A Algorithm Pseudocode

In this section we present the pseudocode for consistency model training (Algorithm 1), consistency model sampling (Algorithm 2), and RL training (Algorithm 3).

Algorithm 1 Training

Input: dataset $D = \{(Z_i, u_i)\}_{i \in M}$, where M refers to the number of data points in D, Z_i is composed of x_i and h_i, initial model parameter θ, learning rate η, step schedule $\mathcal{N}(\cdot)$, EMA decay rate schedule $\mu(\cdot)$, $\theta^- \leftarrow \theta$ and $k \leftarrow 0$;
repeat
  Sample $z, u \sim D$, and $n \sim \mathcal{U}[1, N(k) - 1]$;
  Decompose z into x and h;
  Sample $\epsilon_v \sim \mathcal{N}(0, I)$ for $v \in \{x, h\}$;
  Subtract center of gravity from ϵ_x;
  Define $f_\theta^{n,v} \leftarrow f_\theta(v + t_n \cdot \epsilon_v, t_n | u)$ for $v \in \{x, h\}$;
  $L(\theta, \theta^-) \leftarrow \sum_{v \in \{x,h\}} \text{MSE}(f_\theta^{n+1,v}, f_\theta^{n,v})$;
  $\theta \leftarrow \theta - \eta \nabla_\theta L(\theta, \theta^-)$;
  $\theta^- \leftarrow \mu(k)\theta^- + (1 - \mu(k))\theta$;
  $k \leftarrow k + 1$;
until convergence

Algorithm 2 Sampling with Scoring and Selection

Input: Consistency model $f_\theta(\cdot, \cdot) = (f_\theta^x, f_\theta^h)$, sequence of time points $\{\tau_1, \tau_2, \ldots, \tau_{N-1}\}$ where $\tau_1 > \tau_2 > \ldots > \tau_{N-1}$, evaluation starting point m, where $1 \leq m \leq N - 1$
Sample $\epsilon_x \sim \mathcal{N}(0, I)$ and $\epsilon_h \sim \mathcal{N}(0, I)$
Subtract center of gravity from ϵ_x
$\epsilon \leftarrow [\epsilon_x, \epsilon_h]$
$\widetilde{Z}_T \leftarrow \epsilon$
$z \leftarrow f_\theta(\widetilde{Z}_T, T | u)$
Initialize max_score $\leftarrow -\infty$
Initialize $Z_{\text{best}} \leftarrow$ null
for $n = 1$ to $N - 1$ **do**
  Sample $\epsilon_x \sim \mathcal{N}(0, I)$ and $\epsilon_h \sim \mathcal{N}(0, I)$
  Subtract center of gravity from ϵ_x
  $\epsilon \leftarrow [\epsilon_x, \epsilon_h]$
  $\widetilde{Z}_{\tau_n} \leftarrow z + \sqrt{\tau_n - \tau_{n+1}} \cdot \epsilon$
  $z \leftarrow f_\theta(\widetilde{Z}_{\tau_n}, \tau_n | u)$
  if $n \geq m$ **then**
    score $\leftarrow$ CustomScore(z)
    if score $>$ max_score **then**
      max_score $\leftarrow$ score
      $Z_{\text{best}} \leftarrow z$
    end if
  end if
end for
output: Z_{best}

Algorithm 3 Policy Gradient Version of RLCM

Input: Consistency model policy $\pi_\theta = f_\theta(\cdot, \cdot) + \varepsilon$, finetune horizon H, context set set $\mathcal{C}$, batch size b, inference pipeline P
for $i = 1$ **to** M **do**
  Sample b contexts from $\mathcal{C}$, $\boldsymbol{u} \sim \mathcal{C}$.
  $\boldsymbol{Z} \leftarrow P(f_\theta, H, \boldsymbol{u})$ ▷ where $\boldsymbol{Z}$ is the batch of molecules
  Normalize rewards $r(\cdot, \cdot)$ per context
  Split $\boldsymbol{Z}$ into k minibatches.
  for each minibatch **do**
    for $t = 0$ to H **do**
      Accumulate gradients of θ using rule:

$$\nabla_\theta \left[\min \left\{ r(\boldsymbol{x}_0, \boldsymbol{c}) \cdot \frac{\pi_{\theta_{i+1}}(a_t|s_t)}{\pi_{\theta_i}(a_t|s_t)}, r(\boldsymbol{x}_0, \boldsymbol{c}) \cdot \texttt{clip}\left(\frac{\pi_{\theta_{i+1}}(a_t|s_t)}{\pi_{\theta_i}(a_t|s_t)}, 1-\varepsilon, 1+\varepsilon \right) \right\} \right]$$

    end for
    Update parameters based on accumulated gradients.
  end for
end for
Output trained consistency model $f_\theta(\cdot, \cdot)$

B Information on Hyperparameters and Experiment Details

Parameter setting for TurboHopp

Setting	Parameters
TurboHopp	timesteps: 150, 100, 50, 25
Denoiser: GVP	batch size: 256
layers: 6	lr: 1e-4
hidden features: 256	schedule: ReduceLROnPlateau (min: 1e-6, factor: 0.9)
GNN layers: 7	num epochs: 5500
Attention: True	σ_{min}: 0.002
Embedding size: 64	σ_{max}: 80.0
Optimizers: Adam	σ_{data}: 0.5
γ: 1e-3	ρ: 7.0
β: (0.9, 0.995)	
Dataset: PDBBind filtered	
Device: 4x NVIDIA A100 GPUs	

Parameter setting for RLCM

Setting	Parameters
Parameters for Docking Objective and Steric Clashes	gradient accumulation steps: 1
	batch size: 215
Dataset: PDBBind filtered test set	num epochs: 200
Device: 8x NVIDIA A100 GPUs	sample iters: 1
	buffer size: 32
	min count: 16
	train batch size per gpu: 215
	num inner epochs: 1
	lr: 1e-5
	clip range: 1e-4
	max grad norm: 10

C RLCM and Consistency Model Training Curves

For reproducibility, we present the curves from our training runs for both the consistency model and the use of RLCM.

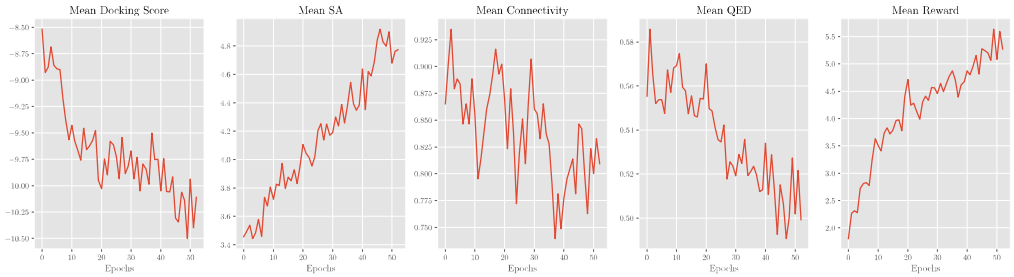

Figure 5: Training curves for the metrics which compose of the loss function. Notice that all either increase or maintain approximately the same value. Connectivity and QED score slightly decrease because we start from a previously RL finetuned checkpoint which optimizes only for connectivity, SA, and QED score.

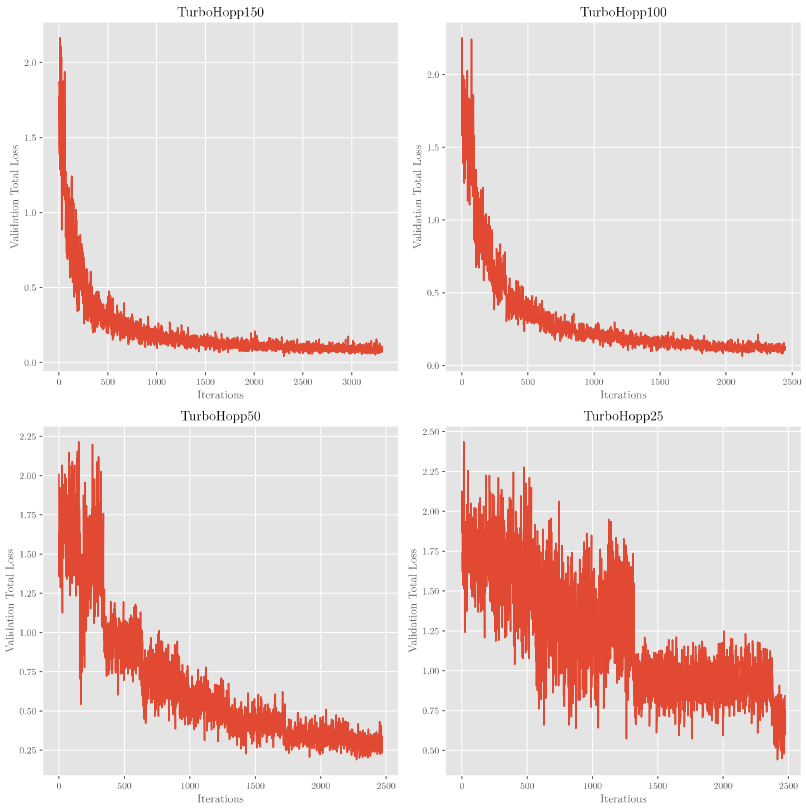

Figure 6: Training curves of total validation loss for different step size variants. We train a number of consistency models to empirically determine the optimal tradeoff between step size fidelity and speed. Turbohopp25 had low validity and proved to be too unstable for consistent generations.t

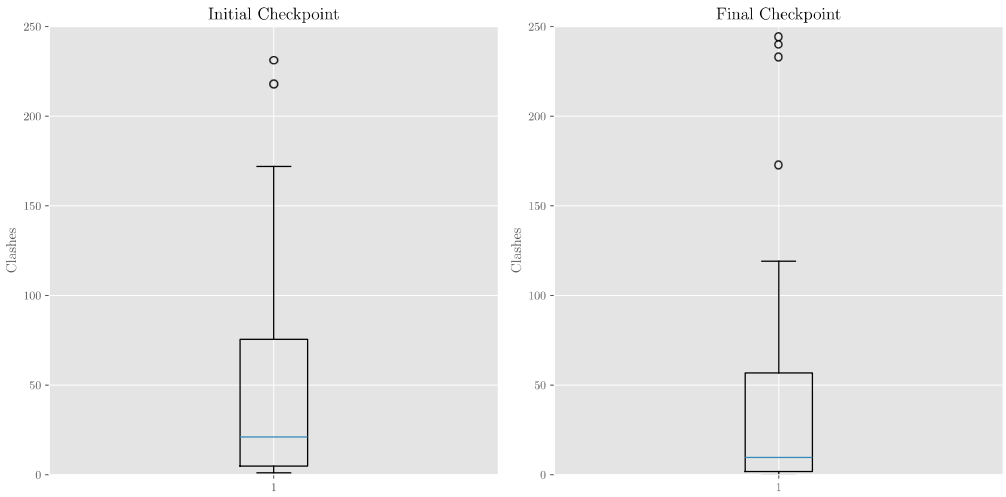

Figure 7: Plot of clashes before and after finetuning with a reward function mentioned in the main text. The initial Turbohopp-100 model was a RLCM finetuned model for connectivity, synthesizability, and QED score. Notice the shift toward smaller number of clashes during training. However, we believe that further iteration of the reward function will lead to more effective finetuning but we leave this to future work.

D Sample Molecules

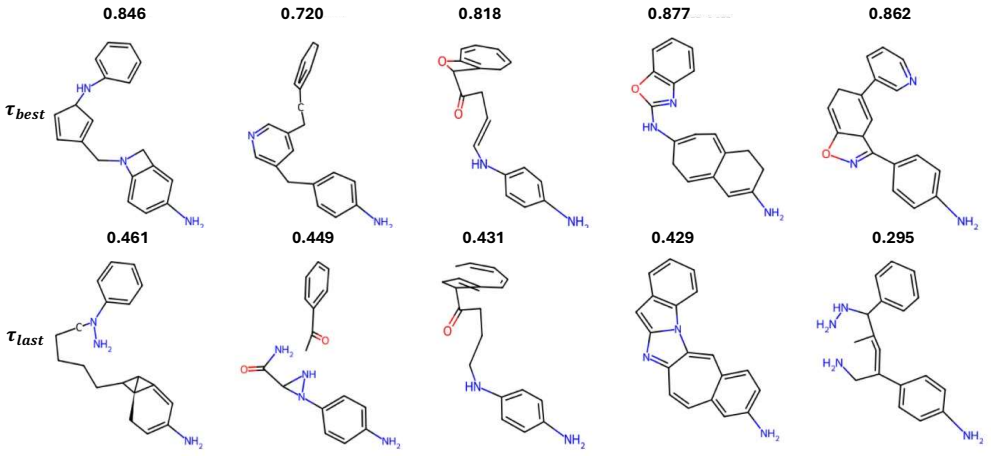

Figure 8: Examples of samples generated for PDB 6QQW with drug-likeliness. 1st row samples collected during multi-step phase with best scores and 2nd row indicates samples from the final step. Connectivity and overall metrics increased when we adopted custom score-based sampling.

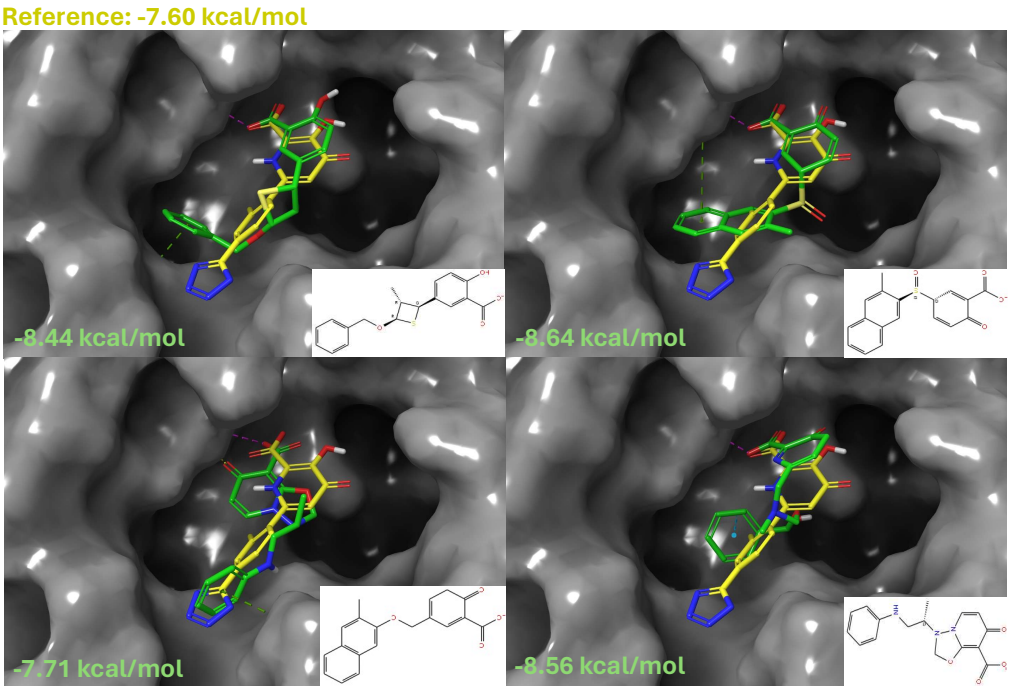

Figure 9: Scaffolds generated for PDB ID 6E6W by TurboHopp-100. Reference molecule is in yellow, while generated molecules are in green. Functional groups are the carboxylic acid and hydroxyl groups in the upper corner. Dotted lines in refer to ligand-receptor intermolecular interactions: green, blue, yellow, purple being pi-cation, pi-pi stacking, hydrogen bonds, halogen bonds respectively. Compared to reference molecule, generated molecules had new interactions and higher binding affinity, while maintaining similar binding pose.

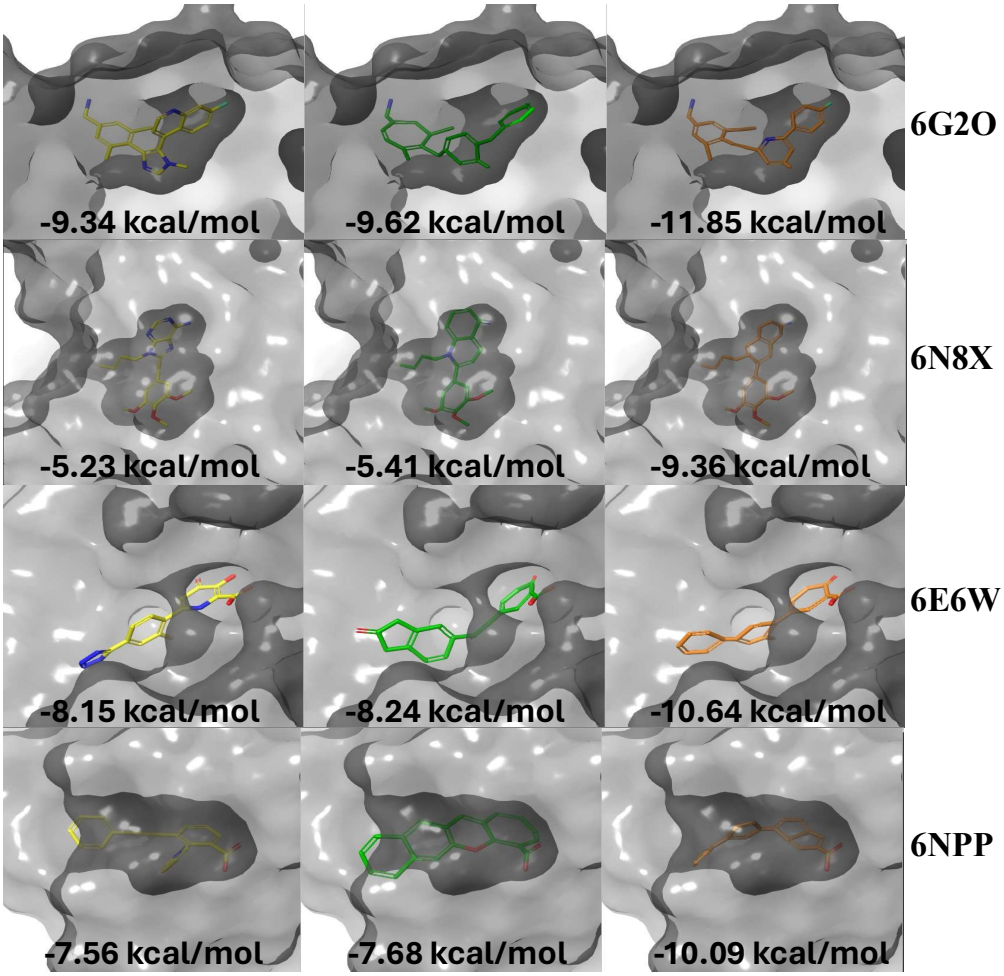

Figure 10: Comparison of Reference molecule(Yellow), and molecules generated by TurboHopp(Green) and TurboHopp-RL(Orange). Notice that Turbohopp and TurboHopp-RL generate molecules that have higher binding affinity with the protein.

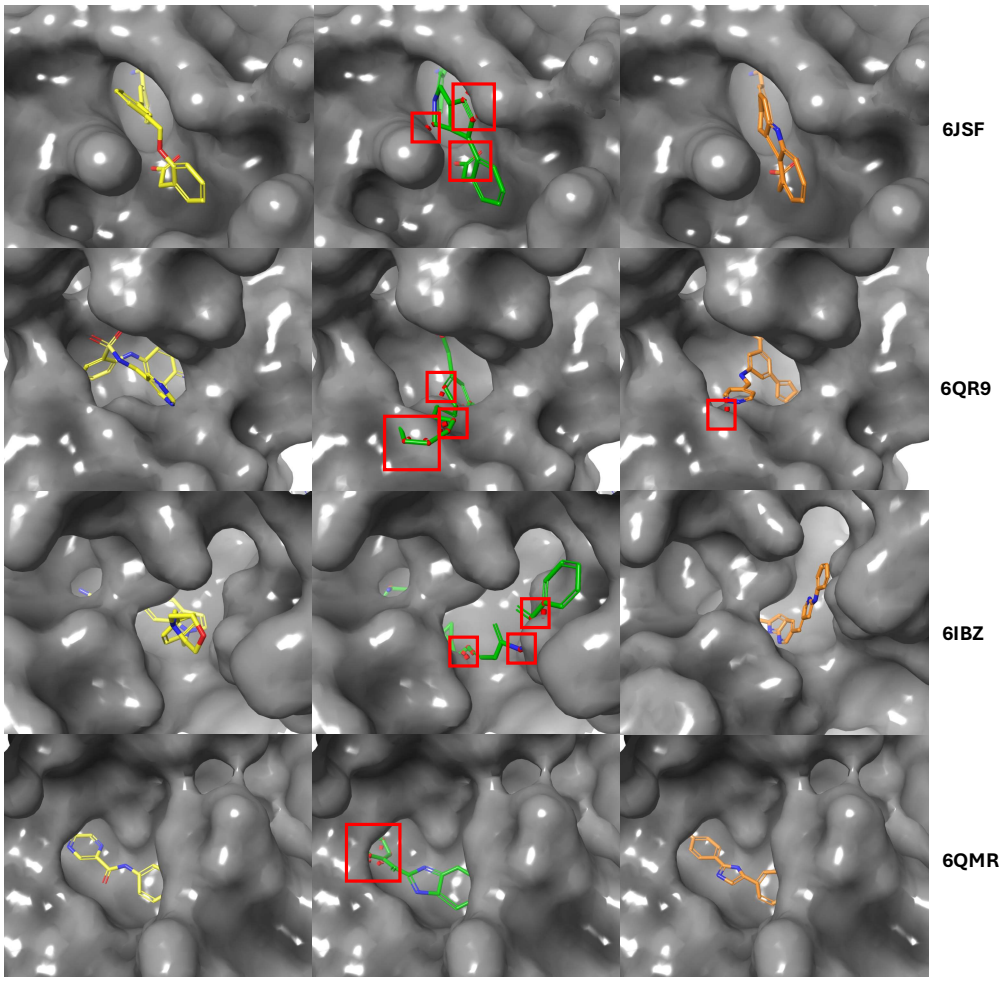

Figure 11: Comparison of Reference molecule(Yellow), and molecules generated by TurboHopp(Green) and TurboHopp-RL(Orange). Red box indicates collision points with protein atoms. TurboHopp-RL generates molecules that has less clashes with the protein.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We have claimed our model's performance and contributions based on empirical findings and experimental results. Our model can be easily adapted to similar frameworks using diffusion models for drug discovery.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We have added limitations and possible further research studies in the conclusion section.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our focus remains on empirical findings and application, bypassing the need for formal theoretical frameworks or validations.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We will share our code/checkpoints upon acceptance. In the paper, we provide the algorithms for training model, RL, and inference. We also share training curves for model training and RL as well as hyperparameters required to reproduce our results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We provide the data we used to train our model, and upon code release, we will specify preprocessing code ,modifications, as well as environment for reproduction.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: We specify the hyperparameters in the appendix. Further, we use a standardized dataset.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: All claims of statistical significance were justified through a number of trials. Multiple seeds are not needed for RLCM since we are applying a method previously shown to work.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We specify the compute resources required in the appendix.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: Yes we read the ethics guidlines and are abiding by them.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Please check the conclusion section of the paper.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: Models are not yet going to be released and we will take into account proper safeguards when releasing models with high risk for misuse.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All terms of use and licences are respected.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: No new assets are released

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: We do not do any crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: We do not do crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

A Algorithm Pseudocode

In this section we present the pseudocode for consistency model training (Algorithm 1), consistency model sampling (Algorithm 2), and RL training (Algorithm 3).

Algorithm 1 Training

Input: dataset $D = \{(Z_i, u_i)\}_{i \in M}$, where M refers to the number of data points in D, Z_i is composed of x_i and h_i, initial model parameter θ, learning rate η, step schedule $\mathcal{N}(\cdot)$, EMA decay rate schedule $\mu(\cdot)$, $\theta^- \leftarrow \theta$ and $k \leftarrow 0$;
repeat
  Sample $z, u \sim D$, and $n \sim \mathcal{U}[1, N(k) - 1]$;
  Decompose z into x and h;
  Sample $\epsilon_v \sim \mathcal{N}(0, I)$ for $v \in \{x, h\}$;
  Subtract center of gravity from ϵ_x;
  Define $f_\theta^{n,v} \leftarrow f_\theta(v + t_n \cdot \epsilon_v, t_n | u)$ for $v \in \{x, h\}$;
  $L(\theta, \theta^-) \leftarrow \sum_{v \in \{x,h\}} \text{MSE}(f_\theta^{n+1,v}, f_\theta^{n,v})$;
  $\theta \leftarrow \theta - \eta \nabla_\theta L(\theta, \theta^-)$;
  $\theta^- \leftarrow \mu(k)\theta^- + (1 - \mu(k))\theta$;
  $k \leftarrow k + 1$;
until convergence

Algorithm 2 Sampling with Scoring and Selection

Input: Consistency model $f_\theta(\cdot, \cdot) = (f_\theta^x, f_\theta^h)$, sequence of time points $\{\tau_1, \tau_2, \ldots, \tau_{N-1}\}$ where $\tau_1 > \tau_2 > \ldots > \tau_{N-1}$, evaluation starting point m, where $1 \leq m \leq N - 1$
Sample $\epsilon_x \sim \mathcal{N}(0, I)$ and $\epsilon_h \sim \mathcal{N}(0, I)$
Subtract center of gravity from ϵ_x
$\epsilon \leftarrow [\epsilon_x, \epsilon_h]$
$\widetilde{Z}_T \leftarrow \epsilon$
$z \leftarrow f_\theta(\widetilde{Z}_T, T | u)$
Initialize max_score $\leftarrow -\infty$
Initialize $Z_{\text{best}} \leftarrow$ null
for $n = 1$ to $N - 1$ **do**
  Sample $\epsilon_x \sim \mathcal{N}(0, I)$ and $\epsilon_h \sim \mathcal{N}(0, I)$
  Subtract center of gravity from ϵ_x
  $\epsilon \leftarrow [\epsilon_x, \epsilon_h]$
  $\widetilde{Z}_{\tau_n} \leftarrow z + \sqrt{\tau_n - \tau_{n+1}} \cdot \epsilon$
  $z \leftarrow f_\theta(\widetilde{Z}_{\tau_n}, \tau_n | u)$
  if $n \geq m$ **then**
    score $\leftarrow$ CustomScore(z)
    if score $>$ max_score **then**
      max_score $\leftarrow$ score
      $Z_{\text{best}} \leftarrow z$
    end if
  end if
end for
output: Z_{best}

Algorithm 3 Policy Gradient Version of RLCM

Input: Consistency model policy $\pi_\theta = f_\theta(\cdot, \cdot) + \varepsilon$, finetune horizon H, context set set $\mathcal{C}$, batch size b, inference pipeline P
for $i = 1$ **to** M **do**
  Sample b contexts from $\mathcal{C}$, $\boldsymbol{u} \sim \mathcal{C}$.
  $\boldsymbol{Z} \leftarrow P(f_\theta, H, \boldsymbol{u})$ ▷ where $\boldsymbol{Z}$ is the batch of molecules
  Normalize rewards $r(\cdot, \cdot)$ per context
  Split $\boldsymbol{Z}$ into k minibatches.
  for each minibatch **do**
    for $t = 0$ to H **do**
      Accumulate gradients of θ using rule:

$$\nabla_\theta \left[\min \left\{ r(\boldsymbol{x}_0, \boldsymbol{c}) \cdot \frac{\pi_{\theta_{i+1}}(a_t|s_t)}{\pi_{\theta_i}(a_t|s_t)}, r(\boldsymbol{x}_0, \boldsymbol{c}) \cdot \texttt{clip}\left(\frac{\pi_{\theta_{i+1}}(a_t|s_t)}{\pi_{\theta_i}(a_t|s_t)}, 1-\varepsilon, 1+\varepsilon \right) \right\} \right]$$

    end for
    Update parameters based on accumulated gradients.
  end for
end for
Output trained consistency model $f_\theta(\cdot, \cdot)$

B Information on Hyperparameters and Experiment Details

Parameter setting for TurboHopp

Setting	Parameters
TurboHopp	timesteps: 150, 100, 50, 25
Denoiser: GVP	batch size: 256
layers: 6	lr: 1e-4
hidden features: 256	schedule: ReduceLROnPlateau (min: 1e-6, factor: 0.9)
GNN layers: 7	num epochs: 5500
Attention: True	σ_{min}: 0.002
Embedding size: 64	σ_{max}: 80.0
Optimizers: Adam	σ_{data}: 0.5
γ: 1e-3	ρ: 7.0
β: (0.9, 0.995)	
Dataset: PDBBind filtered	
Device: 4x NVIDIA A100 GPUs	

Parameter setting for RLCM

Setting	Parameters
Parameters for Docking Objective and Steric Clashes	gradient accumulation steps: 1
	batch size: 215
Dataset: PDBBind filtered test set	num epochs: 200
Device: 8x NVIDIA A100 GPUs	sample iters: 1
	buffer size: 32
	min count: 16
	train batch size per gpu: 215
	num inner epochs: 1
	lr: 1e-5
	clip range: 1e-4
	max grad norm: 10

C RLCM and Consistency Model Training Curves

For reproducibility, we present the curves from our training runs for both the consistency model and the use of RLCM.

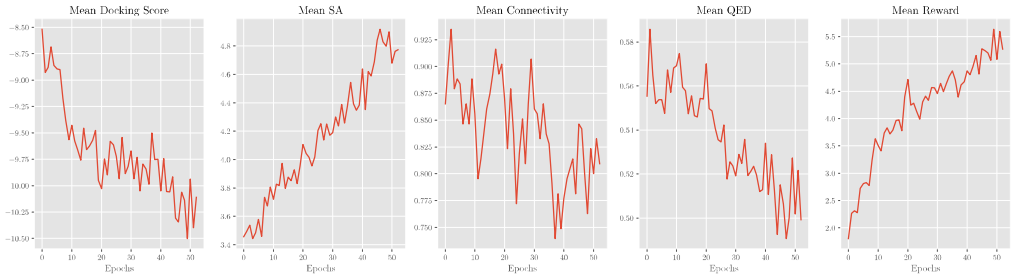

Figure 5: Training curves for the metrics which compose of the loss function. Notice that all either increase or maintain approximately the same value. Connectivity and QED score slightly decrease because we start from a previously RL finetuned checkpoint which optimizes only for connectivity, SA, and QED score.

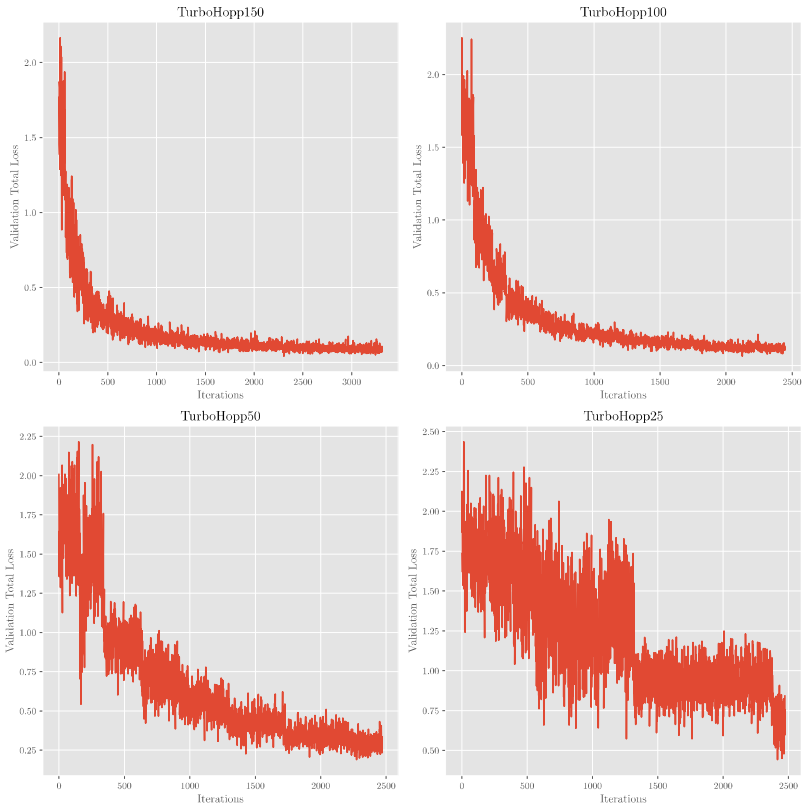

Figure 6: Training curves of total validation loss for different step size variants. We train a number of consistency models to empirically determine the optimal tradeoff between step size fidelity and speed. Turbohopp25 had low validity and proved to be too unstable for consistent generations.t

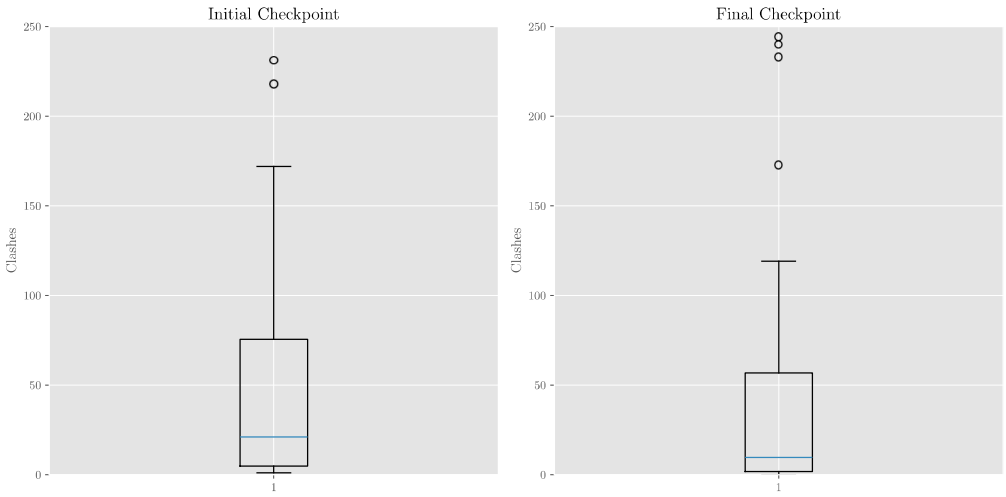

Figure 7: Plot of clashes before and after finetuning with a reward function mentioned in the main text. The initial Turbohopp-100 model was a RLCM finetuned model for connectivity, synthesizability, and QED score. Notice the shift toward smaller number of clashes during training. However, we believe that further iteration of the reward function will lead to more effective finetuning but we leave this to future work.

D Sample Molecules

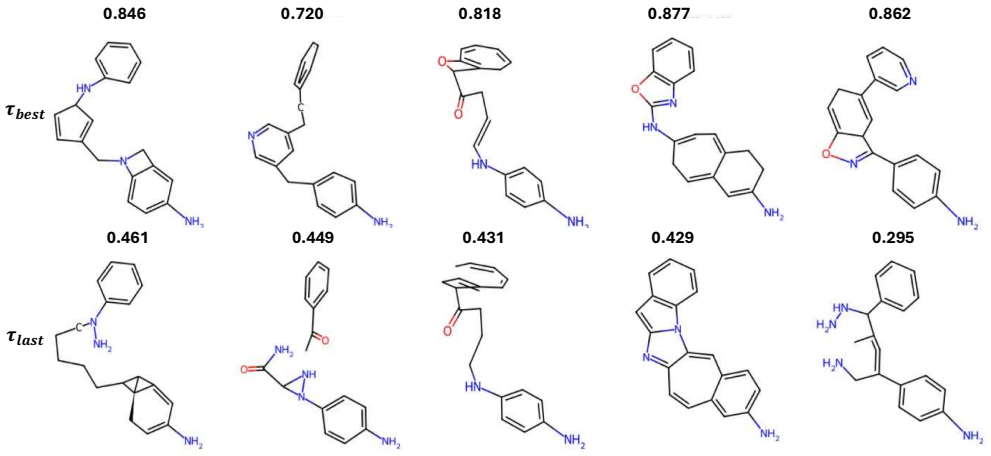

Figure 8: Examples of samples generated for PDB 6QQW with drug-likeliness. 1st row samples collected during multi-step phase with best scores and 2nd row indicates samples from the final step. Connectivity and overall metrics increased when we adopted custom score-based sampling.

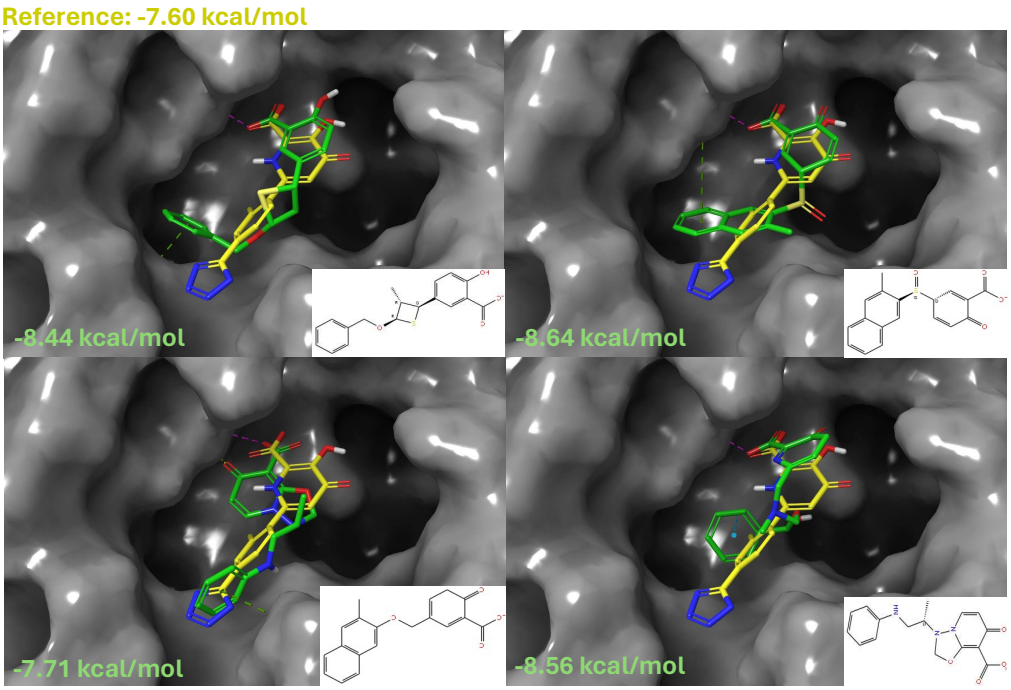

Figure 9: Scaffolds generated for PDB ID 6E6W by TurboHopp-100. Reference molecule is in yellow, while generated molecules are in green. Functional groups are the carboxylic acid and hydroxyl groups in the upper corner. Dotted lines in refer to ligand-receptor intermolecular interactions: green, blue, yellow, purple being pi-cation, pi-pi stacking, hydrogen bonds, halogen bonds respectively. Compared to reference molecule, generated molecules had new interactions and higher binding affinity, while maintaining similar binding pose.

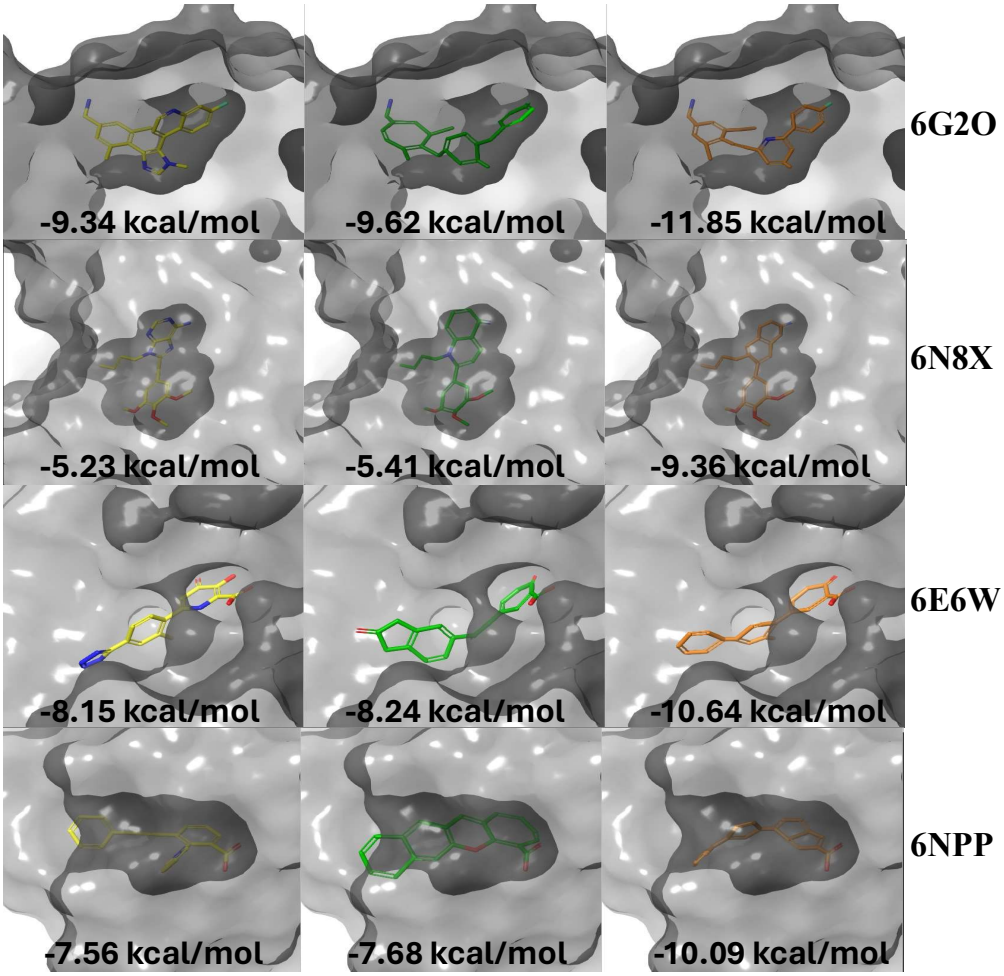

Figure 10: Comparison of Reference molecule(Yellow), and molecules generated by TurboHopp(Green) and TurboHopp-RL(Orange). Notice that Turbohopp and TurboHopp-RL generate molecules that have higher binding affinity with the protein.

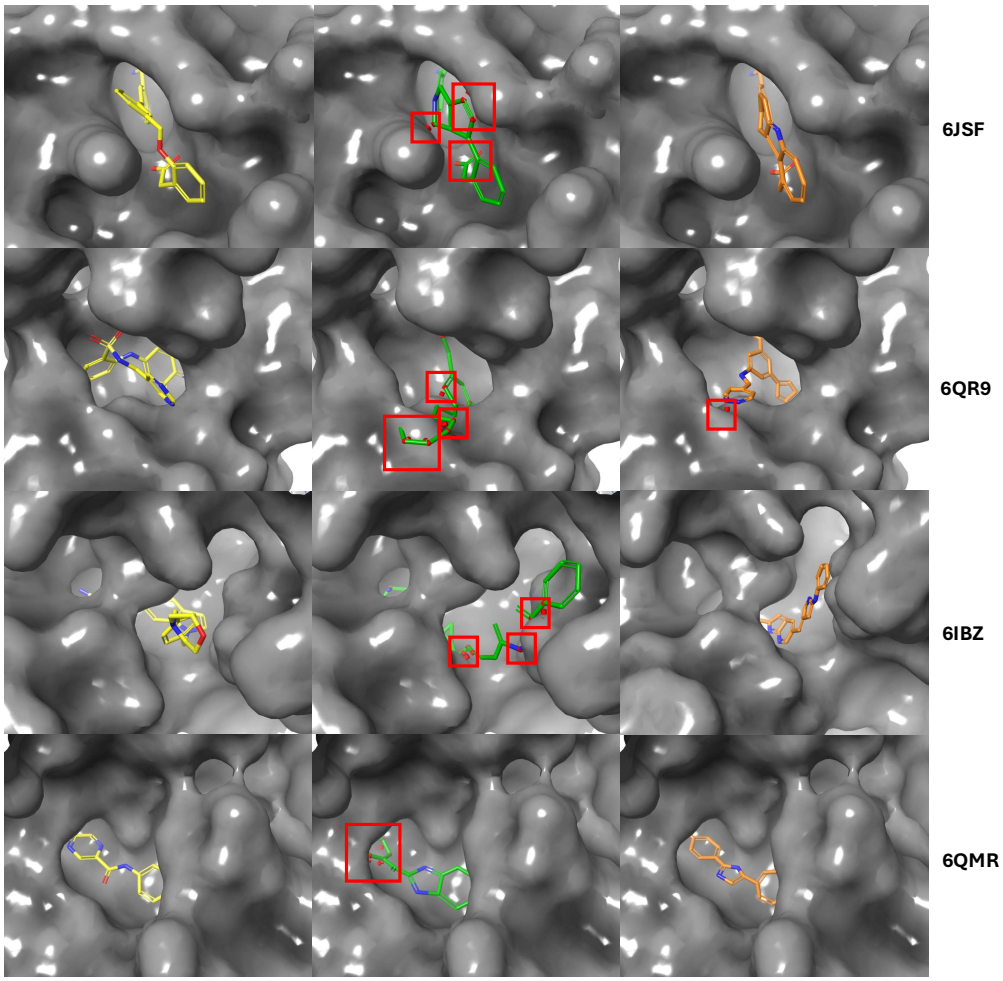

Figure 11: Comparison of Reference molecule(Yellow), and molecules generated by TurboHopp(Green) and TurboHopp-RL(Orange). Red box indicates collision points with protein atoms. TurboHopp-RL generates molecules that has less clashes with the protein.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We have claimed our model's performance and contributions based on empirical findings and experimental results. Our model can be easily adapted to similar frameworks using diffusion models for drug discovery.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We have added limitations and possible further research studies in the conclusion section.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our focus remains on empirical findings and application, bypassing the need for formal theoretical frameworks or validations.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We will share our code/checkpoints upon acceptance. In the paper, we provide the algorithms for training model, RL, and inference. We also share training curves for model training and RL as well as hyperparameters required to reproduce our results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We provide the data we used to train our model, and upon code release, we will specify preprocessing code ,modifications, as well as environment for reproduction.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: We specify the hyperparameters in the appendix. Further, we use a standardized dataset.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: All claims of statistical significance were justified through a number of trials. Multiple seeds are not needed for RLCM since we are applying a method previously shown to work.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We specify the compute resources required in the appendix.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: Yes we read the ethics guidlines and are abiding by them.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Please check the conclusion section of the paper.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: Models are not yet going to be released and we will take into account proper safeguards when releasing models with high risk for misuse.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All terms of use and licences are respected.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: No new assets are released

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: We do not do any crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: We do not do crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.