UNIVERSITY
OF SKÖVDE

The need for Bayesian deep learning models in the search for novel medical drugs

University of Skövde

Recent trends in drug discovery

- As with many fields, cheminformatics is currently going through a deep learning revolution.
- Previously, the most common approach to property prediction was to extract features from each molecule, and then apply machine learning models.
- Now it is common to use graph convolutional neural networks to predict properties of molecules Kearnes et al. (2016); Duvenaud et al. (2015).

Recent trends in drug discovery

- There has also been an effort to use reinforcement learning to generate novel molecules.
- Using a reinforcement learning framework, molecules are generated and then evaluated using a predictive model and the generative model gets rewarded for molecules which have desirable properties.
- Examples of authors using this approach are Olivecrona et al. (2017) and Popova et al. (2018).

Reinforcement framework

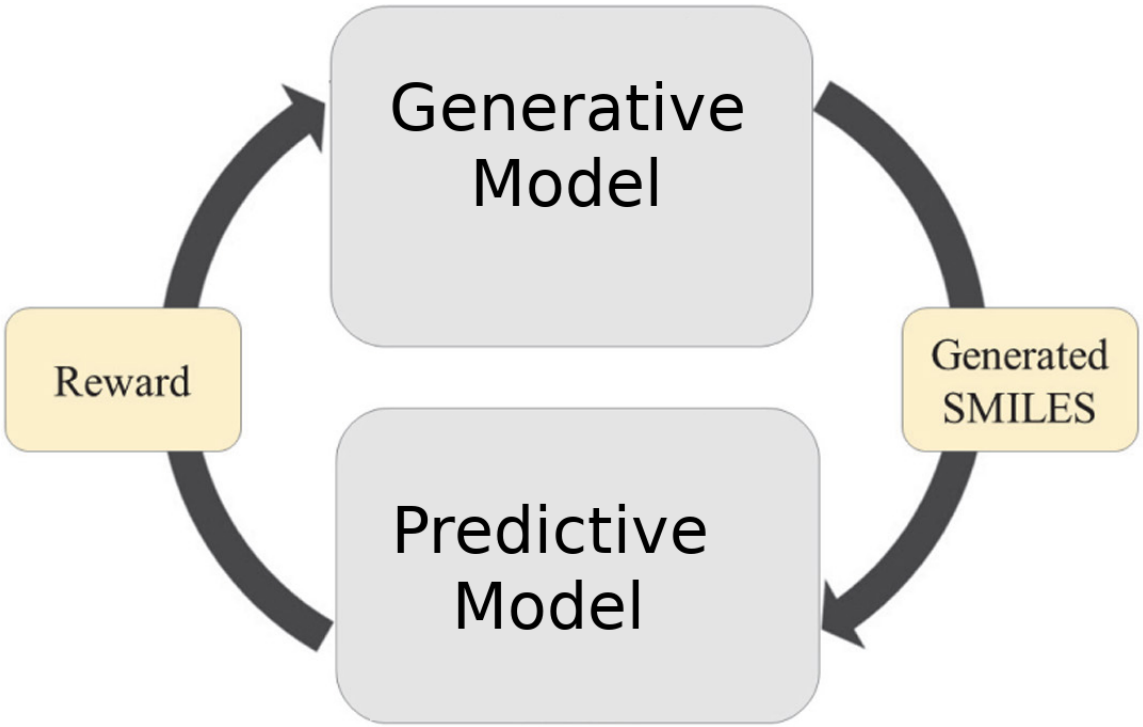

Figure: A reinforcement learning framework for the generation of new drugs.

With this presentation I would like to show:

- Why it is risky to use generative models trained through reinforcement learning, especially when generating new molecules.

- How this risk can be decreased by the use of Bayesian neural networks.

Extrapolation - A short side track

- Which is the next number in the sequence $1, 3, 5, 7, \ldots$?
- My deep neural network says: 217341.
- Because it learned the function:
  $f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 4527773x + 217331$
- Which gives $f(1) = 1$, $f(2) = 3$, $f(3) = 5$ and $f(4) = 7$.

Extrapolation - A short side track

- Which is the next number in the sequence $1, 3, 5, 7, \ldots$?
- My deep neural network says: 217341.
- Because it learned the function:
  $f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 4527773x + 217331$
- Which gives $f(1) = 1$, $f(2) = 3$, $f(3) = 5$ and $f(4) = 7$.

Extrapolation - A short side track

- Which is the next number in the sequence $1, 3, 5, 7, \ldots$?
- My deep neural network says: 217341.
- Because it learned the function:
  $f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 4527773x + 217331$
- Which gives $f(1) = 1$, $f(2) = 3$, $f(3) = 5$ and $f(4) = 7$.

Extrapolation - A short side track

- Which is the next number in the sequence $1, 3, 5, 7, \ldots$?
- My deep neural network says: 217341.
- Because it learned the function:
  $f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 4527773x + 217331$
- Which gives $f(1) = 1$, $f(2) = 3$, $f(3) = 5$ and $f(4) = 7$.

Extrapolation and molecules

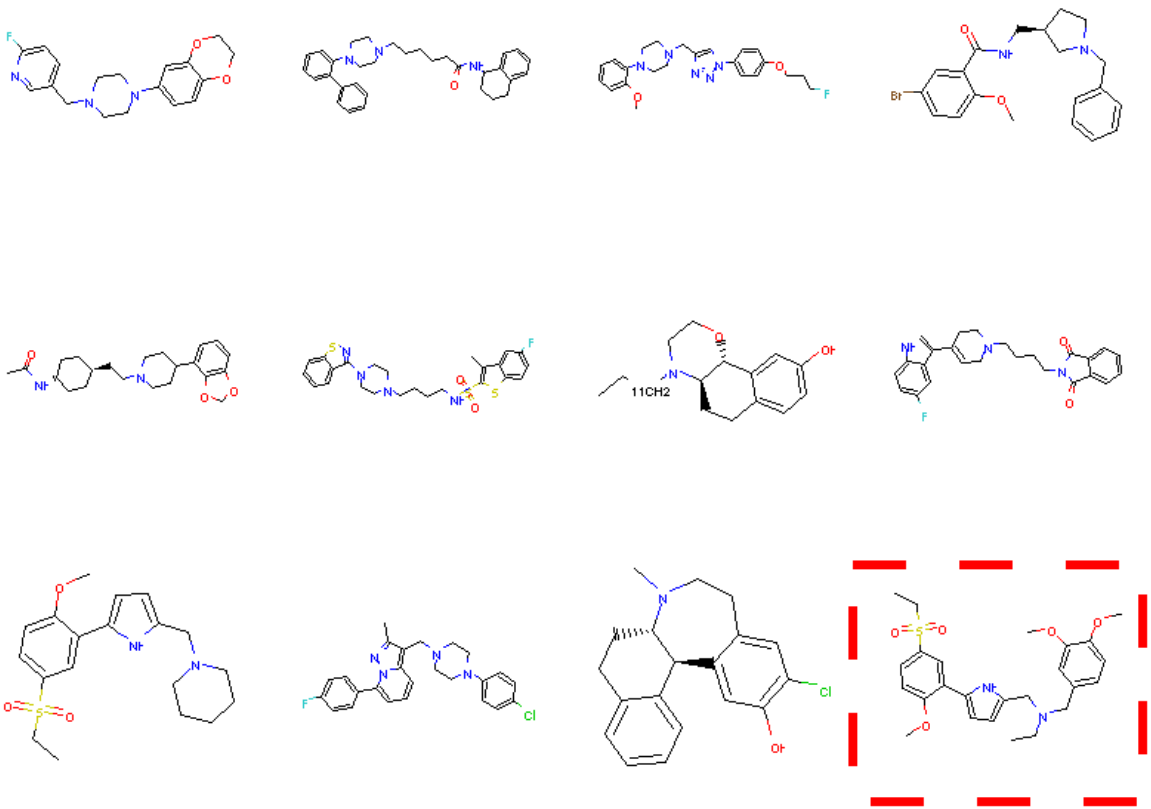

Figure: Will a model, trained on the first eleven molecules, interpolate or extrapolate the value of the last molecule.

Drift in generative models

- The generative model may exploit the extrapolation ability of the predictive model.

- This is undesirable and we would instead want the generative model to make use of what the predictive model has learned from the training data.

Proposed approach

- If we build multiple models using the same training data, they should provide approximately the same predictions close to the training data, while they may extrapolate in different ways.
- Therefore, we use a Bayesian graph convolutional neural network and sample internal weights for the network from the posterior distribution $P(W|Data)$.
- This allows us to measure the uncertainty in predictions, which may be used to prevent the drift for the generative model.
- See Gal and Ghahramani (2016) for more information.

Experiments

- We conducted three separate experiments, concerning molecules with known equilibrium constants with dopamine D2.
- The data were collected from ChEMBLE: https://www.ebi.ac.uk/chembl/
- There are two subtypes of molecules, tricyclic and non-tricyclic molecules.

How the data are splitt

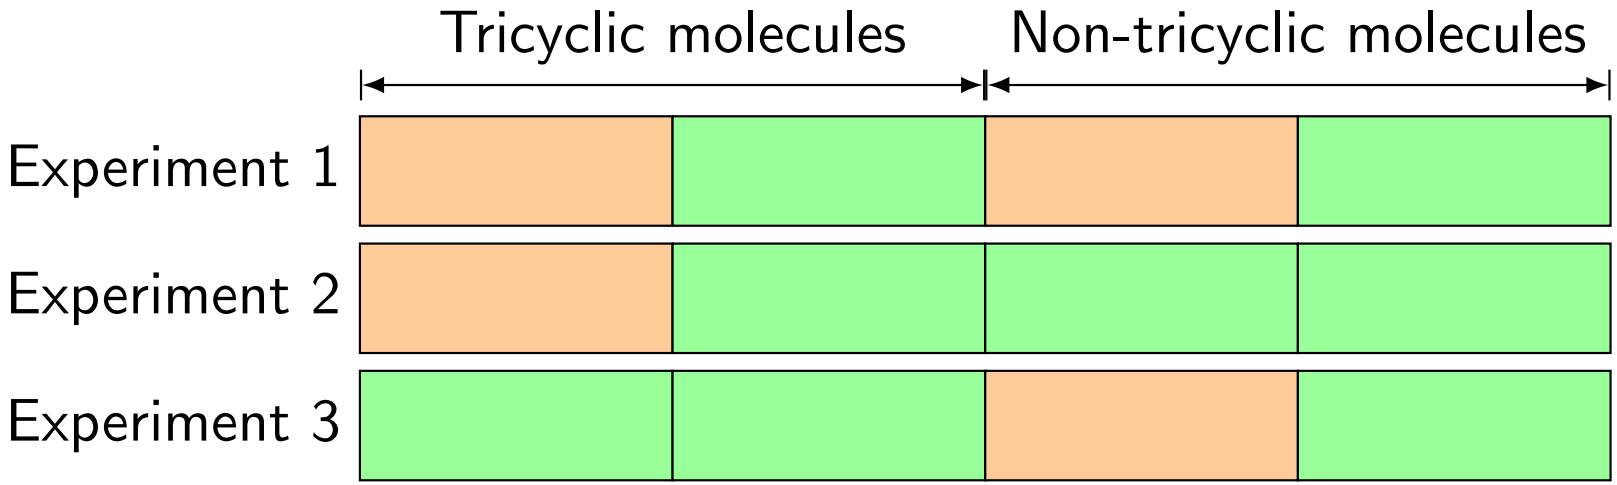

Figure: Orange stands for training data, while green represents training data.

Result from experiment 1

When training on a mix of tricyclic and non-tricyclic molecules:

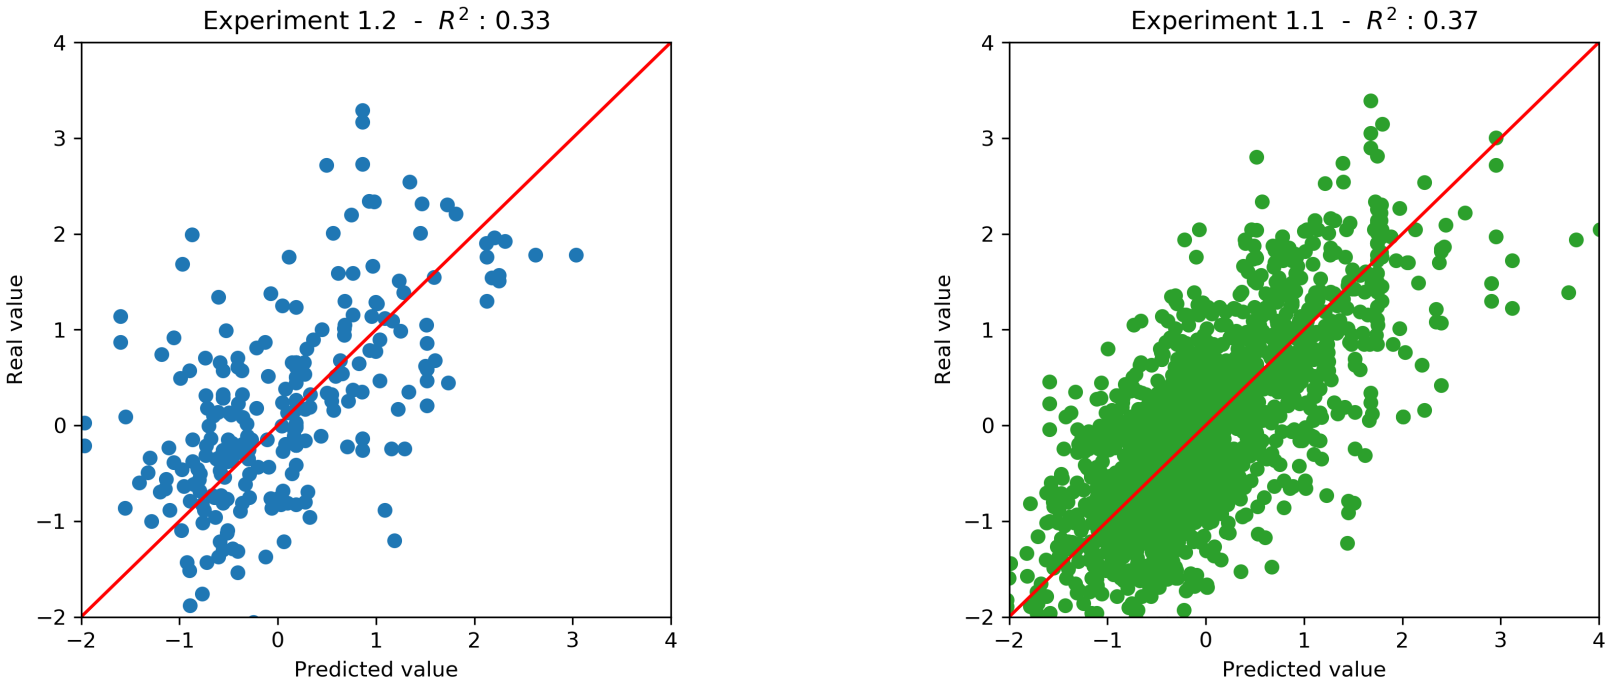

Figure: The real values over the predicted values for the tricyclic test set (blue) and the non-tricyclic test set (green).

Result from experiment 2

When training on only tricyclic molecules:

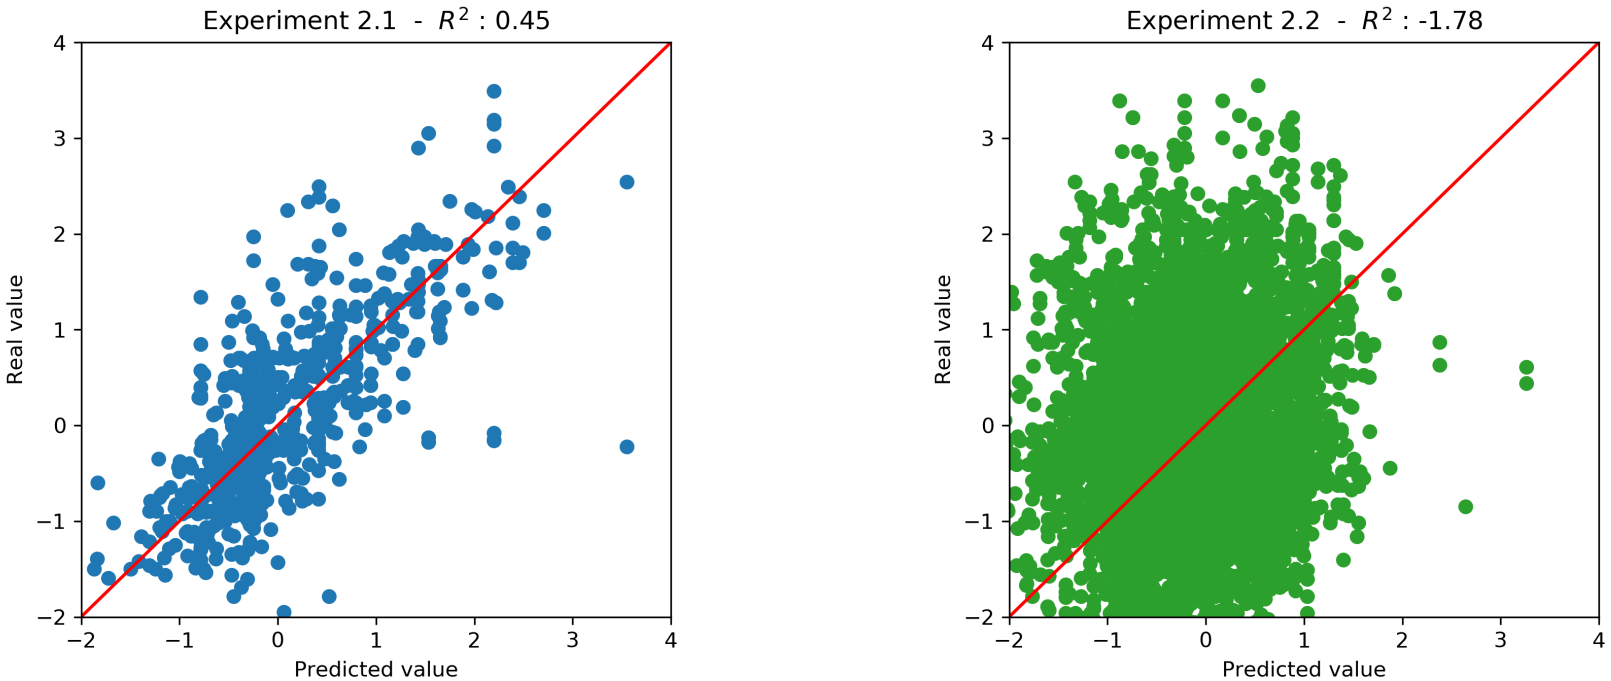

Figure: The real values over the predicted values for the tricyclic test set (blue) and the set of all non-tricyclic molecules (green).

Result from experiment 3

When training on only non-tricyclic molecules:

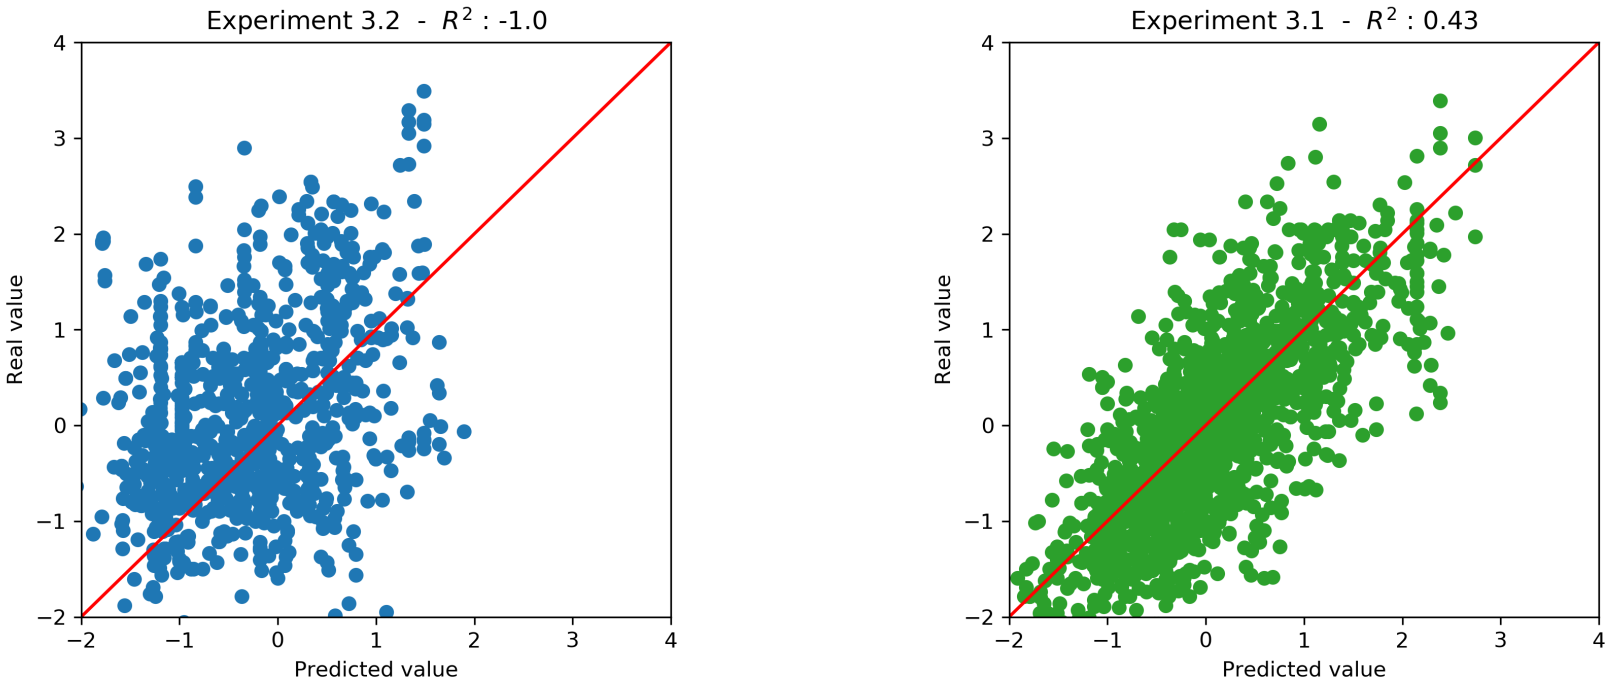

Figure: The real values over the predicted values for the set of all tricyclic molecules (blue) and the non-tricyclic test set (green).

Uncertainty of the predictions

When training on only tricyclic molecules:

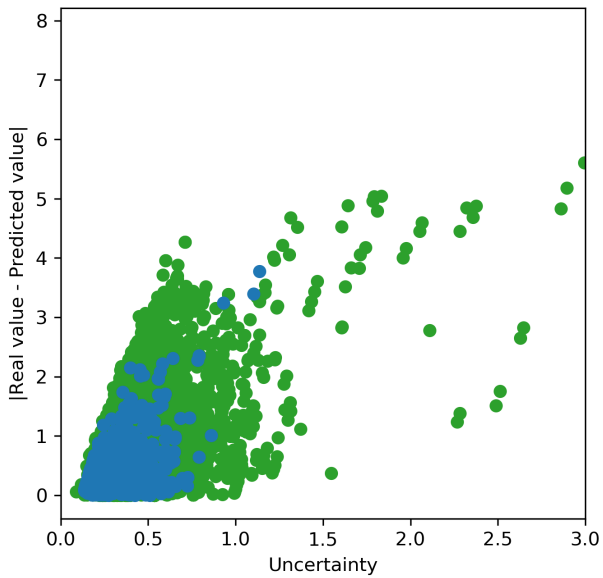

Figure: The absolute distance between the predicted value and the real value plotted over the model uncertainty for the second experiment. Here we can see that the uncertainty grows when the difference between the real value and the predicted value grows.

Conclusions

- Extrapolation is always dangerous, but especially when working with non-euclidean data.
- Bayesian deep learning and averaging predictions over several possible model parameters can allow us to detect when the model is uncertain (and thus inconsistent) about its predictions.
- When using generative models in combination with predictive models it is important to stay in the same domain as the training data.

References I

Duvenaud, D. K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. (2015). Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*, pages 2224–2232.

Gal, Y. and Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059.

Kearnes, S., McCloskey, K., Berndl, M., Pande, V., and Riley, P. (2016). Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8):595–608.

References II

Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. (2017). Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):48.

Popova, M., Isayev, O., and Tropsha, A. (2018). Deep reinforcement learning for de novo drug design. *Science advances*, 4(7):eaap7885.