Accelerated Chemical Reaction Optimization using Multi-Task Learning

Connor J. Taylor[1,2‡]*, Kobi C. Felton[‡3], Daniel Wigh[2,3], Mohammed I. Jeraal[4], Rachel Grainger[1], Gianni Chessari[1], Christopher N. Johnson[1], Alexei A. Lapkin[2,3,4]*

[1]Astex Pharmaceuticals, 436 Cambridge Science Park, Milton Road, Cambridge, CB4 0QA, UK. [2]Innovation Centre in Digital Molecular Technologies, Yusuf Hamied Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge, CB2 1EW, UK. [3]Department of Chemical Engineering and Biotechnology, University of Cambridge, Philippa Fawcett Dr, Cambridge CB3 0AS, UK. [4]Cambridge Centre for Advanced Research and Education in Singapore Ltd., 1 Create Way, CREATE Tower #05-05, 138602, Singapore.

[‡]Authors contributed equally.

*Corresponding authors: Connor.Taylor@astx.com, aal35@cam.ac.uk.

Abstract

Functionalization of C–H bonds is a key challenge in medicinal chemistry, particularly for fragment-based drug discovery (FBDD) where such transformations need to be executed in the presence of polar functionality necessary for fragment-protein binding. New technologies such as high-throughput experimentation and self-optimization have the potential to revolutionize synthetic approaches to challenging target molecules in FBDD. Recent work has shown the effectiveness of Bayesian optimization (BO) for the self-optimization of chemical reactions, however, in all previous cases these algorithmic procedures have started with no prior information about the reaction of interest. In this work, we explore the use of multi-task Bayesian optimization (MTBO) in several *in silico* case studies by leveraging reaction data collected during related historical optimization campaigns to accelerate the optimization of new reactions - this was performed for Suzuki-Miyaura and C–N couplings. This methodology was then translated to real-world, medicinal chemistry applications in the yield optimization of several pharmaceutical intermediates using an autonomous flow-based reactor platform. The use of the MTBO algorithm was shown to be successful in determining optimal conditions (both continuous and categorical variables) of unseen experimental C–H activation reactions with differing substrates, demonstrating up to a 98 % cost reduction when compared to industry-standard process optimization techniques. Our findings highlight the effectiveness of the methodology as an enabling tool in medicinal chemistry workflows, where efficient utilization of precious starting materials is particularly important. This work represents a step-change in the utilization of previously obtained reaction data and machine learning with the ultimate goal of accelerated reaction optimization.

Introduction

In recent years, there has been an increased interest in the use of automated, self-optimizing continuous flow platforms for the optimization of chemical processes.[1-5] These platforms use automated reactors and machine-learning algorithms to learn from previous experiments, and thereby choose future experiments that ultimately maximize yield and/or other process objectives. The use of self-optimizing platforms has arisen from the desire for faster reaction optimization, improved process sustainability and cheaper overall process development.[6-9] The use of these platforms aims to enhance the capabilities of the researcher by removing the need for repetitive and labor-intensive experimentation, allowing them to focus on more challenging tasks. By leveraging algorithms, the platforms utilize only minimal reaction material but gain the most process information possible, making their deployment in fine chemical and pharmaceutical industries very attractive.[10, 11]

Recent work has shown that Bayesian optimization is a particularly powerful tool for self-optimization applications.[12-14] However, in all previous studies, each optimization begins with no *a priori* information about the chemical landscape for the reaction of interest. This protocol therefore requires initial experimental iterations whereby the algorithm is learning about the experimental design space, without any prior information on where the optimal reaction conditions may be. This initial exploration can be expensive in terms of both cost and time, particularly when there may already be data on the broad chemical transformation of interest from previous optimization campaigns. This also opens the methodology to some uncertainty over which initial design strategy to use, such as forms of factorial design or Latin Hypercube sampling (LHS), which will affect the overall experimental budget. *General* well-performing reaction conditions can also be identified for particular reaction classes, as highlighted in recent work by Angello *et al.*,[15] but do not give optimal conditions for specific transformations and cannot account for important parameters such as reactor heterogeneity, solubility, reaction selectivity, differences in substrate functionality or further adjacent objectives (E-factor, purification, downstream processing etc.).

For many medicinal chemistry applications, such as in developing chemistries for the synthesis of potential drug candidates, the use of efficient optimization techniques is paramount due to the minimal quantity and increased preciousness of intermediate reaction materials.[16, 17] This is a particular problem in fragment-based drug discovery (FBDD),[18, 19] as challenging transformations are often required on highly functionalized molecules - including difficult C–C forming reactions utilizing precise C–H activations which ideally must be executable in the presence of polar groups that are required for binding to the target protein.[20] The excessive material consumption when utilizing existing algorithms may also be a reason that medicinal chemists are less attracted to these cutting-edge optimization techniques than process chemists. Our hypothesis is that optimization strategies that can utilize pre-existing chemical knowledge could mitigate unnecessary material use, accelerate process development and present the potential for broader applicability in new synthetic chemistry methods.

This work shows the first real-world examples of leveraging previous reaction optimization data for unseen chemical transformations using multi-task Bayesian optimization (MTBO), with our prior work on MTBO for chemistry only showcasing its use in a simulated setting.[21] The framework of MTBO, first introduced by Swersky *et al.*,[22] replaces the standard probabilistic model in Bayesian optimization with a multi-task model. As these multi-task models can be trained on data from related tasks, we can therefore utilize data from previously conducted similar reactions - both from the laboratory and from the literature. In this work, we first explore and benchmark the use of MTBO in simulated studies, then exploit the methodology to optimize several pharmaceutically relevant C–H activation reactions using an autonomous flow reactor platform. These experimental case studies were chosen to highlight the effectiveness of MTBO in a medicinal chemistry, particularly FBDD, context through efficient material usage. There are many reports of similar automated workflows in the recent literature where a self-optimization protocol is utilized.[23, 24] Our reactor platform is equipped with a liquid handling robot and can optimize both continuous variables (residence time, temperature etc.) *and* categorical variables (solvent, ligand etc.). This ability is seldom reported in the literature (with some notable examples from several research groups[1, 25-27]), likely due to engineering and equipment challenges, but is very important in determining all relevant parameters that influence reaction outcomes. The MTBO algorithm utilized is integrated into the open-source reaction optimization package Summit[28] and represents a powerful data-driven optimization technique that can utilize known reaction data and ultimately lead to savings in material, time and overall cost.

Results and Discussion

Bayesian Optimization to Multi-Task Bayesian Optimization

As shown in Figure 1a, Bayesian optimization (BO) relies on three key components: a probabilistic model, an acquisition function and an optimization algorithm.[29] The probabilistic model is trained using experimental data and acts as a surrogate or "simulation" of the real chemical reaction. Given this probabilistic model, the acquisition function estimates the values of different potential experimental reaction conditions. The optimization algorithm is then used to find the set of experimental conditions that maximizes the acquisition function, and these experimental conditions are hence suggested as the next real experiment to run. The combination of the probabilistic model and the acquisition function enables exploitation of known high performance areas and exploration of new chemical space. By iteratively executing the suggested experimental conditions, retraining the model and optimizing the acquisition function, the BO protocol progressively identifies the best reaction conditions for the output of interest.

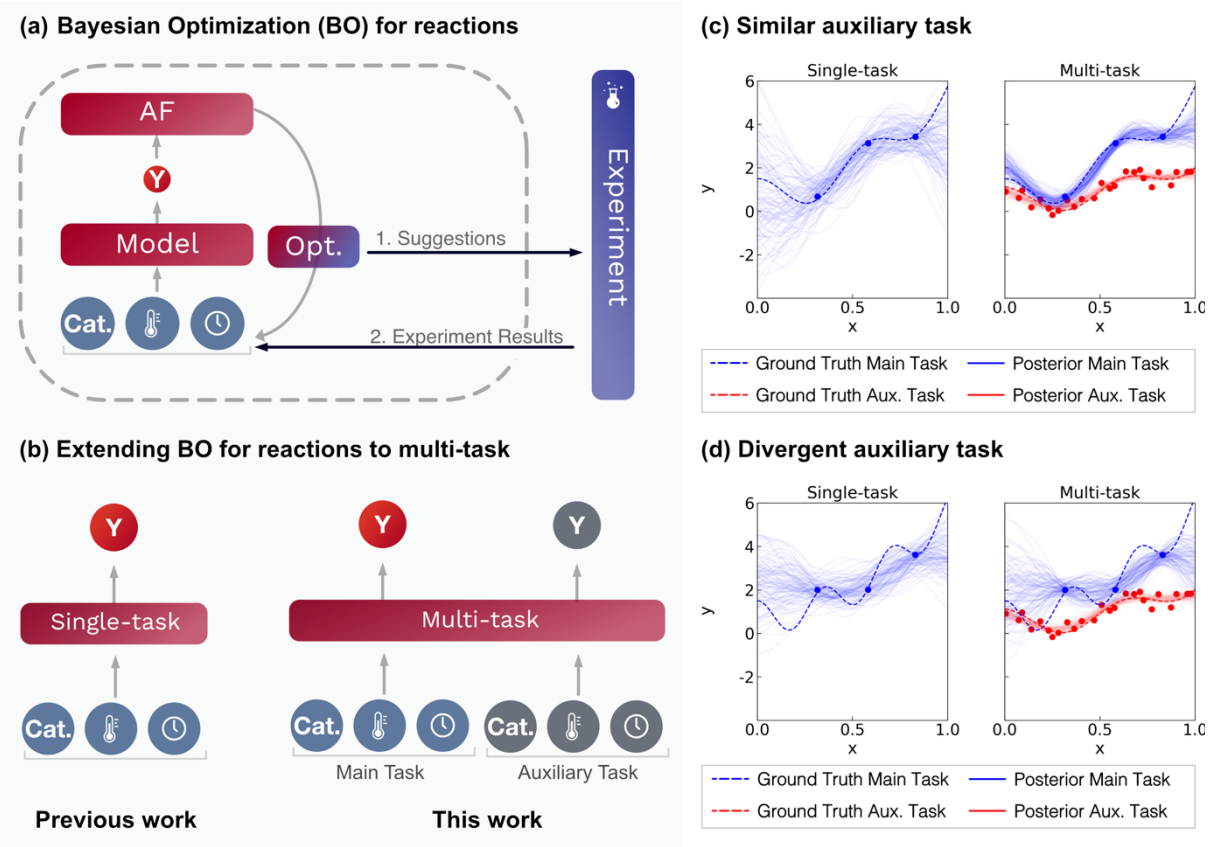

Figure 1: A schematic description of multi-task Bayesian optimization to the context of reaction optimization. **(a)** Bayesian optimization consists of a probabilistic model (typically a Gaussian process) that predicts experiment outputs (e.g. yield) given experiment conditions; an acquisition function (AF) that predicts the value of potential new experiments; and an optimization algorithm (opt). **(b)** Multi-task Bayesian optimization replaces the Gaussian process with a multi-task Gaussian process trained simultaneously on an auxiliary task. In our case, this auxiliary task is a similar reaction to the one being optimized, utilizing previous experimental results. **(c)** When the auxiliary task for a multi-task Gaussian process is similar to the main optimization task, predictions on the main task are improved significantly. **(d)** When the auxiliary task for a multi-task Gaussian process is divergent to the main optimization task, predictions on the main task are similar to what is observed for the baseline single-task Gaussian process.

As shown in Figure 1b, MTBO changes the probabilistic model in BO. Typically, a Gaussian process (GP) is used as the probabilistic model in BO due to the general applicability and efficiency of GPs in the small data regime.[30] MTBO replaces a GP with a multi-task GP that can learn the correlations between different tasks to enable better predictions. In our case, the tasks are chemical

transformations from the same reaction class with varying substrates. A formal definition of GPs and multi-task GPs is in the Methods section.

As a simple illustration of the benefits of multi-task GPs, we created example functions with one input and one output, then trained both a GP and multi-task GP on only three data points. In the multi-task GP case, we also generated 25 data points from an auxiliary task. As shown in Figure 1c, when the main and the auxiliary tasks are similar, the predictions from the multi-task GP (shown as samples from the posterior of the GP) more accurately represent the underlying function than the predictions from the single-task GP. The multi-task GP leverages covariance between the data in the two tasks to improve predictions on the main task, even with limited data for the main task - this is shown formally in the Methods section. As shown in Figure 1d, when the main and the auxiliary tasks are divergent, predictions from the single-task and multi-task GP are highly variable. However, this variability in the multi-task GP is still useful because the BO algorithm will explore to better capture the underlying distribution of the main task.

In Silico Case Studies: Suzuki-Miyaura Couplings

We first executed *in silico* MTBO studies using model chemical reactions as benchmarks. These models were generated using neural networks trained on literature experimental data that predict reaction yield;[28] more detail on these models can be found in the Methods section. The model 'Suzuki B1' was trained using Suzuki cross-coupling data from Baumgartner *et al.*,[31] while the models 'Suzuki R1-4' were trained using data from Reizman *et al.*[32] - these specific transformations and the variables that affect these models are shown in Scheme 1.

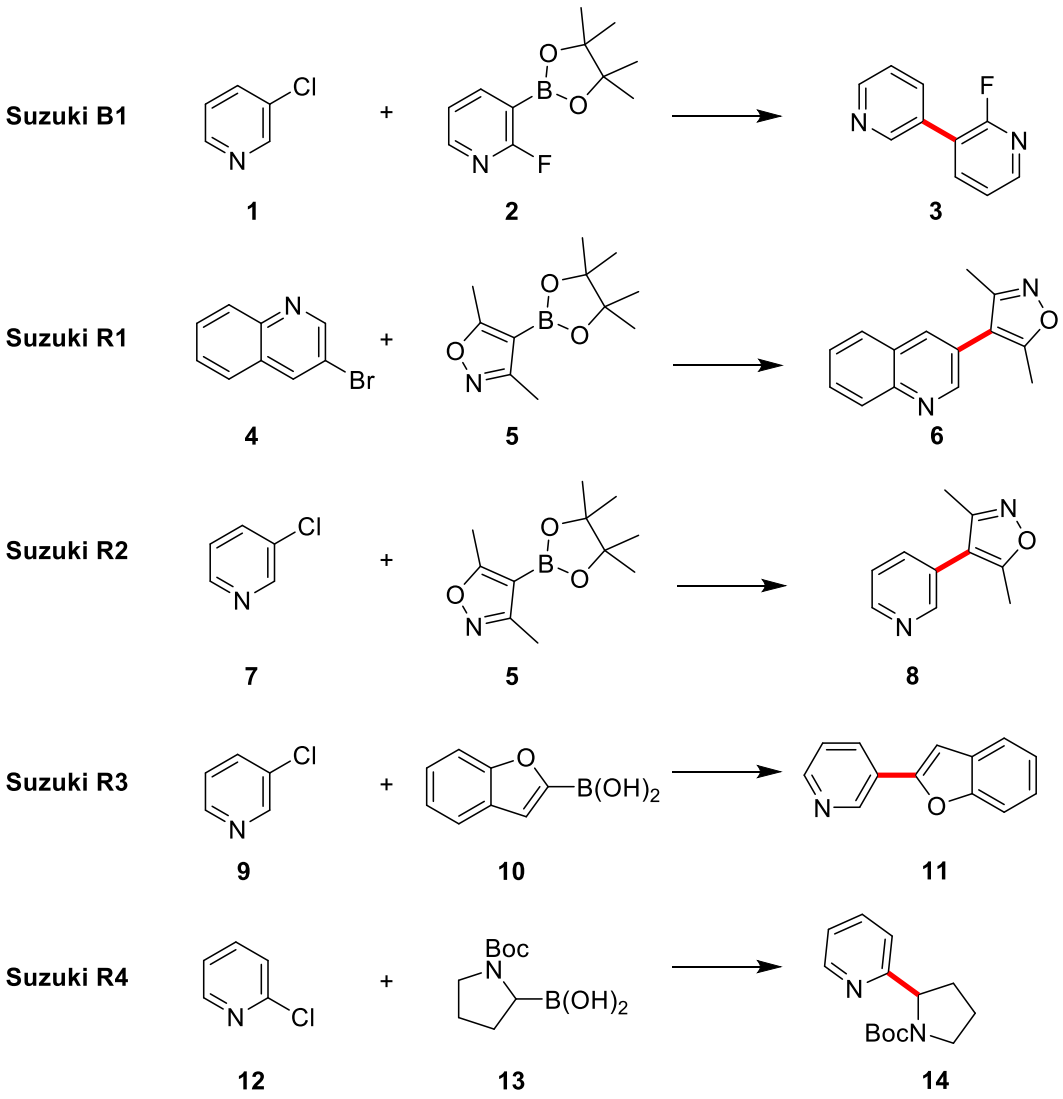

Fixed conditions: DBU (2.0 equiv) | 5:1 THF:H_2O

Continuous Variables

Catalyst Loading (0.5 - 2.5 %) Temperature (30 - 120 °C) Time (1 - 10 min)

Categorical Variables

1: P1-L1 **2:** P2-L1 **3**: P1-L2 **4:** P1-L3 **5**: P1-L4 **6:** P1-L5 **7:** P1-L6 **8:** P1-L7

Pre-catalysts

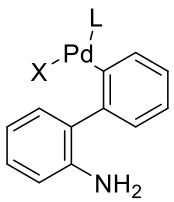

P1 (X = OMs)
P2 (X = Cl)

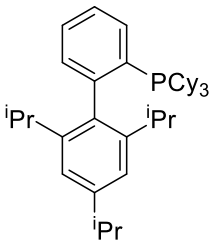

L1 (XPhos)

Ligands

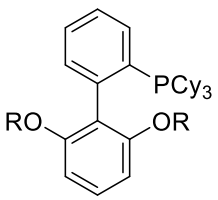

L2 (R = Me, SPhos)
L3 (R = ^{i}Pr, RuPhos)

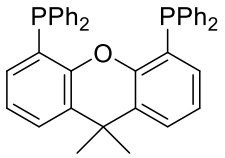

L4 (XantPhos)

L5 PCy_3 **L6** PPh_3
L7 P^tBu_3

Scheme 1: The reactions of interest for the Suzuki-Miyaura coupling *in silico* case studies. The datasets for training the model Suzuki B1 was taken from Baumgartner *et al.*[31] and for Suzuki R1-R4 from Reizman *et al.*[32] In each of these studies, the continuous and categorical variables (with the bounds shown) were optimized for reaction yield.

Four specific case studies are highlighted in Figure 2, each where the main task is Suzuki B1 and the auxiliary training task is one dataset from each of Suzuki R1-4. In each case study, a conventional single-task Bayesian optimization (STBO) benchmark for the Suzuki B1 reaction serves as a comparison. For each MTBO study, 96 datapoints from the auxiliary task were utilized. The average best yield for each algorithm is shown with a 95 % confidence interval over 20 repeated runs.

When leveraging Suzuki R1 as an auxiliary training task, MTBO initially suggests optimal conditions from the training task with P1-L4 (XantPhos). However, these give very low yields (< 25 %) which leads to further exploration of the chemical space, resulting in optimal conditions with P1-L1 (XPhos) and a much higher yield than STBO.

In the second case study, when the auxiliary task is Suzuki R2, MTBO appears to perform poorly - this is likely due to the low reactivity observed in Suzuki R2 and a noisy simulation benchmark (see Figure S8). In this case, the best conditions from the training task also do not perform well on the main task, but the yield is moderate enough that it makes further exploration of the chemical space initially difficult in obtaining a better response. This suggests that MTBO may bias the training data in these circumstances when higher yields are possible but not expected, when given very low-yielding auxiliary tasks.

In the case studies where the auxiliary tasks were Suzuki R3-4, the reactivity of the substrates was much more similar in both the main and the training tasks, leading to similar optimal conditions being found. This means that MTBO achieved better, and much faster, results than STBO in these cases.

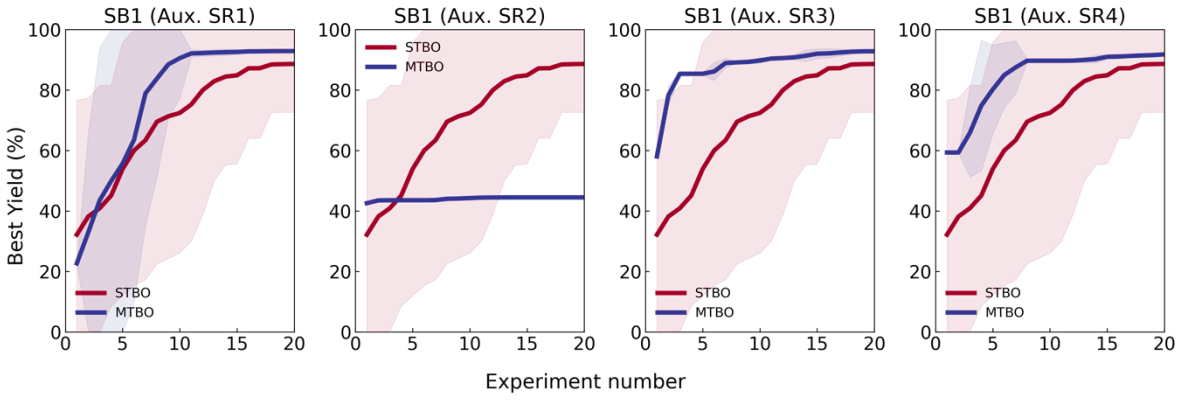

Figure 2: A comparison of the performance of single-task Bayesian optimization (STBO) and multi-task Bayesian optimization (MTBO) of Suzuki B1 with Suzuki R1-R4 as auxiliary tasks. The average best yield with a 95% confidence interval over 20 repeats is shown. The label above each plot refers to the auxiliary (Aux.) task based on the names in Scheme 1, where Suzuki is abbreviated to S.

Performance of MTBO can be greatly improved using *multiple* auxiliary tasks. As shown in Figure 3a, when Suzuki B1 is optimized with Suzuki R1-R4 as auxiliary tasks, the optimal conditions are always found by MTBO in fewer than five experiments. Both P1-L1 and P2-L1 are considered optimal for this reaction,[31] and MTBO selects these two catalysts in over 80 % of experiments during twenty repeats, when compared to < 50 % frequency for STBO - this is highlighted in Figure 3b. As MTBO utilizes optimal regions of chemical space that have been identified in previous tasks with similar reactivity, this allows the algorithm to identify new (and better performing) optimal reaction conditions faster.

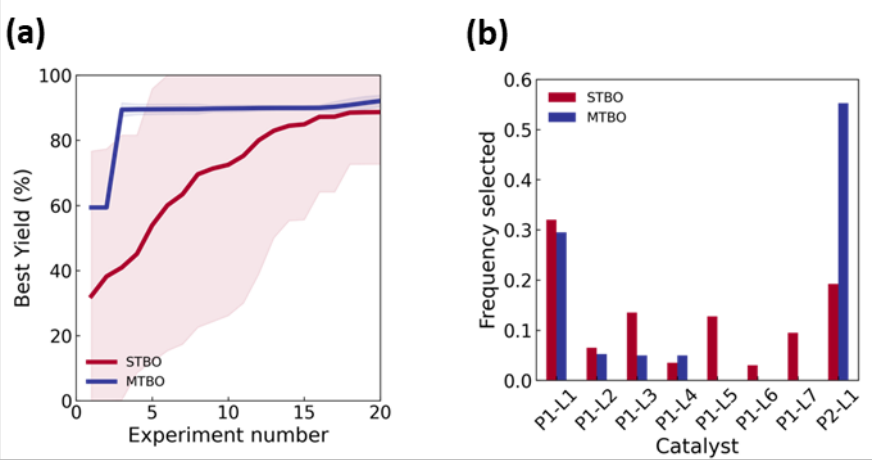

Figure 3: A comparison of the performance of single-task Bayesian Optimization (STBO) and multi-task Bayesian Optimization (MTBO) of Suzuki B1 and all of Suzuki R1-R4 as auxiliary tasks. **(a)** The average best yield with a 95% confidence interval over 20 repeats is shown. **(b)** Frequency of selection of each catalyst in Scheme 1 by STBO and MTBO.

These simulated case studies suggest that the use of MTBO is often beneficial, particularly when not mapping the predicted reactivity differences of the main and auxiliary substrates *a priori*. Initial guesses (optimization starting points) are typically better than random initialization because of previous reaction information, and the rate of 'best yield' improvement is also greater. In the best-case scenario, the reactivity of the new substrate is similar to those of previous datasets and results in a greater yield much faster than standard STBO. In the worst-case scenario, MTBO can fail with one noisy auxiliary case, but we found that using multiple auxiliary tasks helps to mitigate these issues. With these findings, we were confident that MTBO would be effective in real-world case studies where we have experimental datasets from previous optimization campaigns. Further *in silico* case studies for other reaction types, namely Buchwald-Hartwig cross couplings, were also conducted and showed similarly promising results; these studies can be found in the Supporting Information.

Experimental Case Studies: C–H Activation

The reaction class that we targeted for our experimental MTBO study was the palladium-catalyzed C–H activation reaction, reported by Hennessy and Buchwald,[33] yielding pharmaceutically relevant oxindoles (**16**) from their corresponding chloroacetanilides (**15**), as shown in Scheme 2. Each case study is shown in Table 1 and is highlighted if it is forming a potential bioactive fragment or active pharmaceutical ingredient (API) intermediate. The rationale behind these studies is two-fold: firstly, these oxindoles are closely related to many known bioactive molecules and hence medicinal chemistry projects, and secondly, when considering optimal growth vectors for bioactive molecular fragments to grow into more potent drug candidates (such as in FBDD),[18] the most beneficial transformations are often exploiting C–H bonds on the fragment to form new C–C bonds.[20] Therefore, using MTBO, we aimed to rapidly optimize several transformations using different starting materials with unique functionalities to yield structurally diverse oxindole products by forming valuable sp^2-sp^3 C–C bonds. Then, for future optimization campaigns requiring oxindole syntheses, this model can be employed to expediate reaction optimization and process development for new substrates.

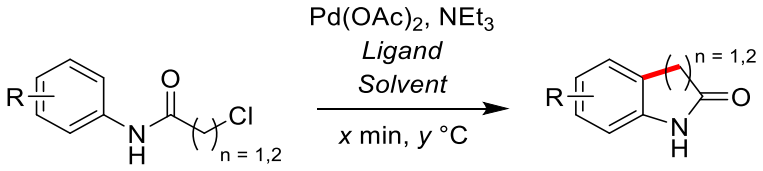

Scheme 2: The reaction class of interest for the MTBO study, where the substituted chloroacetanilide, **15**, reacts to form the corresponding oxindole, **16**. $Pd(OAc)_2$ and NEt_3 remain constant in each experiment, but the ligand, solvent, catalyst concentration, residence time and reaction temperature are optimized for each case study.

Table 1: Each experimental case study explored in this work, including the starting material used, the product formed and the API structure that the product is linked to. These reactions, and references to their known API structure, are highlighted in Schemes 3 - 6.

Case study	Starting material	Product	API structure
1	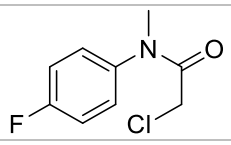	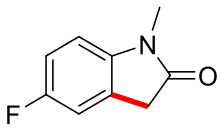	-
2	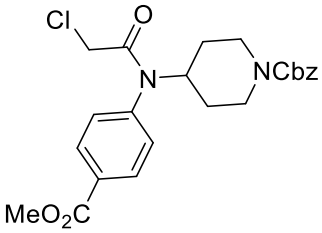	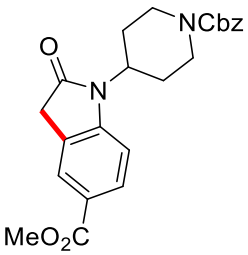	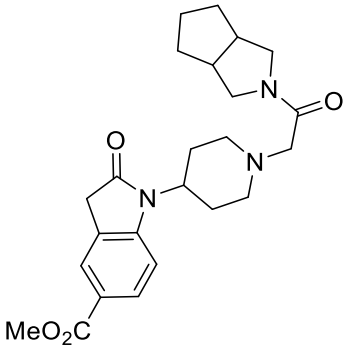
3	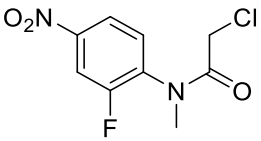	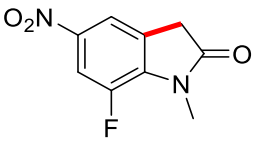	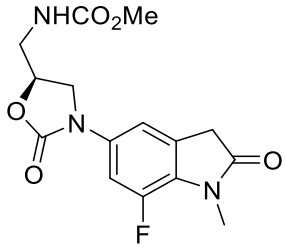
4	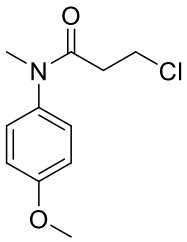	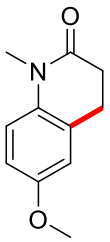	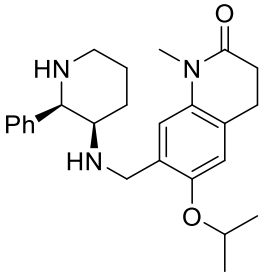

For all experimental work conducted during this study, a self-optimizing flow reactor platform was utilized with a control interface previously disclosed by our group.[34] This platform employs an autonomous optimization workflow, where all experiments are conducted and analyzed without human intervention. All initial training experiments are planned using LHS, then the results from these automated experiments (the yield of the product) are exported using on-line HPLC sampling. Based

on these reaction data, and previously conducted auxiliary tasks, the MTBO algorithm then determines the most beneficial reaction conditions to execute in order to maximize product yield. The actual product yield obtained from this reaction is then communicated back to the algorithm, where the experimental feedback loop is closed, as the algorithm suggests conditions for the next optimization iteration (as shown in in Figure 4). Furthermore, only minimal amounts of reaction material are consumed in each experiment by using reaction slugs;[35] this is an important miniaturization consideration relevant to medicinal chemistry settings, but could potentially be miniaturized further. The minimum slug length is determined on the basis of dispersion in laminar flow such that sampling from a slug is consistent between slugs in repeated tests - the volume of the slugs used in these studies is 4 mL. The aim of this experimental methodology is to accelerate the optimization timeline by requiring fewer experiments and less reaction material consumption. More information on the reactor setup can be found in the Methods section.

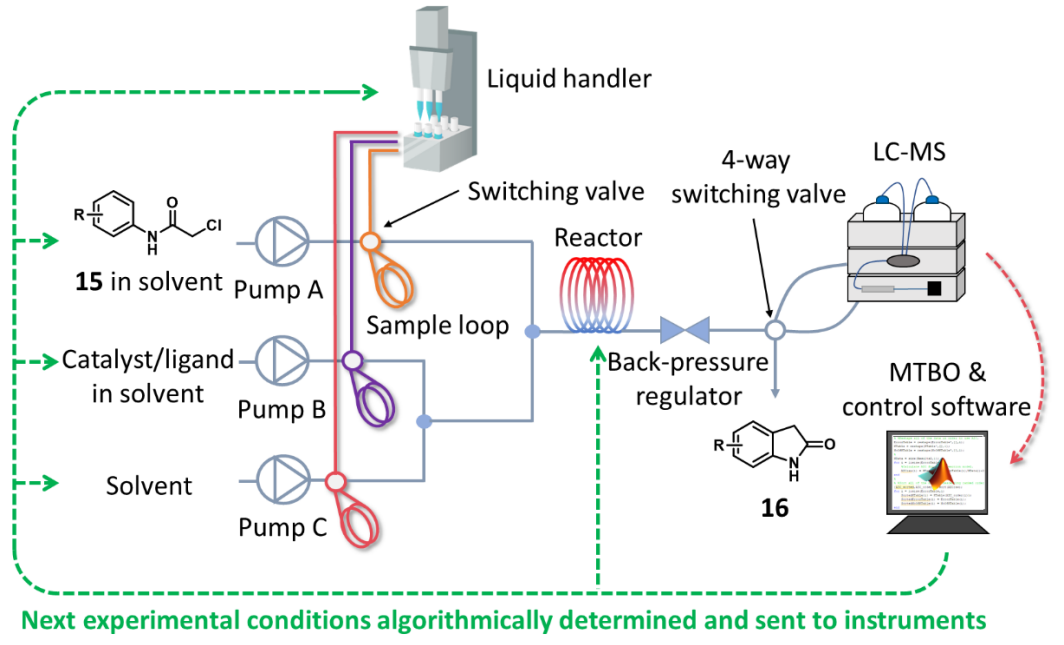

Figure 4: A schematic diagram of the experimental setup and protocol we used for the MTBO self-optimization studies.

For each case study, we optimized the continuous parameters: residence time (5 - 60 minutes), reaction temperature (50 - 150 °C) and catalyst concentration (1 - 10 mol%), and the categorical parameters: solvent (toluene, DMA, acetonitrile, DMSO, NMP) and ligand (JohnPhos, SPhos, XPhos, DPEPhos), for the maximum product yield output. While it is possible to represent these categorical variables in numerous ways, the simplest representation (one-hot encoding) proved sufficient to learn from.[36, 37] The first case study, as shown in Scheme 3, utilizes only single-task Bayesian optimization (STBO) as there is no previous data to leverage model understanding for MTBO. The starting material, **17**, reacts to form the molecular fragment (with potential growth vectors for further functionalization), **18**. The optimization was initialized using 16 (2^4) training experiments before the algorithm began to suggest experimental conditions.

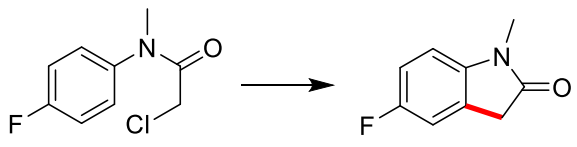

17 **18**

Scheme 3: The first case study explored using STBO, where the substituted chloroacetanilide, **17**, reacts to form the oxindole, **18**. This product is previously unreported via this C–H activation methodology.

After the initial training experiments, the feedback loop (as described in the Methods section) was implemented and 7 further experiments were conducted, finding the optimal reaction parameters of: NMP, XPhos, 53 minutes residence time, 89 °C reactor temperature with 9 mol% catalyst, yielding the product, **18**, in 74.6 % yield. These results are interesting, because with many reported optimization campaigns the optimal conditions are often the most forcing (highest temperature, highest residence time, highest catalyst concentration)[3, 5, 38] However, in this case, the algorithm determines that a moderate reactor temperature is important for a higher yield. This is because the starting material reacts to form other products, leading to a decrease in the desired product yield under more forcing conditions. Furthermore, the optimized conditions reported in the original publication describing these types of reactions feature toluene and JohnPhos,[33] which are different from our optimized parameters for this reaction. However, these reported conditions require reaction times of 2.5 - 6 hours which are difficult to replicate in flow, which could be the reason why the same categorical parameters were not determined to be optimal in our 5 - 60 minutes residence time optimization space. A plot of the experimental data, both training and optimization experiments, and the yields achieved are shown in Figure 5. These 23 experiments required to achieve optimal conditions are also significantly fewer in number than what would be required for current industrial-standard optimization procedures, such as design of experiments (DoE), which would require >750 experiments of efficient design space exploring data points. All reaction data for each case study is reported in full in the ESI, as well as efficiency comparisons with industrial-standard optimization procedures.

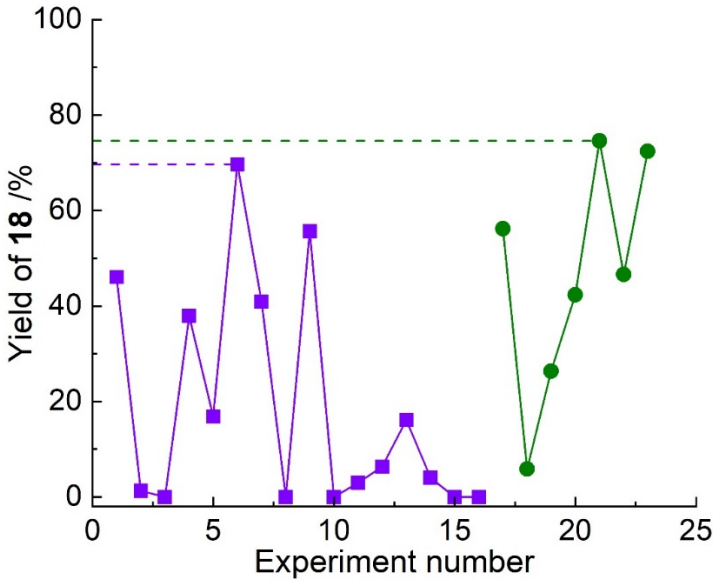

Figure 5: A plot of yield of product, **18**, against experimental number in the STBO campaign. Where: = training experiments and = optimization experiments.

Utilizing the experimental dataset from this optimization campaign, a different substrate was then explored in the second case study. The starting material, **19**, reacts to form a key intermediate for a serine palmitoyl transferase (SPT) inhibitor, **20**, as shown in Scheme 4.[39] As this is a similar transformation, the use of MTBO should hasten optimization and produce optimal reaction conditions much more quickly. The optimization is initiated, and the first suggested experiment deviates only slightly from the previously obtained best parameters, whilst still utilizing NMP and XPhos as the categorical variables but produces a poor yield of the product (14.8 %). As this yield is much lower

than what the underlying multi-task model had predicted, the corresponding weightings to select this area of parameter space for this case study are greatly reduced and thereby the likelihood of exploring this area again during this campaign is reduced. The model then balances the exploration of new parameter space with the exploitation of known favorable conditions, particularly from the previous case study, to iterate through further experiments. The optimal reaction conditions were found in 11 experiments: acetonitrile, JohnPhos, 28 minutes residence time, 127 °C reactor temperature with 5 mol% catalyst, yielding the product, **20**, in 84.9 % yield. It is important to note that this area of parameter space is far from the identified optimum in the previous case study, showing the adaptability of MTBO to similar optimization tasks without simply exploiting near the previously obtained optimal conditions. To identify these process parameters, this entire workflow consumed only 980 mg of the starting material, **19**, and has a much greater throughput (requiring less catalyst loading, cheaper materials and non-complex solvent mixtures) than other reports of this chemistry that yield only 76 % of the desired product.[40] This experimental data is displayed at the end of this section in Figure 5 (red dotted line).

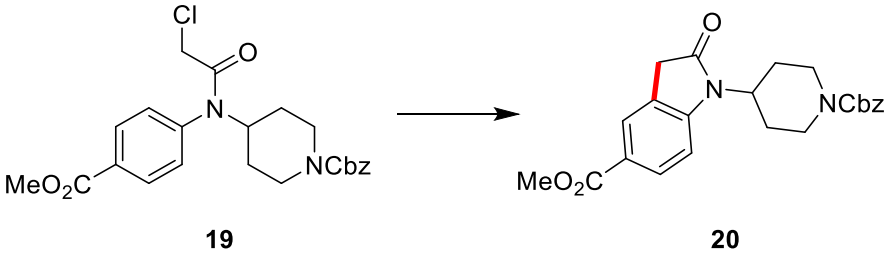

19 **20**

Scheme 4: The second case study explored using MTBO, where the substituted chloroacetanilide, **19**, reacts to form the key intermediate en route to a serine palmitoyl transferase (SPT) inhibitor, **20**. This product is previously unreported via this C–H activation methodology.

With two completed optimization campaigns, these datasets could then be leveraged for the optimization of process parameters for a third case study. This case study features the transformation of **21** into the antibacterial intermediate, **22**, necessary for the synthesis of the oxazolidinone antibiotic Linezolid,[41] as shown in Scheme 4.

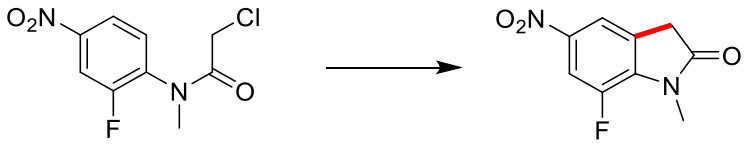

21 **22**

Scheme 5: The third case study explored using MTBO, where the substituted chloroacetanilide, **21**, reacts to form the key intermediate, **22**, for the antibiotic Linezolid. This case study utilized data from the previous two optimization campaigns. This product is previously unreported via this C–H activation methodology.

The initial experiment in MTBO used similar conditions to the optimal conditions from the second case study, with acetonitrile and JohnPhos as categorical variables with 18 minutes residence time with 5 mol% catalyst at 139 °C. This produced a good yield of 71 % but was subsequently improved by using NMP and XPhos, as the MTBO algorithm discovered from the first case study is also a parameter space region of high interest, immediately improving the yield to 83 %. Upon further adjustment of the continuous variables, a yield of 98 % was achieved in only five total experiments. This is the first optimization campaign where one ortho site was blocked for cyclisation, but this variation is seemingly not enough to divert chemical reactivity from what the MTBO algorithm expects, thereby proving the task's applicability to these divergent structures. The entire workflow for

optimizing this process used only 250 mg of the starting material, **21**, which also resulted in a greater yield, throughput and greener process than other reports in the literature (86 % yield in batch, overnight using fluorinated solvents).[42] This experimental data is displayed at the end of this section in Figure 5 (orange solid line).

The next experimental case study features the transformation of the starting material, **23**, into the NK1 receptor antagonist intermediate, **24**, as shown in Scheme 6.[43] In this optimization campaign, the MTBO algorithm leveraged data from the previous three case studies, yet this is the first substrate that forms a 6-membered cyclized ring instead of the typical 5-membered ring in the previous oxindoles. Initial experiments in previously identified well-performing parameter space produced low yields, but the algorithm could thereby determine that further exploration of the parameter space was important as the substrate showed more variability from the previous tasks.

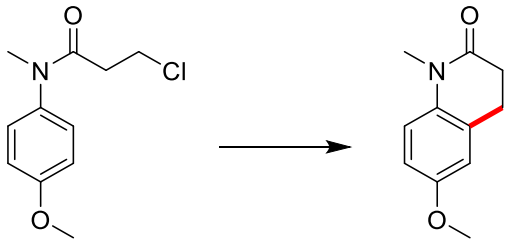

23 **24**

Scheme 6: Another case study explored using MTBO, where the substituted chloroacetanilide, **23**, reacts to form the pharmaceutical intermediate, **24**, for the synthesis of an NK1 receptor antagonist. This case study utilized data from each of the previous optimization campaigns. This product is previously unreported via this C–H activation methodology.

Through further iterations, the categorical variables were exploited: DMSO and DPEPhos, with the most forcing continuous parameters: 60 minutes residence time with 10 mol% catalyst loading at 150 °C. These were determined to be the optimal conditions as found by the self-optimization workflow, giving the product in 82 % yield in 10 experiments using only 450 mg of the starting material, **23**. Despite this functional change, the algorithm was still able to determine the optimal conditions utilizing previous data and quickly found that although a similar reaction task was present, further exploration of the parameter space was necessary. This further shows the adaptability of the MTBO approach to wider substrate scopes with different functionalities. This experimental data is displayed at the end of this section in Figure 5 (green dashed line).

A final case study was then attempted using this workflow, which is the same oxindole-forming C–H activation reaction conducted in every other reaction, but this time featured an electron-rich aromatic ring rather than an electron-deficient ring. The substrate of interest, N-methyl-2-methylchloroacetanilide, also had one ortho position blocked for cyclisation. This study was conducted to further test the limits and the adaptability of the MTBO algorithm, but even with the most forcing conditions possible using our workflow we could only achieve a 29 % product yield. This was also true when using the reported categorical conditions for this substrate in the initial publication[33] - however, the heterogeneity of the reactor systems may have negatively affected the yield outcome, i.e. 6 hour reaction times cannot be achieved easily in flow. Given these observations, we concluded that the reactivity of this species is sufficiently different to previous case studies and therefore cannot be considered as a similar task to the other optimization campaigns. Therefore, for the optimization of these substrates (or any substrates sufficiently different to the tasks of interest) further MTBO campaigns must be conducted for the models to encapsulate these differences to efficiently optimize any case study of interest. With the addition of DFT and reaction similarity scoring,

all substrates of interest can be categorized *a priori* into their respective task bins, avoiding the necessity for additional experimentation. It may also be appropriate in such cases that promising upper bounds leading to full conversion of starting materials are identified, potentially avoiding wasteful experiments in inaccurately defined parameter spaces. Further experimental information on this case study can be found in the ESI.

For each of these C–H activation case studies, generally fewer experiments were necessary to achieve an optimal set of reaction conditions for the highest process yields - this is illustrated in Figure 6. This is because there was an increasing data-density that detailed optimal areas of parameter space for similar tasks (reactions of similar substrates), allowing for a progressively more efficient optimization workflow. In each case study, only minimal amounts (for our specific reaction system) of starting materials were consumed to find optimal reaction conditions, which is very important in early-stage medicinal chemistry development applications when preservation of precious starting materials and speed of optimization are paramount. For these C–H activation case studies, utilizing this workflow with MTBO to find optimal reaction conditions resulted in a material saving of 132 g (£25k) when compared with kinetic studies, and a material saving of 167 g (£32k) when compared with traditional DoE optimization studies (see ESI for details). These costly figures indicate why intermediate reaction optimization in medicinal chemistry applications are often difficult or infeasible to conduct, and why MTBO can serve as an enabling tool for these scenarios. Other common optimization strategies, such as traditional one-factor at a time (OFAT) approaches, may provide modest process improvements in these scenarios but have been shown repeatedly to underperform when compared with statistical-based techniques.[1, 44, 45] This methodology has therefore proven to be effective in real-world pharmaceutical applications for material- and cost-efficiency, with the bonus of full automation that allows scientists to use their human resource to focus on other areas of chemical development. Although these experimental studies focused on C–C bond formation by targeting C–H activation, these techniques can be utilized for other transformations to ultimately accelerate optimization.

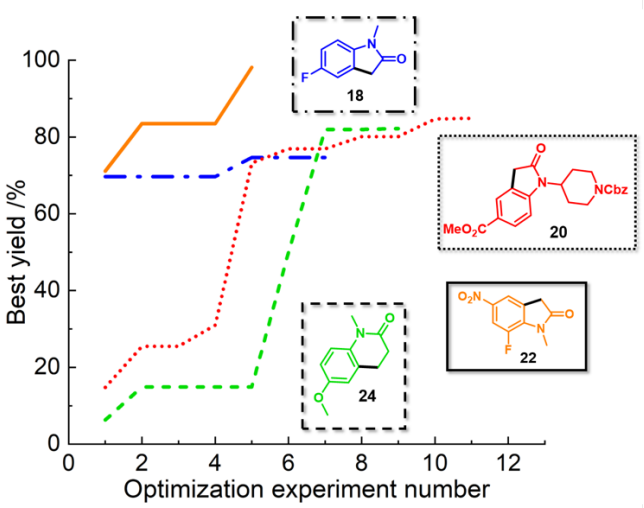

Figure 6: A plot of best yield of the products in each case study against the optimization experimental number in each campaign. The color and dash type of the graph correspond to each product molecule: case study 1 (blue dash-dot), case study 2 (red dot), case study 3 (orange solid) and case study 4 (green dash).

Conclusions

The studies performed in this work, both *in silico* and in real-world chemical applications, represent the first use of datasets from similar reactions to expediate current optimization campaigns with

multi-task Bayesian optimization. This methodology drastically shortens optimization timelines for pharmaceutically relevant transformations, whereas other traditional process optimization techniques (i.e. design of experiments, kinetics studies) would require a significantly higher investment in starting materials, time and cost. This would likely make their optimization infeasible in medicinal fragment-to-lead/FBDD workflows and early-stage process development, unless using intuition-based optimization techniques (such as OFAT) that are unlikely to obtain optimal results.[1] By introducing more miniaturization technology, including smaller reactors/slugs and plate-based screening, there is the added opportunity to reduce material consumption even further using these automated platforms.

With the increasing density of chemical reaction data, both in the literature and in private data storage, there is a wealth of information that can be leveraged for building task-specific models to further increase the efficiency of future reaction optimizations. When using these multi-task learning approaches, it is possible to generate sets of models for specific reaction classes (e.g. Buchwald-Hartwig, Suzuki etc.) and subsets of those models (electron-rich, sterically hindered etc.) to rapidly optimize any transformation likely to be encountered. This is a particularly powerful technique in cases where starting materials are sparse and the reaction is poorly understood, yet suitable quantities of product is required for further molecular design, functionalization and biological testing. Similarly, this importance is echoed in early process development when scale-up of a novel synthetic intermediate is required from the milligram scale to multi-gram or kilogram scale.

The primary challenge when using multi-task Bayesian optimization is its tendency to bias towards the best conditions found in a single auxiliary task, as shown in our *in silico* studies. However, our results demonstrate that additional useful auxiliary tasks can reduce the impact of a noisy, low-yielding auxiliary task. Interestingly, we also noticed that the benefit of MTBO use over STBO increases as the number of categorical options increases. Indeed, our experimental case studies each had 20 categorical options, and we found MTBO to perform very well as we proceeded to add multiple auxiliary tasks. Future work could use a more exploratory acquisition function in combination with the multi-task model to strike the right balance between biasing towards the auxiliary task data and exploring untested conditions.

The multi-task Bayesian optimization algorithm used in this study is open-source and is released as a package within the Summit framework previously reported by our group.[28] This step towards utilizing machine learning and previous reaction data for future optimization campaigns will ultimately result in faster and more efficient optimizations, thereby serving as a broadly applicable enabling tool with relevance to medicinal chemistry and FBDD settings, where industry-standard process optimization techniques are impractical or even impossible to implement.

Methods

Flow Reactor Platform

The reactor platform consists of two Vapourtec R2 modules for controlling flow rates, a Vapourtec R4 reactor module for controlling reactor temperature, a Gilson GX-271 liquid handler for dispensing and collecting reaction material and LC-MS analytical equipment (Shimadzu/Waters) for reaction outcome determination. For each reaction, with the experimental conditions determined through LHS or algorithmically, the liquid handler dispenses 2 mL slugs of the starting material (in this case, the chloroacetanilide **15**) pre-dissolved in the selected solvent into the sample loop for pump A - this solution also contains biphenyl as an internal standard. The selected catalyst/ligand combination in the same solvent is then loaded into the sample loop for pump B, and the solvent of interest is loaded

into pump C for dilution. The reaction is conducted with a constant 0.09 M reactor concentration, yielding the corresponding product (in this case, the oxindole **16**), which is thereby analyzed utilizing a 4-way switching valve[46] for on-line LC-MS. Using this methodology, experiments can be run using only minimal amounts of reaction material for each experiment as we are utilizing reaction slugs. This experimental workflow is illustrated in Figure 4.

Gaussian Processes

For single-task Bayesian optimization, we leverage a Gaussian process (GP) as the probabilistic model in BO due to its excellent performance in the limit of small of data.[30] A GP is a stochastic process characterized by a mean $\mu_\theta(x)$ and covariance function $k_\theta(x, x')$. The covariance function is often called a kernel, which is the term we will use henceforth.

$$f(x) = \mathcal{GP}\big(\mu(x), k_\theta(x, x')\big)$$

where $\boldsymbol{\theta}$ are referred to as hyperparameters of the kernel. Given a finite set of N inputs $X = \{x_1, x_2, \dots, x_N \,|x_i \in \mathbb{R}^m\}$ that correspond with outputs $y = \{y_1, y_2, \dots, y_N | y_i \in \mathbb{R}\}$ the GP is a multivariate Gaussian distribution:

$$f(\boldsymbol{X}) \sim \mathcal{N}(\mu_\theta(\boldsymbol{X}), k_\theta(\boldsymbol{X}, \boldsymbol{X}'))$$

The mean function and kernel act as a prior on the GP. $\mu_\theta(x)$ is usually set to zero because the kernel $k_\theta(x, x')$ fully expresses any arbitrary function. In this work, we use the Matérn 5/2 kernel, with hyperparameters $\boldsymbol{\theta} = \{\sigma, \boldsymbol{L}\}$. $\sigma \in \mathbb{R}$ is the scaling hyperparameter and $\boldsymbol{L} \in \mathbb{R}^m$ is a lengthscale that indicates the significance of each input feature.

$$k_\theta(x, x') = \sigma^2 \left(1 + \sqrt{5} d_\theta(x, x') + \frac{5}{3} d_\theta(x, x')^2\right) \exp(-\sqrt{5}\, d_\theta(x, x'))$$

where d is the Euclidean distance weighted by the length scale:

$$d_\theta(x, x') = \left\| \frac{x - x'}{L} \right\|_2$$

Inference on the GP is done by calculating the posterior of the GP. The posterior of the GP is also a Gaussian distribution:

$$\tilde{f}(\boldsymbol{X}) \sim \mathcal{N}(\tilde{\mu}(\boldsymbol{X}), \tilde{\sigma}(\boldsymbol{X}))$$

$$\tilde{\mu}_\theta(x) = k_\theta(x, \boldsymbol{X}) k_\theta(\boldsymbol{X}, \boldsymbol{X}')^{-1} \boldsymbol{y}$$

$$\tilde{k}_\theta(x, x') = k_\theta(x, x') - k_\theta(\boldsymbol{X}, x) k_\theta(\boldsymbol{X}, \boldsymbol{X})^{-1} k_\theta(x, \boldsymbol{X})$$

where $\tilde{\sigma}(x)$ are the diagonals of the covariance matrix calculated using $\tilde{k}_\theta(x, x')$. To train the GP, the log likelihood is maximized, which is the probability that the model predicts the training outputs given the inputs and hyperparameters. The log likelihood avoids overfitting by trading off accuracy of fit to the training data and complexity of the model. The optimal hyperparameters $\boldsymbol{\theta}^*$ are found by maximizing the log likelihood of the outputs y given the inputs $\boldsymbol{X}$ and the hyperparameters θ[47] (where $\Sigma_\theta = \tilde{k}_\theta(\boldsymbol{X}, \boldsymbol{X}')$):

$$\log p\,(y|\boldsymbol{X}, \theta) = \underbrace{-\frac{1}{2}\big(y - \tilde{\mu}_\theta(\boldsymbol{X})\big)^T \Sigma_\theta^{-1} (y - \tilde{\mu}_\theta(\boldsymbol{X}))}_{\text{Data fit}} \underbrace{-\frac{1}{2}\log|\Sigma_\theta| - \frac{d}{2}\log 2\pi}_{\text{Complexity penalty}}$$

Multi-task Gaussian Processes

Multi-task GPs can be used on multioutput functions $f: \chi \rightarrow R^T$, where each of the T outputs can be seen as solutions to unique regression tasks. The key idea is to use a kernel that can extend to multiple tasks. As detailed in the work by Bonilla *et al.*,[48] we use the intrinsic model of coregionalization, which transforms a latent function to yield the outputs:

$$k_\theta^{ICM}(x, x') = k_\theta^t \otimes \ k_\theta \ (x, x')$$

The task kernel k_θ^t is a $T \ \times T$ matrix of trainable parameters where T is the number of tasks. These parameters represent the inter-task correlation.

Bayesian Optimization

Bayesian optimization aims to solve the optimization problem:

$$max_x y(x)$$

where $y(x)$ is the underlying function that we observe via experiments. We use the expected improvement (EI) acquisition function for *in silico* experiments[49] or q-noisy expected improvement (qNEI) acquisition function for flow chemistry experiments.[50]

In BO with EI as an acquisition function, the aim is to choose the point that is expected to improve the most upon the existing best observed point $y^* \geq y(x_i) \forall i \in (1, \dots, t)$ where t is the number of observations thus far. Therefore, we create an improvement function $I(x)$ describing the improvement of the posterior of the GP over the best observed point. If there is no improvement, $I(x) = 0$.

$$I(x) = \ \max\left(\tilde{f}(\boldsymbol{X}) - y^*, 0\right)$$

After several manipulations, a closed form of EI can be found:

$$\mathrm{EI}(x) = \underbrace{(\tilde{\mu}_\theta(x) - \mathrm{y}^*)\Phi(\mathrm{Z}^*)}_{\text{Exploitation}} + \underbrace{\tilde{\sigma}(\mathrm{x})\phi(\mathrm{Z}^*)}_{\text{Exploration}}$$

where $Z^* = \frac{y^* - \mu_\theta(\mathrm{x})}{\tilde{\sigma}(x)}$.

EI suffers from issues with noisy experiments due to its reliance on the best observed point y^*, which is a biased estimate, especially in the low data regime. qNEI aims to overcome this issue by using the maximum of posterior of the GP over the observed inputs: [50]

$$qNEI(x) = \ \mathbb{E}[(\max \xi - \max \xi_{obs})_+]$$

where $\xi_{obs} \sim \tilde{f}(x)$ and $\xi_{obs} \sim \tilde{f}(\boldsymbol{X})$ are samples from the posterior of the GP. We use BOtorch for implementations of GPs and Bayesian optimization.[50] For the experimental C–H activation case studies shown in Scheme 4-6, the qNEI acquisition function was used, while EI was used in the simulation case studies due to computational limitations.

Benchmarks

Prior to real experimentation, we wanted to understand the performance of MTBO in simulated studies. We examined two literature reports that contain experimental results from Suzuki-Miyaura

coupling reactions,[31, 32] and one report with results from a Buchwald-Hartwig cross-coupling[51] (demonstrated in the Supporting Information), building a predictive model for the reaction yield to behave as the ground-truth for simulated optimization studies. Buchwald-Hartwig and Suzuki-Miyaura couplings are ubiquitous in the pharmaceutical and fine chemicals industries as they allow rapid construction of aromatic scaffolds through reactions with few impurities.[52] We therefore chose these reaction classes because of their high value and applicability to real-world scenarios. More details on benchmark training can be found in the Supporting Information.

Acknowledgements

CJT is a Sustaining Innovation Postdoctoral Research Associate at Astex Pharmaceuticals and would like to thank Astex Pharmaceuticals for funding, Suzi Cowan for NMR guidance, Stuart Whibley and Shirley Chen for analytical guidance, and David Twigg, Mark Wade and David Rees for their support. KCF has received PhD funding from the Marshall Scholarship, Cambridge Trust and BASF SE. DW received PhD funding from UCB Pharma. This project was co-funded by European Regional Development Fund via the project "Innovation Centre in Digital Molecular Technologies", UKRI via project EP/S024220/1 'EPSRC Centre for Doctoral Training in Automated Chemical Synthesis Enabled by Digital Molecular Technologies', and Pharma Innovation Partnership in Singapore (PIPS) via project 'C4 Development of multi-step processes in Pharma (PI)'.

Conflicts of Interest

There are no conflicts to declare.

Code Availability

The code for this project can be found at https://github.com/sustainable-processes/multitask.

References

1. Reizman, B.J. and K.F. Jensen, *Feedback in flow for accelerated reaction development.* Accounts of chemical research, 2016. **49**(9): p. 1786-1796.
2. Fabry, D., E. Sugiono, and M. Rueping, *Online monitoring and analysis for autonomous continuous flow self-optimizing reactor systems.* Reaction Chemistry & Engineering, 2016. **1**(2): p. 129-133.
3. Fitzpatrick, D.E., C. Battilocchio, and S.V. Ley, *A novel internet-based reaction monitoring, control and autonomous self-optimization platform for chemical synthesis.* Organic Process Research & Development, 2016. **20**(2): p. 386-394.
4. Hall, B.L., et al., *Autonomous optimisation of a nanoparticle catalysed reduction reaction in continuous flow.* Chemical Communications, 2021. **57**(40): p. 4926-4929.
5. Cortés-Borda, D., et al., *An autonomous self-optimizing flow reactor for the synthesis of natural product carpanone.* The Journal of organic chemistry, 2018. **83**(23): p. 14286-14299.
6. Henson, A.B., P.S. Gromski, and L. Cronin, *Designing algorithms to aid discovery by chemical robots.* ACS central science, 2018. **4**(7): p. 793-804.
7. Houben, C. and A.A. Lapkin, *Automatic discovery and optimization of chemical processes.* Current opinion in chemical engineering, 2015. **9**: p. 1-7.
8. Taylor, C.J., et al., *Rapid, automated determination of reaction models and kinetic parameters.* Chemical Engineering Journal, 2021. **413**: p. 127017.
9. Taylor, C.J., et al., *Flow chemistry for process optimisation using design of experiments.* Journal of Flow Chemistry, 2021. **11**(1): p. 75-86.
10. Clayton, A.D., et al., *Automated self-optimisation of multi-step reaction and separation processes using machine learning.* Chemical Engineering Journal, 2020. **384**: p. 123340.

11. Clayton, A.D., et al., *Algorithms for the self-optimisation of chemical reactions.* Reaction Chemistry & Engineering, 2019. **4**(9): p. 1545-1554.
12. Amar, Y., et al., *Machine learning and molecular descriptors enable rational solvent selection in asymmetric catalysis.* Chemical science, 2019. **10**(27): p. 6697-6706.
13. Schweidtmann, A.M., et al., *Machine learning meets continuous flow chemistry: Automated optimization towards the Pareto front of multiple objectives.* Chemical Engineering Journal, 2018. **352**: p. 277-282.
14. Shields, B.J., et al., *Bayesian reaction optimization as a tool for chemical synthesis.* Nature, 2021. **590**(7844): p. 89-96.
15. Angello, N.H., et al., *Closed-loop optimization of general reaction conditions for heteroaryl Suzuki-Miyaura coupling.* Science, 2022. **378**(6618): p. 399-405.
16. Grainger, R., et al., *Enabling synthesis in fragment-based drug discovery by reactivity mapping: photoredox-mediated cross-dehydrogenative heteroarylation of cyclic amines.* Chemical science, 2019. **10**(8): p. 2264-2271.
17. Buitrago Santanilla, A., et al., *Nanomole-scale high-throughput chemistry for the synthesis of complex molecules.* Science, 2015. **347**(6217): p. 49-53.
18. Murray, C.W. and D.C. Rees, *The rise of fragment-based drug discovery.* Nature chemistry, 2009. **1**(3): p. 187-192.
19. Congreve, M., et al., *Recent developments in fragment-based drug discovery.* Journal of medicinal chemistry, 2008. **51**(13): p. 3661-3680.
20. Chessari, G., et al., *C–H functionalisation tolerant to polar groups could transform fragment-based drug discovery (FBDD).* Chemical science, 2021. **12**(36): p. 11976-11985.
21. Felton, K., D. Wigh, and A. Lapkin, *Multi-task Bayesian optimization of chemical reactions.* 2021.
22. Swersky, K., J. Snoek, and R.P. Adams, *Multi-task bayesian optimization.* Advances in neural information processing systems, 2013. **26**.
23. Sans, V. and L. Cronin, *Towards dial-a-molecule by integrating continuous flow, analytics and self-optimisation.* Chemical Society Reviews, 2016. **45**(8): p. 2032-2043.
24. Mateos, C., M.J. Nieves-Remacha, and J.A. Rincón, *Automated platforms for reaction self-optimization in flow.* Reaction Chemistry & Engineering, 2019. **4**(9): p. 1536-1544.
25. Kershaw, O.J., et al., *Machine learning directed multi-objective optimization of mixed variable chemical systems.* Chemical Engineering Journal, 2023. **451**: p. 138443.
26. Perera, D., et al., *A platform for automated nanomole-scale reaction screening and micromole-scale synthesis in flow.* Science, 2018. **359**(6374): p. 429-434.
27. Kreutz, J.E., et al., *Evolution of catalysts directed by genetic algorithms in a plug-based microfluidic device tested with oxidation of methane by oxygen.* Journal of the American Chemical Society, 2010. **132**(9): p. 3128-3132.
28. Felton, K.C., J.G. Rittig, and A.A. Lapkin, *Summit: benchmarking machine learning methods for reaction optimisation.* Chemistry-Methods, 2021. **1**(2): p. 116-122.
29. Shahriari, B., et al., *Taking the human out of the loop: A review of Bayesian optimization.* Proceedings of the IEEE, 2015. **104**(1): p. 148-175.
30. Snoek, J., H. Larochelle, and R.P. Adams, *Practical bayesian optimization of machine learning algorithms.* Advances in neural information processing systems, 2012. **25**.
31. Baumgartner, L.M., et al., *Optimum catalyst selection over continuous and discrete process variables with a single droplet microfluidic reaction platform.* Reaction Chemistry & Engineering, 2018. **3**(3): p. 301-311.
32. Reizman, B.J., et al., *Suzuki–Miyaura cross-coupling optimization enabled by automated feedback.* Reaction chemistry & engineering, 2016. **1**(6): p. 658-666.
33. Hennessy, E.J. and S.L. Buchwald, *Synthesis of substituted oxindoles from α-chloroacetanilides via Palladium-catalyzed C− H functionalization.* Journal of the American Chemical Society, 2003. **125**(40): p. 12084-12085.

34. Jeraal, M.I., S. Sung, and A.A. Lapkin, *A Machine Learning-Enabled Autonomous Flow Chemistry Platform for Process Optimization of Multiple Reaction Metrics.* Chemistry-Methods, 2021. **1**(1): p. 71-77.
35. Guidi, M., P.H. Seeberger, and K. Gilmore, *How to approach flow chemistry.* Chemical Society Reviews, 2020. **49**(24): p. 8910-8932.
36. Wigh, D.S., J.M. Goodman, and A.A. Lapkin, *A review of molecular representation in the age of machine learning.* Wiley Interdisciplinary Reviews: Computational Molecular Science, 2022: p. e1603.
37. Pomberger, A., et al., *The effect of chemical representation on active machine learning towards closed-loop optimization.* Reaction Chemistry & Engineering, 2022.
38. Cortés-Borda, D., et al., *Optimizing the Heck–Matsuda reaction in flow with a constraint-adapted direct search algorithm.* Organic Process Research & Development, 2016. **20**(11): p. 1979-1987.
39. Hanada, K., *Serine palmitoyltransferase, a key enzyme of sphingolipid metabolism.* Biochimica et Biophysica Acta (BBA)-Molecular and Cell Biology of Lipids, 2003. **1632**(1-3): p. 16-30.
40. Kiser, E.J., et al., *Kilogram-lab-scale oxindole synthesis via palladium-catalyzed C–H functionalization.* Organic Process Research & Development, 2012. **16**(2): p. 255-259.
41. Brickner, S.J., et al., *Synthesis and antibacterial activity of U-100592 and U-100766, two oxazolidinone antibacterial agents for the potential treatment of multidrug-resistant gram-positive bacterial infections.* Journal of medicinal chemistry, 1996. **39**(3): p. 673-679.
42. Choy, A., et al., *Development of a synthesis for a long-term oxazolidinone antibacterial.* Organic Process Research & Development, 2008. **12**(5): p. 884-887.
43. Wakabayashi, H. and M. Ikunaka, *Substituted benzolactam compounds as substance P antagonists*. 2001, Google Patents.
44. Lendrem, D.W., et al., *Lost in space: design of experiments and scientific exploration in a Hogarth universe.* Drug Discovery Today, 2015. **20**(11): p. 1365-1371.
45. Peris-Díaz, M.D., M.A. Sentandreu, and E. Sentandreu, *Multiobjective optimization of liquid chromatography–triple-quadrupole mass spectrometry analysis of underivatized human urinary amino acids through chemometrics.* Analytical and Bioanalytical Chemistry, 2018. **410**(18): p. 4275-4284.
46. Taylor, C.J., et al., *Modern advancements in continuous-flow aided kinetic analysis.* Reaction Chemistry & Engineering, 2022. **7**(5): p. 1037-1046.
47. Williams, C.K. and C.E. Rasmussen, *Gaussian processes for machine learning*. Vol. 2. 2006: MIT press Cambridge, MA.
48. Bonilla, E.V., K. Chai, and C. Williams, *Multi-task Gaussian process prediction.* Advances in neural information processing systems, 2007. **20**.
49. Jones, D.R., M. Schonlau, and W.J. Welch, *Efficient global optimization of expensive black-box functions.* Journal of Global optimization, 1998. **13**(4): p. 455-492.
50. Balandat, M., et al., *BoTorch: a framework for efficient Monte-Carlo Bayesian optimization.* Advances in neural information processing systems, 2020. **33**: p. 21524-21538.
51. Baumgartner, L.M., et al., *Use of a droplet platform to optimize Pd-catalyzed C–N coupling reactions promoted by organic bases.* Organic Process Research & Development, 2019. **23**(8): p. 1594-1601.
52. Brown, D.G. and J. Bostrom, *Analysis of past and present synthetic methodologies on medicinal chemistry: where have all the new reactions gone? Miniperspective.* Journal of medicinal chemistry, 2016. **59**(10): p. 4443-4458.

Supporting Information for Accelerated Chemical Reaction Optimization using Multi-Task Bayesian Optimization

Connor J. Taylor[1,2‡]*, Kobi C. Felton[‡3], Daniel Wigh[2,3], Mohammed I. Jeraal[4], Rachel Grainger[1], Gianni Chessari[1], Christopher N. Johnson[1], Alexei A. Lapkin[2,3,4]*

[1]Astex Pharmaceuticals, 436 Cambridge Science Park, Milton Road, Cambridge, CB4 0QA, UK. [2]Innovation Centre in Digital Molecular Technologies, Yusuf Hamied Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge, CB2 1EW, UK. [3]Department of Chemical Engineering and Biotechnology, University of Cambridge, Philippa Fawcett Dr, Cambridge CB3 0AS, UK. [4]Cambridge Centre for Advanced Research and Education in Singapore Ltd., 1 Create Way, CREATE Tower #05-05, 138602, Singapore.

[‡]Authors contributed equally.

*Corresponding authors: Connor.Taylor@astx.com, aal35@cam.ac.uk.

Contents

1. Benchmarking using Literature Data

We used literature data to benchmark our Bayesian optimization strategies *in silico*. The challenge is that chemical data is often recorded in unstructured text documents such as journal publications and patents. While there are some conventions, each author is free to express the details of a chemical reaction as they wish. Such unstructured information is not amenable to most machine learning algorithms. Therefore, we designed a custom data extraction workflow that we think is a model for how to apply transfer learning in chemistry.

1.1 Data extraction workflow

As shown in Figure S1, we developed a benchmarking workflow that converts unstructured data into benchmarks that can used for comparing various strategies. We leveraged the Open Reaction Database (ORD) format as a common representation of reactions.[1] We wrote converters from spreadsheet formats to ORD. Once the data was transformed into ORD, the data was loaded into local storage on disk for featurization. Subsequently, the featurization step turned the ORD schema into a set of features that can be used for training a benchmark or a GP for Bayesian Optimization. We used one-hot encodings to represent the categorical variables.

We utilize data from publications on Suzuki couplings (see main text Scheme 1) and C-N couplings (see Scheme S1).[2, 3]

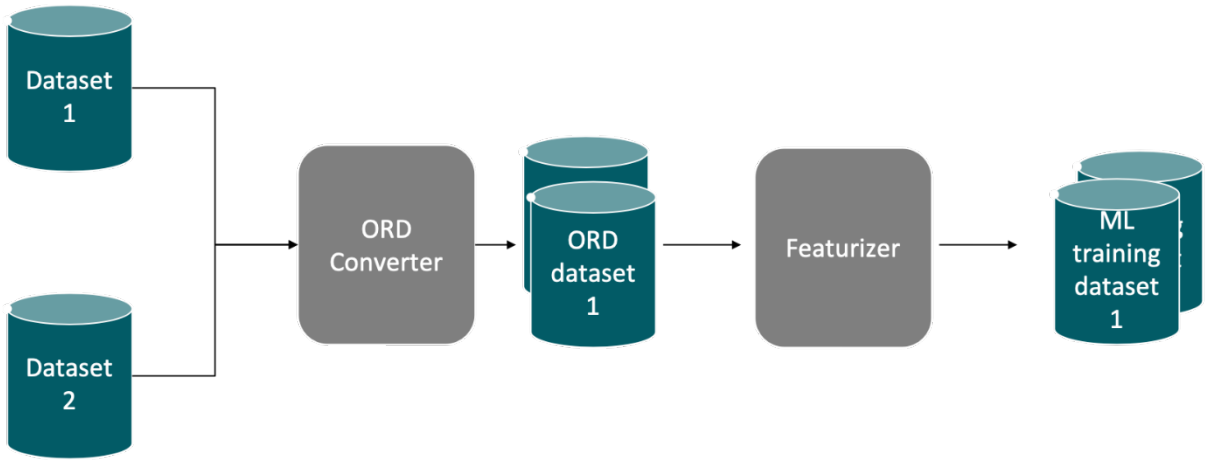

Figure S1: Workflow for converting data in spreadsheets into ORD format and subsequently training datasets for machine leaning.

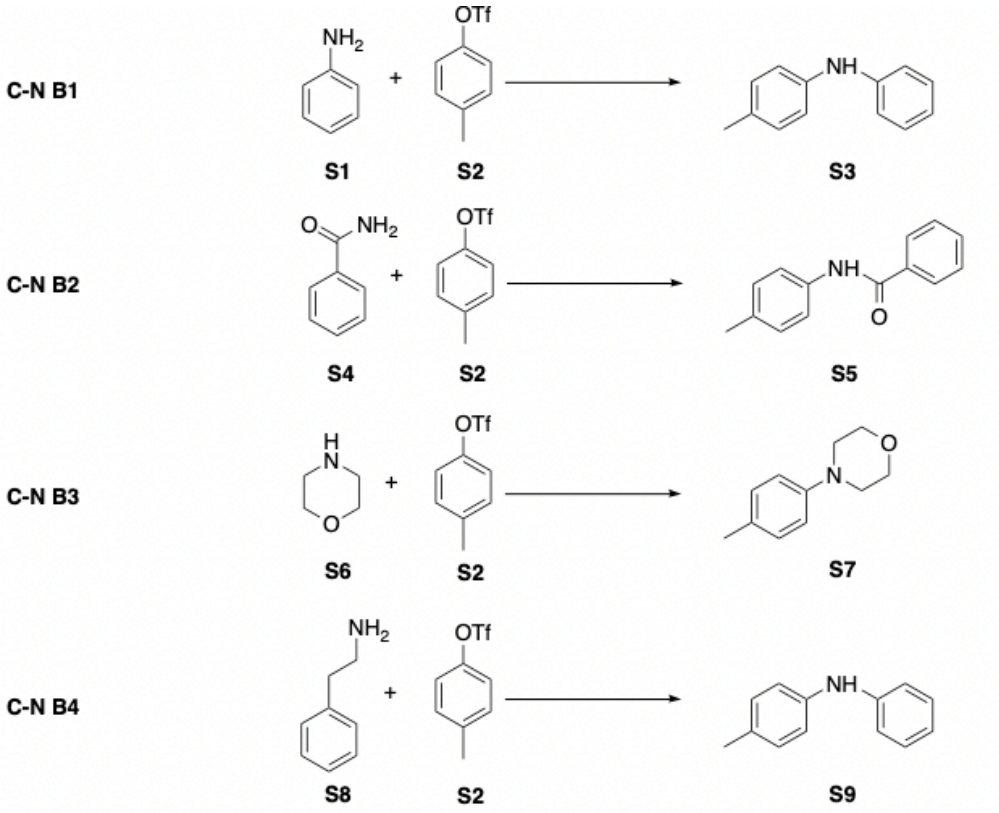

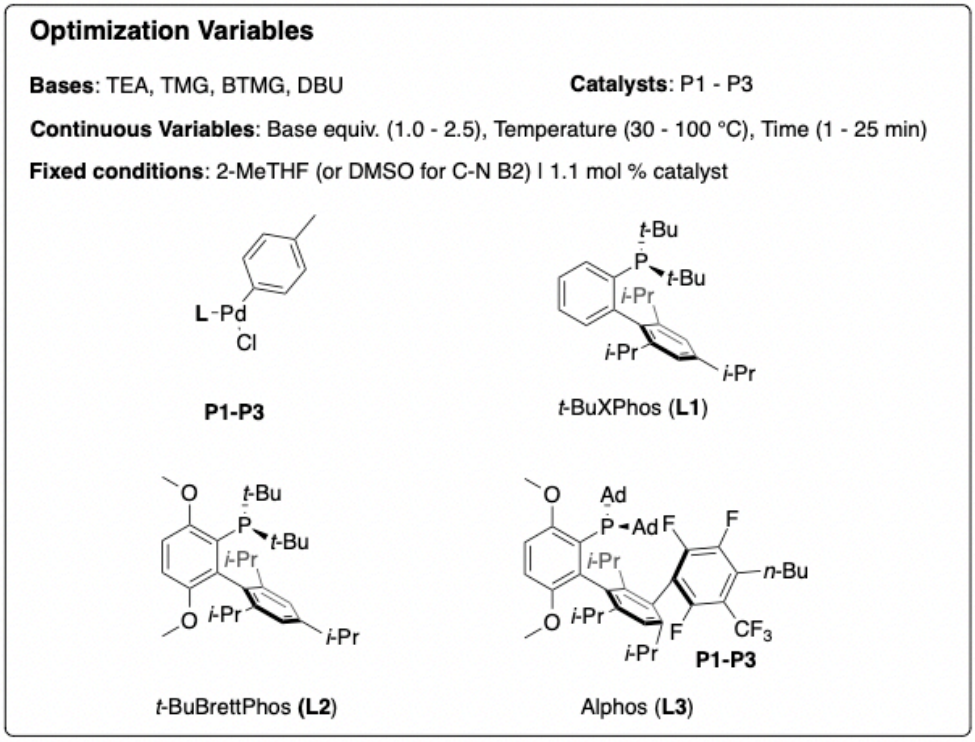

Scheme S1: Benchmark examples of C-N cross coupling (C-N B1-B4) of nitrogen functionalized aromatics (**S1, S4, S6, S8**) with *p*-tolyl triflate. Data is based on data Baumgartner *et al.*[2] The base equivalents, temperature and reaction time, base and catalyst are varied to maximize yield.

Figure S2: Workflow for training an *ExperimentalEmulator* to act as a benchmark.

1.2 Benchmark Training

We leveraged Summit[4] to build the predictive models from the literature reports. As shown in Figure S2, we utilized the *ExperimentalEmulator* feature in Summit, which creates a benchmark based on experimental data. The regressor used was a neural network with one hidden layer of 512 units with a ReLu activation function. A one-hot encoding was used for the pre-catalyst and ligand combinations. The neural networks were trained by five-fold cross validation over 1000 epochs using stochastic gradient descent. Figures S3 - S10 show the parity plots for the benchmarks.

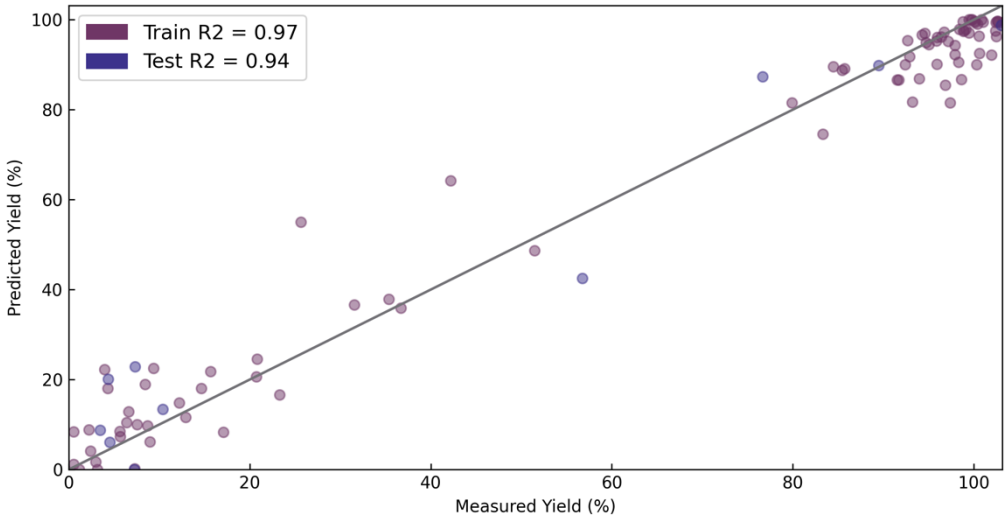

Figure S3: Parity plot of benchmark training for prediction of reaction yield for C-N B1.

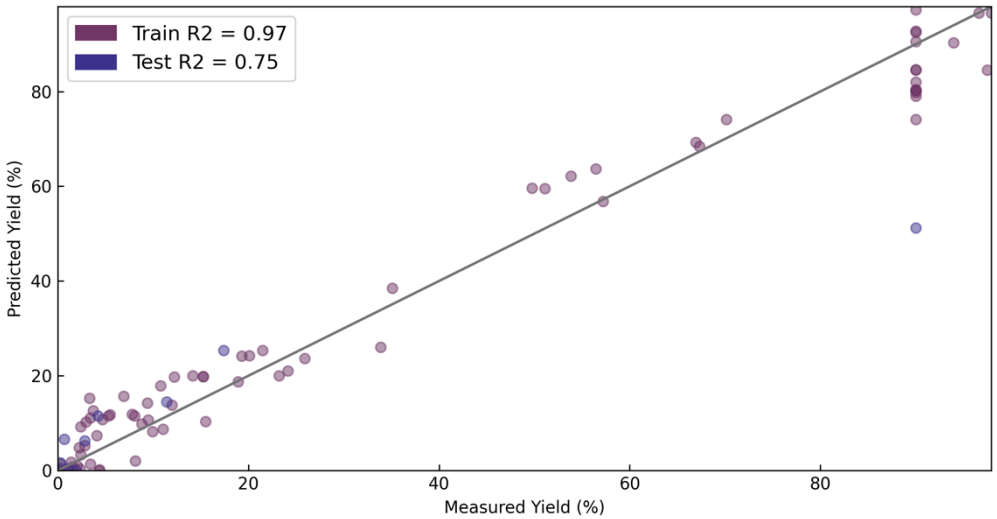

Figure S4: Parity plot of benchmark training for prediction of reaction yield for C-N B2.

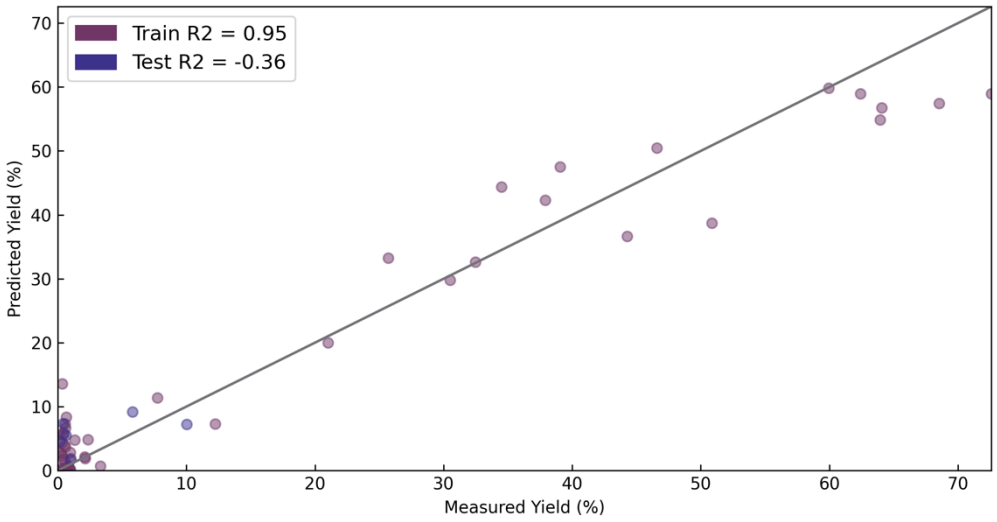

Figure S4: Parity plot of benchmark training for prediction of reaction yield for C-N B3.

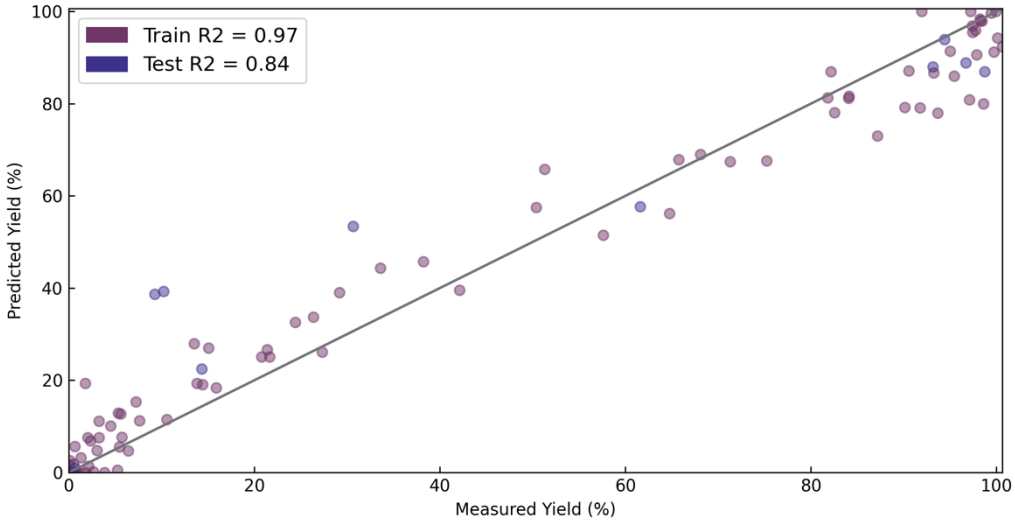

Figure S5: Parity plot of benchmark training for prediction of reaction yield for C-N B4.

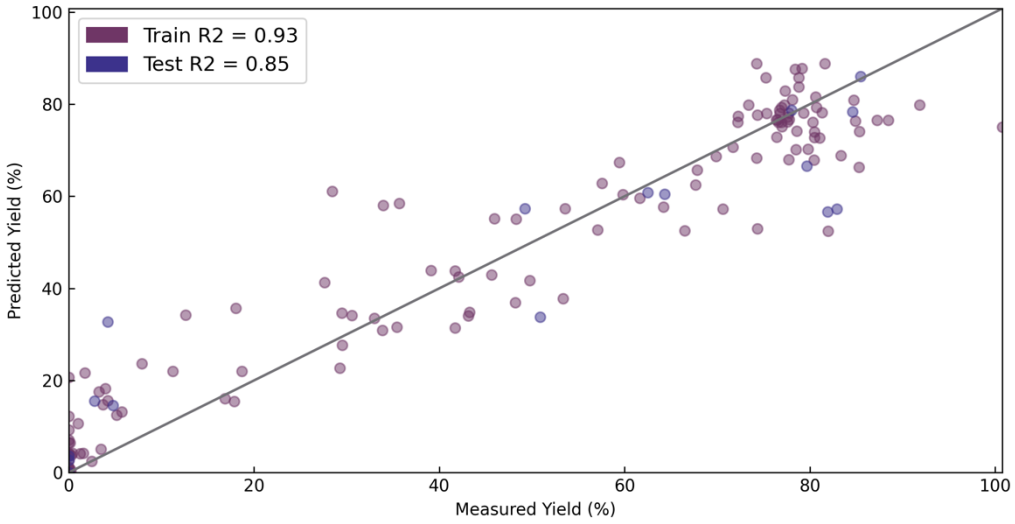

Figure S6: Parity plot of benchmark training for prediction of reaction yield for Suzuki B1.

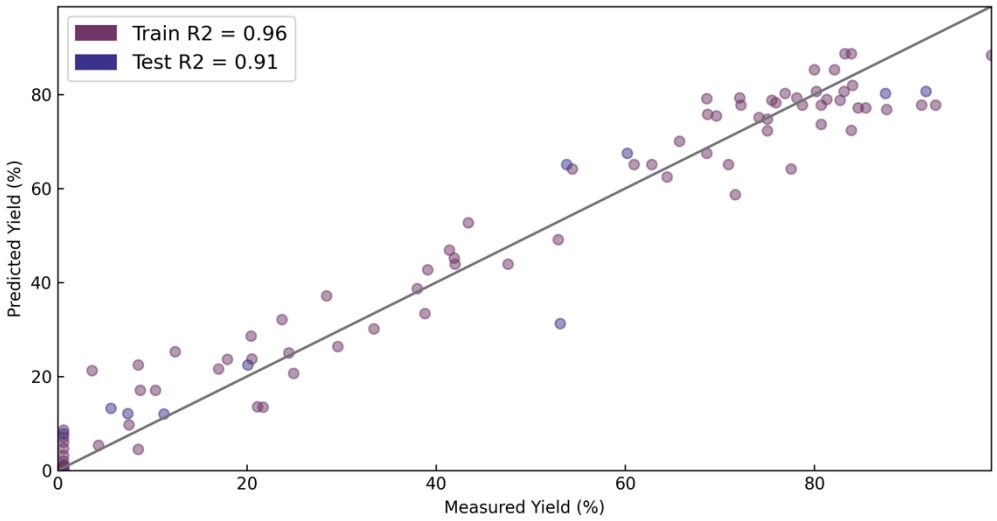

Figure S7: Parity plot of benchmark training for prediction of reaction yield for Suzuki R1.

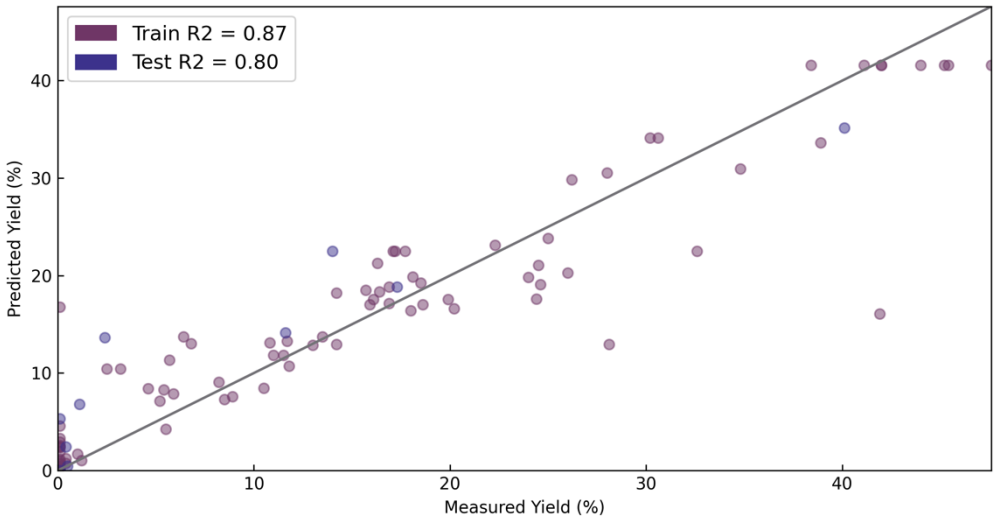

Figure S8: Parity plot of benchmark training for prediction of reaction yield for Suzuki R2.

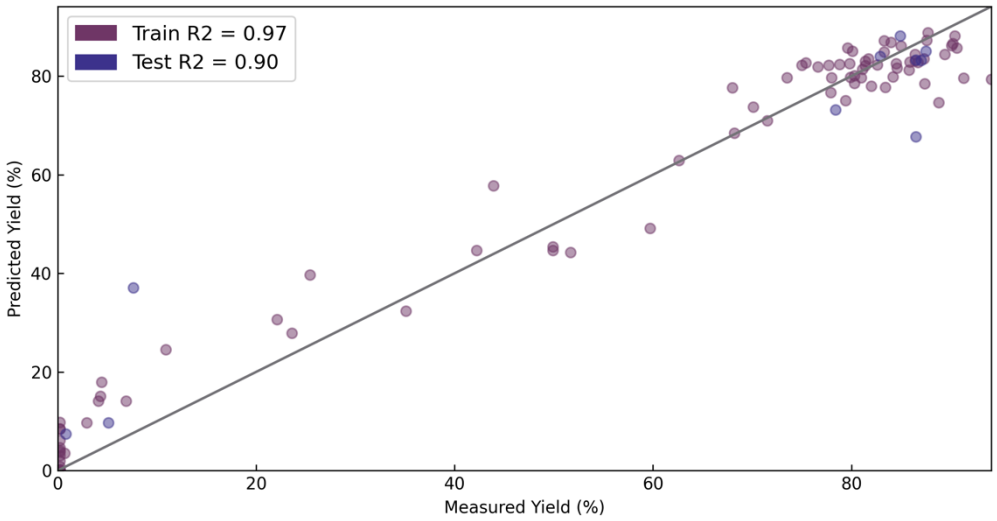

Figure S9: Parity plot of benchmark training for prediction of reaction yield for Suzuki R3.

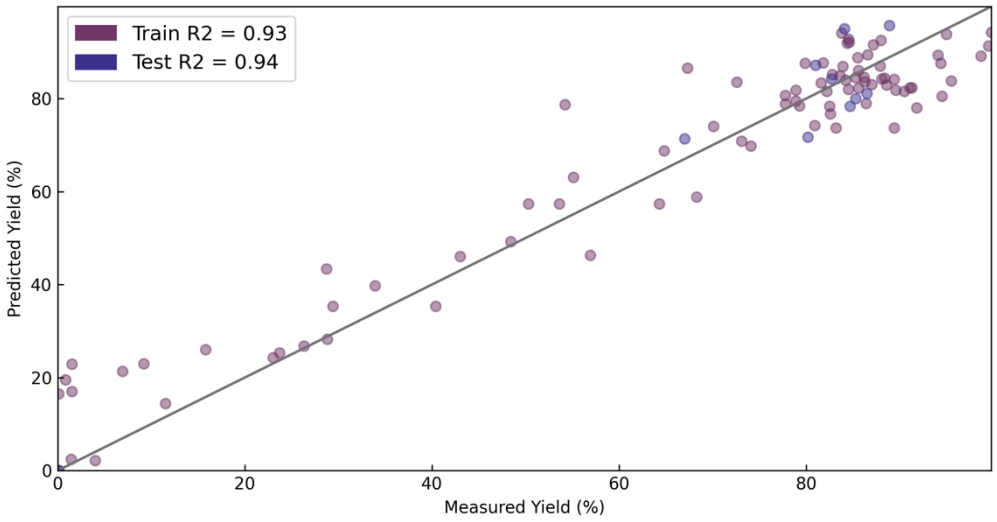

Figure S10: Parity plot of benchmark training for prediction of reaction yield for Suzuki R4.

1.3 Benchmarking Simulation Details

All benchmark simulations were executed on an Amazon Web Services instance with an Nvidia T4 GPU via Lightning AI. For each configuration of optimization and task and auxiliary task, 20 repeats were completed. Figures show the average performance and the 95% confidence interval at each interval. For multitask benchmarks, the first experiment executed was the highest yielding condition from the auxiliary task(s).

1.4 Suzuki Benchmarks

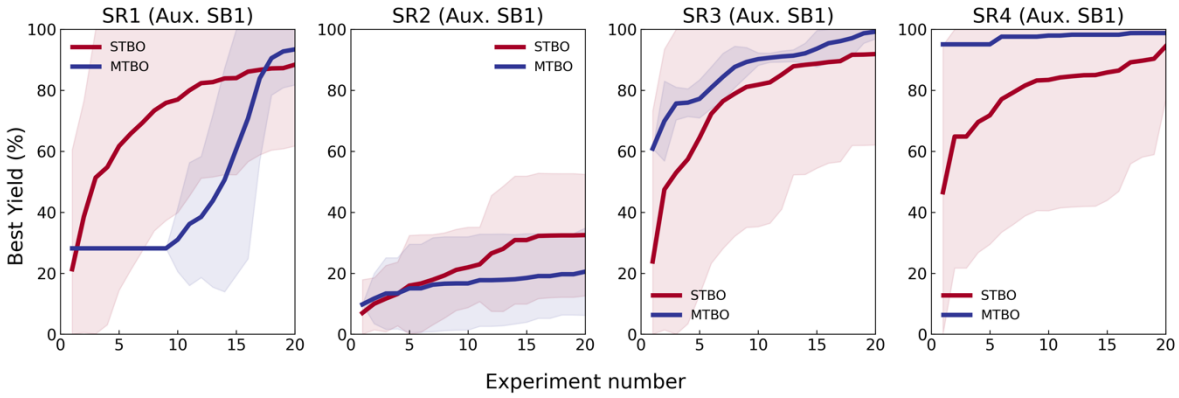

Figure S11: Comparison of the performance of single-task Bayesian optimization (STBO) and multi-task Bayesian optimization (MTBO) on Suzuki reactions R1-R4 with auxiliary data from Suzuki B1.

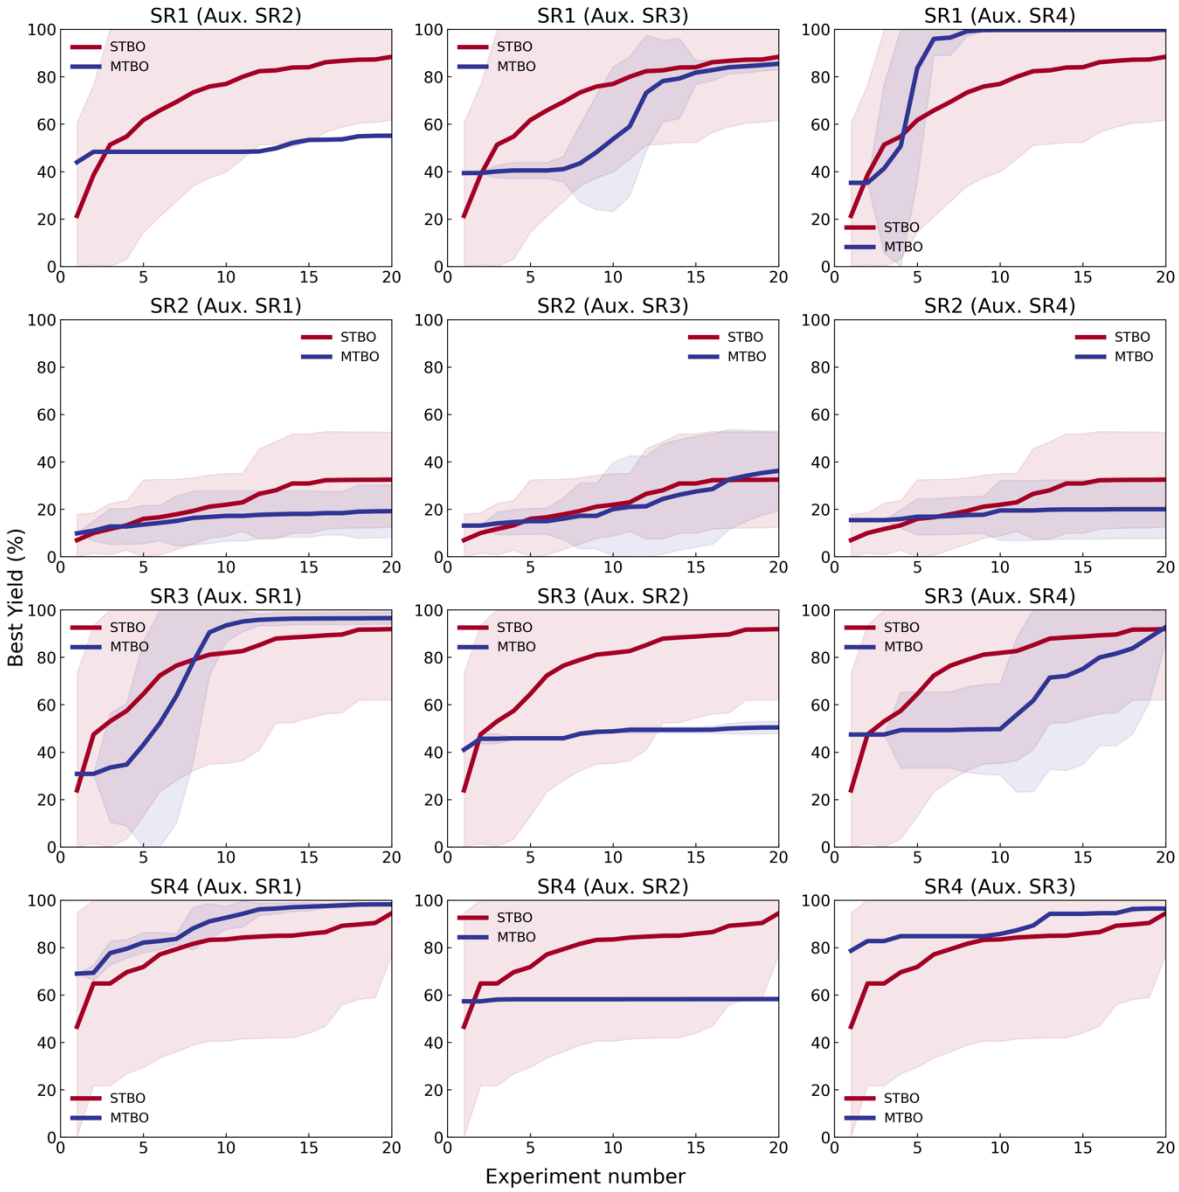

Figure S12: Comparison of the performance of single-task Bayesian optimization (STBO) and multi-task Bayesian optimization (MTBO) on Suzuki R1-R4 with Suzuki R1-R4 as auxiliary tasks. The text above the plot represents the data used as an auxiliary task.

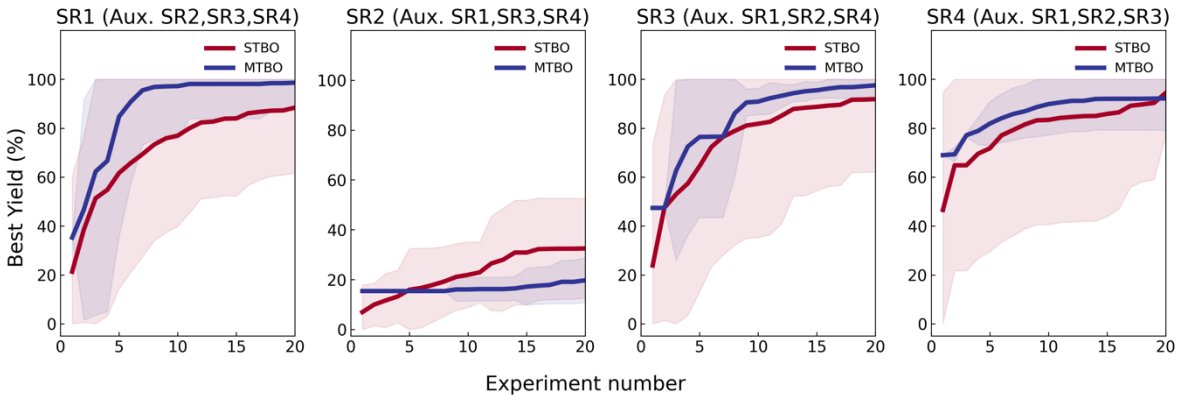

Figure S13: Comparison of the performance of single-task Bayesian optimization (STBO) and multi-task Bayesian optimization (MTBO) on Suzuki R1-R4 with Suzuki R1-R4 as auxiliary tasks. The text above the plot represents the data used as an auxiliary task.

1.5 C-N Benchmarks

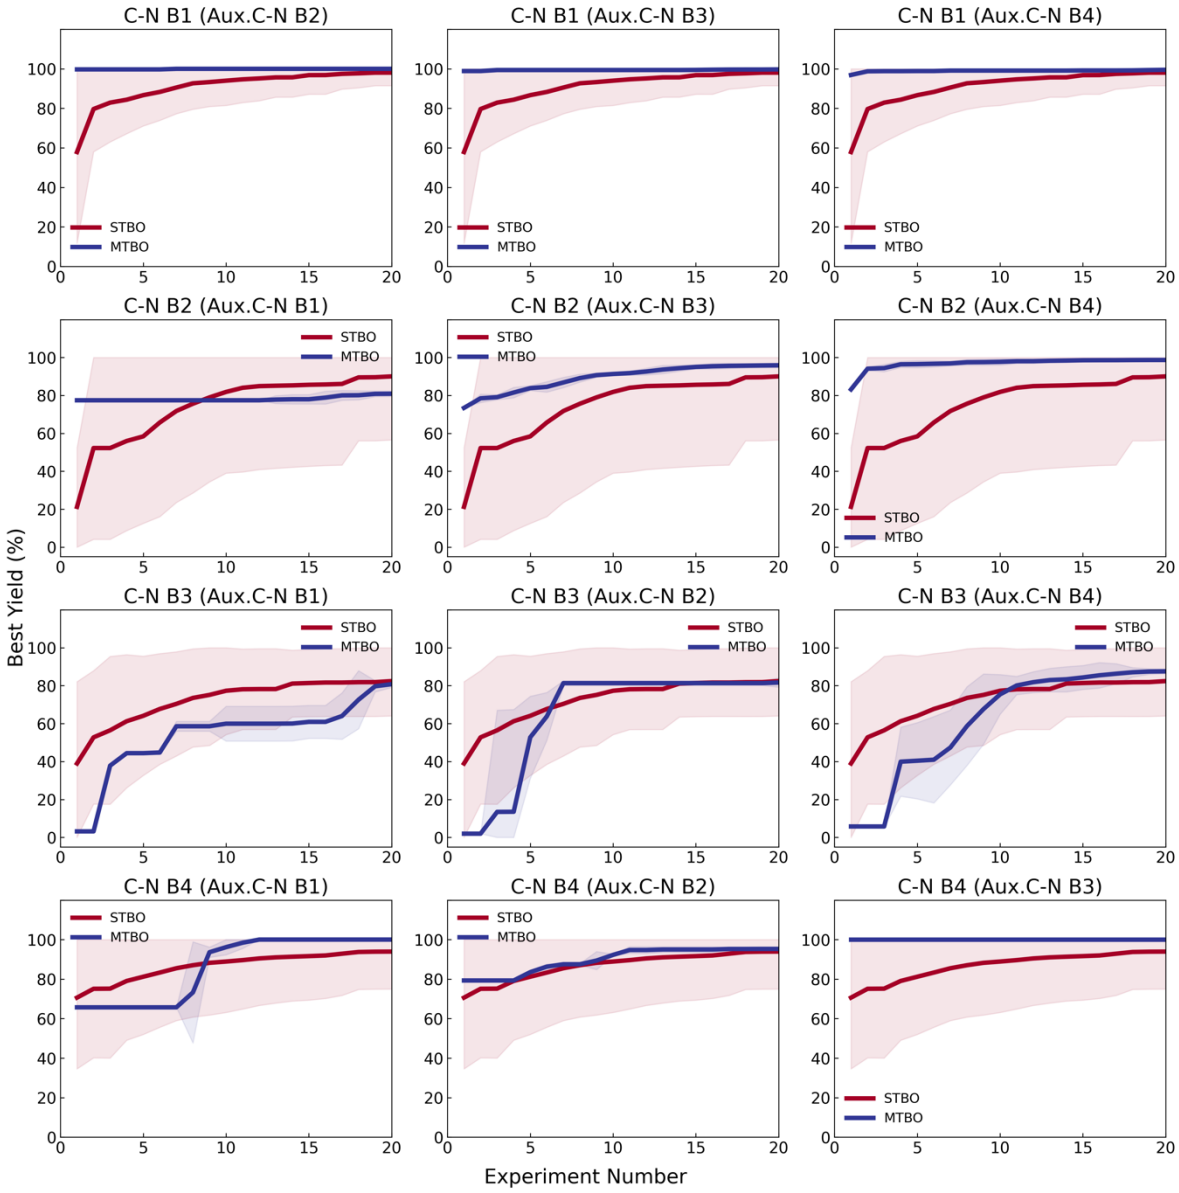

Figure S14: Comparison of the performance of single-task Bayesian optimization (STBO) and multi-task Bayesian optimization (MTBO) on C-N B1-B4 with C-N B1-B4 as auxiliary tasks. The text above the plot represents the data used as an auxiliary task.

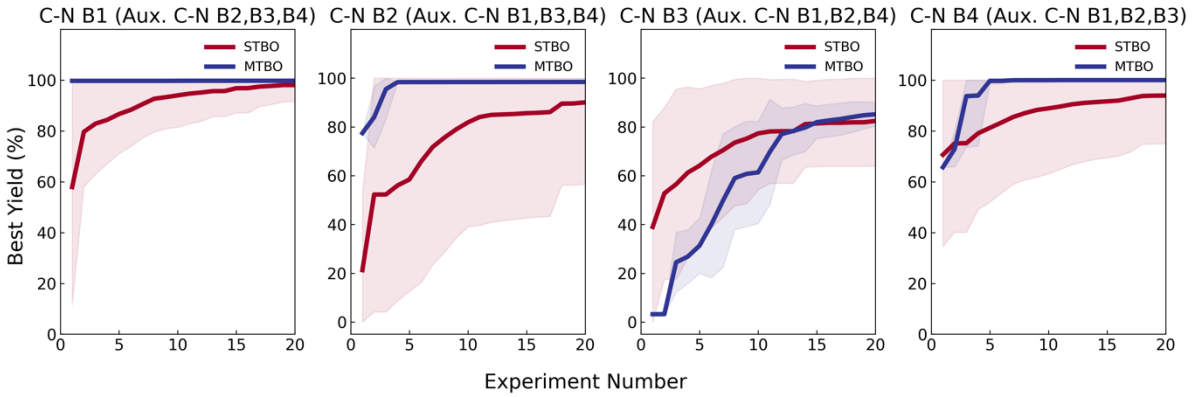

Figure S15: Comparison of the performance of single-task Bayesian optimization (STBO) and multi-task Bayesian optimization (MTBO) on C-N B1-B4 with all remaining C-N tasks for auxiliary training. The text above the plot represents the data used as an auxiliary task.

2. C-H Activation Case Studies

2.1 General procedure for synthesis of product analytical standards

The general procedure as highlighted by Hennessy and Buchwald[5] was followed. An oven-dried 50 mL round-bottomed flask was equipped with a magnetic stir bar and air condenser. Palladium acetate (2 mol%), JohnPhos (4 mol%) and the chloroacetanilide substrate (5 mmol) was then charged, and the flask was evacuated and backfilled with nitrogen 3 times. Anhydrous triethylamine (1.05 mL, 10 mmol) was added, followed by anhydrous toluene (5 mL). The reaction mixture was heated at 80 °C and ran overnight, then diluted with ethyl acetate (50 mL). The mixture was filtered through a plug of celite, concentrated on a rotary evaporator, then purified by silica gel chromatography to give the oxindole product. For each case study, the starting material was purchased from J&H Chemical in 95 % purity. All other chemicals were purchased from Sigma Aldrich unless otherwise stated.

2.2 Experimental setup

The following experimental setup was used, as pictured in Figure S16, as described in the Methods section of the paper.

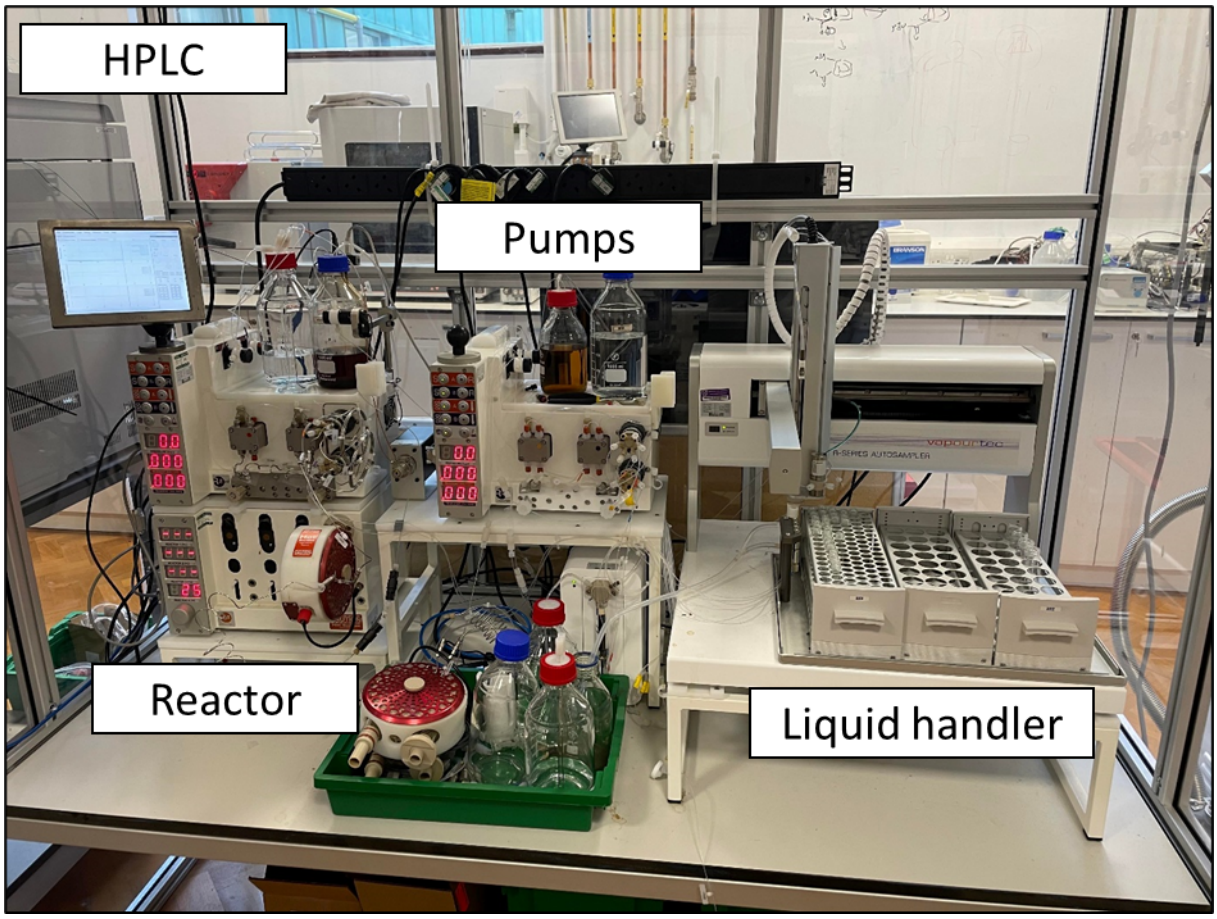

Figure S16: A photo of the experimental setup used for all C-H activation case studies.

Solubility studies were initially conducted to determine the maximum concentrations of each starting material and product in different solvents. This information would then be used to conclude the ideal concentration of the starting material in each reaction (0.09 M) and the solvents of choice to help to prevent reactor clogging. In each study, a pre-weighed vial and a fixed amount of material was dosed with small amounts of solvent and stirred vigorously until the material was fully dissolved - the vial was then weighed, and the resulting concentration calculated. These maximum concentrations are reported in Table S1:

Table S1: The maximum concentration observed for each starting material and product from each case study in a number of solvents, as determined from solubility studies.

Case study 1								
	H_2O /M	Acetonitrile /M	THF /M	Methanol /M	Acetone /M	Chloroform /M	IPA /M	Toluene /M
17	0.06	5.44	4.65	4.22	4.88	5.47	2.06	4.15
18	0.03	0.76	0.85	0.30	0.73	2.43	0.31	0.46
Case study 2								
	H_2O /M	Acetonitrile /M	THF /M	Methanol /M	Acetone /M	Chloroform /M	IPA /M	Toluene /M
19	0.01	0.66	1.60	0.05	0.71	1.89	0.07	0.54
20	0.02	0.10	0.20	0.03	0.06	1.29	0.01	0.13
Case study 3								

	H_2O /M	Acetonitrile /M	THF /M	Methanol /M	Acetone /M	Chloroform /M	IPA /M	Toluene /M
21	0.02	2.62	2.29	0.28	2.51	2.24	0.24	0.70
22	0.01	0.08	0.08	0.03	0.06	0.19	0.01	0.06
Case study 4								
	H_2O /M	Acetonitrile /M	THF /M	Methanol /M	Acetone /M	Chloroform /M	IPA /M	Toluene /M
23	0.01	5.01	3.56	4.01	4.33	5.65	1.75	3.67
24	0.01	0.82	0.71	0.23	1.10	2.56	0.21	0.49

The weighing of the individual components into the sample vials differed for each of the case studies - these vials were used by the liquid handler as solution reservoirs. The vials 1 - 5 contained these differing components, whilst vials 6 - 25 remained constant throughout. 10 mL of reaction solution was present in each vial, which was each purged with nitrogen. The triethylamine and all solvents are anhydrous, and biphenyl was used as an internal standard. These concentrations were calculated so that a constant 0.09 M concentration of starting material in the reactor could be obtained through dilution, to enable a fair comparison between all conditions where mol% of catalyst was varied. The weights of all components are displayed in Table S2 and S3:

Table S2: Individual component weighing for each 'starting material' sample vial for the C-H activation case studies.

Case study 1				
Vial	*Solvent*	***17** mass /g*	*NEt_3 mass /g*	*Biphenyl mass /g*
1	Toluene	0.40	0.30	0.03
2	DMA	0.40	0.30	0.03
3	Acetonitrile	0.40	0.30	0.03
4	DMSO	0.40	0.30	0.03
5	NMP	0.40	0.30	0.03
Case study 2				
Vial	*Solvent*	***19** mass /g*	*NEt_3 mass /g*	*Biphenyl mass /g*
1	Toluene	0.89	0.30	0.03
2	DMA	0.89	0.30	0.03
3	Acetonitrile	0.89	0.30	0.03
4	DMSO	0.89	0.30	0.03
5	NMP	0.89	0.30	0.03
Case study 3				
Vial	*Solvent*	***21** mass /g*	*NEt_3 mass /g*	*Biphenyl mass /g*
1	Toluene	0.49	0.30	0.03
2	DMA	0.49	0.30	0.03
3	Acetonitrile	0.49	0.30	0.03
4	DMSO	0.49	0.30	0.03
5	NMP	0.49	0.30	0.03
Case study 4				
Vial	*Solvent*	***23** mass /g*	*NEt_3 mass /g*	*Biphenyl mass /g*
1	Toluene	0.45	0.30	0.03
2	DMA	0.45	0.30	0.03
3	Acetonitrile	0.45	0.30	0.03
4	DMSO	0.45	0.30	0.03
5	NMP	0.45	0.30	0.03

Table S3: Individual component weighing for each 'reagent' sample vial for the C-H activation case studies.

All case studies				
Vial	*Solvent*	*Ligand*	*$Pd(OAc)_2$ mass /g*	*Ligand mass /g*
6	Toluene	JohnPhos	0.045	0.119
7	DMA	JohnPhos	0.045	0.119
8	Acetonitrile	JohnPhos	0.045	0.119
9	DMSO	JohnPhos	0.045	0.119
10	NMP	JohnPhos	0.045	0.119

11	Toluene	SPhos	0.045	0.164
12	DMA	SPhos	0.045	0.164
13	Acetonitrile	SPhos	0.045	0.164
14	DMSO	SPhos	0.045	0.164
15	NMP	SPhos	0.045	0.164
16	Toluene	XPhos	0.045	0.194
17	DMA	XPhos	0.045	0.194
18	Acetonitrile	XPhos	0.045	0.194
19	DMSO	XPhos	0.045	0.194
20	NMP	XPhos	0.045	0.194
21	Toluene	DPEPhos	0.045	0.215
22	DMA	DPEPhos	0.045	0.215
23	Acetonitrile	DPEPhos	0.045	0.215
24	DMSO	DPEPhos	0.045	0.215
25	NMP	DPEPhos	0.045	0.215

2.3 Case study 1

An analytical standard for **18** was obtained using the general procedure, where LCMS analysis gave product m/z 166.15 and the ^{1}H NMR showed conversion to product as shown in Figure S17.

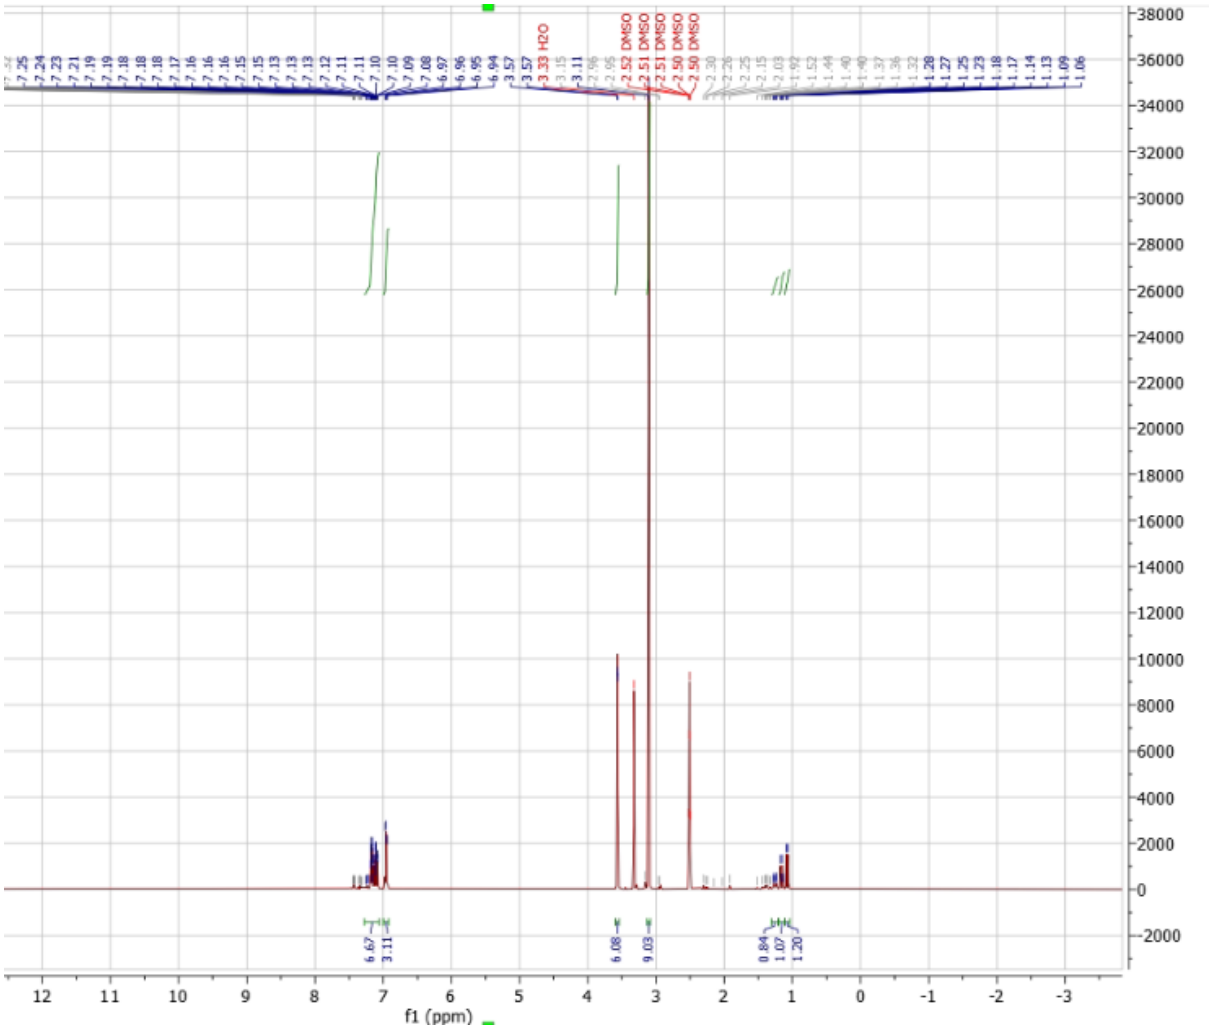

Figure S17: ^{1}H NMR for oxindole product **18**.

Table S4: All reaction data for the C-H activation case study 1.

	Conditions					
Type	Solvent	Ligand	Residence Time /min	Temp /°C	Mol%	Yield /%
Training	DMSO	DPEPhos	60	140	10	46.09
Training	DMSO	JohnPhos	45	70	3	1.31
Training	Toluene	DPEPhos	42	61	8	0.00
Training	DMA	JohnPhos	9	135	5	37.91
Training	MeCN	JohnPhos	11	134	2	16.87
Training	MeCN	JohnPhos	19	80	5	69.70
Training	NMP	JohnPhos	47	143	4	40.92
Training	DMSO	SPhos	29	122	4	0.00
Training	NMP	XPhos	52	112	8	55.69
Training	NMP	DPEPhos	26	53	6	0.00
Training	Toluene	XPhos	6	83	3	3.01
Training	MeCN	XPhos	39	90	9	6.33
Training	DMA	XPhos	51	104	9	16.11
Training	DMA	DPEPhos	22	67	4	4.04
Training	DMSO	XPhos	33	115	7	0.00
Training	Toluene	SPhos	19	57	8	0.00
Optimization	NMP	XPhos	60	102	10	56.23
Optimization	MeCN	JohnPhos	21	66	6	5.89
Optimization	Toluene	JohnPhos	53	141	8	26.36
Optimization	NMP	XPhos	60	123	9	42.37
Optimization	NMP	XPhos	53	89	9	74.64
Optimization	NMP	XPhos	46	88	9	46.62
Optimization	NMP	XPhos	60	85	8	72.40

2.4 Case study 2

An analytical standard for **20** was obtained using the general procedure, where LCMS analysis gave product m/z 409.30 and the ^{1}H NMR showed conversion to product in Figure S18. A pure product could not be obtained for this analytical standard and HPLC calibration was performed after quantitative NMR assays to determine %purity of product.

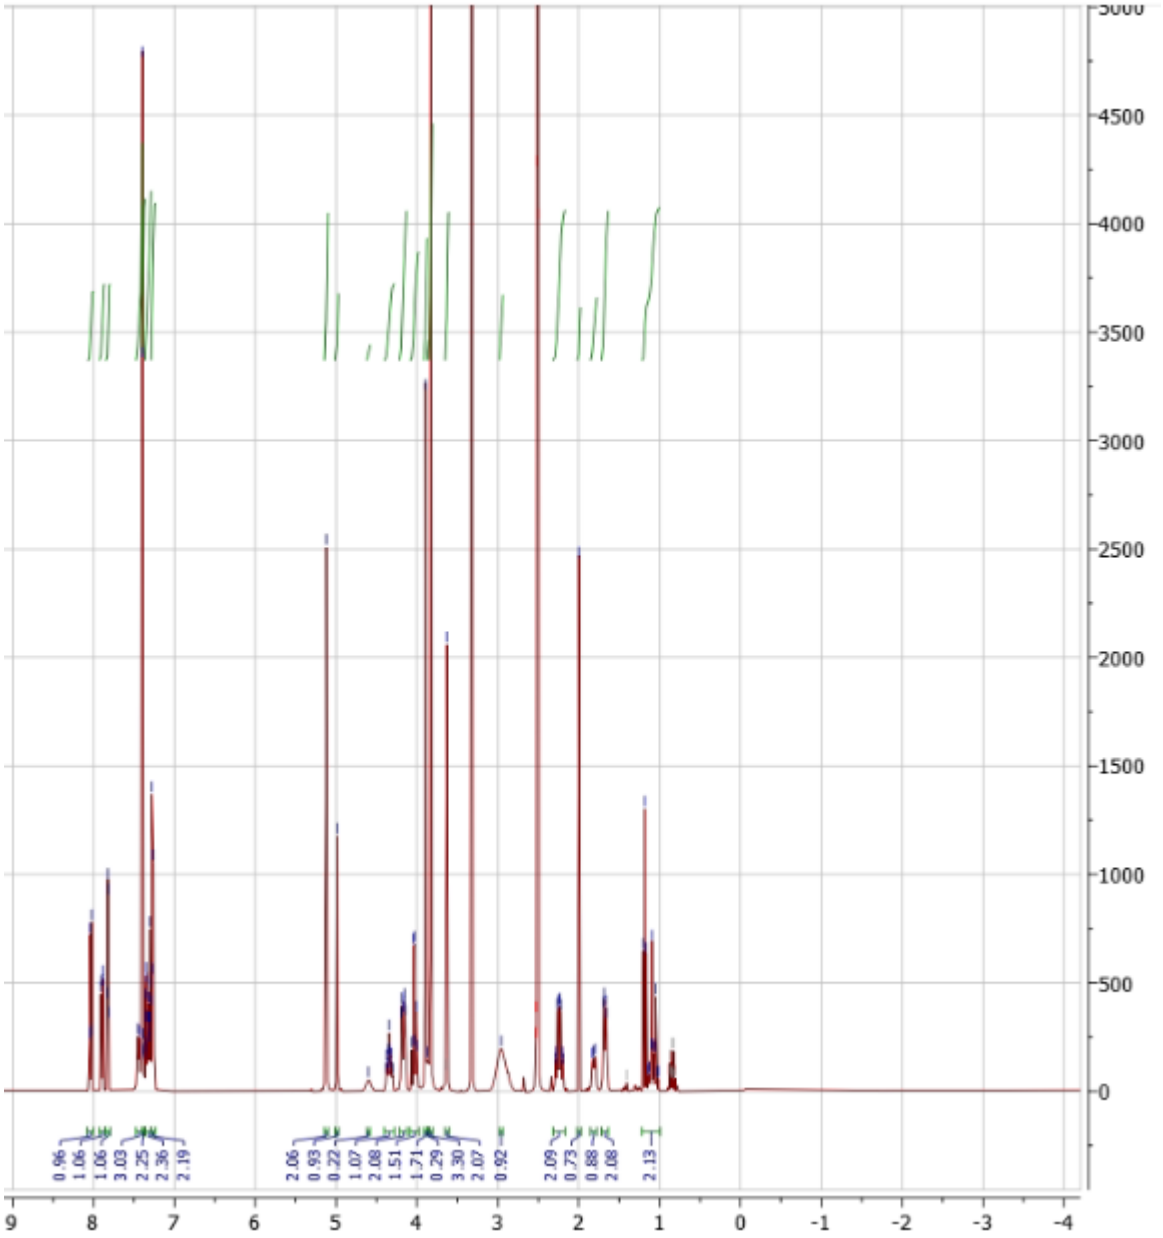

Figure S18: ^{1}H NMR for oxindole product **20**.

Table S5: All reaction data for the C-H activation case study 2.

	Conditions					
Type	Solvent	Ligand	Residence Time /min	Temp /°C	Mol%	Yield /%
Optimization	NMP	XPhos	19	75	4	14.78
Optimization	Toluene	DPEPhos	50	121	8	25.48
Optimization	MeCN	JohnPhos	21	59	7	11.10
Optimization	NMP	XPhos	58	91	9	30.98
Optimization	MeCN	JohnPhos	18	88	4	73.24
Optimization	MeCN	JohnPhos	27	93	4	76.92

Optimization	MeCN	JohnPhos	20	105	4	72.38
Optimization	MeCN	JohnPhos	40	100	8	80.11
Optimization	MeCN	JohnPhos	15	141	6	79.49
Optimization	MeCN	JohnPhos	19	117	5	84.70
Optimization	MeCN	JohnPhos	28	127	5	84.90

2.5 Case study 3

An analytical standard for **22** was obtained using the general procedure, where LCMS analysis gave product m/z 211.05 and the ^{1}H NMR showed conversion to product in Figure S19.

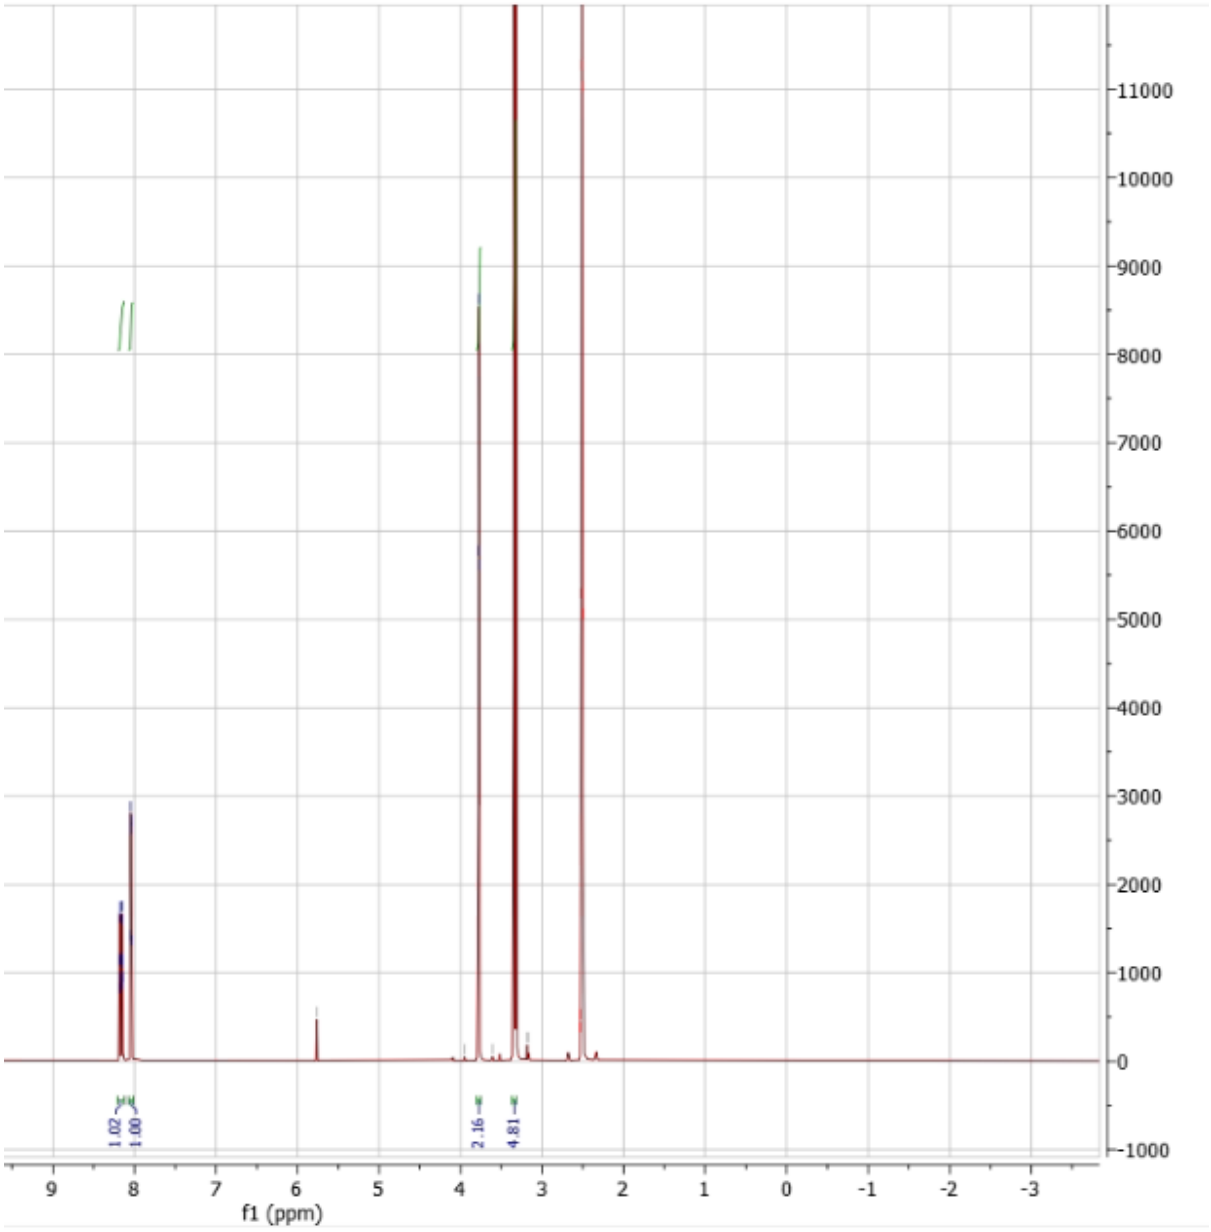

Figure S19: ^{1}H NMR for oxindole product **22**.

Table S6: All reaction data for the C-H activation case study 3.

	Conditions					
Type	Solvent	Ligand	Residence Time /min	Temp /°C	Mol%	Yield /%
Optimization	MeCN	JohnPhos	18	139	5	71.07

Optimization	NMP	XPhos	60	72	9	83.49
Optimization	NMP	XPhos	60	72	7	77.10
Optimization	NMP	XPhos	60	50	9	45.90
Optimization	NMP	XPhos	60	96	9	98.16

2.6 Case study 4

An analytical standard for **24** was purchased from Sigma Aldrich.

Table S7: All reaction data for the C-H activation case study 4.

	Conditions					
Type	Solvent	Ligand	Residence Time /min	Temp /°C	Mol%	Yield /%
Optimization	NMP	XPhos	60	96	9	6.29
Optimization	MeCN	JohnPhos	18	139	5	14.91
Optimization	Toluene	XPhos	28	125	3	1.63
Optimization	Toluene	SPhos	34	144	9	8.50
Optimization	Toluene	JohnPhos	55	140	6	9.95
Optimization	MeCN	JohnPhos	44	147	10	49.90
Optimization	DMSO	DPEPhos	59	143	9	81.97
Optimization	DMSO	DPEPhos	57	148	7	56.49
Optimization	DMSO	DPEPhos	60	150	10	82.21

2.7 Case study 5

An analytical standard for 1,7-dimethylindolin-2-one was obtained using the general procedure, where LCMS analysis gave product m/z 162.31 and the ^{1}H NMR showed conversion to product in Figure S20. A pure product could not be obtained for this analytical standard and HPLC calibration was performed after quantitative NMR assays to determine %purity of product.

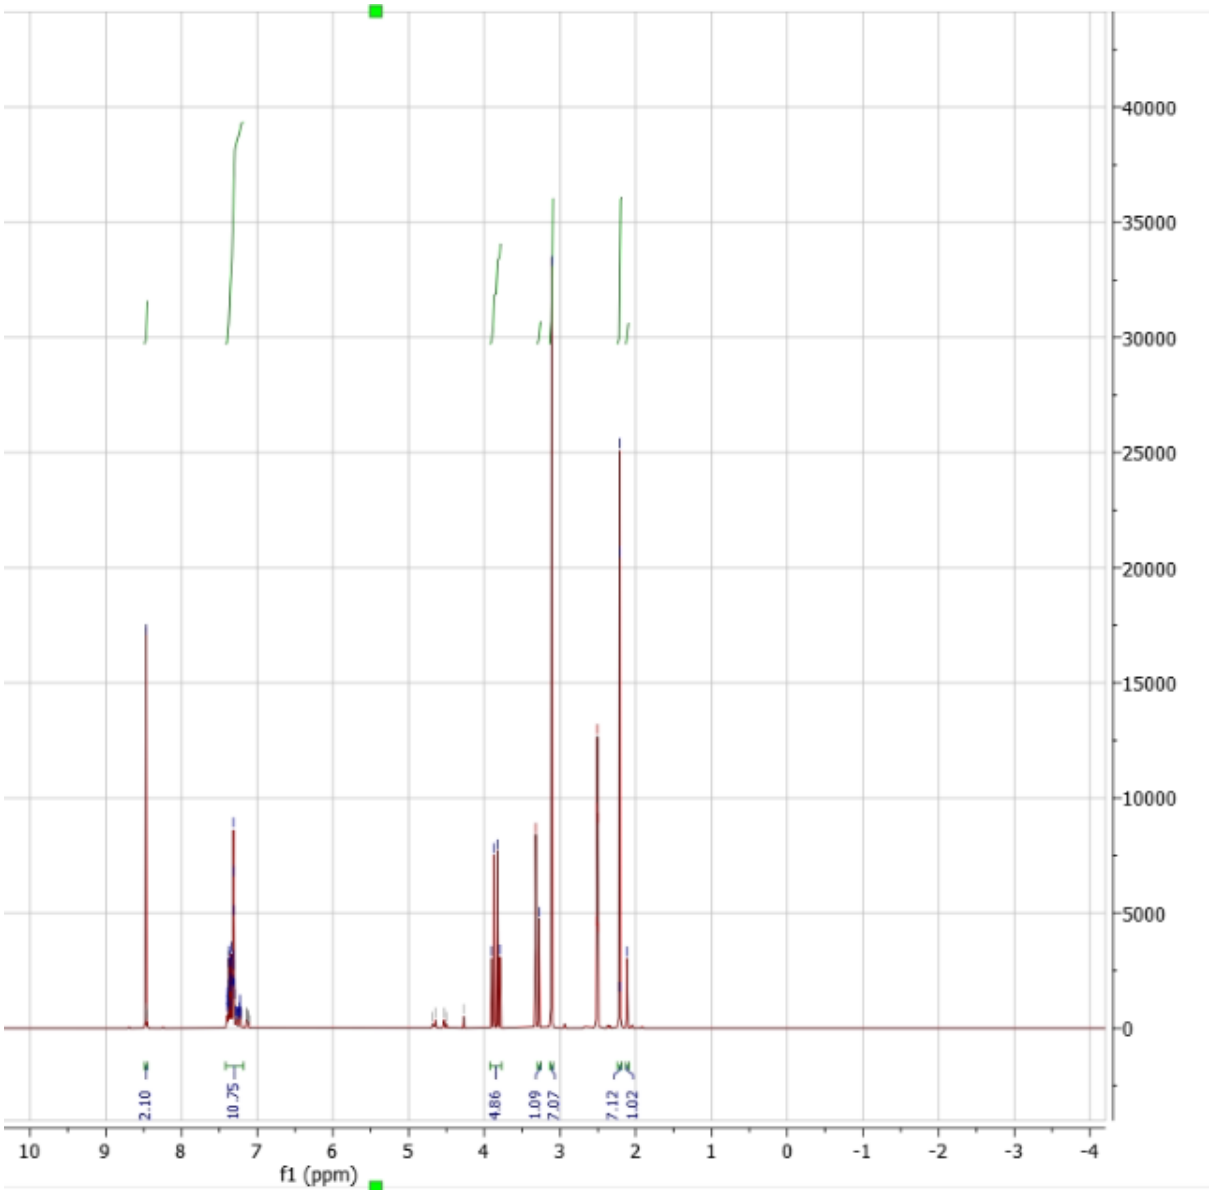

Figure S20: ^{1}H NMR for oxindole product 1,7-dimethylindolin-2-one.

Table S8: All reaction data for the C-H activation case study 5.

	Conditions					
Type	Solvent	Ligand	Residence Time /min	Temp /°C	Mol%	Yield /%
Optimization	MeCN	JohnPhos	18	139	5	17.23
Optimization	NMP	XPhos	60	72	9	3.85
Optimization	MeCN	JohnPhos	42	50	7	0.20
Optimization	NMP	DPEPhos	5	50	3	0.00
Optimization	Toluene	DPEPhos	54	150	4	1.56
Optimization	MeCN	JohnPhos	23	150	6	21.70
Optimization	MeCN	JohnPhos	25	132	6	21.50
Optimization	NMP	SPhos	36	141	9	13.70

Optimization	MeCN	JohnPhos	30	149	6	27.00
Optimization	MeCN	JohnPhos	37	150	6	29.30

2.8 Comparison with industry-standard optimization techniques

Using our experimental setup, outlined in the Methods section, and the current pricing for each of the starting material used for this work (highlighted below), it has been shown that the MTBO approach used less material and was much cheaper to implement than other standard industrial optimization techniques. The materials used are valued as follows:

- Case study 1: (2-chloro-N-(4-fluorophenyl)-N-methylacetamide) (**17**) - £392/g.[6]
- Case study 2: benzyl 4-(2-chloro-N-(4-(methoxycarbonyl)phenyl)acetamido)-piperidine-1-carboxylate (**19**) - £1960/50g.[7]
- Case study 3: 2-chloro-N-(2-fluoro-4-nitrophenyl)-N-methylacetamide (**21**) - £303/g.[8]
- Case study 4: 3-chloro-N-(4-methoxyphenyl)-N-methylpropanamide (**23**) - £189/g.[9]

As the starting material is the most expensive reaction component, this was the only factor taken into consideration - other costs were not considered, but would also lead to a significant increase in cost: catalyst, ligands, solvents, energy etc. It is assumed that a 2 mL reaction slug of 0.09 M starting material solution is used for each experiment in a flow reactor. For kinetic experiments in flow, this would also be roughly equivalent to the amount of reaction material ideal for one batch experiment with multiple time-point sampling. The addition of skilled labor and cost of time have also not been considered, but would incur a significantly higher cost for these other techniques as this scales with the number of experiments conducted. More process understanding could be gained from these techniques and the additional datapoints, but for medicinal chemistry applications of intermediate reaction optimization, the costs are substantially cheaper using MTBO.

Technique	**What is required**	**Material usage /g**				**Cost /£**	**Saving /%**
		(**17**)	(**19**)	(**21**)	(**23**)		
MTBO	Use of multi-task Bayesian optimization and experiments as required until optimal results are achieved. 2^3 initial training experiments for the first case study.	0.93	0.98	0.25	0.45	563	-
Design of experiments study	One screening design (11 experiments) and one optimization design (27 experiments) to consider the factors of: temperature, residence time and catalyst mol%. This is repeated for each combination of solvent (toluene, DMA, acetonitrile, DMSO, NMP) and ligand (JohnPhos, SPhos, XPhos, DPEPhos).	30.64	67.62	37.48	34.61	32560	98.3
Kinetic studies	3 sets of 10 time-point experiments, varying concentration and temperature at the same time in each set. This is repeated for each combination of every solvent and ligand.	24.19	53.39	29.59	27.32	25704	97.8

3. Safety statement

During experimentation, there were no unexpected or unusually high safety hazards that were encountered.

4. References

1. Kearnes, S.M., et al., *The open reaction database.* Journal of the American Chemical Society, 2021. **143**(45): p. 18820-18826.
2. Baumgartner, L.M., et al., *Use of a droplet platform to optimize Pd-catalyzed C–N coupling reactions promoted by organic bases.* Organic Process Research & Development, 2019. **23**(8): p. 1594-1601.
3. Reizman, B.J., et al., *Suzuki–Miyaura cross-coupling optimization enabled by automated feedback.* Reaction chemistry & engineering, 2016. **1**(6): p. 658-666.
4. Felton, K.C., J.G. Rittig, and A.A. Lapkin, *Summit: benchmarking machine learning methods for reaction optimisation.* Chemistry-Methods, 2021. **1**(2): p. 116-122.
5. Hennessy, E.J. and S.L. Buchwald, *Synthesis of substituted oxindoles from α-chloroacetanilides via Palladium-catalyzed C− H functionalization.* Journal of the American Chemical Society, 2003. **125**(40): p. 12084-12085.
6. *Sigma Aldrich*. https://www.sigmaaldrich.com/GB/en/product/aldrich/tmt00604. 2022.
7. J&H Chemical. Individual Enquiry. 2021.
8. *Carbosynth*. https://www.carbosynth.com/p/CAA65316/2653-16-9-2-chloro-n-methyl-n-4-nitrophenylace. 2022.
9. *AK Scientific*. https://aksci.com/item_detail.php?cat=7997AD. 2022.