nature portfolio

Peer Review File

ClickGen: Directed Exploration of Synthesizable Chemical Space via Modular Reactions and Reinforcement Learning

Corresponding Author: Professor Chang-Yu Hsieh

This manuscript has been previously reviewed at another journal. This document only contains reviewer comments, rebuttal and decision letters for versions considered at Nature Communications.

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors present a novel deep learning-based molecular generation framework that is fundamentally based on chemical reactions. Distinctively, this study integrates an inpainting technique to overcome prevalent limitations in analogous models, thereby substantially enhancing the uniqueness and diversity of the generated molecules. This innovative approach has been successfully leveraged to design and synthesize two compounds with nanomolar activity against PARP1, notably with one exhibiting superior inhibitory efficacy and low toxicity compared to the commercially available drug, Rucaparib. This work introduces an innovative method, supported by promising validation outcomes, that merits publication consideration. Despite that this work looks nice, several questions concerning the experimental methodology and the clarity of the textual content need to be addressed to fully qualify my recommendation for its publication.

Experimental aspects:
1. In Figure 2, the DENSITY of molecules generated by the method is significantly lower than that of molecules in the REAL database. Can it be inferred that many invalid molecules were generated? The distribution of M-mol properties is illustrated, yet the distributions for l-mol and r-mol are absent. Why is this the case?
2. In Table 3, the authors employ various metrics to evaluate the quality of the generated molecule set. The calculation methods for these metrics should be detailed in the Methods section or Supplementary Materials. Nonetheless, some metrics might cause confusion, such as Diver. Since there are multiple ways to calculate diversity, the authors should clarify the specific method used, or alternatively, present calculations for various diversity metrics.
3. SYBA, a machine learning (deep learning)-based synthesizability metric, should not be directly compared with diversity metrics like SA-score and SC-score. The authors should consider adding another machine learning-based metric for comprehensive validation.
4. The absence of baseline model testing beyond Figure 4 results in a disjointed article structure. The authors are requested to provide clarification or additional experiments.
5. The model was tested using ROCK1 and SARS-Cov2 targets, yet PARP1 was chosen for the final synthesis and activity experiments. Can the authors explain this decision? Also, while the study focuses on the SARS-CoV2 and ROCK1 targets, it would be beneficial to discuss the potential applicability of ClickGen to other types of targets. A brief analysis or discussion on the model's generalizability to diverse biological targets would enhance the manuscript's impact.
6. Is the Safety Index an original evaluation metric created by the authors? I cannot find relevant literature on this, and an explanation from the authors would be appreciated.
7. The Monte Carlo Tree Search (MCTS) method described in the Reinforcement learning model section appears non-reproducible. The calculation of Q(st, a) mentioned in the SELECTION cannot be performed based on the provided method. The authors should revise this section.
8. The authors ultimately selected three molecules for synthesis. A detailed explanation of the selection method in the Methods section would aid readers in replicating the entire process. Furthermore, the authors should discuss the novelty of the generated molecule scaffolds.
Textual and Content Considerations:
1. Given that the method is fundamentally based on fragment-based molecular generation, it is advisable for the authors to include a relevant methodological introduction in the Introduction section.

2. The primary innovation of this approach seems to be its application in real-world drug design scenarios. It is recommended that this aspect be emphasized in both the abstract and introduction, rather than merely highlighting the algorithmic innovation.
3. The analysis in Table 3 refers to "illustrates that both ClickGen models……homolog production." Could the authors clarify what is meant by "homolog production"?
4. The claim in the text that Figure 4 demonstrates ClickGen models “generate more novel scaffolds” requires further substantiation. How is this proven? Additional details are warranted.
5. What is the intended demonstration with the experimental results in Figure 5? Detailed explanations should be added to the main text for clarity.
6. The visualization in Figure 6A is too small to discern the colors of the points representing molecules generated by the ClickGen model. A 2x2 arrangement might enhance readability.
7. Although the authors provided plausible synthetic routes for the generated molecules in the "Feasibility of the ClickGen Model in Real-World Drug Design" section, a similar analysis regarding the ease of sourcing reactants, as done in the “Design PARP1 inhibitors with ClickGen” section, would be beneficial.
8. The approach results in the generation of an extensive number of molecules (up to one hundred thousand), which then undergo tiered screening. Given that the molecules are designed with both synthesizability and affinity in mind, theoretically making them "perfect," why is there a need to generate and then screen such a large quantity? Would users of this software need to generate a similar number of molecules for practical drug design purposes? As the paper involves the practical application of the method, it is crucial to provide a detailed explanation of this aspect.
Overall, the manuscript presents a significant contribution to the field of AI-assisted drug discovery with the development of the ClickGen model. The innovative approach of combining modular reactions with reinforcement learning to enhance synthesizability is a promising direction for future research. With the incorporation of the suggested revisions, the manuscript will be significantly improved and will provide valuable insights to researchers in the field.

Reviewer #2

(Remarks to the Author)
The authors proposed a reinforcement learning (RL)-based molecular generative model, namely ClickGen, that explicitly uses reaction rules adopted in building the Enamine REAL library to produce highly synthesizable molecules. The model is composed of three parts: Reaction Combiner, Inpaiting Generator, and RL. The first part assembles prepared synthons according to the two pre-defined reaction rules: click chemistry and amidation reaction. The second part generates a novel scaffold to connect between two given fragments like the inpaiting technique used in image generation. The last part is employed to generate molecules with high binding scores to the given target protein by using a reward function based on the VINA score. The authors assessed the performances of the proposed models, ClickGen and ClickGen-inpainting, in various aspects such as the validity, novelty, uniqueness, synthesizability, etc of generated molecules, with comparison to those of two benchmark models with similar concepts, BBAR and Synnet. Finally, the authors demonstrated the applicability of the proposed model to lead generation in drug discovery through experimental validation with a target, PARP1. They could identify two lead compounds with nanomolar activities against the target in bioactivity assays. Overall, this manuscript was well written with a significant value for potential publication. However, the authors need to address the following issues for further clarification.

Q1. Proof of superior synthesizability.
Throughout the manuscript, the authors put significant emphasis on the method's ability to generate synthesizable compounds. However, from a statistical point of view, the demonstration in the current text (for example, figure 3) is insufficient to ensure such a claim. For instance, the authors need to:
Perform more quantitative comparisons or statistical tests between the distributions in Figure 3.
Validate the synthesizability of at least 30 selected molecules in the final experiment, for example, by showing a possible synthetic route of each compound.

Q2. Availability of the inpainter's novel synthons.
While ClickGen selects synthons from a set explicitly derived from the REAL database, ClickGen-inpainting generates novel molecular fragments. Then, how can we ensure that the newly generated fragments are also readily available via, for example, chemical suppliers? This is critical because their availability is directly related to synthesizability.

Q3. Role of Reaction Combiner.
The role of Reaction Combiner is not clear enough. Given reaction templates, it is obvious to know whether two synthons can be combined or not. The authors' code already reflects this point: in the combine_syns function in run.py, the authors use the SMARTS patterns of amidation and click chemistry to select possible products deterministically. Then, what's the advantage of training Reaction Combiner and using its probabilistic outputs? Maybe the authors further elucidate the definition of "compatible" and "non-compatible" synthons used during training to clarify this point.

Q4. Is diversity advantageous or disadvantageous?
In Figure 4, the authors claim via the t-SNE analysis that ClickGen's chemical-space coverage is superior to the REAL database. Subsequently, using Figures 6 and S2, they claim that ClickGen is superior because it matches more closely to known inhibitors' distribution than the benchmark models. Suppose that the generated molecules exhibit no significant

deviation in the binding affinity and binding mode of reference inhibitors. Will it be advantageous for the generated molecules to have structures similar to the reference, even in the aspect of "innovativeness"?

In addition, Figure 4 is used to demonstrate that both ClickGen models offer a more diverse set of molecules than that of the REAL database. The authors performed three rounds of virtual screening within the Enamine REAL database to serve as a comparative benchmark for the ClickGen models, ultimately selecting the top 10,000 molecules based on Glide XP. These selected molecules, characterized by high binding affinity as assessed by three distinct docking programs, naturally cover a constrained fraction of the chemical space due to the stringent selection criteria focused on the binding affinity. The author directly compared this highly enriched subset from the REAL database to the generated samples with a reward function based on VINA. This comparison could potentially overestimate the diversity of the generated samples by including molecules with low potency. Therefore, the authors need to apply the same virtual screening criteria to the generated molecules for fair comparison.

The authors should describe how exactly they computed the t-SNE plots.

Q5. Benchmark results in Table 3 and Figure S3
Table 3 and Figure S3 show that BBAR has a validity of 0.83 and an average docking score of ~6 kcal/mol against SARS-Cov-2, respectively. However, these results are notably different from the previous reports on BBAR, which showed a validity of 1.0 in various conditions and an average docking score of less than -8 kcal/mol against SARS-CoV-2.

Considering that BBAR is a conditional generative model, whereas SynNet and ClickGen are optimization methods, the authors should report the training and generation conditions for BBAR used in this work, because the condition setting for conditional generative models can significantly impact their performances.

Q6. Effect of ligand size on docking scores.
Similarly, in the analysis of Figures 6 and S2, the authors are encouraged to compute ligand efficiency to see whether the comparison is still valid. It is well known that computed binding affinities like the VINA score can be readily increased (to more negative values) by increasing the number of ligand atoms.

Q7. RMSD calculation.
In Figures 7 and S3, how exactly the RMSDs were computed? If only the overlapping substructures were used and if ClickGen(-inpainting), BBAR, and SynNet showed meaningful differences in the overlapping ratio between generated and reference molecules, the authors should verify that such differences leave the comparison of RMSD distributions still valid.

Q8. Advantage of modularization.
In the authors' code, the binding score only affects MCTS, and modules like Combiner and Inpainting operate independently. What's the advantage of the current setting over making the two modules more directly used, for example, the protein pocket information? It seems that the current implementation operates by iterating many generation and docking cycles, where the two are independent of each other. The authors mention that one inference of the whole model takes 5-8 hours, and such a long computation time might originate from the independence.

Reviewer #3

(Remarks to the Author)
Introduction [48 – 146]
The authors highlight, with a thoroughly cited summary, the advances in de novo generative models for drug-discovery. They highlight several documented caveats with loss of feature information when using on 2D representations, validation of models is only complete with experimental testing and the synthetic tractability challenges of generative follow-up compounds. To maximise the synthetic success, the authors focus on using two reactions (Click and amidation chemistry) for ClickGen, the generative reaction-based model they describe in this work, for exploring SARS-Cov2/Mpro and PARP1. It would be helpful to the reader if the reaction tractability claims are made sooner in the introduction.

Results and discussion
[148 - 172] – The authors walk through the three components of ClickGen, using customised synthons as inputs for reaction combining, the authors claim the properties of the protein pocket are captured via the docking scores used in the reinforcement model and the inpainting of the synthons, retaining the synthon parent as a core in synthetically tractable click and amide chemistry, again coupled with reinforcement learning leads to greater diversity.

[174 – 210] The authors describe how they trained and validated the reaction combiner reinforcement mod. Ratios of at least 10 x incorrect to correct synthons was required to get the best reaction validity. The authors do not describe how they determined reaction validity, was it a SMARTS pattern validity check?

[212 - 245] In this section, the authors describe the checks they completed to confirm the in-painting component maximised novelty and uniqueness. Methods to assign both novelty and uniqueness are not described in the text or referenced in the SI. Physicochemical properties, MW, H-donor and H-acceptor) of the in-painted molecules, using REAL compounds for training data, were also compared, using kernel density estimation distributions, with good overlay with the REAL

compounds.

[247 – 294] ClickGen's performance is evaluated against other reaction-based generative models, BBAR and Synnet. Using synthetic scoring methods, SA, SC and RA, the authors claim ClickGen produces more tractable molecules. In Figure 3 [293], the authors highlight the differences in scores, with the median values of the distributions for ClickGen indicating better synthetic tractability. Methods for acquiring novelty and uniqueness of the molecules compared between the different models are not described.

[296 - 335] ClickGen's physiochemical properties are compared to the REAL DB using diversity and chemical space heuristics, with the ClickGen in-painting compounds showing greater space exploration via t_SNE plots.

[337 – 356] Physiochemical properties, using docking scores and Tanimoto similarity to known inhibitors, of ClickGen are compared to BBAR and SynNet towards ROCK1 and SARS-Cov2. Besides for the three known inhibitors used for comparison, the source of the other known inhibitors used is not referenced.

[358 - 432] ClickGen's molecules are compared, using RMSD of docked structures and two (His41 and Cys145) interactions, to known inhibitors of SARS-Cov2. Some of ClickGen's compounds are shown to have some similar cores to known inhibitors. It would be helpful if the authors describe in more detail the interactions recapitulated (compared to the known inhibitors) by the ClickGen compounds and the new interactions the ClickGen compounds make to the SARS-target. Please check Figure 7, the inpainting compound's conformation, used to for comparison to L-26, appears to deviate from the known inhibitor. This retrospective analysis shows ClickGen can potentially yield compounds similar to known inhibitors, how would these compounds be selected in a real-world discovery campaign?

[434 – 493] In this section the author's use ClickGen in conjunction with docking, clustering and pharmacophoric matching (Shrodinger) to filter down a 100 000 ClickGen molecules to 3 compounds that were synthesised and tested, with MTT cell, enzymatic and inhibition assays, against PARP1. The results are very impressive, but it is not clear if the lead compounds were a result of the several filtering steps post ClickGen?

Reviewer #4

(Remarks to the Author)
I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
I have carefully reviewed the revisions made by the authors to the manuscript titled "ClickGen: Directed Exploration of Synthesizable Chemical Space via Modular Reactions and Reinforcement Learning," in response to the comments provided by myself and other reviewers. The authors have addressed the concerns raised with detailed explanations and additional data, which have significantly strengthened the manuscript.

Given the improvements and the authors' diligent work in response to the feedback, I believe the manuscript is now in an acceptable form for publication in Nature Communications. It represents a valuable contribution to the field of AI-assisted drug discovery, and I am confident that it will be well-received by the scientific community.

Reviewer #2

(Remarks to the Author)
This reviewer acknowledges that the authors have provided detailed responses to the previous questions and revised the main text and supporting information accordingly. Most of the concerns seem resolved, yet this reviewer would like to leave a few comments and clarification on some of the previous questions.

Qs 1, 2, 3, 4.1, 5, 8, 9.
This reviewer appreciates the authors' detailed responses to the questions. For Q1, in particular, the thorough addition of synthesis guidelines to the generated compounds is appreciated.

Q4.1 Comment on "Is diversity advantageous or disadvantageous?"
This reviewer understands what the authors meant by their response. An additional minor comment is that since the authors used ECFP for t-SNE, the distance in the resulting plots correlates more directly to the structural similarity of molecules than their physicochemical properties. In this regard, mentions like "In TSNE analysis, molecules within the same cluster are deemed to have similar physicochemical properties" and "It is important to note that similar physicochemical properties within a cluster inevitably increases the likelihood of similar chemical structures" may have been written more appropriately.

Q7. RMSD calculation.
It seems the previous question misled the authors. The question was how the authors computed RMSD between molecules with different 2D structures. The active molecules and generated molecules must have different molecular structures; moreover, ClickGen, BBAR, and SynNet may have generated different molecular structures between one another. Because RMSD cannot be defined between two different sets of atoms, one possible choice would be to use the maximum common substructure (MCS). However, the three methods may generate molecules showing MCSs (to active molecules) with different sizes. For example, suppose ClickGen generates more novel molecules than the other two. In that case, the MCSs of the generated molecules w.r.t. active molecules should be smaller, and thus the RMSD can have smaller values more easily.

Reviewer #3

(Remarks to the Author)
Thank you for your constructive responses to my comments and subsequent improvements made to the text.

Reviewer #4

(Remarks to the Author)
I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Version 2:

Reviewer comments:

Reviewer #2

(Remarks to the Author)
This reviewer believes that the authors have clarified all previous issues and revised the manuscript accordingly; therefore, this reviewer recommends the paper for publication.