
Multi-task Bayesian Optimization of Chemical Reactions

Kobi C. Felton

Department of Chemical Engineering
University of Cambridge
Cambridge, UK
kcmf2@cam.ac.uk

Daniel Wigh

Department of Chemical Engineering
University of Cambridge
Cambridge, UK
dsw46@cam.ac.uk

Alexei A. Lapkin *

Department of Chemical Engineering
University of Cambridge
Cambridge, UK
aal35@cam.ac.uk

Abstract

Recent work has shown how Bayesian optimization (BO) is an efficient method for optimizing expensive experiments such as chemical reactions. However, in previous studies, each optimization has been started from scratch with no information about previous or similar chemical optimization studies. Therefore, BO can still require more iterations than many experimental budgets provide. Here, we overcome this challenge using multi-task BO. Through *in silico* benchmarking studies, we show how past experimental data can be leveraged to improve the quality and speed of reaction optimization.

1 Introduction

Finding the optimal conditions for a chemical reaction is an essential task in the fine chemicals industry. Whether it is optimizing the recipe for a new advanced material or the conditions for a reaction in a multi-step pharmaceutical synthesis, this task is equally ubiquitous as it is difficult. Previous estimates have shown that the parameter space for a single catalytic reaction is upwards of 50 million potential combinations [1]. Given that the constraints of real experimental systems are in the hundreds to thousands of experiments per day [2], chemists and engineers must judiciously select experiments from this large parameter space to find optimal conditions.

Previous work has shown that Bayesian optimization (BO) has the potential to quickly find optimal reaction conditions. Schweidtmann et al. were the first to demonstrate this using the BO algorithm TSEMO to optimize several reactions using a flow chemistry reactor [3]. Since then, several case studies have shown the applicability of BO [4, 5], including a recent preprint, which conducted a systematic study and showed that BO performs well on several benchmarks for chemical reactions [6].

However, in each of these studies, BO was initialized without incorporation of prior knowledge about the reaction being optimized. This is the opposite of a chemist’s intuition about reaction optimization; they would bring to bear results from previous optimization campaigns they had run. Here, we ask if this same prior knowledge can be used to accelerate BO algorithms.

*Corresponding Author




Figure 1: Red is for Task one, and blue is for Task two. Dotted lines are the ground truth. Dots are training observations, and translucent lines are posterior samples of the trained model for each task. (a) Training a multi-task model with five observations for both tasks (b) Training a multi-task model with twenty observations for Task one and five observations for Task two. The extra observations in Task one improve the posterior predictions for Task two.

We rely on the framework of multi-task BO, first introduced by Swersky et al [7]. The idea is to replace the standard probabilistic model in BO with a multi-task model. As illustrated in a simple example in Figure 1, a multi-task model trained on data from a related task is better able to predict the outcomes of a new task than the equivalent single task model. Since the model is arguably the most important part of a BO algorithm [8], this improved prediction accuracy is likely to lead to better optimization performance.

In the following sections, we first review the multi-task BO framework before presenting results from three benchmarking studies that illustrate the advantage of using multi-task BO in reaction optimization. We investigate the influence of the number of data points from previous tasks on the speed of optimization, and we examine whether using a dissimilar reaction for pre-training will have a negative influence on the reaction optimization.

2 Multi-task Bayesian optimization

Bayesian optimization (BO) is a strategy for optimizing black-box functions where a probabilistic model can be used as a surrogate for a function in which evaluation/point generation is expensive, e.g. performing a chemical experiment. Within a BO strategy, a Gaussian process (GP) model is commonly used for specifying prior distributions over functions $f : \chi \rightarrow \mathbb{R}$ given that any finite set of N points $\mathbf{X} = \{\mathbf{x}_n \in \chi\}_{n=1}^N$ induces a Gaussian distribution on \mathbb{R}^N . The GP is defined by its mean function $m : \chi \rightarrow \mathbb{R}$ and a positive definite covariance, or kernel function $\mathbf{K} : \chi \times \chi \rightarrow \mathbb{R}$. Using these two functions, the posterior is parameterized by:

$$\mu(\mathbf{x}; \{\mathbf{x}_n, y_n\}, \theta) = \mathbf{K}(\mathbf{X}, \mathbf{x})^\top \mathbf{K}(\mathbf{X}, \mathbf{X})^{-1}(\mathbf{y} - m(\mathbf{X})), \quad (1)$$

$$\Sigma(\mathbf{x}, \mathbf{x}'; \{\mathbf{x}_n, y_n\}, \theta) = K(\mathbf{x}, \mathbf{x}') - \mathbf{K}(\mathbf{X}, \mathbf{x})^\top \mathbf{K}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{K}(\mathbf{X}, \mathbf{x}'). \quad (2)$$

To determine which experiments to perform next, an acquisition function $a(\mathbf{x}; \{\mathbf{x}_n, y_n\}, \theta)$ is used, which specifies the quality (information gain) of a suggested experiment given a model and its parameters. By optimizing the acquisition function, a myopic estimate of the best next point is found [8]. The expected improvement criterion (EI) is the acquisition function used within this work:

$$a_{EI}(\mathbf{x}; \{\mathbf{x}_n, y_n\}, \theta) = \sqrt{\Sigma(\mathbf{x}, \mathbf{x}; \{\mathbf{x}_n, y_n\}, \theta)}(\gamma(\mathbf{x})\Phi(\gamma(\mathbf{x})) + \mathcal{N}(\gamma(\mathbf{x}); 0, 1)), \quad (3)$$

$$\gamma(\mathbf{x}) = (y_{best} - \mu(\mathbf{x}; \{\mathbf{x}_n, y_n\}, \theta)) / \sqrt{\Sigma(\mathbf{x}, \mathbf{x}; \{\mathbf{x}_n, y_n\}, \theta)}. \quad (4)$$

Multi-task GPs can be used on vector-valued functions $f : \chi \rightarrow \mathbb{R}^T$, where each of the T outputs can be seen as solutions to unique regression tasks. A covariance function, or kernel, $K((\mathbf{x}, t), (\mathbf{x}', t'))$




Figure 2: Results of simulated optimization of reactions using single task (STBO) and multi-task (MTBO) strategies. Each strategy was run for ten iterations, and the mean and 95% confidence interval of ten random repeats are shown. (a) Comparison on a benchmark where MTBO is jointly trained on prior data from a reaction with the same mechanism as the one being optimized. n represents the number of auxiliary data points. (b) Comparison on a benchmark where MTBO is jointly trained on prior data from a reaction with a different mechanism to the one being optimized. n represents the number of auxiliary data points. (c) Comparison of a single task (STBO) and multi-task (MTBO) optimization of Pd-catalyzed C-N cross-coupling reaction [10]. Both optimizations have similar performance.

between input-task pairs must be defined to understand the relationship across inputs and tasks. The intrinsic model of coregionalization transforms a latent function to yield the outputs, and equation 5 is one example of such a function.

$$K_{multi}((\mathbf{x}, t), (\mathbf{x}', t')) = K_t(t, t') \otimes K_x(\mathbf{x}, \mathbf{x}') \quad (5)$$

\otimes is the Kronecker product, K_x is a kernel representing the relationship between inputs, while K_t is a kernel that represents the relationship between tasks. Transferring BO to apply to additional related tasks can simply be done by restricting new observations to only belong to the new task (chemical reaction) and conducting optimization of the acquisition function w.r.t. the task at hand [7].

3 Results

We optimize *in silico* benchmarks that represent real reactions. The first two benchmarks are a series of hand-crafted kinetic models designed as tests for different optimization strategies [9]. The third benchmark is based on data published by Baumgartner et al. for a Pd-catalyzed C-N cross-coupling [10]. This data contains optimizations of the same reaction with different substrates, providing a perfect chance to test how multi-task BO works with real data. To build the benchmark, a Bayesian neural network is trained to act as an oracle that the optimization strategies can query [6, 11]. For all benchmarks, the decision variable domain has both continuous variables (e.g., temperature) and categorical variables (e.g., catalysts). To handle the categorical variables, we use the "basic" method from Garrido-Merchán et al., which rounds the outputs of the internal acquisition function optimization to a one-hot encoding [12].

Figures 2a and 2b compare the performance of a single task (STBO) and a multi-task (MTBO) optimization in finding the optima of the hand-crafted kinetic model. The benchmark has three continuous variables (catalyst concentration, reaction time, reactor temperature) and one categorical (choice of catalyst) variable available to find the maximum yield. Latent hypercube sampling is performed to acquire the initial auxiliary dataset for MTBO [13]. Figure 2a shows a case where the auxiliary and optimised tasks (i.e., reactions) for MTBO have the same mechanism with different parameters. MTBO clearly improves the optimization speed, and adding more observations to the auxiliary dataset both improves the speed and reliability of the optimization (i.e., lower variance). A similar trend is seen in Figure 2b where the mechanism differs slightly between the auxiliary and optimization tasks; the optimization task introduces a side reactions that consumes one of the reactants. Despite this change in mechanism, the MTBO still has a significant speed and reliability advantage over STBO.

Figure 2c illustrates what might happen when the mechanism changes significantly from the auxiliary task to the optimization task. Here, the substrate changes from an aniline to amide, and as the original paper notes, this changes the optimal catalyst and residence time [10]. Therefore, the auxiliary data does not provide much benefit in improving the quality of the model. Despite this, we observe that the difference between STBO and MTBO is small in this case, indicating that negative transfer does not occur.

4 Related Work

In this work, we show that a model optimised for one task, specifically a chemical reaction, may be able to generalise the information within the task and subsequently perform better on a different task. Notably, we were able to achieve good performance using relatively few datapoints. This contrasts to reinforcement learning approach to optimising reactions, which requires thousands of pretraining iterations for each combination of decision variables and objectives [14]. Other work has built multi-task models for prediction of site selectivity but not incorporated reaction conditions or optimization[15]. Our work also differs from previous optimization work in that it aims generalize across changes in in mechanism[16].

5 Conclusions

Reaction optimization is an essential task in the fine chemicals industry. Yet, it is often time- and labor-intensive. Here, we demonstrate that multi-task Bayesian optimization (BO) enables closed-loop optimization of reactions by leveraging data from past experiments. Our benchmarking suggests that using past data will have a positive or, at worst, null effect on the quality and speed of the optimization. We therefore foresee multitask BO as an effective way to overcome the problem of limited experimental resources in chemical laboratories. Both academic and industrial labs have historical data, often going back years, and our technique would allow them to utilize that data to automate and accelerate reaction optimization.

References

- [1] Paul M. Murray, Simon N. G. Tyler, and Jonathan D. Moseley. Beyond the Numbers: Charting Chemical Reaction Space. *Org. Process Res. Dev.*, 17(1):40–46, jan 2013. ISSN 1083-6160. doi: 10.1021/op300275p. URL <http://pubs.acs.org/doi/10.1021/op300275p>.
- [2] Damith Perera, Joseph W Tucker, Shalini Brahmabhatt, Christopher J Helal, Ashley Chong, William Farrell, Paul Richardson, and Neal W Sach. A platform for automated nanomole-scale reaction screening and micromole-scale synthesis in flow. *Science*, 359(6374):429–434, jan 2018. ISSN 0036-8075. doi: 10.1126/science.aap9112. URL <http://www.sciencemag.org/lookup/doi/10.1126/science.aap9112>.
- [3] Artur M. Schweidtmann, Adam D. Clayton, Nicholas Holmes, Eric Bradford, Richard A. Bourne, and Alexei A. Lapkin. Machine learning meets continuous flow chemistry: Automated optimization towards the Pareto front of multiple objectives. *Chem. Eng. J.*, 352:277–282, nov 2018. doi: 10.1016/J.CEJ.2018.07.031. URL <https://www.sciencedirect.com/science/article/pii/S1385894718312634>.
- [4] Yehia Amar, Artur M. Schweidtmann, Paul Deutsch, Liwei Cao, and Alexei Lapkin. Machine learning and molecular descriptors enable rational solvent selection in asymmetric catalysis. *Chem. Sci.*, 2019. ISSN 2041-6520. doi: 10.1039/C9SC01844A.
- [5] Adam D. Clayton, Artur M. Schweidtmann, Graeme Clemens, Jamie A. Manson, Connor J. Taylor, Carlos G. Niño, Thomas W. Chamberlain, Nikil Kapur, A. John Blacker, Alexei A. Lapkin, and Richard A. Bourne. Automated self-optimisation of multi-step reaction and separation processes using machine learning. *Chemical Engineering Journal*, 384:123340, mar 2020. ISSN 13858947. doi: 10.1016/j.cej.2019.123340.
- [6] Kobi Felton, Jan Rittig, and Alexei Lapkin. Summit: Benchmarking Machine Learning Methods for Reaction Optimisation. *ChemRxiv*, 9 2020. doi: 10.26434/chemrxiv.12939806.v1.

URL https://chemrxiv.org/articles/preprint/Summit_Benchmarking_Machine_Learning_Methods_for_Reaction_Optimisation/12939806.

- [7] Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-Task Bayesian Optimization. In *NIPS Proceedings*, pages 2004–2012, 2013. URL <http://papers.neurips.cc/paper/5086-multi-task-bayesian-optimization>.
- [8] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization, jan 2016. ISSN 15582256.
- [9] Lorenz M. Baumgartner, Connor W. Coley, Brandon J. Reizman, Kevin W. Gao, and Klavs F. Jensen. Optimum catalyst selection over continuous and discrete process variables with a single droplet microfluidic reaction platform. *React. Chem. Eng.*, 3(3):301–311, jun 2018. ISSN 2058-9883. doi: 10.1039/C8RE00032H. URL <http://xlink.rsc.org/?DOI=C8RE00032H>.
- [10] Lorenz M. Baumgartner, Joseph M. Dennis, Nicholas A. White, Stephen L. Buchwald, and Klavs F. Jensen. Use of a Droplet Platform to Optimize Pd-Catalyzed C-N Coupling Reactions Promoted by Organic Bases. *Org. Process Res. Dev.*, 23(8):1594–1601, aug 2019. ISSN 1520586X. doi: 10.1021/acs.oprd.9b00236.
- [11] Florian Häse, Loïc M. Roch, and Alán Aspuru-Guzik. Gryffin: An algorithm for Bayesian optimization for categorical variables informed by physical intuition with applications to chemistry. *arXiv preprint*, page arXiv: 2003.12127, 2020.
- [12] Eduardo C. Garrido-Merchán and Daniel Hernández-Lobato. Dealing with Categorical and Integer-valued Variables in Bayesian Optimization with Gaussian Processes. *Neurocomputing*, 380:20–35, may 2018. doi: 10.1016/j.neucom.2019.11.004. URL <http://arxiv.org/abs/1805.03463><http://dx.doi.org/10.1016/j.neucom.2019.11.004>.
- [13] R J Beckman, W J Conover, and M D McKay. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 1979.
- [14] Zhenpeng Zhou, Xiaocheng Li, and Richard N. Zare. Optimizing Chemical Reactions with Deep Reinforcement Learning. *ACS Central Science*, 3(12):1337–1344, December 2017. ISSN 2374-7943. doi: 10.1021/acscentsci.7b00492. URL <https://doi.org/10.1021/acscentsci.7b00492>.
- [15] Thomas J. Struble, Connor W. Coley, and Klavs F. Jensen. Multitask prediction of site selectivity in aromatic C–H functionalization reactions. *Reaction Chemistry & Engineering*, 5(5):896–902, May 2020. ISSN 2058-9883. doi: 10.1039/D0RE00071J. URL <http://pubs.rsc.org/en/content/articlelanding/2020/re/d0re00071j>. Publisher: The Royal Society of Chemistry.
- [16] Robert W. Epps, Michael S. Bowen, Amanda A. Volk, Kameel Abdel-Latif, Suyong Han, Kristofer G. Reyes, Aram Amassian, and Milad Abolhasani. Artificial Chemist: An Autonomous Quantum Dot Synthesis Bot. *Advanced Materials*, 32(30):2001626, 2020. ISSN 1521-4095. doi: 10.1002/adma.202001626. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202001626>.