Cost-Informed Bayesian Reaction Optimization

Alexandre A. Schoepfer,[†,‡,¶,§] Jan Weinreich,[†,¶,§] Ruben Laplaza,[†,¶] Jerome Waser,[*,‡,¶] and Clemence Corminboeuf[*,†,¶]

†*Laboratory for Computational Molecular Design, Institute of Chemical Sciences and Engineering, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland*

‡*Laboratory of Catalysis and Organic Synthesis, Institute of Chemical Sciences and Engineering, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland*

¶*National Center for Competence in Research-Catalysis (NCCR-Catalysis), École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

§*These authors contributed equally to this work*

E-mail: jerome.waser@epfl.ch; clemence.corminboeuf@epfl.ch

Abstract

Bayesian optimization (BO) is an increasingly popular method for optimization and development of chemical reactions. Although effective in guiding experimental design, BO does not account for experimentation costs: testing readily available reagents under different conditions might be more cost and time-effective than synthesizing or buying additional ones. To address this issue, we present cost-informed BO (CIBO), an approach tailored for the rational planning of chemical experimentation that prioritizes the most cost-effective experiments. Reagents are used only when their anticipated improvement in reaction performance sufficiently outweighs their costs. Our algorithm tracks the available reagents, including recently acquired ones, and dynamically updates their cost during the optimization. Using literature data of Pd-catalyzed reactions, we show that CIBO reduces the cost of reaction optimization by up to 90% compared to standard BO. Our approach is compatible with any type of cost, *e.g.*, the cost of buying equipment or compounds, waiting time, and environmental or security concerns. We believe CIBO supersedes BO in chemistry and envision applications in both traditional and self-driving laboratories for experiment planning.

1 Introduction

Reaction optimization is a challenging task that is often tackled "one factor at a time" by sequentially optimizing individual parameters such as catalyst, temperature, or additives. While this strategy simplifies the problem significantly, it remains time and resource-intensive and might disregard promising combinations of parameters (*e.g.*, an additive and ligand that were discarded for their lacking individual performance may yield optimal results when combined).

As an alternative, data-driven computational tools, such as machine learning (ML), have recently been used to guide experimental efforts towards the best possible performance by predicting reaction yield or selectivity from substrates, catalysts, and reaction conditions.[1–6]

Among the different ML frameworks, Bayesian optimization (BO) is ideally suited for this task.[7,8] Given some initial data, BO leverages predictions and their corresponding uncertainties to suggest the next most promising experiments to conduct. BO-driven reaction optimization has seen significant success in the last few years, especially in the automated laboratory and high-throughput experimentation (HTE) setting.[3,9–16] Therein, all necessary materials (*i.e.*, substrates, catalysts, additives, solvents) to be considered are typically procured prior to experimentation, and BO is used to find the best reagents and reaction conditions.[1,3,17,18]

Yet, the implementation of BO and other ML frameworks in traditional laboratories is still limited.[19–22] In this setting, defining and acquiring all necessary materials for the optimization beforehand is not ideal, especially when dealing with unexplored chemistry. Furthermore, classical BO methods usually attribute the same cost to all suggested experiments. In reaction optimization, where *e.g.*, catalyst ligands and reaction conditions have to be adjusted simultaneously, this assumption is unsuitable. Depending on whether a ligand is already available in the laboratory, commercially available, reported in the literature, or has never been synthesized before, the cost of an experiment (in terms of money, time investment, or risk) is significantly different. Thus, the experiments suggested by BO may not be practical or even feasible. Testing a known, available ligand with different reaction conditions may yield a comparable improvement at a much lower cost.

To overcome this limitation, here we introduce cost-informed Bayesian optimization (CIBO), a BO framework that incorporates cost into the decision-making process for practical and rational batch experimentation planning. Note that we do not aim to optimize the overall cost of individual reactions nor to include constraints, as pursued in other works,[16,18,23,24]

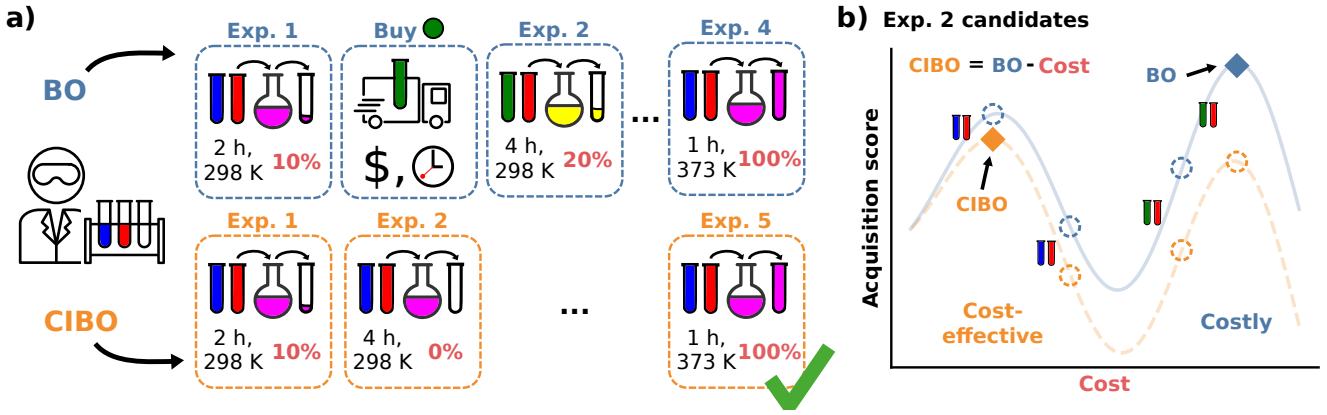

Figure 1: Overview of standard BO (blue) *vs.* cost-informed Bayesian optimization (CIBO, orange) for yield optimization. (a): BO recommends purchasing more materials. Meanwhile, CIBO balances purchases with their expected improvement of the experiment, at the cost of performing more experiments (here five *vs.* four). (b): A closer look at the two acquisition functions of BO and CIBO for the selection of experiment two. In CIBO, the BO acquisition function is modified to account for the cost by subtracting the latter. Following the blue BO curve, the next experiment to perform uses green and red reactants (corresponding to the costly maximum on the right). Subtracting the price of the experiments results in the orange CIBO curve, which instead suggests the more cost-effective experiment on the left (blue and red reactants).

but rather expedite reaction optimization by proposing experiments that are as informative and promising, yet as cost-effective, as possible. CIBO's acquisition function prioritizes experiments with a high benefit-to-cost ratio, enabling minimization of both the number of experiments to carry out *and* their total cost (Figure 1). Moreover, CIBO keeps a digital inventory that tracks what is available in the laboratory at all times, which is used to update the cost of experiments at each iteration (*e.g.*, once a ligand is bought, it might be available for several experiments).

Previous methods that account for costs in BO include contextual BO,[25] addressing environmental effects, as well as direct modifications of the acquisition function.[26–28] Alternatively, one can rely on human-in-the-loop strategies.[29] Different cost-aware methods have been developed that favor low-cost experiments following a given budget.[30–34] In all these examples, costs are kept fixed during the optimization. A different but related problem that has been investigated in the literature is resource management optimization, albeit not applied to BO.[35–37] Finally, recent works have focused on the cost of changing between different experimental setups to account for the associated expenses.[38–45]

To our knowledge, a framework accounting both for the cost of experiments and the fact that these costs change over time has not been discussed before for reaction optimization. We demonstrate the performance of CIBO using two HTE datasets of Pd-catalyzed reactions and find that, despite occasionally requiring additional experiments to match standard BO, the overall cost of the optimization campaign is significantly reduced. Our benchmarks evaluate cost using the price of commercially available reagents. However, CIBO is compatible with any cost definition, such as the number of steps or estimated time required for synthesizing reported ligands, as well as sustainability metrics for solvents or compounds. In the latter case, the platform will prioritize options with lower environmental impact.[46–48]

Overall, CIBO promises an efficient and sustainable alternative to existing design-of-experiment methods in the traditional lab setting, as well as in future self-driving laboratories.[49,50]

2 Methods

2.1 Cost-informed Bayesian optimization

BO is an effective method for optimizing noisy functions that are expensive or time-consuming to evaluate.[51] At its core lies a surrogate model that predicts the value of the function (mean μ) and the uncertainty of the prediction (standard deviation σ). Acquisition functions are then used to identify promising experiments in the design space.[52] This is done by considering all possible experiments and weighting their potential to maximize μ. In addition to performing one experiment at a time, batch BO proposes a joined set (*i.e.*, a *batch* of experiments) which provides the largest expected improvement when performed in parallel.

Given the previous definition, BO does not account for the varying costs of the resources involved in the optimization. Instead of buying or synthesizing additional substances, chemists may first choose to vary easily controllable conditions (*e.g.*, temperature, reaction time), resulting in lower costs and a better-informed decision before acquiring additional compounds.

Finding the best experimental conditions for the smallest budget is different from identifying the best value–cost trade-off in the optimized reaction. The latter is relevant when large amounts of the compounds involved need to be acquired repeatedly: *e.g.*, large-scale synthesis. The former, and current case, is relevant when the budget (in terms of cost, time, or other) for the experimentation campaign itself is important.

Our method, cost-informed Bayesian optimization (CIBO), balances minimizing experimentation cost with maximizing measured improvement (see Figure 1). It results in experiments that mimic more closely the optimization process in a chemistry lab, only acquiring a compound if the expected improvement justifies the costs. CIBO stands out from standard BO by promoting the search for more cost-effective experiments while not constraining its search space (*i.e.*, an expensive ligand may still be selected if its expected improvement justifies it).

We account for experimentation costs by including them in the acquisition function,[53,54] scaling the cost adjustment with the expected improvement values. Given the set of all possible experiments $\{e\}$, we use the batch noisy expected improvement (qNEI) acquisition function $\{\alpha_e\}$, computed from the predicted μ and σ from the surrogate model to determine the next batch $\mathcal{B} := \{e_1, \ldots, e_5\} \subset \{e\}$ for each iteration.[55,56] Here e_j is the j-th experiment in batch $\mathcal{B}$ with acquisition function value $\alpha_j \equiv \alpha_{e_j}$. Note there is exactly one α_j per experiment e_j. For a batch of experiments, here $N_e = 5$ per batch, we consider the current cost of each compound involved. As user input, only the compound prices per gram, per mol, per bottle, or other user-defined costs p_j are required.

For simplicity, here we cover the case of one compound j per experiment e_j. Batches are ordered with respect to the norm of $\mathcal{B}$, defined by the sum of the acquisition function values α_i in each batch,

$$|\mathcal{B}| := \sum_{j=1}^{N_e=5} \alpha_j. \tag{1}$$

In standard batch BO, the batch with the highest rank (*i.e.*, the highest expected improvement, represented by the blue line in Figure 1b) is chosen for the next set of experiments, as it offers the best combination of expected improvement and cost. In CIBO, we modify qNEI of each experiment e_j by subtracting a term[44,53,54] proportional to the cost $\mathcal{C}(e_j)$ as follows:

$$\tilde{\alpha}_j = \alpha_j - \mathcal{C}(e_j) \cdot S(\{\alpha\}), \tag{2}$$

where $\mathcal{C}$ probes if the compound j was bought at price p_j in a previous iteration or added to the same batch before, *i.e.*,

$$\mathcal{C}(e_j) = \begin{cases} 0 & \text{if } j \text{ bought or already in } \mathcal{B}, \\ p_j & \text{otherwise.} \end{cases} \tag{3}$$

This means that the cost is zero when j was obtained in some previous iteration or appeared before in the same batch $\mathcal{B}$ but under differ-

ent experimental conditions. For simplicity, throughout this work we will assume that compounds in the inventory are never exhausted. However, our framework could be set to deduct exact quantities from the inventory until nothing remains.

Finally in Equation 2, we have introduced a scaling function S that depends on all acquisition function values $\{\alpha\}$ evaluated on all experiments not yet included in the surrogate model and the current prices $\{p\}$ for each ligand,

$$S := \max\{\alpha\}/\max\{p\}, \tag{4}$$

to update the magnitude of the prices that enter Equation 2 such that the maximum value of the subtracted term has the same value as the largest acquisition function value. The purpose of this rescaling is to balance cost with the exploration–exploitation trade-off defined by the original qNEI acquisition function term α. Note that the scaling function is updated after each iteration to ensure adapting to the current costs and acquisition function values. Additionally, the scaling function removes the cost units.

After computing the modified acquisition values for each potential experiment in a batch $\tilde{\alpha}_1, \ldots, \tilde{\alpha}_5$, we evaluate the updated norm value,

$$|\tilde{\mathcal{B}}| := \sum_{j=1}^{N_e=5} \tilde{\alpha}_j. \tag{5}$$

This corresponds to the orange line in Figure 1b, which differs from the blue line depending on the scaled cost term of Equation 2. The batch with maximal value $\tilde{\mathcal{B}}$, offering the best cost–benefit ratio, is then selected for the next iteration.

Note that our approach does not penalize exploration *per se*, but rather favors the most inexpensive ways to explore before committing to higher-cost experiments. Further details of our implementation are available in the SI Section S1.

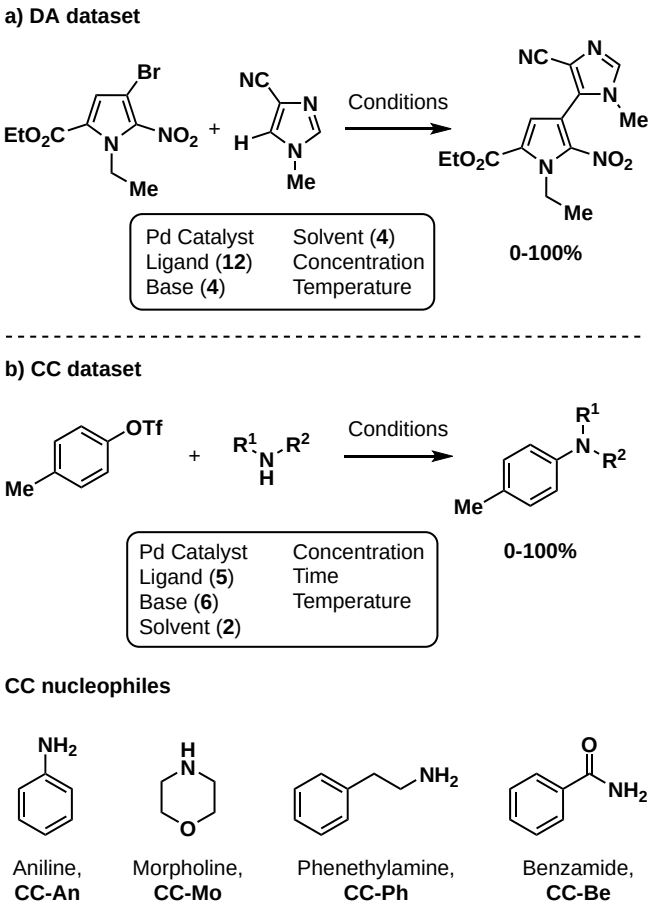

Figure 2: Reaction schemes of the two datasets used in this work. (a): Direct arylation (**DA**) with yields ranging from 0–100%.[3] (b): Cross-coupling (**CC**) with yields ranging from 0–100%.[17] The four nucleophiles explored in the latter dataset are depicted below, each leading to a subset (**CC-An**, **CC-Mo**, **CC-Ph**, **CC-Be**).

2.2 Datasets and models

To demonstrate the potential of CIBO, we attempt to maximize reaction yield as cost-efficiently as possible in two literature datasets: Pd-catalyzed direct C–H arylation (**DA**, Figure 2a)[3] and Pd-catalyzed Buchwald–Hartwig cross-couplings of amine nucleophiles using a droplet platform (**CC**, Figure 2b).[17] Since ligands are the most expensive elements of the optimization, CIBO should avoid using unnecessary ones to reach a high yield.

In both cases, we use Gaussian process regression (see SI Section S2 for more details) as the surrogate model in the optimization. The batch size is always fixed to five,[3,18] but we note

any other value could be used according to the preferred experimental setup. As cost, we use the price per gram of material (dollars per gram \$/g, see SI Section S3 for details), since chemical suppliers provide most chemicals gram-wise.

For consistency and to make the task challenging, we start the optimization using models trained only on a small subset of experiments performed with chemicals with low overall cost or performance (iteration zero, *vide infra*). The cost of the initial experiments is included in the total cost. For simplicity, we assume that we never run out of chemicals in the inventory, which applies both to the reagents used in the initial experiments and to newly acquired ones. The results are averaged over 5 separate runs using the same initialization.[57]

The **DA** dataset consists of 1728 measurements where the monophosphine ligand, base and solvent, concentration, and temperature are varied to optimize the formation of one product. The cost of chemicals was already reported previously[18] and was converted to dollars per gram for this work. We considered that concentration and temperature can be varied without additional cost. In general, the costs of the base and solvent are negligible compared to the ligand. The optimization begins with the surrogate model trained exclusively on the 144 experiments in the dataset that use the dimethylphenylphosphine ligand, which is inexpensive and the worst-performing ligand overall.

For the **CC** dataset, four different amine nucleophiles, aniline (**CC-An**), morpholine (**CC-Mo**), phenethylamine (**CC-Ph**), and benzamide (**CC-Be**), are each coupled with *p*-tolyl triflate yielding different products over 363 reactions (around 90 per nucleophile). Each is considered an individual reaction in which the ligand (precatalyst), base (concentration and equivalents), solvent, temperature, and time all change.[23] The prices for all chemicals involved were obtained from supplier websites and converted to dollars per gram if necessary. Concentrations, equivalents, temperature, and time are varied without additional cost. On average, nine experiments are used at the start of each optimization. For the **CC-Ph** and **CC-Be** subsets of **CC** we initialize using tBuXPhos as a ligand (for the precatalyst) and DBU as a base, the cheapest combination. For **CC-An**, we chose to initialize with EPhos and TEA as the cheapest initialization would have resulted in perfect yield. For **CC-Mo** no measurements with tBuXPhos were performed. Thus, we initialize with tBuBrettPhos and DBU. For the **CC** experiments, the cost of the solvents and bases were also taken into account, since the dataset contains few different ligands per nucleophile. A full list of chemical abbreviations is presented in Table S1.

3 Results

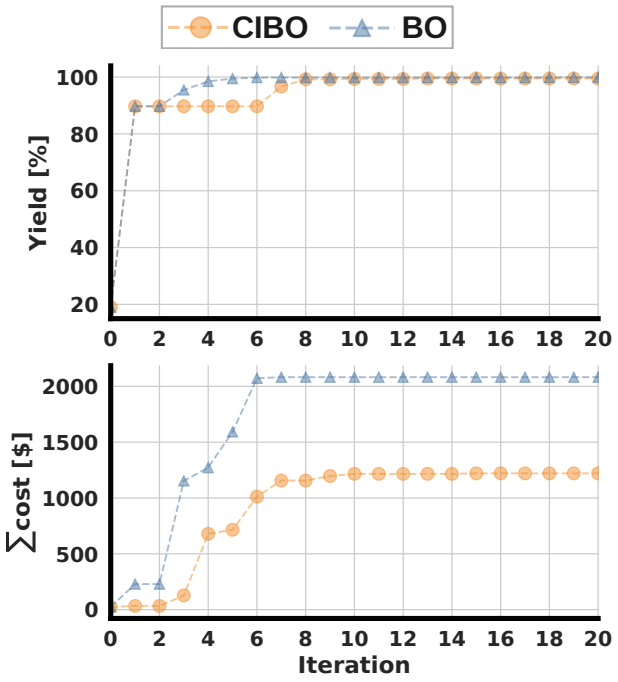

Figure 3: Comparison of cost-informed Bayesian optimization (CIBO, orange) and standard BO (blue) for the **DA** dataset. Average curves over five runs are shown. The best-obtained yield in each batch iteration is shown in the top panel, and the sum spent to acquire the ligands is shown below.

3.1 Direct Arylation

Following the original publication, the optimization was run for 20 total iterations (100 ex-

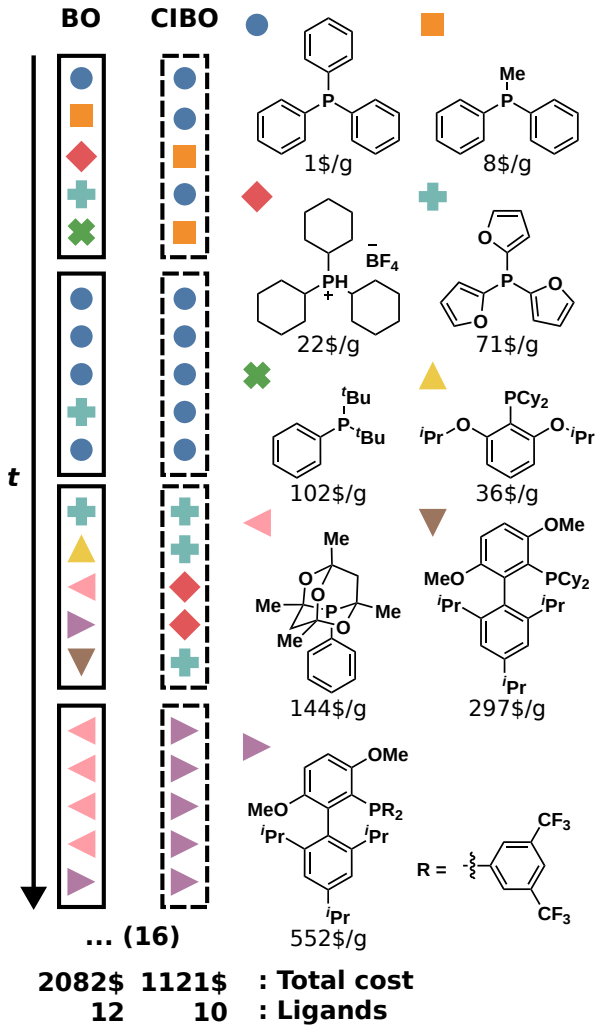

Figure 4: Composition of the four first batches of ligands with different reaction conditions for the **DA** dataset. We compare the total cost and number of ligands after 20 (16 more) iterations for BO and CIBO. Every colored symbol corresponds to one ligand shown on the right respectively.

periments).[3] The best-observed yield for several batched iterations, as well as the total amount spent, are shown in Figure 3. CIBO achieves a yield of over 99% after eight iterations, compared to five iterations in standard BO. However, the total cost of running the five iterations (25 experiments) suggested by BO is on average \$1592, whereas the cost of running the eight iterations proposed by CIBO (40 experiments) is on average of \$1156, 38% less.

After 20 iterations, standard BO bought all possible ligands (12), amounting to a cost of \$2082, while CIBO avoided buying two less (10), summing up to \$1221, saving more than 40% in comparison.

Considering the error of the surrogate model (see SI Section S1) it is a prudent strategy not to buy additional ligands after reaching a yield above 99%. Following this logic, past that point in the tenth iteration, CIBO suggests optimizing reaction conditions instead of buying expensive ligands.

To study the difference in experimentation planning between BO and CIBO, we visualize the ligand batch composition of the first four iterations (see Figure 4) for one of the five repetitions. CIBO batches are less diverse than BO in terms of ligands: CIBO suggests at most two different ligands per batch. In the first iteration, BO suggests acquiring five more distinct ligands besides dimethylphenylphosphine, compared to two for CIBO.

3.2 Cross-coupling

The dataset is split into subsets of the four different amine nucleophiles (see Figure 2b) resulting in four distinct optimization problems. As before, we consider the best yet achieved yield and total cost per iteration. The total amount spent, listed in the first row of Table 1, depends on the nucleophile subset, since not all combinations of ligands (precatalysts) and bases were tested experimentally for each nucleophile. Due to the limited amount of data available, the optimization is continued until the experimental data is exhausted, and the results are averaged over five different runs.

As shown in the bottom row of Figure 5 standard BO suggests buying all available ligands and bases combinations immediately after the first iteration – irrespective of their costs. CIBO acquires less expensive reagents first and recommends experiments under varying conditions before finally buying all reagents – if no other experiments are left in the dataset. By inspecting the amount spent versus the number of iterations, Figure 5, it is apparent that CIBO suggests more expensive molecules only as a last option. If possible, CIBO optimizes the yield based on combinations of reagents and conditions that can be performed at small or

Table 1: Comparison of the amounts spent with CIBO and classical BO for each nucleophile subset of **CC** experiments. The first row lists the total cost of all compounds that could be acquired, which equals the amount spent by classical BO. The subsequent rows show the amount spent with CIBO and the amount and percentage saved compared to classical BO to achieve a yield above 70%. The last row indicates how many compounds did not have to be obtained out of the available ones.

	CC-An	**CC-Mo**	**CC-Ph**	**CC-Be**
Cost all reagents	\$2474	\$2105	\$2170	\$2144
CIBO spent	\$1124	\$2105	\$142	\$690
Saved wrt. BO	\$1350 (**54%**)	0	\$2028 (**93%**)	\$1454 (**68%**)
Saved reagents	5/9	0/7	4/8	3/8

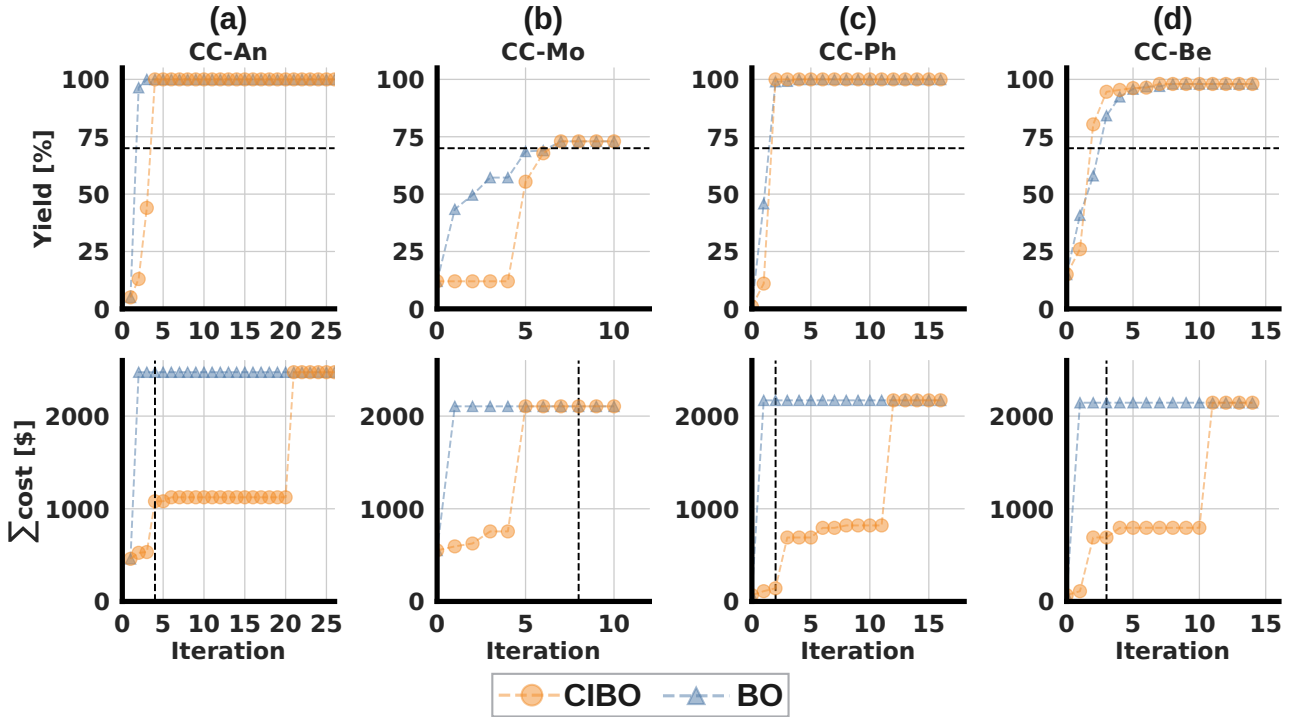

Figure 5: Yield optimization for the **CC** dataset with four different nucleophiles (**CC-An**, **CC-Mo**, **CC-Ph**, **CC-Be**). We compare cost-informed Bayesian optimization (CIBO, orange) to Bayesian optimization (BO, blue). Average curves over five runs are shown. The top row shows the best yield found as a function of the batch iteration, and the bottom row displays the cumulative costs. The black dotted line in the top row indicates the target yield at 70%. The resulting terminal iteration (vertical black line) in the bottom row indicates the total budget spent with CIBO.

zero additional costs. For instance, in the case of **CC-Be**, after iteration four, over 30 experiments are performed without acquiring any additional reagents, resulting in a cost plateau. Similar observations are made for **CC-An** and **CC-Ph** (Figure 5a,c) where the optimization achieved a yield of 100% after a few iterations, followed by iterations where CIBO does not acquire additional reagents.

The goal of every experimentation series always depends on the context. We define a stopping criterion for the experimentation such that at least 70% yield should be achieved, as this was quoted as a *high* yield in the original **CC** publication.[17] In Table 1 we show the cost and reagent savings compared to standard BO when using this criterion. In all cases, except for **CC-Mo**, the cost is reduced by over 50%, and fewer

compounds were acquired.

For **CC-Mo**, CIBO performs as well as BO in terms of yield optimization for the same cost, but only because the most performant ligand is the most expensive one and has to be acquired to achieve a high yield. For **CC-Ph**, CIBO acquires not a single additional ligand but manages to find a perfect yield by buying a much cheaper base (BTTP) thus saving 93% compared to the standard BO experiments that demand acquiring all ligands in a single step.

4 Conclusion

We introduced cost-informed Bayesian optimization (CIBO), a variant of Bayesian optimization that balances cost and ease of experimentation with a global optimization objective. The algorithm retains the flexibility to identify the most promising experiments of BO but takes a more cost-efficient optimization path. Updating the acquisition function according to the current inventory status is necessary to guide the optimization, with substantial savings in terms of acquisition costs as well as the number of purchases.

While in this work we focused on the economic cost of buying the compounds needed for an experiment, our framework is general and amenable to any user-defined kind of cost, including logistical availability, synthesizability, safety, time, environmental impact, and sustainability. Similarly, the number of steps for a synthesis, if the compounds were already reported or commercially available, structural complexity, time required to run an experiment, or the running cost of a laboratory could also directly be taken into account. Our proposed approach only plans one experiment ahead (myopic). CIBO may be extended and improved by planning multiple steps and accounting for resource management.[33,45,58]

CIBO is far more applicable than standard BO in contexts where the cost of each experiment are very dissimilar, such as state-of-the-art fine chemistry where the difference in price between two ligands may be considerable. Ultimately, we envision our method to be useful both for traditional and future self-driving laboratories in the context of reaction optimization and development, as well as in other similar tasks in chemistry and beyond where the cost of performing the actual optimization matters.

5 Code availability

CIBO is available at `https://github.com/lcmd-epfl/cibo`.

6 Acknowledgement

The authors thank EPFL for computational resources. This publication was created as part of NCCR Catalysis (grant number 180544), a National Centre of Competence in Research funded by the Swiss National Science Foundation. We also thank Puck van Gerwen and Dr. Ksenia Briling for fruitful discussions.

References

(1) Ahneman, D. T.; Estrada, J. G.; Lin, S.; Dreher, S. D.; Doyle, A. G. Predicting reaction performance in C–N cross-coupling using machine learning. *Science* **2018**, *360*, 186–190.

(2) Zahrt, A. F.; Henle, J. J.; Rose, B. T.; Wang, Y.; Darrow, W. T.; Denmark, S. E. Prediction of higher-selectivity catalysts by computer-driven workflow and machine learning. *Science* **2019**, *363*, eaau5631.

(3) Shields, B. J.; Stevens, J.; Li, J.; Parasram, M.; Damani, F.; Alvarado, J. I. M.; Janey, J. M.; Adams, R. P.; Doyle, A. G. Bayesian reaction optimization as a tool for chemical synthesis. *Nature* **2021**, *590*, 89–96.

(4) Schwaller, P.; Vaucher, A. C.; Laplaza, R.; Bunne, C.; Krause, A.; Corminboeuf, C.; Laino, T. Machine intelligence for chemical reaction space. *WIREs Comput. Mol. Sci.* **2022**, *12*, e1604.

(5) Shim, E.; Tewari, A.; Cernak, T.; Zimmerman, P. M. Machine Learning Strategies for Reaction Development: Toward the Low-Data Limit. *J. Chem. Inf. Model.* **2023**, *63*, 3659–3668.

(6) Ranković, B.; Schwaller, P. BoChemian: Large language model embeddings for Bayesian optimization of chemical reactions. NeurIPS 2023 Workshop on Adaptive Experimental Design and Active Learning in the Real World. 2023.

(7) Guo, J.; Ranković, B.; Schwaller, P. Bayesian optimization for chemical reactions. *CHIMIA* **2023**, *77*, 31–38.

(8) Ranković, B.; Griffiths, R.-R.; Moss, H. B.; Schwaller, P. Bayesian optimisation for additive screening and yield improvements – beyond one-hot encoding. *Digital Discovery* **2024**, Advance Article.

(9) Schweidtmann, A. M.; Clayton, A. D.; Holmes, N.; Bradford, E.; Bourne, R. A.; Lapkin, A. A. Machine learning meets continuous flow chemistry: Automated optimization towards the Pareto front of multiple objectives. *Chem. Eng. J.* **2018**, *352*, 277–282.

(10) Clayton, A. D.; Schweidtmann, A. M.; Clemens, G.; Manson, J. A.; Taylor, C. J.; Niño, C. G.; Chamberlain, T. W.; Kapur, N.; Blacker, A. J.; Lapkin, A. A.; Bourne, R. A. Automated self-optimisation of multi-step reaction and separation processes using machine learning. *Chem. Eng. J.* **2020**, *384*, 123340.

(11) Angello, N. H.; Rathore, V.; Beker, W.; Wołos, A.; Jira, E. R.; Roszak, R.; Wu, T. C.; Schroeder, C. M.; Aspuru-Guzik, A.; Grzybowski, B. A.; Burke, M. D. Closed-loop optimization of general reaction conditions for heteroaryl Suzuki-Miyaura coupling. *Science* **2022**, *378*, 399–405.

(12) McCullough, K. E.; King, D. S.; Chheda, S. P.; Ferrandon, M. S.; Goetjen, T. A.; Syed, Z. H.; Graham, T. R.; Washton, N. M.; Farha, O. K.; Gagliardi, L.; Delferro, M. High-throughput experimentation, theoretical modeling, and human intuition: Lessons learned in metal–organic-framework-supported catalyst design. *ACS Cent. Sci.* **2023**, *9*, 266–276.

(13) Taylor, C. J.; Felton, K. C.; Wigh, D.; Jeraal, M. I.; Grainger, R.; Chessari, G.; Johnson, C. N.; Lapkin, A. A. Accelerated chemical reaction optimization using multi-task learning. *ACS Cent. Sci.* **2023**, *9*, 957–968.

(14) Seifrid, M.; Pollice, R.; Aguilar-Granda, A.; Morgan Chan, Z.; Hotta, K.; Ser, C. T.; Vestfrid, J.; Wu, T. C.; Aspuru-Guzik, A. Autonomous chemical experiments: Challenges and perspectives on establishing a self-driving lab. *Acc. Chem. Res.* **2022**, *55*, 2454–2466.

(15) Hickman, R. J.; Bannigan, P.; Bao, Z.; Aspuru-Guzik, A.; Allen, C. Self-driving laboratories: A paradigm shift in nanomedicine development. *Matter* **2023**, *6*, 1071–1081.

(16) Ramirez, A.; Lam, E.; Gutierrez, D. P.; Hou, Y.; Tribukait, H.; Roch, L. M.; Copéret, C.; Laveille, P. Accelerated exploration of heterogeneous CO_2 hydrogenation catalysts by Bayesian-optimized high-throughput and automated experimentation. *Chem Catal.* **2024**, *4*, 100888.

(17) Baumgartner, L. M.; Dennis, J. M.; White, N. A.; Buchwald, S. L.; Jensen, K. F. Use of a droplet platform to optimize Pd-catalyzed C–N coupling reactions promoted by organic bases. *Org. Process Res. Dev.* **2019**, *23*, 1594–1601.

(18) Torres, J. A. G.; Lau, S. H.; Anchuri, P.; Stevens, J. M.; Tabora, J. E.; Li, J.; Borovika, A.; Adams, R. P.; Doyle, A. G.

A multi-objective active learning platform and web app for reaction optimization. *J. Am. Chem. Soc.* **2022**, *144*, 19999–20007.

(19) Strieth-Kalthoff, F.; Sandfort, F.; Kühnemund, M.; Schäfer, F. R.; Kuchen, H.; Glorius, F. Machine learning for chemical reactivity: The importance of failed experiments. *Angew. Chem. Int. Ed.* **2022**, *61*, e202204647.

(20) Schleinitz, J.; Langevin, M.; Smail, Y.; Wehnert, B.; Grimaud, L.; Vuilleumier, R. Machine learning yield prediction from NiCOlit, a small-size literature data set of nickel catalyzed C–O couplings. *J. Am. Chem. Soc.* **2022**, *144*, 14722–14730.

(21) Maloney, M. P.; Coley, C. W.; Genheden, S.; Carson, N.; Helquist, P.; Norrby, P.-O.; Wiest, O. Negative data in data sets for machine learning training. *Org. Lett.* **2023**, *25*, 2945–2947.

(22) Voinarovska, V.; Kabeshov, M.; Dudenko, D.; Genheden, S.; Tetko, I. V. When yield prediction does not yield prediction: An overview of the current challenges. *J. Chem. Inf. Model.* **2023**, *64*, 42–56.

(23) Felton, K. C.; Rittig, J. G.; Lapkin, A. A. Summit: Benchmarking machine learning methods for reaction optimisation. *Chem. Methods* **2021**, *1*, 116–122.

(24) Hickman, R. J.; Aldeghi, M.; Häse, F.; Aspuru-Guzik, A. Bayesian optimization with known experimental and design constraints for chemistry applications. *Digital Discovery* **2022**, *1*, 732–744.

(25) Krause, A.; Ong, C. Contextual Gaussian process bandit optimization. Adv. Neural Inf. Process. Syst. 2011; pp 2447–2455.

(26) Snoek, J.; Larochelle, H.; Adams, R. P. Practical Bayesian Optimization of Machine Learning Algorithms. Adv. Neural Inf. Process. Syst. 2012; pp 2951–2959.

(27) Berk, J.; Nguyen, V.; Gupta, S.; Rana, S.; Venkatesh, S. Exploration enhanced expected improvement for Bayesian optimization. Joint European Conference on Machine Learning and Knowledge Discovery in Databases. 2018; pp 621–637.

(28) Tachibana, R.; Zhang, K.; Zou, Z.; Burgener, S.; Ward, T. R. A Customized Bayesian Algorithm to Optimize Enzyme-Catalyzed Reactions. *ACS Sustain. Chem. Eng.* **2023**, *11*, 12336–12344.

(29) Biswas, A.; Liu, Y.; Creange, N.; Liu, Y.-C.; Jesse, S.; Yang, J.-C.; Kalinin, S. V.; Ziatdinov, M. A.; Vasudevan, R. K. A dynamic Bayesian optimized active recommender system for curiosity-driven partially Human-in-the-loop automated experiments. *npj Comput. Mater.* **2024**, *10*, 29.

(30) Dolatnia, N.; Fern, A.; Fern, X. Bayesian optimization with resource constraints and production. *Proceedings of the ICAPS* **2016**, *26*, 115–123.

(31) Lee, E. H.; Perrone, V.; Archambeau, C.; Seeger, M. Cost-aware Bayesian optimization. ICML 2020 Workshop on AutoML. 2020.

(32) Luong, P.; Nguyen, D.; Gupta, S.; Rana, S.; Venkatesh, S. Adaptive cost-aware Bayesian optimization. *Knowl.-Based Syst.* **2021**, *232*, 107481.

(33) Lee, E. H.; Eriksson, D.; Perrone, V.; Seeger, M. A nonmyopic approach to cost-constrained Bayesian optimization. Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence. 2021; pp 568–577.

(34) Belakaria, S.; Doppa, J. R.; Fusi, N.; Sheth, R. Bayesian optimization over iterative learners with structured responses: A budget-aware planning approach. Proceedings of The 26th International Conference on Artificial Intelligence and Statistics. 2023; pp 9076–9093.

(35) Rodriguez Diaz, P.; Killian, J. A.; Xu, L.; Suggala, A. S.; Taneja, A.; Tambe, M. Flexible budgets in restless bandits: A primal-dual algorithm for efficient budget allocation. *Proceedings of the AAAI Conference on Artificial Intelligence* **2023**, *37*, 12103–12111.

(36) Zuo, J.; Joe-Wong, C. Combinatorial multi-armed bandits for resource allocation. 2021 55th Annual Conference on Information Sciences and Systems (CISS). 2021; pp 1–4.

(37) Thananjeyan, B.; Kandasamy, K.; Stoica, I.; Jordan, M.; Goldberg, K.; Gonzalez, J. Resource allocation in multi-armed bandit exploration: Overcoming sublinear scaling with adaptive parallelism. Proceedings of the 38th International Conference on Machine Learning. 2021; pp 10236–10246.

(38) Vellanki, P.; Rana, S.; Gupta, S.; Rubin, D.; Sutti, A.; Dorin, T.; Height, M.; Sanders, P.; Venkatesh, S. Process-constrained batch Bayesian optimisation. Adv. Neural Inf. Process. Syst. 2017; pp 3414–3423.

(39) Folch, J. P.; Zhang, S.; Lee, R.; Shafei, B.; Walz, D.; Tsay, C.; van der Wilk, M.; Misener, R. SnAKe: Bayesian optimization with pathwise exploration. Adv. Neural Inf. Process. Syst. 2022; pp 35226–35239.

(40) Ramesh, S. S.; Sessa, P. G.; Krause, A.; Bogunovic, I. Movement penalized Bayesian optimization with application to wind energy systems. Adv. Neural Inf. Process. Syst. 2022; pp 27036–27048.

(41) Yang, A. X.; Aitchison, L.; Moss, H. B. MONGOOSE: Path-wise smooth Bayesian optimisation via meta-learning. *arXiv preprint* **2023**, arXiv:2302.11533.

(42) Folch, J. P.; Odgers, J. A. C.; Zhang, S.; Lee, R. M.; Shafei, B.; Walz, D.; Tsay, C.; van der Wilk, M.; Misener, R. Practical path-based Bayesian optimization. NeurIPS 2023 Workshop on Adaptive Experimental Design and Active Learning in the Real World. 2023.

(43) Liu, P.; Wang, H.; Qiyu, W. Bayesian optimization with switching cost: Regret analysis and lookahead variants. Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23. 2023; pp 4011–4018.

(44) Liang, R.; Hu, H.; Han, Y.; Chen, B.; Yuan, Z. CAPBO: A cost-aware parallelized Bayesian optimization method for chemical reaction optimization. *AIChE J.* **2024**, *70*, e18316.

(45) Folch, J. P.; Tsay, C.; Lee, R. M.; Shafei, B.; Ormaniec, W.; Krause, A.; van der Wilk, M.; Misener, R.; Mutný, M. Transition constrained Bayesian optimization via Markov decision processes. *arXiv preprint* **2024**, arXiv:2402.08406.

(46) Tulus, V.; Pérez-Ramírez, J.; Guillén-Gosálbez, G. Planetary metrics for the absolute environmental sustainability assessment of chemicals. *Green Chem.* **2021**, *23*, 9881–9893.

(47) Byrne, F. P.; Jin, S.; Paggiola, G.; Petchey, T. H. M.; Clark, J. H.; Farmer, T. J.; Hunt, A. J.; McElroy, C. R.; Sherwood, J. Tools and techniques for solvent selection: green solvent selection guides. *Sustain. Chem. Process.* **2016**, *4*, 7.

(48) Mitchell, S.; Martín, A. J.; Guillén-Gosálbez, G.; Pérez-Ramírez, J. The Future of Chemical Sciences is Sustainable. *Angew. Chem. Int. Ed.* **2024**,

(49) Häse, F.; Roch, L. M.; Aspuru-Guzik, A. Next-generation experimentation with self-driving laboratories. *Trends Chem.* **2019**, *1*, 282–291.

(50) MacLeod, B. P. et al. Self-driving laboratory for accelerated discovery of thin-film materials. *Sci. Adv.* **2020**, *6*, eaaz8867.

(51) Rasmussen, C. E. *Gaussian Processes in Machine Learning*; Lecture Notes in Computer Science; Springer, 2004; Vol. 3176.

(52) Kaelbling, L. P.; Littman, M. L.; Moore, A. W. Reinforcement learning: A survey. *J. Artif. Intell. Res.* **1996**, *4*, 237–285.

(53) Garnett, R. *Bayesian Optimization*; Cambridge University Press, 2023.

(54) Raiffa, H.; Schlaifer, R. *Applied Statistical Decision Theory*; Division of Research, Graduate School of Business Administration, Harvard University, 1961; Chapter 4.

(55) Jones, D. R.; Schonlau, M.; Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **1998**, *13*, 455–492.

(56) Balandat, M.; Karrer, B.; Jiang, D. R.; Daulton, S.; Letham, B.; Wilson, A. G.; Bakshy, E. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. Adv. Neural Inf. Process. Syst. 2020; pp 21524–21538.

(57) CIBO is based on the batch noisy expected improvement (qNEI) function, which uses random initialization. Results change negligibly over the 5 independent runs, but we present the averaged results for completeness.

(58) Yue, X.; Kontar, R. A. Why Non-myopic Bayesian Optimization is Promising and How Far Should We Look-ahead? A Study via Rollout. Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics. 2022; pp 2808–2818.