
Bayesian optimisation for additive screening and yield improvements in chemical reactions – beyond one-hot encodings

Bojana Ranković

Laboratory of Artificial Chemical Intelligence (LIAC) &
National Centre of Competence in Research (NCCR) Catalysis
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
bojana.rankovic@epfl.ch

Ryan-Rhys Griffiths

Department of Physics
University of Cambridge, UK
rrg27@cam.ac.uk

Henry B. Moss

Secondmind.ai
Cambridge, UK
henry.moss@secondmind.ai

Philippe Schwaller

Laboratory of Artificial Chemical Intelligence (LIAC) &
National Centre of Competence in Research (NCCR) Catalysis
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
philippe.schwaller@epfl.ch

Abstract

Reaction additives play a significant role in controlling the reactivity and outcome of chemical reactions. For example, a recent high-throughput additive screening identified a phthalimide ligand additive Ni-catalysed photoredox decarboxylative arylations. This discovery enabled a 4-fold yield improvement by stabilising oxidative addition complexes and breaking up deactivated catalyst aggregates. However, such large-scale screenings are currently inaccessible to most research groups. This work demonstrates how these discoveries can be made under much lower experimental budgets using Bayesian optimisation. We consider a unique reaction screening setting with 720 additives which forces us to go beyond simple one-hot encoding of the reaction components. We investigate a range of molecular representations and demonstrate convincing improvements over baselines. Our approach is not limited to Ni-catalysed reactions but can be generally applied to, for example, achieving yield improvements in diverse cross-coupling reactions or unlocking access to new chemical spaces of interest to the chemical and pharmaceutical industries.

1 Introduction

Artificial intelligence holds the promise to accelerate chemical sciences¹⁻³. In the last decade, we witnessed ground-breaking advances in machine learning for de novo molecular design⁴⁻⁸, synthesis planning⁹⁻¹⁵, and reaction outcome prediction¹⁶⁻²³. More recently, sequential model-based optimisation algorithms have been investigated to efficiently find optimal conditions for chemical reactions²⁴⁻³². As demonstrated in the space of chemical reactions by Shields et al.²⁹, Bayesian optimisation (BO) is particularly well suited for trading exploration and exploitation in the low data

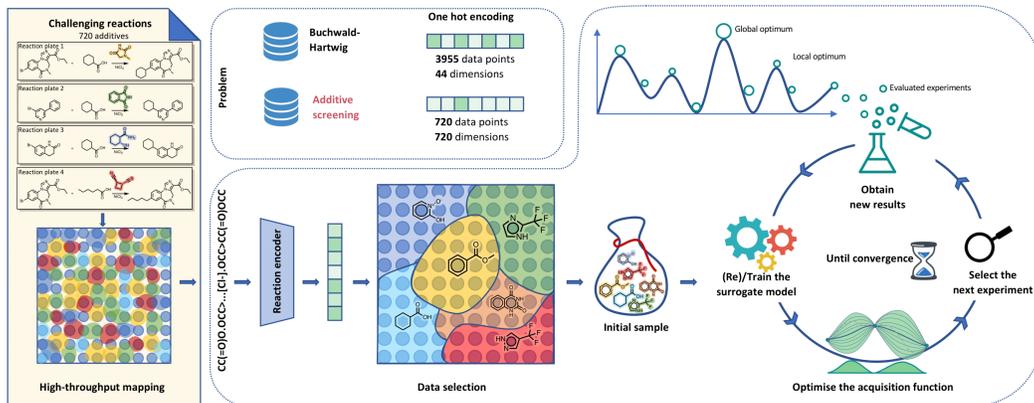


Figure 1: A visualisation of Bayesian optimisation for additive screening. Starting from the HTE dataset³³, we propagate the data through the reaction encoder to obtain suitable reaction representations. Using these representations, we organise the latent space and select the initial data points to set up the Gaussian process surrogate model. The BO loop then runs for a selected number of iterations during which we reach the global optimum in terms of the highest yield.

regime. Most BO studies report one-hot encodings (OHE) of the reaction components, where no chemical information is used, to perform remarkably well^{28,32} even compared to more elaborate quantum mechanical (QM) descriptors²⁹. While small, OHE still incorporate enough information for BO to read and exploit, whereas more complex representations bring additional noise.

In this work, we apply BO to explore an additive screening dataset containing 720 additives³³. The dataset consists of four plates with different Ni-catalysed photoredox decarboxylative arylations reactions. Additives are key for altering the reactivity and outcome of chemical reactions^{34,35}. While Prieto Kullmer et al.³³ used those additives through high-throughput experimentation, not all laboratories have access to robotic platforms. Still, synthetic chemists could benefit from efficiently finding the best additives through BO, i.e. improving a reaction without having to run all possible reactions (Figure 1). Compared to existing applications of BO to chemical reactions (e.g., Buchwald-Hartwig reactions with 44 OHE dimensions³⁶), the additive dataset is substantially more challenging. Firstly, using OHE is not suitable as it has a dimensionality of 720, with only one datapoint corresponding to each additive. Indeed, applying BO to this representation fails to improve over random search. Secondly, the additives are structurally more diverse than the components screened in other HTE studies, which makes the computation of human-labelled local QM descriptors more laborious. We overcome those limitations by using reaction fingerprints as a representation and a maximal diversity initialisation scheme. Using as little as 5/10/20 initialisation reactions, we demonstrate that BO efficiently finds high-yielding reactions in less than 100 single-point optimisation iterations.

2 Methods

2.1 Reaction representations

The most straightforward reaction representation is the OHE encoding of the individual reaction components. For instance, if the design space contains 15 aryl halides, 23 additives, 4 catalysts, and 3 bases³⁶, the components present in a reaction would be represented by a 1 in a binary reaction vector, and as a 0 otherwise. More chemically meaningful representations are mixtures of molecular and atomic QM descriptors as used by Ahneman et al.³⁶ and Shields et al.²⁹. However, the local atomic QM descriptors are limited to molecules with similar functional groups. Hence, those representations are not well suited for the additive screening dataset³³. Instead, we could represent reactions by isolating the additives and encoding them using molecular fingerprints³⁷⁻³⁹. Because the additives are the only varying part of the chemical reaction in our dataset, they uniquely represent each data point. The optimisation process remains the same, but by separating them from the rest of the reaction, we lose information intertwined in the reaction components.

Simplified molecular-input line-entry system (SMILES)⁴⁰ have emerged as one of the leading methods of representing the chemical space of diverse molecules and reactions. Using a molecular graph representation of the different components in a reaction, they impose a highly informative reaction depiction suitable for advanced machine-learning models based on textual representation⁴¹. Additionally, their easy conversion to a vectorised form makes them appealing to various computational models⁴²⁻⁴⁴. For instance, the reaction fingerprint by Schneider et al.⁴² is computed by subtracting the molecular fingerprints^{37,38} of the reactants from the ones of the products. More recently, Schwaller et al.⁴³ derived data-driven reaction fingerprints directly from the reaction SMILES (RXNFP) and Probst et al.⁴⁴ introduced the differential reaction fingerprint (DRFP), which is based on the symmetric difference of two sets containing the circular molecular n -grams generated from the molecules listed left (reactants and reagents) and right (products) from the reaction arrow. We use the RXNFP and DRFP in our experiments.

2.2 Bayesian optimisation

Many problems in scientific discovery may be framed as global optimisation problems of the form

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (1)$$

where $f : \mathcal{X} \rightarrow \mathbb{R}$ is a function over a design space \mathcal{X} . In the case of molecular representations, the design space is discrete and structured, containing graphs, fingerprint vectors and strings. Equation 1 is a black-box optimisation problem as we do not know the analytic form of f or its gradients and may only query f pointwise. Furthermore, evaluations of f require laboratory experiments and so are high-cost and time-consuming. Lastly, our observations of f are subject to a (potentially heteroscedastic⁴⁵) noise process. BO is an adaptive strategy that has recently emerged as a powerful solution method for black-box optimisation problems with proven success in applications including machine learning hyperparameter optimisation⁴⁶, chemical reaction optimisation²⁹, protein design⁴⁷, and as a subcomponent in AlphaGo⁴⁸. The Appendix A.4 provides pseudocode for BO and more details on the algorithm.

2.3 Data initialisation

The BO algorithm relies on sample data to initialise its surrogate model and accelerate the discovery of the needed outcome. Using Gaussian processes as the surrogate models allows us to work in the low data regime due to the well-calibrated uncertainty estimates they provide. See⁴⁹ or⁵⁰ for a description of Gaussian process in the context of structural inputs. However, the common machine learning doctrine where more data brings more clarity still stands. When initiated with more data points, the surrogate model has a better overview of the underlying function of the data, so the uncertainty measures are more precise. On the other hand, the chemist’s incentive is to start the optimisation process early and preserve the time needed for running reactions in the laboratory. These two viewpoints collide, and there is a need for a balancing solution. Instead of randomly choosing experiments for initializing the surrogate model, we explore our search space of reactions and select a diverse sample of points achieving good coverage of the space. In particular, we select a random initial point in the latent space and look for its matching pair in terms of the maximal distance (i.e., Jaccard⁵¹). We repeat the process for the identified least similar pair, ensuring we removed its previous match from the distance matrix. This policy allows us to diversify the initial sample of data points and expand the set coverage.

3 Results & Discussion

Our results show that utilizing BO substantially accelerates the discovery of the optimal additives for the highest yield improvement. Most importantly, we can achieve these results with novel priors like the DRFP, an NLP-inspired reaction fingerprint⁴⁴.

Figure 2 provides an overview of BO performance compared to the random search on different reaction plates. We used the first reaction plate as a testing ground for model configurations and their influence on the yield. We ran the experiments on the remaining plates and demonstrated the compelling influence of BO for finding the best reaction conditions. While both DRFP and RXNFP

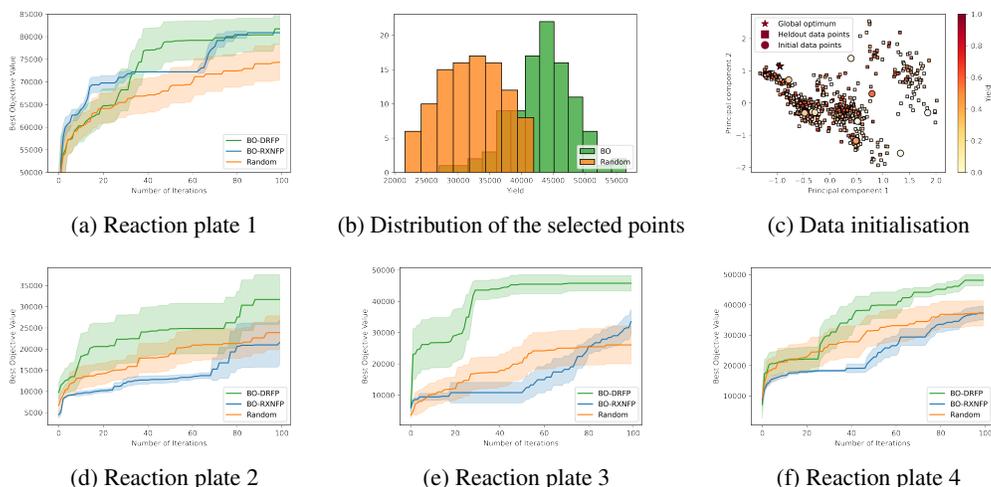


Figure 2: An overview of the results with different reaction representations. Figures a, d, e, f show the optimisation path towards the optimal parameters with mean observed value at each iteration over twenty different-seed trials (Gaussian Process surrogate model, 10 initialisation points). The shaded regions represent the 95% confidence interval. Figure b shows the distribution of the selected yield between random and BO search. Figure c depicts the latent space of DRFP reaction representations using principal component analysis and emphasises the selected initial data points together with the global optimum.

representations perform well on the first reaction plate, there is an evident divergence in their success when employed on the remaining reactions. Note that DRFP achieves the best results over all reaction plates. More importantly, we clearly differentiate the path that the BO with DRFP takes to reach the optimum, in contrast to the random search. While we cannot equate a chemist’s intuition with random exploration, relevant research²⁹ acknowledges that scientists show bias due to their expert knowledge. BO algorithms do not suffer from this favouritism and analyse the search space efficiently, offering practical suggestions to chemists. Figure 2b shows us that, on average, BO models choose parameters with higher output over all iterations compared to the random search. The distribution of BO-selected values is centred around the larger yield increase.

We define a top- n neighbourhood as a set of n reactions with the highest yield for each reaction plate. Searching for the best reaction inside this set is the main goal of the optimisation process. However, it is equally important to uncover a variety of reactions that bring a comparable yield improvement. Table 1 displays the percentage of the top- n yield reactions discovered over all twenty trials for different reaction plates (10 initialisation reactions). A more comprehensive table with additional parameters is presented in the A.4. The table shows that BO performs much better than the random search. The results suggest that BO can help chemists discover relevant conditions for reaction optimisation even in challenging settings such as this additive design space.

Reaction plate	Optimisation	top 1 [%]	top 3 [%]	top 5 [%]	top 10 [%]
Plate 1	BO	55	61.67	50	38
	Random	25	18.33	15	16.5
Plate 2	BO	65	61.67	58	39
	Random	5	10	15	14.5
Plate 3	BO	95	95	76	67
	Random	10	10	10	18
Plate 4	BO	80	76.67	76	54
	Random	10	20	19	19.5

Table 1: [%] of the top n -highest yields discovered during the 100 iterations of optimisation, starting from the 10 initial data points.

4 Conclusion

Bayesian optimisation is a powerful optimisation method that steers the exploration of the search space towards promising regions. It is especially valuable in the domain of chemistry as it saves time and resources and allows uncovering of unexplored but high-yielding chemical reactions²⁹. The main goal of this paper was to determine whether BO provides insights into a challenging additive screening HTE dataset³³, where OHE and QM descriptors are not well suited. We demonstrated that exploiting BO using an elaborate initialisation scheme and reaction fingerprints efficiently reveals key additives that lead to yield improvement in chemical reactions.

Acknowledgments

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.

References

- [1] Connor W Coley, Natalie S Eyke, and Klavs F Jensen. Autonomous discovery in the chemical sciences part i: Progress. *Angewandte Chemie International Edition*, 59(51):22858–22893, 2020.
- [2] Kjell Jorner, Anna Tomberg, Christoph Bauer, Christian Sköld, and Per-Ola Norrby. Organic reactivity from mechanism to machine learning. *Nature Reviews Chemistry*, 5(4):240–255, 2021.
- [3] Philippe Schwaller, Alain C Vaucher, Ruben Laplaza, Charlotte Bunne, Andreas Krause, Clemence Corminboeuf, and Teodoro Laino. Machine intelligence for chemical reaction space. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, page e1604, 2022.
- [4] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.
- [5] Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.
- [6] Marwin HS Segler, Thierry Kogej, Christian Tyrchan, and Mark P Waller. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS central science*, 4(1):120–131, 2018.
- [7] Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science*, 11(2):577–586, 2020.
- [8] Francesca Grisoni, Berend JH Huisman, Alexander L Button, Michael Moret, Kenneth Atz, Daniel Merk, and Gisbert Schneider. Combining generative artificial intelligence and on-chip synthesis for de novo drug design. *Science Advances*, 7(24):eabg3338, 2021.
- [9] Marwin HS Segler, Mike Preuss, and Mark P Waller. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018.
- [10] Tomasz Klucznik, Barbara Mikulak-Klucznik, Michael P McCormack, Heather Lima, Sara Szymkuć, Manishabrata Bhowmick, Karol Molga, Yubai Zhou, Lindsey Rickershauser, Ewa P Gajewska, et al. Efficient syntheses of diverse, medicinally relevant targets planned by computer and executed in the laboratory. *Chem*, 4(3):522–532, 2018.
- [11] Connor W Coley, Dale A Thomas III, Justin AM Lummiss, Jonathan N Jaworski, Christopher P Breen, Victor Schultz, Travis Hart, Joshua S Fishman, Luke Rogers, Hanyu Gao, et al. A robotic platform for flow synthesis of organic compounds informed by ai planning. *Science*, 365(6453):eaax1566, 2019.

- [12] Philippe Schwaller, Riccardo Petraglia, Valerio Zullo, Vishnu H Nair, Rico Andreas Haeuselmann, Riccardo Pisoni, Costas Bekas, Anna Iuliano, and Teodoro Laino. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical science*, 11(12):3316–3325, 2020.
- [13] Amol Thakkar, Thierry Kogej, Jean-Louis Reymond, Ola Engkvist, and Esben Jannik Bjerrum. Datasets and their influence on the development of computer assisted synthesis planning tools in the pharmaceutical domain. *Chemical science*, 11(1):154–168, 2020.
- [14] Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist, and Esben Bjerrum. Aizynthfinder: a fast, robust and flexible open-source software for retrosynthetic planning. *Journal of cheminformatics*, 12(1):1–9, 2020.
- [15] Barbara Mikulak-Klucznik, Patrycja Gołębiowska, Alison A Bayly, Oskar Popik, Tomasz Klucznik, Sara Szymkuć, Ewa P Gajewska, Piotr Dittwald, Olga Staszewska-Krajewska, Wiktor Beker, et al. Computational planning of the synthesis of complex natural products. *Nature*, 588(7836):83–88, 2020.
- [16] Jennifer N Wei, David Duvenaud, and Alán Aspuru-Guzik. Neural networks for the prediction of organic chemistry reactions. *ACS central science*, 2(10):725–732, 2016.
- [17] Marwin HS Segler and Mark P Waller. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry—A European Journal*, 23(25):5966–5971, 2017.
- [18] Connor W Coley, Regina Barzilay, Tommi S Jaakkola, William H Green, and Klavs F Jensen. Prediction of organic reaction outcomes using machine learning. *ACS central science*, 3(5):434–443, 2017.
- [19] Philippe Schwaller, Theophile Gaudin, David Lanyi, Costas Bekas, and Teodoro Laino. “found in translation”: predicting outcomes of complex organic chemistry reactions using neural sequence-to-sequence models. *Chemical science*, 9(28):6091–6098, 2018.
- [20] Connor W Coley, Wengong Jin, Luke Rogers, Timothy F Jamison, Tommi S Jaakkola, William H Green, Regina Barzilay, and Klavs F Jensen. A graph-convolutional neural network model for the prediction of chemical reactivity. *Chemical science*, 10(2):370–377, 2019.
- [21] Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A Hunter, Costas Bekas, and Alpha A Lee. Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction. *ACS central science*, 5(9):1572–1583, 2019.
- [22] 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, 2020.
- [23] Philippe Schwaller, Alain C Vaucher, Teodoro Laino, and Jean-Louis Reymond. Prediction of chemical reaction yields using deep learning. *Machine learning: science and technology*, 2(1):015016, 2021.
- [24] 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. *Chemical Engineering Journal*, 352:277–282, 2018.
- [25] Natalie S Eyke, William H Green, and Klavs F Jensen. Iterative experimental design based on active machine learning reduces the experimental burden associated with reaction screening. *Reaction Chemistry & Engineering*, 5(10):1963–1972, 2020.
- [26] Kobi C Felton, Jan G Rittig, and Alexei A Lapkin. Summit: benchmarking machine learning methods for reaction optimisation. *Chemistry-Methods*, 1(2):116–122, 2021.
- [27] 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, 2021.

- [28] Alexander Pomberger, AA Pedrina McCarthy, Ahmad Khan, Simon Sung, CJ Taylor, MJ Gaunt, Lucy Colwell, David Walz, and AA Lapkin. The effect of chemical representation on active machine learning towards closed-loop optimization. *Reaction Chemistry & Engineering*, 2022.
- [29] 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, 2021.
- [30] Pia Müller, Adam D Clayton, Jamie Manson, Samuel Riley, Oliver S May, Norman Govan, Stuart Notman, Steven V Ley, Thomas W Chamberlain, and Richard A Bourne. Automated multi-objective reaction optimisation: which algorithm should i use? *Reaction Chemistry & Engineering*, 7(4):987–993, 2022.
- [31] Jose Garrido Torres, Sii Hong Lau, Pranay Anchuri, Jason Stevens, Jose Tabora, Jun Li, Alina Borovika, Ryan Adams, and Abigail Doyle. A multi-objective active learning platform and web app for reaction optimization. 2022.
- [32] Riley Hickman, Jurgis Ruža, Loic Roch, Hermann Tribukait, and Alberto García-Durán. Equipping data-driven experiment planning for self-driving laboratories with semantic memory: case studies of transfer learning in chemical reaction optimization. 2022.
- [33] Cesar N Prieto Kullmer, Jacob A Kautzky, Shane W Krska, Timothy Nowak, Spencer D Dreher, and David WC MacMillan. Accelerating reaction generality and mechanistic insight through additive mapping. *Science*, 376(6592):532–539, 2022.
- [34] Ana Bellomo, Jiadi Zhang, Nisalak Trongsiwat, and Patrick J Walsh. Additive effects on palladium-catalyzed deprotonative-cross-coupling processes (dccp) of sp³ c–h bonds in diaryl-methanes. *Chemical Science*, 4(2):849–857, 2013.
- [35] Julien C Vantourout, Haralampos N Miras, Albert Isidro-Llobet, Stephen Sproules, and Allan JB Watson. Spectroscopic studies of the chan–lam amination: A mechanism-inspired solution to boronic ester reactivity. *Journal of the American Chemical Society*, 139(13):4769–4779, 2017.
- [36] 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, 2018.
- [37] Raymond E Carhart, Dennis H Smith, and R Venkataraghavan. Atom pairs as molecular features in structure-activity studies: definition and applications. *Journal of Chemical Information and Computer Sciences*, 25(2):64–73, 1985.
- [38] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- [39] Robin Winter, Floriane Montanari, Frank Noé, and Djork-Arné Clevert. Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations. *Chemical science*, 10(6):1692–1701, 2019.
- [40] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- [41] Philippe Schwaller, Benjamin Hoover, Jean-Louis Reymond, Hendrik Strobelt, and Teodoro Laino. Extraction of organic chemistry grammar from unsupervised learning of chemical reactions. *Science Advances*, 7(15):eabe4166, 2021.
- [42] Nadine Schneider, Daniel M Lowe, Roger A Sayle, and Gregory A Landrum. Development of a novel fingerprint for chemical reactions and its application to large-scale reaction classification and similarity. *Journal of chemical information and modeling*, 55(1):39–53, 2015.
- [43] Philippe Schwaller, Daniel Probst, Alain C Vaucher, Vishnu H Nair, David Kreutter, Teodoro Laino, and Jean-Louis Reymond. Mapping the space of chemical reactions using attention-based neural networks. *Nature Machine Intelligence*, 3(2):144–152, 2021.

- [44] Daniel Probst, Philippe Schwaller, and Jean-Louis Reymond. Reaction classification and yield prediction using the differential reaction fingerprint drfp. *Digital discovery*, 1(2):91–97, 2022.
- [45] Ryan-Rhys Griffiths, Alexander A Aldrick, Miguel Garcia-Ortegon, Vidhi Lalchand, et al. Achieving robustness to aleatoric uncertainty with heteroscedastic Bayesian optimisation. *Machine Learning: Science and Technology*, 3(1):015004, 2021.
- [46] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- [47] Henry Moss, David Leslie, Daniel Beck, Javier Gonzalez, and Paul Rayson. Boss: Bayesian optimization over string spaces. *Advances in neural information processing systems*, 33:15476–15486, 2020.
- [48] Yutian Chen, Aja Huang, Ziyu Wang, Ioannis Antonoglou, Julian Schrittwieser, David Silver, and Nando de Freitas. Bayesian optimization in AlphaGo. *arXiv preprint arXiv:1812.06855*, 2018.
- [49] Henry B. Moss and Ryan-Rhys Griffiths. Gaussian Process Molecule Property Prediction with FlowMO. *arXiv e-prints*, art. arXiv:2010.01118, October 2020.
- [50] Ryan-Rhys Griffiths, Leo Klarner, Henry Moss, Aditya Ravuri, Sang T. Truong, Bojana Rankovic, Yuanqi Du, Arian Rokkum Jamasb, Julius Schwartz, Austin Tripp, Gregory Kell, Anthony Bourached, Alex Chan, Jacob Moss, Chengzhi Guo, Alpha Lee, Philippe Schwaller, and Jian Tang. GAUCHE: A library for Gaussian processes in chemistry. In *ICML 2022 2nd AI for Science Workshop*, 2022. URL <https://openreview.net/forum?id=i9MKI7zrWa1>.
- [51] Paul Jaccard. The distribution of the flora in the alpine zone. 1. *New phytologist*, 11(2):37–50, 1912.
- [52] Christopher K Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.
- [53] James Wilson, Frank Hutter, and Marc Deisenroth. Maximizing acquisition functions for Bayesian optimization. *Advances in Neural Information Processing Systems*, 31:9884–9895, 2018.
- [54] Artur M Schweidtmann, Dominik Bongartz, Daniel Grothe, Tim Kerkenhoff, Xiaopeng Lin, Jaromil Najman, and Alexander Mitsos. Global optimization of Gaussian processes. *arXiv preprint arXiv:2005.10902*, 2020.
- [55] Antoine Grosnit, Alexander Imani Cowen-Rivers, Rasul Tutunov, Ryan-Rhys Griffiths, Jun Wang, and Haitham Bou-Ammar. Are we forgetting about compositional optimisers in Bayesian optimisation? *J. Mach. Learn. Res.*, 22:160–1, 2021.
- [56] Henry Moss, David Leslie, Daniel Beck, Javier Gonzalez, and Paul Rayson. BOSS: Bayesian optimization over string spaces. *Advances in neural information processing systems*, 33:15476–15486, 2020.
- [57] Greg Landrum, Paolo Tosco, Brian Kelley, Ric, sriniker, gedec, Riccardo Vianello, NadineSchneider, David Cosgrove, Eisuke Kawashima, Andrew Dalke, Dan N, Gareth Jones, Brian Cole, Matt Swain, Samo Turk, AlexanderSavelyev, Alain Vaucher, Maciej Wójcikowski, Ichiru Take, Daniel Probst, Kazuya Ujihara, Vincent F. Scalfani, guillaume godin, Axel Pahl, Francois Berenger, JLVarjo, strets123, JP, and DoliathGavid. rdkit/rdkit: 2022_09_1b1 (q3 2022) release. October 2022. doi: 10.5281/zenodo.7179566. URL <https://doi.org/10.5281/zenodo.7179566>.
- [58] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*,

A.3 Reproducibility

We implemented the code in Python. We used PyTorch⁵⁸, PyTorch Lightning⁵⁹, Gauche⁵⁰ and BoTorch⁶⁰. Different trials were set using a set of seeds ranging from 1 to 20 with PyTorch Lightning’s seed_everything function. The surrogate model used is a Gaussian process implemented with SingleTaskGP from Botorch. We used the Tanimoto kernel^{49,61} from Gauche. Reaction representations were defined with DRFP⁴⁴, RXNFP⁴³ and CDDD³⁹ libraries. The acquisition function is UpperConfidenceBound from BoTorch. We tested beta values [0.1, 0.5, 0.9]. The results presented in the main paper use UpperConfidenceBound with beta equaling to 0.1. We will release the code in open source.

A.4 Additional results for initialisation sample sizes

The Table A.4 shows additional results for different initialisation sample sizes.

Reactions	Init size	Optimisation	Beta	top-1 [%]	top-3 [%]	top-5 [%]	top-10 [%]	
Plate 1	5	BO	0.1	35	40	35	33	
			0.9	45	51.67	44	36.5	
		Random		5	13.33	14	15.5	
		10	BO	0.1	55	61.67	50	38
				0.9	55	61.67	50	37.5
		Random		20	20	20	18	
	20	BO	0.1	10	10	9	15	
			0.9	15	15	14	16.5	
		Random		10	20	18	16.5	
	Plate 2	5	BO	0.1	33.33	40.73	47.78	36.11
				0.9	25	33.33	43	34
Random				10.53	10.53	20	18.95	
10			BO	0.1	65	61.67	58	39
				0.9	55	50	48	33
Random				5	10	15	14.5	
20		BO	0.1	55	43.33	45	32	
			0.9	55	46.67	47	33.5	
		Random		10.53	10.53	20	18.95	
Plate 3		5	BO	0.1	60	58.33	51	47
				0.9	60	61.67	53	48
	Random			15	10	15	17.5	
	10		BO	0.1	95	95	76	67
				0.9	35	26.67	25	28
	Random			10	10	10	18	
	20	BO	0.1	75	75	60	53.5	
			0.9	75	75	61	54	
		Random		25	18.33	13	20	
	Plate 4	5	BO	0.1	80	46.67	40	28
				0.9	80	46.67	40	28.5
Random				10	21.67	16	13.5	
10			BO	0.1	80	76.67	76	54
				0.9	50	50	49	40
Random				10	20	19	19.5	
20		BO	0.1	65	28.33	21	24.5	
			0.9	60	26.67	20	24.5	
		Random		20	26.67	22	18.5	

Table 2: [%] of the top n -highest yields discovered during the 100 iterations of BO with different initial sample size and beta values using DRFP fingerprints. Comparison with the random search. Lower beta values correspond to data exploitation while higher values define an exploration strategy.