Exploring Optimal Reaction Conditions Guided by Graph Neural Networks and Bayesian Optimization

Youn-Suk Choi (ysuk.choi@samsung.com)
Samsung Advanced Institute of Technology

Youngchun Kwon
Seoul National University https://orcid.org/0000-0001-7911-3670

Dongseon Lee
Samsung Advanced Institute of Technology

Jin Woo Kim
Samsung Advanced Institute of Technology

Sun Kim
Seoul National University

Article

Keywords:

Posted Date: June 7th, 2022

DOI: https://doi.org/10.21203/rs.3.rs-1621254/v1

Abstract

The optimization of organic reaction conditions to obtain the target product in high yield is crucial to avoid expensive and time-consuming chemical experiments. Advancements in artificial intelligence have enabled various data-driven approaches to predict suitable chemical reaction conditions. However, for many novel syntheses, experimental optimization of the reaction conditions is inevitable. Bayesian optimization (BO), an iterative global optimization algorithm, demonstrated exceptional performance to identify suitable reagents compared to synthesis experts. However, BO requires several initial randomly selected experimental results (yields) to train yield prediction surrogate models (approximately 10 experimental trials). In addition, parts of this process, such as the cold-start problem in recommender systems, are inefficient. Here, we present an efficient BO algorithm for reaction optimization with a message passing neural network trained on vast amounts of experimental data. We demonstrated the effectiveness of the proposed algorithm on benchmarks and in-house datasets with reference to optimization methods and human experts.

Introduction

Substantial effort has been dedicated over the past few years to develop various technologies to optimize chemical reaction conditions. Traditionally, depending on the particular scientific or engineering discipline, optimization was accomplished against a variety of criteria, for example finding the lowest energy state of a chemical structure, identifying the factors that most closely relate the molecular shape with the properties, or searching for the optimal set of conditions to increase the efficiency of experimental procedures. Many optimization technologies have been specifically developed for chemistry. For example, chemical reaction conditions can be optimized with systematic methods such as the design of experiments (DOE)[1]. More recent optimization procedures based on computational methods were designed to assist chemists to identify chemical derivatives of known drugs to best treat a given disease[2], pinpoint candidates for organic photovoltaics, predict organic reaction paths, or conduct automated experimentation[3–10] without human intervention. Often, these applications are subject to multiple local optima, and involve costly evaluations of the proposed conditions in terms of the required experimentation or extensive computation. Recently proposed optimization algorithms capable of efficiently finding local optima are gradient-based algorithms, such as gradient descent[11], conjugate gradient[12], or the more sophisticated Broyden Fletcher Goldfarb Shanno algorithm (BFGS)[13]. Bayesian optimization (BO) approaches have emerged as popular optimization solutions to search for the global optimal point[14–19]. BO schemes consist of two major steps: First, an approximation (surrogate) to the merit landscape of the conditions is constructed. Second, based on this surrogate, a new set of conditions is proposed for the next evaluation to identify the global optimum. As such, BO speculates about the experimental outcome using all previously conducted experiments, and verifies its speculations by requesting the evaluation of a new set of conditions. Several different models have been suggested for approximating the objective function areas, ranging from random forests[20] (RF), over Gaussian processes[17,18] (GP), to active learning models[21]. However, these models require numerous evaluations of data generated in the form of laboratory experiments or computations, and are thus not well suited for solving optimization problems in chemistry. This is because evaluations of the objective are often costly and materials synthesis is another major barrier in materials development as it is still carried out laboriously by human researchers. Lately, data-driven approaches have been employed to recommend conditions for specific types of reactions. The application of powerful machine-learning techniques to large datasets of organic reactions such as the Reaxys database[22] has led to major advances both in searching for possible retrosynthetic pathways[23–32] and in evaluating the feasibility of the proposed reactions[33–39] and synthetic environments. Unfortunately, the disadvantage of data-driven methods is their limited predictive performance based on data that completely deviate from the training data distribution. In particular, the data extracted from most successful studies involving chemical experiments are likely to have been biased to one side. In this study, our attempts to overcome the limited ability of advanced approaches to find the globally optimal reaction conditions led us to propose the hybrid-type dynamic reaction optimizations (HDO) method, which complements the previous two methodologies as data-driven

approaches in that it is based on a message passing neural network (MPNN) with BO. With respect to its strategy, HDO not only utilizes the accumulated experimental data but dynamically transforms the decisions according to the yield results obtained from on-going experiments. This approach enables us to efficiently explore the optimal combination of conditions compared with previous studies. Modern advances in high-throughput experimentation (HTE)[7,40,41] enabled the construction of three named datasets (Suzuki-Miyaura Reaction, Buchwald-Hartwig Reaction and Arylation Reaction) that contain different types of chemical reaction data. These data include all the capabilities of a collection consisting of a few thousand data points under a limited set of conditions in datasets from Shields's study[42,43]. In addition, we validated our proposed algorithm using additional reactions ("Ullmann Reaction" and "Chan-Lam Reaction") with synthetic experts by using our own HTE facilities, details of which are provided in Table 1.

Methods

Overview of HDO

The HDO was designed as an iterative process capable of automatically updating the acquisition function (AF_{HDO}) for calculating the priority of the next experimental reaction condition given the reaction materials, as shown in Fig. 1. The condition classifier of the MPNN[47] is trained on a given reaction scheme, expressed in the form of graph-type molecular structures, by searching for the chemical conditions within a pre-defined space, as illustrated in Fig. 1a. Considering that all the reaction conditions are not efficient and could be unnecessarily costly, HDO narrows the search space using MPNN models (f_{MPNN}) that are able to predict the chemical context (the catalyst(s), base(s) solvent(s), and ligand(s) can be changed depending on the type of named reaction) most suitable for any particular organic reactions given the reactants and product structures. Trained on nearly 10 million examples from the Reaxys database, each context condition area has an independent prediction model, which is a multi-label classification model to define several candidate sets of reaction condition materials.

Combinations of conditions, selected from the narrowed area of candidates and expected to deliver the target yield, should be chosen for the experiment. When sampling the initial conditions, in the narrowed search space, our aim is to maintain a balance between exploitation and exploration. Thus, we adapted the candidate conditions predicted by the MPNN as exploitation and selected the Maximin-Latin Hypercube sampling method[48] to ensure an effective distribution of exploration (Fig. 1b). With the software architecture in hand, our next objective was to optimize its performance by tuning the algorithm components critical to maximizing the yield of reaction optimization. Subsequent to experimentation with the initial conditions selected on the basis of HTE (Fig. 1c), the objective function for the f_{GP} model is trained. For the acquisition function, AF_{GP_UCB}, of BO we adapted upper confidence bounds (UCB), which ranked the priority of the next combination of conditions (Fig. 1d), as detailed in a separate subsection. Finally, HDO calculates the priority of the next candidates in the form of an ensemble by considering the historical results, MPNN (f_{MPNN}), and the optimization model GP (Fig. 1e). Depending on the outcome, the search space could be expanded to include additional reaction conditions (Fig. 1f). For maximum efficiency, HDO is designed to make comprehensive judgments using not only the results predicted by the MPNN but also the experimental results, frequency of past experiments, and uncertainties in the objective function of the predictive model. The data preprocessing and model formulation steps are described in detail in subsequent subsections. HDO proceeds iteratively until it finds the global optimal combination of conditions that produce the desired target yield and updates the objective functions whenever the complete results of each experiment are known. The proposed approach offers a platform on which fully automated organic synthesis experiments can be carried out by using robots and management software.

Dataset and graph-type representation for training MPNN

The dataset with all of the reactions with their conditions for training MPNN was extracted from the Reaxys reaction database consisting of 53 million reaction records. The data include structural expressions of the reactants and single products, conditions, and isolated yields. We used the structural expressions of the reactants and each single product, the Reaxys chemical ID, and simplified molecular-input line-entry system (SMILES) notation (if available) or the name of the reaction. Each chemical reaction is labeled with the reagents that participate in the reaction. Each instance is represented as *(R, P, y)*, where $R = \{G_{r1}, G_{r2}\}_k$ and $P = \{G_p\}_k$ are the set *k* of the two reactants with the product structures in the reaction, respectively, and *y* is the one-hot vector of the reaction conditions such as catalysts, bases, solvents, and ligands. Owing to the different circumstances of each synthesis experiment, we did not include the reaction condition datasets from the Reaxys database based on a pre-defined list of reagents in each experiment. In addition, the number of reaction sets *k* and classes of conditions can be different for each type of condition and optimization task, respectively. Apart from this, we restrict our scope to include single-product and single-step reactions to ensure a closer alignment with the application to computer-aided synthesis planning. We also note the ambiguous labeling of certain catalysts and reagents in the Reaxys database, in which many catalysts are recorded as reagents, causing the data to be sparser for catalysts and increasing the number of distinct reagents. This issue can hardly be eliminated completely because a strict separation between reagents and catalysts would be difficult to achieve.

Figure 2. **Illustration of process of MPNN models to predict suitable reaction conditions given graph-type reaction representations** G.

Message Passing Neural Networks for predicting suitable reagents

A graph is data structure that presents a powerful non-Euclidean method for establishing the extent to which features (nodes) are connected to their relationships (edges). We adapted the message passing neural network (MPNN) to process each molecular graph *G* in $r_{1,2}$ and *p*. The MPNN is designed to accept *G* as input and to return the graph representation vector *r* as output:

$q = M_\theta(G)$.

The MPNN uses six message-passing steps with an edge network as the message function and a GRU network as the update function to produce node representation vectors, the dimensionality of which is set to 64. Then, a set2set model, which uses six processing steps, is employed as the readout function for global pooling over the node representation vectors to obtain graph-level embedding, which is invariant to the order of the nodes. The embedding is further processed by a fully connected layer of 512 ReLUs[49], resulting in the graph representation vector *q*. The use of the MPNN ensures that the representation is invariant to graph isomorphism. We summate the respective graph representation vectors regarding $R = \{G_{r,1}, G_{r,2}\}$ and $P = \{G_p\}$. In this manner, the representation becomes invariant with respect to the order of reactants and products. The two summated vectors of reactants are concatenated to produce a reaction representation vector h (Eq.(1)):

$$h = \left[\sum_{l=1}^{k} qR, l, \sum_{l=1}^{i} qP, l\right] \quad (1)$$

The reaction graph embedding vector h is further processed by a feed-forward neural network (FNN) having four fully connected layers, each of which contains 512 ReLUs. The output functions of the FNN are equal to the length of the one-hot vectors by each condition types ($\hat{y}$). The predictive model MPNN is trained as independent models according to the types of reaction conditions and the predicted catalyst, ligand, base, and solvent as $\hat{y}_{category}$, respectively. The value of the priority is calculated by generating all possible combinations of conditions using the weight value of the one-hot

vector predicted by the model for each of the conditions. Therefore, the priority $\hat{y}_{priority}$ of condition combination x is defined as (Eq. (2)):

$$\text{Priority (y)} = f_{MPNN}(x_k) \quad (2)$$

The types of conditions vary among experiments and have different ranges. Figure 2 illustrates the training process for the condition prediction model MPNN.

Bayesian Optimizations in the HDO

The purpose of the BO method is to reduce the number of objective evaluations that need to be performed to solve the optimization problem. For this, they iteratively suggest, in a careful and intelligent way, an input location in which the objective that is being optimized should be evaluated for each of the experiments. For this, at each iteration $N = 1, 2, 3 \ldots$ of the optimization process, the BO method fits a probabilistic model, a Gaussian process (GP) in our case, to the collected observations of the objective. The uncertainty in the potential values of the objective is provided by the predictive distribution of the GP. This uncertainty is used to generate an acquisition function AF_{GP_UCB}, whose value at each input location indicates the expected utility of evaluating f_{GP} there. The upper-confidence-bound algorithm (UCB)[50] is adapted to calculate the priority of the next best combinations of conditions as (Eq. (3)).

$$AF_{GP\_UCB}\left(x_k\right) = argmax(f_{GP}\left(x_k\right) + \sqrt{\frac{2In(t)}{N\left(x_k\right)}}) \quad (3)$$

The next point x_k at which to evaluate f_{GP} is the one that maximizes AF_{GP_UCB}. After collecting this observation, the process is repeated. When sufficient data have been collected, the GP predictive mean value f_{GP} can be optimized to find the solution of the problem. Considering that the acquisition function AF_{GP_UCB} is calculated by using only the results currently being experimented with, it is possible to correct the process of navigating in the wrong direction resulting from the inaccurate prediction results of the MPNN model.

Acquisition function of the HDO and rules for expanding the search space

In this section, we present our technique, named the learning to acquisition function, for efficient reaction optimization. To determine the next iterate x_k based on the belief about f_{MPNN} and AF_{GP_UCB} given the history H_k, a sampling strategy is defined as follows:

$$AF_{HDO}\left(x_k\right) = \frac{N\left\{f_{MPNN}\left(x_k\right)\right\}}{\ln(t)} + N\{AF_{GP\_UCB}\left(x_k\right)\} \quad (4)$$

In Eq. (4), the AF_{HDO} was designed to efficiently optimize by combining the current experiment-driven priority model AF_{GP_UCB} with the priority f_{MPNN}, obtained by using a vast number of experimental documents for training. The normalization N is intended to prevent bias to one side by the values of the two acquisition functions and is calculated as Eq. (5) by min-max feature scaling.

$$N = \frac{X - X_{min}}{X_{max} - X_{min}} \quad (5)$$

Under the influence of the weight of AF_{GP_UCB}, AF_{HDO} is induced to increase with the number of trials t to more closely reflect the results of the current experiments H. To dynamically expand the search space that was initially narrowed by

MPNN models, if t exceeds 20 experimental trials and the accumulated maximum conversion yield is less than 10%, the search space is expanded by 10% every five experiments.

The MPNN model was implemented using PyTorch in Python. The BO module is facilitated by “scikit-optimize python library” and we used the “skopt.gp_minimize” function for AF_{GP_UCB}. The results of the experimental investigations are reported and discussed in the following section.

Performance Benchmarking Results

The effectiveness of the HDO method is demonstrated using two types of optimizing tasks on datasets from Shields’s study (Task 1) and on in-house data accumulated using HTE equipment (Task 2).

Task 1: A dataset comprising reaction conditions pertaining to the Suzuki-Miyaura reaction (1a) was released by Shields. They conducted high-throughput experiments on the class of Suzuki-Miyaura cross-coupling reactions. Twelve couplings of electrophiles and nucleophiles across the combinations of 11 ligands, 7 bases, and 4 solvents were considered, thereby resulting in combinations for a total of 3,696 reactions with a product. Buchwald-Hartwig reaction (2a-2e): They also conducted high-throughput experiments on the class of Pd-catalyzed Buchwald-Hartwig C-N cross-coupling reactions. They experimented with combinations of 3 aryl halides, 4 catalyst, 3 bases, and 22 additives for a total of 792 reactions per target product, of which there were five types. Arylation (3a): They studied the arylation of imidazoles, a key step in the commercial synthesis of the JAK2 inhibitor BMS-911543[44,45]. They selected a subspace consisting of 1,728 reactions including 12 ligands, 4 bases, 4 solvents, 3 temperatures, and 3 concentrations as a tractable set of experiments to be used as the ground truth. The data for the arylation reaction (3a) included results contributed by 50 expert chemists and engineers from academia and industry, who played the reaction optimization game. All of the data used in Task 1 are accessible from https://github.com/b-shields/edbo/ .

Task 2: Because HDO provides more general evaluations compared with skilled synthetic experts, a total of 22 experiments were conducted on four named reactions (details of Task 2 appear in Table 1). In this task, HDO was evaluated by comparing the yield results of the reaction conditions proposed by five experts in organic synthesis. A search space for each of the Suzuki-Miyaura (4a-4j), Buchwald-Hartwig (5a-5h), Ullmann (6a, 6b), and Chan-Lam reactions (7a, 7b) was defined by the experts and details thereof are provided in Table 1. Details are described in Supplement Information (S1 and S2).

Table 1. Details of the two performance benchmarking tasks

Summary of task sets with definition of search space

		Task1 (vs. baselines, 50 humans)			Task2 (vs. only 5 humans)			
	Target conditions /Reaction type	Suzuki-Miyaura (1a)	Buchwald-Hartwig (2a-2e)	Arylation (3a)	Suzuki-Miyaura (4a-4j)	Buchwald-Hartwig (5a-5h)	Ullmann Reaction (6a, 6b)	Chan-lam Reaction (7a, 7b)
Target number of conditions	Reactant1	3	-	-	-	-	-	-
	Reactant2	4	3	-	-	-	-	-
	Additive	-	22	-	-	-	-	-
	Catalyst	-	4	-	9	9	7	7
	Base	7	3	4	8	9	13	3
	Solvent	4	-	4	4	5	5	4
	Ligand	11	-	12	12	12	9	-
	Concentration	-	-	3	-	-	-	-
	Temperature	-	-	3	-	-	-	-
Search space area		3,696	792	1,728	3,456	4,860	4,095	84
> 95% (yields)		1.92%	0.48%	0.58%	-	-	-	-
Number of target products		1	5	1	10	8	2	2
Reactions for training		1,227,756	8,541	49,625	158,605	17,705	10,518	2,694

Task 1: Optimization of reaction conditions to benchmark the performance

We investigated the effectiveness of the proposed method for optimizing the chemical conditions on the Suzuki-Miyaura, Buchwald-Hartwig, and Arylation reactions. We first targeted the number of trials (NT) to identify the top 1, 5, and 10% target yields in the entire search space of reaction conditions as a performance measurement. We calculated the average improvement rate (AIR) over a random selection of NT^rs as performance indicator to allow for comparison with other baselines (Eq. (6)). Here, N is the number of epochs required for the overall optimization process and we used 50 epochs for all the models. In each optimization model, NT_n^c is the NT in the c category of baseline models (Random Forest (RF)[46], BO[42], humans[42], MPNN[47], and HDO), respectively. AIR represents the average improvement in the performance of a model over the naïve approach that one might use as a starting point for experimental optimization. In other words, AIR is the extent (percentage) to which the model differs from the naïve model. AIR 0 means the same number of trials with random selection. When AIR has a negative value, the number of trails to find the combination of conditions with a target yield exceeds the random selection.

Table 2. Comparison of reaction optimization performance with the baselines on Task 1.

Table 2. Summary of evaluation

Reaction	target	Top%	Target yield	RandomSearch	RF	BO	MPNN	Human (50 experts)	HDO (Ours)
				ANT^r	Average Improvement Rate over RS (AIR)				
Suzuki-Miyaura	1(a)	1	*96.20*	*98.15*	0.1844	0.5885	0.6589	-	**0.6685**
		5	*92.52*	*19.55*	0.0925	0.3750	**0.4740**	-	**0.4740**
		10	*88.19*	*9.51*	0.2460	0.4440	**0.6900**	-	**0.6900**
Buchwald-Hartwig	2(a)	1	*52.67*	*99.15*	0.2244	**0.7585**	-0.1985	-	0.6642
		5	*47.94*	*18.54*	0.1175	**0.4245**	-0.1725	-	0.3945
		10	*44.65*	*8.99*	0.1060	**0.4850**	0.3850	-	0.4420
	2(b)	1	*83.09*	*97.54*	0.2046	0.5846	0.5149	-	**0.6885**
		5	*79.06*	*19.54*	0.0745	0.4395	0.3220	-	**0.5505**
		10	*73.58*	*9.15*	0.1850	0.4550	0.3850	-	**0.5850**
	2(c)	1	*94.34*	*98.01*	0.2844	0.5545	-0.0846	-	**0.5585**
		5	*86.76*	*19.44*	0.0675	**0.3725**	-0.0670	-	0.3425
		10	*81.27*	*8.95*	0.1890	**0.5420**	-0.2510	-	0.3850
	2(d)	1	*65.46*	*98.18*	0.2549	0.6185	-0.1785	-	**0.6745**
		5	*52.56*	*19.24*	0.2175	0.1940	0.1440	-	**0.2210**
		10	*49.15*	*8.97*	0.2850	0.5010	0.3500	-	**0.5850**
	2(e)	1	*97.56*	*99.54*	0.3085	0.6785	-0.2485	-	**0.6986**
		5	*91.05*	*19.54*	0.2245	0.3945	0.1245	-	**0.4125**
		10	*86.25*	*8.98*	0.1850	0.3850	0.1200	-	**0.4850**
Arylation	3(a)	1	*91.21*	*98.66*	0.1846	**0.7485**	0.4785	0.6382	0.7285
		5	*76.54*	*19.81*	0.2420	0.4210	0.3345	0.4225	**0.4560**
		10	*59.11*	*9.43*	0.4440	0.5850	0.4210	0.6750	**0.6950**

Bold numbers indicate the highest performing model for each reaction target.

$$AIR = 1 / N \sum_{n=1}^{n} \frac{\left(NT_n^r - NT_n^c \right)}{NT_n^r}$$

(6)

The results in Table 2 indicate that HDO significantly outperformed the base models across the categories for the top 1, 5, and 10% optimization tasks. While MPNN actually produced negative top 1% AIRs over RS in some cases, these scores showed the limitations of an optimization approach based on only exploitation. In the case of 2(a) and 2(c), HDO has difficulties with the top 1, 5, and 10% tests over BO, and this could be the result of poor performance by MPNN; excluding this target case, the HDO outperformed the baseline. A few interesting reaction type trends can be seen across the model types. For instance, in the Suzuki-Miyaura reaction type, the training dataset for MPNN contains approximately 4,200,000

reaction conditions, whereas the Buchwald-Hartwig and Arylation datasets contain approximately 40,600. Therefore, MPNN models the found best combination faster than others. Similarly, the HDO model also performed well in the Suzuki-Miyaura reaction over BO, which still has a cold start problem. Likewise, for the top 5 and 10% yield searching cases in Arylation reaction 3(a), the AIR of HDO was higher than that of the human experts and BO, although BO is the best for the top 1%. Overall, the proposed model HDO found the best combination of conditions for a high yield compare to other optimization algorithms and human experts. (The results of 3(a) are shown in detail in Supplementary Information S4 sector)

Task 2: Validation of the HDO against five human chemists

This task was designed to additionally verify the performance of HDO in terms of general organic synthesis experiments compared to chemists. Task 1 involved a comparison of various optimization models and repeat simulations on the basis of Shields's dataset, which includes yield results **Fig. 3** Average cumulative maximum observed yields by the HDO (blue curve) and the average yield of the combination of conditions proposed by five experts (black dotted line). HDO required an average of 4.7 experiments to find the conditions with the same average yield obtained with the combination of conditions proposed by the five experts for the 22 reactions.

corresponding to all combinations of reaction conditions. However, it is difficult to obtain results from 1,000–4,000 different combinations of the conditions for one reaction. We therefore conducted experiments to create various additional target molecules to assess the stability of the HDO performance in comparison with that of humans. In Task 2, we used 22 different reactions to determine the number of experimental trials required to reach the yield results referenced by five specialists in the field of chemical synthesis. In these experiments, the four named reactions (the Suzuki-Miyaura, Buchwald-Hartwig, Chan-Lam, and Ullmann reactions) were used to synthesize the 22 target products. In addition, all the synthesis experiments included in Task 2 were also conducted on the HTE platform. In these tests, we focused on optimizing four aspects of the reaction context (catalyst, solvent, base, ligands) for each individual reaction. The results are plotted in Fig. 3. For the all 22 reactions, the best combination of reagents proposed by the five synthesis experts had a yield value of 64.48 on average. In comparison, HDO achieved this yield value in 4.7 experiments (i.e., by using 4.7 combinations of the reaction conditions) and average yield values of 82.18 were obtained by repeating the experiments 50 times (Fig. 3). The performance varied for each of the four aforementioned named reactions. Because the Reaxys database contains 158,605 training data for the Suzuki-Miyaura reaction, the HDO based on MPNN quickly identified the reaction conditions that delivered the yield of the combination of reaction conditions recommended by experts in an average of 4.22 times for 10 reactions. However, as is evident from the cumulative number of experiments, synthesis with the Suzuki-Miyaura reaction is easier than with the other named reactions. Therefore, the experts also tended to swiftly determine an effective combination of conditions for this reaction. In the eight experiments based on the Buchwald-Hartwig reaction, both HDO and the experts experienced difficulties to identify reaction conditions with high yield. Except for the reaction shown in 5(c) in Fig. 5, all of the reactions yielded poor results, yet even in these difficult situations, HDO found expert-level yields after 1.9 trails on average. For the Ullmann and Chan-Lam reactions, HDO required an average of 7.15 and 3.84 attempts, respectively, to identify the combination of conditions proposed by the experts. Details are provided in Fig. 4.

Finally, Fig. 5 presents examples of fragments of each of the named reactions. In the Suzuki-Miyaura reaction in 4(a), HDO proposed the same combination of conditions with the same yield value as the experts, and required a single experiment. In examples 5(c) and 7(a) for the Buchwald-Hartwig and Ullmann reactions, respectively, HDO obtained higher yield results than the experts with different combinations of reaction conditions and requiring two experimental trials in both cases. These are examples of the optimal combination of conditions found in the initial five experiments proposed by the MPNN models. In these cases, similar experiments were included in the training data, which are well-predicted examples. In marked contrast, the Chan-Lam coupling reaction in 6(b) required 20 experimental trails to identify reaction conditions with a conversion yield similar to that of the reference. Nevertheless, the proposed HDO algorithm

demonstrated optimization performance comparable to that of the experts and was shown to have reliable and efficient navigation capabilities for a variety of combinations of reaction conditions.

Discussion

In this paper, we proposed the HDO algorithm based on graph neural networks and Bayesian optimization to efficiently determine the optimal reaction conditions for organic synthesis. In order to train the yield prediction model, Shields's study[42,43], which required experiments to be conducted under the conditions selected as random for the initial 10 times, is inefficient in the case of a relatively easy reaction. Conversely, the prediction results of the MPNN model, which simply trains experimental data extracted from the papers, is not working in an entirely novel reaction experiment cases. To represent the numerical difficulty of a reaction experiment is a challenge. However, we designed to efficiently change the process of dynamically finding optimal conditions to match the results of the on-going experiment. The proposed algorithm is designed to utilize vast amounts of synthetic information to rapidly find the optimal experimental conditions as well as to reliably search for the best conditions under which to conduct experiments when developing novel reactions. That is based on an inexpensive acquisition function we formulated by balancing the explorative and exploitative behavior of the algorithm. This approach enables an intuitive sampling policy to efficiently accomplish global optimization. As a result, in experimental simulations HDO could find the optimal conditions that satisfied the target yield faster than other algorithms that rely only on using Bayesian optimization.

In addition, we verified the time required by the HDO algorithms to find the optimal combination of conditions proposed by the specialists in organic synthesis. The HDO approach also met the target yield by more swiftly identifying a combination of reaction conditions that are either the same or similar to those proposed by the synthesis experts (requiring approximately 5–10 times less time) for four named reactions. Ultimately, we expect this method to serve as an enabling tool for searching for promising chemical species and optimizing the structures of materials for various applications in the field of materials discovery.

Declarations

Data availability

Processed reaction outcome data for task1 is available at https://github.com/b-shields/edbo and task2's optimization details with search space is in Supplementary Information. The additional data that support the findings of this study are available from the corresponding author upon reasonable request.

Supporting Information

The Supporting Information is available.

Author contributions

Y.K., Y.S.C., and S.K. designed the study, implemented the methodology, and wrote the manuscript. D. L, J.W.K defined search space and validated proposed method. Y.S.C and S.K. supervised the research. All authors participated in drafting the manuscript and approved the final version.

Competing interests

The authors declare to competing interests.

References

1. I., J. O. The Design of Experiments. *Nature* **137**, 252–254 (1936).
2. 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).
3. Nikolaev, P., Hooper, D., Perea-López, N., Terrones, M. & Maruyama, B. Discovery of Wall-Selective Carbon Nanotube Growth Conditions via Automated Experimentation. *ACS Nano* **8**, 10214–10222 (2014).
4. Fitzpatrick, D. E., Battilocchio, C. & Ley, S. V. A Novel Internet-Based Reaction Monitoring, Control and Autonomous Self-Optimization Platform for Chemical Synthesis. *Org. Process Res. Dev.* **20**, 386–394 (2016).
5. Nikolaev, P. *et al.* Autonomy in materials research: a case study in carbon nanotube growth. *npj Comput. Mater.* **2**, 16031 (2016).
6. Duros, V. *et al.* Human versus Robots in the Discovery and Crystallization of Gigantic Polyoxometalates. *Angew. Chemie Int. Ed.* **56**, 10815–10820 (2017).
7. Coley, C. W. *et al.* A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science (80-.).* **365**, (2019).
8. Gao, W., Raghavan, P. & Coley, C. W. Autonomous platforms for data-driven organic synthesis. *Nat. Commun.* **13**, 1075 (2022).
9. Pyzer-Knapp, E. O. *et al.* Accelerating materials discovery using artificial intelligence, high performance computing and robotics. *npj Comput. Mater.* **8**, 84 (2022).
10. Huo, H. *et al.* Semi-supervised machine-learning classification of materials synthesis procedures. *npj Comput. Mater.* **5**, 62 (2019).
11. Ruder, S. An overview of gradient descent optimization algorithms. (2016).
12. Hestenes, M. R. & Stiefel, E. Methods of conjugate gradients for solving linear systems. *J. Res. Natl. Bur. Stand. (1934).* **49**, 409 (1952).
13. Fletcher, R. *Practical Methods of Optimization*. (John Wiley \& Sons, 1987).
14. Kushner, H. J. A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *J. Basic Eng.* **86**, 97–106 (1964).
15. Močkus, J. On bayesian methods for seeking the extremum. in 400–404 (1975). doi:10.1007/3-540-07165-2_55.
16. Mockus, J. B. & Mockus, L. J. Bayesian approach to global optimization and application to multiobjective and constrained problems. *J. Optim. Theory Appl.* **70**, 157–172 (1991).
17. Snoek, J., Swersky, K., Zemel, R. S. & Adams, R. P. Input Warping for Bayesian Optimization of Non-stationary Functions. (2014).
18. Snoek, J., Larochelle, H. & Adams, R. P. Practical Bayesian Optimization of Machine Learning Algorithms. (2012).
19. Srinivas, N., Krause, A., Kakade, S. M. & Seeger, M. W. Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting. *IEEE Trans. Inf. Theory* **58**, 3250–3265 (2012).
20. Hutter, F., Hoos, H. H. & Leyton-Brown, K. Sequential Model-Based Optimization for General Algorithm Configuration. in 507–523 (2011). doi:10.1007/978-3-642-25566-3_40.
21. Settles, B. Active Learning. *Synth. Lect. Artif. Intell. Mach. Learn.* **6**, 1–114 (2012).
22. Goodman, J. Computer Software Review: Reaxys. *J. Chem. Inf. Model.* **49**, 2897–2898 (2009).
23. Szymkuć, S. *et al.* Computer-Assisted Synthetic Planning: The End of the Beginning. *Angew. Chemie Int. Ed.* **55**, 5904–5937 (2016).
24. Segler, M. H. S., Preuss, M. & Waller, M. P. Learning to Plan Chemical Syntheses. (2017) doi:10.1038/nature25978.

25. Law, J. *et al.* Route Designer: A Retrosynthetic Analysis Tool Utilizing Automated Retrosynthetic Rule Generation. *J. Chem. Inf. Model.* **49**, 593–602 (2009).
26. Liu, B. *et al.* Retrosynthetic Reaction Prediction Using Neural Sequence-to-Sequence Models. *ACS Cent. Sci.* **3**, 1103–1113 (2017).
27. Bøgevig, A. *et al.* Route Design in the 21st Century: The IC SYNTH Software Tool as an Idea Generator for Synthesis Prediction. *Org. Process Res. Dev.* **19**, 357–368 (2015).
28. Coley, C. W., Rogers, L., Green, W. H. & Jensen, K. F. Computer-Assisted Retrosynthesis Based on Molecular Similarity. *ACS Cent. Sci.* **3**, 1237–1245 (2017).
29. Segler, M. H. S., Preuss, M. & Waller, M. P. Planning chemical syntheses with deep neural networks and symbolic AI. *Nature* **555**, 604–610 (2018).
30. Kim, E., Lee, D., Kwon, Y., Park, M. S. & Choi, Y.-S. Valid, Plausible, and Diverse Retrosynthesis Using Tied Two-Way Transformers with Latent Variables. *J. Chem. Inf. Model.* **61**, 123–133 (2021).
31. Tetko, I. V., Karpov, P., Van Deursen, R. & Godin, G. State-of-the-art augmented NLP transformer models for direct and single-step retrosynthesis. *Nat. Commun.* **11**, 5575 (2020).
32. Ucak, U. V., Ashyrmamatov, I., Ko, J. & Lee, J. Retrosynthetic reaction pathway prediction through neural machine translation of atomic environments. *Nat. Commun.* **13**, 1186 (2022).
33. Kayala, M. A., Azencott, C.-A., Chen, J. H. & Baldi, P. Learning to Predict Chemical Reactions. *J. Chem. Inf. Model.* **51**, 2209–2222 (2011).
34. Kayala, M. A. & Baldi, P. ReactionPredictor: Prediction of Complex Chemical Reactions at the Mechanistic Level Using Machine Learning. *J. Chem. Inf. Model.* **52**, 2526–2540 (2012).
35. Jin, W., Coley, C. W., Barzilay, R. & Jaakkola, T. Predicting Organic Reaction Outcomes with Weisfeiler-Lehman Network. (2017).
36. Coley, C. W., Barzilay, R., Jaakkola, T. S., Green, W. H. & Jensen, K. F. Prediction of Organic Reaction Outcomes Using Machine Learning. *ACS Cent. Sci.* **3**, 434–443 (2017).
37. Maser, M. R. *et al.* Multilabel Classification Models for the Prediction of Cross-Coupling Reaction Conditions. *J. Chem. Inf. Model.* **61**, 156–166 (2021).
38. Gao, H. *et al.* Using Machine Learning To Predict Suitable Conditions for Organic Reactions. *ACS Cent. Sci.* **4**, 1465–1476 (2018).
39. Kwon, Y., Lee, D., Choi, YS. *et al.* Uncertainty-aware prediction of chemical reaction yields with graph neural networks. *J Cheminform* **14**, 2 (2022).
40. Perera, D. *et al.* A platform for automated nanomole-scale reaction screening and micromole-scale synthesis in flow. *Science (80-.).* **359**, 429–434 (2018).
41. Plehiers, P. P. *et al.* Artificial Intelligence for Computer-Aided Synthesis In Flow: Analysis and Selection of Reaction Components. *Front. Chem. Eng.* **2**, (2020).
42. Shields, B. J. *et al.* Bayesian reaction optimization as a tool for chemical synthesis. *Nature* **590**, 89–96 (2021).
43. Fu, Z. *et al.* Optimizing chemical reaction conditions using deep learning: a case study for the Suzuki–Miyaura cross-coupling reaction. *Org. Chem. Front.* **7**, 2269–2277 (2020).
44. 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).
45. 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).
46. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. (2012).

47. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. & Dahl, G. E. Neural Message Passing for Quantum Chemistry. (2017) doi:https://doi.org/10.48550/arXiv.1704.01212.
48. Stein, M. Large Sample Properties of Simulations Using Latin Hypercube Sampling. *Technometrics* **29**, 143 (1987).
49. Agarap, A. F. Deep Learning using Rectified Linear Units (ReLU). (2018).
50. Carpentier, A. *et al.* Upper-Confidence-Bound Algorithms for Active Learning in Multi-Armed Bandits. (2015).

Figures

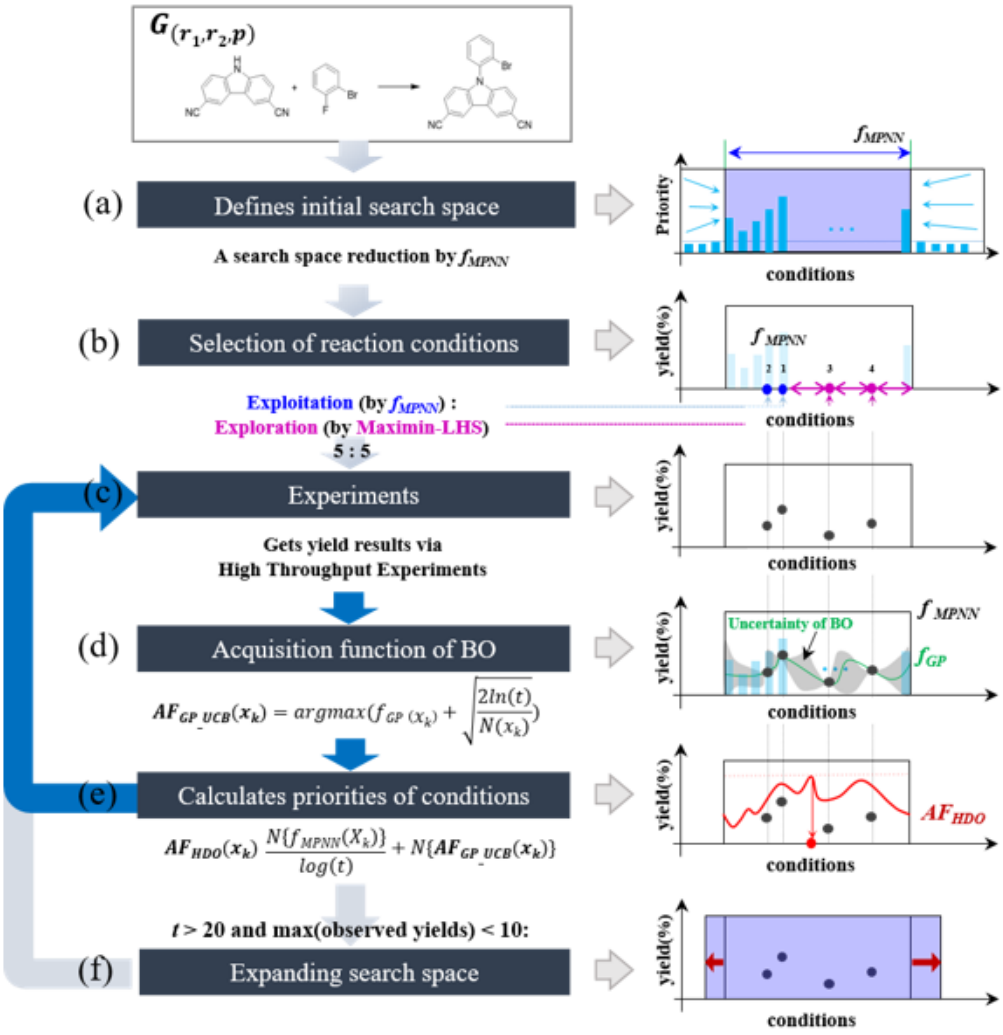

Figure 1

Process of reaction optimization by HDO. (a) Given a reaction representation in the form of a graph-type molecular structure, HDO specifies a search space using the best combination of conditions predicted by MPNN. (b) Initial experimental conditions are selected by adapting balanced methods that consider the trade-off between exploration and exploitation. (c) Reaction yields acquired via HTE to experiment with a selected combination of initial conditions. (d) The surrogate model (f_{GP}) of BO is trained using the initial experimental results (yield) and calculates the acquisition function (AF_{GP_UCB}) of BO. (e) The priority is calculated (AF_{HDO}) and the method determines whether to continue experimenting or to expand the search space. (f) If the number of experiments exceeds 20 trials and the maximum yield is less than 10, the initially narrowed range is expanded step-by-step (details in the text).

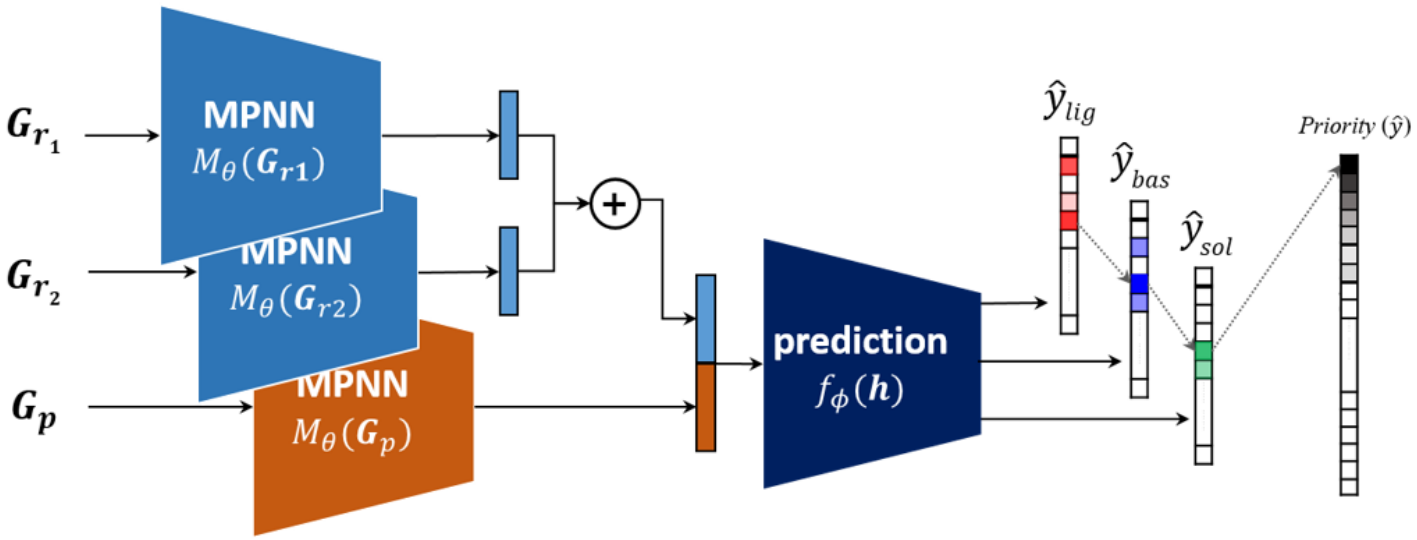

Figure 2

Illustration of process of MPNN models to predict suitable reaction conditions given graph-type reaction representations *G*.

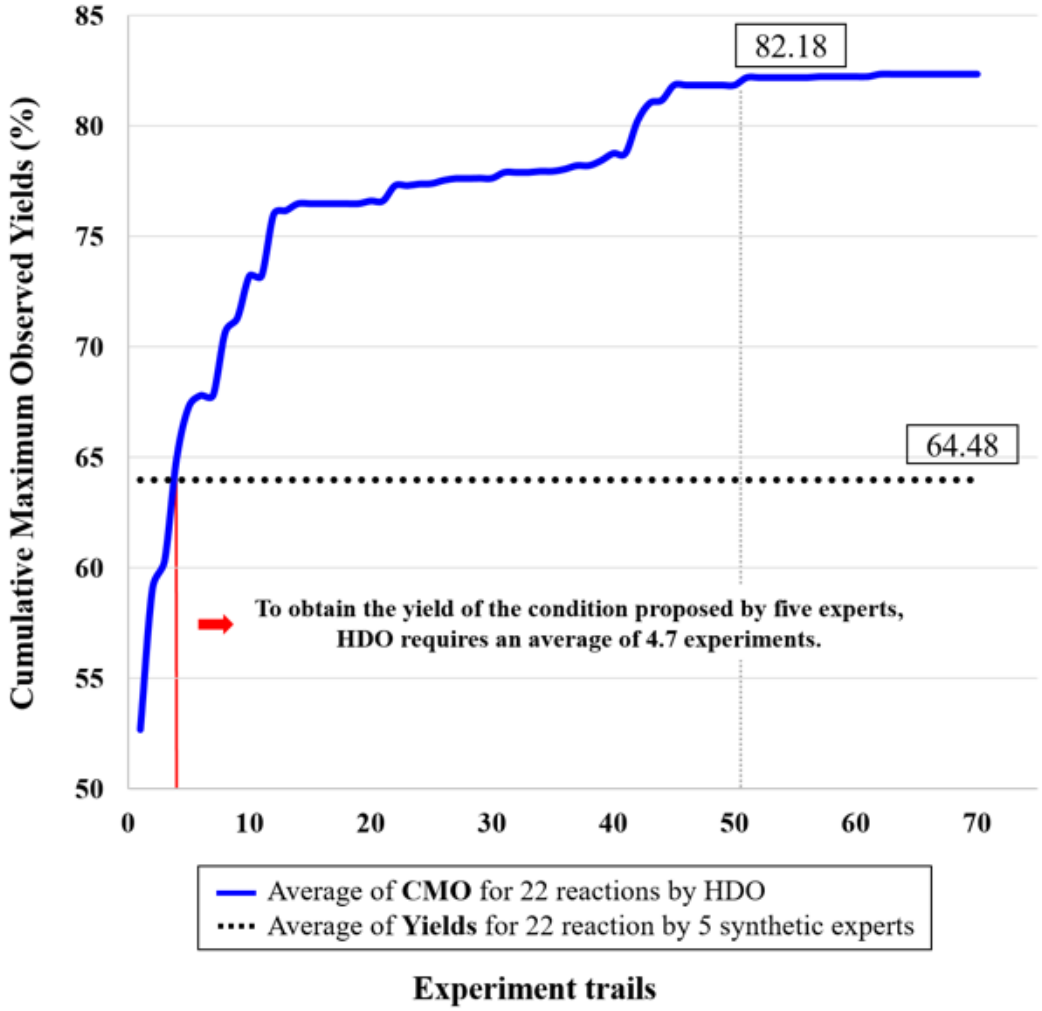

Figure 3

Average cumulative maximum observed yields by the HDO (blue curve) and the average yield of the combination of conditions proposed by five experts (black dotted line). HDO required an average of 4.7 experiments to find the conditions with the same average yield obtained with the combination of conditions proposed by the five experts for the 22 reactions.

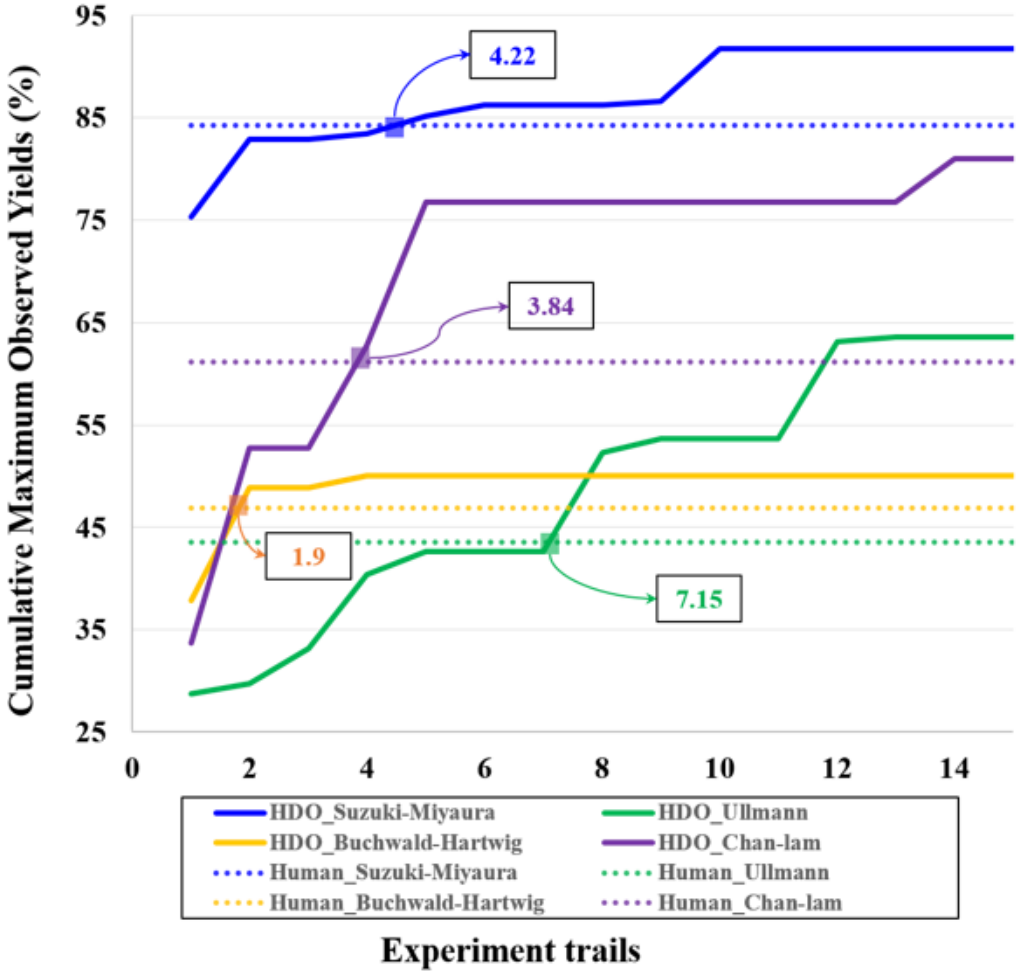

Figure 4

Performance of HDO with respect to the four named reactions. Comparison of the yield results of HDO with the conditions proposed by the experts for the Suzuki-Miyaura, Buchwald-Hartwig, Chan-Lam, and Ullmann reactions. On average, HDO identified suitable conditions after only 4.22, 1.9, 3.84, and 7.15 experimental trails, respectively.

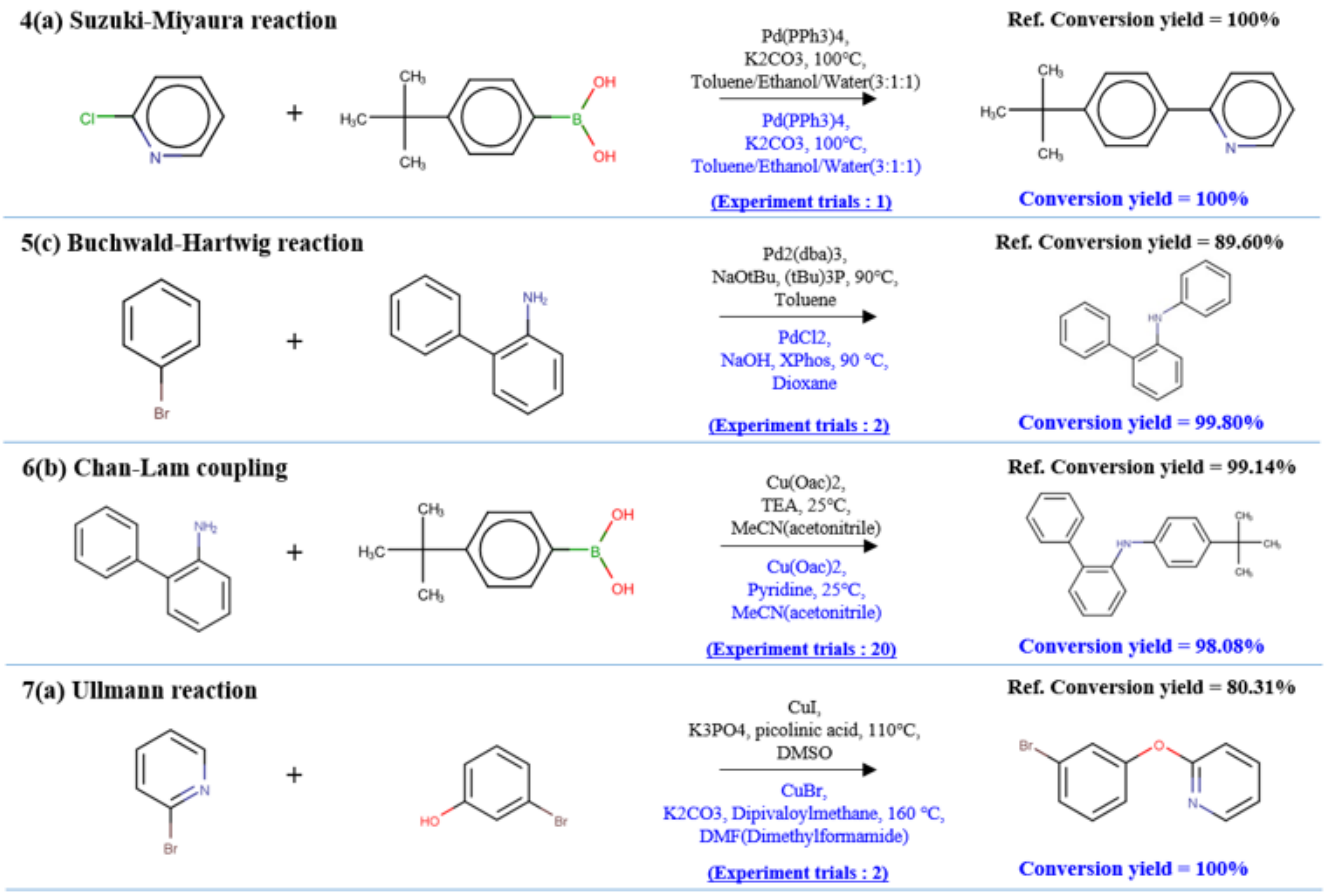

Figure 5

Examples of proposed reaction conditions details for the HDO with references.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- manuscriptHDOsupplementaryinformation0.1npjcommat.docx