

1 TeachOpenCADD goes Deep Learning: 2 Open-source Teaching Platform 3 Exploring Molecular DL Applications

4 Michael Backenköhler^{1†}, Paula Linh Kramer^{1,2†}, Joschka Groß^{2†}, Gerrit
5 Großmann^{2,3†}, Roman Joeres^{4,5†}, Azat Tagirdzhanov^{6,7†}, Dominique Sydow⁸,
6 Hamza Ibrahim¹, Floriane Odje¹, Verena Wolf^{2,3}, Andrea Volkamer^{1,8*}

*For correspondence:

volkamer@cs.uni-saarland.de (AV)

†These authors contributed
equally to this work

7 ¹Data Driven Drug Design, Center for Bioinformatics, Saarland University; ²Modeling
8 and Simulation, Saarland University; ³German Research Center for Artificial Intelligence
9 (DFKI), Saarbrücken; ⁴Drug Bioinformatics, Helmholtz Institute for Pharmaceutical
10 Research Saarland; ⁵Drug Bioinformatics, Center for Bioinformatics, Saarland
11 University; ⁶Human-Microbe Systems Bioinformatics, Helmholtz Institute for
12 Pharmaceutical Research Saarland; ⁷Clinical Bioinformatics, Center for Bioinformatics,
13 Saarland University; ⁸Structural Bioinformatics and in silico Toxicology, Institute of
14 Physiology, Universitätsmedizin Berlin

16 **Abstract** TeachOpenCADD is a free online platform that offers solutions to common
17 computer-aided drug design (CADD) tasks using Python programming and open-source data and
18 packages. The material is presented through interactive Jupyter notebooks, accommodating
19 users from various backgrounds and programming levels.

20 Due to the tremendous impact of deep learning (DL) methods in drug design, the
21 TeachOpenCADD platform has been expanded to include an introduction to molecular DL tasks.
22 This edition provides an overview of DL and its application in drug design, highlighting the usage
23 of diverse molecular representations in this field. The platform introduces various neural
24 network architectures, including graph neural networks (GNNs), equivariant graph neural
25 networks (EGNNs), and recurrent neural networks (RNNs). It demonstrates how to use these
26 architectures for developing predictive models for molecular property and activity prediction,
27 exemplified by the Quantum Machine 9 (QM9), ChEMBL, and Kinase Inhibitor BioActivity (KiBA)
28 data sets. The DL edition covers methods for evaluating the performance of neural networks
29 using uncertainty estimation. Furthermore, it introduces an application of GNNs for
30 protein-ligand interaction predictions, incorporating protein structure and ligand information.
31 The TeachOpenCADD platform is continuously updated with new content and is open to
32 contributions, bug reports, and questions from the community through its GitHub repository
33 (github.com/volkamerlab/teachopencadd). It can be used for self-study, classroom instruction, and
34 research applications, accommodating users from beginners to advanced levels.

36 Introduction

37 CADD in the deep learning era

38 The process of discovering new drugs remains both expensive and time-consuming. The approval
39 of a single drug typically takes between 10 and 15 years, with average costs exceeding one bil-

40 lion US dollars (*Scannell et al., 2012*). Computer-aided drug design (CADD) has become a crucial
41 component in the drug development process, offering data-driven guidance in the search for or op-
42 timization of innovative compounds. Over the last decade, the immense growth of freely available
43 chemical databases such as ChEMBL (*Gaulton et al., 2017*) and Protein Data Bank (PDB) (*Berman
44 et al., 2000*) has further stimulated the development and application of data-driven approaches
45 such as machine and deep learning (DL). The latter has brought about significant advancements
46 in various fields in recent years, as evidenced by innovations like ChatGPT (*Brown et al., 2020*) and
47 AlphaFold (*Jumper et al., 2021; Wu et al., 2022*).

48 In the realm of drug discovery, DL has demonstrated immense potential (*Volkamer et al., 2023*)
49 due to its ability to process and learn from large and complex data sets (*Lavecchia, 2019*). Here,
50 we propose a learning pipeline based on Jupyter notebooks for chemists, biologists, and computer
51 scientists alike. Previous training material is available introducing cheminformatics and DL but
52 with a different scope and setup (*Menke et al., 2023*), or a stronger computer science background
53 (*Ramsundar et al., 2019*). We start from scratch by explaining the theoretical foundations and
54 show practical examples in Python, solving real-world molecular problems using widely known
55 DL methods. The learning pipeline is based on the well-established TeachOpenCADD framework
56 (*Sydow et al., 2019, 2022; Kimber et al., 2021*).

57 Molecular deep learning in a nutshell

58 In the field of drug discovery, DL has been applied to many different problem settings, such as
59 molecular activity, and toxicity prediction (*Coley et al., 2017; Unke and Meuwly, 2019; Wu et al.,
60 2018; Mayr et al., 2016; Coley et al., 2017*). Moreover, several docking approaches based on DL
61 have been published reporting promising results (*Corso et al., 2022; Ganea et al., 2021; Stärk et al.,
62 2022*), as well as generative models for *de novo* drug design (*Jin et al., 2020; Hoogbeem et al.,
63 2022*).

64 A DL network typically consists of multiple, connected layers with non-linear, parameterized
65 transformations. The data is provided to the input layer, which then gets processed through a pre-
66 defined number of hidden layers, and finally, an output layer generating the prediction (see Figure
67 1 for some drug design examples) (*Goodfellow et al., 2016*). In the process of training a network,
68 the parameters are adjusted to distill large data sets down to relevant features and patterns asso-
69 ciated with the prediction task. Neural networks can be trained for a variety of tasks. They can be
70 used for classification tasks, such as determining whether a molecule is toxic or not, or regression
71 tasks, like predicting binding affinity. Depending on the input data, there are many different classes
72 of neural networks suited for handling molecular data, each having different (dis)advantages. For
73 instance, graph neural networks (GNNs) offer a natural architecture for molecular graphs that has
74 several advantages: They capture both atom and bond information, as well as the connectivity
75 between atoms while being invariant to the nodes' input order. They can handle molecules of
76 varying sizes and complexities and learn both local and global features of molecular structures.
77 Convolutional neural networks (CNNs) are often used for image data, while recurrent neural net-
78 works (RNNs) and transformers are designed to handle sequential data (such as text). Some of
79 these architectures will be covered in our tutorials.

80 TeachOpenCADD: Scope and DL extension

81 As of September 2022, TeachOpenCADD (*Sydow et al., 2022*) contained 28 talktorials covering
82 diverse topics in the broader area of CADD. Most talktorials are exemplified by compound and
83 structural data available for the EGFR kinase (*Herbst, 2004*). The platform contains talktorials in-
84 troducing the following topics: (i) Cheminformatics basics, e.g. molecular filtering, clustering, and
85 substructure search, as well as similarity search and machine learning models for activity predic-
86 tion; (ii) chemical database queries, e.g. ChEMBL (*Gaulton et al., 2017*), PDB (*Berman et al., 2000*),
87 PubChem (*Kim et al., 2022*), and KLIFS queries (*Kanev et al., 2020*); (iii) structural bioinformatics,
88 e.g. binding site detection and comparison, docking, protein-ligand interaction profiling, as well as

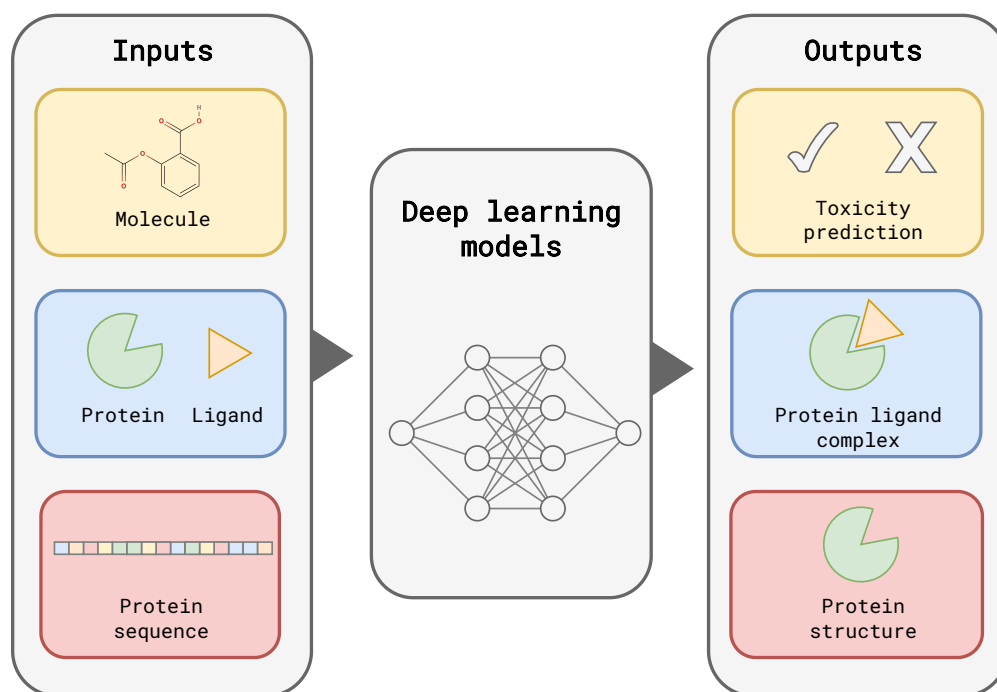


Figure 1. Some applications of DL in the field of drug discovery.

89 molecular dynamics simulations; and (iv) kinase similarity assessment including different perspec-
 90 tives, e.g. sequence, structure, interaction, and profiling data (Kimber *et al.*, 2021).

91 With the *TeachOpenCADD-DL* edition, we introduce the concepts of DL applied to molecules in
 92 six new talktorials. The topics are summarized in Figure 2. As an introduction, we discuss various
 93 methods of representing molecules to facilitate their processing by neural networks. For each of
 94 the representations, we introduce a class of neural networks: (i) GNNs with molecules represented
 95 as a graph, (ii) RNNs where molecules are represented as a SMILES string (Weininger, 1988), and (iii)
 96 equivariant graph neural networks (EGNNs) which process molecules as point clouds. Each neu-
 97 ral network is trained to perform a regression task with the objective of predicting the quantum-
 98 mechanical properties of small molecules. In addition to the network architectures, we also cover
 99 uncertainty estimations to evaluate the performance of a trained model using molecular finger-
 100 prints as input. Finally, we describe an important application of DL for protein-ligand interaction
 101 prediction.

102 Data

103 In this section, we describe the three molecular data sets used to exemplify the different architec-
 104 tures to solve diverse prediction tasks.

105 Quantum Machines 9 (QM9) Data Set

106 QM9 is a public data set that consists of 130k small, organic molecules with up to 9 heavy atoms (Ra-
 107 makrishnan *et al.*, 2014). Each molecule is annotated with various geometric, energetic, electronic,
 108 and thermodynamic properties. QM9 is part of MoleculeNet (Wu *et al.*, 2018), a widely adopted
 109 property prediction benchmark in the molecular machine learning community, e.g., see (Schütt
 110 *et al.*, 2017; Gilmer *et al.*, 2017; Gasteiger *et al.*, 2020). PyTorch Geometric (Fey and Lenssen, 2019)
 111 provides pre-implemented classes and methods for working with the QM9 data set in a molecular
 112 ML setting.

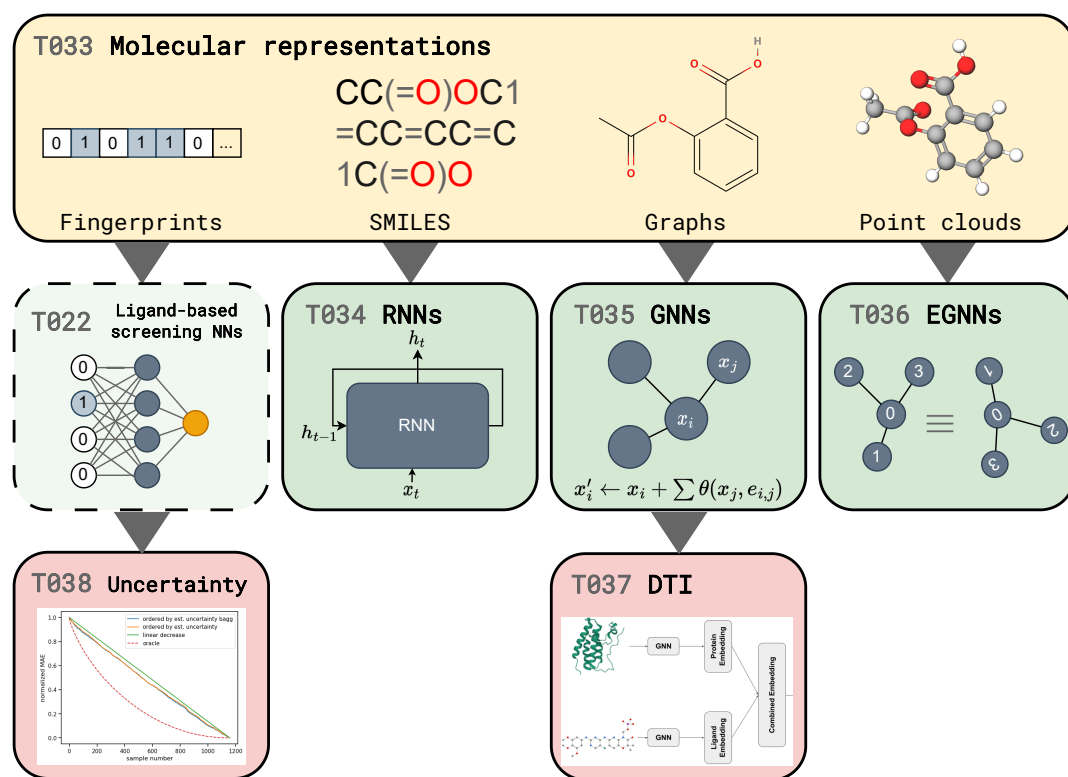


Figure 2. DL-talktorials: The newly contributed talktorials cover molecular representations for machine learning (T033), corresponding deep learning architectures for processing them (T034-36), and more involved topics such as concrete applications (T037) and uncertainty analysis (T038).

113 ChEMBL EGFR Subset

114 In the uncertainty estimation talktorial, we make use of activity data available for the EGFR kinase
 115 from ChEMBL (*Gaulton et al., 2017*). Protein kinases play a central role in many stages of a cell's life
 116 cycle. Dysfunctional signaling of EGFR kinase, e.g., has been associated with cancer progression
 117 (*Chen et al., 2016*). The activity data we use is extracted from the public ChEMBL database (*Gaulton*
 118 *et al., 2017*), version 25. Only IC50 data from binding assays (assay_type="B") and exact measure-
 119 ments (standard_type="") were kept. The data set contains ~3900 compounds with activities from
 120 binding assays available as IC50 values.

121 Kinase Inhibitor BioActivity Data Set

122 The Kinase Inhibitor BioActivity (KiBA) data set has been assembled from diverse published kinase
 123 profiling data sets to provide a large benchmark set for kinase drug-target activity. It is a collection
 124 of 467 kinases, 52, 498 ligands, and 246, 088 KiBA scores thereof. The KiBA scores are computed
 125 to combine data acquired through different bioactivity experiments and measurements such as
 126 IC50, K(i), and K(d) (*Tang et al., 2014*).

127 For the protein-ligand interaction talktorial (see Section Protein-Ligand Interaction Prediction),
 128 we selected a subset of KiBA in order to speed up the training process, reduce memory consump-
 129 tion, and make it trainable on average CPUs in a reasonable time. This is done in two steps: First,
 130 all ligands measured against less than 200 kinases are discarded. Second, from the remaining
 131 data points, all kinases with data available for less than 10 ligands are removed. Furthermore, a
 132 pipeline was provided to scrape the matching PDB structure per kinase starting from UniProt IDs
 133 (*Consortium, 2022*) and enforcing some structure quality filters. This resulted in 76 kinases, 275
 134 ligands, and 20, 475 KiBA scores thereof.

Table 1. Summary of the topics covered in the TeachOpenCADD-DL edition.

Topic	Description	Mol. input
Molecular representations	Introduction to molecules and their representations	All below
Recurrent neural networks (RNNs)	RNNs and Gated Recurrent Unit (GRU) for molecular property prediction	SMILES
Graph neural networks (GNNs)	Convolutional and isomorphism GNNs for molecular property prediction	Graph
$E(3)$ -invariant graph neural networks (EGNNs)	EGNNs compared to standard GNNs for molecular property prediction	Point clouds
Uncertainty estimation	Methods for model uncertainty estimation	Fingerprints
Protein-ligand interaction prediction	Applying GNNs to predict protein-ligand interactions	SMILES & PDB

135 Talktorials

136 In this section, we describe the six novel topics covered in the TeachOpenCADD-DL edition (see
137 Table 1). Note that all talktorials serve as teaching or starting examples, thus, the architectures
138 were intentionally kept simple and no parameters are tuned to optimize prediction performance.

139 Molecular Representations

140 Molecules are intricate, dynamic, three-dimensional (3D) entities composed of atoms, interacting
141 with each forming covalent as well as non-covalent bonds. It is essential to represent molecules
142 in a computer-readable form that corresponds to the information processed through a neural
143 network. In this talktorial, we cover popular molecular representations and discuss their unique
144 implications and (dis-)advantages. This will provide the foundation for subsequent talktorials.

145 Representing molecules as *graphs* allows for an intuitive and comprehensive representation of
146 their structure. In a graph-based representation, atoms are represented as (labeled) nodes, and
147 bonds are represented as (labeled) edges. However, to represent a graph, we need node ordering.
148 This node ordering, while necessary, is arbitrary, and ideally, a DL predictor should yield the same
149 output regardless of the node order chosen. GNNs address this issue by inherently ensuring this
150 so-called *permutation invariance* by design (Atz et al., 2021).

151 Molecular *fingerprints* are fixed-length, permutation-invariant representations of the molecular
152 graphs. Unlike GNNs, which learn task-specific representations, molecular fingerprints are task-
153 independent. They can be generated based on the occurrences of specific sub-graphs (i.e., molec-
154 ular fragments or atom environments) (Rogers and Hahn, 2010). Generally, it is not feasible to
155 reverse-engineer a fingerprint back to the original molecular graph. Due to their fixed length, fin-
156 gerprints are compatible with machine learning methods that require a constant input size, such
157 as Multi-Layer Perceptrons (MLPs).

158 *Text-based* representations (like SMILES (Weininger, 1988), SELFIES (Krenn et al., 2020), or InChI
159 (Grisoni, 2023)) traverse the molecular graph and convert it into a sequence of characters. How-
160 ever, ambiguity can occur due to the possibility of multiple strings mapping to the same molecule,
161 depending on the order of traversal. To reduce ambiguity, *canonical* SMILES can be used, although
162 what counts as canonicalized SMILES string is not standardized and may differ based on the soft-
163 ware package in use. A text-based molecular representation is well-suited for machine learning
164 (ML) models capable of handling sequences with varying lengths. Specifically, they have been suc-
165 cessfully used as input to language models (Wang et al., 2019; Chithrananda et al., 2020).

166 *Point cloud* representations annotate atoms with their 3D coordinates, corresponding to a sin-
167 gle conformation. A molecular conformation (conformer) is a specific spatial arrangement of atoms
168 within a molecule, reflecting a single energetically favorable configuration of its 3D structure. Like

169 in GNNs, this necessitates a special type of invariance for DL methods that take point clouds as
170 input. Our specific goal is to attain invariance to Euclidean space transformations (e.g., the output
171 of the neural network model should remain unchanged when the entire molecule is rotated). Point
172 cloud representations are especially advantageous, as they encompass more comprehensive infor-
173 mation. In particular, they capture the relative atomic positions, which reflect the collective effect
174 of all forces acting within a molecule, beyond just covalent bonds (*Atz et al., 2021*).

175 In our talktorial, we discuss the different molecular representations in more detail and demon-
176 strate how to generate and utilize them in Python.

177 **Recurrent Neural Networks**

178 In recent years, DL-based natural language processing (NLP) has made significant progress, with
179 RNNs and transformers among the most successful models. These models proved to be good at
180 capturing text semantics and, when applied to molecular data, can capture the molecular structure
181 in its textual representation. As a result, NLP models have become a powerful tool in numerous
182 drug discovery applications, including *de novo* drug design (*Gupta et al., 2018*), virtual screening
183 (*Karimi et al., 2019*), and molecular property prediction (*Bjerrum, 2017*).

184 RNNs were originally developed to handle sequential data (*Elman, 1990*). These models can
185 process variable-length sequences of inputs and propagate the information through the sequence
186 using their internal state. In this talktorial, we focus on applying RNNs to SMILES strings. We briefly
187 cover the usual preprocessing steps that transform SMILES into numerical form and discuss two
188 RNN architectures in detail, starting with the Elman network, also known as a simple RNN (*Elman,*
189 *1990*). This architecture is suitable for demonstrating the basic principles of RNNs, but in practice, it
190 struggles with learning long-term dependencies in the data. This problem is addressed in the more
191 advanced Gated Recurrent Unit (GRU) (*Cho et al., 2014*) architecture. GRU selectively updates its
192 internal state using gating mechanisms, allowing the model to learn to identify and retain the most
193 important information while discarding irrelevant information.

194 We implement RNN- and GRU-based regression models and apply them to molecular property
195 prediction using the QM9 data set. As a regression task, we have chosen to predict the dipole
196 moment μ , which is a measure of a molecule's polarity. Our results show that the GRU model
197 learns faster and achieves better performance than the simple RNN model.

198 **Graph Neural Networks**

199 The most natural representation for molecules are graphs spanned by their atoms and bonds.
200 Thus, one intuitive way to apply DL techniques to molecular data is using GNNs. GNNs are widely
201 used in drug discovery, for example for property prediction (*Wu et al., 2018; Wieder et al., 2020*)
202 and *de novo* drug design (*Xia et al., 2019; Tong et al., 2021*).

203 Instead of the fully connected layers commonly used in standard neural networks, GNNs have
204 message-passing layers, that collect information about the neighboring nodes in the graph (*Kipf*
205 *and Welling, 2016*). For each node in the graph, all the information from the neighbors is gathered
206 and aggregated using an aggregation function such as the sum. One important property of a GNN
207 is the permutation invariance. This means that changing the arbitrary order of nodes in the graph
208 should not have an effect on the outcome. On the other hand, GNNs should ideally also be able to
209 distinguish between similar graphs.

210 In our talktorial, we present two commonly used GNN architectures in more detail: one of the
211 simplest GNNs, namely the graph convolutional neural network (GCN (*Kipf and Welling, 2016*)),
212 and a more powerful GNN called the graph isomorphism network (GIN (*Xu et al., 2018*)). GINs
213 are better at distinguishing similar, non-identical graphs compared to GCNs, which often leads to
214 better performance. We demonstrate how to implement GNNs and how to train them using the
215 QM9 data set (see Section Quantum Machines 9 (QM9) Data Set) to predict one quantum-mechanic
216 property of small molecules. We predict the same molecular property as in the previous talktorial
217 (see Section Recurrent Neural Networks).

218 **E(3)-invariant Graph Neural Network**

219 Reasoning about molecular properties is often easier when 3D information (e.g. in the form of con-
220 formations) is available. Some tasks may also strictly require the use of molecular representations
221 that include 3D information. Examples of this are binding pose predictions of ligand-protein com-
222 plexes (*Corso et al., 2022*) or force predictions for molecular dynamics simulations (*Doerr et al.,*
223 *2021*). It is widely accepted that GNNs which process molecules based on their point cloud repre-
224 sentation (see Section Molecular Representations) should satisfy certain invariance or equivariance
225 properties with respect to global *Euclidean transformations* such as translations or rotations.

226 The Euclidean group that corresponds to these transformations in three dimensions is denoted
227 by $E(3)$. $E(3)$ -invariance implies that the output of a GNN is unaffected by rotations or translations
228 of its input point cloud. For example, when predicting binding affinity based on the structure of
229 a ligand-protein complex, this prediction should remain unchanged if the entire complex is trans-
230 lated or rotated. $E(3)$ -equivariance implies that rotating or translating the GNN's input should in-
231 duce an equivalent transformation of its output. For example, when predicting the binding pose of
232 a ligand-based on a given protein structure, rotating the latter should give rise to an equivalently
233 rotated pose prediction.

234 This talktorial discusses these concepts in more detail in the theory part. It demonstrates how to
235 implement $E(3)$ -invariant graph neural networks for property prediction based on the point cloud
236 representation of the molecules included in the QM9 data set. The practical part concludes by
237 training and evaluating such a model in comparison to a plain GNN. The application shows that
238 the theoretical advantages mentioned above also lead to better results in practice.

239 **Uncertainty Estimation**

240 Often researchers pay a lot of attention to the overall accuracy of their predictions. However, when
241 implementing any predictive method in practice, it is equally important to understand the level of
242 confidence in a given estimation. The uncertainty can stem from both the experiments themselves
243 (epistemic) and/or the predictive model (aleatoric). In the former case, the uncertainty of the model
244 arises typically due to a lack of training data while the latter case refers to inherent randomness
245 such as measurement noise (*Der Kiureghian and Ditlevsen, 2009*). Thus, it would be beneficial to
246 obtain not only a point estimate of the prediction but also an indication of how certain we can
247 be about that estimate. The certainty is often modeled by replacing the point estimate with a
248 distributional estimate (*Gawlikowski et al., 2021*). For example, instead of a number as a prediction
249 of an IC50 value, one obtains a distribution of the predicted values.

250 In this talktorial, we showcase uncertainty estimation on a practical example. We start our
251 demonstration by creating a simple model ensemble. This means we train the same model multi-
252 ple times with a varying random seed. At test time, we evaluate all models and use the mean as a
253 predictor. The variance across the ensemble serves as a variance estimate for that prediction. We
254 discuss the calibration of this estimator, which – as is typical – under-estimates the actual variance.

255 In the second step, we improve our ensemble by not only varying the random seed during
256 training but also the data itself. This variation is achieved by bootstrapping the training data. This
257 helps to more accurately estimate uncertainty.

258 Finally, we showcase test time data augmentation as an alternative to the modification of our
259 predictive model. In this technique, we create variants for each query point in our test set. The
260 variants are created by applying random flips to a fingerprint datum. This way, we get an ensemble
261 of predictions out of a single model, without the need to modify the model itself.

262 **Protein-Ligand Interaction Prediction**

263 Protein-ligand interaction prediction is an important field in drug development, e.g. to screen for
264 novel drug candidates. Classical methods to predict drug-target interactions are based on docking
265 (*de Azevedo Jr et al., 2003; M Bernhardt Levin et al., 2017*), biological networks (*AY et al., 2007;*
266 *Chen et al., 2012*), and many more (*Zhao et al., 2022*). More recently, models use DL encoders

267 such as MLPs, i.e. CNNs and GNNs, to compute latent space representations, also called embed-
268 dings, of biochemical molecules (Öztürk *et al.*, 2018; Nguyen *et al.*, 2021). While in classical docking
269 methods, the complex structure is generated and then scored, in these works the two interaction
270 partners are treated separately. The embeddings are combined for each pair of potentially inter-
271 acting molecules, usually concatenated, and then fed into an MLP to predict the output variable.
272 The variable can either be a proxy value for binding affinity or a classification value separating
273 binding and non-binding pairs of protein and ligand.

274 The goal of this talktorial is to introduce the reader to the field of protein-ligand interaction
275 prediction using GNNs for proteins and ligands independently. In contrast to previous works in
276 which the protein was encoded as sequence and a CNN was used for the embedding, (Öztürk
277 *et al.*, 2018; Nguyen *et al.*, 2021), GNNs are used for both, proteins and ligands. Ligands are rep-
278 resented as graphs constructed from the SMILES string. Representing proteins is more complex
279 and done using Residue Interaction Networks (RINs) (Doncheva *et al.*, 2011). These are graphs
280 where nodes represent amino acids and edges represent covalent and non-covalent interactions
281 between amino acids. To compute those, RINminer (Keller *et al.*, 2020) can be used or a distance
282 threshold between amino acids in the three-dimensional space as a surrogate of such. The talk-
283 torial exemplifies this task of predicting interactions between proteins and ligands using the KiBA
284 subset (see Section Kinase Inhibitor BioActivity Data Set) and shows that predicting interaction on
285 the KiBA dataset is possible with little effort and simple GNNs.

286 Prerequisites and technical information

287 Target audience

288 The talktorials were developed to support researchers who are interested in the topics and are
289 new to the field. The covered scope is intended to further bridge the fields of CADD and DL. The
290 talktorials are recommended for biologists, medicinal chemists as well as computer scientists; and
291 should enable the user to apply the techniques in their own work. Since the talktorials form an
292 extension to the TeachOpenCADD platform, they serve as teaching material in the field of structural
293 bio- and cheminformatics.

294 Background knowledge

295 The tutorials are meant to be an introduction to DL and its application to the field of drug discovery.
296 In each talktorial, we first present the theoretical background for the biological and chemical ba-
297 sics as well as the computer science fundamentals. Secondly, we provide thoroughly documented
298 Python code to illustrate the application of DL. However, some proficiency in Python and Jupyter
299 would be helpful.

300 Software requirements

301 All talktorials are written in Python and make use of well-known open-source packages such as
302 Pandas (McKinney, 2011), NumPy (Harris *et al.*, 2020), Matplotlib (Hunter, 2007), SciPy (Virtanen
303 *et al.*, 2020), RDKit (Landrum, 2006). The novel DL talktorials make heavy use of PyTorch (Paszke
304 *et al.*, 2019) and PyTorch Geometric (Fey and Lenssen, 2019). The user only needs to install the
305 teachopencadd conda-forge package, which will install all relevant packages and save a copy of
306 all TeachOpenCADD notebooks on the user's local machine. A read-only mode of the talktorials is
307 accessible via the TeachOpenCADD website at projects.volkamerlab.org/teachopencadd/.

308 Structure of the talktorials

309 The talktorials serve a teaching purpose and are structured as follows: Each Jupyter notebook is
310 split into two parts. We first explain the underlying theory of each topic. We explain the problem
311 setting, give relevant references, and list possible applications. The second part is focusing on the
312 actual implementation in Python. We explain and document each step in the code. We want to
313 make it easy to follow and give the user the chance to extend this to different applications in the
314 field.

315 Conclusion

316 This study provides an insightful introduction to DL important for and applied to molecular predic-
317 tion tasks. We presented six talktorials covering topics such as commonly used representations of
318 molecules and proteins, graph and recurrent neural networks, uncertainty measures, and protein-
319 ligand interaction predictions. Through these talktorials, users can gain a better understanding of
320 DL and its potential applications in drug discovery. We believe that these methods can be used as
321 a starting point and can be adapted for different molecular data sets and more complex questions.

322 Author Contributions

323 MB, PK, JG, GG, AT, and RJ implemented the new notebooks. MB, HI, and DS integrated the new
324 material and maintained the repository. All authors reviewed individual talktorials. AV conceptual-
325 ized the study. VW and AV supervised the project. All authors contributed to writing and reviewing
326 the manuscript.

327 Acknowledgements

328 We thank all former and future supporters and users of the TeachOpenCADD platform. And we are
329 grateful for the fruitful discussions around the DL topics with all our colleagues from the NextAID
330 project (see nextaid.cs.uni-saarland.de/) at Saarland University.

331 Funding Information

332 We are thankful for funding from Saarland University for the NextAID project through which the
333 positions of MB, PK, JG, GG, and AT are supported. RJ was partially funded by the HelmholtzAI
334 project XAI-Graph.

335 References

- 336 **Atz K**, Grisoni F, Schneider G. Geometric deep learning on molecular representations. *Nature Machine Intelli-*
337 *gence*. 2021; 3(12):1023–1032. <https://doi.org/10.1038/s42256-021-00418-8>.
- 338 **AY M**, Goh KI, Cusick ME, Barabasi AL, Vidal M, et al. Drug–target network. *Nature biotechnology*. 2007;
339 25(10):1119–1127. <https://doi.org/10.1038/nbt1338>.
- 340 **de Azevedo Jr WF**, dos Santos GC, dos Santos DM, Olivieri JR, Canduri F, Silva RG, Basso LA, Renard G, da Fonseca
341 IO, Mendes MA, et al. Docking and small angle X-ray scattering studies of purine nucleoside phosphorylase.
342 *Biochemical and Biophysical Research Communications*. 2003; 309(4):923–928. [https://doi.org/10.1016/j.](https://doi.org/10.1016/j.bbrc.2003.08.093)
343 [bbrc.2003.08.093](https://doi.org/10.1016/j.bbrc.2003.08.093).
- 344 **Berman HM**, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE. The protein data
345 bank. *Nucleic acids research*. 2000; 28(1):235–242. <https://doi.org/10.1093/nar/28.1.235>.
- 346 **Bjerrum EJ**. SMILES enumeration as data augmentation for neural network modeling of molecules. *arXiv*
347 *preprint arXiv:170307076*. 2017; <https://doi.org/10.48550/arXiv.1703.07076>.
- 348 **Brown T**, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A,
349 et al. Language models are few-shot learners. *Advances in neural information processing systems*. 2020;
350 33:1877–1901. <https://doi.org/10.48550/arXiv.2005.14165>.
- 351 **Chen J**, Zeng F, Forrester SJ, Eguchi S, Zhang MZ, Harris RC. Expression and Function of the Epidermal Growth
352 Factor Receptor in Physiology and Disease. *Physiological Reviews*. 2016; 96:1025–1069. [doi: 10.1152/phys-](https://doi.org/10.1152/physrev.00030.2015)
353 [rev.00030.2015](https://doi.org/10.1152/physrev.00030.2015).
- 354 **Chen X**, Liu MX, Yan GY. Drug–target interaction prediction by random walk on the heterogeneous network.
355 *Molecular BioSystems*. 2012; 8(7):1970–1978. <https://doi.org/10.1039/C2MB00002D>.
- 356 **Chithrananda S**, Grand G, Ramsundar B. Chemberta: Large-scale self-supervised pretraining for molecular
357 property prediction. *arXiv preprint arXiv:201009885*. 2020; [doi: 10.48550/arXiv.2010.09885](https://doi.org/10.48550/arXiv.2010.09885).

- 358 **Cho K**, Van Merriënboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, Bengio Y. Learning phrase rep-
359 resentations using RNN encoder-decoder for statistical machine translation. arXiv preprint arXiv:14061078.
360 2014; <https://doi.org/10.48550/arXiv.1406.1078>.
- 361 **Coley CW**, Barzilay R, Green WH, Jaakkola TS, Jensen KF. Convolutional embedding of attributed molecular
362 graphs for physical property prediction. *Journal of chemical information and modeling*. 2017; 57(8):1757-
363 1772. <https://doi.org/10.1021/acs.jcim.6b00601>.
- 364 **Consortium TU**. UniProt: the Universal Protein Knowledgebase in 2023. *Nucleic Acids Research*. 2022 11;
365 51(D1):D523–D531. <https://doi.org/10.1093/nar/gkac1052>, doi: 10.1093/nar/gkac1052.
- 366 **Corso G**, Stärk H, Jing B, Barzilay R, Jaakkola T. Diffdock: Diffusion steps, twists, and turns for molecular docking.
367 arXiv preprint arXiv:221001776. 2022; <https://doi.org/10.48550/arXiv.2210.01776>.
- 368 **Der Kiureghian A**, Ditlevsen O. Aleatory or epistemic? Does it matter? *Structural safety*. 2009; 31(2):105–112.
- 369 **Doerr S**, Majewski M, Pérez A, Kramer A, Clementi C, Noe F, Giorgino T, De Fabritiis G. Torchmd: A deep learning
370 framework for molecular simulations. *Journal of chemical theory and computation*. 2021; 17(4):2355–2363.
371 <https://doi.org/10.1021/acs.jctc.0c01343>.
- 372 **Doncheva NT**, Klein K, Domingues FS, Albrecht M. Analyzing and visualizing residue networks of protein struc-
373 tures. *Trends in biochemical sciences*. 2011; 36(4):179–182. <https://doi.org/10.1016/j.tibs.2011.01.002>.
- 374 **Elman JL**. Finding structure in time. *Cognitive science*. 1990; 14(2):179–211. [https://doi.org/10.1016/0364-0213\(90\)90002-E](https://doi.org/10.1016/0364-0213(90)90002-E).
- 376 **Fey M**, Lenssen JE. Fast Graph Representation Learning with PyTorch Geometric. *CoRR*. 2019; abs/1903.02428.
377 <https://doi.org/10.48550/arXiv.1903.02428>.
- 378 **Ganea OE**, Huang X, Bunne C, Bian Y, Barzilay R, Jaakkola T, Krause A. Independent se (3)-equivariant models
379 for end-to-end rigid protein docking. arXiv preprint arXiv:211107786. 2021; <https://doi.org/10.48550/arXiv.2111.07786>.
- 381 **Gasteiger J**, Groß J, Günnemann S. Directional message passing for molecular graphs. arXiv preprint
382 arXiv:200303123. 2020; <https://doi.org/10.48550/arXiv.2003.03123>.
- 383 **Gaulton A**, Hersey A, Nowotka M, Bento AP, Chambers J, Mendez D, Mutowo P, Atkinson F, Bellis LJ, Cibrián-
384 Uhalte E, et al. The ChEMBL database in 2017. *Nucleic acids research*. 2017; 45(D1):D945–D954. <https://doi.org/10.1093/nar/gkw1074>.
- 386 **Gawlikowski J**, Tassi CRN, Ali M, Lee J, Humt M, Feng J, Kruspe A, Triebel R, Jung P, Roscher R, et al. A survey of
387 uncertainty in deep neural networks. arXiv preprint arXiv:210703342. 2021; <https://doi.org/10.48550/arXiv.2107.03342>.
- 389 **Gilmer J**, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural message passing for quantum chemistry. In:
390 *International conference on machine learning* PMLR; 2017. p. 1263–1272. <https://doi.org/10.48550/arXiv.1704.01212>.
- 392 **Goodfellow I**, Bengio Y, Courville A. *Deep Learning*. MIT Press; 2016. <http://www.deeplearningbook.org>.
- 393 **Grisoni F**. Chemical language models for de novo drug design: Challenges and opportunities. *Current Opinion*
394 *in Structural Biology*. 2023; 79:102527. <https://doi.org/10.1016/j.sbi.2023.102527>.
- 395 **Gupta A**, Müller AT, Huisman BJ, Fuchs JA, Schneider P, Schneider G. Generative recurrent networks for de
396 novo drug design. *Molecular informatics*. 2018; 37(1-2):1700111. <https://doi.org/10.1002/minf.201700111>.
- 397 **Harris CR**, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith
398 NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-
399 Marchant P, et al. Array programming with NumPy. *Nature*. 2020; 585(7825):357–362. <https://doi.org/10.1038/s41586-020-2649-2>, doi: 10.1038/s41586-020-2649-2.
- 401 **Herbst RS**. Review of epidermal growth factor receptor biology. *International Journal of Radiation Oncol-*
402 *ogy*Biolog*Physics*. 2004; 59:S21–S26. doi: <https://doi.org/10.1016/j.ijrobp.2003.11.041>.
- 403 **Hoogeboom E**, Satorras VG, Vignac C, Welling M. Equivariant diffusion for molecule generation in 3D. In:
404 *International Conference on Machine Learning* PMLR; 2022. p. 8867–8887. <https://doi.org/10.48550/arXiv.2203.17003>.
- 405

406 **Hunter JD.** Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*. 2007; 9(3):90–95. doi:
407 [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55).

408 **Jin W, Barzilay R, Jaakkola T.** Hierarchical generation of molecular graphs using structural motifs. In: *International conference on machine learning* PMLR; 2020. p. 4839–4848. <https://doi.org/10.48550/arXiv.2002.03230>.

410 **Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Žídek A, Potapenko A, et al.** Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021; 596(7873):583–589. <https://doi.org/10.1038/s41586-021-03819-2>.

413 **Kanev GK, de Graaf C, Westerman BA, de Esch IJP, Kooistra AJ.** KLIFS: an overhaul after the first 5 years of supporting kinase research. *Nucleic Acids Research*. 2020 10; 49(D1):D562–D569. <https://doi.org/10.1093/nar/gkaa895>, doi: 10.1093/nar/gkaa895.

416 **Karimi M, Wu D, Wang Z, Shen Y.** DeepAffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks. *Bioinformatics*. 2019; 35(18):3329–3338. <https://doi.org/10.1093/bioinformatics/btz111>.

419 **Keller S, Miettinen P, Kalinina OV.** Frequent subgraph mining for biologically meaningful structural motifs. *BioRxiv*. 2020; p. 2020–05. <https://doi.org/10.1101/2020.05.14.095695>.

421 **Kim S, Chen J, Cheng T, Gindulyte A, He J, He S, Li Q, Shoemaker BA, Thiessen PA, Yu B, Zaslavsky L, Zhang J, Bolton EE.** PubChem 2023 update. *Nucleic Acids Research*. 2022 10; 51(D1):D1373–D1380. <https://doi.org/10.1093/nar/gkac956>, doi: 10.1093/nar/gkac956.

424 **Kimber TB, Sydow D, Volkamer A.** Kinase Similarity Assessment Pipeline for Off-Target Prediction [Article v1.0]. *Living Journal of Computational Molecular Science*. 2021; 3(1):1599–1599. <https://livecomsjournal.org/index.php/livecoms/article/view/v3i1e1599>, doi: 10.33011/livecoms.3.1.1599.

427 **Kipf TN, Welling M.** Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:160902907*. 2016; <https://doi.org/10.48550/arXiv.1609.02907>.

429 **Krenn M, Häse F, Nigam A, Friederich P, Aspuru-Guzik A.** Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Machine Learning: Science and Technology*. 2020; 1(4):045024. doi: 10.1088/2632-2153/aba947.

432 **Landrum G,** RDKit: Open-source cheminformatics; 2006. <https://www.rdkit.org>, accessed: 2023-05-22.

433 **Lavecchia A.** Deep learning in drug discovery: opportunities, challenges and future prospects. *Drug Discovery Today*. 2019 Oct; 24(10):2017–2032. <https://www.sciencedirect.com/science/article/pii/S135964461930282X>, doi: 10.1016/j.drudis.2019.07.006.

436 **M Bernhardt Levin N, Oliveira Pintro V, Boff de Avila M, Boldrini de Mattos B, Filgueira DAJ, et al.** Understanding the structural basis for inhibition of cyclin-dependent kinases. New pieces in the molecular puzzle. *Current drug targets*. 2017; 18(9):1104–1111. doi: 10.2174/1389450118666161116130155.

439 **Mayr A, Klambauer G, Unterthiner T, Hochreiter S.** DeepTox: toxicity prediction using deep learning. *Frontiers in Environmental Science*. 2016; 3:80. <https://doi.org/10.3389/fenvs.2015.00080>.

441 **McKinney W.** pandas: a foundational Python library for data analysis and statistics. *Python for High Performance and Scientific Computing*. 2011; 14. <https://pandas.pydata.org/>.

443 **Menke J, Homberg S, Koch O.** Introduction to artificial intelligence and deep learning using interactive electronic programming notebooks. *Archiv der Pharmazie*. 2023; p. e2200628. <https://doi.org/10.1002/ardp.202200628>.

445 **Nguyen T, Le H, Quinn TP, Nguyen T, Le TD, Venkatesh S.** GraphDTA: predicting drug–target binding affinity with graph neural networks. *Bioinformatics*. 2021; 37(8):1140–1147. <https://doi.org/10.1093/bioinformatics/btaa921>.

448 **Öztürk H, Özgür A, Ozkirimli E.** DeepDTA: deep drug–target binding affinity prediction. *Bioinformatics*. 2018; 34(17):i821–i829. <https://doi.org/10.1093/bioinformatics/bty593>.

450 **Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, et al.** Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*. 2019; 32. <https://doi.org/10.48550/arXiv.1912.01703>.

452

453 **Ramakrishnan R**, Dral PO, Rupp M, Von Lilienfeld OA. Quantum chemistry structures and properties of 134
454 kilo molecules. Scientific data. 2014; 1(1):1–7. <https://doi.org/10.1038/sdata.2014.22>.

455 **Ramsundar B**, Eastman P, Walters P, Pande V. Deep Learning for the Life Sciences. O'Reilly Media; 2019.

456 **Rogers D**, Hahn M. Extended-Connectivity Fingerprints. Journal of Chemical Information and Modeling. 2010;
457 50(5):742–754. <https://doi.org/10.1021/ci100050t>, doi: 10.1021/ci100050t.

458 **Scannell JW**, Blanckley A, Boldon H, Warrington B. Diagnosing the decline in pharmaceutical R&D efficiency.
459 Nature reviews Drug discovery. 2012; 11(3):191–200. <https://doi.org/10.1038/nrd3681>.

460 **Schütt K**, Kindermans PJ, Saucedo Felix HE, Chmiela S, Tkatchenko A, Müller KR. Schnet: A continuous-filter con-
461 volutional neural network for modeling quantum interactions. Advances in neural information processing
462 systems. 2017; 30. <https://doi.org/10.48550/arXiv.1706.08566>.

463 **Stärk H**, Ganea O, Pattanaik L, Barzilay R, Jaakkola T. Equibind: Geometric deep learning for drug binding
464 structure prediction. In: *International Conference on Machine Learning* PMLR; 2022. p. 20503–20521. <https://doi.org/10.48550/arXiv.2202.05146>.

466 **Sydow D**, Morger A, Driller M, Volkamer A. TeachOpenCADD: a teaching platform for computer-aided drug
467 design using open source packages and data. Journal of Cheminformatics. 2019; 11(1):29. <https://doi.org/10.1186/s13321-019-0351-x>, doi: 10.1186/s13321-019-0351-x.

469 **Sydow D**, Rodríguez-Guerra J, Kimber T, Schaller D, Taylor C, Chen Y, Leja M, Misra S, Wichmann M, Aria-
470 majd A, Volkamer A. TeachOpenCADD 2022: open source and FAIR Python pipelines to assist in struc-
471 tural bioinformatics and cheminformatics research. Nucleic Acids Research. 2022 Jul; 50:W753–W760.
472 <https://doi.org/10.1093/nar/gkac267>, doi: 10.1093/nar/gkac267.

473 **Tang J**, Szwajda A, Shakyawar S, Xu T, Hintsanen P, Wennerberg K, Aittokallio T. Making sense of large-scale
474 kinase inhibitor bioactivity data sets: a comparative and integrative analysis. Journal of Chemical Information
475 and Modeling. 2014; 54(3):735–743. <https://doi.org/10.1021/ci400709d>.

476 **Tong X**, Liu X, Tan X, Li X, Jiang J, Xiong Z, Xu T, Jiang H, Qiao N, Zheng M. Generative models for De Novo drug
477 design. Journal of Medicinal Chemistry. 2021; 64(19):14011–14027. <https://doi.org/10.1021/acs.jmedchem.1c00927>.

479 **Unke OT**, Meuwly M. PhysNet: A neural network for predicting energies, forces, dipole moments, and partial
480 charges. Journal of chemical theory and computation. 2019; 15(6):3678–3693. <https://doi.org/10.1021/acs.jctc.9b00181>.

482 **Virtanen P**, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser
483 W, Bright J, van der Walt SJ, Brett M, Wilson J, Millman KJ, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E,
484 Carey C, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods. 2020;
485 17(3):261–272. doi: 10.1038/s41592-019-0686-2.

486 **Volkamer A**, Riniker S, Nittinger E, Lanini J, Grisoni F, Evertsson E, Rodríguez-Pérez R, Schneider N. Machine
487 Learning for Small Molecule Drug Discovery in Academia and Industry. Artificial Intelligence in the Life Sci-
488 ences. 2023; p. 100056. <https://doi.org/10.1016/j.ailsci.2022.100056>.

489 **Wang S**, Guo Y, Wang Y, Sun H, Huang J. SMILES-BERT: large scale unsupervised pre-training for molecular
490 property prediction. In: *Proceedings of the 10th ACM international conference on bioinformatics, computational*
491 *biology and health informatics*; 2019. p. 429–436. doi: 10.1145/3307339.3342186.

492 **Weininger D**. SMILES, a chemical language and information system. 1. Introduction to methodology and en-
493 coding rules. Journal of Chemical Information and Computer Sciences. 1988; 28(1):31–36. <https://doi.org/10.1021/ci00057a005>, doi: 10.1021/ci00057a005.

495 **Wieder O**, Kohlbacher S, Kuenemann M, Garon A, Ducrot P, Seidel T, Langer T. A compact review of molecular
496 property prediction with graph neural networks. Drug Discovery Today: Technologies. 2020; 37:1–12. <https://doi.org/10.1016/j.ddtec.2020.11.009>.

498 **Wu R**, Ding F, Wang R, Shen R, Zhang X, Luo S, Su C, Wu Z, Xie Q, Berger B, et al. High-resolution de novo structure
499 prediction from primary sequence. BioRxiv. 2022; p. 2022–07. <https://doi.org/10.1101/2022.07.21.500999>.

500 **Wu Z**, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, Leswing K, Pande V. MoleculeNet: a
501 benchmark for molecular machine learning. Chemical science. 2018; 9(2):513–530. <https://doi.org/10.1039/C7SC02664A>.

- 503 Xia X, Hu J, Wang Y, Zhang L, Liu Z. Graph-based generative models for de Novo drug design. Drug Discovery
504 Today: Technologies. 2019; 32:45–53. <https://doi.org/10.1016/j.ddtec.2020.11.004>.
- 505 Xu K, Hu W, Leskovec J, Jegelka S. How powerful are graph neural networks? arXiv preprint arXiv:181000826.
506 2018; <https://doi.org/10.48550/arXiv.1810.00826>.
- 507 Zhao L, Zhu Y, Wang J, Wen N, Wang C, Cheng L. A brief review of protein-ligand interaction prediction. Com-
508 putational and Structural Biotechnology Journal. 2022; <https://doi.org/10.1016/j.csbj.2022.06.004>.