

Bayesian reaction optimization as a tool for chemical synthesis

<https://doi.org/10.1038/s41586-021-03213-y>

Received: 24 June 2020

Accepted: 11 December 2020

Published online: 3 February 2021

 Check for updates

Benjamin J. Shields¹, Jason Stevens², Jun Li², Marvin Parasram¹, Farhan Damani³, Jesus I. Martinez Alvarado¹, Jacob M. Janey², Ryan P. Adams³✉ & Abigail G. Doyle¹✉

Reaction optimization is fundamental to synthetic chemistry, from optimizing the yield of industrial processes to selecting conditions for the preparation of medicinal candidates¹. Likewise, parameter optimization is omnipresent in artificial intelligence, from tuning virtual personal assistants to training social media and product recommendation systems². Owing to the high cost associated with carrying out experiments, scientists in both areas set numerous (hyper)parameter values by evaluating only a small subset of the possible configurations. Bayesian optimization, an iterative response surface-based global optimization algorithm, has demonstrated exceptional performance in the tuning of machine learning models³. Bayesian optimization has also been recently applied in chemistry^{4–9}; however, its application and assessment for reaction optimization in synthetic chemistry has not been investigated. Here we report the development of a framework for Bayesian reaction optimization and an open-source software tool that allows chemists to easily integrate state-of-the-art optimization algorithms into their everyday laboratory practices. We collect a large benchmark dataset for a palladium-catalysed direct arylation reaction, perform a systematic study of Bayesian optimization compared to human decision-making in reaction optimization, and apply Bayesian optimization to two real-world optimization efforts (Mitsunobu and deoxyfluorination reactions). Benchmarking is accomplished via an online game that links the decisions made by expert chemists and engineers to real experiments run in the laboratory. Our findings demonstrate that Bayesian optimization outperforms human decisionmaking in both average optimization efficiency (number of experiments) and consistency (variance of outcome against initially available data). Overall, our studies suggest that adopting Bayesian optimization methods into everyday laboratory practices could facilitate more efficient synthesis of functional chemicals by enabling better-informed, data-driven decisions about which experiments to run.

Optimization of a chemical reaction is a complex, multidimensional challenge that requires experts to evaluate various reaction parameters, such as substrate, catalyst, reagent, additive, solvent, concentration, temperature and reactor type (Fig. 1a). Yet in a typical laboratory, bench chemists can evaluate only a small subset of these conditions during a standard optimization campaign owing to time and material limitations. Modern advances in high-throughput experimentation (HTE) have extended experimental capabilities to the collection of a few thousand data points under a limited set of conditions¹⁰. Thus, the chemist's art is to differentiate between millions of plausible configurations using a laboratory equipped to run only a tiny fraction of the possibilities. To do so, chemists typically carry out their experiments by scouring the chemical literature for similar reactions and intuiting the most influential dimensions (that is, reaction parameters) for reaction

success on the basis of experience, mechanistic understanding, empirical data and simple heuristics (Fig. 1b).

Chemists also commonly utilize systematic, model-driven approaches to reaction optimization^{11–13}. For example, design of experiments (DOE) seeks to sample experimental conditions that facilitate modelling of reaction parameters and deconvolution of interactions (Fig. 1b)^{14,15}. In conjunction with a response surface model, DOE enables the exploitation of knowledge gained from previous evaluations to guide the selection of future experiments. However, the exploration of reaction space is typically left in the hands of pre-defined optimal designs, sensitivity analysis, literature precedence and the operator's intuition¹. In addition, although a typical reaction requires the fine-tuning of numerous discrete parameters, screening requirements grow exponentially with the number of categorical components using

¹Department of Chemistry, Princeton University, Princeton, NJ, USA. ²Chemical Process Development, Bristol-Myers Squibb, New Brunswick, NJ, USA. ³Department of Computer Science, Princeton University, Princeton, NJ, USA. ✉e-mail: rpa@princeton.edu; agdoyle@princeton.edu




Fig. 1 | Bayesian reaction optimization. **a**, Example of chemical process optimization in pharmaceutical development: synthesis of BMS-911543. **b**, Typical approaches to reaction optimization. Mechanistic studies establish important reaction pathways, design of experiments enables systematic identification of important parameters (typically continuous variables such as time, t , temperature, T and concentration, C), high-throughput experimentation facilitates more rapid screening of reaction conditions, and

optimal designs. Thus, in practice some variables may be held constant to enable optimization on a fixed experimental budget^{16,17}.

The fundamental challenge associated with reaction optimization is not unique to chemistry. In machine learning, the development and study of computer algorithms that learn from data, practitioners are often tasked with finding the model hyperparameters that optimize performance². This meta-challenge has driven the development of automated approaches to algorithm optimization^{2,18}. Bayesian optimization, an uncertainty-guided response surface method used for the optimization of computationally expensive objective functions, has shown excellent performance, in many instances outperforming expert practitioners and other state-of-the-art global optimization algorithms (Fig. 1c)^{3,19}. Bayesian optimization is designed to balance the exploration of areas of uncertainty and the exploitation of available information, leading to high-quality configurations in fewer evaluations. Importantly, Bayesian optimization algorithms can be applied to diverse search spaces that include arbitrary parameterized reaction domains and enables the selection of multiple experiments in parallel. Accordingly, this approach is well suited to the optimization of chemical processes^{4,20–23}. However, Bayesian optimization has only recently become of interest to the chemical community. Select applications include automatic chemical design⁵, high-throughput virtual screening²⁴ and programmed flow chemistry^{6,25}. Although researchers have begun to explore the application of machine learning methods to reaction optimization^{4,6–8,25–27}, these efforts have targeted a limited subset of synthetic chemistry that includes only continuous process parameters. That is, to the best of our knowledge, there are (1) no applications to typical batch chemistry, (2) no general-purpose

response surface models predict high-yielding conditions from available data (the example surface is colour-coded according to objective values).

c, Graphical overview of Bayesian optimization. One-dimensional example depicting a Gaussian process surrogate model fitted to data collected from an unknown objective and the corresponding expected utility surface, which is maximized to select the next experiment. The surrogate model is plotted as the posterior mean, with the shaded region showing 2σ units.

software platforms that are easily accessible to non-experts and (3) no systematic comparisons to the performance of expert chemists²⁸.

Herein we report the development of a modular framework for Bayesian reaction optimization and accompanying open-source software that is compatible with automated systems (for example, computer experiments) and human-in-the-loop experimentation (for example, bench-scale screening). Our approach is designed to integrate with existing synthetic-chemistry practices, is applicable to arbitrary search spaces that include continuous and categorically encoded reactions, and enables the inclusion of physics and domain expertise. Moreover, this method is arbitrarily parallelizable (that is, any number of experiments can be selected per batch of iterative experiments) and thus can facilitate both rapid screening and direct translation to large-scale process conditions. First, we describe the development of our Bayesian optimization platform using experimental reaction data mined from the literature. Next, to test our approach, we undertake a systematic investigation of Bayesian optimization compared to human decision-making in the optimization of a new reaction—inspired by a key Pd-catalysed direct arylation from the synthesis of BMS-911543 (Fig. 1a). Finally, we demonstrate the general applicability of Bayesian optimization to the optimization of additional distinct chemical reactions relevant to medicinal chemistry.

Optimizer development

For a given search space, Bayesian reaction optimization begins by collecting initial reaction outcome data via an experimental design

(for example, DOE or at random) or by drawing from existing results (Fig. 1c). These data are used to train a probabilistic surrogate model, which is constructed by combining previous observations with a prior over functions that captures our assumptions about the reaction's response surface (such as smoothness and experimental noise), making it possible to reason about the position of the global optimum. After training the surrogate model, new experiments in the reaction space are sequentially chosen by optimizing an acquisition function that maximizes the expected utility of candidate experiments for the next evaluation (Fig. 1c). Finally, the proposed experiments are carried out, the results are appended to the record and the surrogate model posterior is updated. This process continues iteratively until the reaction yield is maximized, the resources are depleted or the space is explored to a degree that finding improved conditions is improbable.

We began our studies by developing a flexible python package for Bayesian reaction optimization named Experimental Design via Bayesian Optimization (EDBO). In the development of EDBO, we emphasized a simple but modular interface, the use of arbitrary user-defined reaction spaces and applicability to human-in-the-loop or automated optimization. We provide the source code and examples as an open resource for the chemical community (see 'Data availability' and 'Code availability' sections).

With the software architecture in hand, our next objective was to optimize its performance by tuning algorithm components critical to maximizing yield in reaction optimization. In particular, we sought an optimizer configuration that offered both good average performance and outcomes with low variance with respect to the initial data available to the optimizer. Accordingly, we elected to carry out reaction optimizations with different random initial starting data and to select the optimizer configuration that gave low average loss, low variance in outcome and low worst-case loss. With these metrics in mind, we found that good optimization performance could be achieved with the available reaction data using a reaction space encoded with density functional theory (DFT; auto-qchem—see 'Data availability' and 'Code availability' sections), a Gaussian process surrogate model²⁹ and parallel expected improvement³⁰ as an acquisition function (see below).

To meet our experimental design, we curated reaction data from the literature for optimizer development and evaluation (Fig. 2). We selected palladium-catalysed cross-coupling data for Suzuki–Miyaura (1)³¹ and Buchwald–Hartwig reactions (2a–2e)³², in which the objective is to optimize the yield of the desired product with respect to a combinatorial set of hundreds or thousands of possible reaction conditions.

Although chemists have a superb capacity to understand reactions described graphically, Bayesian optimization requires a structured representation. We explored the use of chemical-descriptor fingerprint-based reaction encodings based on the quantum chemical properties of reaction components computed via DFT³², cheminformatics descriptors generated using open-source libraries³³ and binary one-hot-encoded (OHE) representations. After tuning the optimizer for each data type independently, we found that the average loss for parallel reaction optimization using each encoding was largely indistinguishable ($p > 0.05$, where p refers to the p-value associated with the Welch t -test for unequal variance—see Methods; Extended Data Table 1). However, DFT-encoded descriptors gave the most consistent results in terms of worst-case loss ($\leq 5\%$ yield for all reactions, versus $\leq 15\%$ and $\leq 8\%$ for Mordred and OHE, respectively). Thus, we elected to carry out the remainder of our experiments using DFT descriptors. Importantly, our findings suggest that acceptable performance can be achieved in the wild with a number of reaction encodings.

Next, we investigated the performance of various surrogate models. The most basic requirements of an effective surrogate model are the ability to make predictions and to estimate variance. Thus, in principle Bayesian optimization can be implemented with many different models, depending on the problem. For example, for optimizations over discrete or mixed domains, the response surface may be better characterized by

a random forest model³⁴. For continuous domains, it is typical to assume that the unknown function is sampled from a Gaussian process³⁵. Whatever the case, building an efficient optimizer is itself an optimization problem. That is, to achieve good performance, the surrogate model must be optimized. Here, we selected surrogate model parameters based on regression performance for reactions 1 and 2a–e. Thus, we elected to take ourselves out of the loop and turn the optimizer on itself. Over the course of our studies we found that a Gaussian process model²⁹ with the Matérn52 kernel offered superior optimization performance over a random forest model³⁶ in terms of mean loss, outcome variance and worst-case loss (Extended Data Table 1).

After training the surrogate model, new experiments in the reaction space are sequentially chosen by optimizing an acquisition function (Fig. 1c). The central tenet of Bayesian optimization (and active learning methods³⁷ in general) is the utilization of both information and uncertainty to drive optimization. Thus, it is instructive to consider the limiting cases. As interpolation methods, constant mean Gaussian processes tend to predict the highest yield in the neighbourhood around the current best observed value. Accordingly, for a non-concave reaction surface (containing local maxima) we anticipate that pure exploitation, exemplified by an acquisition function that selects the point of highest predicted yield for evaluation, could become trapped in local maxima. By contrast, a pioneering acquisition function (pure exploration), exemplified by selecting the point of greatest predictive uncertainty for evaluation, will tend to investigate the entire response surface more thoroughly. Although the explorer will not necessarily find a global maximum without evaluating the entire search space, we expect it to achieve the best global understanding of the reaction surface. To demonstrate this dichotomy, we tracked the exploiter's and explorers' decisions²⁵, after initialization at the same point, in a two-dimensional representation of reaction 1 (Fig. 3a). Indeed, over the first ten evaluations, the exploiter remained in a single cluster while the explorer traversed the entire space, visiting all five larger groups identified via k -means clustering. Next, we tracked the understanding of the surrogate model associated with each acquisition function by measuring its fit to the overall space (Fig. 3b). Over the course of 50 experiments the scores of the explorer and exploiter diverge, with the explorer offering a better fit to the reaction surface. Finally, we considered the yield of each reaction investigated by the two algorithms (Fig. 3b), finding that most points selected by the explorer and exploiter gave low and high yields, respectively.

In practice, acquisition functions derived from a utility that balances exploration and exploitation often give improved performance in non-concave optimization. Importantly, this utility can be a written function of the model posterior distribution, which makes it inexpensive to evaluate and enables information gathering to be explicitly incorporated in candidate selections. For example, expected improvement³⁰, one of several frequently employed analytic acquisition functions, can be maximized to select a candidate reaction that, on average, is expected to improve upon the current best result to the greatest extent³⁸. Alternatively, as a Gaussian process represents a distribution over functions, one can draw and maximize candidate models that fit the data where information is available and vary according to function shape and estimated variance in areas where little is known—a process called Thompson sampling^{39,40}. Maximizing the expected improvement or sampling from the distribution over optima implied by the Gaussian process posterior naturally balances the exploitation of areas of high predicted mean response and the exploration of areas with high marginal variance. Indeed, in comparison to pure exploitation and exploration, the path of the expected improvement in optimizing reaction 1 visits three of the five clusters, and the corresponding surrogate model achieves an intermediate fit to the reaction response surface (Fig. 3a, b).

Bayesian optimization is typically formulated as a sequential problem^{19,41,42}. However, for many reaction optimization problems, running




Fig. 2 | Training data used to select Bayesian optimizer parameters. Suzuki reaction dataset (1) consisting of the reaction yield measured as a function of boronic acid derivative, aryl halide, ligand, base and solvent. Buchwald–Hartwig reaction datasets (2a–e) consisting of the reaction yield measured as a

function of aryl halide, Pd precatalyst, additive and base. Parameters that define the reaction space are highlighted in blue. Heatmaps show reaction yield for experiments in each dataset.

experiments in parallel is critical because time presents a substantial cost (many reactions take hours or days to run to completion). Thus, we next sought to investigate methods for the selection of arbitrary batches of experiments to run in parallel. Thompson sampling naturally lends itself to selecting batches of experiments via the sampling of N candidate response surfaces from the posterior predictive distribution of a Gaussian process surrogate model (Supplementary Fig. 29). However, to achieve parallel decision-making for analytic acquisition functions, we iteratively predict the experiments that maximize the acquisition function, at each step taking the surrogate model of the previous iterations on faith, and factoring its prediction of the most likely outcome into the next selected experiment (Supplementary Fig. 30). This approach, sometimes called the Kriging believer algorithm⁴¹, can be thought of as a chess player thinking ahead multiple moves, at each step inferring the probable consequences of their actions. We found that these basic algorithms offer promising performance in the parallel setting (Fig. 3, Supplementary Figs. 38, 39). Notably, for reactions 1 and 2a–e it was observed that parallel optimization (batch size, 5) performed equally well on average ($p > 0.05$) to sequential optimization (batch size, 1) with a 50-experiment budget (Extended Data Table 1).

Having established effective acquisition and batching strategies, we evaluated parallel optimization performance of various acquisition functions designed to balance exploration and exploitation (Fig. 3; see Supplementary Information for additional acquisition functions). Overall, we found that both parallel expected improvement and Thompson sampling offered excellent performance, and that their average outcomes were statistically indistinguishable ($p > 0.05$ for all six reactions). However, the variance in outcome across simulation runs and the worst-case loss were greater with Thompson sampling. Importantly, the performance of expected improvement was remarkably consistent; over 30 random initializations, it converged within a narrow margin of the optimal solution (worst-case loss $\leq 5\%$ yield).

Performance benchmarking

We next sought to statistically evaluate the performance of DOE methods compared to Bayesian optimization (Extended Data Table 1;

see Supplementary Information for details). Although DOE is most frequently used for the optimization of continuous parameters, we identified two designs that have been employed effectively for the optimization of chemical processes with categorical variables: generalized subset designs (GSD) and D-optimal designs^{16,43}. For each of the reactions in the development set, these DOE-based optimizations deviated from Bayesian optimization in both mean outcome ($p < 0.05$), standard deviation (Bayesian optimization, ≤ 1.9 ; GSD, ≤ 6.9 ; D-optimal, ≤ 3.3) and worst-case loss (Bayesian optimization, ≤ 5 ; GSD, ≤ 16 ; D-optimal, ≤ 15). Thus, we conclude that, other things equal, Bayesian optimization is both more straightforward to apply and provides superior performance in reaction optimization with categorical variables.

With our Bayesian optimization framework tuned for reaction optimization, we next sought to statistically test its performance in a new reaction space. Palladium-catalysed C–H functionalization has garnered increasing interest in pharmaceutical development for its ability to generate molecular complexity without the need for pre-functionalized starting material^{44–46}. The direct functionalization of heterocycles represents a particularly attractive reaction owing to their prevalence in bioactive compounds⁴⁷. However, the functionalization of a given heteroarene substrate often requires refinement of the reaction conditions to achieve optimal reactivity and selectivity. Here we investigate the direct arylation of imidazoles, exemplified by reaction 3 (Fig. 4), which is related to a key step in the commercial synthesis of the JAK2 inhibitor BMS-911543 (Fig. 1a)^{48,49}.

Reaction optimization truly begins by defining the search space. To facilitate the exhaustive evaluation of experimental conditions for statistical validation, we first considered a larger set of plausible experiments, then quantified similarities between potential reaction conditions via unsupervised learning and selected those that we expected would give a satisfactory distribution across the larger search space (see Supplementary Information for details)⁵⁰. We anticipated that choosing an appropriate ligand, base, solvent, temperature and concentration would be critical to achieving an optimal reaction yield based on variable importance for historical Bristol–Myers–Squibb (BMS) direct arylation studies. In this study, we chose 12 ligands from a larger set of 70 potential phosphines (Fig. 4). Importantly, the final selection of




Fig. 3 | Balancing exploration and exploitation in reaction optimization. **a**, t -Distributed stochastic neighbour embedding (t-SNE) of the search space of reaction 1, showing the first ten decisions made in sequential optimization by exploitation, exploration and Bayesian optimization (BO) with expected improvement (EI). Colours denote larger clusters identified via k -means clustering. t-SNE1 and t-SNE2 correspond to the first and second dimensions of the embedding, respectively. **b**, Model fit score for predicting the entire response surface and yields for experiments selected by the optimizers in **a**. **c–h**, Average ($N = 30$) cumulative maximum observed yield for optimizations

with DFT encodings, Gaussian process surrogate models, initial experiment chosen at random and five experiments selected per round. Reactions 1 and 2a–e are shown. Inset tables summarize the outcome distribution statistics, including: Welch's t -test p -values for the null hypothesis of equal average performance to expected improvement (p), standard deviation in maximum achieved yield (σ) and worst-case loss (ζ). Error bars on the expected improvement denote the 0.5σ interval. ϵ -Greedy randomization probability, 0.1. TS, Thompson sampling.

reaction components was based on valuable expert knowledge rather than machine learning. Overall, we selected a subspace consisting of 1,728 reactions including 12 ligands, four bases, four solvents, three temperatures and three concentrations (Fig. 4) as a tractable set of experiments to be used as ground truth.

Next, we collected experimental results for the entire search space via HTE (Fig. 4). Then, to benchmark the performance of Bayesian optimization to that of human experts, we developed a game^{25,51} that would track the decisions made by chemists of different backgrounds and levels of experience when optimizing reaction 3 (see Supplementary Information). Although the game was intended to simulate reaction optimization on a fixed experimental budget, the data were real. Each experiment 'run' returned the actual result of the corresponding experiment in the laboratory. In the game, the participants had 'one month' to find optimal conditions for reaction 3 with the capacity to run one batch of five experiments 'per workday'. Participants 'ran' their

experiments via a web application that returned results and tracked their decisions.

In total, 50 expert chemists and engineers from academia and industry played the reaction optimization game (Fig. 4c). Accordingly, the Bayesian reaction optimizer also played the game 50 times (Fig. 4b), each time starting with a different random initialization. The first point of comparison between human participants and the machine learning optimizer was their raw maximum observed yield at each step during the optimization. Humans made significantly ($p < 0.05$) better initial choices than random selection, on average discovering conditions that had 15% higher yield in their first batch of experiments. However, even with random initialization, within three batches of five experiments the average performance of the optimizer surpassed that of the humans. Notably, in contrast to human participants, Bayesian optimization achieved >99% yield 100% of the time within the experimental budget. Moreover, Bayesian optimization tended to discover globally optimal




Fig. 4 | Statistical validation of Bayesian reaction optimization. **a**, Reaction space and experimental yields collected via HTE for the direct arylation reaction (3). Data were collected via HTE without replication. **b**, Machine learning performance; Bayesian optimization (BO) curves for random initializations (dashed) and average (solid). **c**, Human performance;

optimization curves for individual players (dashed) and average (solid). **d**, Best and worst-case imputation of human performance. **e**, Hypothesis testing, with $p < 0.05$ highlighted in grey. μ_{human} , human average maximum observed yield; μ_{BO} , Bayesian optimizer average maximum observed yield.

statistically different. For the raw data and lower bound, we then infer that on average after the fifth batch of experiments the performance of the optimizer is better than that of humans. By contrast, for the upper bound we found that there is no statistically significant difference between the two central tendencies. Thus, we conclude that in the optimization of reaction 3, Bayesian reaction optimization on average outperformed human experts, tracing the unrealistic imputed upper bound of the recorded data from the games.

statistically different. For the raw data and lower bound, we then infer that on average after the fifth batch of experiments the performance of the optimizer is better than that of humans. By contrast, for the upper bound we found that there is no statistically significant difference between the two central tendencies. Thus, we conclude that in the optimization of reaction 3, Bayesian reaction optimization on average outperformed human experts, tracing the unrealistic imputed upper bound of the recorded data from the games.

Applications

Having validated our approach statistically, we next sought to carry out real-world test cases of Bayesian optimization for the optimization of reactions relevant to pharmaceutical development. Importantly, we chose to apply Bayesian optimization over larger reaction spaces in which exhaustive collection of experimental data via HTE is not possible.

The prevalence of aliphatic alcohols makes them ideal building blocks for the synthesis of complex molecules. Accordingly, heteroatom substitution reactions that directly utilize alcohols play an essential role in medicinal chemistry⁵². The Mitsunobu reaction⁵³ in particular is frequently used owing to the diverse set of nucleophiles that can undergo stereospecific coupling with aliphatic alcohols⁵⁴. However, the standard conditions typically afford only modest yields. Thus, the well defined yet expansive array of potential reagents make the Mitsunobu reaction an ideal test case for Bayesian optimization.




Fig. 5 | Applications of Bayesian reaction optimization. Optimization of a Mitsunobu reaction (4) and a deoxyfluorination reaction (5). The standard reaction conditions (entry 1, red; average of two replicate experiments) and three top-yielding conditions identified by Bayesian optimization (entries 2–4;

single experiments) are listed in the tables. Plots depict the cumulative best observed yield (black), individual experiment outcomes (grey) and the yield for standard reaction conditions (red dashed) as a function of experiment batch (batch sizes: Mitsunobu, 10; fluorination, 5).

We selected the coupling of methyl 3-bromo-1H-indole-6-carboxylate and benzyl alcohol with a reaction space defined by a combinatorial set of six azadicarboxylates, 12 phosphines and five solvents (reaction 4; Fig. 5). In addition, we identified a grid of substrate concentrations, azadicarboxylate equivalents, phosphine equivalents and temperatures as continuous process parameters to give a reaction space consisting of 180,000 possible configurations (see Supplementary Information for a list of reaction parameters).

With the search space in hand, we next carried out controls using reaction conditions most frequently employed at BMS: 1.1 equiv. DIAD, 1.1 equiv. PPh₃, THF 0.1 M, and 25 °C. These standard reaction parameters gave an average of 60% yield over two replicate experiments (59% and 60%). We next carried out Bayesian reaction optimization using DFT encoding, a Gaussian process surrogate model and expected improvement as an acquisition function, running ten experiments in parallel per batch, with initial experiments chosen at random. Notably, we found that the optimizer quickly surpassed the benchmark result, identifying three distinct sets of reaction conditions that produced the desired product in 99% yield in only four rounds of ten experiments (Fig. 5). The top-yielding experiments used unconventional conditions, including P(Ph)₂Me, high concentrations and elevated temperatures. Thus, the optimizer identified quantitative conditions in areas of reaction space that would not typically be searched.

Fluorination of organic compounds has an important role in drug discovery owing to the unique properties of carbon–fluorine bonds^{55,56}. In this context, deoxyfluorination of alcohols is one of the most widely employed methods for the synthesis of aliphatic fluorides⁵⁷. In a previous study, the Doyle group reported that the tuning of reagent structure can enable the efficient fluorination of numerous complex alcohols using sulfonyl fluorides⁵⁸. Thus, this reaction presented an ideal test case for Bayesian optimization. We selected the fluorination of 1-(6-nitrobenzo[d][1,3]dioxol-5-yl)ethan-1-ol with a reaction space defined by a combinatorial set of ten sulfonyl fluorides, ten organic bases, five solvents and a grid of continuous parameters (substrate concentrations, sulfonyl fluoride equivalents, base equivalents and temperatures) to give a reaction space consisting of 312,500 possible configurations (see Supplementary Information for a list of reaction parameters).

We next carried out controls using reaction conditions typically used for the commercial reagent PyFluor (1.1 equiv. PyFluor, 1.1 equiv. DBU, THF 0.5 M and 20 °C)⁵⁹. These standard conditions gave an average of 36% yield over two replicate experiments (35% and 36%). We next carried out Bayesian reaction optimization using DFT encoding, a Gaussian process surrogate model and expected improvement as an acquisition function, running five experiments in parallel per batch, with initial experiments chosen at random. We found that within three rounds

of five experiments the optimizer surpassed the benchmark result, ultimately identifying reaction conditions that produced the desired product in 69% yield in ten rounds of experiments (Fig. 5).

Importantly, in both test reactions, Bayesian optimization identified sets of experimental conditions with largely distinct parameter settings from the standard conditions. In addition, the optimizer delivered multiple configurations, which varied in most dimensions but delivered equivalent good results.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-021-03213-y>.

- Carlson, R. *Design and Optimization in Organic Synthesis* (Elsevier, 1992).
- Luo, G. A review of automatic selection methods for machine learning algorithms and hyper-parameter values. *Netw. Model. Anal. Health Inform. Bioinform.* **5**, 18 (2016).
- Snoek, J., Larochelle, H. & Adams, R. P. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems* Vol. 25 (eds Pereira, F. et al.) 2951–2959 (Curran Associates Inc., 2012).
- Häse, F., Roch, L. M., Kreisbeck, C. & Aspuru-Guzik, A. Phoenix: a Bayesian Optimizer for Chemistry. *ACS Cent. Sci.* **4**, 1134–1145 (2018).
- Griffiths, R.-R. & Hernández-Lobato, J. M. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chem. Sci.* **11**, 577–586 (2020).
- Schweidtmann, A. M. et al. Machine learning meets continuous flow chemistry: automated optimization towards the Pareto front of multiple objectives. *Chem. Eng. J.* **352**, 277–282 (2018).
- Burger, B. et al. A mobile robotic chemist. *Nature* **583**, 237–241 (2020).
- Häse, F., Roch, L. M. & Aspuru-Guzik, A. Gryffin: an algorithm for Bayesian optimization for categorical variables informed by physical intuition with applications to chemistry. Preprint at <https://arxiv.org/abs/2003.12127> (2020).
- Negoescu, D. M., Frazier, P. I. & Powell, W. B. The knowledge-gradient algorithm for sequencing experiments in drug discovery. *INFORMS J. Comput.* **23**, 346–363 (2011).
- Santanilla, A. B. et al. Nanomole-scale high-throughput chemistry for the synthesis of complex molecules. *Science* **347**, 49–53 (2014).
- Clayton, A. D. et al. Algorithms for the self-optimisation of chemical reactions. *React. Chem. Eng.* **4**, 1545–1554 (2019).
- Häse, F., Roch, L. M. & Aspuru-Guzik, A. Next-generation experimentation with self-driving laboratories. *Trends Chem.* **1**, 282–291 (2019).
- Weissman, S. A. & Anderson, N. G. Design of experiments (DoE) and process optimization. A review of recent publications. *Org. Process Res. Dev.* **19**, 1605–1633 (2015).
- Lee, R. Statistical design of experiments for screening and optimization. *Chem. Ing. Tech.* **91**, 191–200 (2019).
- Murray, P. M. et al. The application of design of experiments (DoE) reaction optimisation and solvent selection in the development of new synthetic chemistry. *Org. Biomol. Chem.* **14**, 2373–2384 (2016).
- Hsieh, H.-W., Coley, C. W., Baumgartner, L. M., Jensen, K. F. & Robinson, R. I. Photoredox iridium–nickel dual-catalyzed decarboxylative arylation cross-coupling: from batch to continuous flow via self-optimizing segmented flow reactor. *Org. Process Res. Dev.* **22**, 542–550 (2018).
- Mateos, C., Nieves-Remacha, M. J. & Rincón, J. A. Automated platforms for reaction self-optimization in flow. *React. Chem. Eng.* **4**, 1536–1544 (2019).
- Feurer, M. & Hutter, F. In *Automated Machine Learning: Methods, Systems, Challenges* (eds Hutter, F. et al.) 3–33 (Springer, 2019).
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. & de Freitas, N. Taking the human out of the loop: a review of Bayesian optimization. *Proc. IEEE* **104**, 148–175 (2016).
- Maceiczky, R. M. & deMello, A. J. Fast and reliable metamodelling of complex reaction spaces using Universal Kriging. *J. Phys. Chem. C* **118**, 20026–20033 (2014).
- Rogers, A. & Ierapetritou, M. Feasibility and flexibility analysis of black-box processes part 1: surrogate-based feasibility analysis. *Chem. Eng. Sci.* **137**, 986–1004 (2015).
- Boukouvala, F. & Ierapetritou, M. G. Feasibility analysis of black-box processes using an adaptive sampling Kriging-based method. *Comput. Chem. Eng.* **36**, 358–368 (2012).
- Olofsson, S., Hebing, L., Niedenfür, S., Deisenroth, M. P. & Misener, R. GPdoemd: a Python package for design of experiments for model discrimination. *Comput. Chem. Eng.* **125**, 54–70 (2019).
- Krivák, R., Hoksza, D. & Škoda, P. Improving quality of ligand-binding site prediction with Bayesian optimization. In *2017 IEEE International Conference on Bioinformatics and Biomedicine* 2278–2279 (2017).
- Reker, D., Hoyt, E. A., Bernardes, G. J. L. & Rodrigues, T. Adaptive optimization of chemical reactions with minimal experimental information. *Cell Rep. Phys. Sci.* **1**, 100247 (2020).
- Zhou, Z., Li, X. & Zare, R. N. Optimizing chemical reactions with deep reinforcement learning. *ACS Cent. Sci.* **3**, 1337–1344 (2017).
- Kondo, M. et al. Exploration of flow reaction conditions using machine-learning for enantioselective organocatalyzed Rauhut–Currier and [3+2] annulation sequence. *Chem. Commun.* **56**, 1259–1262 (2020); correction **56**, 12256–12256 (2020).
- Ueno, T., Rhone, T. D., Hou, Z., Mizoguchi, T. & Tsuda, K. COMBO: an efficient Bayesian optimization library for materials science. *Mater. Discov.* **4**, 18–21 (2016).
- Gardner, J., Pleiss, G., Weinberger, K. Q., Bindel, D. & Wilson, A. G. GPYtorch: blackbox matrix–matrix Gaussian process inference with GPU acceleration. In *Advances in Neural Information Processing Systems* Vol. 31 (eds Bengio, S. et al.) 7576–7586 (Curran Associates Inc., 2018).
- Mockus, J. On the Bayes methods for seeking the extremal point. *IFAC Proc.* **8**, 428–431 (1975).
- Perera, D. et al. A platform for automated nanomole-scale reaction screening and micromole-scale synthesis in flow. *Science* **359**, 429–434 (2018).
- Ahnenman, 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* **360**, 186–190 (2018).
- Moriwaki, H., Tian, Y.-S., Kawashita, N. & Takagi, T. Mordred: a molecular descriptor calculator. *J. Cheminform.* **10**, 4 (2018).
- Biau, G. Analysis of a random forests model. *J. Mach. Learn. Res.* **13**, 1063–1095 (2012).
- Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (MIT Press, 2006).
- Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011); <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestRegressor.html>
- Reker, D. & Schneider, G. Active-learning strategies in computer-assisted drug discovery. *Drug Discov. Today* **20**, 458–465 (2015).
- Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).
- Kandasamy, K., Krishnamurthy, A., Schneider, J. & Póczos, B. Parallelised Bayesian optimisation via Thompson sampling. In *International Conference on Artificial Intelligence and Statistics* 133–142 (2018).
- Hernández-Lobato, J. M., Requeima, J., Pyzer-Knapp, E. O. & Aspuru-Guzik, A. Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. Preprint at <https://arxiv.org/abs/1706.01825> (2017).
- Ginsbourger, D., Le Riche, R. & Carraro, L. In *Computational Intelligence in Expensive Optimization Problems* (eds Tenne, Y. & Goh, C.-K.) 131–162 (Springer, 2010).
- Wang, J., Clark, S. C., Liu, E. & Frazier, P. I. Parallel Bayesian global optimization of expensive functions. *Oper. Res.* **68**, 1850–1865 (2020).
- Surowiec, I. et al. Generalized subset designs in analytical chemistry. *Anal. Chem.* **89**, 6491–6497 (2017).
- Davies, H. M. L. & Morton, D. Recent advances in C–H functionalization. *J. Org. Chem.* **81**, 343–350 (2016).
- Lyons, T. W. & Sanford, M. S. Palladium-catalyzed ligand-directed C–H functionalization reactions. *Chem. Rev.* **110**, 1147–1169 (2010).
- Alberico, D., Scott, M. E. & Lautens, M. Aryl–aryl bond formation by transition-metal-catalyzed direct arylation. *Chem. Rev.* **107**, 174–238 (2007).
- Vitaku, E., Smith, D. T. & Njardarson, J. T. Analysis of the structural diversity, substitution patterns, and frequency of nitrogen heterocycles among U.S. FDA approved pharmaceuticals. *J. Med. Chem.* **57**, 10257–10274 (2014).
- Fox, R. J. et al. C–H Arylation in the formation of a complex pyrrolopyridine, the commercial synthesis of the potent JAK2 inhibitor, BMS-911543. *J. Org. Chem.* **84**, 4661–4669 (2019).
- Ji, Y. et al. Mono-oxidation of bidentate bis-phosphines in catalyst activation: kinetic and mechanistic studies of a Pd/xantphos-catalyzed C–H functionalization. *J. Am. Chem. Soc.* **137**, 13272–13281 (2015).
- Durand, D. J. & Fey, N. Computational ligand descriptors for catalyst design. *Chem. Rev.* **119**, 6561–6594 (2019).
- Duros, V. et al. Human versus robots in the discovery and crystallization of gigantic polyoxometalates. *Angew. Chem. Int. Ed.* **56**, 10815–10820 (2017).
- Swamy, K. C. K., Kumar, N. N. B., Balaraman, E. & Kumar, K. V. P. Mitsunobu and related reactions: advances and applications. *Chem. Rev.* **109**, 2551–2651 (2009).
- Mitsunobu, O. & Yamada, M. Preparation of esters of carboxylic and phosphoric acid via quaternary phosphonium salts. *Bull. Chem. Soc. Jpn* **40**, 2380–2382 (1967).
- Fletcher, S. The Mitsunobu reaction in the 21st century. *Org. Chem. Front.* **2**, 739–752 (2015).
- Gillis, E. P., Eastman, K. J., Hill, M. D., Donnelly, D. J. & Meanwell, N. A. Applications of fluorine in medicinal chemistry. *J. Med. Chem.* **58**, 8315–8359 (2015).
- Hagmann, W. K. The many roles for fluorine in medicinal chemistry. *J. Med. Chem.* **51**, 4359–4369 (2008).
- Hu, W.-L., Hu, X.-G. & Hunter, L. Recent developments in the deoxyfluorination of alcohols and phenols: new reagents, mechanistic insights, and applications. *Synthesis* **49**, 4917–4930 (2017).
- Nielsen, M. K., Ahnenman, D. T., Riera, O. & Doyle, A. G. Deoxyfluorination with sulfonyl fluorides: navigating reaction space with machine learning. *J. Am. Chem. Soc.* **140**, 5004–5008 (2018).
- Nielsen, M. K., Ugaz, C. R., Li, W. & Doyle, A. G. PyFluor: a low-cost, stable, and selective deoxyfluorination reagent. *J. Am. Chem. Soc.* **137**, 9571–9574 (2015).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2021

Methods

Reaction encodings

For each reaction, a numerical encoding was generated by concatenating descriptor vectors for each chemical component and continuous variable. For example, for experiment i in reaction 3 the corresponding descriptors are $\hat{\mathbf{d}}_i = \hat{\mathbf{d}}_{\text{ligand}} \oplus \hat{\mathbf{d}}_{\text{base}} \oplus \hat{\mathbf{d}}_{\text{solvent}} + T + C$ (where \oplus denotes concatenation). Then the encoded representation of the N -experiment reaction space is the matrix of descriptors, with each row corresponding to an experiment: $R = \langle \hat{\mathbf{d}}_1, \hat{\mathbf{d}}_2, \dots, \hat{\mathbf{d}}_N \rangle$. The descriptor matrix was then preprocessed by removing highly correlated features (Pearson correlation coefficient > 0.95) and normalizing each feature on the unit hypercube. The numbers of descriptors before and after preprocessing for each reaction are displayed in Extended Data Table 2.

For chemical reagents, the open-source molecular descriptor calculation software Mordred was used to calculate two- and three-dimensional chemical informatics-based descriptors³³. OHE descriptors were computed by representing the presence or absence of a given reaction component with 1 or 0, respectively.

To generate DFT encodings, initial molecular geometry optimization and conformational analysis were carried out using Open Babel 2.4.1⁶⁰. Quantum mechanical modelling was then carried out using Gaussian 16 A.03⁶¹. For all DFT calculations, the B3LYP hybrid exchange-correlation functional was used. Gas-phase geometry optimization, frequency- and time-dependent DFT calculations were carried out using the 6-31G* basis set and the LANL2DZ basis set for heavy atoms not supported by 6-31G*. Electronic and steric descriptors for the global structure (for example, dipole moment and molar volume), atomic properties (for example, charge and buried volume) of overlapping substructures and atoms of minimum/maximum charge were then extracted or computed from the computational output. When conformational analysis was carried out (reactions 1 and 3, free ligands; 4, phosphines, azadicarboxylates; 5, sulfonyl fluorides, bases), descriptors corresponding to the minimum energy conformer, maximum energy conformer and standard deviation of the conformers were computed. Then for a given descriptor d , the weighted average (\bar{d}) of the conformer descriptors was computed according to Maxwell-Boltzmann statistics:

$$\bar{d} = \sum_{i=1}^M \frac{e^{-G_i/k_B T}}{s} d_i \quad (1)$$

where M is the number of conformers, G_i is the Gibbs free energy associated with conformer i , k_B is Boltzmann's constant, T is the absolute temperature and $s = \sum_{j=1}^N e^{-G_j/k_B T}$.

Surrogate model design

A Gaussian process surrogate model³⁵, implemented in python using GPyTorch²⁹, was employed for Bayesian optimization. $\text{GP}(\mu, k_\theta)$ represents a distribution over functions characterized by a prior mean and kernel $k_\theta(\mathbf{x}_1, \mathbf{x}_2)$. We model chemical reaction outcomes, standardized to zero mean and unit variance, using a constant mean Gaussian process and the Matérn52 kernel:

$$k_{\text{Matérn52}}(r) = \alpha \left(1 + \frac{\sqrt{5}r}{l} + \frac{5r^2}{3l^2} \right) e^{-\frac{\sqrt{5}r}{l}}, \quad (2)$$

where α is an output scale parameter, $r = \sqrt{\sum_{i=1}^n (\mathbf{x}_{1i} - \mathbf{x}_{2i})^2}$ is the distance between points \mathbf{x}_1 and \mathbf{x}_2 , and l is a length scale parameter. When conditioned on experimental observations, the Gaussian process posterior distribution mean μ at point \mathbf{x} is given by:

$$\mu(\mathbf{x}) = k_\theta(\mathbf{x})^T (K_\theta + \sigma_n^2 I)^{-1} \mathbf{y} \quad (3)$$

where $k_\theta(\mathbf{x})$ is the vector of covariances between \mathbf{x} and the training points, K_θ is the matrix of covariances between all training points, σ_n^2 is the estimated noise variance, I is the identity matrix and \mathbf{y} is a vector of responses corresponding to the training data. Then, the variance in the Gaussian process posterior distribution at \mathbf{x} is given by:

$$\sigma^2(\mathbf{x}) = k_\theta(\mathbf{x}, \mathbf{x}) - k_\theta(\mathbf{x})^T (K_\theta + \sigma_n^2 I)^{-1} k_\theta(\mathbf{x}), \quad (4)$$

which only depends on the estimated noise and covariance.

In Gaussian process regression, the identity of the kernel determines the general shape of its function distribution. Then, learned parameters for the length scale set the relative variation per dimension, the amplitude calibrates the magnitude of the changes, and the noise captures the variation in measurements. From the placement of l in the denominators of (2), it follows that the length scale acts to marginalize the distance between points. That is, small and large values of l tend to imply larger and smaller variations on a given interval, respectively. Accordingly, we found that learning the length scale per dimension via automatic relevance determination³⁵ improved model performance. However, in the absence of sufficient data, the inclusion of a parameter per dimension can lead to overfitting. Thus, we assume that most of the dimensions are irrelevant, impose this belief via a gamma prior centred on longer length scales, and estimate parameters by maximizing the marginal likelihood:

$$\log p(\mathbf{y}|\theta) = -\frac{1}{2} \mathbf{y}^T (K_\theta + \sigma_n^2 I)^{-1} \mathbf{y} - \frac{1}{2} \log |K_\theta + \sigma_n^2 I| - \frac{n}{2} \log(2\pi) \quad (5)$$

with respect to the kernel parameters and noise amplitude using stochastic gradient descent⁶² as implemented in PyTorch⁶³. Hyperparameter priors were included in the estimation by adding the logarithmic probability for the hyperparameter value specified by the prior distribution to (5). This procedure affords models that capture general trends in the low-data regime and reaction nuances when more evidence is available. The Matérn kernel shape parameter (5/2) and priors were selected via Bayesian optimization using EDBO.

Acquisition function

Acquisition functions are derived from a utility function that maps a candidate experiment \mathbf{x} , the corresponding outcome y and the model hyperparameters to a measurement of the value of the experiment. By using available data from experiments, we marginalize out the unknown result y and obtain the expected utility of evaluating to corresponding experiment \mathbf{x} . We identified the expected improvement³⁰ as a good acquisition function for reaction optimization based on simulations with reactions 1 and 2a-e. The expected improvement is derived from the improvement utility:

$$\mathbb{I}(\mathbf{x}) = \begin{cases} f(\mathbf{x}) - f^+ & f(\mathbf{x}) > f^+ \\ 0 & f(\mathbf{x}) \leq f^+ \end{cases} \quad (6)$$

which is the increase in the objective function value $f(\mathbf{x})$ over the current best observed reaction outcome f^+ . Then, the expectation value of \mathbb{I} for a given experiment \mathbf{x} has the form:

$$\mathbb{E} \mathbb{I}(\mathbf{x}) = \begin{cases} I(\mathbf{x}) \Phi \left(\frac{I(\mathbf{x})}{\sigma(\mathbf{x})} \right) + \sigma(\mathbf{x}) \varphi \left(\frac{I(\mathbf{x})}{\sigma(\mathbf{x})} \right) & \sigma(\mathbf{x}) > \delta \\ 0 & \sigma(\mathbf{x}) \leq \delta \end{cases} \quad (7)$$

where $I(\mathbf{x}) = \mu(\mathbf{x}) - f^+ - \delta$ is the improvement of the surrogate mean prediction $\mu(\mathbf{x})$ diminished by δ , an empirical exploration parameter, $\sigma(\mathbf{x})$ is the surrogate standard deviation, and Φ and φ are the cumulative distribution function and probability density function of the standard normal distribution, respectively. The parameter δ was set to a value of 0.01 on the basis of simulations with reactions 1 and 2a-e.

Selecting the next experiment for evaluation is done by optimizing $El(\mathbf{x})$. Because the experiments are carried out on a finite grid, the resulting reaction space X is finite, and optimization is carried out by explicitly computing $\text{argmax}_{\mathbf{x} \in X} El(\mathbf{x})$. Parallel acquisition is carried out using the Kriging believer algorithm by iteratively computing $\mathbf{x}_i \leftarrow \text{argmax}_{\mathbf{x} \in X} El(\mathbf{x})$, appending the Gaussian process posterior mean $\mu(\mathbf{x}_i)$ to the known data and updating the Gaussian process posterior.

Summary statistics

The reported p-values correspond to the null hypothesis $H_0: \bar{y}_1 = \bar{y}_2$ that the central tendencies \bar{y}_j of two optimization outcome populations are identical. Specifically, p corresponds to the Welch t -test for unequal variance:

$$t = \frac{\bar{y}_1 - \bar{y}_2}{\sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}}$$

where the degrees of freedom associated with sample variance σ_j^2 is computed with the Welch–Satterthwaite equation and n_j is the sample size for sample j . The average loss during optimization:

$$\mu_{\text{loss}} = \frac{1}{n} \sum_{i=1}^n (y^* - y_i)$$

is the mean difference between the global maximum in yield y^* and the highest yielding experiment sampled by the optimizer y_i for each of the n simulations. The reported standard deviation:

$$\sigma = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2}$$

is a measure of the variation in optimization outcome over n simulations. The reported worst-case loss:

$$\xi = \max(y^* - \mathbf{Y})$$

is the maximum difference between the global maximum in yield and the best configuration sampled by the optimizer in each simulation, here written as a vector of simulation results, \mathbf{Y} .

Materials and synthetic procedures

Standard glovebox techniques were employed for handling air-sensitive reagents. All reagents and materials were acquired from commercial suppliers or prepared using published procedures. Data for reaction 3 were collected via HTE using 96-well plates for screening. Experiments were carried out in a glovebox in which bases were automatically dispensed to each 96-well plate using Unchained Labs Powder Protégé, and the remainder of the reagents were dispensed by hand using stock solutions. The reaction yield was determined via UHPLC-MS using 4,4'-di-*tert*-butylbiphenyl as an internal standard. Data for reaction 4 were collected in batches of ten via standard screening practices. Experiments were set up in the glovebox and then transferred to the benchtop, where they were heated and stirred. Data for reaction 5 were

collected in batches of five via standard benchtop screening practices. Experiments were set up on the benchtop without the exclusion of air. The reaction yield was determined via ^{19}F nuclear magnetic resonance using 1-fluoronaphthalene as an external standard.

Data availability

Quantum mechanical computation data and Gaussian output files used to parameterize reactions 1–5 are available at <https://github.com/b-shields/auto-QChem>. Processed reaction outcome data for reactions 1–5 are available at <https://github.com/b-shields/edbo> and in our published Code Ocean capsule at <https://doi.org/10.24433/CO.3864629.v1>. Tabulated player data for the reaction optimization game are available at <https://github.com/b-shields/EvML>.

Code availability

Two software packages and one web application were written to support this work. The first, *auto-qchem*, was written to facilitate high-throughput computational chemistry and reaction featurization. This package is freely available at <https://github.com/b-shields/auto-QChem>. The second, *EDBO*, was written as a user-friendly implementation of Bayesian optimization. This package is freely available at <https://github.com/b-shields/edbo> and in our published Code Ocean capsule at <https://doi.org/10.24433/CO.3864629.v1>. The web application, *EvML*, was written to collect user data for comparison of Bayesian optimization with human expert performance. This package is freely available at <https://github.com/b-shields/EvML>.

- O'Boyle, N. M. et al. Open Babel: an open chemical toolbox. *J. Cheminform.* **3**, 33 (2011).
- Frisch, M. J. et al. *Gaussian 16* Revision A.03 (Gaussian, Inc., 2016).
- Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. Preprint at <https://arxiv.org/abs/1412.6980> (2017).
- Paszke, A. et al. PyTorch: an imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* Vol. 32 (eds Wallach, H. et al.) 8026–8037 (Curran Associates Inc., 2019).

Acknowledgements Financial support was provided by Bristol-Myers Squibb, the Princeton Catalysis Initiative, the NSF under the CCI Center for Computer Assisted Synthesis (CHE-1925607) and the DataX Program at Princeton University through support from the Schmidt Futures Foundation. We thank A. Żurański and J. Ash for discussions. We thank all the participants in the reaction optimization game for their time and effort in contributing to this study. We thank B. Hao for help with HTE protocols.

Author contributions B.J.S. designed the overall research project with A.G.D., R.P.A. and J.J. providing guidance. B.J.S. wrote and ran the software with the assistance of F.D. and input from J.L.; J.J. and J.S. carried out the initial investigation to select the test reaction; J.S. designed and carried out HTE experiments with the assistance of M.P.; J.L. wrote the web application for the reaction optimization game with the assistance of J.S., J.J. and B.J.S.; and B.J.S. carried out data experiments and modelling with input from J.L. and F.D. J.S. and J.M.A. carried out Mitsunobu and deoxyfluorination reaction optimizations. B.J.S. wrote the manuscript with input from all authors.

Competing interests The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-021-03213-y>.

Correspondence and requests for materials should be addressed to R.P.A. or A.G.D.

Peer review information *Nature* thanks Jason Hein and Tiago Rodrigues for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Extended Data Table 1 | Simulation outcome summary for reactions 1 and 2a–e

Reaction Encodings												
Reaction	DFT				Mordred				OHE			
	ρ	μ_{loss}	σ	ξ	ρ	μ_{loss}	σ	ξ	ρ	μ_{loss}	σ	ξ
1	1.000	1	0.9	4	0.425	1	1.2	4	0.413	1	0.7	3
2a	1.000	0	0.5	1	0.241	0	0.7	2	0.044	1	0.7	2
2b	1.000	0	0.1	0	0.006	3	4.6	12	0.174	0	0.4	1
2c	1.000	1	1.1	3	0.899	1	1.1	4	0.269	1	0.9	3
2d	1.000	0	0.5	2	0.124	0	0.5	2	0.409	0	0.8	5
2e	1.000	2	1.9	5	0.881	2	3.0	15	0.089	1	1.8	8

Sequential Optimization and Surrogate Models												
Reaction	Batch Size = 1				Random Forest				Bayesian Linear			
	ρ	μ_{loss}	σ	ξ	ρ	μ_{loss}	σ	ξ	ρ	μ_{loss}	σ	ξ
1	0.259	2	1.3	6	0.053	2	1.2	5	0.000	4	2.7	10
2a	0.046	0	0.1	0	0.197	0	0.7	2	0.002	2	2.4	9
2b	0.346	0	1.9	11	0.059	1	2.3	11	0.001	5	6.8	19
2c	0.112	2	2.0	11	0.409	1	1.3	3	0.000	4	4.1	17
2d	0.445	0	0.4	1	0.025	0	0.3	1	0.000	2	2.5	10
2e	0.805	2	1.9	5	0.740	2	1.8	4	0.000	5	3.4	17

Design of Experiments												
Reaction	Random				Generalized Subset				D-Optimal			
	ρ	μ_{loss}	σ	ξ	ρ	μ_{loss}	σ	ξ	ρ	μ_{loss}	σ	ξ
1	0.004	3	2.4	9	0.000	4	1.8	8	0.003	4	3.3	15
2a	0.003	2	1.7	5	0.001	2	2.2	9	0.000	2	1.8	7
2b	0.004	3	4.2	14	0.000	8	6.9	16	0.066	1	2.8	10
2c	0.006	3	2.5	12	0.002	4	3.0	12	0.012	2	1.8	6
2d	0.004	2	2.0	8	0.001	2	2.0	8	0.004	2	1.5	6
2e	0.004	5	4.1	17	0.000	5	3.4	16	0.000	4	1.8	6

Distribution ($N = 30$) of cumulative maximum observed yields for optimizations using 50 experiments. Reaction encoding simulations were carried out using a Gaussian process surrogate model, expected improvement as an acquisition function, initial experiments chosen at random and five experiments selected per round (batch size, 5). Sequential optimization simulations were carried out with DFT encoding, a Gaussian process surrogate model, expected improvement, the initial experiment chosen at random and batch size 1. Surrogate model simulations were carried out using expected improvement, initial experiments chosen at random and batch size 5. DOE experiments were carried out $N = 20$ times by shuffling the ordering of categorical variables, using OHE descriptors and a polynomial response surface model, with initial experiments selected according to the experimental design and the remainder of the experiments selected by evaluating the top predicted yields. Statistics: Welch's t -test p -values (ρ) for the null hypothesis of equal average performance to Bayesian optimization, with DFT encoding, a Gaussian process surrogate model, expected improvement, the initial experiment chosen at random, and batch size 5; mean loss (ground truth—best achieved yield) at the end of the optimization (μ_{loss}); standard deviation in maximum achieved yield (σ); and worst-case loss (ξ).

Article

Extended Data Table 2 | Summary of reaction encodings

Reaction	Dimensions	Search Space Size	DFT	Mordred	OHE
1	5	3696	1325 (180)	4008 (381)	30 (30)
2a	4	790	1170 (158)	3768 (399)	33 (33)
2b	4	792	1166 (158)	3788 (396)	33 (33)
2c	4	790	1165 (159)	3788 (397)	33 (33)
2d	4	792	1161 (158)	3746 (397)	33 (33)
2e	4	791	1152 (158)	3746 (396)	33 (33)
3	5	1728	1759 (291)	-	-
4	7	180000	1081 (278)	-	-
5	7	312500	876 (259)	-	-

Dimensions denotes the number of optimization dimensions; search space size is the number of experiments in the reaction space; and DFT, Mordred and OHE show the number of descriptors in each encoding before (after) preprocessing.