiv:2212.04450 [cond-mat, physics:physics].
- [25] Jeff Guo, Bojana Ranković, and Philippe Schwaller. Bayesian Optimization for Chemical Reactions. *CHIMIA*, 77(1/2):31–38, February 2023. ISSN 2673-2424. doi: 10.2533/chimia.2023.31. URL https://www.chimia.ch/chimia/article/view/2023_31. Number: 1/2.
- [26] S. S. Gupta and K. J. Miescke. Bayesian look ahead one stage sampling allocations for selecting the largest normal mean. *Statistical Papers*, 35(1):169–177, December 1994. ISSN 1613-9798. doi: 10.1007/BF02926410. URL <https://doi.org/10.1007/BF02926410>.
- [27] Andreas T. Güntner, Sebastian Abegg, Karsten Königstein, Philipp A. Gerber, Arno Schmidt-Trucksäss, and Sotiris E. Pratsinis. Breath Sensors for Health Monitoring. *ACS Sensors*, 4(2):268–280, February 2019. doi: 10.1021/acssensors.8b00937. URL <https://doi.org/10.1021/acssensors.8b00937>. Publisher: American Chemical Society.
- [28] Jeremy J. Henle, Andrew F. Zahrt, Brennan T. Rose, William T. Darrow, Yang Wang, and Scott E. Denmark. Development of a Computer-Guided Workflow for Catalyst Optimization. Descriptor Validation, Subset Selection, and Training Set Analysis. *Journal of the American Chemical Society*, 142(26):11578–11592, July 2020. ISSN 0002-7863. doi: 10.1021/jacs.0c04715. URL <https://doi.org/10.1021/jacs.0c04715>. Publisher: American Chemical Society.
- [29] José Miguel Hernández-Lobato, James Requeima, Edward O. Pyzer-Knapp, and Alán Aspuru-Guzik. Parallel and Distributed Thompson Sampling for Large-scale Accelerated Exploration of Chemical Space. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1470–1479. PMLR, July 2017. URL <https://proceedings.mlr.press/v70/hernandez-lobato17a.html>. ISSN: 2640-3498.

- [30] Florian Häse, Matteo Aldeghi, Riley J. Hickman, Loïc M. Roch, Melodie Christensen, Elena Liles, Jason E. Hein, and Alán Aspuru-Guzik. Olympus: a benchmarking framework for noisy optimization and experiment planning. *Machine Learning: Science and Technology*, 2(3):035021, July 2021. ISSN 2632-2153. doi: 10.1088/2632-2153/abcdc8. URL <https://dx.doi.org/10.1088/2632-2153/abcdc8>. Publisher: IOP Publishing.
- [31] Shali Jiang, Henry Chai, Javier Gonzalez, and Roman Garnett. BINOCULARS for efficient, nonmyopic sequential experimental design. In *Proceedings of the 37th International Conference on Machine Learning*, pages 4794–4803. PMLR, November 2020. URL <https://proceedings.mlr.press/v119/jiang20b.html>. ISSN: 2640-3498.
- [32] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492, December 1998. ISSN 1573-2916. doi: 10.1023/A:1008306431147. URL <https://doi.org/10.1023/A:1008306431147>.
- [33] Leslie Pack Kaelbling. Associative Reinforcement Learning: A Generate and Test Algorithm. *Machine Learning*, 15(3):299–319, June 1994. ISSN 1573-0565. doi: 10.1023/A:1022642026684. URL <https://doi.org/10.1023/A:1022642026684>.
- [34] Leslie Pack Kaelbling. Associative Reinforcement Learning: Functions in k-DNF. *Machine Learning*, 15(3):279–298, June 1994. ISSN 1573-0565. doi: 10.1023/A:1022689909846. URL <https://doi.org/10.1023/A:1022689909846>.
- [35] Johannes Kirschner, Tor Lattimore, and Andreas Krause. Information Directed Sampling for Linear Partial Monitoring, February 2020. URL <http://arxiv.org/abs/2002.11182>. arXiv:2002.11182 [cs, stat].
- [36] Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, and Vincent Fortuin. Promises and Pitfalls of the Linearized Laplace in Bayesian Optimization, July 2023. URL <http://arxiv.org/abs/2304.08309>. arXiv:2304.08309 [cs, stat].
- [37] Agustinus Kristiadi, Felix Strieth-Kalthoff, Marta Skreta, Pascal Poupart, Alán Aspuru-Guzik, and Geoff Pleiss. A Sober Look at LLMs for Material Discovery: Are They Actually Good for Bayesian Optimization Over Molecules?, May 2024. URL <http://arxiv.org/abs/2402.05015>. arXiv:2402.05015 [cs].
- [38] Agustinus Kristiadi, Felix Strieth-Kalthoff, Sriram Ganapathi Subramanian, Vincent Fortuin, Pascal Poupart, and Geoff Pleiss. How Useful is Intermittent, Asynchronous Expert Feedback for Bayesian Optimization?, June 2024. URL <http://arxiv.org/abs/2406.06459>. arXiv:2406.06459 [cs].
- [39] Tor Lattimore. Minimax Regret for Partial Monitoring: Infinite Outcomes and Rustichini’s Regret, February 2022. URL <http://arxiv.org/abs/2202.10997>. arXiv:2202.10997 [cs, math].
- [40] Tor Lattimore and Andras Gyorgy. Mirror Descent and the Information Ratio. In *Proceedings of Thirty Fourth Conference on Learning Theory*, pages 2965–2992. PMLR, July 2021. URL <https://proceedings.mlr.press/v134/lattimore21b.html>. ISSN: 2640-3498.
- [41] Tor Lattimore and Csaba Szepesvári. An Information-Theoretic Approach to Minimax Regret in Partial Monitoring, May 2019. URL <http://arxiv.org/abs/1902.00470>. arXiv:1902.00470 [cs, math, stat].
- [42] Tor Lattimore and Csaba Szepesvári. Cleaning up the neighborhood: A full classification for adversarial partial monitoring. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, pages 529–556. PMLR, March 2019. URL <https://proceedings.mlr.press/v98/lattimore19a.html>. ISSN: 2640-3498.
- [43] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 1 edition, July 2020. ISBN 978-1-108-57140-1 978-1-108-48682-8. doi: 10.1017/9781108571401. URL <https://www.cambridge.org/core/product/identifier/9781108571401/type/book>.

- [44] Yucen Lily Li, Tim G. J. Rudner, and Andrew Gordon Wilson. A Study of Bayesian Neural Network Surrogates for Bayesian Optimization, May 2024. URL <http://arxiv.org/abs/2305.20028>. arXiv:2305.20028 [cs, stat].
- [45] Qiaohao Liang, Aldair E. Gongora, Zekun Ren, Armi Tiihonen, Zhe Liu, Shijing Sun, James R. Deneault, Daniil Bash, Flore Mekki-Berrada, Saif A. Khan, Kedar Hippalgaonkar, Benji Maruyama, Keith A. Brown, John Fisher III, and Tonio Buonassisi. Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains. *npj Computational Materials*, 7(1):1–10, November 2021. ISSN 2057-3960. doi: 10.1038/s41524-021-00656-9. URL <https://www.nature.com/articles/s41524-021-00656-9>. Publisher: Nature Publishing Group.
- [46] H. L. Morgan. The Generation of a Unique Machine Description for Chemical Structures-A Technique Developed at Chemical Abstracts Service. *Journal of Chemical Documentation*, 5(2):107–113, May 1965. ISSN 0021-9576. doi: 10.1021/c160017a018. URL <https://doi.org/10.1021/c160017a018>. Publisher: American Chemical Society.
- [47] J. Moćkus. On Bayesian Methods for Seeking the Extremum. In G. I. Marchuk, editor, *Optimization Techniques IFIP Technical Conference: Novosibirsk, July 1–7, 1974*, pages 400–404. Springer, Berlin, Heidelberg, 1975. ISBN 978-3-662-38527-2. doi: 10.1007/978-3-662-38527-2_55. URL https://doi.org/10.1007/978-3-662-38527-2_55.
- [48] William Muldrew, Peter Hayes, Mingtian Zhang, and David Barber. Active Preference Learning for Large Language Models, June 2024. URL <http://arxiv.org/abs/2402.08114>. arXiv:2402.08114 [cs].
- [49] AkshatKumar Nigam, Robert Pollice, Mario Krenn, Gabriel dos Passos Gomes, and Alán Aspuru-Guzik. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (STONED) algorithm for molecules using SELFIES. *Chemical Science*, 12(20):7079–7090, May 2021. ISSN 2041-6539. doi: 10.1039/D1SC00231G. URL <https://pubs.rsc.org/en/content/articlelanding/2021/sc/d1sc00231g>. Publisher: The Royal Society of Chemistry.
- [50] Cesar N. Prieto Kullmer, Jacob A. Kautzky, Shane W. Krska, Timothy Nowak, Spencer D. Dreher, and David W. C. MacMillan. Accelerating reaction generality and mechanistic insight through additive mapping. *Science*, 376(6592):532–539, April 2022. doi: 10.1126/science.abn1885. URL <https://www.science.org/doi/full/10.1126/science.abn1885>. Publisher: American Association for the Advancement of Science.
- [51] Siriphan Promjun and Nopadon Sahachaisaeree. Factors Determining Athletic Footwear Design: A Case of Product Appearance and Functionality. *Procedia - Social and Behavioral Sciences*, 36:520–528, January 2012. ISSN 1877-0428. doi: 10.1016/j.sbspro.2012.03.057. URL <https://www.sciencedirect.com/science/article/pii/S1877042812005241>.
- [52] Debanjan Rana, Philipp M. Pflüger, Niklas P. Hölter, Guangying Tan, and Frank Glorius. Standardizing Substrate Selection: A Strategy toward Unbiased Evaluation of Reaction Generality. *ACS Central Science*, 10(4):899–906, April 2024. ISSN 2374-7943. doi: 10.1021/acscentsci.3c01638. URL <https://doi.org/10.1021/acscentsci.3c01638>. Publisher: American Chemical Society.
- [53] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [54] Jonas Rein, Soren D. Rozema, Olivia C. Langner, Samson B. Zacate, Melissa A. Hardy, Juno C. Siu, Brandon Q. Mercado, Matthew S. Sigman, Scott J. Miller, and Song Lin. Generality-oriented optimization of enantioselective aminoxyl radical catalysis. *Science*, 380(6646):706–712, May 2023. doi: 10.1126/science.adf6177. URL <https://www.science.org/doi/10.1126/science.adf6177>. Publisher: American Association for the Advancement of Science.
- [55] Aldo Rustichini. Minimizing Regret: The General Case. *Games and Economic Behavior*, 29(1):224–243, October 1999. ISSN 0899-8256. doi: 10.1006/game.1998.0690. URL <https://www.sciencedirect.com/science/article/pii/S089982569890690X>.

- [56] Frederik Sandfort, Felix Strieth-Kalthoff, Marius Kühnemund, Christian Beecks, and Frank Glorius. A Structure-Based Platform for Predicting Chemical Reactivity. *Chem*, 6(6):1379–1390, June 2020. ISSN 2451-9294. doi: 10.1016/j.chempr.2020.02.017. URL <https://www.sciencedirect.com/science/article/pii/S2451929420300851>.
- [57] Oliver Schilter, Daniel Pacheco Gutierrez, Linnea M. Folkmann, Alessandro Castrogiovanni, Alberto García-Durán, Federico Zipoli, Loïc M. Roch, and Teodoro Laino. Combining Bayesian optimization and automation to simultaneously optimize reaction conditions and routes. *Chemical Science*, 15(20):7732–7741, May 2024. ISSN 2041-6539. doi: 10.1039/D3SC05607D. URL <https://pubs.rsc.org/en/content/articlelanding/2024/sc/d3sc05607d>. Publisher: The Royal Society of Chemistry.
- [58] Stefan P. Schmid, Leon Schlosser, Frank Glorius, and Kjell Jorner. Catalysing (organo-)catalysis: Trends in the application of machine learning to enantioselective organocatalysis. *Beilstein Journal of Organic Chemistry*, 20(1):2280–2304, September 2024. ISSN 1860-5397. doi: 10.3762/bjoc.20.196. URL <https://www.beilstein-journals.org/bjoc/articles/20/196>. Publisher: Beilstein-Institut.
- [59] Tobias Schnitzer, Martin Schnurr, Andrew F. Zahrt, Nader Sakhaee, Scott E. Denmark, and Helma Wennemers. Machine Learning to Develop Peptide Catalysts: Successes, Limitations, and Opportunities. *ACS Central Science*, 10(2):367–373, February 2024. ISSN 2374-7943. doi: 10.1021/acscentsci.3c01284. URL <https://doi.org/10.1021/acscentsci.3c01284>. Publisher: American Chemical Society.
- [60] Benjamin J. Shields, Jason Stevens, Jun Li, Marvin Parasram, Farhan Damani, Jesus I. Martinez Alvarado, Jacob M. Janey, Ryan P. Adams, and Abigail G. Doyle. Bayesian reaction optimization as a tool for chemical synthesis. *Nature*, 590(7844):89–96, February 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03213-y. URL <https://www.nature.com/articles/s41586-021-03213-y>. Publisher: Nature Publishing Group.
- [61] Felix Strieth-Kalthoff, Frederik Sandfort, Marius Kühnemund, Felix R. Schäfer, Herbert Kuchen, and Frank Glorius. Machine Learning for Chemical Reactivity: The Importance of Failed Experiments. *Angewandte Chemie International Edition*, 61(29):e202204647, 2022. ISSN 1521-3773. doi: 10.1002/anie.202204647. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/anie.202204647>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/anie.202204647>.
- [62] Connor J. Taylor, Alexander Pomberger, Kobi C. Felton, Rachel Grainger, Magda Barecka, Thomas W. Chamberlain, Richard A. Bourne, Christopher N. Johnson, and Alexei A. Lapkin. A Brief Introduction to Chemical Reaction Optimization. *Chemical Reviews*, 123(6):3089–3126, March 2023. ISSN 0009-2665. doi: 10.1021/acs.chemrev.2c00798. URL <https://doi.org/10.1021/acs.chemrev.2c00798>. Publisher: American Chemical Society.
- [63] William R. Thompson. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3/4):285–294, 1933. ISSN 0006-3444. doi: 10.2307/2332286. URL <https://www.jstor.org/stable/2332286>. Publisher: [Oxford University Press, Biometrika Trust].
- [64] Gary Tom, Stefan P. Schmid, Sterling G. Baird, Yang Cao, Kouros Darvish, Han Hao, Stanley Lo, Sergio Pablo-García, Ella M. Rajaonson, Marta Skreta, Naruki Yoshikawa, Samantha Corapi, Gun Deniz Akkoc, Felix Strieth-Kalthoff, Martin Seifrid, and Alán Aspuru-Guzik. Self-Driving Laboratories for Chemistry and Materials Science. *Chemical Reviews*, 124(16):9633–9732, August 2024. ISSN 0009-2665. doi: 10.1021/acs.chemrev.4c00055. URL <https://doi.org/10.1021/acs.chemrev.4c00055>. Publisher: American Chemical Society.
- [65] Corin C. Wagen, Spencer E. McMinn, Eugene E. Kwan, and Eric N. Jacobsen. Screening for generality in asymmetric catalysis. *Nature*, 610(7933):680–686, October 2022. ISSN 1476-4687. doi: 10.1038/s41586-022-05263-2. URL <https://www.nature.com/articles/s41586-022-05263-2>. Publisher: Nature Publishing Group.
- [66] Jason Y. Wang, Jason M. Stevens, Stavros K. Kariofillis, Mai-Jan Tom, Dung L. Golden, Jun Li, Jose E. Tabora, Marvin Parasram, Benjamin J. Shields, David N. Primer, Bo Hao, David

Del Valle, Stacey DiSomma, Ariel Furman, G. Greg Zipp, Sergey Melnikov, James Paulson, and Abigail G. Doyle. Identifying general reaction conditions by bandit optimization. *Nature*, 626 (8001):1025–1033, February 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07021-y. URL <https://www.nature.com/articles/s41586-024-07021-y>. Publisher: Nature Publishing Group.

- [67] Jian Wu and Peter Frazier. Practical Two-Step Lookahead Bayesian Optimization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://papers.nips.cc/paper_files/paper/2019/hash/2e6d9c6052e99fcdfa61d9b9da273ca2-Abstract.html.
- [68] Andrew F. Zahrt, Jeremy J. Henle, Brennan T. Rose, Yang Wang, William T. Darrow, and Scott E. Denmark. Prediction of higher-selectivity catalysts by computer-driven workflow and machine learning. *Science*, 363(6424):eaau5631, January 2019. doi: 10.1126/science.aau5631. URL <https://www.science.org/doi/10.1126/science.aau5631>. Publisher: American Association for the Advancement of Science.

A Appendix

A.1 Bayesian Optimization for Generality

A.1.1 Aggregation Functions

The aggregation function is a user-defined property that determines how the ‘‘set optimum’’ is calculated across objective functions. Through the choice of the set optimum, prior knowledge and preferences about the specific optimization problem at hand can be included. In this work, the following aggregation functions are evaluated:

Mean Aggregation

$$\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W}) = \frac{1}{|\mathcal{W}|} \sum_{\mathbf{w} \in \mathcal{W}} f(\mathbf{x}; \mathbf{w}) = \frac{1}{n} \sum_{i=1}^n f(\mathbf{x}; \mathbf{w}_i) \quad (3)$$

Threshold Aggregation

$$\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W}) = \sum_{\mathbf{w} \in \mathcal{W}} \sigma(f(\mathbf{x}; \mathbf{w}) - f_{\text{thr}}) = \sum_{i=1}^n \sigma(f(\mathbf{x}; \mathbf{w}_i) - f_{\text{thr}}) \quad (4)$$

Conceivably, other aggregation functions also have practical use-cases, for example:

Mean Squared Error (MSE) Aggregation

$$\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W}) = -\frac{1}{|\mathcal{W}|} \sum_{\mathbf{w} \in \mathcal{W}} (f_{\text{opt}}(\mathbf{x}; \mathbf{w}) - f(\mathbf{x}; \mathbf{w}))^2 = -\frac{1}{n} \sum_{i=1}^n (f_{\text{opt},i} - f(\mathbf{x}; \mathbf{w}_i))^2 \quad (5)$$

Minimum Aggregation

$$\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W}) = \min_{\mathbf{w}_i \in \mathcal{W}} f(\mathbf{x}; \mathbf{w}_i) \quad (6)$$

The above definitions assume that all $f(\mathbf{x}; \mathbf{w}_i)$ have the same range, and that the optimization problem is formulated as maximization problem.

A.1.2 Acquisition Functions and the Sample Average Approximation

For the evaluation of posterior distributions, and the calculation of acquisition function values, we use the sample-average approximation, as introduced by Balandat et al. [7]. From a posterior distribution at time point k $p(g_k(\mathbf{x}))$, M posterior samples $\zeta_m(\mathbf{x}) \sim p(g_k(\mathbf{x}))$ are drawn. These posterior samples can be used to estimate the posterior distribution, and to calculate acquisition function values as expectation values \mathbb{E}_M over all M samples.

Herein, we use the following common acquisition functions:

- Upper Confidence Bound: $\text{UCB}(\mathbf{x}) = \mathbb{E}_M(\zeta_m(\mathbf{x})) + \beta \cdot \mathbb{E}_M(\zeta_m(\mathbf{x}) - \mathbb{E}_M(\zeta_m(\mathbf{x})))$.
- Thompson Sampling: $\text{TS}(\mathbf{x}) = \zeta(\mathbf{x})$ for the case of $M = 1$.
- Posterior Variance: $\text{PV}(\mathbf{x}) = \mathbb{E}_M(\zeta_m(\mathbf{x}) - \mathbb{E}_M(\zeta_m(\mathbf{x})))$, where $\sigma(\mathbf{x})$ describes the standard deviation over all k posterior samples.
- Random Selection, where the acquisition function value is a random number.

Moreover, we evaluate the optimization performance using a primitive implementation of two-step lookahead acquisition functions α^* (see Algorithm 2). The acquisition function value of α^* at a location \mathbf{x}_0 is estimated as follows: For each of the M posterior samples $\zeta_m(\mathbf{x}_0) \sim p(g_k(\mathbf{x}_0))$, a fantasy posterior distribution $p'(g_{k+1}(\mathbf{x}_0))$ is generated by conditioning the posterior on the new observation $(\mathbf{x}_0, \zeta_M(\mathbf{x}_0))$. From this fantasy posterior distribution, the values of the inner acquisition function α' can be computed and optimized over $\mathbf{x} \in \mathcal{X}$. The final value of the two-step lookahead acquisition function is returned as $\alpha^*(\mathbf{x}_0) = \max_{\mathbf{x} \in \mathcal{X}} \alpha'(\mathbf{x})$.

Algorithm 2 Primitive implementation of a two-step lookahead acquisition function using the sample average approximation.

Input:

- input space \mathcal{X}
 - location \mathbf{x}_0 at which to evaluate the two-step lookahead acquisition function
 - posterior distribution $p(g_k(\mathbf{x}) \mid \mathcal{D})$
 - one-step lookahead acquisition function $\alpha(\mathbf{x})$
- 1: draw M posterior samples $\zeta_m(\mathbf{x}_0) \sim p(g_k(\mathbf{x}_0))$
 - 2: empty set of fantasy acquisition function values $\mathcal{A} = \{\}$
 - 3: **for** $m = 1, \dots, M$ **do**
 - 4: compute fantasy posterior $p'(\mathbf{x}) = p(g_{k+1}(\mathbf{x}) \mid (\mathcal{D} \cup (\mathbf{x}_0, \zeta_m(\mathbf{x}_0))))$
 - 5: optimize one-step-lookahead acquisition function $\alpha_m = \max_{\mathbf{x} \in \mathcal{X}} \alpha(p'(\mathbf{x}))$
 - 6: update $\mathcal{A} = \mathcal{A} \cup \{\alpha_m\}$
 - 7: **end for**
 - 8: **return** $\alpha^*(\mathbf{x}_0) = \frac{1}{M} \sum_{m=1}^M \alpha_m$
-

A.1.3 Benchmarked optimization strategies for selecting \mathbf{x}_{next} and \mathbf{w}_{next}

Herein, we outline the use of the benchmarked optimization strategies for generality-oriented optimization. The discussed optimization strategies describe different variations of how to pick the next experiments \mathbf{x}_{next} and \mathbf{w}_{next} .

Following the SAA outlined above, we estimate the predictive posterior distribution $p(\phi(\mathbf{x}) \mid \mathcal{D})$ as follows: For each $\mathbf{w}_i \in \mathcal{W}$, M (typically $M = 512$) samples $\zeta_{im}(\mathbf{x}) \sim p(g_k(\mathbf{x}, \mathbf{w}_i))$ are drawn from the posterior distribution of the surrogate model. Aggregating over all \mathbf{w}_i yields M samples $\zeta_m(\mathbf{x}) \sim p(\phi(\mathbf{x}) \mid \mathcal{D})$ from the posterior distribution over $\phi(\mathbf{x})$, which can be used for calculating the acquisition function values using the sample-based acquisition function logic, as described in Appendix A.1.2. With this, we implement and benchmark the acquisition policies in Table 2.

The sequential acquisition is described in Algorithm 3 and refers to a strategy in which \mathbf{x}_{next} and \mathbf{w}_{next} are selected sequentially. In the first step, \mathbf{x}_{next} is selected by optimizing an \mathbf{x} -specific acquisition function α_x over $\mathbf{x} \in \mathcal{X}$. With the selected \mathbf{x}_{next} in hand, \mathbf{w}_{next} is then selected by optimizing an independent, \mathbf{w} -specific acquisition function over $\mathbf{w} \in \mathcal{W}$. With $\alpha_x = \text{PI}$ (Probability of Improvement) and $\alpha_w = \text{PV}$, this would correspond to the strategy described in [6].

Algorithm 3 Sequential Acquisition Policy

Input:

- posterior distribution $p(g_k(\mathbf{x}, \mathbf{w}) \mid \mathcal{D})$
 - aggregation function $\phi(f(\mathbf{x}; \mathbf{w}), \mathcal{W})$
 - acquisition function α_x
 - acquisition function α_w
- 1: compute posterior distribution $p(\phi(\mathbf{x}) \mid \mathcal{D}) = p(\phi(g_k(\mathbf{x}, \mathbf{w}), \mathcal{W}) \mid \mathcal{D})$
 - 2: acquire $\mathbf{x}_{\text{next}} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \alpha_x(p(\phi(\mathbf{x}) \mid \mathcal{D}))$
 - 3: acquire $\mathbf{w}_{\text{next}} = \operatorname{argmax}_{\mathbf{w} \in \mathcal{W}} \alpha_w(p(g_k(\mathbf{x}_{\text{next}}, \mathbf{w}) \mid \mathcal{D}))$
 - 4: **return** $\mathbf{x}_{\text{next}}, \mathbf{w}_{\text{next}}$
-

Table 2: Benchmarked optimization strategies

Optimization strategy	Description
CURRYBO 2LA-RAND	Sequential two-step lookahead acquisition policy, using Thompson Sampling acquisition function to select \mathbf{x}_{next} and a Random selection of \mathbf{w}_{next} with a greedy final recommendation
CURRYBO 1LA-VAR	Sequential acquisition policy, using an UCB acquisition function with $\beta = 0.5$ to select \mathbf{x}_{next} and a posterior variance acquisition function to select \mathbf{w}_{next} with a greedy final recommendation. This strategy resembles the algorithm described by Angello et al. [6]
CURRYBO 1LA-RAND	Sequential acquisition policy, using an UCB acquisition function with $\beta = 0.5$ to select \mathbf{x}_{next} and a random acquisition of \mathbf{w}_{next} with a greedy final recommendation.
BANDIT	Multi-armed bandit algorithm as outlined and implemented by Wang et al. [66]
CURRYBO 2LA SINGLE SUBSTRATE	Sequential two-step lookahead acquisition policy, using Thompson Sampling to select \mathbf{x}_{next} . For \mathbf{w}_{next} , always the same substrate is selected.
CURRYBO 2LA COMPLETE MONITORING	Sequential two-step lookahead acquisition policy, using Thompson Sampling to select \mathbf{x}_{next} . To select \mathbf{w}_{next} , all substrates are selected in a complete monitoring scenario. After every experiment, the a greedy final recommendation is performed.
CURRYBO 1LA SINGLE SUBSTRATE	Sequential acquisition policy, using Thompson Sampling to select \mathbf{x}_{next} . For \mathbf{w}_{next} , always the same substrate is selected.
CURRYBO 1LA COMPLETE MONITORING	Sequential acquisition policy, using Thompson Sampling to select \mathbf{x}_{next} . To select \mathbf{w}_{next} , all substrates are selected in a complete monitoring scenario. After every experiment, the a greedy final recommendation is performed.
RANDOM GREEDY	Sequential acquisition policy, using a random acquisition function to select \mathbf{x}_{next} and a random acquisition function to select \mathbf{w}_{next} with a greedy final recommendation
RANDOM	Random acquisition policy of \mathbf{x}_{next} and \mathbf{w}_{next} and final random recommendation

A.2 Benchmark Problem Details

A.2.1 Original Benchmark Problems

Three chemical reaction benchmarks have been considered in this work: Product quantity optimization for Pd-catalyzed C–heteroatom couplings [10], and enantioselectivity optimization for a N,S-Acetal formation [68]. In the following, the benchmark problems are described briefly.

Pd-catalyzed carbon-heteroatom coupling

The Pd-catalyzed carbon-heteroatom coupling benchmark is concerned with the reaction of different nucleophiles with 3-bromopyridine (Figure 7). In total, 16 different nucleophiles were tested in a nanoscale high-throughput experimentation platform. As reaction conditions, bases (six different bases) and catalysts (16 different catalysts) were varied. In total, the benchmark consists of 1536 different experiments, for which the conversion is reported.

The average conversion is 2.05%, whereas the maximum conversion is 39.81% (Figure 8). The average of the average conversion of each condition is 2.05%, while the maximum of the average conversion of the catalysts is 7.60% (Figure 8). The catalyst-base combination with the highest average conversion is shown in Figure 8.

With respect to the threshold aggregation function, the chosen threshold was 7.50%. The average number of substrates with a conversion above this threshold are 1.615, while the maximum number




Figure 7: Reaction diagram of the Pd-catalyzed carbon-heteroatom coupling, where 3-bromopyridine reacts with a nucleophile. Reaction conditions include a catalyst and a base. The numbers indicate the amount of different species in the benchmark.




Figure 8: Top left: Distribution of the conversion for the Pd-catalyzed carbon-heteroatom coupling in the original benchmark. Top right: Distribution of the average conversion for each catalyst-base combination for the Pd-catalyzed carbon-heteroatom coupling in the original benchmark. Bottom: Catalyst-base combination with the highest average conversion in the original benchmark.

of substrates is 7 (Figure 9). The catalyst-base combination with the highest number of substrates with a conversion above the threshold is the same as shown in Figure 8.




Figure 9: Left: Distribution of the conversion for the Pd-catalyzed carbon-heteroatom coupling in the original benchmark. Right: Distribution of the number of substrates with a conversion above the specified threshold for each catalyst-base combination for the Pd-catalyzed carbon-heteroatom coupling in the original benchmark.

N,S-Acetal formation

The N,S-Acetal formation benchmark is concerned with the nucleophilic addition of different thiols to imines, catalyzed by chiral phosphoric acids (CPAs) (see Figure 10). In total, five different imines and five different thiols were tested in manual experiments. As reaction conditions, 43 different CPA catalysts were considered. In total, the benchmark consists of 1075 different experiments, for which $\Delta\Delta G^\ddagger$, as a measure of the enantioselectivity, is reported.




Figure 10: Reaction diagram of the N,S-Acetal formation, where an imine reacts with a thiol. Reaction conditions include a catalyst. The numbers indicate the amount of different species in the benchmark.

The average $\Delta\Delta G^\ddagger$ is 0.988 kcal/mol, whereas the maximum $\Delta\Delta G^\ddagger$ is 3.135 kcal/mol (see Figure 11). The average of the average $\Delta\Delta G^\ddagger$ for each condition is 0.988 kcal/mol, while the maximum of the average $\Delta\Delta G^\ddagger$ for all conditions is 2.395 kcal/mol (see Figure 11). The catalyst with the highest average $\Delta\Delta G^\ddagger$ is shown in Figure 11.




Figure 11: Top left: Distribution of $\Delta\Delta G^\ddagger$ for the N,S-Acetal formation in the original benchmark. Top right: Distribution of the average $\Delta\Delta G^\ddagger$ for each catalyst for the N,S-Acetal formation in the original benchmark. Bottom: Catalyst with the highest average $\Delta\Delta G^\ddagger$ in the original benchmark.

With respect to the threshold aggregation function, the chosen threshold was 2.0 kcal/mol. The average number of substrates with $\Delta\Delta G^\ddagger$ above this threshold are 1.907, while the maximum number of substrates is 17 (Figure 12). The catalyst with the highest number of substrates with $\Delta\Delta G^\ddagger$ above the threshold is the same as shown in Figure 11.




Figure 12: Left: Distribution of $\Delta\Delta G^\ddagger$ for the N,S-Acetal formation in the original benchmark. Right: Distribution number of substrates with a $\Delta\Delta G^\ddagger$ above the specified threshold for each catalyst for the N,S-Acetal formation in the original benchmark.

A.2.2 Augmentation

Since the described benchmarks consist of a high number of high-outcome experiments (the respective search spaces were rationally designed by expert chemists), we augment them with more negative examples to make them more relevant to real-world optimization campaigns. New substrates are generated by mutating the originally reported substrates via the STONED algorithm [49]. In a first filtering step, new substrates were removed if they had a Tanimoto similarity to the original substrate smaller than 0.75 and if they did not possess the functional groups required for the reaction. To ensure that the benchmark is augmented with negative examples, random forests are fitted to the original benchmarks (see above). The mean absolute errors of the random forest regressors fitted to the original benchmarks are shown in Table 3. Newly generated substrates were incorporated if the average reaction outcome over all reported reaction conditions is below a defined threshold. For the Pd-catalyzed carbon-heteroatom coupling, and N,S-Acetal formation the thresholds are 1.0%, and 0.7 kcal/mol, respectively. If a substrate passed these filters, the reactions with all different reported conditions were added, with reaction outcomes being taken from as predicted from the random forest emulator.

Table 3: MAE and MSE of fitted random forest regressors to the original benchmark problems.

Benchmark problem	MAE	MSE
Pd-catalyzed carbon-heteroatom coupling	3.16×10^{-3}	7.65×10^{-5}
N,S-Acetal formation	4.95×10^{-2} kcal/mol	5.46×10^{-3} kcal/mol

A.2.3 Augmented Benchmark Problems

Pd-catalyzed carbon-heteroatom coupling

Augmentation increases the number of different nucleophiles from 16 to 31 (see Figure 13). Combined with the 96 reported reaction condition combinations, the augmented dataset consists of 2976 reactions, for which the conversion is reported.



Figure 13: Reaction diagram of the Pd-catalyzed carbon-heteroatom coupling, where 3-bromopyridine reacts with a nucleophile. Reaction conditions include a catalyst and a base. The numbers indicate the amount of different species in the augmented benchmark.

Augmentation decreased the average conversion from 2.05% to 1.34%, whereas the maximum conversion remained the same at 39.81% (see Figure 14). The average of the average conversion of each condition is decreased from 2.05% to 1.34%, and the maximum of the average conversion of

each condition is also decreased from 7.60% to 6.00% (see Figure 14). The catalyst-base combination with the highest average conversion is unaffected by the augmentation.




Figure 14: Top left: Distribution of the conversion for the Pd-catalyzed carbon-heteroatom coupling in the augmented benchmark. Top right: Distribution of the average conversion for each catalyst-base combination for the Pd-catalyzed carbon-heteroatom coupling in the augmented benchmark. Bottom: Catalyst-base combination with the highest average conversion in the augmented benchmark.

With respect to the threshold aggregation function, the chosen threshold was 7.50%. The average number of substrates with a conversion above this threshold are 1.646, while the maximum number of substrates is 8 (Figure 15). The catalyst-base combination with the highest number of substrates with a conversion above the threshold is the same as shown in Figure 14.




Figure 15: Left: Distribution of the conversion for the Pd-catalyzed carbon-heteroatom coupling in the augmented benchmark. Right: Distribution of the number of substrates with a conversion above the specified threshold for each catalyst-base combination for the Pd-catalyzed carbon-heteroatom coupling in the augmented benchmark.

N,S-Acetal formation

Augmentation increases the number of thiols from five to 13, while the number of imines remained constant at five (see Figure 16). Combined with the 43 reported reaction conditions, the augmented benchmark consists of 2795 reactions, for which $\Delta\Delta G^\ddagger$ is reported.

Augmentation decreased the average $\Delta\Delta G^\ddagger$ from 0.988 kcal/mol to 0.757 kcal/mol, whereas the maximum $\Delta\Delta G^\ddagger$ was slightly decreased from 3.135 kcal/mol to 3.114 kcal/mol (see Figure 17). This decrease is due to the fact that values are taken as predicted by the random forest emulator. Through augmentation, the average of the average $\Delta\Delta G^\ddagger$ of each condition decreased from 0.988 kcal/mol to 0.757 kcal/mol, while the maximum of the average $\Delta\Delta G^\ddagger$ of all conditions decreased as well from




Figure 16: Reaction diagram of the N,S-Acetal formation, where an imine reacts with a thiol. Reaction conditions include a catalyst. The numbers indicate the amount of different species in the augmented benchmark.

2.395 kcal/mol to 1.969 kcal/mol (see Figure 17). The catalyst with the highest average $\Delta\Delta G^\ddagger$ is unaffected by the augmentation.




Figure 17: Top left: Distribution of $\Delta\Delta G^\ddagger$ for the N,S-Acetal formation in the augmented benchmark. Top right: Distribution of the average $\Delta\Delta G^\ddagger$ for each catalyst for the N,S-Acetal formation in the augmented benchmark. Bottom: Catalyst with the highest average $\Delta\Delta G^\ddagger$ in the augmented benchmark.

With respect to the threshold aggregation function, the chosen threshold was 2.0 kcal/mol. The average number of substrates with $\Delta\Delta G^\ddagger$ above this threshold are 1.814, while the maximum number of substrates is 16 (Figure 18). The catalyst with the highest number of substrates with $\Delta\Delta G^\ddagger$ above the threshold is the same as shown in Figure 17.




Figure 18: Left: Distribution of $\Delta\Delta G^\ddagger$ for the N,S-Acetal formation in the augmented benchmark. Right: Distribution number of substrates with a $\Delta\Delta G^\ddagger$ above the specified threshold for each catalyst for the N,S-Acetal formation in the augmented benchmark.

A.3 Grid Search for Analyzing Benchmark Problems

To analyse the utility of considering multiple substrates in an optimization campaign, we performed exhaustive grid search on the described benchmark problems. For each problem, the substrates were split into an initial train and test set among the substrates. In total, thirty different train/test splits were performed. The obtained train set was further subsampled into smaller training sets with varying sizes to investigate the influence on the number of substrates. Sampling among the substrates in the train set was performed either through random sampling, random sampling or "Average Sampling", where the required number of substrates was chosen as the substrates with the highest average Tanimoto similarity to all other train substrates. For each subsampled training set, the most general conditions were identified. The general reaction outcome, as specified by the aggregation function, is evaluated for these conditions on the held-out test set. Further, this general reaction outcome was scaled from 0 to 1, where 0 is the worst possible general reaction outcome for the given test set and 1 is the best possible general reaction outcome for the test set. Hence, this score should be maximized. For the different benchmark problems, we report this generality score, where we also compare the behaviour of the original and augmented problems. Below, the results of the described data analysis are shown for the benchmark problems not shown in the main text.

A.4 Details on BO for generality benchmarking

To identify whether BO for generality, as described above, can efficiently identify the general optima, we conducted several benchmarking runs on the described benchmark problems. On each problem, we perform benchmarking for multiple optimization strategies, as listed in Table 2.

In each optimization campaign, we used a single-task GP regressor, as implemented in *GPyTorch*, with a TanimotoKernel as implemented in *Gauche* [24]. Molecules were represented using Morgan Fingerprints [46] with 1024 bits and a radius of 2. Fingerprints were generated using RDKit [2]. It is notable that, while such a representation was chosen due to its suitability for broad chemical spaces, more specific representations such as descriptors might be able to improve the optimization performance.

The acquisition policies were benchmarked on all benchmark problems with differently sampled substrates for the optimization. For each benchmark, we selected the train set randomly, consisting of twelve nucleophiles in the Pd-catalyzed carbon-heteroatom coupling benchmark, three imines and three thiols in the N,S-Acetal formation benchmark. Thirty independent optimization campaigns were performed for each. The generality of the proposed general conditions at each step during the optimization is shown.

A.5 Details on bandit algorithm benchmarking

The benchmarking of BANDIT [66] was performed across the benchmark problems using their proposed UCB1TUNED algorithm with differently sampled substrates for the optimization. For each benchmark, we selected the train set randomly, consisting of twelve nucleophiles in the Pd-catalyzed carbon-heteroatom coupling benchmark, and three imines and three thiols in the N,S-Acetal formation benchmark. Thirty independent optimization campaigns were performed for each. To ensure fair comparison with CURRYBO-based strategies, the ground truth was set to be the proxy function calculated for each dataset. To select the optimum x value at each step t , we relied on the authors definition of the best arm as the most sampled arm at step t .

A.6 Additional Results and Discussion

A.6.1 Additional results on the dataset analysis for utility of generality-oriented optimization

In addition to analysing the utility of generality-oriented optimization for ϕ as the Mean Aggregation, which is shown in Figure 3, we also performed a similar analysis for ϕ as the Threshold aggregation, where the chosen thresholds are as described in Appendix A.2. The results of this analysis is shown in Figure 19. Similar to the case where ϕ is the mean aggregation, we observed that in the case of the augmented benchmark, generality-oriented optimization yields more general reaction conditions. While in the case of the N,S-Acetal formation benchmarks the obtained generality score decreases slightly compared to the original benchmark, the generality score remains high in both cases. For




Figure 19: *Top*: Generality score as determined by exhaustive grid search for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems and ϕ as the threshold aggregation. Average and standard error are taken from thirty different train/test substrates splits. *Bottom*: Generality score as determined by exhaustive grid search for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems and ϕ as the threshold aggregation. Average and standard error are taken from thirty different train/test substrates splits.

the Pd-catalyzed carbon-heteroatom benchmark, we notice worse generalizability scores than for the mean aggregation, which demonstrates that it is harder to obtain general reaction conditions for the threshold aggregation.

Furthermore, we studied how different sampling techniques among the train set substrates influence the obtained generality scores. As sampling techniques, we used random sampling, farthest point sampling and “average sampling”, as outlined in Appendix A.3. For ϕ as the mean aggregation, the results for the two different benchmarks are shown in Figure 20 and Figure 21 for the Pd-catalyzed carbon-heteroatom coupling benchmark and the N,S-Acetal formation benchmark, respectively. For ϕ as the threshold aggregation, the results for the two different benchmarks are shown in Figure 22 and Figure 23 for the Pd-catalyzed carbon-heteroatom coupling benchmark and the N,S-Acetal formation benchmark, respectively. Throughout the different benchmarks and aggregation functions, we observe that the generality score obtained through using the sampled train substrates are highly similar and no method clearly outperforms the others. It is particularly notable that farthest point sampling did not outperform other sampling techniques, as this strategy is commonly used to select chemicals to broadly cover chemical space [20, 21, 28, 59]. We hypothesize that this method insensitivity is due to the low number of substrates for chosen for the train set, which was chosen to still reflect realistic experimental cases.

A.6.2 Additional results on the benchmarking for the Mean Aggregation function

In addition to the experiments shown in the main text, we benchmarked CURRYBO 1LA-VAR, CURRYBO 1LA-RAND and CURRYBO 2LA-RAND against their analogons when a single substrate is selected (Figure 24) as well as the complete monitoring case (Figure 25). For CURRYBO 1LA-VAR and CURRYBO 1LA-RAND, we observe a more efficient optimization behaviour when the generality-oriented optimization considers multiple substrates, as shown in Figure 24 and Figure 27. For CURRYBO 2LA-RAND, we observe similar performance. This is driven by a very efficient optimization of the two-step lookahead strategy for the single substrate in the N,S-Acetal formation reaction benchmark. Further, as single substrates already show a very high generality score for the N,S-Acetal formation benchmark (see Figure 3), this results in a superior performance of single substrate generality-optimization. However, due to this dependence on the benchmark, we expect that this behaviour would change with more available benchmarks.

Figure 25 demonstrates a clear advantage of generality-oriented optimization, which is equipped to deal with the partial monitoring scenario, against a complete monitoring baseline. We hypothesize that this is because more conditions (values of x) can be evaluated, as not every substrate (which




Figure 20: Generality score as determined by exhaustive grid search for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems for the mean aggregation as ϕ . Average and standard error are taken from thirty different train/test substrates splits.

is 9 and 12 for the two benchmarks) has to be tested for a specific set of reaction conditions. This underlines the utility of improved and efficient decision-making algorithms in complex optimization scenarios.




Figure 21: Generality score as determined by exhaustive grid search for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems for the mean aggregation as ϕ . Average and standard error are taken from thirty different train/test substrates splits.

A.6.3 Additional results on the benchmarking for the Threshold Aggregation function

Whereas data-driven approaches have considered the average yield as the only generality metric [6, 66], the need for multiple definitions for generality has been emphasized [9]. Such a flexible definition allows to find the conditions that fulfill the need for specific real-world applications. To showcase the flexibility of generality-oriented optimization as outlined in this work, we implemented several aggregation functions (see Appendix A.1.1). Herein, we perform benchmarking on the Threshold Aggregation function, that strives to find conditions that provide a desirable reaction outcome (i.e. above the threshold), for the highest number of substrates. The benchmarked optimization strategies are as described in Table 2 and already demonstrated for the Mean Aggregation.

In general, we notice a highly similar optimization performance for the CURRYBO 1LA-RAND, CURRYBO 1LA-VAR and CURRYBO 2LA-RAND strategies. These results align with observations for the mean aggregation function, as discussed in Section 4. This shows a strong performance of BO-based strategies in generality-oriented optimization, as efficient optimizations are observed for multiple generality metrics. Similar to optimization for the mean aggregation, CURRYBO 1LA-RAND outperforms the optimization when only considering a single substrate, while it does not outperform CURRYBO 2LA-RAND. Further, both CURRYBO 1LA-VAR and CURRYBO 2LA-RAND outperform generality-oriented optimization with complete monitoring. This again underlines the importance of tailored decision-making algorithms in complex tasks.




Figure 22: Generality score as determined by exhaustive grid search for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems for the threshold aggregation as ϕ . Average and standard error are taken from thirty different train/test substrates splits.




Figure 23: Generality score as determined by exhaustive grid search for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems for the threshold aggregation as ϕ . Average and standard error are taken from thirty different train/test substrates splits.




Figure 24: Top: GAP metric [31] for the generality-oriented optimization over all train substrates, averaged over all benchmark sets on the original (left) and augmented (right) problems. Middle: GAP metric [31] for the generality-oriented optimization over all train substrates for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems. Bottom: GAP metric [31] for the generality-oriented optimization over all train substrates for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems. Optimization strategies are described in Table 1.




Figure 25: Top: GAP metric [31] for the generality-oriented optimization over all train substrates, averaged over all benchmark sets on the original (left) and augmented (right) problems. Middle: GAP metric [31] for the generality-oriented optimization over all train substrates for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems. Bottom: GAP metric [31] for the generality-oriented optimization over all train substrates for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems. Optimization strategies are described in Table 1.




Figure 26: GAP metric [31] for the generality-oriented optimization over all train substrates, averaged over all benchmark sets on the original (left) and augmented (right) problems. The threshold aggregation was considered as a generality metric. Optimization strategies are described in Table 1.




Figure 27: Generality-oriented optimization trajectories of different algorithms, averaged over all benchmark problems on the original (left) and augmented (right) search spaces with the threshold aggregation. Optimization algorithms are described in Table 1.




Figure 28: Top: GAP metric [31] for the generality-oriented optimization over all train substrates, averaged over all benchmark sets on the original (left) and augmented (right) problems. The threshold aggregation was considered as a generality metric. Middle: GAP metric [31] for the generality-oriented optimization over all train substrates for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems. The threshold aggregation was considered as a generality metric. Bottom: GAP metric [31] for the generality-oriented optimization over all train substrates for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems. The threshold aggregation was considered as a generality metric. Optimization algorithms are described in Table 2.




Figure 29: Top: GAP metric [31] for the generality-oriented optimization over all train substrates, averaged over all benchmark sets on the original (left) and augmented (right) problems. The threshold aggregation was considered as a generality metric. Middle: GAP metric [31] for the generality-oriented optimization over all train substrates for the Pd-catalyzed carbon-heteroatom coupling benchmark on the original (left) and augmented (right) problems. The threshold aggregation was considered as a generality metric. Bottom: GAP metric [31] for the generality-oriented optimization over all train substrates for the N,S-Acetal formation benchmark on the original (left) and augmented (right) problems. The threshold aggregation was considered as a generality metric. Optimization algorithms are described in Table 2.