

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

Youngchun Kwon, Dongseon Lee, Jin Woo Kim, Youn-Suk Choi,* and Sun Kim

Cite This: *ACS Omega* 2022, 7, 44939–44950

Read Online

ACCESS |



Metrics & More

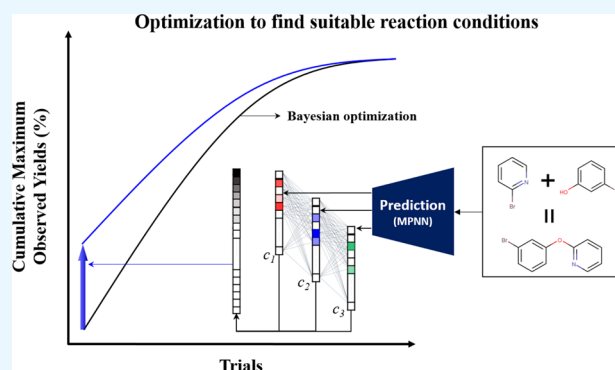


Article Recommendations



Supporting Information

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, the process to determine good reaction conditions is inevitable. Bayesian optimization (BO), an iterative optimization algorithm, demonstrates exceptional performance to identify reagents compared to synthesis experts. However, BO requires several initial randomly selected experimental results (yields) to train a surrogate model (approximately 10 experimental trials). Parts of this process, such as the cold-start problem in recommender systems, are inefficient. Here, we present an efficient optimization algorithm to determine suitable conditions based on BO that is guided by a graph neural network (GNN) trained on a million organic synthesis experiment data. The proposed method determined 8.0 and 8.7% faster high-yield reaction conditions than state-of-the-art algorithms and 50 human experts, respectively. In 22 additional optimization tests, the proposed method needed 4.7 trials on average to find conditions higher than the yield of the conditions recommended by five synthesis experts. The proposed method is considered in a situation of having a reaction dataset for training GNN.



INTRODUCTION

Substantial effort has been dedicated over the past few years to develop various technologies for optimizing 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.^{1–3} The optimization algorithms capable of efficiently finding local optima are gradient-based algorithms, such as gradient descent,⁴ conjugate gradient,⁵ or the more sophisticated Broyden–Fletcher–Goldfarb–Shanno algorithm (BFGS).⁶ Many optimization technologies have been specifically developed for chemistry. For example, chemical reaction conditions can be optimized using systematic methods such as the design of experiments (DOE).⁷ 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,⁸ pinpoint candidates for organic photovoltaics, predict organic reaction paths, and conduct automated experimentation^{9–16} 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 computa-

tion. Bayesian optimization (BO) approaches have emerged as popular optimization solutions to search for the efficient global optimum.^{17–22} BO schemes consist of two major steps: First, an approximation (surrogate model) to the objective function 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 predicts 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^{2,3} (RF), over Gaussian processes^{20,21} (GP), to active learning models.²⁴ However, these models require numerous evaluations (depending on the predefined search space) 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

Received: August 11, 2022

Accepted: November 18, 2022

Published: December 2, 2022



Table 1. Details of the Two Performance Benchmarking Tasks

		summary of task sets with definition of search space						
		task 1 (vs baselines, 50 humans)			task 2 (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)
no. target conditions	reactant 1	3						
	reactant 2	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		3696	792	1728	3456	4860	4095	84
>95 (yields)		1.92%	0.48%	0.58%				
no. target products		1	5	1	10	8	2	2
no. reactions to train GNN models		1,227,756	8541	49,625	158,605	17,705	10,518	2694

often costly, and material synthesis is another major barrier in material development because 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,²⁵ has led to major advances both in searching for possible retrosynthetic pathways^{26–37} and in evaluating the feasibility of the proposed reactions^{38–45} 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 (a lack of negative data).

In this study, our attempts to overcome the limited ability of advanced approaches to determine the optimal reaction conditions led us to propose the hybrid-type dynamic reaction optimization (HDO) method, which complements the previous two methodologies as data-driven approaches in that it is based on a graph neural network (GNN) with BO. This approach enables us to efficiently explore the optimal combination of conditions compared with previous studies. Modern advances in high-throughput experimentation (HTE)^{13,43} 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 of the capabilities of a collection consisting of a few thousand data points under a limited set of conditions in datasets from Shields's study.⁴⁶ In addition, we validated the proposed algorithm using additional reaction experiment ('Ullmann reaction' and 'Chan–Lam reaction') with five synthetic experts using our own HTE facilities, details of which are provided in Table 1.

■ PERFORMANCE BENCHMARKING RESULTS

The ultimate goal of this study is to rapidly determine suitable conditions, given reactions with predefined search space. Optimizing reaction conditions is the process of exploring various types of reaction parameters, such as reagent, solvent,

base, catalyst, concentration, and temperature. The number of their combinations could vary depending on the required parameter ranges. Owing to enormous cost and time, it is impossible to conduct all experiments for a combination of possible conditions. Therefore, an optimization process to define a reasonable scope of reaction conditions and to verify models that are rapidly navigable within that space is essential. In Shields's study,⁴⁶ the authors provide experimental yield results on all possible combinations of conditions in seven different reaction search spaces of three types of named reactions (Suzuki–Miyaura reaction, Buchwald–Hartwig reaction, and Arylation reaction). Thanks to reasonable search scope with experimental yield results from through HTE, our proposed method was verified through that optimization dataset (Task 1). Moreover, 22 additional experiments were conducted using our HTE equipment to check how rapidly the proposed algorithm finds optimal conditions compared with organic synthetic experts (Task 2). The details of task and search spaces for Tasks 1 and 2 are as follows:

Task 1: Entire optimization dataset including search spaces with yield results of reactions are from Shields's study. In addition, methods are also used as a baseline for verifying with previous works. Previous studies conducted high-throughput experiments on the class of Suzuki–Miyaura cross-coupling reactions. Twelve couplings of three electrophiles (Reactant 1) and four nucleophiles (Reactant 2) across the combinations of 11 ligands, seven bases, and four solvents were considered, thereby resulting in combinations for a total of 3696 reactions with a product. Buchwald–Hartwig reaction (2a–2e): They conducted high-throughput experiments on the class of Pd-catalyzed Buchwald–Hartwig C–N cross-coupling reactions. They experimented with combinations of three aryl halides, four catalysts, three bases, and 22 additives for a total of 792 reactions per target product, of which there were five reactions. Arylation (3a): They studied the arylation of imidazoles, a key step in the commercial synthesis of the JAK2 inhibitor BMS-911543.^{47,48} They selected a subspace consisting of 1,728 reactions including 12 ligands, four bases, four solvents, three temperatures, and three 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,

Table 2. Comparison of Reaction Optimization Performance with the Baselines on Task 1^a

reaction type	target	top- <i>k</i>	target yield	ANT ^{RS}					
				random search	RF	BO	MPNN	human (50 experts)	HDO (ours)
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

^aBold numbers indicate the highest-performing model for each reaction target.

who played the reaction optimization game. All of the reaction data and expert information in reaction 3a 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. 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. Further details of search scope with reaction structures are described in Supporting Information 1 and 2.

Task 1: Optimization of Reaction Conditions to Benchmark the Performance. We targeted a number of experimental trials (NT) to achieve the top 1, 5, and 10% yields in the entire search space of reaction conditions as a performance measurement.

$$\text{AIR}^c = 1/i \sum_{i=1}^k \frac{(\text{ANT}^{\text{RS}} - \text{NT}_i^c)}{\text{ANT}^{\text{RS}}} \quad (1)$$

We calculated the average improvement rate (AIR) over a ‘random search’ of NT^{RS} as the performance indicator to allow for comparison with other baselines (eq 1). Here, *i* is the number of iterations required for the overall optimization process, and we used 1000 iterations for all of the models. In each optimization model, NT_{*i*}^{*c*} is the NT in the *c* category of baseline models (Random Forest (RF),⁴⁹ BO,⁴⁶ humans,⁴⁶ and message passing neural network (MPNN),⁵⁰ and HDO). In Table 2, ANT^{RS} represents the average of NT results from *i* iterations of random search, and we determined *k* to be 1000 in this task. AIR^{*c*} represents the average improvement in the performance of baseline models (RF, BO, MPNN, and Human) and the proposed model (HDO) over the naïve

approach ANT^{RS}. AIR 0 indicates the same average number of trials as the random selection approaches. When AIR has a negative value, the number of trials to determine the combination of conditions with a target yield is more attempted than ANT^{RS}. Table 2 shows the top 1, 5, and 10% target yields in the entire search space and the AIR for the RF and BO models as reported by Shields et al. In message passing neural network (MPNN), the GNN-type-condition prediction models are trained separately for each condition type (e.g., solvent, catalyst, etc.), and the highest-rank conditions inferred for each condition type were combined and chosen as first combination for an experiment. Furthermore, the yield results were not reflected, and the following conditions were selected considering only the inferred priority using MPNN models.

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. HDO exhibited a stable optimization for the suitable combination of conditions in the case of the Suzuki–Miyaura reaction, which had abundant training data, including the Buchwald–Hartwig reaction, which had a small reaction dataset for training models. The basic strategy of HDO is designed to be efficiently explored by conducting experiments of conditional combinations initially recommended by the MPNN but quickly modifying the weight of the conditional combination chosen by BO if the yield results are not better than those expected. In the optimization task of Suzuki–Miyaura reaction 1(a), both HDO and MPNN trained approximately 1.2 million reaction data, finding optimal conditions very rapidly compared with others. However, in the case of Buchwald–Hartwig reaction 2(a–e), because of few training data, MPNN exhibited even less AIR than random search in a method recommended only by a combination of highly inferred conditions without optimization techniques. In case 2(a), BO achieved good performance and the distribution of yield values of all reactions within the search space was

Box plot for performance comparison between algorithms

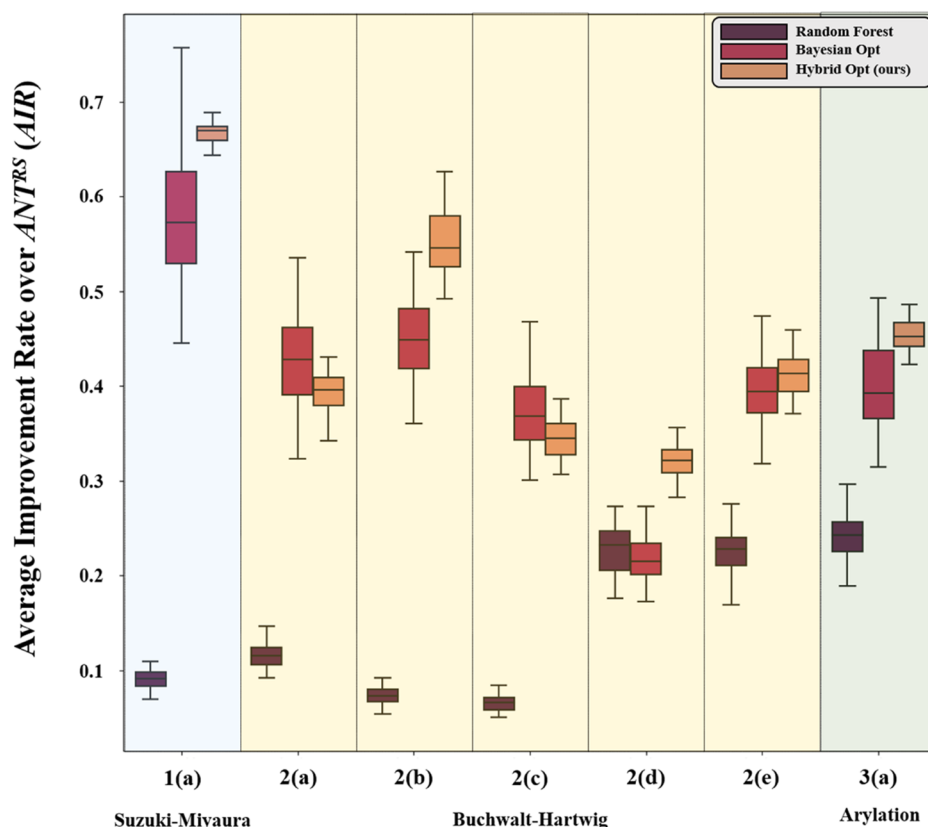


Figure 1. Box-and-whisker plot for comparison of performance (top 5%). The proposed HDO exhibited good performance in terms of median values excluding 2(a, c). In particular, in the Suzuki–Miyaura reaction, because the number of training data of the MPNN prediction model used for HDO was approximately 143 times more than Buchwalt–Hartwig reaction, the prediction performance is relatively high, and the optimal conditions are found rapidly in the beginning steps where the variance is small. Owing to the random selection of BO’s initial 10 experiments, the variance appears higher than that for the HDO for exploration.

relatively lower than others (see the HTE Results section in the Supporting Information and Figure 2). Reaxys database, which we used for training MPNN models, tends to be biased because the data are extracted from relatively successful research papers (without negative data). Therefore, MPNN models trained using the corresponding data do not yield good results for the optimization task with low yield distributions. These scores indicate the limitations of an optimization approach based only on exploitation. Similarly, the HDO model also performed well in the Suzuki–Miyaura reaction over BO, which 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 those of the human experts and BO although BO is the best for the top 1%. Overall, the proposed model HDO determined the best combination of conditions for a high yield compared to other optimization algorithms and human experts. It is necessary to determine a strategy that stably and efficiently searches for the conditions via repeated optimization experiments. Figure 1 shows the results of AIR^c, which were repeated 100 times, with a box-and-whisker plot to determine a condition combination with a yield value or more corresponding to the top 5% of the total search area. HDO stably determined suitable conditions in various optimization tests. In Suzuki–Miyaura reaction optimization test, where training data were abundant, the variance appears to be relatively small because the HDO rapidly determined the optimal conditions predicted by the MPNN initially embedded

in the HDO. This was additionally verified to efficiently determine the optimal conditions in the case of a relatively large number of training data in Task 2. In the Buchwalt–Hartwig case, which has little training reaction data, HDO stably exhibited best AIR^c results for the three reaction types (2b, c, and e) of the five reactions. The BO including the initial 10 random selections could be lucky to find the best conditions rapidly, but it occurs less frequently and has a high variance value. The HDO is designed for stable and efficient optimization by immediately increasing the priority weight of BO in HDO when unsatisfactory yields appear in initial experiments chosen by the MPNN (the weight conversion process is discussed in detail in the Methods Section).

Furthermore, HDO is designed to solve the cold-start problem of BO. Figure 2 shows the cumulative maximum observed (CMO) yield according to the number of experiment trials compared with BO and 50 experts in the Arylation reaction. The proposed HDO discovered high-yield conditions in the early stages compared to state-of-the-art algorithms and 50 synthetic experts (details of human information are described in Shields’s study⁴⁶). However, the CMO of HDO appears to have stagnated search performance in approximately 42 experiment trials, where conditional combinations with yields of nearly 98% or more were found.

Task 2: Validation of the HDO Compared to Five Human Chemists. In Task 2, the performance of the proposed HDO was compared with five organic synthetic

Comparison of the performance in Arylation reaction

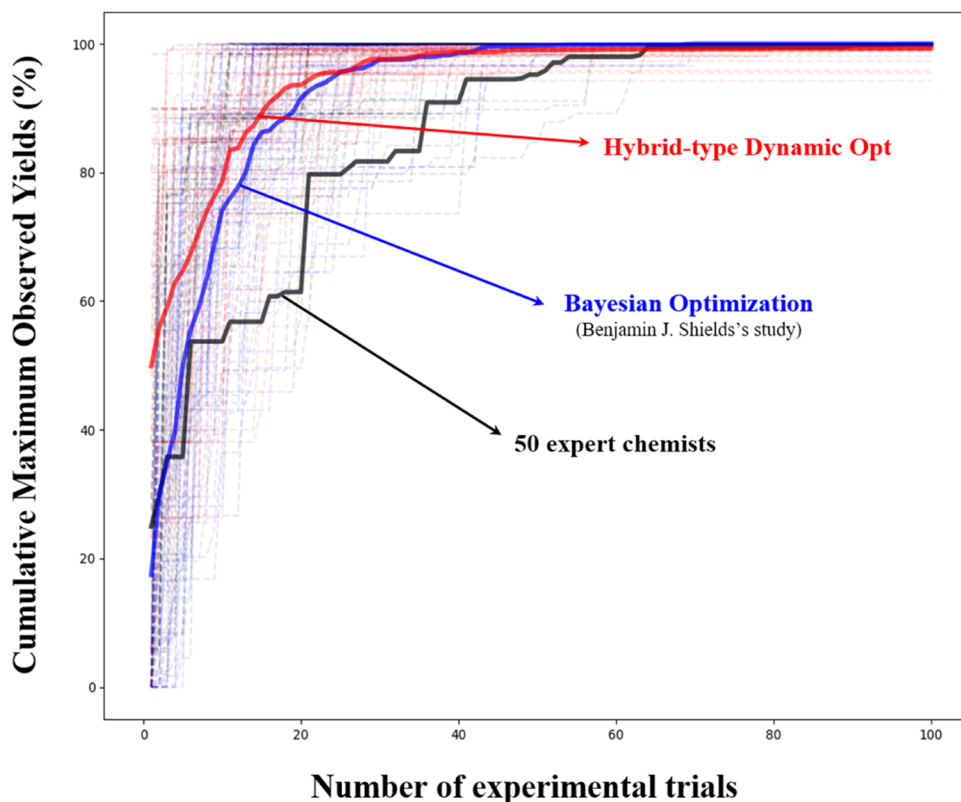


Figure 2. Performance comparison of the BO, 50 expert chemists, and HDO for the Arylation reaction. The results of 50 experiments with HDO and BO are represented by a dotted line, and the solid line represents their average values. Likewise, the different optimization results with 50 humans are represented by dotted lines, and their average values are represented by black solid lines.

chemists in 22 additional optimization experiments. The purpose of this task was to determine the number of trials required for the HDO to determine the condition combination recommended by a group of five chemists. The five experts for Task 2 had doctorates with more than 10 years of experience in organic synthesis.

The reaction experiments in Task 2 were 22 different named reactions of four types (the Suzuki–Miyaura, Buchwald–Hartwig, Chan–Lam, and Ullmann reactions). As described in Table 1, we prepared 20 reactions to optimize around reagents, with a combination of 3456–4860 conditions, and two experiments of Chan–Lam reaction with only 84 search spaces were also tested. All of the experiments were conducted using our own HTE system including autonomous robots for synthesizing and calculating conversion yield values using liquid chromatography equipment. First, five chemists were given a search space with reactants and product structures and the combination of conditions expected to have the highest yield was recommended. We conducted experiments using the proposed combinations of conditions and averaged the yield values of each result. The average yield for the 22 reactions calculated by the five chemists was 64.48, as described in Figure 3. Likewise, given reactants and product structures, HDO recommended the six highest-priority conditional combinations in the search space for the initial experiment. When the experimental results are output, the yield value is reflected in the HDO objective function and the next optimal condition combination is recommended. We conducted up to 70 experiments per reaction and stopped the experiment early

when conditions with a yield value of 95% or more were found. The average cumulative maximum observed yields of HDO are shown in Figure 3 with the yield recommended by the five chemists. On average, HDO determined the suitable conditions with the yield value of condition combinations recommended by the five experts in 4.7 experiments.

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 the 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 trials 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 in identifying reaction conditions with a high yield. Except for the reaction shown in Figure 5b, all of the reactions yielded poor results, yet, even in these difficult situations, HDO found expert-level yields after 1.9 trials 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. The corresponding details are provided in Figure 4.

Finally, Figure 5 shows the examples of each named reaction. In the Suzuki–Miyaura reaction (4(a) in Figure 5),

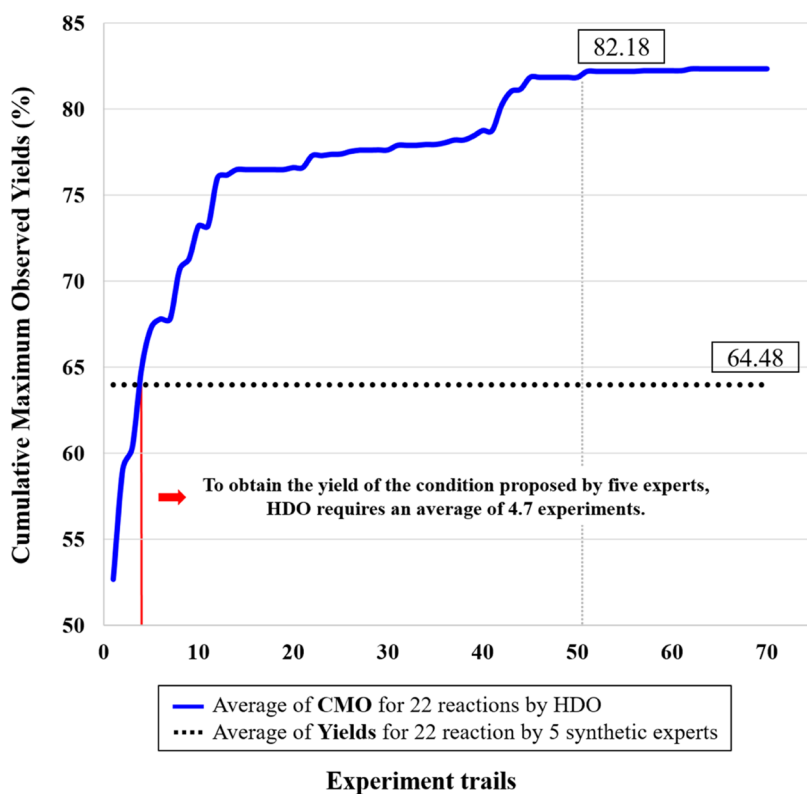


Figure 3. Average cumulative maximum observed yields using the HDO (blue curve) and the average yield of the combination of conditions proposed by five experts (black dotted line). HDO required an average of 4.7 experiments to determine the conditions with the same average yield obtained with the combination of conditions proposed by the five experts for the 22 reactions.

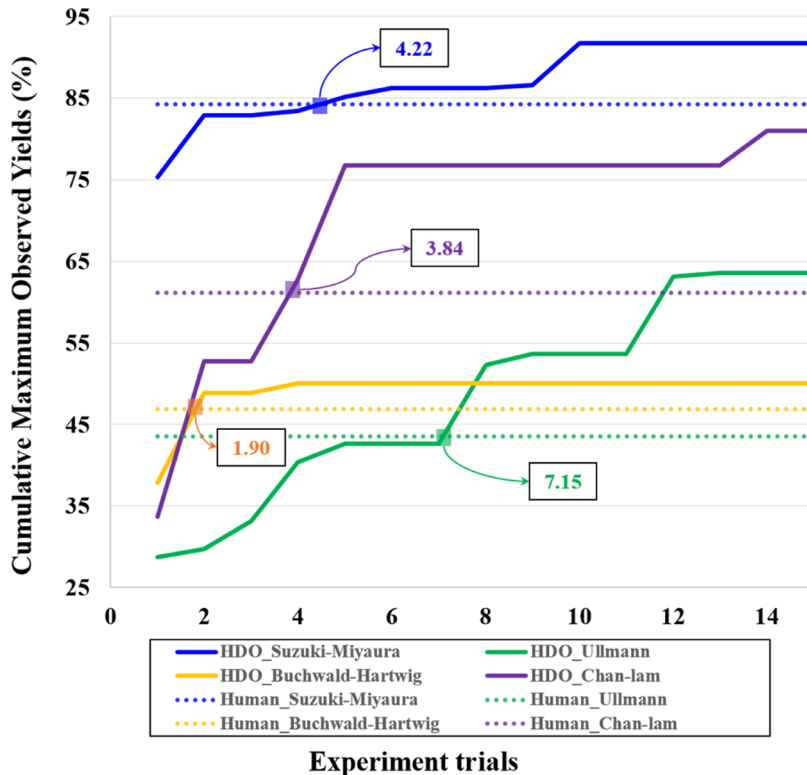


Figure 4. Performance of HDO for the four named reactions. Comparison of the yield results of HDO with the conditions proposed by the experts for the Suzuki–Miyaura, Buchwald–Hartwig, Chan–Lam, and Ullmann reactions. On average, HDO identified suitable conditions after only 4.22, 1.90, 3.84, and 7.15 experimental trials for these four reactions, respectively.

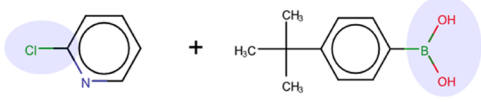

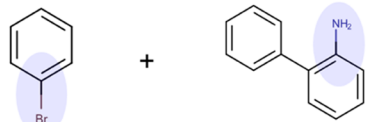
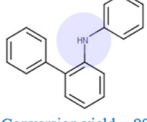
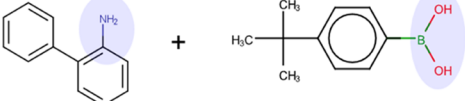
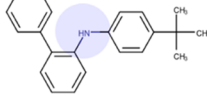
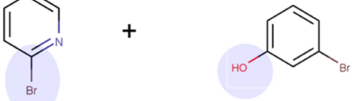
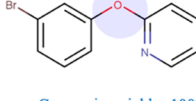
Named reaction/ Reactants	Ref. conditions/ HDO conditions (trials)	Product/ Yield
4(a) 'Suzuki Miyaura coupling' 	Ref. conditions: Pd(PPh ₃) ₄ , K ₂ CO ₃ , 100°C, Toluene/Ethanol/Water(3:1:1) HDO conditions (trials): Pd(PPh ₃) ₄ , K ₂ CO ₃ , 100°C, Toluene/Ethanol/Water(3:1:1) (Experiment trials : 2)	Conversion yield = 100%  Conversion yield = 100%
5(b) 'Buchwald-Hartwig amination reaction' 	Ref. conditions: Pd(dba) ₃ , NaOtBu, (tBu) ₃ P, 90°C, Toluene HDO conditions (trials): PdCl ₂ , NaOH, XPhos, 90 °C, Dioxane (Experiment trials : 1)	Conversion yield = 89.6%  Conversion yield = 99.8%
6(c) 'Chan-Lam reaction' 	Ref. conditions: Cu(Oac) ₂ , TEA, 25°C, MeCN(acetonitrile) HDO conditions (trials): Cu(Oac) ₂ , Pyridine, 25°C, MeCN(acetonitrile) (Experiment trials : 17)	Conversion yield = 99.14%  Conversion yield = 98.08%
7(d) 'Ullmann reaction' 	Ref. conditions: CuI, K ₃ PO ₄ , picolinic acid, 110°C, DMSO HDO conditions (trials): CuBr, K ₂ CO ₃ , Dipivaloylmethane, 160 °C, DMF(Dimethylformamide) (Experiment trials : 3)	Conversion yield = 80.31%  Conversion yield = 100%

Figure 5. Examples of the list of conditions recommended by five experts for four different types of named reactions proposed by the HDO in task 2. 4(a) is an example of the Suzuki–Miyaura coupling reaction. The HDO found the same conditions as five experts in two trials. 5(b) showed that in just one experiment, finding better lists of conditions than 5 experts in Buchwald–Hartwig amination reaction. In 6(c), similar performance condition lists to those of experts were found in 17 trials in the Chan–Lam reaction that lacked reaction data for the training model. On the other hand, in the 7(d) experiment of the Ullmann reaction, HDO found a combination of high-yield conditions compared to experts in three trials.

HDO determined the same combination of conditions with the same yield value as the experts and required a single experimental trial. In examples 5(b) and 7(a) in Figure 5 for the Buchwald–Hartwig and Ullmann reactions, respectively, HDO obtained higher yield results than the experts with different combinations of reaction conditions and required two experimental trials in both cases. These are the examples of the optimal condition combinations determined using the MPNN models for the initial five experiments. 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 trials to identify the 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 demonstrated reliable and efficient navigation capabilities for a variety of combinations of reaction conditions.

METHODS

Overview of HDO. For efficient exploration, HDO consists of an MPNN, which was trained using approximately 1 million experimental reaction data to predict suitable conditions, and a BO model, which explores conditions based on ongoing experimental results. We designed the optimization direction to be dynamically modified for two experimental results by adjusting the weights of the above two models based on the

obtained yield results. The overall condition optimization process is described below.

Considering that all of the reaction conditions are not efficient and could be unnecessarily costly, HDO narrows the search space using MPNN models that are able to predict the chemical context most suitable for any particular organic reaction. 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. Therefore, we adapted the candidate conditions predicted by the MPNN as exploitation and selected the Maximin–Latin Hypercube sampling method⁵¹ to ensure effective distribution of exploration (Figure 6b). Moreover, we experimented with initial conditions and trained the surrogate model of BO by obtaining the yield result (Figure 6c). For the acquisition function of BO, we adopted upper-confidence bounds (UCBs), which ranked the priority of the next combination of conditions (Figure 6d), 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, and BO (Figure 6e). Depending on the outcome, the search space could be expanded to include additional reaction conditions (Figure 6f). For maximum efficiency, HDO was designed to perform comprehensive judgments using not only the results predicted by the MPNN but also the experimental

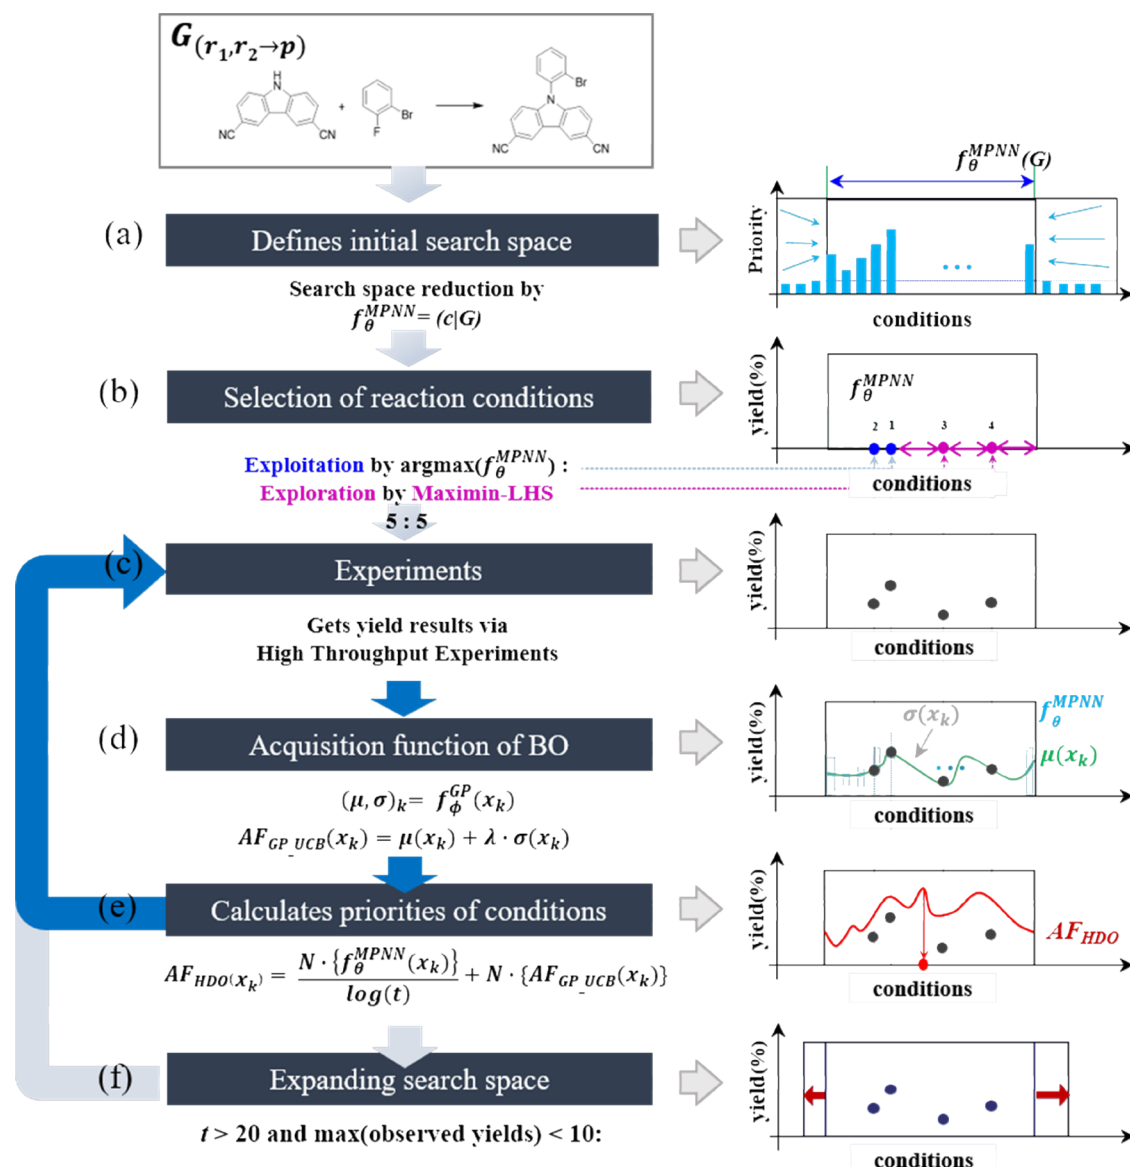


Figure 6. Process of reaction optimization using HDO. (a) Given a reaction representation, HDO specifies a search space using the best combination of conditions predicted by MPNN. (b) Initial experimental conditions are selected by adopting balanced methods that consider the trade-off between exploration and exploitation. (c) Reaction yields acquired via HTE to experiment with a selected combination of initial conditions. (d) Surrogate model (f_{θ}^{GP}) of BO trained using the initial experimental results (yield) that 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 and the maximum yield is less than 10, the initially narrowed range is expanded step by step (details in the text).

results, frequency of past experiments, and uncertainties in the objective function of the predictive model.

The data preprocessing and model formulation steps are detailed in subsequent subsections. HDO proceeds iteratively until it determines the global optimal combination of conditions that produce the desired target yield, and it 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 conducted using robots and management software.

Dataset and Graph-Type Representation for Training MPNN. The dataset for all of the reactions with their conditions for training the MPNN was extracted from the Reaxys reaction database consisting of 53 million reaction

records. The data include structural expressions of the reactants, products, and conditions. 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, c) , where $R = \{G_{r1}, G_{r2}\}_i$ and $P = \{G_p\}_i$ represent the set i of the two reactants with the product structures in the reaction, respectively, and c 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 predefined list of reagents in the experiments. In addition, the number of reaction sets i and




Figure 7. Illustration of the process of MPNN models to predict suitable reaction conditions given graph-type reaction representations G (\oplus denotes the summation of q^R vectors, and \otimes represents the calculation of combinations among weights of one-hot vectors under different conditions c).

classes of conditions can be different for each type of condition and optimization task, respectively. Furthermore, we restricted 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 noted the ambiguous labeling of certain catalysts, solvents, and reagents in the Reaxys, 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 completely eliminated because a strict separation between reagents and catalysts is difficult to achieve.

Message Passing Neural Networks for Predicting Suitable Reagents. A chemical reaction consisting of two reactants was set as $R_{1,2}$ and a single product as P . This chemical reaction is labeled with its reaction conditions c , denoted by $f_{\theta}^{\text{MPNN}}(c | R_{1,2}, P)$. A graph is a data structure that presents a powerful non-Euclidean method for establishing the extent to which features (nodes) are connected to their relationships (edges). We defined each molecule in $R_{1,2}$ and P as an undirected graph $G = (V, E)$, where V and E represent the sets of nodes and edges vectors, respectively. The node feature vector $v^j \in V$ and edge feature vectors $e^{j,k} \in E$ were assumed to consist of heavy atoms (e.g., C, N, O, and F) and their bonds (e.g., single, double, triple, and aromatic), respectively. Hydrogen atoms were not considered. The node feature v^j is a vector indicating the atom type, formal charge, degree, hybridization, number of hydrogen atoms, valency, chirality, whether it accepts or donates electrons, whether it is aromatic, whether it is in a ring, and associated ring sizes. For the bond between atoms j and k , $e^{j,k}$ is a vector indicating the bond type, stereochemistry, whether it is in a ring, and whether it is conjugated. The MPNN is designed to accept $G = (V, E)$ as the input and to return the graph representation vector q as output as follows

$$q = M_{\theta}(G) \quad (2)$$

The MPNN uses six message passing steps with an edge network as the message function and a gated recurrent unit (GRU) network as the update function to produce node-representation vectors, whose dimensionality was 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, as described in Figure 7. The embedding is further processed by a fully connected layer of 512 ReLUs,⁵² 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 to the order of reactants and products. The two summated vectors of reactants are concatenated to produce a reaction representation vector h (3), as follows

$$h = \left[\sum_{l=1}^m q^{R,l}, q^P \right] \quad (3)$$

Here, m is the number of reactants. 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 is represented as $p_{\phi}(h)$. The output functions of the FNN are equal to the length of the one-hot vectors by each condition type (c). The predictive model MPNN is trained as an independent model according to the type of reaction conditions and the predicted catalyst, ligand, base, and solvent as c_{category} , respectively. Once trained, the prediction model f_{θ}^{MPNN} is used for predicting the conditions of new chemical reactions. Given a query reaction (R^*, P^*), we predict conditions vector c as $c_{\text{solvent}}, c_{\text{catalyst}}, c_{\text{base}}$, and c_{etc} (the area category of condition types is dependent on the reaction task). Then, the value of the priority is calculated by generating all possible combinations of conditions using the weight value of the one-hot vector c predicted by the model for each of the conditions. Therefore, the priority of condition combination I is defined as

$$I = \sum_{k=0}^t N \cdot (f_{\theta}^{\text{MPNN}}(R, P)) \quad (4)$$

The types of conditions vary for different experiments and have different ranges. In eq 4, t represents the type of condition. Each condition prediction model f_{θ}^{MPNN} is normalized for reducing gab of weight values. Finally, the priority of combination of predicted conditions is represented by I . Moreover, the priority is represented as $I = N \cdot C_{\text{lig}} + N \cdot C_{\text{bas}} + N \cdot C_{\text{sol}}$, where N is the normalization to reflect the equivalent weights for each category of conditions. Figure 7 illustrates the training process for the condition prediction models. The details of the performances of MPNN models are described in “Supporting Information S. Note 5: Top- k accuracy of the MPNNs for predicting conditions.”

Bayesian Optimization in HDO. The BO method is used to reduce the number of objective evaluations that need to be performed to solve an optimization problem. To achieve this, they iteratively suggest, in a careful and intelligent manner, an

input location in which the objective that is being optimized should be evaluated for each experiment. At each iteration $N = 1, 2, 3 \dots$ of the optimization process, the BO method fits a probabilistic model, a Gaussian process surrogate model (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. We modeled chemical reaction outcomes; GP is defined in Shield's study,⁴⁶ and the Matérn52 kernel was used and represented as f_{ϕ}^{GP} . For each reaction, a numerical encoding was generated by concatenating descriptor vectors for each condition category and continuous variable. For example, the Suzuki–Miyaura reaction in Task 2 involved four categories: catalysts, bases, solvents, and ligands. The input description is described as $\mathbf{d}_i = \mathbf{d}_{\text{catalyst}} \oplus \mathbf{d}_{\text{base}} \oplus \mathbf{d}_{\text{solvent}} \oplus \mathbf{d}_{\text{ligand}}$ (where \oplus denotes concatenation). Here, i is the number of all possible combinations of conditions, that is, the search space. The target was the yield. The target of BO objective function is the experimental yield results. In Gaussian process regression, the surrogate model determines the general shape of its function distribution. The trained parameters for the length scale set the relative variation per dimension, the amplitude calibrates the magnitude of the changes, and the noise captures the variation in measurements. In the acquisition function, upper-confidence-bound (UCB)⁵³ algorithm was adopted to calculate the priority of the next best combination of conditions as (5). Uncertainty was used to generate an acquisition function $\text{AF}_{\text{GP_UCB}}$, whose value at each input location indicates the expected utility of evaluating f_{ϕ}^{GP} at this location. The upper-confidence-bound (UCB)⁵³ algorithm was adopted to calculate the priority of the next best combinations of conditions as

$$\text{AF}_{\text{GP_UCB}}(x_k; \lambda) = \mu(x_k) + \lambda \cdot \sigma(x_k) \quad (5)$$

With UCB, the exploitation vs exploration trade-off is straightforward and easy to tune via the parameter λ . Concretely, UCB is a weighted sum of the expected performance captured by $\mu(x_k)$ of the Matérn52 kernels and of the uncertainty $\sigma(x_k)$, captured by the standard deviation of the f_{ϕ}^{GP} . The next point x_k at which f_{ϕ}^{GP} is evaluated is the one that maximizes $\text{AF}_{\text{GP_UCB}}$. After collecting this observation, the process is repeated. When sufficient data are collected, the GP predictive mean value f_{ϕ}^{GP} can be optimized to find the solution of the problem. Considering that the acquisition function $\text{AF}_{\text{GP_UCB}}$ is calculated using only the results currently being experimented with, it is possible to correct the process of navigating in the wrong direction due to the inaccurate prediction results of the MPNN model.

Acquisition Function for 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 \mathbf{x}_k based on the belief about f_{θ}^{MPNN} and $\text{AF}_{\text{GP_UCB}}$, given the history \mathbf{H}_k , a sampling strategy is defined as follows

$$\text{AF}_{\text{HDO}}(\mathbf{x}_k) = \frac{N \cdot \{f_{\theta}^{\text{MPNN}}(\mathbf{x}_k)\}}{\ln(t)} + N \cdot \{\text{AF}_{\text{GP_UCB}}(\mathbf{x}_k)\} \quad (6)$$

In eq 6, AF_{HDO} was designed for efficient optimization combining the current experiment-driven priority model $\text{AF}_{\text{GP_UCB}}$ with the priority of f_{θ}^{MPNN} , obtained 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 (7) using min–max feature scaling.

$$N = \frac{X - X_{\min}}{X_{\max} - X_{\min}} \quad (7)$$

Under the influence of the weight of $\text{AF}_{\text{GP_UCB}}$, 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% after every five experiments.

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

RESULTS AND DISCUSSION

Searching for optimal high-yield reaction conditions requires a lot of resources and is time-consuming. In this paper, we proposed a method for efficiently exploring suitable synthetic reaction conditions, given reaction structures (reactants and target a product) with a condition search space. As a baseline, BO is a progressive solution to determine suitable conditions, but initial experimental results are inevitably required for training surrogate models to infer conditions. This causes a cold-start problem. Moreover, data-driven approaches such as GNNs can determine optimal combinations of conditions in the early stages of the optimization if they have sufficient chemical reaction data, such as that of the Suzuki–Miyaura reaction. However, due to the lack of training data and novel reactions, the predictive performance of the GNN model can be poor.

Therefore, we designed the hybrid-type dynamic optimization (HDO) method to compensate for the above-mentioned shortcomings while utilizing the advantages of the two approaches. Given reaction structures, GNN models based on MPNN predict appropriate reaction conditions. We utilized priorities of combinations of predictive conditions using a GNN for initial experiments. The experimental results (yield) were used for BO surrogate model training to select the next combination of conditions, and the optimization direction was dynamically modified based on the number of trials and observed yields. This approach enables an intuitive sampling policy to efficiently accomplish global optimization.

As a result, in experimental simulations, HDO could determine the optimal conditions that satisfied the target yield faster than other baselines. To further investigate the performance of HDO, we additionally prepared condition optimization tasks for synthesizing 22 target products and confirmed the number of experiments for HDO to attain the level of five specialists in organic synthesis. The HDO approach also met the target yield by swiftly identifying a combination of reaction conditions that were 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 promising chemical species and optimizing the structures of materials for various applications in the field of materials discovery.

■ ASSOCIATED CONTENT

Data Availability Statement

Processed reaction outcome data for Task 1 is available at <https://github.com/b-shields/edbo>, and the optimization details of Task 2 with search space are presented in the [Supporting Information](#). The additional data that support the findings of this study are available from the corresponding author upon reasonable request.

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acsomega.2c05165>.

Benchmarked reaction details (Figures S1–S3 and Tables S1–S7); optimization details in task 1 (Figures S4 and S5 and Table S8); optimization details in task 2 (Figure S6); distribution of reaction conditions in Reaxys database (Figure S7); performance of MPNN condition prediction models (Tables S9 and S10); and Suzuki–Miyaura coupling reaction optimization processes in task 2 (Figure S8) ([PDF](#))

■ AUTHOR INFORMATION

Corresponding Author

Youn-Suk Choi – *Autonomous Materials Development LAB, Samsung Advanced Institute of Technology, Suwon, Gyeonggi-do 16678, South Korea*; orcid.org/0000-0001-7119-8788; Email: ysuk.choi@samsung.com

Authors

Youngchun Kwon – *Department of Computer Science and Engineering, Seoul National University, Seoul 08826, South Korea*; *Autonomous Materials Development LAB, Samsung Advanced Institute of Technology, Suwon, Gyeonggi-do 16678, South Korea*; orcid.org/0000-0001-7911-3670

Dongseon Lee – *Autonomous Materials Development LAB, Samsung Advanced Institute of Technology, Suwon, Gyeonggi-do 16678, South Korea*

Jim Woo Kim – *Autonomous Materials Development LAB, Samsung Advanced Institute of Technology, Suwon, Gyeonggi-do 16678, South Korea*

Sun Kim – *Department of Computer Science and Engineering, Seoul National University, Seoul 08826, South Korea*

Complete contact information is available at:

<https://pubs.acs.org/doi/10.1021/acsomega.2c05165>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

This work was supported by Samsung Advanced Institute of Technology.

■ REFERENCES

- (1) Biegler, L. T.; Grossmann, I. E. Retrospective on Optimization. *Comput. Chem. Eng.* **2004**, *28*, 1169–1192.
- (2) Yin, X.; Gounaris, C. E. Search Methods for Inorganic Materials Crystal Structure Prediction. *Curr. Opin. Chem. Eng.* **2022**, *35*, No. 100726.
- (3) Wang, K.; Dowling, A. W. Bayesian Optimization for Chemical Products and Functional Materials. *Curr. Opin. Chem. Eng.* **2022**, *36*, No. 100728.
- (4) Ruder, S. An Overview of Gradient Descent Optimization Algorithms, 2016. arXiv:1609.04747. arXiv.org e-Print archive. <https://arxiv.org/abs/1609.04747>.
- (5) Hestenes, M. R.; Stiefel, E. Methods of Conjugate Gradients for Solving Linear Systems. *J. Res. Natl. Bur. Stand.* **1952**, *49*, No. 409.
- (6) Fletcher, R. *Practical Methods of Optimization*, 2nd ed.; John Wiley & Sons: New York, NY, USA, 1987.
- (7) I, J. O. The Design of Experiments. *Nature* **1936**, *137*, 252–254.
- (8) Negoescu, D. M.; Frazier, P. I.; Powell, W. B. The Knowledge-Gradient Algorithm for Sequencing Experiments in Drug Discovery. *INFORMS J. Comput.* **2011**, *23*, 346–363.
- (9) Nikolaev, P.; Hooper, D.; Perea-López, N.; Terrones, M.; Maruyama, B. Discovery of Wall-Selective Carbon Nanotube Growth Conditions via Automated Experimentation. *ACS Nano* **2014**, *8*, 10214–10222.
- (10) 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.* **2016**, *20*, 386–394.
- (11) Nikolaev, P.; Hooper, D.; Webber, F.; Rao, R.; Decker, K.; Krein, M.; Poleski, J.; Barto, R.; Maruyama, B. Autonomy in Materials Research: A Case Study in Carbon Nanotube Growth. *npj Comput. Mater.* **2016**, *2*, No. 16031.
- (12) Duros, V.; Grizou, J.; Xuan, W.; Hosni, Z.; Long, D.-L.; Miras, H. N.; Cronin, L. Human versus Robots in the Discovery and Crystallization of Gigantic Polyoxometalates. *Angew. Chem., Int. Ed.* **2017**, *56*, 10815–10820.
- (13) Coley, C. W.; Thomas, D. A.; Lummiss, J. A. M.; Jaworski, J. N.; Breen, C. P.; Schultz, V.; Hart, T.; Fishman, J. S.; Rogers, L.; Gao, H.; Hicklin, R. W.; Plehiers, P. P.; Byington, J.; Piotti, J. S.; Green, W. H.; Hart, A. J.; Jamison, T. F.; Jensen, K. F. A Robotic Platform for Flow Synthesis of Organic Compounds Informed by AI Planning. *Science* **2019**, *365*, No. eaax1566.
- (14) Gao, W.; Raghavan, P.; Coley, C. W. Autonomous Platforms for Data-Driven Organic Synthesis. *Nat. Commun.* **2022**, *13*, No. 1075.
- (15) Pyzer-Knapp, E. O.; Pitera, J. W.; Staar, P. W. J.; Takeda, S.; Laino, T.; Sanders, D. P.; Sexton, J.; Smith, J. R.; Curioni, A. Accelerating Materials Discovery Using Artificial Intelligence, High Performance Computing and Robotics. *npj Comput. Mater.* **2022**, *8*, No. 84.
- (16) Huo, H.; Rong, Z.; Kononova, O.; Sun, W.; Botari, T.; He, T.; Tshitoyan, V.; Ceder, G. Semi-Supervised Machine-Learning Classification of Materials Synthesis Procedures. *npj Comput. Mater.* **2019**, *5*, No. 62.
- (17) Kushner, H. J. A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *J. Basic Eng.* **1964**, *86*, 97–106.
- (18) Moćkus, J. On Bayesian Methods for Seeking the Extremum. In *Optimization Techniques IFIP Technical Conference*; Springer, 1975; pp 400–404.
- (19) Mockus, J. B.; Mockus, L. J. Bayesian Approach to Global Optimization and Application to Multiobjective and Constrained Problems. *J. Optim. Theory Appl.* **1991**, *70*, 157–172.
- (20) Snoek, J.; Swersky, K.; Zemel, R. S.; Adams, R. P. Input Warping for Bayesian Optimization of Non-Stationary Functions, 2014. arXiv:1402.0929. arXiv.org e-Print archive. <https://arxiv.org/abs/1402.0929>.
- (21) Snoek, J.; Larochelle, H.; Adams, R. P. Practical Bayesian Optimization of Machine Learning Algorithms, 2012. arXiv:1206.2944. arXiv.org e-Print archive. <https://arxiv.org/abs/1206.2944>.
- (22) 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* **2012**, *58*, 3250–3265.
- (23) Hutter, F.; Hoos, H. H.; Leyton-Brown, K. Sequential Model-Based Optimization for General Algorithm Configuration. In *Learning and Intelligent Optimization*; Springer, 2011; pp 507–523.

- (24) Settles, B. *Synthesis Lectures on Artificial Intelligence and Machine Learning*. In *Active Learning*; Springer, 2012; Vol. 6, pp 1–114.
- (25) Goodman, J. Computer Software Review: Reaxys. *J. Chem. Inf. Model.* **2009**, *49*, 2897–2898.
- (26) Szymkuć, S.; Gajewska, E. P.; Klucznik, T.; Molga, K.; Dittwald, P.; Startek, M.; Bajczyk, M.; Grzybowski, B. A. Computer-Assisted Synthetic Planning: The End of the Beginning. *Angew. Chem., Int. Ed.* **2016**, *55*, 5904–5937.
- (27) Segler, M. H. S.; Preuss, M.; Waller, M. P. Learning to Plan Chemical Syntheses, 2017. arXiv:1708.04202. arXiv.org e-Print archive. <https://arxiv.org/abs/1708.04202>.
- (28) Law, J.; Zsoldos, Z.; Simon, A.; Reid, D.; Liu, Y.; Khew, S. Y.; Johnson, A. P.; Major, S.; Wade, R. A.; Ando, H. Y. Route Designer: A Retrosynthetic Analysis Tool Utilizing Automated Retrosynthetic Rule Generation. *J. Chem. Inf. Model.* **2009**, *49*, 593–602.
- (29) Liu, B.; Ramsundar, B.; Kawthekar, P.; Shi, J.; Gomes, J.; Nguyen, Q. L.; Ho, S.; Sloane, J.; Wender, P.; Pande, V. Retrosynthetic Reaction Prediction Using Neural Sequence-to-Sequence Models. *ACS Cent. Sci.* **2017**, *3*, 1103–1113.
- (30) Bøgevig, A.; Federsel, H.-J.; Huerta, F.; Hutchings, M. G.; Kraut, H.; Langer, T.; Löw, P.; Oppawsky, C.; Rein, T.; Saller, H. Route Design in the 21st Century: The IC SYNTH Software Tool as an Idea Generator for Synthesis Prediction. *Org. Process Res. Dev.* **2015**, *19*, 357–368.
- (31) Coley, C. W.; Rogers, L.; Green, W. H.; Jensen, K. F. Computer-Assisted Retrosynthesis Based on Molecular Similarity. *ACS Cent. Sci.* **2017**, *3*, 1237–1245.
- (32) Segler, M. H. S.; Preuss, M.; Waller, M. P. Planning Chemical Syntheses with Deep Neural Networks and Symbolic AI. *Nature* **2018**, *555*, 604–610.
- (33) 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.* **2021**, *61*, 123–133.
- (34) 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.* **2020**, *11*, No. 5575.
- (35) Ucak, U. V.; Ashyrmamatov, I.; Ko, J.; Lee, J. Retrosynthetic Reaction Pathway Prediction through Neural Machine Translation of Atomic Environments. *Nat. Commun.* **2022**, *13*, No. 1186.
- (36) Mo, Y.; Guan, Y.; Verma, P.; Guo, J.; Fortunato, M. E.; Lu, Z.; Coley, C. W.; Jensen, K. F. Evaluating and Clustering Retrosynthesis Pathways with Learned Strategy. *Chem. Sci.* **2021**, *12*, 1469–1478.
- (37) Wang, X.; Qian, Y.; Gao, H.; Coley, C. W.; Mo, Y.; Barzilay, R.; Jensen, K. F. Towards Efficient Discovery of Green Synthetic Pathways with Monte Carlo Tree Search and Reinforcement Learning. *Chem. Sci.* **2020**, *11*, 10959–10972.
- (38) Kayala, M. A.; Azencott, C.-A.; Chen, J. H.; Baldi, P. Learning to Predict Chemical Reactions. *J. Chem. Inf. Model.* **2011**, *51*, 2209–2222.
- (39) Kayala, M. A.; Baldi, P. ReactionPredictor: Prediction of Complex Chemical Reactions at the Mechanistic Level Using Machine Learning. *J. Chem. Inf. Model.* **2012**, *52*, 2526–2540.
- (40) Jin, W.; Coley, C. W.; Barzilay, R.; Jaakkola, T. Predicting Organic Reaction Outcomes with Weisfeiler-Lehman Network, 2017. arXiv:1709.04555. arXiv.org e-Print archive. <https://arxiv.org/abs/1709.04555>.
- (41) 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.* **2017**, *3*, 434–443.
- (42) Maser, M. R.; Cui, A. Y.; Ryou, S.; DeLano, T. J.; Yue, Y.; Reisman, S. E. Multilabel Classification Models for the Prediction of Cross-Coupling Reaction Conditions. *J. Chem. Inf. Model.* **2021**, *61*, 156–166.
- (43) Gao, H.; Struble, T. J.; Coley, C. W.; Wang, Y.; Green, W. H.; Jensen, K. F. Using Machine Learning To Predict Suitable Conditions for Organic Reactions. *ACS Cent. Sci.* **2018**, *4*, 1465–1476.
- (44) Kwon, Y.; Lee, D.; Choi, Y.-S.; Kang, S. Uncertainty-Aware Prediction of Chemical Reaction Yields with Graph Neural Networks. *J. Cheminform.* **2022**, *14*, No. 2.
- (45) Coley, C. W.; Jin, W.; Rogers, L.; Jamison, T. F.; Jaakkola, T. S.; Green, W. H.; Barzilay, R.; Jensen, K. F. A Graph-Convolutional Neural Network Model for the Prediction of Chemical Reactivity. *Chem. Sci.* **2019**, *10*, 370–377.
- (46) Shields, B. J.; Stevens, J.; Li, J.; Parasram, M.; Damani, F.; Alvarado, J. I. M.; Janey, J. M.; Adams, R. P.; Doyle, A. G. Bayesian Reaction Optimization as a Tool for Chemical Synthesis. *Nature* **2021**, *590*, 89–96.
- (47) Fox, R. J.; Cuniere, N. L.; Bakrania, L.; Wei, C.; Strotman, N. A.; Hay, M.; Fanfair, D.; Regens, C.; Beutner, G. L.; Lawler, M.; Lobben, P.; Soumeillant, M. C.; Cohen, B.; Zhu, K.; Skliar, D.; Rosner, T.; Markwalter, C. E.; Hsiao, Y.; Tran, K.; Eastgate, M. D. C–H Arylation in the Formation of a Complex Pyrrolopyridine, the Commercial Synthesis of the Potent JAK2 Inhibitor, BMS-911543. *J. Org. Chem.* **2019**, *84*, 4661–4669.
- (48) Ji, Y.; Plata, R. E.; Regens, C. S.; Hay, M.; Schmidt, M.; Razler, T.; Qiu, Y.; Geng, P.; Hsiao, Y.; Rosner, T.; Eastgate, M. D.; Blackmond, D. G. 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.* **2015**, *137*, 13272–13281.
- (49) Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Müller, A.; Nothman, J.; Louppe, G.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; Duchesnay, E. Scikit-Learn: Machine Learning in Python, 2012. arXiv:1201.0490. arXiv.org e-Print archive. <https://arxiv.org/abs/1201.0490>.
- (50) Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; Dahl, G. E. Neural Message Passing for Quantum Chemistry, 2017. arXiv:1704.01212. arXiv.org e-Print archive. <https://arxiv.org/abs/1704.01212>.
- (51) Stein, M. Large Sample Properties of Simulations Using Latin Hypercube Sampling. *Technometrics* **1987**, *29*, 143–151.
- (52) Agarap, A. F. Deep Learning Using Rectified Linear Units (ReLU), 2018. arXiv:1803.08375. arXiv.org e-Print archive. <https://arxiv.org/abs/1803.08375>.
- (53) Carpentier, A.; Lazaric, A.; Ghavamzadeh, M.; Munos, R.; Auer, P.; Antos, A. Upper-Confidence-Bound Algorithms for Active Learning in Multi-Armed Bandits, 2015. arXiv:1507.04523. arXiv.org e-Print archive. <https://arxiv.org/abs/1507.04523>.