Towards Integrating Structure-Based Design with Deep Generative Models

Morgan Thomas
Centre for Molecular Informatics
Department of Chemistry
Cambridge, CB2 1EW
United Kingdom
mct50@cam.ac.uk

Rob Smith
Sosei Heptares
Steinmetz Building, Granta Park
Cambridge, CB21 6DG
United Kingdom
rob.smith@soseieptares.com

Noel M. O'Boyle
Sosei Heptares
Steinmetz Building, Granta Park
Cambridge, CB21 6DG
United Kingdom
noel.o'boyle@soseiheptares.com

Chris De Graaf
Sosei Heptares
Steinmetz Building, Granta Park
Cambridge, CB21 6DG
United Kingdom
chris.degraaf@soseiheptares.com

Andreas Bender
Centre for Molecular Informatics
Department of Chemistry
Cambridge, CB2 1EW
United Kingdom
ab454@cam.ac.uk

Abstract

Deep generative models are by now able to devise both valid and novel chemistry in a *de novo* molecule generation setting, which could significantly accelerate the identification of bioactive compounds. Most current models, however, use ligand-based predictive methods, to guide molecule generation to a bioactive chemical space. This restricts their application to relatively data-rich targets, neglecting those where little data is available to sufficiently train a predictor. In this work, we now assess the ability of using molecular docking *via* Glide – a structure-based approach – as a scoring function to guide the deep generative model REINVENT, and compare model behaviour and performance to a ligand-based scoring function. We show that the model improves predicted ligand affinity beyond known active molecules. We also show that the structure-based approach learns to satisfy crucial residue interactions (information only relevant when utilizing protein structure). This approach has potential application where *de novo* molecule generation either has no prior ligand knowledge available (early hit finding), or should not be biased by it (novelty-focused).

1 Introduction

Deep generative models (DGMs) are a recent and promising class of *de novo* molecule generation algorithms, which utilise advances in deep neural networks. They are able to generate valid chemical structures by either learning from a dataset of example molecules, or learning appropriate actions to take given a set of symbolic rules. Although these models vary greatly in method [1], *de novo*

Machine Learning for Molecules Workshop at NeurIPS 2020. https://ml4molecules.github.io

molecule generation typically has a single goal, to generate molecules towards a desired property space - one such crucial requirement in drug design is bioactivity. In order to achieve this property, most DGMs use ligand-based approaches. Such approaches include using known bioactive molecules as training data to bias generation towards a similar property space[2–4], or using machine learning (ML) models trained on known bioactive molecules to predict *de novo* molecule bioactivity - the DGM can then be optimized to maximize this value e.g. using reinforcement learning[5–8], Bayesian optimization[9] or particle swarm optimization[10].

However, the use of ligand-based ML models comes with its limitations. Namely, ligand-based ML models are restricted by their applicability domain i.e. they perform well on 'in distribution' data, but struggle to extrapolate to 'out of distribution' data, which is often poorly reflected in model validation[11, 12]. This means that ML models will score molecules similar to those observed in the training data more accurately[13]. Incorporation with DGMs biases *de novo* molecule generation to a similar property space as the scoring function training data, which is usually just one of many possible desirable property spaces. This contributes to the lack of diversity observed in deep generative models[14, 15]. Further, Renz et al.[16] observed that DGMs bias molecule generation towards ML parameters as well as training data *via* a set of ML controls. This bias limits the underlying principle of exploring novel chemical space, which is favourable for finding optimal space with respect to other required properties (such as pharmacokinetic profile) or for avoiding competing intellectual property. This lack of novelty has been commented on in the literature[17]. Hence, the use of ligand-based approaches in DGM scoring functions restricts practical use cases of DGMs.

This work We explored the idea that structure-based scoring functions, in a practical implementation using docking, may mitigate some of the limitations observed with ligand-based scoring functions. In concrete terms, we utilized the REINVENT[5] algorithm and optimized *de novo* molecules to minimize the docking score returned by Glide[18]. Further, to understand the differences between ligand- and structure-based scoring functions, we compare the resulting *de novo* molecules to those generated by an equivalent DGM optimized to maximize the predicted probability of activity *via* a support vector machine (SVM). As a case study, we chose affinity for Dopamine Receptor D2 (DRD2). This receptor has a wealth of associated ligand bioactivity data allowing retrospective validation and has been commonly used in deep generative model publications[5, 9, 19, 8, 15]. To the best of our knowledge, neither this particular approach nor such a comparison has been published in the literature.

2 Methods

2.1 Datasets

The training data used to train the DGM used in this work was extracted from the ZINC database[20] by following the curation workflow described by MOSES[21]. We also removed any known DRD2 actives. This resulted in a database of 3,579,885 canonical, non-isomeric SMILES. As for the set of bioactive compounds, we extracted molecules with known DRD2 bioactivity from ExCAPE-DB[22]. This resulted in 4,613 active and 343,075 inactive molecules against human DRD2. Lastly, for use as a reference baseline, 10,000 random molecules were extracted from ChEMBL26[23].

2.2 REINVENT

We used the previously published REINVENT framework[5] as a proof of concept deep generative model. The ZINC subset was used to train the Prior network for a total of 5 epochs with a batch size of 128 using the Adam optimizer[24] with a learning rate of 0.001. The Agent was then trained for 3,000 steps using a batch size of 64 and the Adam optimizer with a learning rate of 0.0005 and a value for the scalar coefficient (σ) of 60. All neural network training was conducted on an NVIDIA RTX 2080 Ti GPU.

2.3 Scoring function and retrospective performance

Ligand-based We reused the SVM model previously published by Olivecrona et al.[5], where the authors trained it on 7,218 active and 100,000 inactive DRD2 molecules, also extracted from

ExCAPE-DB. The performance was evaluated on an undisclosed held-out test set. Resulting in an accuracy of 98%, precision of 97% and recall of 82%.

Structure-based We used the DRD2 X-ray crystal structure 6CM4 from the PDB and prepared it using the Schrodinger Protein Preparation Wizard[25]. A grid was defined using the centroid of the co-crystallised ligand (Risperidone) as the centre. Before docking, molecules were prepared using LigPrep[26], enumerating unspecified stereocentres, tautomers and ionization states (up to 8 variants per molecule). All respective variants were then docked using Glide standard precision (GlideScore-SP) with default settings. The lowest (best) docking score of all variants was used as the resulting value for respective molecules. To make this task more computationally tractable, we implemented a script that parallelized the docking protocol across a compute cluster using the python library Dask[27]. Using between 36 and 50 CPUs, the wall time required for 3,000 iterations was approximately 7 days. All known DRD2 active molecules and a random subset of 10,000 DRD2 inactive molecules were docked following the protocol described. A docking score classification threshold of -8.5 led to an accuracy of 74%, precision of 82% and recall of 12%. This example threshold was chosen so as to minimize false-positives among positive predictions.

3 Results and discussion

3.1 Optimization of docking score

We first investigated whether the Agent was able to optimize the respective properties evaluated by the two scoring functions. We were able to reproduce the ability to maximize predicted probability of DRD2 activity by the 'SVM-Agent' - as reported previously[5]. We next evaluated the optimization of DRD2 docking score by the 'Glide-Agent', which is shown in Figure 1a. Put into the context of reference datasets (Figure 1b), the trained Glide-Agent shows enriched docking score ($\mu = -8.25, \sigma = 0.90$) compared to Prior molecules ($\mu = -6.23, \sigma = 1.03$) and even known DRD2 actives ($\mu = -7.34, \sigma = 1.04$). Based on the assumption that a docking score of -8.5 corresponds to 82% precision (section 2.3), sampled *de novo* molecules have a hit rate of 32.72%, 6.84% and 0.67% for the Glide-, SVM-Agent and Prior, respectively. In other words, the Glide-Agent shows 48-fold hit rate enrichment over the Prior.

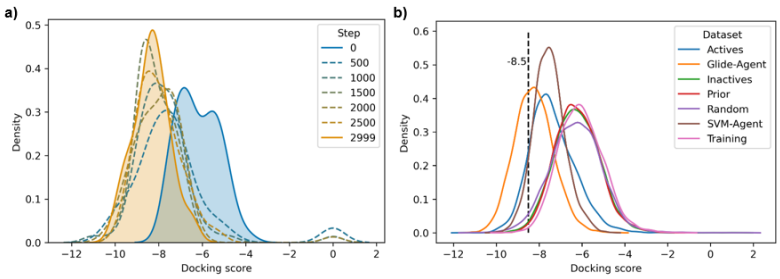

Figure 1: **(a)** Optimization of docking score distribution over training steps by the Glide-Agent and **(b)** comparison of docking score distribution of trained Agent *de novo* molecules and reference datasets.

3.2 Comparison to known DRD2 active molecules

To evaluate the *relevance* of *de novo* chemistry, we next assessed recovery of known DRD2 active molecules *via* probability of recovery (Table 1). Firstly, the Prior has an inherent bias to generative inactive molecules due to the removal of known actives from the training data. The Glide-Agent shifts this bias towards active molecules 68-fold and the SVM-Agent 3,940-fold, however, this is predominantly attributable to the SVM-Agents ability to avoid recovering known inactive molecules. While the probability of recovering known active molecules is comparable between the Glide- and SVM-Agents ($21.54\text{x}10^{-6}$ vs $22.68\text{x}10^{-6}$, respectively). In summary, both Agents have the ability to recover known DRD2 active molecules, albeit with different underlying behaviour.

Table 1: Probability of recovering known DRD2 active or inactive molecules, averaged across three sample sizes of 10,000, 100,000 and one million structures.

Origin of dataset	Probability of generating active molecule (x10^{-6})	Probability of generating inactive molecule (x10^{-3})	Activity bias (fold change from Prior)
Prior	1.17	3.60	0.0003(1)
Glide-Agent	21.54	0.98	0.0221 (68)
SVM-Agent	22.68	0.02	1.2837 (3,940)

3.3 Chemical Space

To aid visualization of the chemical differences between *de novo* molecules, we used Uniform Manifold Approximation and Projection (UMAP)[28] to project the molecules onto two dimensions; either using Morgan fingerprints (Figure 2a) or physicochemical descriptors (Figure 2b). It can be seen in Figure 2a that the Glide- and SVM-Agent *de novo* molecules occupy different regions of chemical space, neither of which display a perfect overlap with known DRD2 active molecules. In-fact, not one of the 10,000 *de novo* molecules sampled from either Agent is co-generated by both, highlighting the complementarity between these two approaches. While Figure 2b shows that the Glide-Agent *de novo* molecules occupy more diverse areas of physicochemical space than known actives and SVM-Agent *de novo* molecules. This indicates that the Glide-Agent is not biased towards the physicochemical properties of known bioactive molecules.

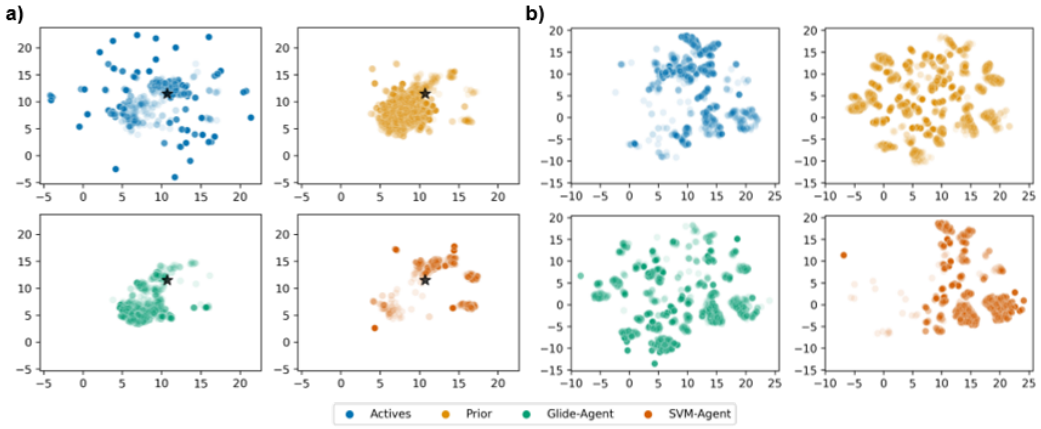

Figure 2: **(a)** UMAP of Morgan fingerprints based on a radius of 2 (equiv. ECFP4), and **(b)** of 13 physicochemical descriptors (all calculated using RDKit[29]).

3.4 Protein-ligand interactions

The highly conserved aspartic acid residue D^{3x32} provides a crucial interaction point required for bioactivity among aminergic receptors [30, 31]. By inspecting docked poses (Figure 3a-d) we hence observed more frequent satisfaction of D^{3x32} interactions by the Glide-Agent *de novo* molecules (Figure 3c). This was corroborated (Figure 3e) by using Structure Interaction Fingerprints (SIFts)[32] to analyze the ratio of molecules satisfying D^{3x32} interactions according to dataset. This shows that the Glide-Agent *de novo* molecules more often also form a hydrogen bond with D^{3x32} in DRD2. This could be an empirical explanation of how the Glide-Agent learns to optimize docking score.

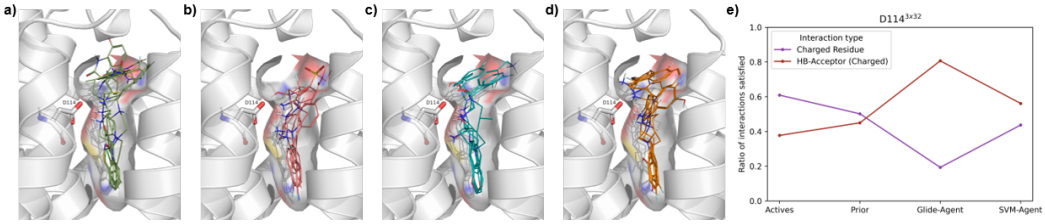

Figure 3: Docked pose of the centroids of the five best scoring clusters for **(a)** active, **(b)** Prior, **(c)** Glide- and **(d)** SVM-Agent molecules. **(e)** Ratio of D114^{3x32} interactions satisfied by *de novo* molecules according to dataset.

4 Conclusion

In this work we integrated generative molecular *de novo* design with structure-based molecular docking and compared results to ligand-based scoring functions. We show on a commonly used benchmark dataset that this approach results in chemically sensible molecules, which are able to improve docking scores beyond that of known receptor ligands, while exhibiting increased physicochemical diversity compared to ligand-based scoring. The setup presented here facilitates the use of deep generative models in settings also where no ligand data is available, or novelty is of particular interest (provided an X-ray crystal structure or a suitable homology model is available). Further validation on a variety of protein targets is both required and ongoing. Future work is also intended to further investigate the impact of incorporating prior structural knowledge, such as particular water/residue interactions that can affect selectivity.

References

[1] Daniel C. Elton, Zois Boukouvalas, Mark D. Fuge, and Peter W. Chung. Deep learning for molecular design—a review of the state of the art. *Molecular Systems Design & Engineering*, 4:828–849, 2019.

[2] Marwin H.S. Segler, Thierry Kogej, Christian Tyrchan, and Mark P. Waller. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Central Science*, 4:120–131, 1 2018.

[3] Alex Zhavoronkov, Yan A. Ivanenkov, Alex Aliper, Mark S. Veselov, Vladimir A. Aladinskiy, Anastasiya V. Aladinskaya, Victor A. Terentiev, Daniil A. Polykovskiy, Maksim D. Kuznetsov, Arip Asadulaev, Yury Volkov, Artem Zholus, Rim R. Shayakhmetov, Alexander Zhebrak, Lidiya I. Minaeva, Bogdan A. Zagribelnyy, Lennart H. Lee, Richard Soll, David Madge, Li Xing, Tao Guo, and Alán Aspuru-Guzik. Deep learning enables rapid identification of potent ddr1 kinase inhibitors. *Nature Biotechnology*, 37:1038–1040, 9 2019.

[4] Oleksii Prykhodko, Simon Viet Johansson, Panagiotis Christos Kotsias, Josep Arús-Pous, Esben Jannik Bjerrum, Ola Engkvist, and Hongming Chen. A de novo molecular generation method using latent vector based generative adversarial network. *Journal of Cheminformatics*, 11:74, 12 2019.

[5] Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9, 9 2017.

[6] Mariya Popova, Olexandr Isayev, and Alexander Tropsha. Deep reinforcement learning for de novo drug design. *Science Advances*, 4, 7 2018.

[7] Sai Krishna Gottipati, Boris Sattarov, Sufeng Niu, Yashaswi Pathak, Haoran Wei, Shengchao Liu, Karam M. J. Thomas, Simon Blackburn, Connor W. Coley, Jian Tang, Sarath Chandar, and Yoshua Bengio. Learning to navigate the synthetically accessible chemical space using reinforcement learning. *arXiv*, 4 2020.

[8] Julien Horwood and Emmanuel Noutahi. Molecular design in synthetically accessible chemical space via deep reinforcement learning. *arXiv*, 2020.

[9] Wengong Jin, Kevin Yang, Regina Barzilay, and Tommi Jaakkola. Learning multimodal graph-to-graph translation for molecular optimization. *arXiv*, 12 2018.

[10] Robin Winter, Floriane Montanari, Andreas Steffen, Hans Briem, Frank Noé, and Djork Arné Clevert. Efficient multi-objective molecular optimization in a continuous latent space. *Chemical Science*, 10:8016–8024, 8 2019.

[11] Ann E. Cleves and Ajay N. Jain. Effects of inductive bias on computational evaluations of ligand-based modeling and on drug discovery. *Journal of Computer-Aided Molecular Design*, 22:147–159, 3 2008.

[12] Izhar Wallach and Abraham Heifets. Most ligand-based classification benchmarks reward memorization rather than generalization. *Journal of Chemical Information and Modeling*, 58:916–932, 5 2018.

[13] Robert P. Sheridan, Bradley P. Feuston, Vladimir N. Maiorov, and Simon K. Kearsley. Similarity to molecules in the training set is a good discriminator for prediction accuracy in qsar. *Journal of Chemical Information and Computer Sciences*, 44:1912–1928, 11 2004.

[14] Silvia Amabilino, Peter Pogány, Stephen D. Pickett, and Darren V. S. Green. Guidelines for recurrent neural network transfer learning-based molecular generation of focused libraries. *Journal of Chemical Information and Modeling*, 7 2020.

[15] Thomas Blaschke, Ola Engkvist, Jürgen Bajorath, and Hongming Chen. Memory-assisted reinforcement learning for diverse molecular de novo design. *ChemRxiv*, 7 2020.

[16] Philipp Renz, Van Rompaey, Kurt Wegner, Sepp Hochreiter, Gunter Klambauer, Dries Van Rompaey, Jorg Kurt Wegner, and G Unter Klambauer. On failure modes of molecule generators and optimizers. *ChemRxiv*, 4 2020.

[17] W. Patrick Walters and Mark Murcko. Assessing the impact of generative ai on medicinal chemistry. *Nature Biotechnology*, 38:143–145, 1 2020.

[18] Richard A. Friesner, Jay L. Banks, Robert B. Murphy, Thomas A. Halgren, Jasna J. Klicic, Daniel T. Mainz, Matthew P. Repasky, Eric H. Knoll, Mee Shelley, Jason K. Perry, David E. Shaw, Perry Francis, and Peter S. Shenkin. Glide: A new approach for rapid, accurate docking and scoring. 1. method and assessment of docking accuracy. *Journal of Medicinal Chemistry*, 47:1739–1749, 3 2004.

[19] Panagiotis-Christos Kotsias, Josep Arús-Pous, Hongming Chen, Ola Engkvist, Christian Tyrchan, and Esben Jannik Bjerrum. Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks. *Nature Machine Intelligence*, 2:254–265, 5 2020.

[20] Teague Sterling and John J. Irwin. Zinc 15 - ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55:2324–2337, 11 2015.

[21] Daniil Polykovskiy, Alexander Zhebrak, Benjamin Sanchez-Lengeling, Sergey Golovanov, Oktai Tatanov, Stanislav Belyaev, Rauf Kurbanov, Aleksey Artamonov, Vladimir Aladinskiy, Mark Veselov, Artur Kadurin, Simon Johansson, Hongming Chen, Sergey Nikolenko, Alán Aspuru-Guzik, and Alex Zhavoronkov. Molecular sets (moses): A benchmarking platform for molecular generation models. *arXiv*, 2019.

[22] Jiangming Sun, Nina Jeliazkova, Vladimir Chupakhin, Jose-Felipe Golib-Dzib, Ola Engkvist, Lars Carlsson, Jörg Wegner, Hugo Ceulemans, Ivan Georgiev, Vedrin Jeliazkov, Nikolay Kochev, Thomas J. Ashby, and Hongming Chen. Excape-db: an integrated large scale dataset facilitating big data analysis in chemogenomics. *Journal of Cheminformatics*, 9:17, 12 2017.

[23] David Mendez, Anna Gaulton, A Patrícia Bento, Jon Chambers, Marleen De Veij, María Paula MagariñosMagari, Juan F Mosquera, Prudence Mutowo, Michał Nowotka, María Gordillo-Marã, Fiona Hunter, Laura Junco, Grace Mugumbate, Milagros Rodriguez-Lopez, Francis Atkinson, Nicolas Bosc, Chris J Radoux, Aldo Segura-Cabrera, Anne Hersey, and Andrew R Leach. Chembl: towards direct deposition of bioassay data. *Nucleic Acids Research*, 47, 2019.

[24] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. International Conference on Learning Representations, ICLR, 12 2015.

[25] Schrödinger Release 2019-4. Protein preparation wizard.

[26] Schrödinger Release 2019-4. Ligprep.

[27] Dask Development Team. Dask: Library for dynamic task scheduling, 2016.

[28] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv*, 2 2018.

[29] RDKit. Open-source cheminformatics.

[30] Albert J. Kooistra, Henry F. Vischer, Daniel McNaught-Flores, Rob Leurs, Iwan J.P. De Esch, and Chris De Graaf. Function-specific virtual screening for gpcr ligands using a combined scoring method. *Scientific Reports*, 6:1–21, 6 2016.

[31] Márton Vass, Albert J Kooistra, Tina Ritschel, Rob Leurs, Iwan Jp De Esch, and Chris De Graaf. Molecular interaction fingerprint approaches for gpcr drug discovery. *Current Opinion in Pharmacology*, 30:59–68, 2016.

[32] Zhan Deng, Claudio Chuaqui, and Juswinder Singh. Structural interaction fingerprint (sift): A novel method for analyzing three-dimensional protein-ligand binding interactions. *Journal of Medicinal Chemistry*, 47:337–344, 1 2004.