

---

# Molecular representation learning with language models and domain-relevant auxiliary tasks

---

**Benedek Fabian**

**Thomas Edlich**

**H el ena Gaspar**

**Marwin Segler**

**Joshua Meyers**

**Marco Fiscato**

**Mohamed Ahmed**

**BenevolentAI**

4-8 Maple St, Bloomsbury, London W1T 5HD

<firstname.lastname>@benevolent.ai

## Abstract

We apply a Transformer architecture, specifically BERT, to learn flexible and high quality molecular representations for drug discovery problems, and study the impact of using different combinations of self-supervised tasks for pre-training. Our results on established Virtual Screening and QSAR benchmarks show that: i) The selection of appropriate self-supervised task(s) for pre-training has a significant impact on performance in subsequent downstream tasks such as Virtual Screening. ii) Using auxiliary tasks with more domain relevance for Chemistry, such as learning to predict calculated molecular properties, increases the fidelity of our learnt representations. iii) Finally, we show that molecular representations learnt by our model ‘MOLBERT’ improve upon the current state of the art on the benchmark datasets.

## 1 Introduction

Molecular representations underpin predictive, generative and analytical tasks in drug discovery [1]. The choice of a suitable representation can drastically impact the efficiency of discovering a novel drug candidate. For instance, applications such as *Virtual Screening* and *Quantitative Structure-Activity-Relationship (QSAR)* modeling rely on the availability of effective molecular representations [2].

Language models have been applied to text-based molecular representations such as Simplified Molecular Input Line Entry System (SMILES) [3]. They show impressive performance across a range of domain applications including molecular property [4, 5] and reaction prediction problems [6, 7], as well as generative tasks [8]. Numerous strategies have been explored to encourage learning of high quality representations with language models including input reconstruction [9–11], whereby a model learns to predict masked or corrupted tokens; and input translation [5, 12], where the goal is to translate the input to another modality or representation. Further improvements have been made by incorporating calculated molecular properties into the representation, either by concatenating with the learnt representations [13], or through devising pre-training schemes [14]. Finally, a range of model architectures have been explored including autoencoders [5], RNNs [9] and transformers [10, 11].

Aside from the modal limitations of representing molecules as strings, a drawback to learning from text-based molecular representations is introduced by the ambiguity of linearizing the molecular graph [15]. In the case of SMILES, many valid sequences may represent the same molecule depending on the traversal path of the molecular graph. This ambiguity has led to the development of canonicalization algorithms [16, 17] which, while practical, introduce artifacts to linearized SMILES such that a language model may be distracted by the rules of canonicalization. Previous works have shown the benefits of learning using permutations of SMILES [18].

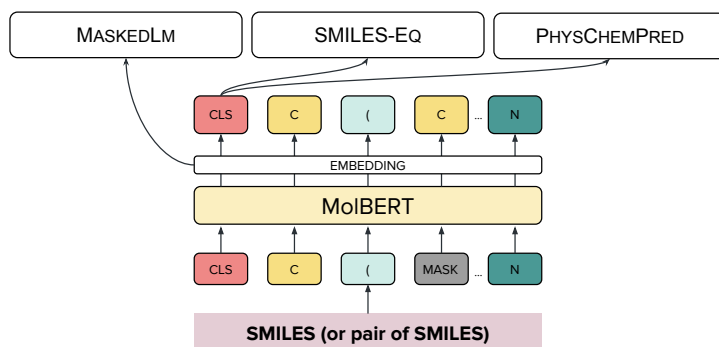


Figure 1: Diagram of MoLBERT illustrating the various auxiliary tasks utilized for pre-training.

In this work, we evaluate the application of the widely used Bidirectional Encoder Representations from Transformers (BERT) [19] architecture for the generation of molecular representations. We explore the impact of employing a range of domain-relevant auxiliary tasks during pre-training and evaluate the produced learnt representations on downstream Virtual Screening and QSAR benchmarks. Code and pre-trained models are available at <https://github.com/BenevolentAI/MoLBERT>.

## 2 MoLBERT

MoLBERT, as depicted in Figure 1, is a bidirectional language model that uses the BERT architecture [19]. To understand the impact of pre-training with different domain-relevant tasks on downstream applications, we experiment with the following set of self-supervised tasks:

**Masked language modeling (MASKEDLM):** The canonical task proposed by BERT, whereby the model is trained to predict the true identity of masked tokens. The task is optimized using the cross-entropy loss between the sequence output and the masked tokens of the input.

**SMILES equivalence (SMILES-EQ):** Given an input of two SMILES where the second is, with equal probability, either randomly sampled from the training set or a synonymous permutation of the first SMILES, the task is to predict whether the two inputs represent the same molecule. This is a binary classification task, which may be traditionally solved by comparing the canonicalized molecular graphs of the inputs. It is optimized using the cross-entropy loss.

**Calculated molecular descriptor prediction (PHYSCHEMPRED):** Using RDKit [20] we are able to calculate a simple set of real-valued descriptors of chemical characteristics and composite models of physicochemical properties for each molecule (see Appendix C). The goal of this task is to predict the normalized [21] set of descriptors for each molecule. The task is optimized using the mean squared error over all predicted values.

The final loss is given by the arithmetic mean of all individual task losses.

### Fine-tuning

The representations learnt after pre-training MoLBERT may be applied to downstream tasks in several manners. Here we explore: i) using representations directly for similarity search, ii) training a downstream model; in this case, following Winter et al. [5], we use a Support Vector Machine (SVM), and finally iii), following [19], we add an explicit downstream task head in the form of a newly initialized network which is then optimized.

### Evaluation methodology

We evaluate MoLBERT on two downstream applications: *Virtual Screening* and *QSAR*. In Virtual Screening we are typically interested in selecting from a pool of candidate compounds the ones that best satisfy a property of interest, such as the likelihood of binding to a particular drug target. Alternatively, in QSAR we are interested in learning to predict a given molecular property, such as the binding affinity to a target of interest. We evaluate MoLBERT when used directly and when specialized for downstream applications using two established benchmarks:

**Virtual Screening:** we use the filtered virtual screening subset (version 1.2) from the RDKit benchmarking platform [22, 23]. The benchmark is made up of 69 datasets, each corresponding to an individual protein target. Each dataset consists of a small number of active molecules amongst a larger number of target specific decoys. The benchmark measures how well-suited molecular

			w/ permutation		w/o permutation	
MaskedLM	PhysChemPred	SMILES-Eq	AUROC	BEDROC20	AUROC	BEDROC20
✓	✓	✓	0.685 ± 0.069	0.246 ± 0.041	0.707 ± 0.059	0.280 ± 0.042
✓	✓	✗	<b>0.738 ± 0.060</b>	<b>0.323 ± 0.071</b>	<b>0.740 ± 0.066</b>	<b>0.322 ± 0.065</b>
✓	✗	✓	0.483 ± 0.092	0.092 ± 0.069	0.493 ± 0.068	0.108 ± 0.070
✗	✓	✓	0.476 ± 0.077	0.064 ± 0.034	0.514 ± 0.165	0.084 ± 0.014
✓	✗	✗	0.696 ± 0.058	0.283 ± 0.077	0.676 ± 0.060	0.250 ± 0.073
✗	✓	✗	0.719 ± 0.057	0.293 ± 0.071	0.716 ± 0.061	0.290 ± 0.076
✗	✗	✓	0.129 ± 0.067	0.005 ± 0.037	0.508 ± 0.068	0.048 ± 0.035

Table 1: The impact of pre-training MOLBERT with different auxiliary task combinations on the Virtual Screening benchmark performance. Best values are in bold, higher values are better.

	AUROC	BEDROC20
MOLBERT(100 epochs)	<b>0.743 ± 0.062</b>	<b>0.344 ± 0.062</b>
CDDD	0.725 ± 0.057	0.310 ± 0.080
RDKit descriptors	0.633 ± 0.027	0.217 ± 0.000
ECFC4	0.603 ± 0.056	0.170 ± 0.079

Table 2: Results for Virtual Screening using the RDKit benchmarking platform.

representations are to retrieving active compounds from the pool, given a fixed number of query molecules ( $n = 5$ ) – for details see Riniker and Landrum [22].

Retrieval is performed according to the cosine distance between MOLBERT embeddings. Following Riniker and Landrum [22], we report results using: i) the Area Under the Curve of the Receiver Operating Characteristic (AUROC) and ii) the Boltzmann-Enhanced Discrimination of ROC (BEDROC,  $\alpha = 20$ ) – a widely used early enrichment metric that assigns higher weight to the first  $\alpha\%$  of molecules retrieved from the pool according to the Boltzmann distribution [24].

**QSAR:** we use a subset of the MoleculeNet benchmark suite [25], which consists of datasets of varying size for a range of QSAR problems. Specifically we report regression results for the ESOL, FreeSolv and Lipophilicity datasets, as well as classification results for the BACE, BBBP and HIV datasets. Since MoleculeNet does not provide explicit training, validation and test folds for all datasets, we use folds provided by ChemBench [26] to present reproducible results.

### 3 Experimental Evaluation

In this section, we first present an ablation on MOLBERT using the different pre-training tasks using the Virtual Screening datasets, then report our results in the Virtual Screening and QSAR benchmarks.

**Pre-training dataset:** All models are pre-trained using the GuacaMol benchmark dataset [27] consisting of  $\sim 1.6$ M compounds curated from ChEMBL [28]. We use the training and validation splits which consist of 80% and 5% of the data, respectively.

**Models:** MOLBERT models are implemented using the Hugging Face transformers library [29], and use the BERT-Base architecture (12 attention heads, 12 layers, 768 dimensional hidden layer,  $\sim 85$ M parameters). We use the Adam optimizer [30] with a learning rate of  $3 \times 10^{-5}$ , and train for 20 epochs, except for the final model, which we train for 100 epochs. The average pre-training time for MOLBERT was  $\sim 40$  hours using 2 GPUs and 16 CPUs.

**Fine-tuning:** All experiments use the same base model, the best model from the ablation trained for 100 epochs. To ease reproducibility we use the SVM parameters published in Winter et al. [5] ( $C = 5.0$ , RBF kernel) and all fine-tuning networks are a single linear layer attached to the pooled output. To generate molecular embeddings we use the pooled output, except where MASKEDLM is the only task used for pre-training, in which case we use the average of the sequence output, since the pooled output has no dependent tasks.

MOLBERT is trained using a fixed vocabulary of 42 tokens and a sequence length of 128 characters. In line with [19], all tasks are trained by masking 15% of the tokenized input. To support the use of arbitrary length input SMILES at inference time, we use relative positional embeddings as described by Dai et al. [31].

#### Results

**Ablation study:** We first analyze the utility of using different combinations of tasks for pre-training. From Table 1, the main observations are that: i) The PHYSCHEMPRED task has the highest impact on

Regression Datasets: RMSE					
	RDKit (norm)	ECFC4	CDDD	MOLBERT	MOLBERT (finetune)
ESOL	0.687 ± 0.08	0.902 ± 0.06	0.567 ± 0.06	0.552 ± 0.07	<b>0.531 ± 0.04</b>
FreeSolv	1.671 ± 0.45	2.876 ± 0.38	1.456 ± 0.43	1.523 ± 0.66	<b>0.948 ± 0.33</b>
Lipop	0.738 ± 0.04	0.770 ± 0.03	0.669 ± 0.02	0.602 ± 0.01	<b>0.561 ± 0.03</b>
Classification Dataset: AUROC					
	RDKit (norm)	ECFC4	CDDD	MolBERT	MolBERT (finetune)
BACE	0.831 ± 0.00	0.845 ± 0.00	0.833 ± 0.00	0.849 ± 0.00	<b>0.866 ± 0.00</b>
BBBP	0.696 ± 0.00	0.678 ± 0.00	0.761 ± 0.00	0.750 ± 0.00	<b>0.762 ± 0.00</b>
HIV	0.708 ± 0.00	0.714 ± 0.00	0.753 ± 0.00	0.747 ± 0.00	<b>0.783 ± 0.00</b>

Table 3: QSAR results on regression and classification tasks from MoleculeNet. Best values are in bold.  $\pm$  indicates standard deviation over cross-validation splits. First four columns are generated using the SVM, while the last column refers to fine-tuning a new task head.

the performance metrics, with an average BEDROC20 of 0.292 when using the PHYSCHEMPRED task alone (with and without permutations) *versus* 0.266 for the MASKEDLM alone. ii) Although the best performing model is trained on both the PHYSCHEMPRED and MASKEDLM, the additive gain from the MASKEDLM task is relatively minor; +0.031 on average for the BEDROC20 metric. iii) The addition of the SMILES-EQ task slightly but consistently decreases performance.

Given the effectiveness of pre-training with the PHYSCHEMPRED task, we explored the impact of grouping the 200 calculated descriptors into disjoint related subsets. Table A5 shows that, although using ALL the descriptors achieves the best overall result, the SURFACE properties of a molecule provide a very competitive supervision task using only 25% of the descriptors.

**Virtual Screening:** Table 2 compares the performance of MOLBERT when trained for 100 epochs with the best performing auxiliary task combination, with three baseline methods. i) CDDD [5], a neural model that achieves the current state-of-the-art results for this benchmark. ii) The RDKit calculated physicochemical descriptors [20] used for the PHYSCHEMPRED task during pre-training. iii) Extended Connectivity Fingerprints with a diameter of 4 (ECFC4), one of the most commonly used descriptors in drug discovery. Our results show that MOLBERT outperforms all other descriptors both in terms of overall classification and early enrichment. A detailed breakdown of *per*-dataset results is given in Figure A1 for the BEDROC20 and AUROC, respectively. Finally, upon closer investigation as to the benefit of input permutation and calculated molecular property prediction, we observed that these strategies enable MOLBERT to organize the learnt embeddings. More concretely, pre-training with the PHYSCHEMPRED task and input permutation leads to models which on average assign a lower pairwise similarity to non-identical compounds; see Appendix B.

**QSAR:** To further evaluate the usefulness of MOLBERT embeddings, we build QSAR models for datasets from MoleculeNet [25] and include models built using CDDD [5], ECFC4 [32] and the normalized RDKit calculated physicochemical descriptors [20] as baselines. As described in Section 2, we train an SVM (implemented using `sklearn` [33]) using each molecular representation, and compare against fine-tuning MOLBERT.

From the results in Table 3, we see that neural models substantially outperform traditional molecular representations (RDKit and ECFC4). Furthermore, finetuned MOLBERT models achieve the best performance in all of the six benchmark datasets. We also observe that MOLBERT representations combined with a SVM outperform the other descriptors on three of the six tasks.

## 4 Conclusions

We have introduced MOLBERT, a language model for learning molecular embeddings using BERT. We investigated the impact of pre-training with domain-relevant auxiliary tasks and found that the choice of self supervision task significantly impacts the performance on downstream tasks. Nevertheless, with the right set of tasks MOLBERT achieves state-of-the-art performance on established *Virtual Screening* and *QSAR* benchmarks. We leave to future work the exploration of how to use MOLBERT for learning representations of other entities such as proteins [34–36], along with further developments in our learning strategies [37].

## References

- [1] Laurianne David, Amol Thakkar, Rocío Mercado, and Ola Engkvist. Molecular representations in AI-driven drug discovery: a review and practical guide. *Journal of Cheminformatics*, 12(1):1–22, 2020.
- [2] Jessica Vamathevan, Dominic Clark, Paul Czodrowski, Ian Dunham, Edgardo Ferran, George Lee, Bin Li, Anant Madabhushi, Parantu Shah, Michaela Spitzer, and Shanrong Zhao. Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18(6):463–477, 2019.
- [3] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- [4] Stanisław Jastrzębski, Damian Leśniak, and Wojciech Marian Czarnecki. Learning to SMILE (S). *arXiv preprint arXiv:1602.06289*, 2016.
- [5] Robin Winter, Floriane Montanari, Frank Noé, and Djork-Arné Clevert. Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations. *Chemical Science*, 10(6):1692–1701, 2019.
- [6] Bowen Liu, Bharath Ramsundar, Prasad Kawthekar, Jade Shi, Joseph Gomes, Quang Luu Nguyen, Stephen Ho, Jack Sloane, Paul Wender, and Vijay Pande. Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS Central Science*, 3(10):1103–1113, 2017.
- [7] Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A Hunter, Costas Bekas, and Alpha A Lee. Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction. *ACS Central Science*, 5(9):1572–1583, 2019.
- [8] Marwin HS Segler, Thierry Kogej, Christian Tyrchan, and Mark P Waller. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Central Science*, 4(1):120–131, 2018.
- [9] Xinhao Li and Denis Fourches. Inductive Transfer Learning for Molecular Activity Prediction: Next-Gen QSAR Models with MolPMoFiT. *Journal of Cheminformatics*, 12(27):1–15, 2020.
- [10] Łukasz Maziarka, Tomasz Danel, Sławomir Mucha, Krzysztof Rataj, Jacek Tabor, and Stanisław Jastrzębski. Molecule Attention Transformer. *arXiv preprint arXiv:2002.08264*, 2020.
- [11] Sheng Wang, Yuzhi Guo, Yuhong Wang, Hongmao Sun, and Junzhou Huang. Smiles-bert: Large scale unsupervised pre-training for molecular property prediction. In *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*, page 429–436, 2019. doi: 10.1145/3307339.3342186.
- [12] Paul Morris, Rachel St. Clair, William Edward Hahn, and Elan Barenholtz. Predicting binding from screening assays with transformer network embeddings. *Journal of Chemical Information and Modeling*, 60(9):4191–4199, 2020.
- [13] Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, Andrew Palmer, Volker Settels, Tommi Jaakkola, Klavs Jensen, and Regina Barzilay. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388, 2019.
- [14] Garrett B. Goh, Charles Siegel, Abhinav Vishnu, and Nathan Hodas. Using Rule-Based Labels for Weak Supervised Learning: A ChemNet for Transferable Chemical Property Prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 302–310, 2018. doi: 10.1145/3219819.3219838.
- [15] Milan Randić. On canonical numbering of atoms in a molecule and graph isomorphism. *Journal of Chemical Information and Computer Sciences*, 17(3):171–180, 1977.
- [16] B. Weisfeiler and A. A. Lehman. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsia*, 2(9), 1968.
- [17] Nadine Schneider, Roger A. Sayle, and Gregory A. Landrum. Get your atoms in order—an open-source implementation of a novel and robust molecular canonicalization algorithm. *Journal of Chemical Information and Modeling*, 55(10):2111–2120, 2015.
- [18] Esben Jannik Bjerrum and Boris Sattarov. Improving Chemical Autoencoder Latent Space and Molecular De Novo Generation Diversity with Heteroencoders. *Biomolecules*, 8(4):131, 2018.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

- [20] G. A. Landrum. RDKit: Open-Source Cheminformatics Software, 2020. URL <http://www.rdkit.org>.
- [21] Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, Andrew Palmer, Volker Settels, Tommi Jaakkola, Klavs Jensen, and Regina Barzilay. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388, 2019.
- [22] Sereina Riniker and Gregory A Landrum. Open-source platform to benchmark fingerprints for ligand-based virtual screening. *Journal of Cheminformatics*, 5(1):26, 2013.
- [23] Sereina Riniker, Nikolas Fechner, and Gregory A Landrum. Heterogeneous classifier fusion for ligand-based virtual screening: or, how decision making by committee can be a good thing. *Journal of Chemical Information and Modeling*, 53(11):2829–2836, 2013.
- [24] Jean-François Truchon and Christopher I Bayly. Evaluating virtual screening methods: good and bad metrics for the “early recognition” problem. *Journal of Chemical Information and Modeling*, 47(2): 488–508, 2007.
- [25] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530, 2018.
- [26] Shen Wanxiang. ChemBench: The molecule benchmarks and MolMapNet datasets, 2020. URL <https://github.com/shenwanxiang/ChemBench>.
- [27] Nathan Brown, Marco Fiscato, Marwin H.S. Segler, and Alain C. Vaucher. GuacaMol: Benchmarking Models for de Novo Molecular Design. *Journal of Chemical Information and Modeling*, 59(3):1096–1108, 2019.
- [28] Anna Gaulton, Anne Hersey, Michał Nowotka, A. Patrícia Bento, Jon Chambers, David Mendez, Prudence Mutowo, Francis Atkinson, Louisa J. Bellis, Elena Cibrián-Uhalte, Mark Davies, Nathan Dedman, Anneli Karlsson, María Paula Magariños, John P. Overington, George Papadatos, Ines Smit, and Andrew R. Leach. The ChEMBL database in 2017. *Nucleic Acids Research*, 45(D1):D945–D954, 2017.
- [29] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *arXiv preprint arXiv:1910.03771*, 2019.
- [30] Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [31] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
- [32] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010.
- [33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [34] Martin Simonovsky and Joshua Meyers. Deeplytough: Learning structural comparison of protein binding sites. *Journal of Chemical Information and Modeling*, 60(4):2356–2366, 2020.
- [35] Ethan C Alley, Grigory Khimulya, Surojit Biswas, Mohammed AlQuraishi, and George M Church. Unified rational protein engineering with sequence-based deep representation learning. *Nature Methods*, 16(12): 1315–1322, 2019.
- [36] Paul Kim, Robin Winter, and Djork-Arné Clevert. Deep protein-ligand binding prediction using unsupervised learned representations. *ChemRxiv*, 2020. doi: 10.26434/chemrxiv.11523117.v1.
- [37] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay S. Pande, and Jure Leskovec. Strategies for Pre-training Graph Neural Networks. *arXiv preprint arXiv:1905.12265*, 2019.
- [38] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

## Appendix

### A Virtual Screening benchmark: performance *per target*

Figure A1 shows results for MOLBERT and other baseline methods on the individual Virtual Screening datasets listed in [22]. The datasets are sorted based on the BEDROC20 enrichment metric. We observe that MOLBERT displays superior performance for 45 of 69 targets and performs competitively in all other cases.




Figure A1: BEDROC20 and AUROC performance for each of the target datasets in the Virtual Screening benchmark [22] sorted based on the BEDROC20 metric. We report results for the best performing MOLBERT model in our ablation study (see Table 1) trained for 100 epochs.

### B Auxiliary task induced biases in MOLBERT models

To understand the types of inductive biases introduced by the different auxiliary tasks used in pre-training, we explore the difference in behaviour of the resulting learnt representations.

Table A1 lists three well-studied compounds. We sample ten random permutations of each, along with 1000 randomly selected molecules from ChEMBL, and use this data to compare the distributions of pairwise cosine similarities of the representations learnt with each task combination in Table 1.

Since the benchmark is a nearest neighbor retrieval task, we aim to understand whether learning to finely disambiguate between molecules correlates with benchmark performance. We hypothesize that models that achieve high average pairwise similarity between permutations of the same molecule and a low average pairwise similarity between random molecules achieve improved benchmark performance.

Figure A2 gives the results of our analysis where we observe three broad behaviours: First, MOLBERT models that were pre-trained on task combinations including the physicochemical property prediction (PHYSCHEMPRED) task successfully assign high similarities to permutations of the same molecule, while assigning low average similarity between randomly selected ChEMBL compounds. This suggests that tasks that encourage this large margin property help models organize the embedding space in a more semantically meaningful manner.

The MOLBERT representations resulting from a sample of tasks with this characteristic are given in Table A2. For each task we plot t-SNE [38] projections of MOLBERT representations for the 1000 randomly selected molecules from ChEMBL in gray and each of the three molecules listed in Table A1 along with their ten SMILES permutations. We use the sklearn implementation of t-SNE [33] with parameters: *perplexity* = 30, *early\_exaggeration* = 12, *learning\_rate* = 200.

Second, two task combinations MASKEDLM and MASKEDLM + PERMUTE lead to models that assign very low similarities to both the permutations of the same molecule, and to other molecules. This suggests that these

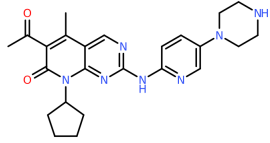
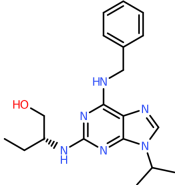
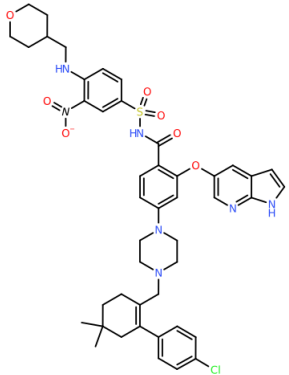
	Seed Molecule	Permutations
Palbociclib	 <p><chem>C1(NC2=NC=C(N3CCNCC3)C=C2)=NC=C2C(C)=C(C(C)=O)N(C3CCCC3)C2=N1</chem>  <chem>C1(NC2=NC=C(N3CCNCC3)C=C2)=NC=C2C(C)=C(C(C)=O)C(=O)N(C3CCCC3)C2=N1</chem></p>	<p><chem>C1(NC2=NC=C(N3CCNCC3)C=C2)=NC=C2C(C)=C(C(C)=O)N(C3CCCC3)C2=N1</chem>  <chem>C1CCC(N2C3=NC(NC4=NC=C(N5CCNCC5)C=C4)=NC=C3C(C)=C(C(=O)C)C2=O)C1</chem>  <chem>C1(C)=C(C(C)=O)C(=O)N(C2CCCC2)C2=NC(NC3=NC=C(N4CCNCC4)C=C3)=NC=C12</chem>  <chem>C1C(N2C3=NC(NC4=NC=C(N5CCNCC5)C=C4)=NC=C3C(C)=C(C(C)=O)C)C2=O)CCC1</chem>  <chem>N1=C(NC2=NC=C3C(C)=C(C(=O)C)C(=O)N(C4CCCC4)C3=N2)C=C(C(N2CCNCC2)=C1</chem>  <chem>C1=CC(N2CCNCC2)=CN=C1NC1=NC=C2C(C)=C(C(=O)C)C(=O)N(C3CCCC3)C2=N1</chem>  <chem>C1NCCN(C2=CN=C(NC3=NC=C4C(C)=C(C(C)=O)C)C(=O)N(C5CCCC5)C4=N3)C=C2)C1</chem>  <chem>C1CCC(N2C3=NC(NC4=NC=C(N5CCNCC5)C=C4)=NC=C3C(C)=C(C(C)=O)C2=O)C1</chem>  <chem>C1CN(C2=CN=C(NC3=NC=C4C(=N3)N(C3CCCC3)C(=O)C(C(=O)C=C4)C=C2)CCN1</chem>  <chem>N1(C2=CN=C(NC3=NC=C4C(=N3)N(C3CCCC3)C(=O)C(C(C)=O)C=C4)C=C2)CCNCC1</chem></p>
Seliciclib	 <p><chem>CC[C@H](CO)NC1=NC(NCC2=CC=CC=C2)=C2C(=N1)-N(C(C)C)C=N2</chem></p>	<p><chem>C(NC1=C2C(=NC(N[C@H](CC)CO)=N1)N(C(C)C)C=N2)C1=CC=CC=C1</chem>  <chem>C(NC1=C2N=CN(C(C)C)C2=NC(N[C@H](CO)CC)=N1)C1=CC=CC=C1</chem>  <chem>C1=CC=C(CNC2=C3C(=NC(N[C@@H](CO)CC)=N2)N(C(C)C)C=N3)C=C1</chem>  <chem>CC(C)N1C=NC2=C(NCC3=CC=CC=C3)N=C(N[C@@H](CO)CC)N=C12</chem>  <chem>C1=CC(CNC2=C3N=CN(C(C)C)C3=NC(N[C@H](CC)CO)=N2)=CC=C1</chem>  <chem>N1(C(C)C)C2=NC(N[C@H](CO)CC)=NC(NCC3=CC=CC=C3)=C2N=C1</chem>  <chem>CC(C)N1C2=NC(N[C@@H](CO)CC)=NC(NCC3=CC=CC=C3)=C2N=C1</chem>  <chem>C12=NC(N[C@H](CC)CO)=NC(NCC3=CC=CC=C3)=C1N=CN2C(C)C</chem>  <chem>[C@H](CO)(CC)NC1=NC(NCC2=CC=CC=C2)=C2C(=N1)N(C(C)C)C=N2</chem>  <chem>C1=C(CNC2=C3N=CN(C(C)C)C3=NC(N[C@H](CO)CC)=N2)=CC=C1</chem></p>
Venetoclax	 <p><chem>CC1(C)CCC(CN2CCN(C3=CC(OC4=CN=C5C(=C4)C=CN5)=C(C(=O)NS(=O))(=O)C4=CC([N+](=O)[O-])=C(NCC5=CCOCC5)C=C4)C=C3)CC2)=C(C2=CC=C(C1)C=C2)C1</chem></p>	<p><chem>C1CN(C2=C(C3=CC=C(C1)C=C3)CC(C)(C)CC2)CCN1C1=CC(OC2=CN=C3C(=C2)-C=CN3)=C(CNS(C2=CC([N+](O-))=O)=C(NCC3CCOCC3)C=C2)(=O)=O)C=C1</chem>  <chem>C1=CC(C(=O)NS(=O))(=O)C2=CC([N+](=O)[O-])=C(NCC3CCOCC3)C=C2)=C(OC2=CN=C3C(=C2)C=CN3)C=C1N1CCN(C2=C(C3=CC=C(C1)C=C3)CC(C)(C)CC2)CC1</chem>  <chem>[N+](=O)(C1=C(NCC2CCOCC2)C=CC(S(NC(=O)C2=CC=C(N3CCN(CC4=C(C5=CC=C(C1)C=C5)CC(C)(C)CC4)CC3)C=C2OC2=CN=C3C(=C2)C=CN3)(=O)=O)C1)[O-]</chem>  <chem>C1(C2=CC=C(C1)C=C2)=C(CN2CCN(C3=CC(OC4=CN=C5NC=CC5=C4)=C(CNS(=O)(C4=CC([N+](O-))=O)=C(NCC5CCOCC5)C=C4)=O)C=C3)CC2)CC(C)(C)C1</chem>  <chem>[O-][N+](=O)C1=C(NCC2CCOCC2)C=CC(S(=O)(NC(=O)C2=CC=C(N3CCN(CC4=C(C5=CC=C(C1)C=C5)CC(C)(C)CC4)CC3)C=C2OC2=CN=C3NC=CC3=C2)=O)C1</chem>  <chem>C1(OC2=CN=C3C(=C2)C=CN3)=CC(N2CCN(CC3=C(C4=CC=C(C1)C=C4)CC(C)(C)C3)CC2)=CC=C1C(NS(C1=CC([N+](=O)[O-])=O)=C(NCC2CCOCC2)C=C1)(=O)=O</chem>  <chem>C1(C2=C(CN3CCN(C4=CC(OC5=CN=C6C(=C5)C=CN6)=C(C(=O)NS(=O)(C5=CC([N+](=O)[O-])=O)=C(NCC6CCOCC6)C=C5)C=C4)CC3)CCC(C)(C)C2)=CC=C(C1)C=C1</chem>  <chem>C1(N2CCN(CC3=C(C4=CC=C(C1)C=C4)CC(C)(C)CC3)CC2)=CC(OC2=CN=C3C(=C2)C=CN3)=C(C(=O)NS(C2=CC([N+](O-))=O)=C(NCC3CCOCC3)C=C2)(=O)=O)C=C1</chem>  <chem>C1C(NC2=C([N+](=O)[O-])C=C(S(=O)(NC(=O)C3=CC=C(N4CCN(CC5=C(C6=CC=C(C1)C=C6)CC(C)(C)CC5)CC4)C=C3OC3=CN=C4NC=CC4=C3)=O)C=C2)-CCOC1</chem>  <chem>C12=CC(OC3=CC(N4CCN(CC5=C(C6=CC=C(C1)C=C6)CC(C)(C)CC5)CC4)=CC=C3-C(NS(=O))(=O)C3=CC([N+](O-))=O)=C(NCC4CCOCC4)C=C3)=O)CN=C1NC=C2</chem></p>

Table A1: Drugs and their color highlighted for the t-SNE plots in Tables A2, A3 and A4.

combinations encourage models to map inputs to a large representation space, but fail to structure this space in a semantically meaningful way – observed in the two sets of similarities being almost equal. See Table A3.

Finally, explicit encouragement of the model to learn to recognise permutations of identical SMILES in the form of the SMILES-EQ task does not seem to be sufficient to enable models to increase the average margin between representations of identical permutations and unrelated molecules. Table A4 shows that the SMILES-EQ + PERMUTE combination is not able to generate a structured representation space. However, when the MASKEDLM task is added, the embeddings of the different permutations are grouped together. This also results in an improved BEDROC from Table 1.






Figure A2: Average pairwise cosine similarity for each group (Seliciclib, Venetoclax, Palbociclib, ChEMBL). Models are sorted from left to right by decreasing BEDROC20 as shown in Table 1.

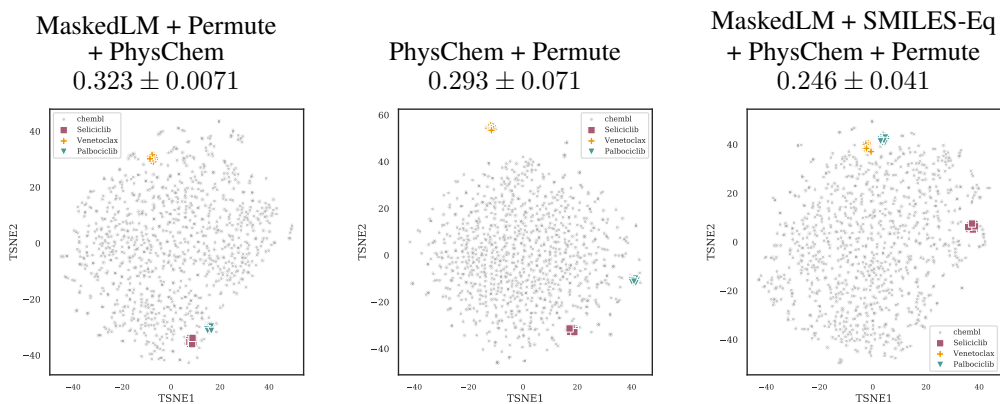


Table A2: t-SNE plots of molecules in Table A1. Examples of tasks which encourage disambiguation between permutations of SMILES and others. Numbers in heading denote the BEDROC performance from Table 1.

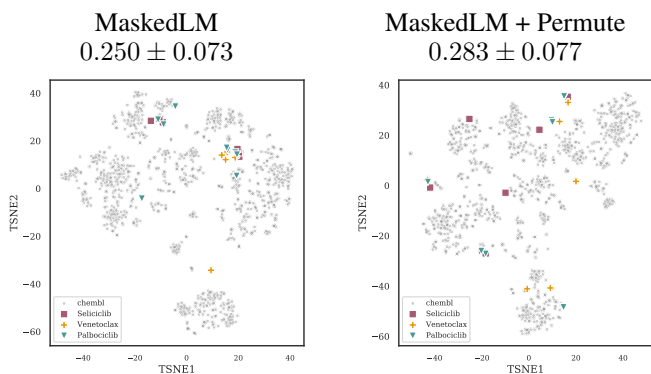


Table A3: t-SNE plots of molecules in Table A1. Examples of task combinations which don't encourage disambiguation between permutations of SMILES and others. Numbers in heading denote the BEDROC performance from Table 1.

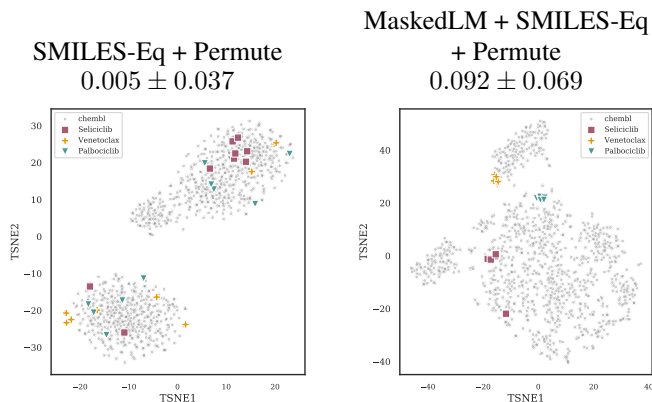


Table A4: t-SNE plots of molecules in Table A1 for two models trained with the SMILES-EQ task. It shows that even though this task is specifically aimed at teaching a model the equivalence of permuted SMILES, it alone is not sufficient to learn a structured representation space. Numbers in heading denote the BEDROC performance from Table 1.

	$n$	AUROC	BEDROC20
ALL	200	<b>0.738 ± 0.060</b>	<b>0.323 ± 0.071</b>
SURFACE	49	<b>0.738 ± 0.061</b>	0.310 ± 0.056
CHARGE	18	0.711 ± 0.063	0.270 ± 0.074
FRAGMENT	101	0.704 ± 0.065	0.277 ± 0.070
SIMPLE	8	0.695 ± 0.065	0.255 ± 0.060
GRAPH	19	0.693 ± 0.066	0.249 ± 0.067
ESTATE	25	0.676 ± 0.063	0.232 ± 0.070
DRUGLIKENESS	24	0.671 ± 0.060	0.214 ± 0.069
LOGP	13	0.651 ± 0.064	0.186 ± 0.064
REFRACTIVITY	11	0.649 ± 0.063	0.193 ± 0.058
GENERAL	12	0.633 ± 0.064	0.201 ± 0.069

Table A5: AUROC and BEDROC20 on the RDKit virtual screening benchmark for MOLBERT trained with various subsets of the RDKit calculated physicochemical properties.

## C PhysChemPred subset ablation

Table A5 gives an ablation of the PHYSCHEMPRED task by grouping the 200 descriptors used into smaller subsets of related descriptors, and repeating the Virtual Screening benchmark.

From the table we observe that in general all physicochemical descriptors do well on the benchmark. Moreover, some subsets which contain very few descriptors (SURFACE  $n = 49$  and CHARGE  $n = 18$ ) are able to achieve almost the same results as using all of physicochemical descriptors.

Finally, the lists of descriptors and their groupings are as follows:

**ALL:** Full set of 200 descriptors from RDKit.

BalabanJ, BertzCT, Chi0, Chi0n, Chi0v, Chi1, Chi1n, Chi1v, Chi2n, Chi2v, Chi3n, Chi3v, Chi4n, Chi4v, EState\_VSA1, EState\_VSA10, EState\_VSA11, EState\_VSA2, EState\_VSA3, EState\_VSA4, EState\_VSA5, EState\_VSA6, EState\_VSA7, EState\_VSA8, EState\_VSA9, ExactMolWt, FpDensityMorgan1, FpDensityMorgan2, FpDensityMorgan3, FractionCSP3, HallKierAlpha, HeavyAtomCount, HeavyAtomMolWt, Ipc, Kappa1, Kappa2, Kappa3, LabuteASA, MaxAbsEStateIndex, MaxAbsPartialCharge, MaxEStateIndex, MaxPartialCharge, MinAbsEStateIndex, MinAbsPartialCharge, MinEStateIndex, MinPartialCharge, MolLogP, MolMR, MolWt, NHOHCount, NOCount, NumAliphaticCarbocycles, NumAliphaticHeterocycles, NumAliphaticRings, NumAromaticCarbocycles, NumAromaticHeterocycles, NumAromaticRings, NumHAcceptors, NumHDonors, NumHeteroatoms, NumRadicalElectrons, NumRotatableBonds, NumSaturatedCarbocycles, NumSaturatedHeterocycles, NumSaturatedRings, NumValenceElectrons, PEOE\_VSA1, PEOE\_VSA10, PEOE\_VSA11, PEOE\_VSA12, PEOE\_VSA13,

PEOE\_VSA14, PEOE\_VSA2, PEOE\_VSA3, PEOE\_VSA4, PEOE\_VSA5, PEOE\_VSA6, PEOE\_VSA7, PEOE\_VSA8, PEOE\_VSA9, RingCount, SMR\_VSA1, SMR\_VSA10, SMR\_VSA2, SMR\_VSA3, SMR\_VSA4, SMR\_VSA5, SMR\_VSA6, SMR\_VSA7, SMR\_VSA8, SMR\_VSA9, SlogP\_VSA1, SlogP\_VSA10, SlogP\_VSA11, SlogP\_VSA12, SlogP\_VSA2, SlogP\_VSA3, SlogP\_VSA4, SlogP\_VSA5, SlogP\_VSA6, SlogP\_VSA7, SlogP\_VSA8, SlogP\_VSA9, TPSA, VSA\_EState1, VSA\_EState10, VSA\_EState2, VSA\_EState3, VSA\_EState4, VSA\_EState5, VSA\_EState6, VSA\_EState7, VSA\_EState8, VSA\_EState9, fr\_Al\_COO, fr\_Al\_OH, fr\_Al\_OH\_noTert, fr\_ArN, fr\_Ar\_COO, fr\_Ar\_N, fr\_Ar\_NH, fr\_Ar\_OH, fr\_COO, fr\_COO2, fr\_C\_O, fr\_C\_O\_noCOO, fr\_C\_S, fr\_HOCCN, fr\_Imine, fr\_NHO, fr\_NH1, fr\_NH2, fr\_N\_O, fr\_Ndealkylation1, fr\_Ndealkylation2, fr\_Nhpyrrole, fr\_SH, fr\_aldehyde, fr\_alkyl\_carbamate, fr\_alkyl\_halide, fr\_allylic\_oxid, fr\_amide, fr\_amidine, fr\_aniline, fr\_aryl\_methyl, fr\_azide, fr\_azo, fr\_barbitur, fr\_benzene, fr\_benzodiazepine, fr\_bicyclic, fr\_diazo, fr\_dihydropyridine, fr\_epoxide, fr\_ester, fr\_ether, fr\_furan, fr\_guanido, fr\_halogen, fr\_hdrzine, fr\_hdrzone, fr\_imidazole, fr\_imide, fr\_isocyan, fr\_isothiocyan, fr\_ketone, fr\_ketone\_Toppliss, fr\_lactam, fr\_lactone, fr\_methoxy, fr\_morpholine, fr\_nitrile, fr\_nitro, fr\_nitro\_ arom, fr\_nitro\_ arom\_nonortho, fr\_nitroso, fr\_oxazole, fr\_oxime, fr\_para\_hydroxylation, fr\_phenol, fr\_phenol\_noOrthoHbond, fr\_phos\_acid, fr\_phos\_ester, fr\_piperdine, fr\_piperzine, fr\_priamide, fr\_prisulfonamd, fr\_pyridine, fr\_quatN, fr\_sulfide, fr\_sulfonamd, fr\_sulfone, fr\_term\_acetylene, fr\_tetrazole, fr\_thiazole, fr\_thiocyan, fr\_thiophene, fr\_unbrch\_alkane, fr\_urea, qed

**SURFACE:** MOE-based surface descriptor subset.

EState\_VSA1, EState\_VSA10, EState\_VSA11, EState\_VSA2, EState\_VSA3, EState\_VSA4, EState\_VSA5, EState\_VSA6, EState\_VSA7, EState\_VSA8, EState\_VSA9, LabuteASA, PEOE\_VSA1, PEOE\_VSA10, PEOE\_VSA11, PEOE\_VSA12, PEOE\_VSA13, PEOE\_VSA14, PEOE\_VSA2, PEOE\_VSA3, PEOE\_VSA4, PEOE\_VSA5, PEOE\_VSA6, PEOE\_VSA7, PEOE\_VSA8, PEOE\_VSA9, SMR\_VSA1, SMR\_VSA10, SMR\_VSA2, SMR\_VSA3, SMR\_VSA4, SMR\_VSA5, SMR\_VSA6, SMR\_VSA7, SMR\_VSA8, SMR\_VSA9, SlogP\_VSA1, SlogP\_VSA10, SlogP\_VSA11, SlogP\_VSA12, SlogP\_VSA2, SlogP\_VSA3, SlogP\_VSA4, SlogP\_VSA5, SlogP\_VSA6, SlogP\_VSA7, SlogP\_VSA8, SlogP\_VSA9, TPSA

**CHARGE:** Partial charge and VSA/charge descriptor subset.

MaxAbsPartialCharge, MaxPartialCharge, MinAbsPartialCharge, MinPartialCharge, PEOE\_VSA1, PEOE\_VSA10, PEOE\_VSA11, PEOE\_VSA12, PEOE\_VSA13, PEOE\_VSA14, PEOE\_VSA2, PEOE\_VSA3, PEOE\_VSA4, PEOE\_VSA5, PEOE\_VSA6, PEOE\_VSA7, PEOE\_VSA8, PEOE\_VSA9

**FRAGMENT:** Count and fragment based descriptor subset.

NHOHCount, NOCount, NumAliphaticCarbocycles, NumAliphaticHeterocycles, NumAliphaticRings, NumAromaticCarbocycles, NumAromaticHeterocycles, NumAromaticRings, NumHAcceptors, NumHDonors, NumHeteroatoms, NumRotatableBonds, NumSaturatedCarbocycles, NumSaturatedHeterocycles, NumSaturatedRings, RingCount, fr\_Al\_COO, fr\_Al\_OH, fr\_Al\_OH\_noTert, fr\_ArN, fr\_Ar\_COO, fr\_Ar\_N, fr\_Ar\_NH, fr\_Ar\_OH, fr\_COO, fr\_COO2, fr\_C\_O, fr\_C\_O\_noCOO, fr\_C\_S, fr\_HOCCN, fr\_Imine, fr\_NHO, fr\_NH1, fr\_NH2, fr\_N\_O, fr\_Ndealkylation1, fr\_Ndealkylation2, fr\_Nhpyrrole, fr\_SH, fr\_aldehyde, fr\_alkyl\_carbamate, fr\_alkyl\_halide, fr\_allylic\_oxid, fr\_amide, fr\_amidine, fr\_aniline, fr\_aryl\_methyl, fr\_azide, fr\_azo, fr\_barbitur, fr\_benzene, fr\_benzodiazepine, fr\_bicyclic, fr\_diazo, fr\_dihydropyridine, fr\_epoxide, fr\_ester, fr\_ether, fr\_furan, fr\_guanido, fr\_halogen, fr\_hdrzine, fr\_hdrzone, fr\_imidazole, fr\_imide, fr\_isocyan, fr\_isothiocyan, fr\_ketone, fr\_ketone\_Toppliss, fr\_lactam, fr\_lactone, fr\_methoxy, fr\_morpholine, fr\_nitrile, fr\_nitro, fr\_nitro\_ arom, fr\_nitro\_ arom\_nonortho, fr\_nitroso, fr\_oxazole, fr\_oxime, fr\_para\_hydroxylation, fr\_phenol, fr\_phenol\_noOrthoHbond, fr\_phos\_acid, fr\_phos\_ester, fr\_piperdine, fr\_piperzine, fr\_priamide, fr\_prisulfonamd, fr\_pyridine, fr\_quatN, fr\_sulfide, fr\_sulfonamd, fr\_sulfone, fr\_term\_acetylene, fr\_tetrazole, fr\_thiazole, fr\_thiocyan, fr\_thiophene, fr\_unbrch\_alkane, fr\_urea

**SIMPLE:** Small set of commonly used descriptors.

FpDensityMorgan2, FractionCSP3, MolLogP, MolWt, NumHAcceptors, NumHDonors, NumRotatableBonds, TPSA

**GRAPH:** Graph descriptor subset (following the grouping found in the *Graph.Descriptors* module in RDKit).

BalabanJ, BertzCT, Chi0, Chi0n, Chi0v, Chi1, Chi1n, Chi1v, Chi2n, Chi2v, Chi3n, Chi3v, Chi4n, Chi4v, HallKierAlpha, Ipc, Kappa1, Kappa2, Kappa3

**ESTATE:** Electrotopolological state (e-state) and VSA/e-state descriptor subset.

EState\_VSA1, EState\_VSA10, EState\_VSA11, EState\_VSA2, EState\_VSA3,

EState\_VSA4, EState\_VSA5, EState\_VSA6, EState\_VSA7, EState\_VSA8, EState\_VSA9,  
MaxAbsEStateIndex, MaxEStateIndex, MinAbsEStateIndex, MinEStateIndex, VSA\_EState1,  
VSA\_EState10, VSA\_EState2, VSA\_EState3, VSA\_EState4, VSA\_EState5, VSA\_EState6,  
VSA\_EState7, VSA\_EState8, VSA\_EState9

**DRUGLIKENESS:** Subset of descriptors commonly used to assess druglikeness.

ExactMolWt, FractionCSP3, HeavyAtomCount, MolLogP, MolMR, MolWt, NHOHCount,  
NOCCount, NumAliphaticCarbocycles, NumAliphaticHeterocycles, NumAliphaticRings,  
NumAromaticCarbocycles, NumAromaticHeterocycles, NumAromaticRings, NumHAcceptors,  
NumHDonors, NumHeteroatoms, NumRotatableBonds, NumSaturatedCarbocycles,  
NumSaturatedHeterocycles, NumSaturatedRings, RingCount, TPSA, qed

**LOGP:** LogP and VSA/LogP descriptor subset.

MolLogP, SlogP\_VSA1, SlogP\_VSA10, SlogP\_VSA11, SlogP\_VSA12, SlogP\_VSA2,  
SlogP\_VSA3, SlogP\_VSA4, SlogP\_VSA5, SlogP\_VSA6, SlogP\_VSA7, SlogP\_VSA8, SlogP\_VSA9

**REFRACTIVITY:** MOE-based refractivity related descriptor subset.

MolMR, SMR\_VSA1, SMR\_VSA10, SMR\_VSA2, SMR\_VSA3, SMR\_VSA4, SMR\_VSA5, SMR\_VSA6,  
SMR\_VSA7, SMR\_VSA8, SMR\_VSA9

**GENERAL:** General descriptor subset (following the grouping found in the *Descriptors* module in RDKit).

ExactMolWt, FpDensityMorgan1, FpDensityMorgan2, FpDensityMorgan3, HeavyAtomMolWt,  
MaxAbsPartialCharge, MaxPartialCharge, MinAbsPartialCharge, MinPartialCharge,  
MolWt, NumRadicalElectrons, NumValenceElectrons