A Systematic Comparison Study on Hyperparameter Optimisation of Graph Neural Networks for Molecular Property Prediction

Yingfang Yuan
School of Mathematical and
Computer Sciences
Heriot-Watt University
Edinburgh, United Kingdom
yyy2@hw.ac.uk

Wenjun Wang
School of Mathematical and
Computer Sciences
Heriot-Watt University
Edinburgh, United Kingdom
wenjun.wang@hw.ac.uk

Wei Pang*
School of Mathematical and
Computer Sciences
Heriot-Watt University
Edinburgh, United Kingdom
w.pang@hw.ac.uk

ABSTRACT

Graph neural networks (GNNs) have been proposed for a wide range of graph-related learning tasks. In particular, in recent years, an increasing number of GNN systems were applied to predict molecular properties. However, a direct impediment is to select appropriate hyperparameters to achieve satisfactory performance with lower computational cost. Meanwhile, many molecular datasets are far smaller than many other datasets in typical deep learning applications. Most hyperparameter optimization (HPO) methods have not been explored in terms of their efficiencies on such small datasets in the molecular domain. In this paper, we conducted a theoretical analysis of common and specific features for two state-of-the-art and popular algorithms for HPO: TPE and CMA-ES, and we compared them with random search (RS), which is used as a baseline. Experimental studies are carried out on several benchmarks in MoleculeNet, from different perspectives to investigate the impact of RS, TPE, and CMA-ES on HPO of GNNs for molecular property prediction. In our experiments, we concluded that RS, TPE, and CMA-ES have their individual advantages in tackling different specific molecular problems. Finally, we believe our work will motivate further research on GNN as applied to molecular machine learning problems in chemistry and materials sciences.

CCS CONCEPTS

• **Computing methodologies** → *Neural networks*; **Search methodologies**; • **Applied computing**; • **General and reference** → *Experimentation*;

KEYWORDS

Graph Neural Networks, Molecular Property Prediction, Hyperparameter Optimisation

ACM Reference Format:

Yingfang Yuan, Wenjun Wang, and Wei Pang. 2021. A Systematic Comparison Study on Hyperparameter Optimisation of Graph Neural Networks for

*Corresponding author

This work is licensed under a Creative Commons Attribution International 4.0 License.
GECCO '21, July 10–14, 2021, Lille, France
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8350-9/21/07.
<https://doi.org/10.1145/3449639.3459370>

Molecular Property Prediction. In *2021 Genetic and Evolutionary Computation Conference (GECCO '21)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3449639.3459370>

1 INTRODUCTION

Graph neural networks (GNNs) are efficient approaches for learning the representations of structured graph data (e.g., molecules, citation networks) [30]. In recent years, several types of GNNs have been proposed for predicting molecular properties, and they have achieved excellent results [8, 10, 14, 18, 19, 29]. Moreover, the application of GNNs accelerated the work in many related domains, including drug discovery [9, 27], biology [6], physics [26], and these GNNs reduced computational cost compared to traditional first-principles methods such as Density Functional Theory [14, 22]. In practice, there are many GNN variants that can be employed for molecular property prediction. Each variant is proposed based upon a distinct idea for feature learning of molecules. For example, GC (graph convolutional network) [8] exploits neural architectures to generalise the chemical operation of the circular fingerprint to extract molecular features. In contrast, Weave [18] and MPNN (message passing neural network) [10] have been proposed to learn molecular features by taking readout operations [28] on atomic features. To learn atomic features, Weave applies global convolution operation, and MPNN uses the message passing process.

However, the hyperparameter selection is a direct impediment for GNNs to achieve excellent results. In general, the process of searching optimal hyperparameters is often a trial-and-error process. Traditionally, people used to adjust hyperparameters manually, but this requires domain experience and intuition. To extricate people from this predicament, random search (RS) has been employed for hyperparameter optimisation (HPO). In brief, RS draw hyperparameter values from uniform distributions within given ranges. The drawn hyperparameter values are evaluated on an objective function, and the one with the best performance will be selected when the given computational resource is exhausted. Although very simple, RS has been proved to be efficient for HPO of neural networks for many problems [4]. In recent years, there is an increasing number of strategies proposed for HPO. TPE [4] and CMA-ES [13] are two state-of-the-art HPO algorithms, and they are proposed to improve the efficiency of search for promising hyperparameters by utilising the experience of previous trials. In this paper, a trial denotes the process of evaluating a hyperparameter setting on an objective function [5].

Research on HPO of GNNs for molecular property prediction is still in its infancy. For example, the pioneering work of GNN presented in [8, 10, 29] did not discuss the problem of HPO in detail. Meanwhile, most HPO methods have not been explored in terms of their efficiency on GNNs when facing this type of problems. Their performance may need to be further investigated because the sizes of molecular datasets vary from hundreds to thousands, which are far less than those of the datasets used for typical deep learning applications (e.g., image recognition, natural language processing). At the same time, predicting molecular properties requires more sophisticated neural architectures to process irregular molecular structures, which is different from image processing problems that have regular spatial patterns within image patches. Therefore, it has become necessary to explore the performance of existing HPO methods on GNNs in molecular domains. This motivates our research, and we conducted methodology comparison and experimental analysis for RS, TPE, and CMA-ES to assess their effects on GNNs as HPO methods. We expect that our research can inform researchers in both molecular machine learning and chemistry and material sciences.

The contributions of our research are summarised as below:

- We conducted systematic experiments to compare and analyse HPO methods including RS, TPE, and CMA-ES for GNN in the molecular domain in terms of their features, computational cost, and performance.
- Our research on HPO for GNN can be applied to a wider range of domains such as physical and social sciences.
- The outcomes of our research will contribute to the development of molecular machine learning [24] as well as HPO for GNNs in general.

The rest of this paper is organized as follows. In Section 2, the related work of RS, TPE, and CMA-ES will be presented. In Section 3, we will conduct a methodology comparison of RS, TPE, and CMA-ES. After that, the design of experiments and detailed experimental results will be described and discussed in Section 4. Finally, in Section 5, conclusions and future work will be given.

2 RELATED WORK

2.1 Random Search

Algorithm 1: Random Search (RS) [5]

```

1 RS( $f, T, U$ ); //  $T$  the total number of trials,  $U$  normal
  distribution
2 for  $t \leftarrow 1$  to  $T$  do
3    $\lambda_t \leftarrow U$ ;
4    $\mathcal{L}_t \leftarrow \text{Evaluate } f(x_{train}, x_{valid}, \lambda_t)$ ;
5    $\mathcal{H} \leftarrow \mathcal{H} \cup (\mathcal{L}_t, \lambda_t)$ ; //  $\mathcal{H}$  historical optimization
    records
6 end
7 return  $\mathcal{H}$ 

```

Random Search (RS) [5] is an approach that uses uniform probability in determining iterative procedures at the price of optimality

[31], and it is helpful for handling many ill-structured global optimization problems with continuous and/or discrete variables [31] such as HPO.

The process of applying RS for HPO is shown in Algorithm 1. In **Line 3** of Algorithm 1, a solution λ_t (i.e. a set of hyperparameter values) is sampled from a uniform distribution U , and then evaluated on objective function $f(x_{train}, x_{valid}, \lambda_t)$ in **Line 4**, which is normally the most expensive step. The evaluation result $\mathcal{L}_t$ and the solution λ_t is paired and recorded in $\mathcal{H}$. The procedures from **Line 3** to **Line 5** are iteratively executed T times. Finally, the best solution is obtained by sorting historic solutions in $\mathcal{H}$ according to their corresponding $\mathcal{L}$ values.

Furthermore, Bergstra. J et al. [5] holds the opinion that RS is the natural baseline for sequential HPO methods. Meanwhile, it is noted that Zelds B. [31] considered that RS is likely to be able to solve large-scale problems efficiently in a way that is not possible for deterministic algorithms. However, when using RS for HPO, the disadvantage is that its performance accompanied by high variance, and it may not produce satisfactory results given a larger search space and limited computational resources.

2.2 TPE

Algorithm 2: TPE [4]

```

1 TPE( $f, T, M_0, S$ ); //  $M$  surrogate model,  $S$  sampling
  function
2 for  $t \leftarrow 1$  to  $T$  do
3    $\lambda_t^* \leftarrow \text{argmax}_{\lambda} S(\lambda, M_{t-1})$ ;
4    $\mathcal{L}_t \leftarrow \text{Evaluate } f(x_{train}, x_{valid}, \lambda_t^*)$ ;
5    $\mathcal{H} \leftarrow \mathcal{H} \cup (\mathcal{L}_t, \lambda_t^*)$ ;
6   Fit a new model  $M_t$  to  $\mathcal{H}$ ; // update  $M$ 
7 end
8 return  $\mathcal{H}$ 

```

The problems of expensive evaluation on fitness function can be solved by sequential model-based global optimisation (SMBO) algorithms [4, 15, 16]. In HPO, a challenge is that the fitness function $f : X \rightarrow \mathbb{R}$ may be expensive to evaluate given a trial of hyperparameters. By using model-based algorithms with a surrogate to approximate f can reduce the evaluation cost. Typically, the core of an SMBO algorithm is to optimise the surrogate for the real fitness function, or some transformation of the surrogate. The two key components of SMBO algorithms are (1) what criterion is defined and optimized to obtain promising solutions given a model (or surrogate) of f , and (2) how f can be approximated via historical trials/solutions.

Tree-structured Parzen Estimator (TPE) [4] is an approach based on SMBO, as shown in Algorithm 2. Compared with RS, it makes the significant change in **Line 3**, in which solutions are sampled by S , instead of uniform distribution U . In S , many candidates λ are drawn according to surrogate model M and the one (λ^*) with the most promising performance evaluated by Expected Improvement (EI, introduce later) is returned [17]. In **Line 4**, λ^* is then evaluated on fitness function f and recorded in $\mathcal{H}$ in **Line 5**. In **Line 6**, the surrogate M is optimised to approximate the real fitness function

by updated $\mathcal{H}$. Finally, the best solution can be obtained by sorting $\mathcal{H}$ after T iterations. In the following paragraphs, we will review the most important work in TPE based on SMBO in detail.

In TPE, EI [17] has been chosen as the criterion to guide the search for optimal solution(s), and it keeps the balance between exploitation and exploration during the search process. The utility function is defined as $u(\lambda) = \max(0, f' - f(\lambda))$, where f' denotes the output of the current best solution, and λ is the solutions we want to find whose $f(\lambda)$ is expected as smaller as possible than f' . The value of the difference between f' and $f(\lambda)$ will be return as a reward. In each iteration, the optimal solution λ^* is given by $EI_{y^*}(\lambda) := \int_{-\infty}^{\infty} \max(y^* - y, 0) p_M(y | \lambda) dy$, where $p_M(y | \lambda)$ is the surrogate of the real fitness function, and y^* represents some quantile of the observed y values.

Meanwhile, modelling of $p_M(y | \lambda)$ is costly, and TPE proposed an indirect way to model $p(\lambda | y)$ (Eq. 1). $p(\lambda | y)$ is defined by Eq. 2, where $\ell(\lambda)$ and $g(\lambda)$ are two density functions modelled by Parzen Estimator [23], a non-parametric method to approximate the probability density function of a random variable. The collected observations are sorted by loss of f , and are divided into two groups based on some quantile. $\ell(\lambda)$ is generated by using the observations $\{\lambda^{(i)}\}$ such that the corresponding loss $f(\lambda)$ was less than y^* and the remaining observations are used to generate $g(\lambda)$. In practice, a number of hyperparameter settings are sampled according to ℓ , evaluated in term of $g(\lambda)/\ell(\lambda)$, and the one that yields the minimum value under $\ell(\lambda)/g(\lambda)$ corresponding to the greatest EI is returned. This solution is then evaluated on the fitness function, and we call this process a trial. In this way, $g(\lambda)$ and $\ell(\lambda)$ are optimized according to the updated observation set, thus the exploration of optimal solutions moves to more promising regions of the whole search space by increasing the densities.

$$p(y | \lambda) = \frac{p(\lambda | y) * p(y)}{p(\lambda)} \quad (1)$$

$$p(\lambda | y) = \begin{cases} \ell(\lambda) & \text{if } y < y^* \\ g(\lambda) & \text{if } y \geq y^* \end{cases} \quad (2)$$

In TPE, Tree-structure means that the hyperparameter space is tree-like, and the value chosen for one hyperparameter determines what hyperparameter will be chosen next and what values are available for it.

2.3 CMA-ES

Covariance matrix adaptation evolution strategy (CMA-ES) [13] is a derivative-free evolutionary algorithm for solving black-box optimization problems, and it has been applied for HPO with large-scale parallel GPU computational resources [4, 21].

The pseudo-code of CMA-ES is shown in Algorithm 3. In **Line 4**, a solution $\lambda_k^{(g)}$ is generated by sampling from a multivariate normal distribution $\mathcal{N}$ until the size of the population is satisfied, where k denotes the index of offspring and g for generation. Thereafter, in **Line 7**, the individuals of λ^g will be evaluated on fitness function f and recorded in $\mathcal{H}$. In **Line 10**, similar to TPE, it exploits the historical information $\mathcal{H}$ to optimise the search process. However, it is noted that CMA-ES optimises $\mathcal{N}$ rather than the surrogate, and we will discuss this in the following paragraphs.

Algorithm 3: CMA-ES [13]

```

1 CMA-ES( $f, G, K, \mathcal{N}$ ) ; //  $G$  the maximum number of
  generations,  $K$  the size of population,  $\mathcal{N}$  normal
  distribution
2 for  $g \leftarrow 1$  to  $G$  do
3   for  $k \leftarrow 1$  to  $K$  do
4      $\lambda_k^g \leftarrow \mathcal{N}$ ;
5   end
6   for  $k \leftarrow 1$  to  $K$  do
7      $\mathcal{L}_k^g \leftarrow \text{Evaluate } f(x_{train}, x_{valid}, \lambda_k^g)$ ;
8      $\mathcal{H} \leftarrow \mathcal{H} \cup (\mathcal{L}_k^g, \lambda_k^g)$ ;
9   end
10  Update  $\mathcal{N}$ 
11 end
12 return  $\mathcal{H}$ 

```

In CMA-ES, the multivariate distribution is re-defined. Specifically, a population of solutions λ^{g+1} (i.e., individuals or offspring) is generated by sampling from a multivariate normal distribution $\mathcal{N}$ (Eq. 3). In Eq. 3, $\mathcal{N}(\mathbf{0}, C^{(g)})$ is a multivariate normal distribution with zero mean and covariance matrix $C^{(g)}$. The latter decides the distribution shape and describes the correlations of the variables. Meanwhile, $\mathbf{m}^{(g)}$ represents the mean value which is the centroid of the distribution, and it determines the search region of the whole search space in generation g . $\sigma^{(g)}$ represents the step size which also decides the global variance; in other words, it controls the size of the region.

$$\mathbf{x}_k^{(g+1)} \sim \mathbf{m}^{(g)} + \sigma^{(g)} \mathcal{N}(\mathbf{0}, C^{(g)}) \quad \text{for } k = 1, \dots, K, \quad (3)$$

To promote the efficiency of sampling, the key is to update $\mathbf{m}^{(g+1)}$, $C^{(g+1)}$, and $\sigma^{(g+1)}$ for the new generation. The mean is updated by the weighted average of μ selected individuals by $\mathbf{m}^{(g+1)} = \mathbf{m}^{(g)} + \sum_{i=1}^{\mu} w_i (\lambda_{i:K}^{(g+1)} - \mathbf{m}^{(g)})$, where w_i means corresponding weights to λ_i . The selection is according to the performance of individuals on the objective function. The novelty of CMA-ES is that it adapts the covariance matrix by combining rank- μ -update and rank-one update [2]. In this way, rank- μ update can efficiently make use of the information from the entire population. At the same time, rank-one update can be used to exploit the information of correlations among generations from the evolution path. This solution keeps the balance between less generations with a large population and more generation with a smaller population. Additionally, CMA-ES introduces the mechanism of step-size control based on evolution path (cumulative step-size adaptation of the global step-size) which aims to approximate the optimal overall step length efficiently by evolution path, because co-variance matrix may not be able to find the optimal overall step length efficiently.

Generally, CMA-ES imitates the biological evolution, assuming that no matter what kind of gene changes, the results (traits) always follow a Gaussian distribution of a variance and zero-mean. Meanwhile, the generated population is evaluated on a objective function. A part of well-performed individuals is selected to guide

evolution, moving to the area where better individuals would be sampled with higher probability.

3 METHODOLOGY COMPARISON

3.1 Common Features

- **Randomness** plays an important role in RS, TPE and CMA-ES. RS is supported by a number of independent uniform distributions with random sampling to explore hyperparameter space and find an optimal solution. TPE and CMA-ES both have exploitation and exploration mechanisms, which means they are given a more specific region of search space to explore compared with RS. TPE draws samples with randomness over the space of density function of $\ell(\lambda)$. Meanwhile, the sampling in CMA-ES is backed by a multivariate distribution.
- **Derivative-free** denotes that the approach does not use derivative information to guide the search for optimal solutions. For RS, TPE, and CMA-ES, as the last paragraph described, they search for the optimal solutions depending on drawing samples with randomness, rather than using gradient information as in the training of neural networks.
- **Termination Condition** As Section 2 shows, RS, TPE, and CMA-ES all have loops which means they need to preset the condition to stop the optimisation. However, this situation might cause a dilemma of balancing computational cost and performance.

3.2 Specific Features

- **Uniform Distribution vs Gaussian Mixture Model vs Multivariate Normal Distribution** The uniform distribution is a symmetric probability distribution that gives a finite number of values with equal probability to be drawn. Meanwhile, in RS, the dimension of hyperparameter solutions corresponds to the required number of uniform distributions, and each uniform distribution is independent. In contrast, in CMA-ES, multivariate distribution is a distribution used to approximately describe a set of correlated real-valued random variables, each of which clusters around a mean value. Furthermore, TPE makes use of the Gaussian mixture model which assumes all the point are generated from a mixture of a number of Gaussian distributions with unknown parameters.
- **Model-based vs Model-free** These are two distinct approaches originally defined in reinforcement learning [12]. RS is a very representative model-free approach that directly searches for better solutions via a process of trial-and-error. In contrast, TPE is a model-based approach that firstly uses density functions $\ell(x)$ and $g(x)$ to model hyperparameter space in terms of the surrogate, then searching solutions over the space of the functions.
- **Bayesian Optimisation vs Evolutionary Strategy** The main idea of applying evolution strategies to black-box optimisation is to search through iterative adjustment of a multivariate normal distribution. The distribution is controlled by the mean and covariance, which are adjusted and moved to the area where better solutions could be sampled with

higher probability. The adjustment generally has four main steps: sampling, evaluation, selecting good individuals, and updating the mean and co-variance by selected individuals. In contrast, in TPE, it starts by Bayesian optimisation to approximate the distribution of hyperparameters and objective function. Instead of using the Gaussian Process to model the distribution, TPE makes use of the Parzen estimator (i.e., kernel density estimation). The posterior distribution is unceasingly updated to approximate the real situation, and an acquisition function (TPE uses EI as the acquisition function) is used to approach the optimal solution.

4 EXPERIMENTAL INVESTIGATION

In this section, we first describe the design of our systematic experiments. We then conducted four sets of experiments in Sections 4.2 and 4.3, to compare RS, TPE, and CMA-ES from the perspective of performance and computational cost.

4.1 Experimental Settings

To investigate the performance of HPO methods for GNN on molecular property prediction, three representative datasets from DeepChem [25] (MoleculeNet)[29] were selected in our experiments: ESOL (1128 records), FreeSolv (642 records), and Lipophilicity (4200 records), which respectively correspond to the tasks of predicting the following molecular properties: water solubility, hydration free energy, and octanol/water distribution coefficient. These properties are crucial in many problems. For example, in drug discovery, lipophilicity is an important property to reflect the affinity of a molecule, and it affects both membrane permeability and solubility [20]. Furthermore, the research presented in [7] analyses molecular solubility data for exploring organic semiconducting materials. Therefore, the above three representative molecular property datasets are worth investigating and will benefit the research of many related problems. Using different sizes of datasets in experiments is helpful for us to conduct more comprehensive analyses. Meanwhile, there are many GNN variants, and we chose graph convolutional network (GC) [8] because it was proposed considering the molecular domain background knowledge. The architectures of GC generalise the chemical operation of circular fingerprint [11] to extract molecular features.

Four hyperparameters of GC are selected for HPO: batch size s_b , learning rate l_r , the size of fully-connected layer s_f , and the size of graph convolution layer s_g . The selection is motivated by the related benchmark work presented in [29] and considers molecular domain knowledge. By default setting, the number of graph convolution layers is two, but their sizes are the same, while the number of fully-connected layer is one. RS, TPE, and CMA-ES are implemented by Optuna [1]. The arguments of HPO methods in our experiments are set empirically or default values offered by Optuna. Meanwhile, considering that most practitioners/researchers do not have sufficient large-scale GPU computational resources as in industry (e.g., DeepMind, FAIR), we would like to assess the performance of HPO given limited resource, so our experiments all are conducted on a single GPU (GeForce GTX 1070), while the MedianPruner technique [1] is used to speed up HPO. We expect our experimental outcomes

would inspire and help other people when they face similar HPO problems and are given limited computational resource.

In our experiments, every dataset is split into training, validation, and test sets with 80%, 10%, and 10%. The training set is used to fit GC given a hyperparameter setting, and the validation set provides an unbiased evaluation of the hyperparameters during the search. The test set is used to evaluate the performance of HPO methods. The evaluation metric is the root mean square error (RMSE) of GC, and the loss function of GC defines the evaluation function. To make our experiments more persuasive, the best hyperparameter settings found by each method are given to GC, and then the GC is run for 30 times independently to calculate the mean of RMSEs on the training, validation, and test datasets. Meanwhile, to statistically analyse the difference between those means, we conducted corresponding t -tests, in which t denotes t -value and h represents the hypothesis. t -test in our all experiments is set with the significance level of $\alpha = 5\%$. When $h = 1$, the equal mean hypothesis is rejected, and a positive t means the latter has better performance than the former; a negative t represents that the former is better than the latter. In contrast, $h = 0$, the equal mean hypothesis is accepted. Empirically, during HPO, we found that the results of evaluating each trial often fluctuated, and to minimize this effect on HPO performance evaluation, we use the mean value of RMSEs from three repeated evaluations of GC as a single evaluation value.

4.2 Computational Cost as a Primary Consideration

In this section, we assess the performance of RS, CMA-ES, and TPE, while considering the computational cost as a priority. In HPO, any discussion and analysis of method performance without considering computational cost is argumentative because naive RS can find optimal solutions if given sufficient time and computational budget, and it is a highly friendly approach for parallel computing. Therefore, any comparison of HPO methods must be performed with acceptable computational cost.

4.2.1 General Experiments. Different HPO methods employ different optimisation strategies. To compare them fairly, we first proposed assigning each of them with a total of 100 trials, assuming all HPO methods have equal computational cost. In other words, RS, TPE, and CMA-ES all have 100 opportunities to evaluate the hyperparameter settings on GC.

Tables 1~3 summary the experiments of RS, TPE, and CMA-ES, respectively on three datasets: ESOL, FreeSolv, and Lipophilicity from DeepChem over the defined search space where batch size s_b and the size of graph convolution layer s_g range from 32 to 256 with the incremental step 32; learning rate l_r is from 0.0001 to 0.0016 with the incremental step of 0.0001; the size of fully-connected layer s_f is from 64 to 512, with the step of 64. This search space has 2^{13} solutions in total. Meanwhile, t -test with $\alpha = 5\%$ is employed to determine if there is a significant difference between the means of RMSEs. $h = 0$ denotes null hypothesis that the performance of GC under the two hyperparameter settings do not have significant difference; $h = 1$ means rejecting the equal mean hypothesis. The t -test results of Tables 1~3 are respectively summarised in Tables 4~6.

In Tables 1 and 2, TPE and CMA-ES did not present significant difference on performance on the test set (see Tables 4 and 5), but

Table 1: The general experimental settings on the ESOL Dataset

Search space: s_b :32~256, $step=32$; l_r :0.0001~0.0016, $step=0.0001$; s_g :32~256, $step=32$; s_f :64~512, $step=64$

HPO Methods	Hyperparameters	Train	Validation	Test	
RS	$s_g=256$	0.2666	0.9067	0.8888	Mean RMSE
	$s_f=64$				
	$l_r=0.0016$	0.0364	0.0542	0.0411	Mean STD
	$s_b=64$				
TPE	$s_g=192$	0.3083	0.8739	0.8667	Mean RMSE
	$s_f=192$				
	$l_r=0.0015$	0.0534	0.0401	0.0476	Mean STD
	$s_b=32$				
CMA-ES	$s_g=256$	0.2939	0.8739	0.8782	Mean RMSE
	$s_f=64$				
	$l_r=0.0016$	0.0458	0.0424	0.0562	Mean STD
	$s_b=32$				

the hyperparameter settings they found have different s_g and s_f in both datasets. The negative value of t means the former is better, for example, the row of RS-TPE in Table 5, the t on the test set is -3.3167, which represents RS has a smaller RMSE value than TPE (i.e., RS is better than TPE); when the value of t is positive, it means the latter is better than the former, for example, the row of RS-TPE in Table 4, the t on the test set is 1.8355, which means the performance of TPE is better than that of RS. A larger absolute value of t indicates a bigger difference. Moreover, RS outperforms the other two methods in Table 2 with significant difference (see Table 5). We consider the problem of FreeSolv may be more complex (the work research in [3] discussed the problem of deviations of calculating hydration free energy), and TPE and CMA-ES are constrained by the size of search space, the number of trials, and the size of the dataset, thus they got stuck in local optima. In contrast, RS uses a completely random strategy which is helpful to deal with this kind of special situation. However, we believe that CMA-ES and TPE would find better solutions if given more trials and a larger search space.

In Table 3, TPE demonstrated better performance than CMA-ES and RS with significant difference (see Table 6). The size of the Lipophilicity dataset is the largest in our experiments; compared with smaller datasets, the evaluation on validation set would return a result with less deviations, which is helpful for TPE and CMA-ES to improve and update their strategies for promising solutions. However, CMA-ES did not show excellent performance in all datasets in this group of experiments, and we considered that CMA-ES is based on the evolution strategy, which means it depends on unceasingly generating new offspring to find solutions, so 100 trials might restrict its performance.

4.2.2 Experiments on One Hour Runtime. The same number of trials may not be able to assign the same computation cost for different HPO methods in practice, because different trials of hyperparameters may incur different computational cost on evaluation. For example, a larger value of s_f/s_g means more trainable parameters, which will take more computational resource for the corresponding trial. Therefore, in this section, we design another set of experiments, in which we assign 1 hour time and the same hardware configuration to different HPO methods on the ESOL dataset with the same search space defined in Table 1, to see which method can find the best solution.

Table 2: The general experimental settings on the FreeSolv Dataset

Search space: s_b :32~256, $step=32$; l_r :0.0001~0.0016, $step=0.0001$; s_g :32~256, $step=32$; s_f :64~512, $step=64$

HPO Methods	Hyperparameters	Train	Validation	Test	
RS	$s_g=256$	0.6197	1.2175	1.1040	Mean RMSE
	$s_f=320$				
	$l_r=0.0015$	0.1248	0.1055	0.0995	Mean STD
	$s_b=32$				
TPE	$s_g=160$	0.6875	1.3425	1.2006	Mean RMSE
	$s_f=448$				
	$l_r=0.0015$	0.1854	0.1711	0.1212	Mean STD
	$s_b=32$				
CMA-ES	$s_g=256$	0.5792	1.2721	1.1967	Mean RMSE
	$s_f=64$				
	$l_r=0.0016$	0.2653	0.1907	0.2128	Mean STD
	$s_b=32$				

Table 3: The general experimental settings on the Lipophilicity Dataset

Search space: s_b :32~256, $step=32$; l_r :0.0001~0.0016, $step=0.0001$; s_g :32~256, $step=32$; s_f :64~512, $step=64$

HPO Methods	Hyperparameters	Train	Validation	Test	
RS	$s_g=96$	0.2682	0.7024	0.6949	Mean RMSE
	$s_f=384$				
	$l_r=0.001$	0.0444	0.0279	0.0248	Mean STD
	$s_b=64$				
TPE	$s_g=224$	0.2475	0.6914	0.6655	Mean RMSE
	$s_f=192$				
	$l_r=0.0015$	0.0328	0.0229	0.0219	Mean STD
	$s_b=32$				
CMA-ES	$s_g=32$	0.3496	0.7191	0.7183	Mean RMSE
	$s_f=64$				
	$l_r=0.0016$	0.0425	0.0309	0.0245	Mean STD
	$s_b=32$				

Table 4: t -Test on the ESOL

	t	h	t	h	t	h
HPO Methods	Train		Valid		Test	
RS - TPE	-3.4671	1	2.6223	1	1.8355	1
RS - CMA-ES	-2.5080	1	2.5666	1	0.7706	0
TPE - CMA-ES	1.0984	0	-0.0007	0	-0.8396	0

Table 5: t -Test on the FreeSolv

	t	h	t	h	t	h
HPO Methods	Train		Valid		Test	
RS - TPE	-1.6328	0	-3.3492	1	-3.3167	1
RS - CMA-ES	0.7435	0	-1.3487	0	-2.1245	1
TPE - CMA-ES	1.8011	1	1.4804	0	0.0855	0

Table 6: t -Test on the Lipophilicity

	t	h	t	h	t	h
HPO Methods	Train		Valid		Test	
RS - TPE	2.0146	0	1.6454	0	4.7676	1
RS - CMA-ES	-7.1251	1	-2.1564	1	-3.6088	1
TPE - CMA-ES	-10.2330	1	-3.8762	1	-8.6244	1

Table 7: Experiments on the ESOL Dataset given one hour running time

Search Space: s_b :32~256, $step=32$; l_r :0.0001~0.0016, $step=0.0001$; s_g :32~256, $step=32$; s_f :64~512, $step=64$

HPO Methods	Number of Trials	Hyperparameters	Train	Validation	Test	
RS	96	$s_g=224$	0.3301	0.8817	0.8994	Mean RMSE
		$s_f=448$				
		$l_r=0.0008$	0.0492	0.0457	0.0544	Mean STD
		$s_b=32$				
TPE	54	$s_g=256$	0.3193	0.8605	0.8634	Mean RMSE
		$s_f=256$				
		$l_r=0.0014$	0.0462	0.0400	0.0408	Mean STD
		$s_b=32$				
CMA-ES	63	$s_g=32$	0.4287	0.9231	0.9688	Mean RMSE
		$s_f=512$				
		$l_r=0.0016$	0.0933	0.0706	0.0845	Mean STD
		$s_b=32$				

Table 8: t -Test on experiments on the ESOL Dataset given one hour running time

	t	h	t	h	t	h
HPO Methods	Train		Valid		Test	
RS - TPE	0.8573	0	1.8774	1	2.8449	1
RS - CMA-ES	-5.0338	1	-2.6530	1	-3.7178	1
TPE - CMA-ES	-5.6546	0	-4.1563	1	-6.0439	1

Within one hour, the best trials of hyperparameters from the HPO methods were selected to configure GC, and this GC will be run 30 times, and the results and t -test are shown in Tables 7 and 8. In Table 7, RS completed the largest number of trials, and the performance is approximately equal to the one shown in Table 1 because it completed almost 100 trials, which is similar to the previous experiment. We believe RS is efficient and stable in such a small search space. Furthermore, TPE obviously showed surprising performance with 54 trials to accomplish almost the same performance as shown in Table 1. Additionally, it is noted that TPE found two different hyperparameter settings respectively shown in Table 1 and 7 but with almost the same performance on the test dataset. Meanwhile, as shown in Table 8, the performance of three HPO methods within 1 hour runtime has significant difference: TPE performs the best. CMA-ES did not reach our expectation that at least it should maintain the similar performance to RS, and we consider that CMA-ES might not be suitable for our particular HPO with insufficient computational cost and relatively small search space. Furthermore, it is noted the under-performing of CMA-ES may be alleviated by further exploring the "meta parameters" of CMA-ES, for example, the population size of CMA-ES. However, this seems to be an even more challenging "meta-HPO" problem, which is beyond the scope of this research. We will explore this in our future work.

4.3 Performance as Primary Consideration

In this section, we design another group of experiments to explore RS, TPE and CMA-ES with performance as a primary consideration by providing as much as possible computational cost.

4.3.1 Experiments on Repeated HPO Runs. In order to purely compare HPO performance, we designed to respectively run three HPO methods on the ESOL dataset for 10 times independently, and each

Table 9: Experiments on the ESOL Dataset with performance as primary consideration

	RS	TPE	CMA-ES
Mean RMSE	0.8529	0.8190	0.8469
Std	0.0169	0.0090	0.0169

Table 10: t -Test on ESOL with performance as primary consideration

HPO Methods	t	h
RS - TPE	5.2891	1
RS - CMA-ES	0.7464	0
TPE - CMA-ES	-4.3625	1

time is assigned with 100 trials and keep the search space the same as that defined in Table 1. The performance is evaluated by calculating the mean of RMSE values of the best trial from every 100 trials. We did not run it on the test data set because the 10 times RMSEs correspond to 10 different hyperparameter settings. Our purpose is to discover the capability of those methods to fit HPO problem, and the results and t -test are shown in Tables 9 and 10. TPE again outperforms RS and CMA-ES, and it also presents more stable performance as with less standard deviation. t -test in Table 10 shows the same outcome on the test dataset to Table 1 that RS and CMA-ES in this problem and search space have similar performance (no significant difference). So, we believe the CMA-ES still has room for improvement for its performance in our experiments.

4.3.2 Experiments on A Larger Search Space. To further investigate the performance of the HPO methods, we increased the search space so that s_b and s_g range from 8 to 512 with the step size of 8; l_r is changed to from 0.0001 to 0.0032 with the step size of 0.0001; s_f is increased from 32 to 1024 with the step size of 32. The new search space has 2^{22} configurations, and the increments become small and the value ranges are increased. The experimental details are shown in Tables 11, 12, and 13; while the corresponding t -test are presented in Tables 14, 15, and 16.

In the three datasets, in general, the RMSEs for RS, TPE, and CMA-ES on the test datasets have improved compared with the experiments in Section 4.2 given the same number of trials. Meanwhile, by observing the results on the validation and test datasets for all three datasets, we do not see over-fitting issues. In ESOL, TPE and CMA-ES have almost the same performance, and both of them are better than RS as indicated by t -test (see Table 14). In addition, in FreeSolv, three HPO methods show no significant difference performance, while TPE and CMA-ES made the improvements compared with the previous experiments (see Table 2). It is sensible that a potentially complex problem should be given a large search space to find the most suitable hyperparameters. In Table 13, the rank of performance of the methods is still that TPE is the best, and RS is better than CMA-ES.

Table 11: Experiments in larger search space on the ESOL Dataset

s_b :8~512, $step$ =8; l_r :0.0001~0.0032, $step$ =0.0001; s_g :8~512, $step$ =8; s_f :32~1024, $step$ =32

HPO Methods	Hyperparameters	Train	Validation	Test	
RS	s_g =384	0.3190	0.8727	0.8479	Mean RMSE
	s_f =160				
	l_r =0.0016				Mean STD
	s_b =24				
TPE	s_g =312	0.5089	0.8203	0.781	Mean RMSE
	s_f =224				
	l_r =0.003				Mean STD
	s_b =8				
CMA-ES	s_g =512	0.5793	0.848	0.7772	Mean RMSE
	s_f =1024				
	l_r =0.0032				Mean STD
	s_b =8				

Table 12: Experiments in larger search space on the FreeSolv Dataset

s_b :8~512, $step$ =8; l_r :0.0001~0.0032, $step$ =0.0001; s_g :8~512, $step$ =8; s_f :32~1024, $step$ =32

HPO Methods	Hyperparameters	Train	Validation	Test	
RS	s_g =200	0.3747	1.2412	1.0880	Mean RMSE
	s_f =64				
	l_r =0.0030				Mean STD
	s_b =48				
TPE	s_g =424	0.6144	1.1288	1.0620	Mean RMSE
	s_f =224				
	l_r =0.0008				Mean STD
	s_b =16				
CMA-ES	s_g =512	0.6973	1.2329	1.0835	Mean RMSE
	s_f =32				
	l_r =0.0032				Mean STD
	s_b =8				

Table 13: Experiments in larger search space on the Lipophilicity Dataset

s_b :8~512, $step$ =8; l_r :0.0001~0.0032, $step$ =0.0001; s_g :8~512, $step$ =8; s_f :32~1024, $step$ =32

HPO Methods	Hyperparameters	Train	Validation	Test	
RS	s_g =312	0.2570	0.6736	0.6552	Mean RMSE
	s_f =32				
	l_r =0.0031				Mean STD
	s_b =32				
TPE	s_g =496	0.2413	0.6786	0.6395	Mean RMSE
	s_f =32				
	l_r =0.0022				Mean STD
	s_b =24				
CMA-ES	s_g =248	0.2442	0.6931	0.6826	Mean RMSE
	s_f =480				
	l_r =0.0015				Mean STD
	s_b =120				

Table 14: t -test on the ESOL in larger search space

	t	h	t	h	t	h
HPO Methods	Train		Valid		Test	
RS - TPE	-7.7384	1	2.7832	1	4.9451	1
RS - CMA-ES	-10.7815	1	1.0330	0	3.5578	1
TPE - CMA-ES	-2.1105	1	-0.9456	0	0.1855	0

Table 15: *t*-test on the FreeSolv in larger search space

	<i>t</i>	<i>h</i>	<i>t</i>	<i>h</i>	<i>t</i>	<i>h</i>
HPO Methods	Train		Valid		Test	
RS - TPE	-11.0149	1	3.6976	1	0.9385	0
RS - CMA-ES	-16.7183	1	0.2580	0	0.1661	0
TPE - CMA-ES	-3.6269	1	-3.2042	1	-0.7475	0

Table 16: *t*-test on the Lipophilicity in larger search space

	<i>t</i>	<i>h</i>	<i>t</i>	<i>h</i>	<i>t</i>	<i>h</i>
HPO Methods	Train		Valid		Test	
RS - TPE	2.7588	1	-0.7815	0	2.7989	1
RS - CMA-ES	1.3964	0	3.0423	1	-5.2825	1
TPE - CMA-ES	-0.3299	0	-2.8299	1	-8.9802	1

5 CONCLUSION AND FUTURE WORK

Overall, our experimental results indicate that TPE is the most suited HPO method for GNN as applied to our molecular property prediction problems given limited computational resources. Meanwhile, RS is the simplest method but can achieve comparable performance against TPE and CMA-ES. In our future work, facing molecular problems on small datasets, the use of CMA-ES also deserves further investigation, and we believe that CMA-ES, RS, and TPE will have very similar performance given more computational budget. Furthermore, as mentioned in Section 4.3.1, the selection of the "meta-parameters" for HPO methods deserve more research; we will investigate the impact of HPO methods' meta-parameter values on their performance.

Finally, we expect that our work will help people from various fields (e.g., machine learning, chemistry, materials science) when they are facing similar type interdisciplinary problems. As the applications of GNNs have been explored in many areas and indeed benefited the research in those areas, we believe that our research outcomes would give them valuable insights to facilitate their research.

ACKNOWLEDGMENTS

This research is supported by the Engineering and Physical Sciences Research Council (EPSRC) funded Project on New Industrial Systems: Manufacturing Immortality (EP/R020957/1). The authors are also grateful to the Manufacturing Immortality consortium.

DATA STATEMENT

All data used in our experiments are from MoleculeNet [29], which are publicly available in <http://moleculenet.ai/datasets-1>.

REFERENCES

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2623–2631.
- [2] Youhei Akimoto, Yuichi Nagata, Isao Ono, and Shigenobu Kobayashi. 2010. Bidirectional relation between CMA evolution strategies and natural evolution strategies. In *International Conference on Parallel Problem Solving from Nature*. Springer, 154–163.
- [3] D Asthagiri, Lawrence R Pratt, and HS Ashbaugh. 2003. Absolute hydration free energies of ions, ion–water clusters, and quasichemical theory. *The Journal of chemical physics* 119, 5 (2003), 2702–2708.
- [4] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In *25th annual conference on neural information processing systems (NIPS 2011)*, Vol. 24. Neural Information Processing Systems Foundation.
- [5] James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *The Journal of Machine Learning Research* 13, 1 (2012), 281–305.
- [6] Pasquale Bove, Alessio Micheli, Paolo Milazzo, and Marco Podda. 2020. Prediction of Dynamical Properties of Biochemical Pathways with Graph Neural Networks.. In *BIOINFORMATICS*. 32–43.
- [7] Duc T Duong, Bright Walker, Jason Lin, Chunki Kim, John Love, Balaji Purushothaman, John E Anthony, and Thuc-Quyen Nguyen. 2012. Molecular solubility and hansen solubility parameters for the analysis of phase separation in bulk heterojunctions. *Journal of Polymer Science Part B: Polymer Physics* 50, 20 (2012), 1405–1413.
- [8] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. *Advances in neural information processing systems* 28 (2015), 2224–2232.
- [9] Evan N Feinberg, Debnil Sur, Zhenqin Wu, Brooke E Husic, Huanghao Mai, Yang Li, Saisai Sun, Jianyi Yang, Bharath Ramsundar, and Vijay S Pande. 2018. PotentialNet for molecular property prediction. *ACS central science* 4, 11 (2018), 1520–1530.
- [10] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212* (2017).
- [11] Robert C Glen, Andreas Bender, Catrin H Arnbj, Lars Carlsson, Scott Boyer, and James Smith. 2006. Circular fingerprints: flexible molecular descriptors with applications from physical chemistry to ADME. *IDrugs* 9, 3 (2006), 199.
- [12] Adrian M Haith and John W Krakauer. 2013. Model-based and model-free mechanisms of human motor learning. In *Progress in motor control*. Springer, 1–21.
- [13] Nikolaus Hansen. 2016. The CMA evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772* (2016).
- [14] Zhongkai Hao, Chengqiang Lu, Zhenya Huang, Hao Wang, Zheyuan Hu, Qi Liu, Enhong Chen, and Cheekong Lee. 2020. ASGN: An Active Semi-supervised Graph Neural Network for Molecular Property Prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 731–752.
- [15] Frank Hutter. 2009. *Automated configuration of algorithms for solving hard computational problems*. Ph.D. Dissertation. University of British Columbia.
- [16] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2011. Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization*. Springer, 507–523.
- [17] Donald R Jones. 2001. A taxonomy of global optimization methods based on response surfaces. *Journal of global optimization* 21, 4 (2001), 345–383.
- [18] Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. 2016. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design* 30, 8 (2016), 595–608.
- [19] George Lamb and Brooks Paige. 2020. Bayesian Graph Neural Networks for Molecular Property Prediction. *arXiv preprint arXiv:2012.02089* (2020).
- [20] Sonia Lobo. 2020. Is there enough focus on lipophilicity in drug discovery? *Expert opinion on drug discovery* 15, 3 (2020), 261–263.
- [21] Ilya Loshchilov and Frank Hutter. 2016. CMA-ES for hyperparameter optimization of deep neural networks. *arXiv preprint arXiv:1604.07269* (2016).
- [22] Chengqiang Lu, Qi Liu, Chao Wang, Zhenya Huang, Peize Lin, and Lixin He. 2019. Molecular property prediction: A multilevel quantum interactions modeling perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1052–1060.
- [23] Emanuel Parzen. 1962. On estimation of a probability density function and mode. *The annals of mathematical statistics* 33, 3 (1962), 1065–1076.
- [24] Philipp M Pflüger and Frank Glorius. 2020. Molecular machine learning: the future of synthetic chemistry? *Angewandte Chemie International Edition* 59, 43 (2020), 18860–18865.
- [25] Bharath Ramsundar, Peter Eastman, Patrick Walters, Vijay Pande, Karl Leswing, and Zhenqin Wu. 2019. *Deep Learning for the Life Sciences*. O'Reilly Media. <https://www.amazon.com/Deep-Learning-Life-Sciences-Microscopy/dp/1492039837>.
- [26] Jonathan Shlomi, Peter Battaglia, and Jean-Roch Vlimant. 2020. Graph neural networks in particle physics. *Machine Learning: Science and Technology* 2, 2 (2020), 021001.
- [27] Oliver Wieder, Stefan Kohlbacher, Méline Kuenemann, Arthur Garon, Pierre Ducrot, Thomas Seidel, and Thierry Langer. 2020. A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies* (2020).
- [28] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE*

- transactions on neural networks and learning systems* (2020).
- [29] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science* 9, 2 (2018), 513–530.
- [30] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018).
- [31] Zeld B Zabinsky. 2010. Random search algorithms. *Wiley encyclopedia of operations research and management science* (2010).