

RESEARCH ARTICLE

Open Access



# MERMAID: an open source automated hit-to-lead method based on deep reinforcement learning

Daiki Erikawa<sup>1</sup>, Nobuaki Yasuo<sup>2</sup> and Masakazu Sekijima<sup>1,2\*</sup>

## Abstract

The hit-to-lead process makes the physicochemical properties of the hit molecules that show the desired type of activity obtained in the screening assay more drug-like. Deep learning-based molecular generative models are expected to contribute to the hit-to-lead process. The simplified molecular input line entry system (SMILES), which is a string of alphanumeric characters representing the chemical structure of a molecule, is one of the most commonly used representations of molecules, and molecular generative models based on SMILES have achieved significant success. However, in contrast to molecular graphs, during the process of generation, SMILES are not considered as valid SMILES. Further, it is quite difficult to generate molecules starting from a certain molecule, thus making it difficult to apply SMILES to the hit-to-lead process. In this study, we have developed a SMILES-based generative model that can be generated starting from a certain molecule. This method generates partial SMILES and inserts it into the original SMILES using Monte Carlo Tree Search and a Recurrent Neural Network. We validated our method using a molecule dataset obtained from the ZINC database and successfully generated molecules that were both well optimized for the objectives of the quantitative estimate of drug-likeness (QED) and penalized octanol-water partition coefficient (PLogP) optimization. The source code is available at <https://github.com/sekijima-lab/mermaid>.

**Keywords:** Molecular generation, Lead Optimization, Hit-to-Lead, Monte Carlo Tree Search, Drug Discovery

## Introduction

Approximately 8000 drugs are currently being developed worldwide [1]. From drug discovery to launch, a new drug takes an average of 10 to 15 years to be developed and costs \$2.6 billion [1, 2]. Among the drug candidates that enter Phase I clinical trials, less than 12 % are approved by the Food and Drug Administration (FDA) [1]. After the target protein of a therapeutic drug for a disease has been determined, high-throughput screening (HTS) is used to exhaustively test the binding affinity of thousands to hundreds of thousands of compounds to the

target protein in the search for hit compounds. Although the number of possible structures of a compound is  $10^{60}$  and depends on the quality of the compound library to be tested, the hit rate of HTS is approximately 0.1 % [3], which provides an opportunity to discover unexpected hit compounds but also highlights the problem of high experimental cost. To reduce the number of compounds to be tested, virtual screening, a computer-aided drug design (CADD) method for selecting new drug candidates, was proposed in the late 1990s [4]. In virtual screening, compounds that have a high potential to bind to a target protein are ranked in order from a database of thousands to millions of compounds using an evaluation function that expresses the binding affinity calculated by a computer. The compounds narrowed down by the virtual screening are verified by biochemical experiments

\*Correspondence: [sekijima@c.titech.ac.jp](mailto:sekijima@c.titech.ac.jp)

<sup>1</sup> Department of Computer Science, Tokyo Institute of Technology, 4259–

J3–23, Nagatsuta-cho, Midori-ku, Yokohama, Japan

Full list of author information is available at the end of the article



© The Author(s) 2021. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

[5–7], and those that are actually determined to be active proceed to the hit-to-lead. Hit-to-Lead is a stage in early drug discovery where small molecule compounds hit by high-throughput screening (HTS) are processed through certain optimizations to identify promising lead compounds [8]. In addition to simulation, machine learning (ML) methods such as random forests and deep learning have been used in virtual screening [9–12]; however, molecular design methods using generative models are expected to be used in hit-to-lead [13].

The significant progress made in ML in recent years, especially in terms of deep learning, has led to a breakthrough in image processing and natural language processing [14]. Subsequently, various ML models have been applied in the field of molecular design and have shown impressive results [15]. Gomez-Bombarelli et al. [16] used a variational autoencoder (VAE) for molecular design. Representing the molecule as a continuous variable enables us to perform gradient-based optimization in latent space. Considering that simplified molecular input line entry system (SMILES) is a string, which is one representation of molecules, it is natural to adopt recurrent neural networks (RNNs), which are suitable for time-series data such as strings, for molecular design. Segler et al. [17] used long short-term memory (LSTM), which is an RNN, for molecule generation. Although this method showed high validity, it is not suitable for the purpose of generating molecules with desirable properties. This is because LSTM training is only optimized to satisfy SMILES grammar, and the generation process does not consider the properties of molecules. Therefore, we have to repeat the generation process incessantly until we generate the desired molecules. So far, a variety of SMILES-based methods for optimizing specific chemical properties have been proposed [18, 19]. Xiufeng Yang et al. [20] used Monte Carlo tree search (MCTS) to generate desirable molecules with better efficiency than random sampling in RNN-based molecule generation. The aforementioned methods are SMILES-based molecular generative models, and they cannot take a specific molecule as a starting point during optimization tasks. This is because, unlike the case of molecular graphs, sub-SMILES cannot be considered as valid molecules owing to the nature of SMILES grammar.

Graph representation, which is called a molecular graph, is also a useful representation of molecules. Graph representation is easier to understand visually, and checking the valence allows all generated molecules to be valid. Various ML models such as generative adversarial network (GAN) [21] and VAE [22] have been applied to the generation of molecular graphs [23–25]. These methods often outperformed SMILES-based methods in terms of optimization for certain chemical properties, as

well as for metrics, such as validity and novelty, for generated molecules. However, handling molecular graphs on a computer is more difficult than SMILES. Because VAE deals with likelihoods explicitly, graph matching is necessary to calculate the loss function, which has a high computational cost [26]. Although GANs do not deal with likelihoods explicitly, the GNN used in the discriminator and generator require a significant computational cost [27]. For these reasons, the molecular graph-based approach can only deal with small molecules [28]. In addition, for both SMILES-based and molecular graph-based approaches, reinforcement learning is used to optimize specific chemical properties; however, the reward model, which is represented by neural networks and is not always accurate (especially in the case of extrapolation) [29], must be retrained for each evaluation function.

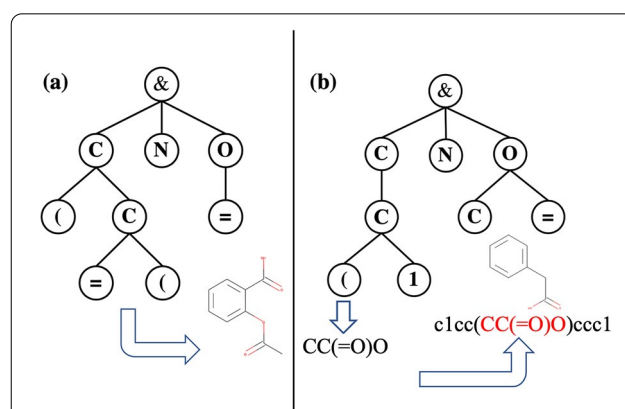
To address these issues, we have introduced *MERMAID*, a generative model based on SMILES using MCTS and RNN. Our model can take a specific molecule as a starting point (Fig. 1), and because we adopt SMILES representation, the restrictions on the size of molecules is not as stringent as the molecular graph-based approach, and our model does not require retraining of the model for each chemical properties.

## Methods

### Optimization of specific molecules

The purpose of this study is to generate the derivatives of specific molecules. The aforementioned MCTS-based generative model cannot start from specific molecules simply, because MCTS only adds nodes to the tail of a search tree. Our proposed method is as follows.

- 1 Extract partial SMILES strings from the initial point of SMILES.



**Fig. 1** Comparison with existing MCTS-based methods **(a)** Existing MCTS-based molecular generative model such as ChemTS [20]. This model generates full SMILES strings through MCTS. **(b)** Our model starting from a specific molecule. It generates partial SMILES and replaces a part of the seed SMILES with the generated one

- 2 Use MCTS to generate a series of characters representing the partial SMILES strings
- 3 Replace the extracted partial SMILES strings with the generated partial SMILES

This method is capable of generating molecules that are obtained by removing or adding a series of zeros or more SMILES characters from SMILES regarded as the starting point (Seed molecule). An overview of our entire methods is given in the Fig. 2. First, the method of (partial) SMILES generation by MCTS will be explained. After that, the RNN and its training procedure will be explained, and details of the substitution procedure used with MCTS will be given.

### MCTS for molecular generation

MCTS [30, 31] is a model-based reinforcement learning approach used to solve large space planning problems by sampling episodes and constructing search trees. The generation of (partial) SMILES by MCTS is illustrated in Fig. 3(a). A node corresponds to a state  $s_i$  and has the value  $Q(s)$ , which represents the evaluation of itself and the number of visits  $N(s)$ . While sampling episodes, MCTS selects a node from a search tree using tree policies and evaluates a new node by rollout. Rollout is the default policy for simulations without adding new nodes to search tree. The details of Rollout are described later.

In molecular generation, a node corresponds to a character of SMILES; therefore, hence, a path from the root node to leaf node corresponds to a SMILES. The MCTS algorithm includes the following four steps and iterates until some convergence condition is satisfied.

### Selection

The purpose of this step is to select a node from a search tree for the expansion of nodes. Starting from the root node, the child node of current node is selected based on a tree policy. The tree policy is discussed later. Node selection is repeated recursively until the leaf node is selected.

### Expansion

The purpose of this step is to expand the current node (the node is selected in Selection step). Some SMILES characters following the current node's SMILES character are selected from predefined vocabulary.

### Simulation

The purpose of this step is to evaluate added nodes. The evaluation of non-terminal nodes is a difficult task in reinforcement learning. Therefore, MCTS evaluates non-terminal nodes using the *Rollout* procedure. The *Rollout* procedure expands recursively from evaluating the nodes until the terminal node appears. When the terminal node appears, the path from the root node to the terminal node corresponds to a complete SMILES. Therefore, We can evaluate the path easily using some metrics such as QED, LogP. We used the score as the non-terminal node.

### Backpropagation

The purpose of this step is to update the value  $Q(s)$  and the number of visit  $N(s)$  of traversed nodes in this episode. Starting from the evaluated node in the Simulation step, the parent node of the current node is updated using the calculated score  $r$  recursively until the root node appears. The update formulas are defined as follows.

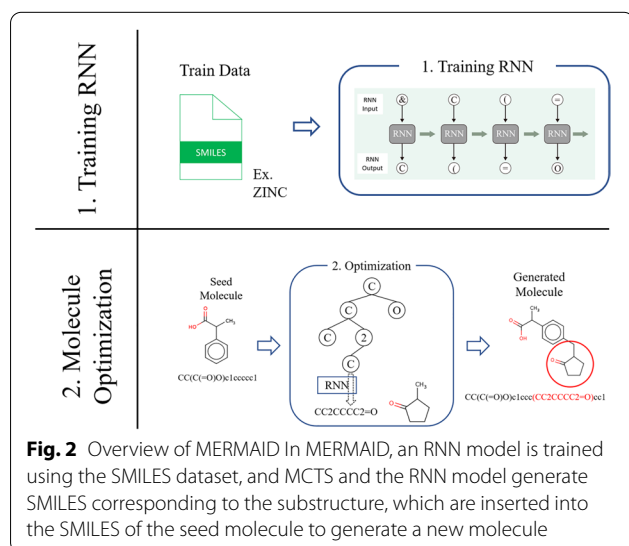
$$Q(s) \leftarrow \frac{Q(s)N(s) + r}{N(s) + 1} \quad (1)$$

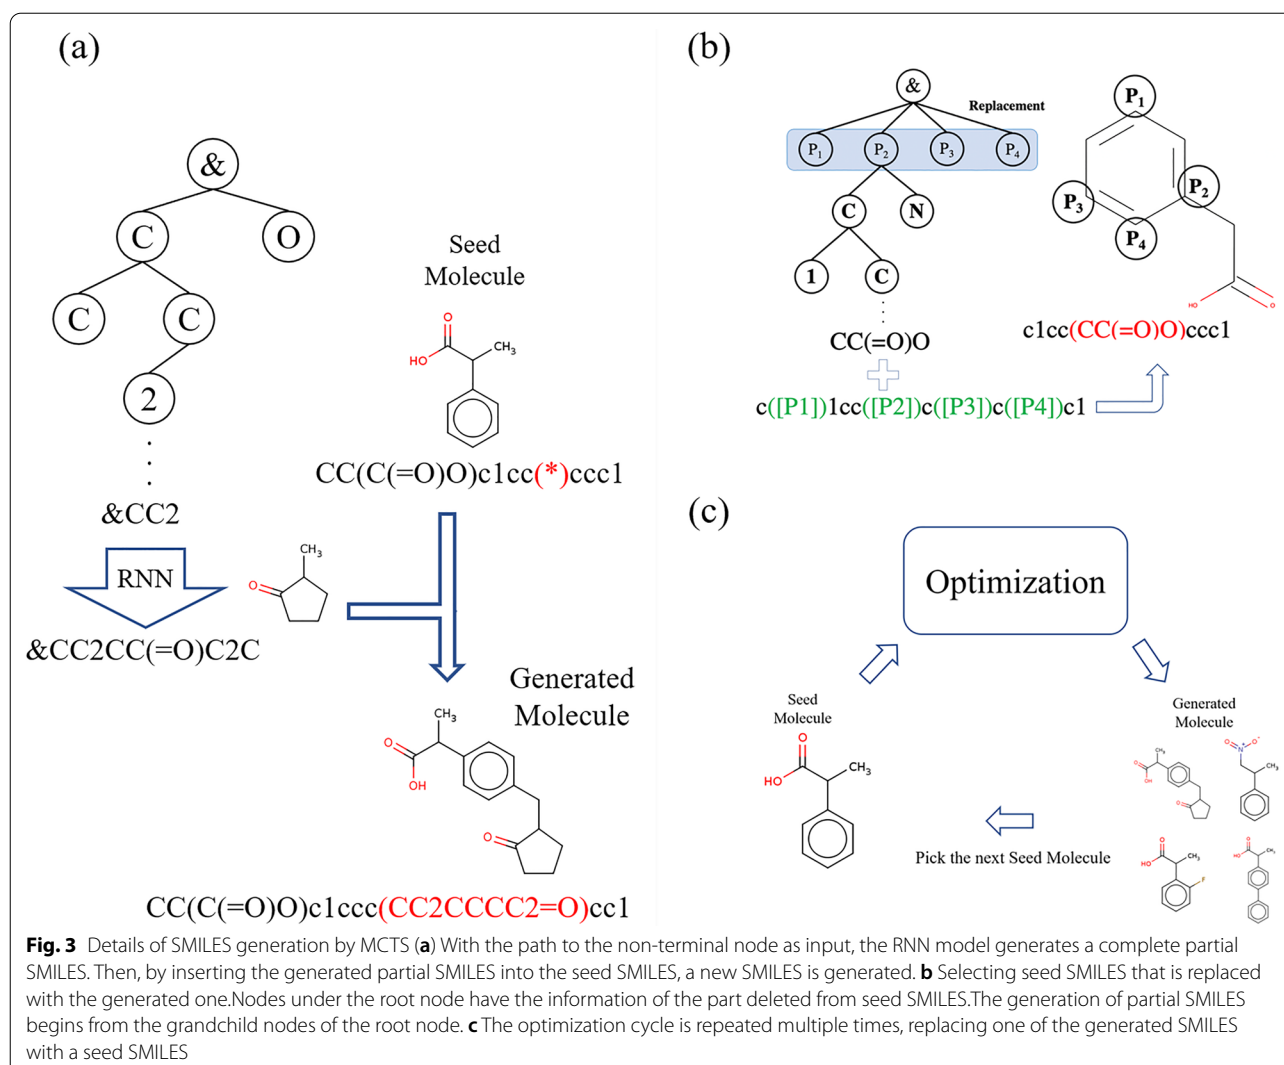
$$N(s) \leftarrow N(s) + 1 \quad (2)$$

Tree policies are important for the performance of MCTS. The upper confidence bound (UCB) score, which is proposed for the multi-armed bandit problem, is often used as the tree policy. Node selection based on the UCB score is as follows.

$$\pi(s) = \arg \max_i \left\{ Q(s_i) + 2C_p \sqrt{\frac{\ln N(s_p)}{N(s_i)}} \right\} \quad (3)$$

where,  $Q(s)$  is the mean estimated value of state  $s$  for its child nodes, and  $N(s)$  is the number of visits to the state  $s$ .  $s_i$  and  $s_p$  are the states of each node  $i$  and the parent node, respectively.  $C_p$  is the hyperparameter of the bias





**Fig. 3** Details of SMILES generation by MCTS **(a)** With the path to the non-terminal node as input, the RNN model generates a complete partial SMILES. Then, by inserting the generated partial SMILES into the seed SMILES, a new SMILES is generated. **(b)** Selecting seed SMILES that is replaced with the generated one. Nodes under the root node have the information of the part deleted from seed SMILES. The generation of partial SMILES begins from the grandchild nodes of the root node. **(c)** The optimization cycle is repeated multiple times, replacing one of the generated SMILES with a seed SMILES

term. The first term corresponds to exploitation, and the second term corresponds to exploration.

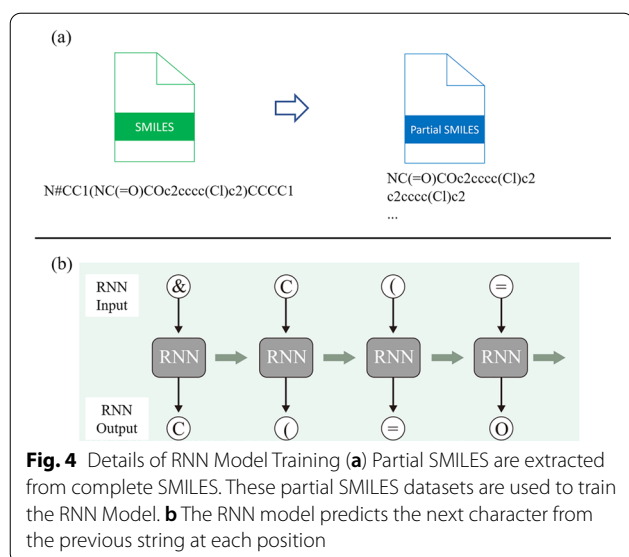
The advantage of using UCB score as the tree policy is that the probability of selecting sub-optimal actions converges to zero as the number of iteration tends to infinity under specific conditions (appropriate  $C_p$  and the value of reward ranges between 0 and 1). However, it is not possible for the number of iterations to tend toward infinity. Furthermore, the performance of rollout-based algorithms degrades similar to that of other algorithms [32]. Since a character following the SMILES string needs to satisfy SMILES grammar, the number of suitable characters that follow the SMILES string is smaller than the number of vocabularies, which is action space. Therefore, a few characters selected by RNN are considered as actions in our approach.  $C_p$  is chosen to be an upper bound of the accumulated reward in practice. However, in this case,  $C_p$  is large to an extent that the state space is

restricted, i.e., exploration is considered more valuable than exploitation. Therefore, we set  $C_p = \frac{1}{\sqrt{2}}$  like other studies [20] that using MCTS.

#### Inference of SMILES using RNN

RNN is a type of neural network that propagates information not only in the direction of layers but also that of time series. In this study, LSTM [33], is a type of RNN superior to normal RNN in terms of longer dependency, is used to capture the features of SMILES grammar.

The role of RNN in MCTS is to select SMILES characters following an incomplete SMILES string in the expansion step and the default policy in the simulation step, as shown in Fig. 4(b). In the expansion step, the incomplete SMILES string  $s_1s_2\dots s_t$  (encoded to  $\mathbf{x}_0\mathbf{x}_1\dots\mathbf{x}_t$ , e.g., one-hot vector) corresponding to path from root to the selected node is the input for RNN. RNN receives the encoded



sequence as input and outputs the probability of selecting characters following the input sequence, which is incomplete SMILES. The probability of selecting a character  $s^i$  based on output of RNN  $y_t$  is as follows.

$$P(s_{t+1}^i | s_1 s_2 \dots s_t) = \frac{\exp(y_t^i)}{\sum_j \exp(y_t^j)} \quad (4)$$

Several characters are selected through a fixed number of samplings from the probability and are added as child nodes to the selected node in the selection step. In the simulation step, a SMILES character that follows the current node is selected recursively until the terminal character is selected in the same way as the expansion step.

### RNN training

The role of MCTS is to generate partial SMILES string as described in above section. Therefore, RNN combined with MCTS should be trained with partial SMILES string rather than full SMILES. In this study, a dataset of partial SMILES strings was generated from a dataset of full SMILES, as shown in Fig. 4(a). Preprocessing was done as follows.

- 1 Partial SMILES strings were extracted exhaustively from full SMILES of the original dataset.
- 2 “Invalid” partial SMILES strings were filtered. Partial SMILES strings are regarded as valid if the SMILES generated by inserting partial SMILES strings into  $C^*C$  or  $C(^*)C$  is valid.

Approximately 250,000 molecules were obtained from the ZINC database to train the RNN model and evaluate

the proposed generative model. 90 % of the molecules were used for training, and the remaining were used to evaluate molecule generation.

The RNN model consists of two layers of LSTM and receives encoded SMILES strings as input and outputs sequences of probability that represents the suitability of a SMILES character following the current input sequence for each position and each SMILES character in vocabulary. This model was trained on the preprocessed training set for 20 epochs using the Adam optimizer [34] to minimize cross entropy loss.

### Replacement of partial SMILES string

In the proposed method, the selection of partial SMILES strings removed from the initial point of SMILES is done in MCTS using the “Replacement” node. Fig. 3(b) shows how to select removed partial SMILES strings. Specifically, the “Replacement” node, which is a child node of the root node, has the values of the starting position of a removed SMILES string and its length of that. A partial SMILES string is generated from the grandchild nodes of the root node, and the generated string is replaced with the part of the initial point of SMILES corresponding to the “Replacement” node.

### Initial point of molecule

This approach is capable of generating new molecules from the original molecule. However, this approach has problem with generated molecules. Because this approach replaces the substructure generated by MCTS with a portion of the original molecule, the generated molecule is a molecule in which only one part of the original molecule has been changed on SMILES. In other words, the generated molecule is not modified in more than one place. Therefore, we propose a method that applies this approach multiple times (Fig. 3(c)). Specifically, for every fixed number of steps, the initial point of SMILES is replaced with the SMILES generated up to that point, and MCTS is performed from the beginning.

For efficient performance, it is important to know how the next initial point of SMILES is selected. In this study, we preferred to optimize SMILES with the maximum reward score, for example QED, LogP, etc. is selected as next initial point. The effect of the difference between fixed or changed initial points is investigated in the Experiment section.

This policy is simple and easy to understand and implement; however, it is possible that the generated molecules flow into a local solution. In addition, it is not necessary that future molecules that have desirable and better properties need to be generated from the best molecule among the generated molecules. Using this policy, good results are obtained in this study; however, the selection



policy for next initial SMILES must be further considered to generate better molecules.

## Experiment

We conducted two experiments to demonstrate the performance of our method. The first experiment involves normal optimization, which modifies a molecule to maximize a single evaluation function. The optimization targets in this experiment were the QED [35] score and penalized LogP. QED is a measure of drug-likeness, and the more drug-like it is, the closer this value is to 1 ranging between 0 and 1. Penalized LogP is defined as follows.

$$PLogP(mol) = LogP(mol) + SAscore(mol) + RingPenalty(mol) \quad (5)$$

The Penalized LogP consists of three terms: the normal LogP (octanol-water partition coefficient), the SA score that penalizes complex structures, and the penalties for large rings. Note that any other target property metrics that can be calculated from SMILES can also be used in this model. The 200 lowest PLogP/QED molecules in the validation data set were selected as the seed in this experiment.

Molecule generation was done in 10,000 steps for each of the test molecules. We analyzed two cases for the proposed approach, depending on whether the seed molecule is fixed. The model in which the seed molecule is fixed is called “Single”, and the other model is called “Multi”. The seed molecule of the “Multi” model is replaced every 2000 steps with the highest scoring (PLogP/QED) molecule generated up to that point.

The second experiment is constrained optimization for comparison with conventional methods, which modifies a molecule to maximize a single evaluation function while satisfying some conditions. In this case, PLogP was optimized with the condition of using the Tanimoto coefficient based on ECFP4 fingerprint. Models perform 4 cycles  $\times$  50 steps = 200 steps of optimization for each of the 800 molecules with the lowest PLogP in the ZINC dataset. Mol-CycleGAN [36] and GCPN [29] were used for comparison.

## Results

### Normal optimization

We evaluate the performance of the optimization task using the best molecular property score and the distribution of generated molecules using validity, uniqueness and novelty. Validity rate is defined as the ratio of SMILES that can be parsed by RDKit to all generated molecules. Uniqueness is the ratio of duplicate molecules to valid molecules, and Novelty is the ratio of molecules that are not included in the training dataset to those included.

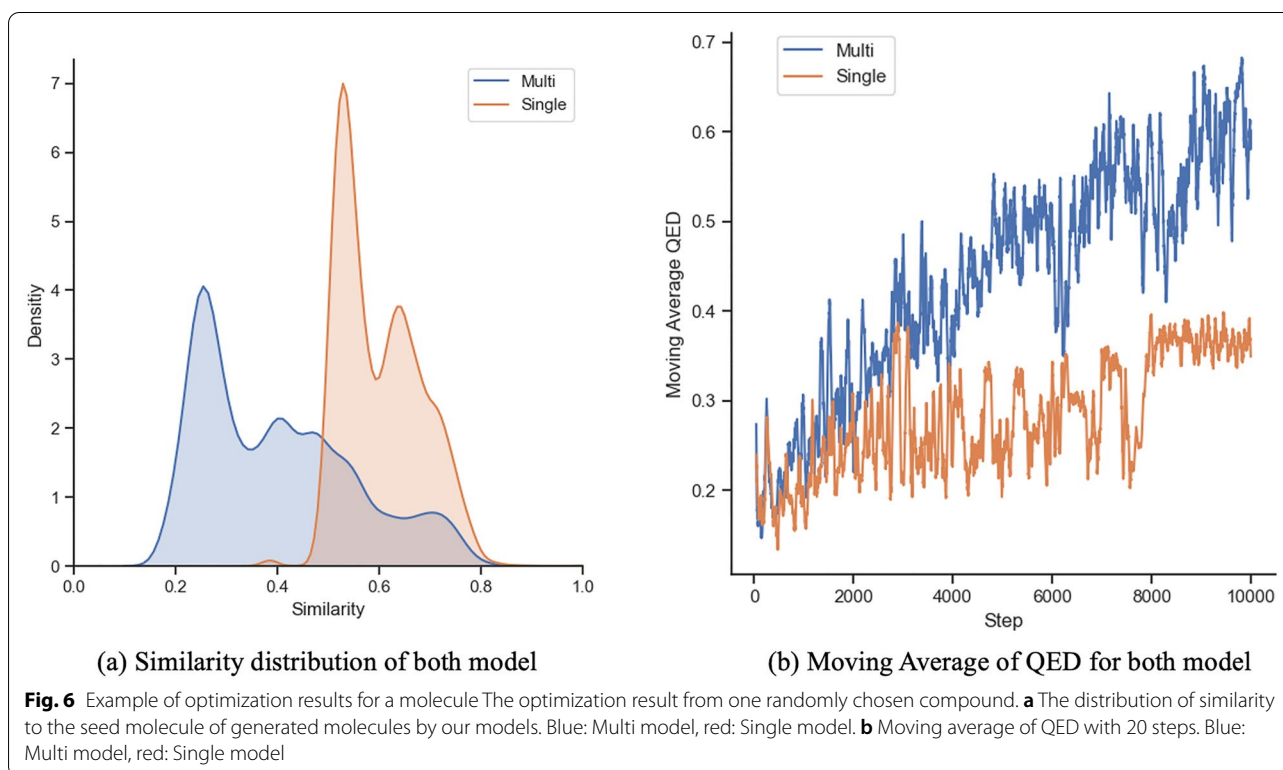
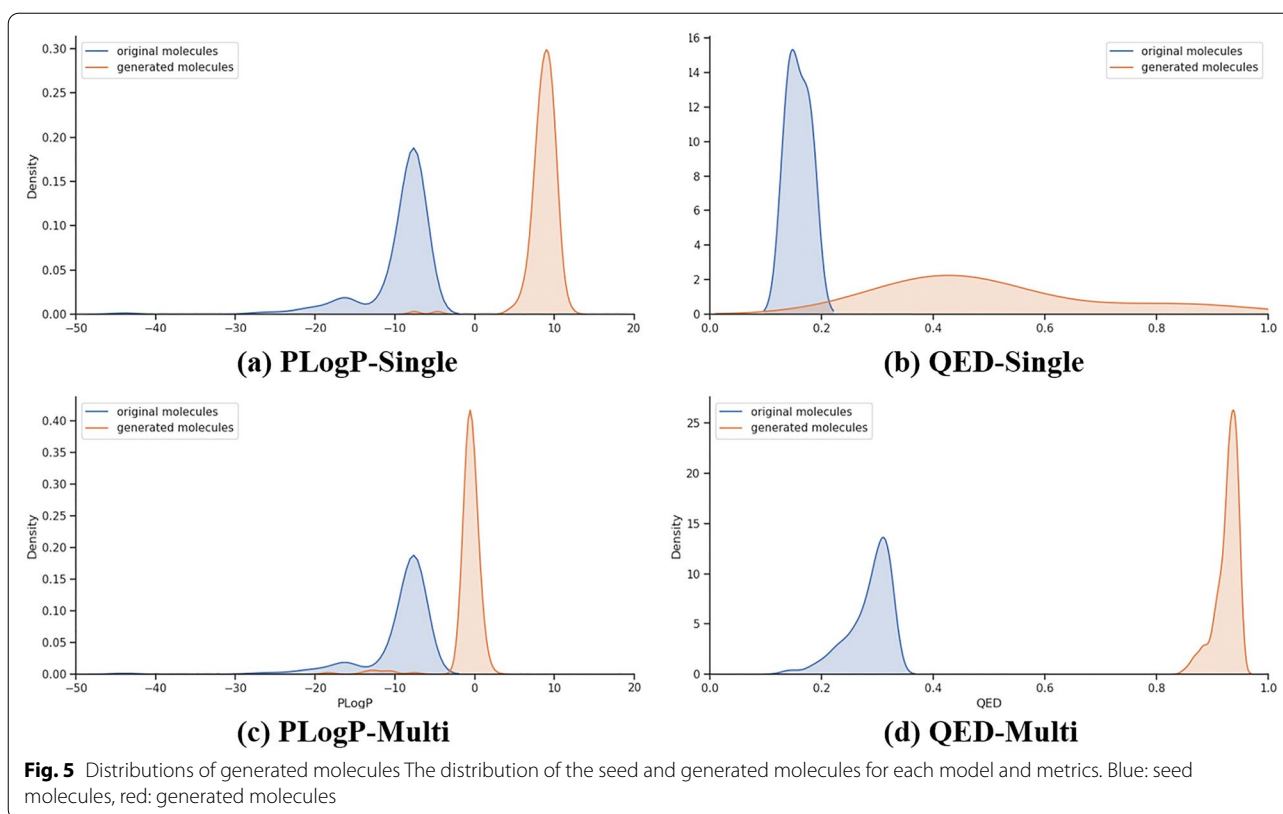
Optimization results are shown in Table 1. Our model shows sufficient results in both the QED and PLogP optimization tasks. In particular, the “Multi” model produces molecules with better scores than the “Single” model. This result is also shown in Fig. 5. The distribution of the QED/PLogP score of seed molecules shifted towards a higher score. This can be confirmed from Fig. 6, which shows that the “Multi” model generates molecules with higher scores as the number of steps increases. The distribution of similarity between the generated molecules and the seed molecules in Fig. 6 shows that the “Multi” model also generates molecules in regions where the “Single” model does not generate. Both the “Single” and “Multi” models generate molecules with relatively high similarity. Additionally, the “Multi” model seeks higher-scoring molecules and expands the chemical space of generated molecules to a lower similarity region. The reason for this is that the “Multi” model can generate a molecule with changes occurring at multiple positions in SMILES as shown in Fig. 7 because the seed is updated, however, the “Single” model cannot generate such a molecule because it inserts partial SMILES at only one position.



The results for the group of all generated molecules are shown in Table 2. Although validity is low, both uniqueness and novelty are high. The specific structures of the generated molecules are shown in Fig. 8. For each starting point molecule, molecules with high scores and different structures are generated. The values of the validity ratios differed considerably depending on the evaluation function. This is because the molecules generated by the PLogP optimization are larger; therefore, the search space is wider, and many invalid molecules are generated.

**Table 1** The results of the optimization tasks

Method	Penalized LogP				QED			
	1st	2nd	3rd	Validity	1st	2nd	3rd	Validity
ZINC	-9.41	—	—	—	0.285	—	—	—
Single	-2.20	-2.39	-2.53	29.8%	0.681	0.671	0.666	62.6%
Multi	11.33	11.20	11.10	31.8%	0.920	0.915	0.912	77.0%

The mean of the top 3 highest scored Penalized LogP and QED scores of the validation set and the validity rate are described





In this manner, the evaluation function itself also affects the performance.

It is observed that the properties of molecules generated by the “Multi” model improve with each step. The number of steps for each iteration is 10,000, and the number of iterations is 5. We obtained good results from this experiment; however, these parameters are not optimum. Therefore, we need to investigate the relation

**Table 2** Mean property of all generated molecules for 200 validation compounds

	Validity	Uniqueness	Novelty
PLogP-Single	0.298	0.981	1.0
PLogP-Multi	0.318	0.991	1.0
QED-Single	0.626	0.945	0.999
QED-Multi	0.770	0.958	0.999

Validity, uniqueness and novelty of all generated molecules are shown for four models. Validity is the ratio of valid SMILES to all generated SMILES. Uniqueness is the ratio of non-duplicate molecules to all valid molecules. Novelty is the ratio of molecules that are not included in the training dataset to all valid molecules

between molecules and the hyperparameters of the “Multi” model, such as the number of steps and selection policy of the next seed molecule, in order to generate better molecules.

### Comparison with conventional methods of constrained optimization

The results of constrained optimization are shown in Table 3. The left column shows the difference in PLogP between the original molecule and the generated molecule with the highest PLogP as Improvement. Our method outperforms others in terms of the properties of molecules. However, success rate, which is the percentage of molecules with similarity above a threshold and improved PLogP, is worse than GCPN. Note that in GCPN, the policy is updated sequentially, while our method has a fixed policy. Thus, our method shows consistent results even when the number of evaluation steps is less (i.e., when optimizing properties that require considerable time to evaluate, such as DFT-based physicochemical values requiring several hours to calculate [37]).

### Discussion

In this experiment, we demonstrated that our method performs well in common tasks (QED and PLogP optimization with ZINC dataset and constrained similarity optimization); however, it can be used for various other tasks as well. To use this method, two areas require modification by the user: the evaluation function and the training data. For the evaluation function, all possible evaluation values that can be calculated from SMILES are available, and the user must design an appropriate evaluation function by combining these based on knowledge of the target task. The training data does not necessarily require modification by the user, as the purpose of training is to capture the features of partial SMILES and because the evaluation values are not directly related to training. Nevertheless, some types of tasks may bias the structure of relevant molecules; in such cases, optimization will improve efficiency if the user prepares the training data in advance.



**Table 3** Results of constrained optimization for 800 validation molecules

$\delta$	GCPN			Mol-CycleGAN			MERMAID		
	Improvement	Similarity	Success (%)	Improvement	Similarity	Success (%)	Improvement	Similarity	Success (%)
0.2	4.12 ± 1.19	0.34 ± 0.11	100	5.79 ± 2.35	0.30 ± 0.11	93.8	9.94 ± 2.74	0.23 ± 0.04	100.0
0.4	2.49 ± 1.30	0.47 ± 0.08	100	2.89 ± 2.08	0.52 ± 0.10	58.8	6.04 ± 2.29	0.42 ± 0.02	100.0
0.6	0.79 ± 0.63	0.68 ± 0.08	100	1.22 ± 1.48	0.69 ± 0.07	19.3	1.99 ± 1.74	0.62 ± 0.02	85.3

The results of GCPN and Mol-CycleGAN are cited from Maziarka et al. [36]. The mean and standard deviation of improvement, similarity, and success rate of generated molecules are shown

As shown in the experimental results, our method produces molecules that improve the evaluation function value. Therefore, it is important to design the evaluation function carefully to avoid generating molecules that deviate from the seed molecule or have chemically unnatural structures.

Future work will focus on several approaches for fundamentally addressing these challenges, such as adding restrictions based on SMILES grammar to the substitution method. Molecular generation, starting from a specific molecule, has an important role in real-world applications; however, it has not been sufficiently studied compared with other types of molecular generation (generation from nothing). Hence, this study aimed to contribute toward filling this gap.

## Conclusions

In this paper, we developed a generative model based on MCTS and RNN to generate derivative molecules starting from a specific molecule. This model generates molecules by generating partial SMILES using MCTS and RNN and replacing it with part of SMILES of the seed molecule. In addition, we propose “Single” and “Multi” models. Unlike the “Single” model, the “Multi” model replaces the seed molecule with one of the generated molecules at a certain number of steps. As a result, it was demonstrated that the “Multi” model is superior to the “Single” model in terms of optimizing the QED score. Additionally, molecules generated by our model have high uniqueness and novelty, and the chemical space consisting of generated molecules is large in terms of similarity and molecular weight.

## Acknowledgements

This work was partially supported by the Platform Project for Supporting Drug Discovery and Life Science Research (Basis for Supporting Innovative Drug Discovery and Life Science Research (BINDS)) from AMED under Grant Number JP20am0101112 and the Japan Society for the Promotion of Science (JSPS) KAKENHI Grant Numbers 20H00620 (To M.S.).

## Authors' contributions

MS conceived the project, DE developed software, and NY supported experiment. DE wrote the manuscript. All authors read and approved the final manuscript.

## Availability of data and materials

MERMAID is freely available at <https://github.com/sekijima-lab/mermaid>.

## Declarations

### Competing interests

The authors declare no competing interests.

### Author details

<sup>1</sup>Department of Computer Science, Tokyo Institute of Technology, 4259–J3–23, Nagatsuta-cho, Midori-ku, Yokohama, Japan. <sup>2</sup>Academy for Convergence of Materials and Informatics (TAC-MI), Tokyo Institute of Technology, S6–23, 2–12–1, Ookayama, Meguro-ku, Tokyo, Japan.

Received: 8 July 2021 Accepted: 15 November 2021

Published online: 27 November 2021

## References

- PhRMA: Biopharmaceuticals in perspective summer 2019 (2019). [https://www.phrma.org/-/media/Project/PhRMA/PhRMA-Org/PhRMA-Org/PDF/P-R/PhRMA\\_2019\\_ChartPack\\_Final.pdf](https://www.phrma.org/-/media/Project/PhRMA/PhRMA-Org/PhRMA-Org/PDF/P-R/PhRMA_2019_ChartPack_Final.pdf) (visited: 2021-3-22)
- Mullard A (2014) New drugs cost US \$2.6 billion to develop. *Nat Rev Drug Discov* 13(12):877
- Varma H, Lo D, Stockwell B (2010) High-throughput and high-content screening for huntington's disease therapeutics. In: *Neurobiology of Huntington's Disease*. CRC Press, Amsterdam, pp. 121–14. <https://doi.org/10.1201/ebk0849390005-c5>
- Schneider G (2010) Virtual screening: an endless staircase? *Nat Rev Drug Discov* 9(4):273–276. <https://doi.org/10.1038/nrd3139>
- Chiba S, Ikeda K, Ishida T, Gromiha MM, Taguchi Y, Iwadata M, Umeyama H, Hsin K-Y, Kitano H, Yamamoto K, Sugaya N, Kato K, Okuno T, Chikenji G, Mochizuki M, Yasuo N, Yoshino R, Yanagisawa K, Ban T, Teramoto R, Ramakrishnan C, Thangakani AM, Velmurugan D, Prathipati P, Ito J, Tsuchiya Y, Mizuguchi K, Honma T, Sekijima M (2015) Identification of potential inhibitors based on compound proposal contest: tyrosine-protein kinase Yes as a target. *Sci Rep* 5:17209
- Chiba S, Ishida T, Ikeda K, Mochizuki M, Teramoto R, Taguchi Y, Iwadata M, Umeyama H, Ramakrishnan C, Thangakani AM, Velmurugan D, Gromiha MM, Okuno T, Kato K, Minami S, Chikenji G, Suzuki SD, Yanagisawa K, Shin W-H, Kihara D, Yamamoto KZ, Moriwaki Y, Yasuo N, Yoshino R, Zozulya S, Borysko P, Stavniichuk R, Honma T, Hirokawa T, Akiyama Y, Sekijima M (2017) An iterative compound screening contest method for identifying target protein inhibitors using the tyrosine-protein kinase yes. *Sci Rep* 7(1):12038
- Chiba S, Ohue M, Gryniukova A, Borysko P, Zozulya S, Yasuo N, Yoshino R, Ikeda K, Shin W-H, Kihara D, Iwadata M, Umeyama H, Ichikawa T, Teramoto R, Hsin K-Y, Gupta V, Kitano H, Sakamoto M, Higuchi A, Miura N, Yura K, Mochizuki M, Ramakrishnan C, Thangakani AM, Velmurugan D, Gromiha MM, Nakane I, Uchida N, Hakariya H, Tan M, Nakamura HK, Suzuki SD, Ito T, Kawatani M, Kudoh K, Takashina S, Yamamoto KZ, Moriwaki Y, Oda K, Kobayashi D, Okuno T, Minami S, Chikenji G, Prathipati P, Nagao C, Mohsen A, Ito M, Mizuguchi K, Honma T, Ishida T, Hirokawa T, Akiyama

- Y, Sekijima M (2019) A prospective compound screening contest identified broader inhibitors for sirtuin 1. *Sci Rep*. <https://doi.org/10.1038/s41598-019-55069-y>
8. Rao V, Srinivas K (2011) Modern drug discovery process: an in silico approach. *J Bioinform Sequence Anal*. 3(5):89–94
  9. Li H, Leung K-S, Wong M-H, Ballester PJ (2015) Improving AutoDock vina using random forest: The growing accuracy of binding affinity prediction by the effective exploitation of larger data sets. *Mol Inform* 34(2–3):115–126. <https://doi.org/10.1002/minf.201400132>
  10. Ragoza M, Hochuli J, Idrobo E, Sunseri J, Koes DR (2017) Protein–ligand scoring with convolutional neural networks. *J Chem Inform Modeling* 57(4):942–957. <https://doi.org/10.1021/acs.jcim.6b00740>
  11. Yasuo N, Sekijima M (2019) Improved method of structure-based virtual screening via interaction-energy-based learning. *J Chem Inform Modeling* 59(3):1050–1061. <https://doi.org/10.1021/acs.jcim.8b00673>
  12. Yasuo N, Nakashima Y, Sekijima M (2018) CoDe-DTI: collaborative deep learning-based drug-target interaction predictor. In: 2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). IEEE, New York, pp. 792–797
  13. Elton DC, Boukouvalas Z, Fuge MD, Chung PW (2019) Deep learning for molecular design—a review of the state of the art. *Mol Syst Design Eng* 4(4):828–849
  14. Elton D, Boukouvalas Z, Fuge M, Chung P (2019) Deep learning for molecular design—a review of the state of the art. *Mol Syst Design Eng*. <https://doi.org/10.1039/C9ME00039A>
  15. Sanchez-Lengeling B, Aspuru-Guzik A (2018) Inverse molecular design using machine learning: generative models for matter engineering. *Science* 361(6400):360. <https://doi.org/10.1126/science.aat2663>
  16. Gómez-Bombarelli R, Wei JN, Duvenaud D, Hernández-Lobato JM, Sánchez-Lengeling B, Sheberla D, Aguilera-Iparraguirre J, Hirzel TD, Adams RP, Aspuru-Guzik A (2018) Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Sci* 4(2):268–276. <https://doi.org/10.1021/acscentsci.7b00572>
  17. Segler MHS, Kogej T, Tyrchan C, Waller MP (2018) Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Central Sci* 4(1):120–131. <https://doi.org/10.1021/acscentsci.7b00512>
  18. Winter R, Montanari F, Steffen A, Briem H, Noé F, Clevert D-A (2019) Efficient multi-objective molecular optimization in a continuous latent space. *Chem Sci*. 10:8016–8024. <https://doi.org/10.1039/C9SC01928F>
  19. Gao K, Nguyen DD, Tu M, Wei G-W (2020) Generative network complex for the automated generation of drug-like molecules. *J Chem Inform Model* 60(12):5682–5698. <https://doi.org/10.1021/acs.jcim.0c00599>
  20. Yang X, Zhang J, Yoshizoe K, Terayama K, Tsuda K (2017) Chemts: an efficient python library for de novo molecular generation. *Sci Technol Adv Mater* 18(1):972–976. <https://doi.org/10.1080/14686996.2017.1401424>
  21. Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. In: Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2. NIPS'14. MIT Press, Cambridge, pp. 2672–2680
  22. Kingma DP, Welling M (2013) Auto-encoding variational Bayes. <http://arxiv.org/abs/1312.6114>. <http://arxiv.org/abs/1312.6114>
  23. Jin W, Barzilay R, Jaakkola T (2018) Junction tree variational autoencoder for molecular graph generation 80:2323–2332
  24. Zhou Z, Kearnes S, Li L, Zare RN, Riley P (2019) Optimization of molecules via deep reinforcement learning. *Sci Rep* 9(1):10752. <https://doi.org/10.1038/s41598-019-47148-x>
  25. Shi C, Xu M, Zhu Z, Zhang W, Zhang M, Tang J (2020) GraphAF: a flow-based autoregressive model for molecular graph generation
  26. Simonovsky M, Komodakis N (2018) Graphvae: towards generation of small graphs using variational autoencoders. In: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4–7, 2018, Proceedings, Part I. pp. 412–422
  27. De Cao N, Kipf T (2018) MolGAN: an implicit generative model for small molecular graphs. ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models
  28. Jin W, Barzilay R, Jaakkola T (2020) Hierarchical generation of molecular graphs using structural motifs
  29. You J, Liu B, Ying R, Pande V, Leskovec J (2018) Graph convolutional policy network for goal-directed molecular graph generation. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems. NIPS'18. Curran Associates Inc., Red Hook, pp. 6412–6422
  30. Coulom R (2006) Efficient selectivity and backup operators in monte-carlo tree search. Proceedings of the 5th international conference on Computers and games, 72–83
  31. Browne CB, Powley E, Whitehouse D, Lucas SM, Cowling PI, Rohlfshagen P, Tavener S, Perez D, Samothrakis S, Colton S (2012) A survey of monte carlo tree search methods. *IEEE Trans Comput Intell AI Games* 4(1):1–43. <https://doi.org/10.1109/TCIAIG.2012.2186810>
  32. Kocsis L, Szepesvári C (2006) Bandit based monte-carlo planning. In: Fürnkranz J, Scheffer T, Spiliopoulou M, eds. Machine Learning: ECML. Springer, Berlin, pp. 282–293
  33. Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput*. 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
  34. Kingma DP, Ba J (2017) Adam: a method for stochastic Opoimization. <http://arxiv.org/abs/1412.6980>. <http://arxiv.org/abs/1412.6980>
  35. Bickerton R, Paolini G, Besnard J, Muresan S, Hopkins A (2012) Quantifying the chemical beauty of drugs. *Nat Chem* 4:90–8. <https://doi.org/10.1038/nchem.1243>
  36. Maziarka L, Pocha A, Kaczmarczyk J, Rataj K, Danel T, Warchol M (2020) Mol-cyclegan: a generative model for molecular optimization. *J Cheminform* 12(1):2. <https://doi.org/10.1186/s13321-019-0404-1>
  37. Senn H, Thiel W (2009) Qm/mm methods for biomolecular systems. *angew chem int ed* 48:1198. *Angewandte Chemie (International ed. in English)* 48, 1198–229. <https://doi.org/10.1002/anie.200802019>

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more [biomedcentral.com/submissions](https://biomedcentral.com/submissions)

