

Supplementary Text for Multi-Objective De Novo Drug Design with Conditional Graph Generative Model

Yibo Li Liangren Zhang* Zhenming Liu*

Summary

1	Extracting structures from ChEMBL dataset	1
2	Extracting scaffolds from drug dataset	1
3	Training of activity models for JNK3 and GSK3β	2
4	Algorithm for $q_\alpha(r G)$	2
5	Additional proof 1	3
6	Additional proof 2	4

1 Extracting structures from ChEMBL dataset

Both graph based and SMILES based models are trained using structures extracted from ChEMBL. Molecule structures from ChEMBL[1] are first standardized using RDKit[2]. This process involves salt removal, molecule neutralization, removing isotopes, and conversion to canonical SMILES strings. We only keep molecules containing less than 50 heavy atoms and whose elements belong to the set $\{H, B, C, N, O, F, P, S, Cl, Br, I\}$. This result in a dataset containing 1.5 million molecules. 5-fold cross validation is performed during evaluation.

2 Extracting scaffolds from drug dataset

As discussed in the method part, the scaffold set S is extracted from the list of approved drugs in DrugBank[3]. Two type of structure is considered during extraction: (1) the Bemis-Murcko scaffolds, and (2) ring assemblies. The extraction process is performed using RDKit. Scaffolds with a molecular weight larger than 300 are removed. Fragments that are tautomer to each other are merged into a single entity, as they are unidentifiable during substructure matching

(since the matching algorithm in RDKit is design to ignore hydrogens). The resulted S contains a total of 1129 scaffold structures.

3 Training of activity models for JNK3 and GSK3 β

For bioactivity dataset, search is performed on ExCAPE-DB[4] to extract activity data for each target. We use the activity flag provided by the database to separate the active and inactive compounds. This results in 3334 active compounds and 300186 inactive compounds for GSK3 β , as well as 923 active compounds and 59412 inactive compounds for JNK3. For each dataset, 80% of data is randomly selected as training set, and the rest as test set. The RF models for each target are implemented using Scikit-learn[5] with the number of estimators (decision trees) set to 100. RDKit is used to calculate the ECFP6.

4 Algorithm for $q_\alpha(r|G)$

Algorithm 1 Sampling r from $q_\alpha(r|G)$ and get the likelihood value

```
1: procedure SAMPLE( $G, \alpha$ ) ▷  $G = (V, E)$ 
2:   Order atoms in  $V$ 
3:    $G_{current} \leftarrow (\emptyset, \emptyset)$ ,  $q \leftarrow 1$ ,  $v \leftarrow \text{null}$ ,  $v_{stack} \leftarrow \text{empty stack}$ 
4:   Mark all atoms and bonds in  $G$  unvisited
5:   while True do
6:     if  $G_{current}$  is empty then
7:        $v^* \leftarrow$  the atom in  $V$  with highest rank
8:       Sample  $choise$  from  $Bernoulli(\alpha)$ 
9:       if  $choise = 1$  then
10:         $v \leftarrow v^*$ ,  $q \leftarrow q \times \alpha$ 
11:      else
12:         $v \leftarrow$  randomly chosen atom from  $V/v^*$ ,  $q \leftarrow q \times \frac{1-\alpha}{|V|-1}$ 
13:      end if
14:      Mark  $v$  visited
15:       $t \leftarrow$  the corresponding transition
16:      Append  $(G_{current}, t)$  to  $r$ , set  $G_{current} = t(G_{current})$ 
17:    else
18:       $V_e \leftarrow$  the set of visited atoms with unvisited connection to  $v$ 
19:      if  $V_e$  is not empty then
20:        Sort  $V_e$  from newest visited to oldest visited
21:        for  $v_e$  in  $V_e$  do
22:           $e \leftarrow$  the bond connecting  $v_e$  and  $v$ 
23:          Mark  $e$  visited
24:           $t \leftarrow$  the corresponding transition
25:          Append  $(G_{current}, t)$  to  $r$ , set  $G_{current} = t(G_{current})$ 
26:        end for
```

```

27:  $V_n \leftarrow$  the set of unvisited neighbors of  $v$ 
28: if  $V_n$  is not empty then
29:    $v_{stack}.push(v)$ 
30:    $v' \leftarrow v$ ,  $v^* \leftarrow$  the atom in  $V_n$  with highest rank
31:   Sample  $choise$  from  $Bernoulli(\alpha)$ 
32:   if  $choise = 1$  then
33:      $v \leftarrow v^*$ ,  $q \leftarrow q \times \alpha$ 
34:   else
35:      $v \leftarrow$  randomly chosen atom in  $V_n/v^*$ ,  $q = q \times \frac{1-\alpha}{|V_n|-1}$ 
36:   end if
37:   Mark  $v$  and the bond connecting  $v$  and  $v'$  visited
38:    $t \leftarrow$  the corresponding transition
39:   Append  $(G_{current}, t)$  to  $r$ , set  $G_{current} = t(G_{current})$ 
40: else
41:   if  $v_{stack}$  is not empty then:
42:      $v \leftarrow v_{stack}.pop()$ 
43:     continue
44:   else
45:      $t \leftarrow$  termination action
46:     Append  $(G_{current}, t)$  to  $r$ 
47:     return  $r, q$ 
48:   end if
49: end if
50: end if
51: end while
52: end procedure

```

5 Additional proof 1

Here, we demonstrate that during the training of conditional generative model, the following objective is being minimized:

$$L(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{c} \sim p(\mathbf{c})} [D_{KL}(p(\mathbf{x}|\mathbf{c}) || q_{\boldsymbol{\theta}})] \quad (1)$$

We can rewrite $L(\boldsymbol{\theta})$ as follows:

$$\begin{aligned}
L(\boldsymbol{\theta}) &= \mathbb{E}_{\mathbf{c} \sim p(\mathbf{c})} [D_{KL}(p(\mathbf{x}|\mathbf{c}) || q_{\boldsymbol{\theta}})] \\
&= \mathbb{E}_{\mathbf{c} \sim p(\mathbf{c})} [\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}|\mathbf{c})} [\log p(\mathbf{x}|\mathbf{c}) - \log q_{\boldsymbol{\theta}}(\mathbf{x})]] \\
&= \mathbb{E}_{(\mathbf{c}, \mathbf{x}) \sim p} [\log p(\mathbf{x}|\mathbf{c})] - \mathbb{E}_{(\mathbf{c}, \mathbf{x}) \sim p} [\log q_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{c})]
\end{aligned} \quad (2)$$

$\mathbb{E}_{(\mathbf{c}, \mathbf{x}) \sim p} [\log p(\mathbf{x}|\mathbf{c})]$ is in fact a constant term with respect to $\boldsymbol{\theta}$. Therefore, it can be safely omitted from $L(\boldsymbol{\theta})$:

$$L(\boldsymbol{\theta}) = -\mathbb{E}_{(\mathbf{c}, \mathbf{x}) \sim p} [\log q_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{c})] \quad (3)$$

This value can be approximated using Monte Carlo sampling:

$$\hat{L}(\boldsymbol{\theta}) = -\frac{1}{N} \sum_{i=1}^N \log q_{\boldsymbol{\theta}}(\mathbf{x}_i | \mathbf{c}_i) \quad (4)$$

Where $(\mathbf{c}_1, \mathbf{x}_1), \dots, (\mathbf{c}_N, \mathbf{x}_N)$ are sampled from the data distribution $p(\mathbf{c}, \mathbf{x})$. It is easy to see that eq. 4 is exactly the negative log-likelihood (NLL) loss used in MLE based training of conditional generative models.

6 Additional proof 2

Here, we show that optimizing the objective function $G(\mathbf{x})$ in the REINVENT method is equivalent to optimizing the KL divergence $D_{KL}(q_{\boldsymbol{\theta}} || p(\mathbf{c}|\mathbf{x}))$. In fact, if we set the score function $\sigma S(\mathbf{x})$ to $\log p(\mathbf{c}|\mathbf{x})$, it can be proved that:

$$\nabla_{\boldsymbol{\theta}} D_{KL}(q_{\boldsymbol{\theta}} || p(\mathbf{x}|\mathbf{c})) = -\mathbb{E}_{\mathbf{x} \sim q_{\boldsymbol{\theta}}} [\nabla_{\boldsymbol{\theta}} G(\mathbf{x})] \quad (5)$$

We expand the term $\nabla_{\boldsymbol{\theta}} D_{KL}(q_{\boldsymbol{\theta}} || p(\mathbf{x}|\mathbf{c}))$ as follows:

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} D_{KL}(q_{\boldsymbol{\theta}} || p(\mathbf{x}|\mathbf{c})) &= \nabla_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{x} \sim q_{\boldsymbol{\theta}}} [\log q_{\boldsymbol{\theta}}(\mathbf{x}) - \log p(\mathbf{x}|\mathbf{c})] \\ &= \nabla_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{x} \sim q_{\boldsymbol{\theta}}} [\log q_{\boldsymbol{\theta}}(\mathbf{x}) - \log p(\mathbf{x}) - \log p(\mathbf{c}|\mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x} \sim q_{\boldsymbol{\theta}}} [(\log q_{\boldsymbol{\theta}}(\mathbf{x}) - \log p(\mathbf{x}) - \log p(\mathbf{c}|\mathbf{x})) \nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}}(\mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x} \sim q_{\boldsymbol{\theta}}} [\nabla_{\boldsymbol{\theta}} (\log q_{\boldsymbol{\theta}}(\mathbf{x}) - \log p(\mathbf{x}) - \log p(\mathbf{c}|\mathbf{x}))^2] \end{aligned} \quad (6)$$

Since we have $G(\mathbf{x}) = -(\log p(\mathbf{x}) + \log p(\mathbf{c}|\mathbf{x}) - \log q_{\boldsymbol{\theta}}(\mathbf{x}))^2$, we can get the following equivalence:

$$\nabla_{\boldsymbol{\theta}} D_{KL}(q_{\boldsymbol{\theta}} || p(\mathbf{x}|\mathbf{c})) = -\mathbb{E}_{\mathbf{x} \sim q_{\boldsymbol{\theta}}} [\nabla_{\boldsymbol{\theta}} G(\mathbf{x})] \quad (7)$$

References

- [1] Gaulton, A., Bellis, L.J., Bento, A.P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S., Michalovich, D., Al-Lazikani, B.: ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Res* **40**(D1), 1100–1107 (2011)
- [2] RDKit: Open Source Cheminformatics. <http://www.rdkit.org/>
- [3] Wishart, D.S., Knox, C., Guo, A.C., Shrivastava, S., Hassanali, M., Stothard, P., Chang, Z., Woolsey, J.: Drugbank: a comprehensive resource for in silico drug discovery and exploration. *Nucleic Acids Res* **34**(Database issue), 668–672 (2006)

- [4] Sun, J., Jeliaskova, N., Chupakhin, V., Golib-Dzib, J.-F., Engkvist, O., Carlsson, L., Wegner, J., Ceulemans, H., Georgiev, I., Jeliaskov, V.: Excape-db: an integrated large scale dataset facilitating big data analysis in chemogenomics. *J Cheminform* **9**(1), 17 (2017)
- [5] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V.: Scikit-learn: Machine learning in python. *J Mach Learn Res* **12**(Oct), 2825–2830 (2011)