Human-in-the-loop active learning for goal-oriented molecule generation

Presented by **Yasmine Nahal**

AIDD – AiChemist School 2024

Goal-oriented molecule generation

- **Optimizing** molecular structures against specific properties

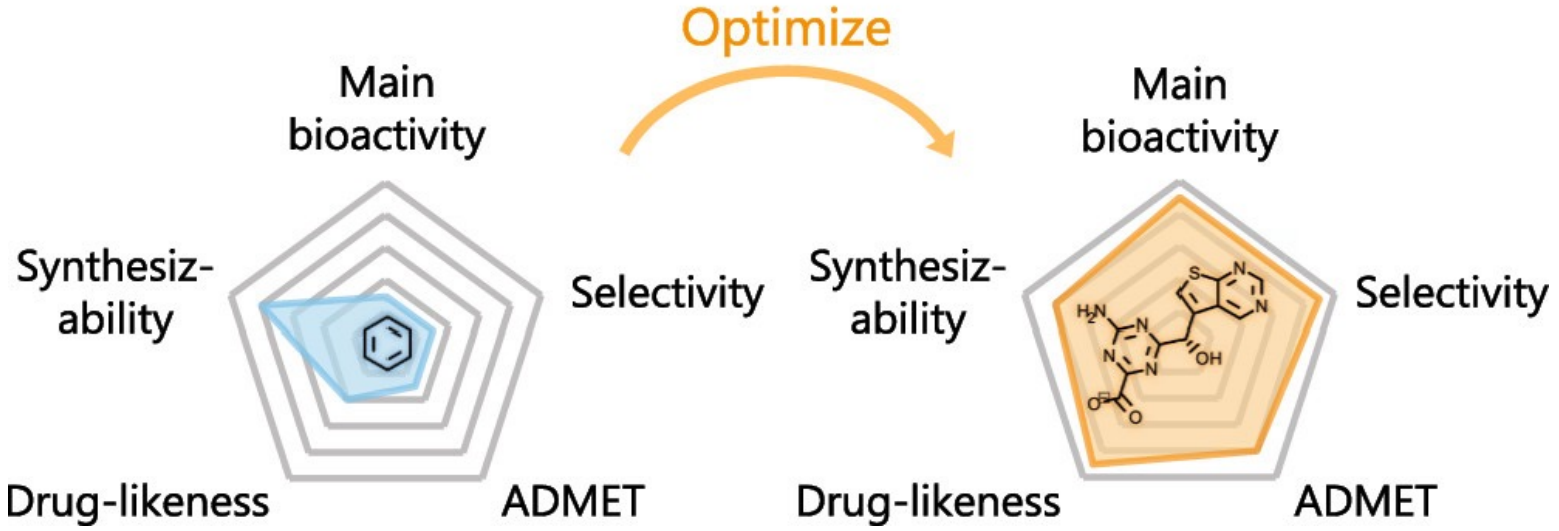

Goal-oriented molecule generation

- **Optimizing** molecular structures against specific properties
- **Complex properties** (e.g., bioactivity) may be modelled from **experimental data**

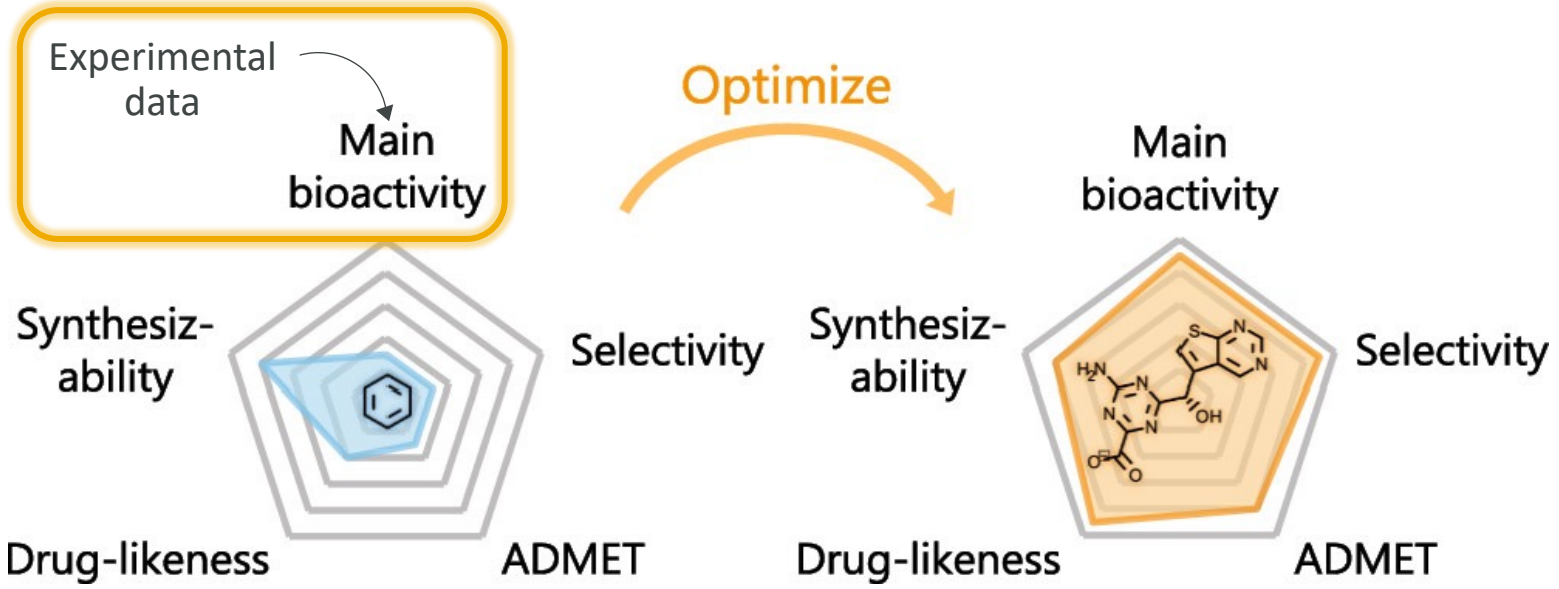

Goal-oriented molecule generation

- **Optimizing** molecular structures against specific properties
- **Complex properties** (e.g., bioactivity) may be modelled from **experimental data**

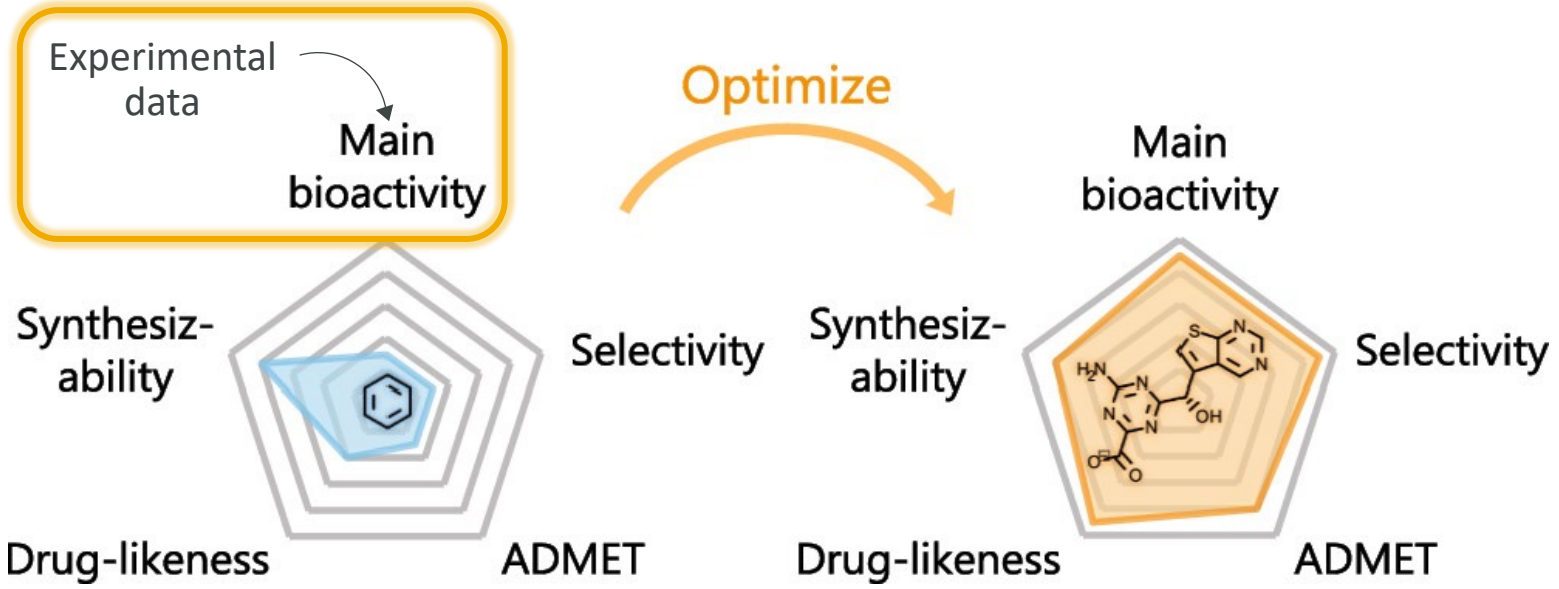

- Molecules optimized against **approximated values for these properties** instead of ground truth values

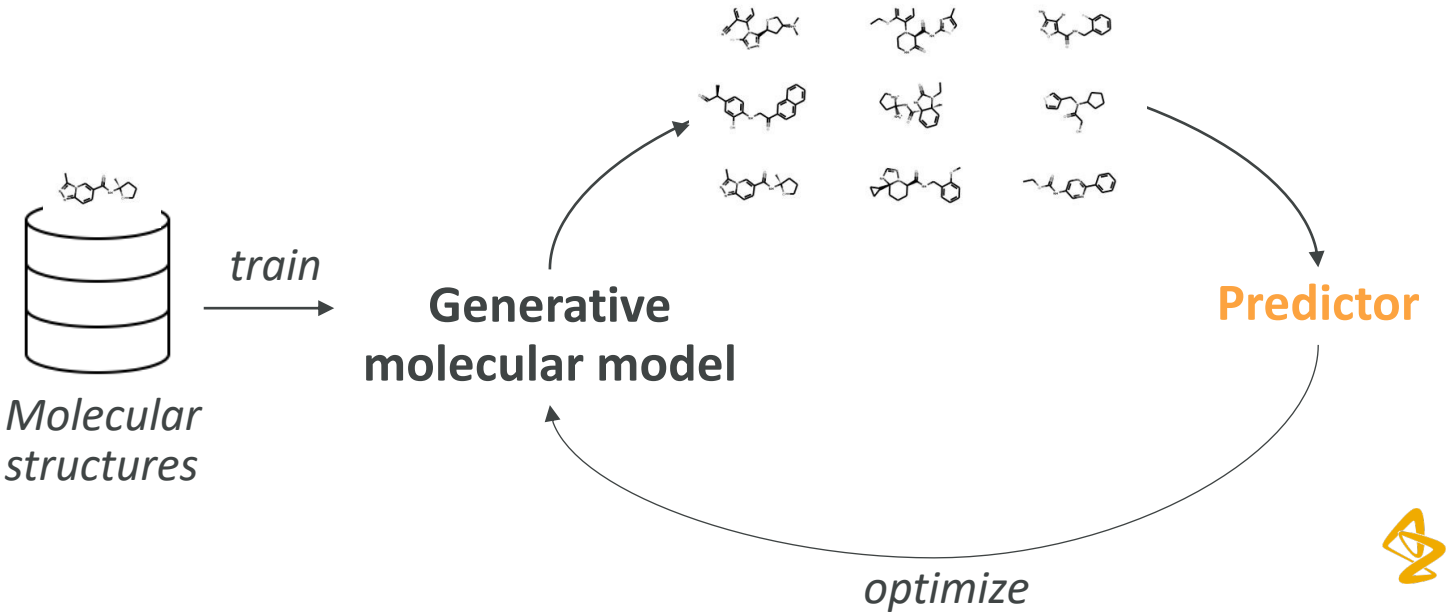

Goal-oriented molecule generation: failure modes

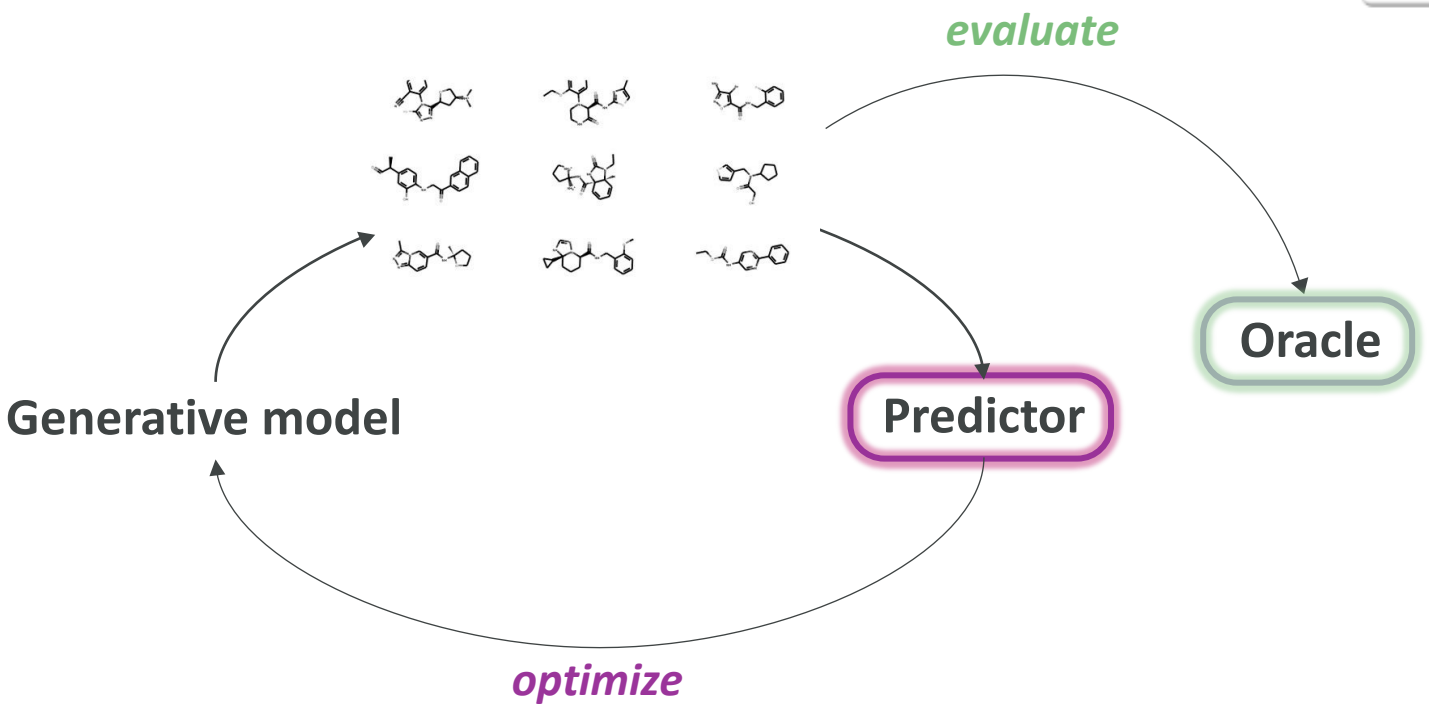

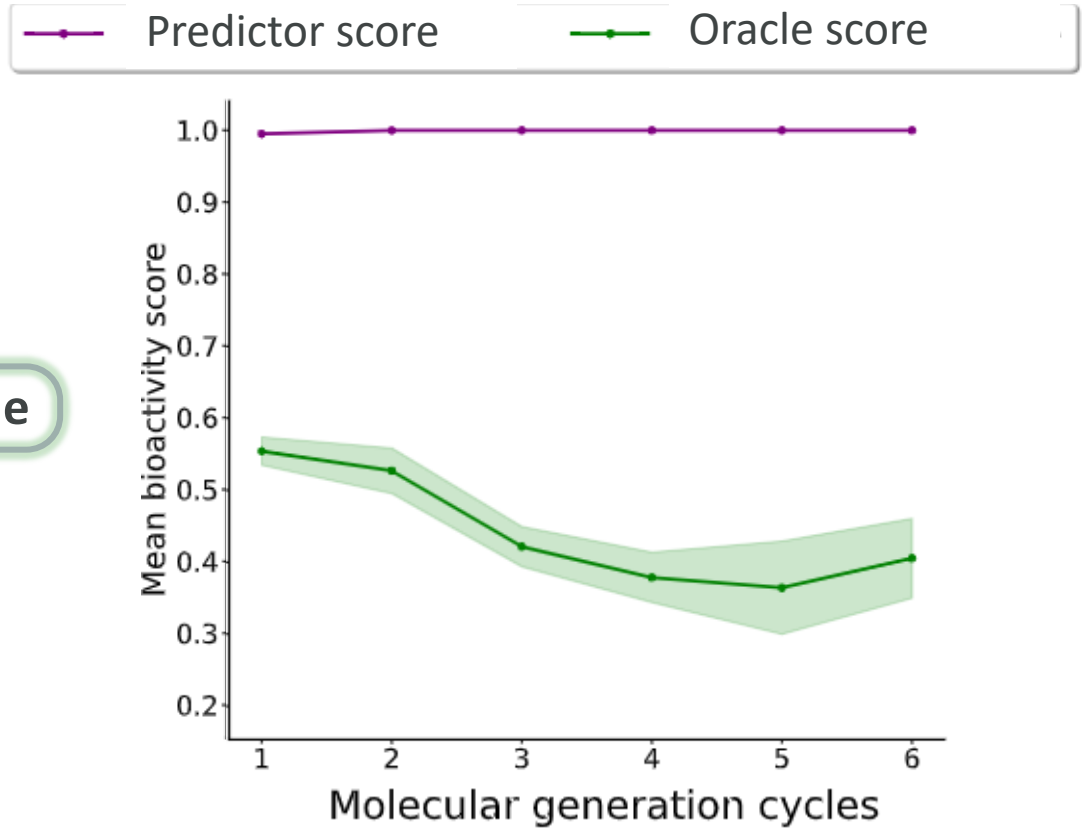

"when molecules predicted to be good are not that good"

- **Predictor hacking** was first described by *Renz et al., 2019*

Goal-oriented molecule generation: failure modes

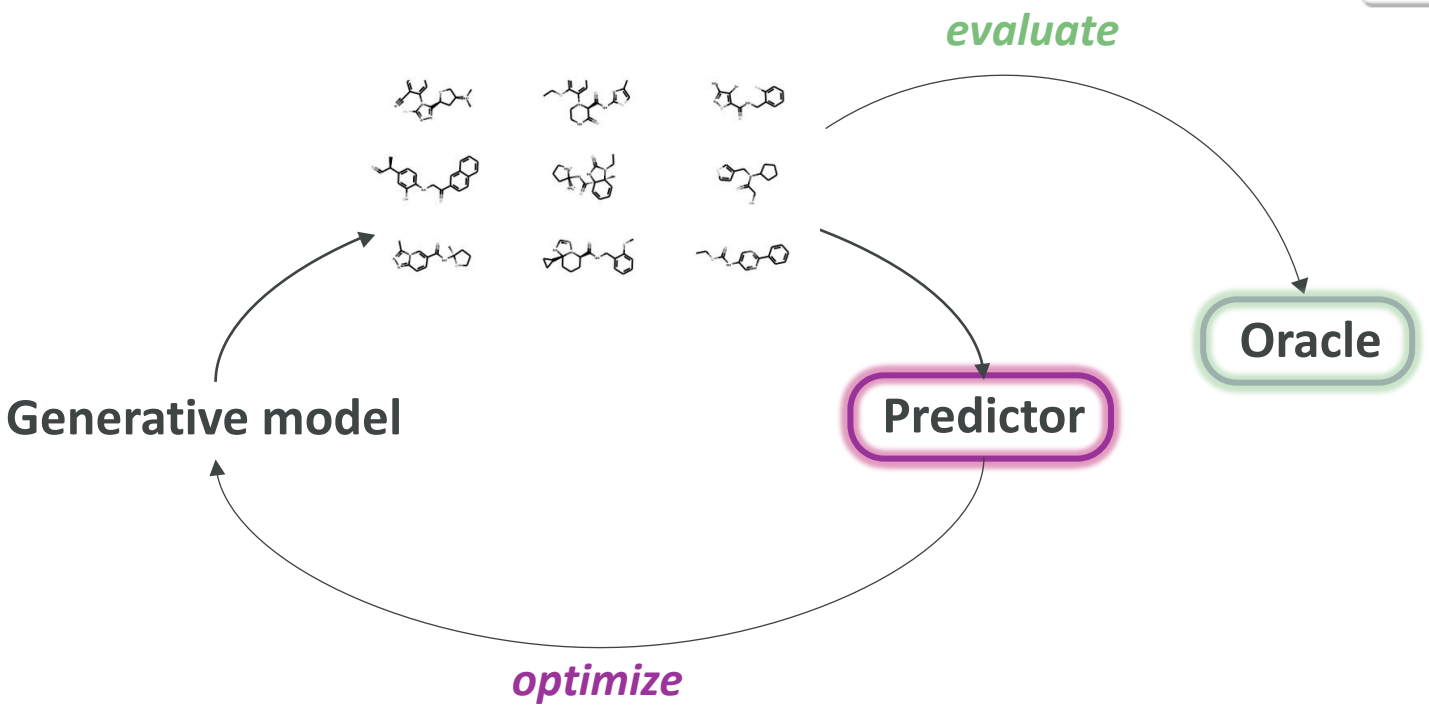

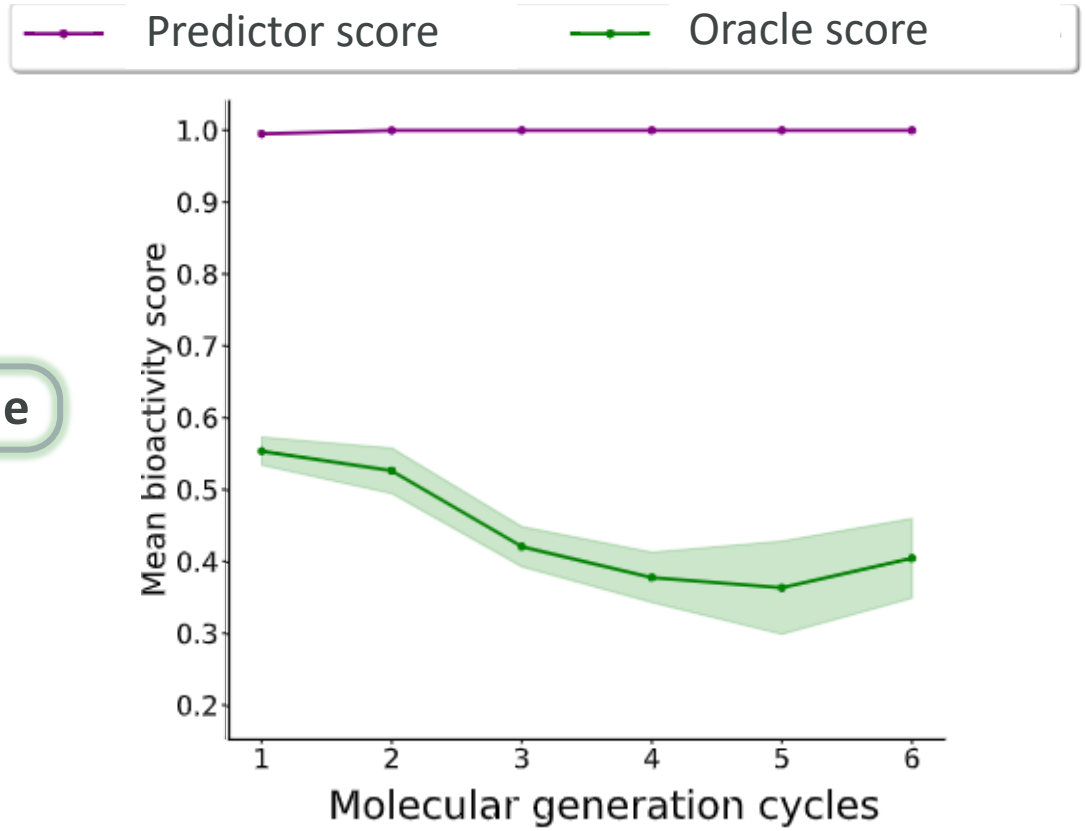

Proposed solutions and perspectives:

- Building more conservative predictors *(Fu et al. 2021)*
- Applying similarity constraints w.r.t. predictor training set *(Griffiths et al. 2022)*
- ***Active learning*** *(Bilodeau et al. 2022)*

How to use Active Learning?

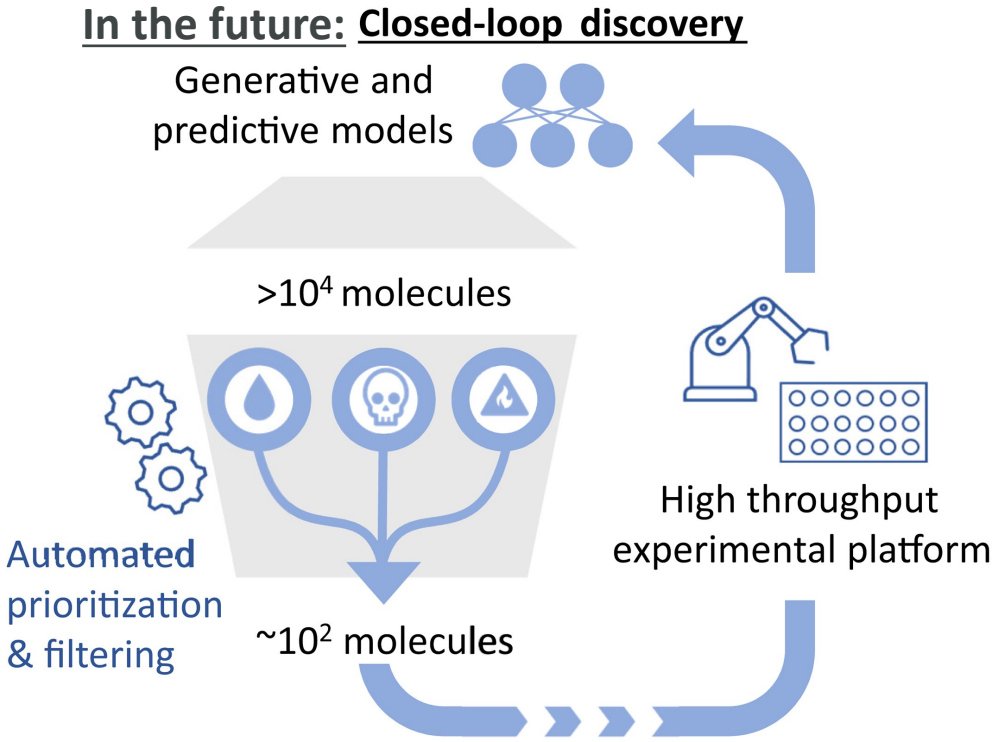

(Bilodeau et al. 2022)

- **Active learning** to fine-tune predictors **requires direct experimental testing** (currently unfeasible)

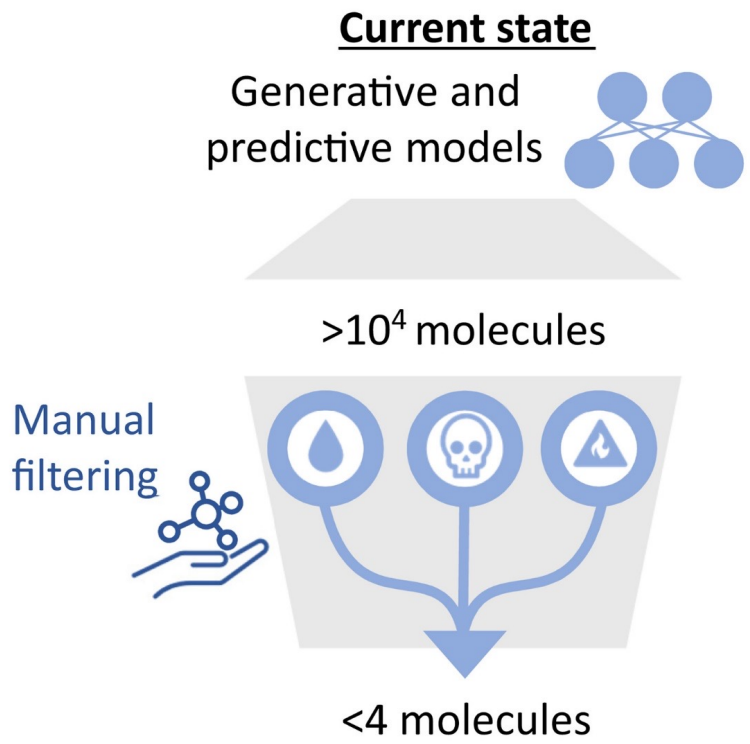

(Bilodeau et al. 2022)

- **Chemists** are often involved in **prioritizing compounds for experimental testing**

Human-in-the-loop Active Learning to assist Molecule Generation

Can we leverage **interactive learning from human knowledge** to reduce manual post-filtering efforts ?

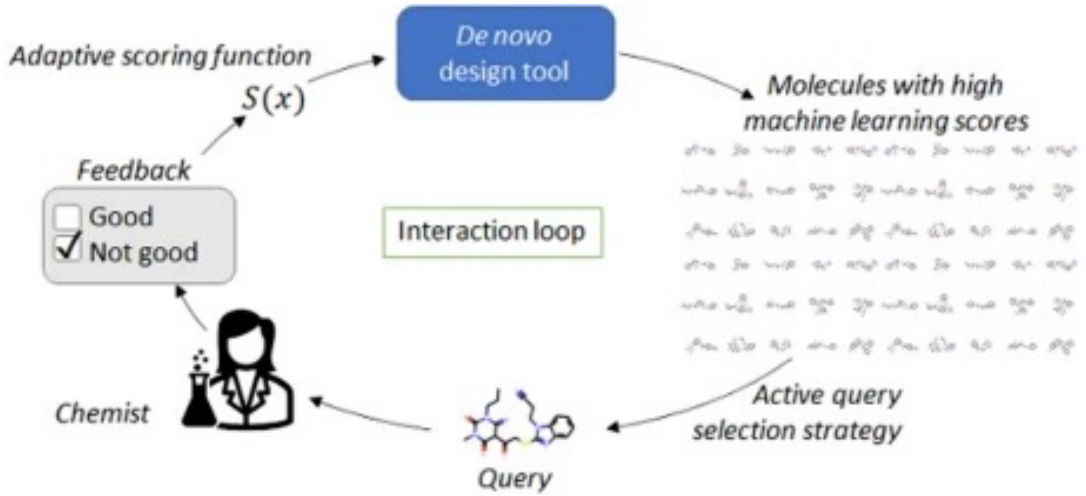

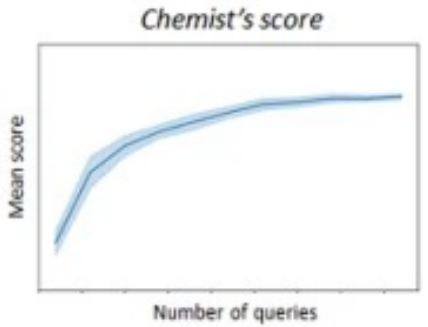

We can learn a user model that efficiently captures human preferences on molecules

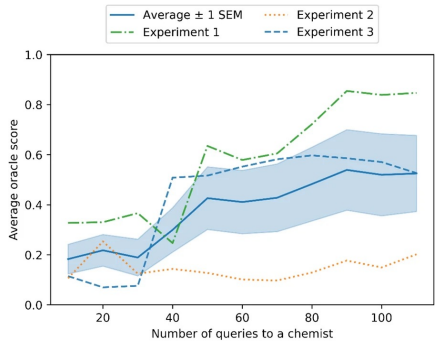

Optimizing molecules against user's DRD2 bioactivity score improves oracle score

Sundin, I., Voronov, A., Xiao, H. *et al.* Human-in-the-loop assisted de novo molecular design. *J Cheminform* 14, 86 (**2022**).

Human-in-the-loop Active Learning to assist Molecule Generation

Can we leverage **interactive learning from human knowledge** to mitigate the predictor hacking issue ?

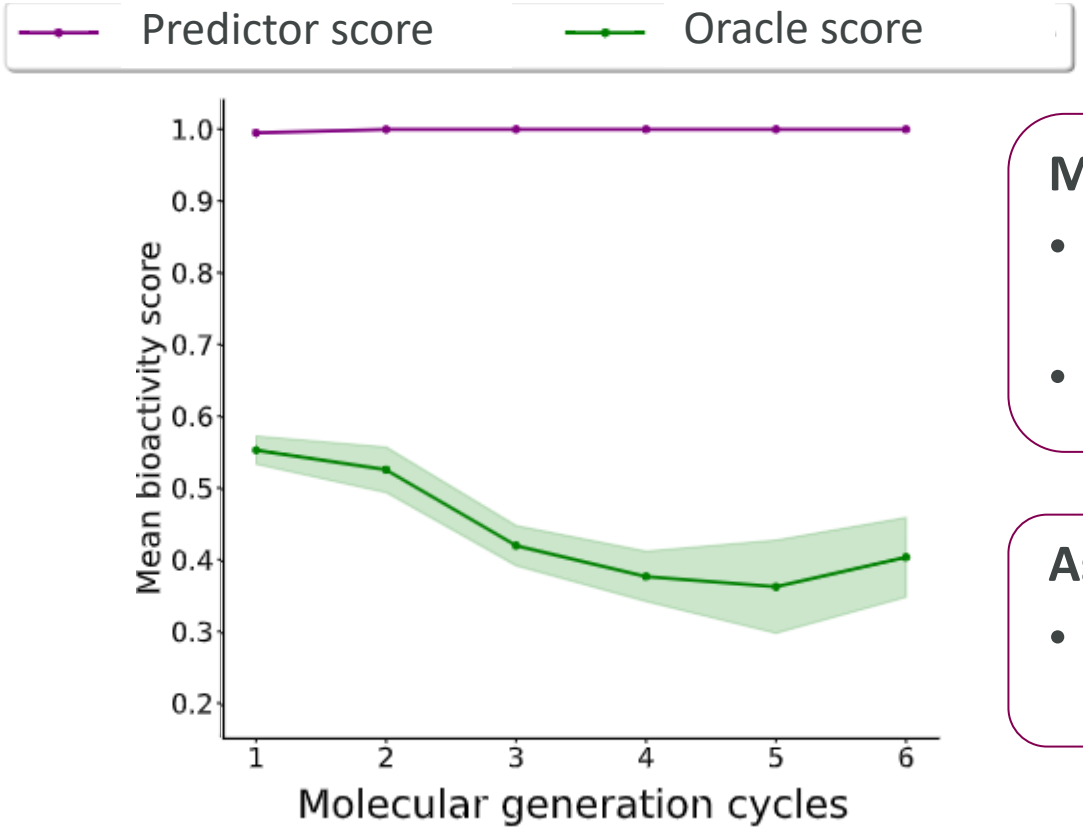

Mitigating predictor hacking:

- Reducing predictor overconfidence via aligning predicted and oracle scores
- Improving TPR within high-scoring molecules

Assumption:

- Humans can provide better calibrated judgements which when optimized are less prone to hacking

Gendreau, P., Turk, J-A., Drizard, N. *et al.* Molecular Assays Simulator to Unravel Predictors Hacking in Goal-Directed Molecular Generations. *Journal of Chemical Information and Modeling* 13, 63 (**2023**).

Method

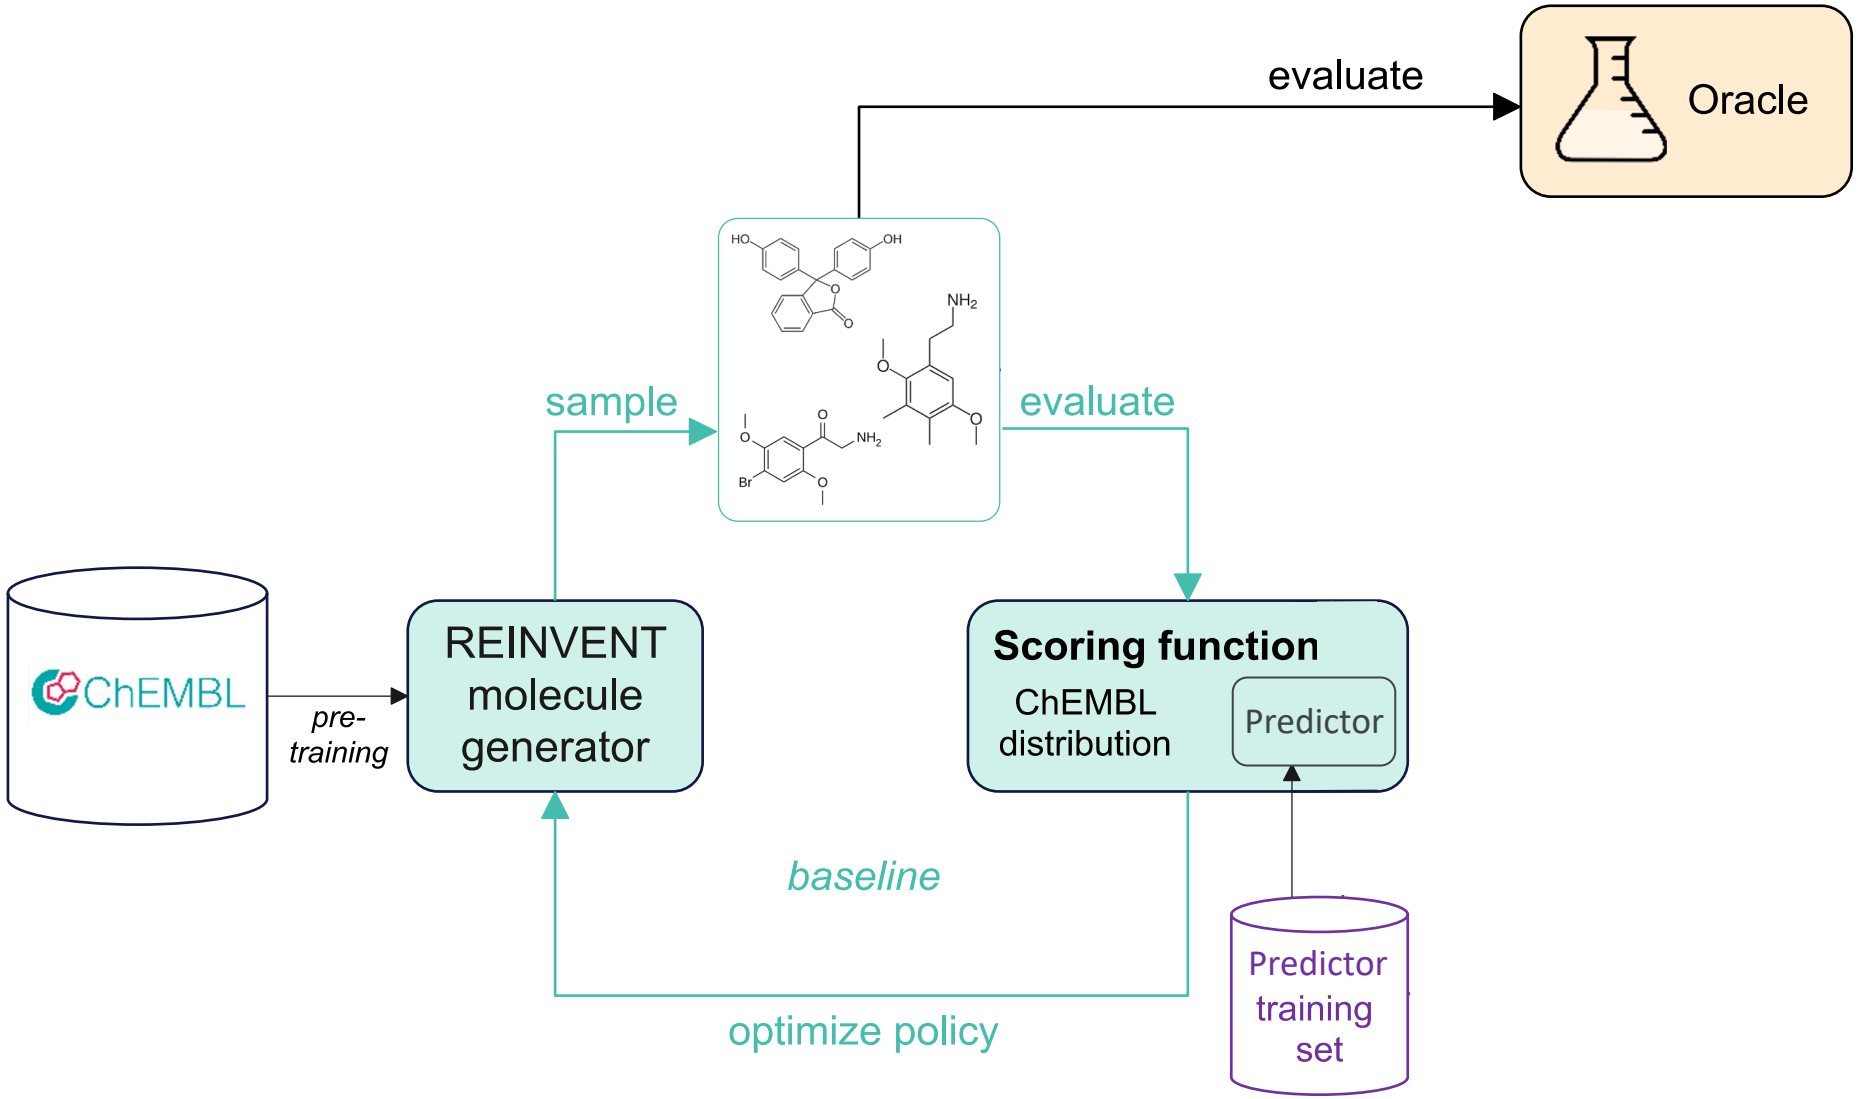
evaluate
Oracle
sample
evaluate
ChEMBL
pre-
training
REINVENT
molecule
generator
Scoring function
ChEMBL
distribution
Predictor
baseline
Predictor
training
set
optimize policy

Method

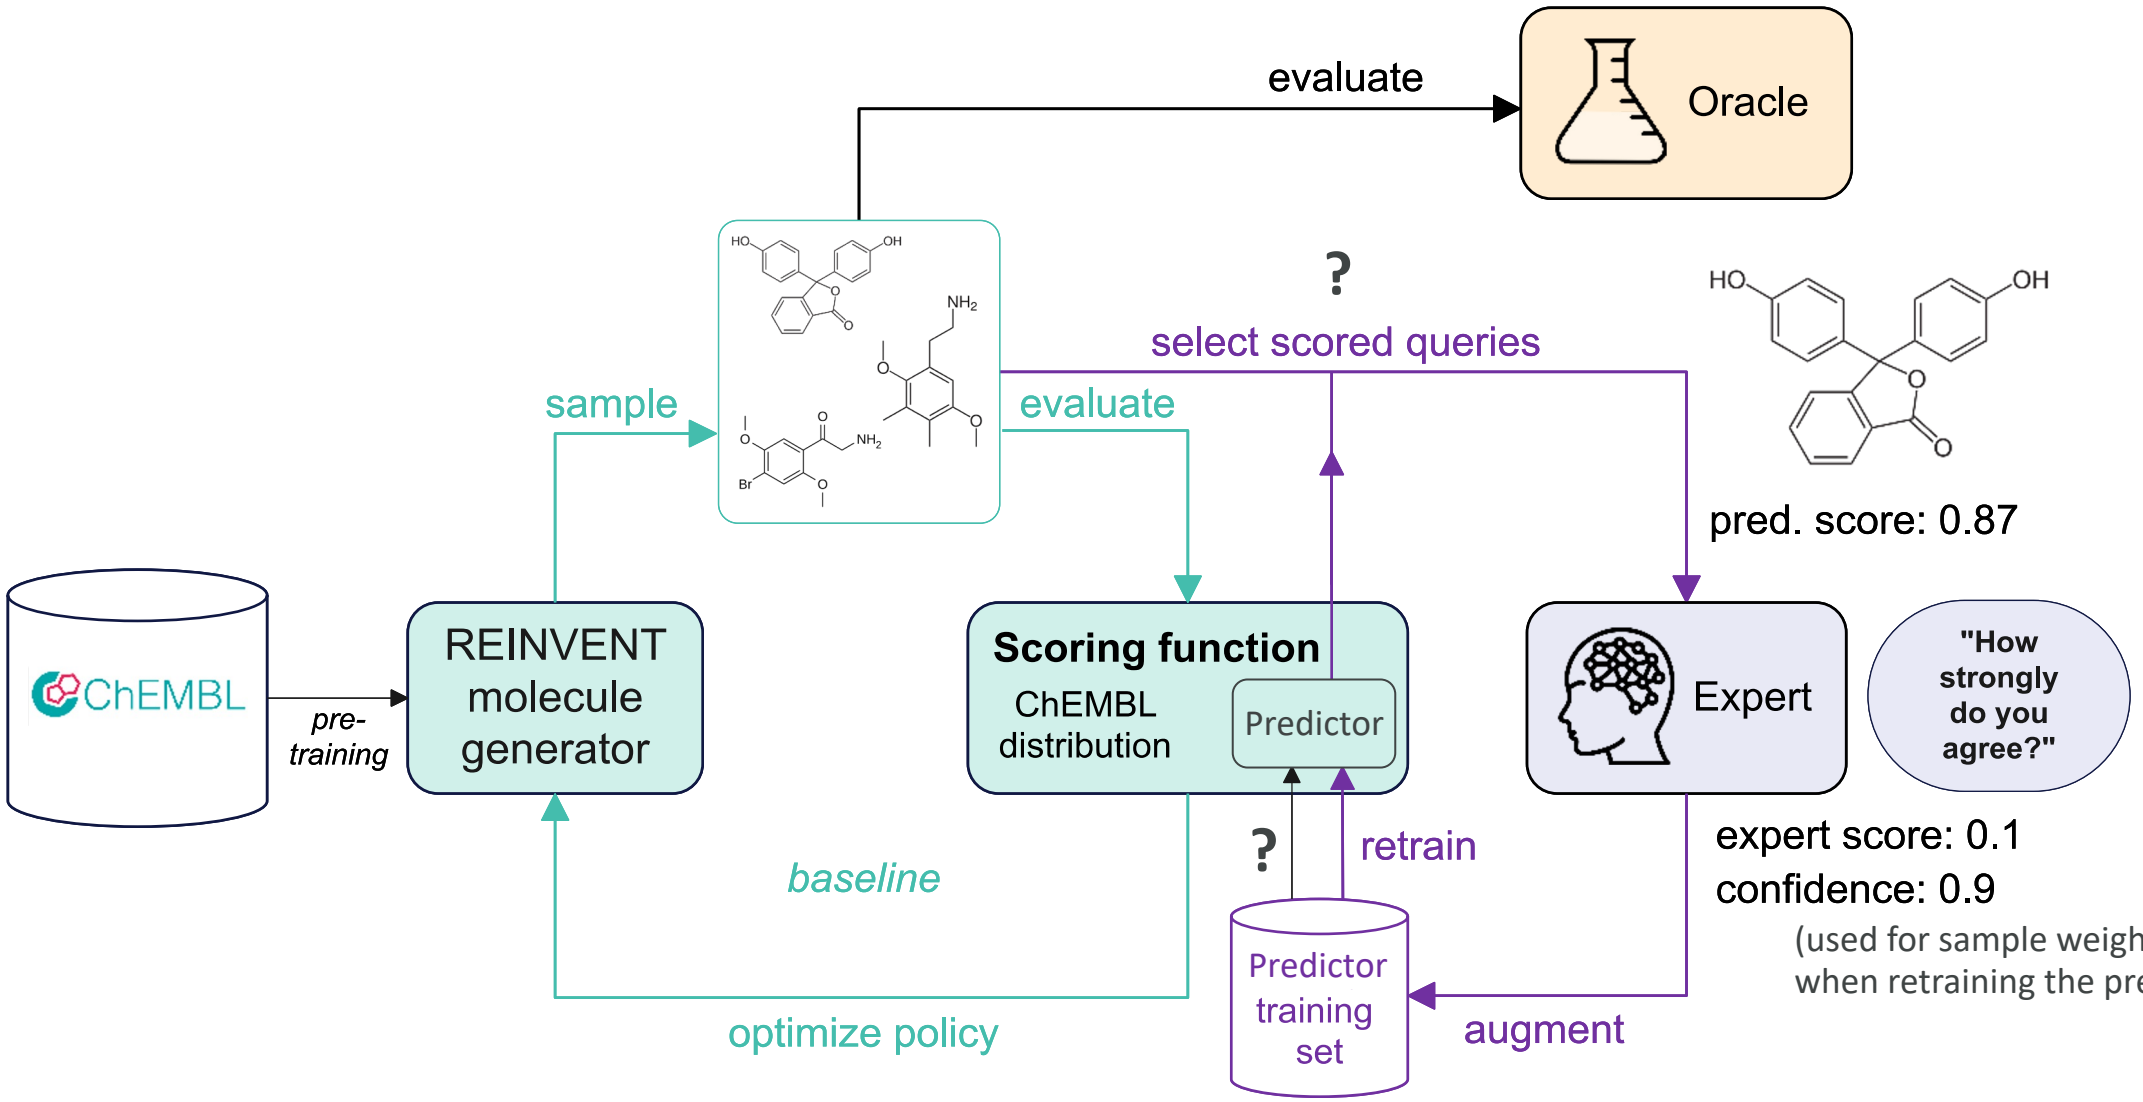
evaluate
Oracle
?
select scored queries
sample
evaluate
pred. score: 0.87
ChEMBL
pre-
training
REINVENT
molecule
generator
Scoring function
ChEMBL
distribution
Predictor
Expert
"How
strongly
do you
agree?"
expert score: 0.1
confidence: 0.9
(used for sample weighting
when retraining the predictor)
?
retrain
baseline
Predictor
training
set
augment
optimize policy

Predictor training

- Drug discovery projects usually start with small available experimental data.

Predictor training

- Drug discovery projects usually start with small available experimental data.

- Random Forest models demonstrate better resilience against hacking by generative agents compared to linear and “almost linear” models.

Ding, Y., Wang, L., Zhang, H., Yi, J., Fan, D., Gong, B. Defending Against Adversarial Attacks Using Random Forest. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops* (**2019**).

Active query selection

Queries to the expert are selected based on the current predictor $\widehat{\boldsymbol{\theta}}$.

1. Expected Predictive Information Gain (EPIG)

Select experiments that are expected to contribute with the greatest reduction in predictive uncertainty at a given target input distribution after being observed

$$(\mathbf{x}_t, y_t) = \underset{\mathbf{x}\in\mathcal{U}}{\operatorname{argmax}} \; \mathbb{E}_{p_\star(\mathbf{x}_\star)p_{\hat{\boldsymbol{\theta}}}(y|\mathbf{x})}\left[\mathrm{H}[p_{\hat{\boldsymbol{\theta}}}(y_\star \mid \mathbf{x}_\star)] - \mathrm{H}[p_{\hat{\boldsymbol{\theta}}}(y_\star \mid \mathbf{x}_\star, \mathbf{x}, y)]\right]$$

2. Greedy

Select instances with the highest predicted scores

$$\mathbf{x}_t = \underset{\mathbf{x}\in\mathcal{U}}{\operatorname{argmax}} \; f_{\hat{\boldsymbol{\theta}}}(\mathbf{x})$$

3. Uncertainty (Query By Committee)

Select instances with the highest entropy within individual predictions

$$\mathbf{x}_t = \underset{\mathbf{x}\in\mathcal{U}}{\operatorname{argmax}} \; \frac{1}{B}\sum_{i=1}^{B} \mathbb{1}\left[f_{\hat{\theta}_i}(\mathbf{x}) = 1\right]$$

4. Uniform random sampling

Bickford Smith, F., Kirsch, A., Farquhar, S., *et al*. Prediction-oriented Bayesian active learning. *International Conference on Artificial Intelligence and Statistics* (**2023**).

Experiments

1. Simulated HITL experiments

Simple use case: optimize molecules against bioavailability ($\mathrm{LogP} \in [2,4]$)

More complex use case: optimize molecules against DRD2 bioactivity

2. Real HITL experiments

Optimize molecules against DRD2 bioactivity in a multi-objective setting
(+ hERG inactivity and QED)

3. Benchmarking and studying the effect of different hyperparameters

Setup for simulated experiments

1. Oracle training

- **Data:** 101039 (1039) molecules from ExCAPE-DB, activity cutoff pIC50 = 7.3
- **Pre-processing:** count ECFP6 of length 2048
- **Predictor:** RF Classifier (5-fold CV to set the HP values)
- **Predictive performance (Train/Test):** scaffold-based splitting, 0.99/0.99 **ROC AUC**, 0.97/0.63 **PR AUC**

2. Predictor training

- **Data:** 240 (62) molecules from ExCAPE-DB reduced to small chemical space (2 topological scaffolds)
- **Pre-processing:** count ECFP6 of length 2048, labels given by the oracle (94% coverage original labels)
- **Predictor:** RF Classifier (5-fold CV to set the HP values)
- **Predictive performance (Train only):** 0.99 **ROC AUC**, 0.99 **PR AUC**

3. Simulated expert

- **Expert score:** $f_{\text{human}}(\mathbf{x}_t) = g\left(f^{\star}(\mathbf{x}_t) + \epsilon\right) \in [0,1], \quad \epsilon \sim \mathcal{N}(0, \sigma_\epsilon)$

Oracle probability of being active

- **Number of selected queries:** 10

Results

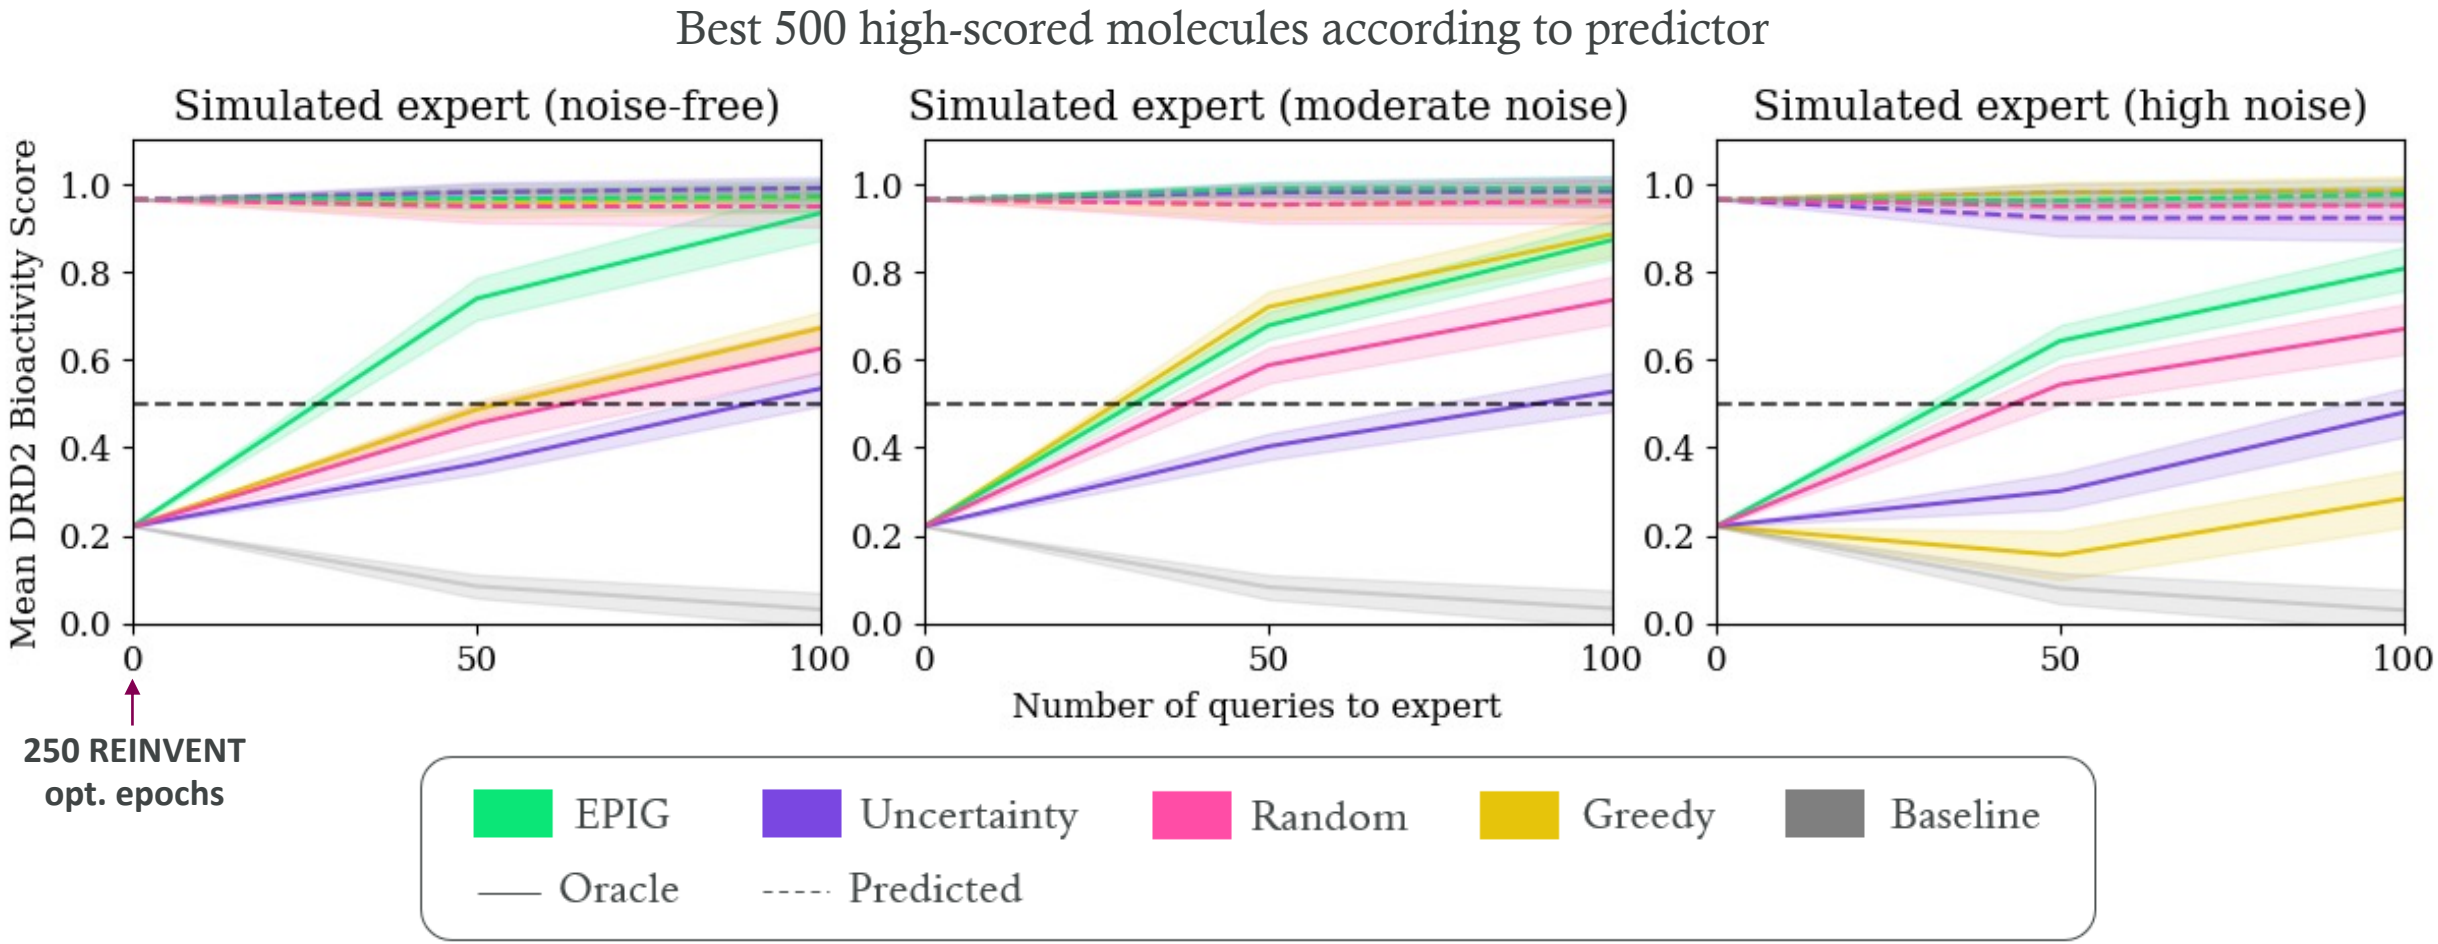

★ Incorporating well-aligned expert feedback via AL results in better alignment between predicted and oracle bioactivity scores and higher true positive rate among predictor- high-scored molecules.

★ Predictor that uses EPIG as a data acquisition strategy is robust in presence of noisy inputs.

Setup for real user experiments

- **Query selection strategy:** EPIG
- **Same settings for REINVENT**, except that initial set of generated molecules is sampled from an agent optimized through 1200 epochs instead of 250 epochs
- REINVENT connected with **Metis GUI**:

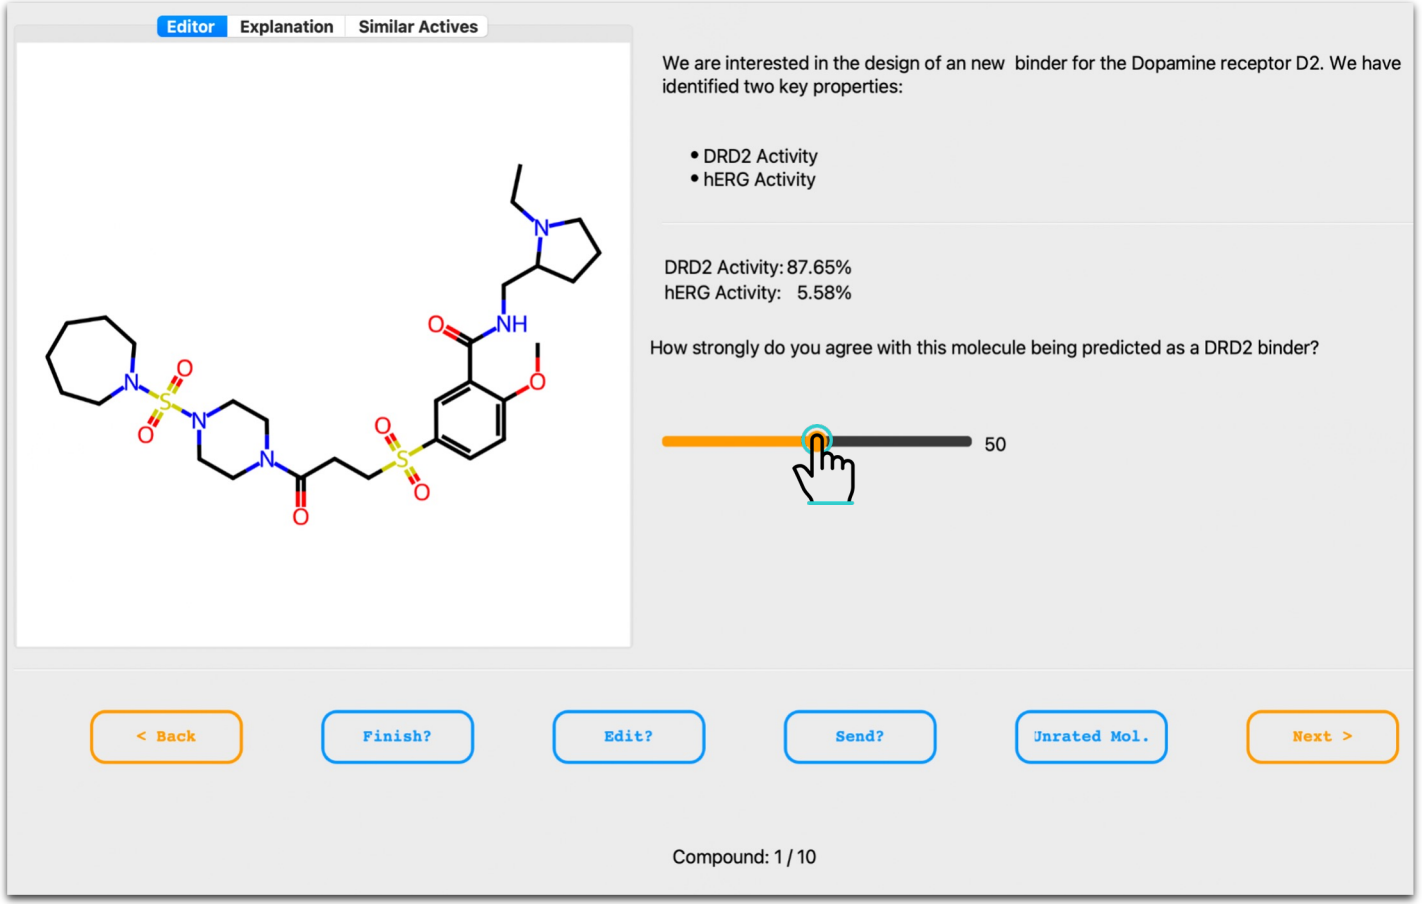

Esben Jannik Bjerrum. An open-source molecular editor based using Python, PySide2 and RDKit

Setup for real user experiments

- **Query selection strategy:** EPIG
- **Same settings for REINVENT**, except that initial set of generated molecules is sampled from an agent optimized through 1200 epochs instead of 250 epochs
- REINVENT connected with **Metis GUI**:

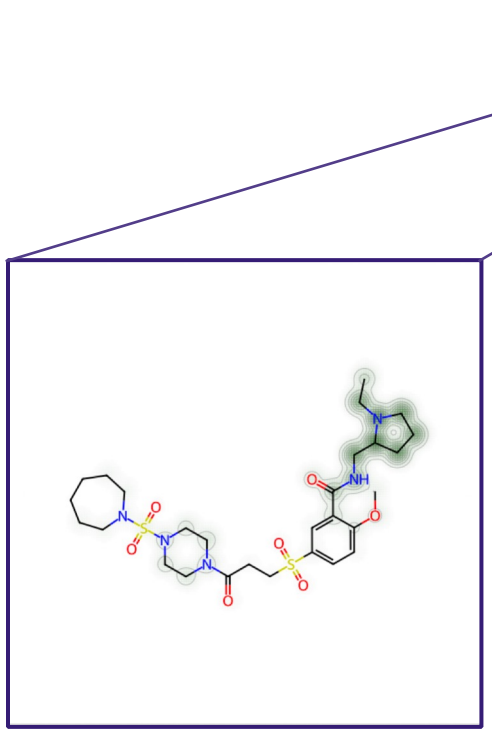

Atom-based heatmaps to explain individual predictions

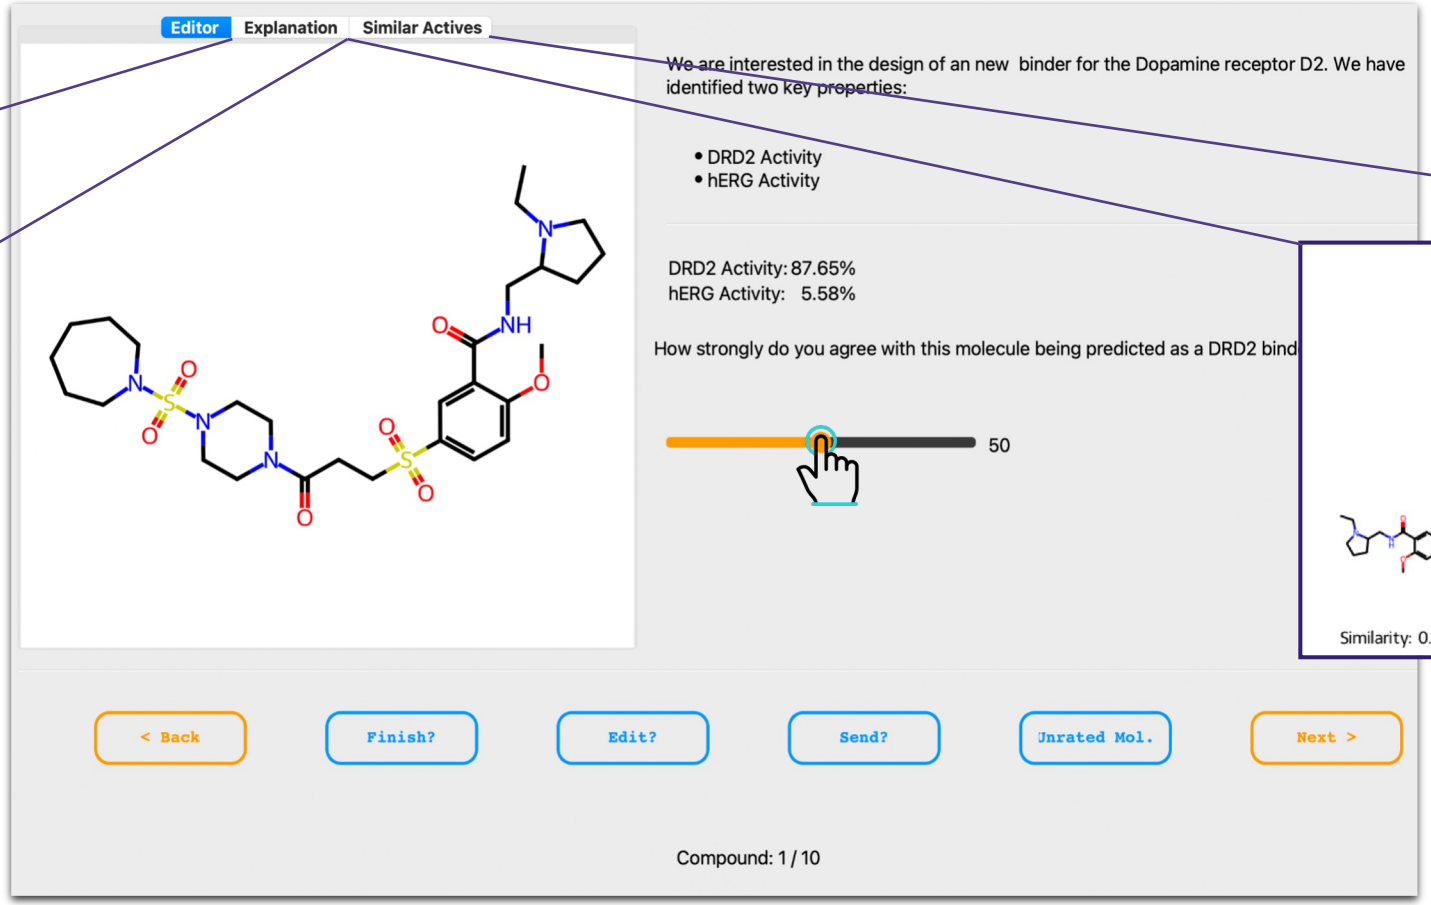

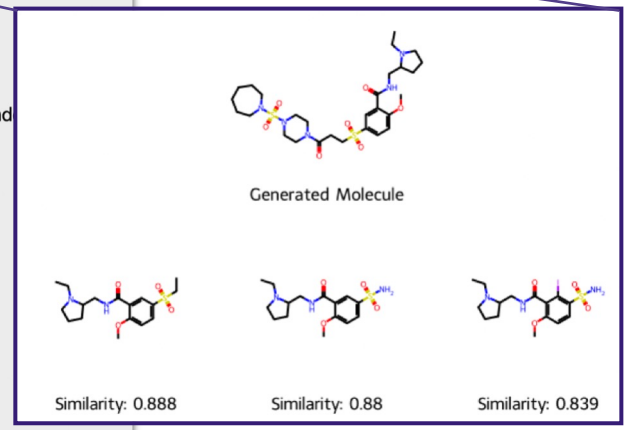

MACCS keys-based search algorithm for displaying the most similar existing actives

Esben Jannik Bjerrum. An open-source molecular editor based using Python, PySide2 and RDKit

Results

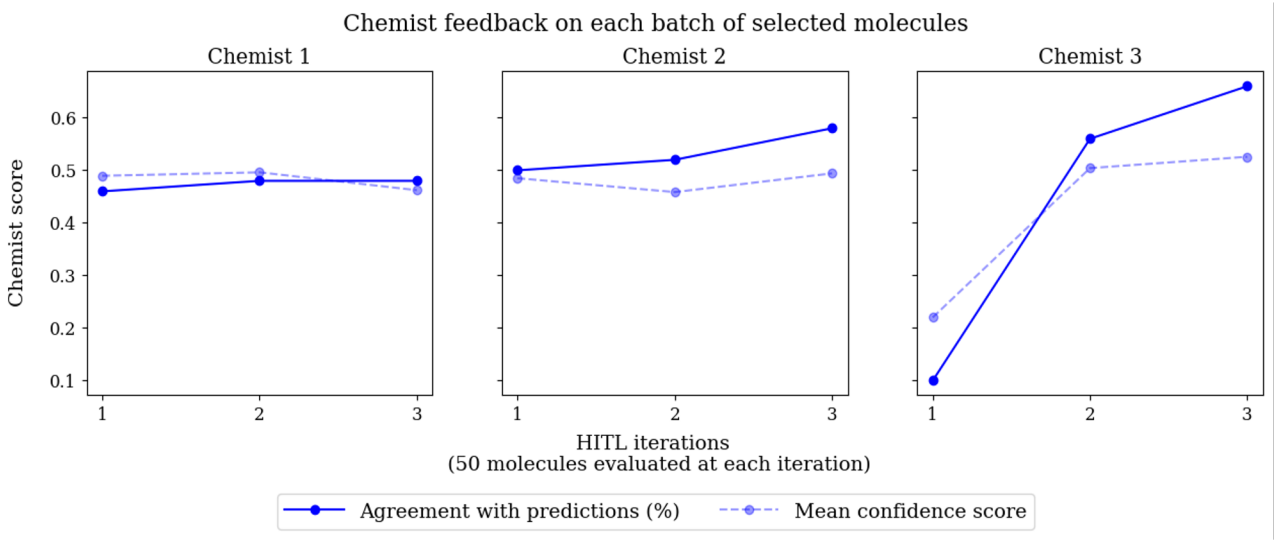

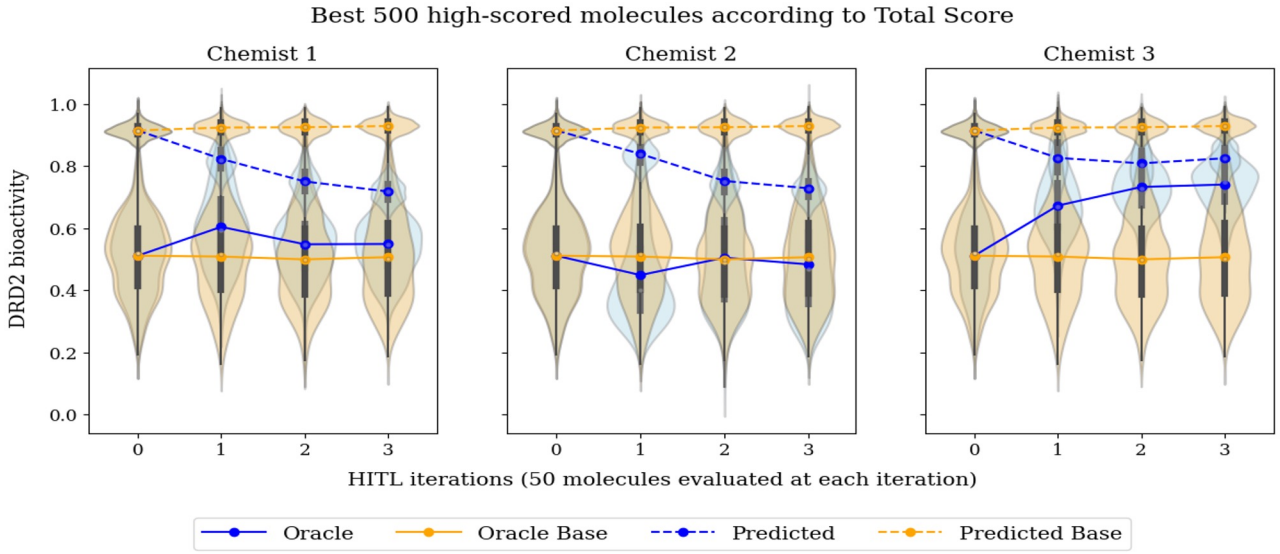

★ Chemists 1 and 2 understood the exercise similarly: **"prioritizing easy-to-make molecules"** due to the question asked that was rather vague

★ The question was reformulated for chemist 3 to focus on the predicted DRD2 bioactivities

★ Reduced predictor overconfidence

★ Significant improvement of TPR with chemist 3 who proves to be better aligned with the DRD2 activity oracle

Results of real user experiments

Final set of generated molecules. Best 500 high-scoring molecules according to Total Score.

Metric (mean)	No feedback	With feedback on DRD2 activity		
		Chemist 1	Chemist 2	Chemist 3
DRD2 Predicted score	0.93	0.80 *	0.81 *	0.84 *
DRD2 Oracle score	0.50	0.55	0.49	**0.74** *
Mean Absolute Error	0.42	0.25 *	0.32 *	**0.10** *
QED score	0.57	0.61 *	0.58	**0.71** *
hERG inactivity score	0.91	0.88	0.90	0.82
Total Oracle score*	0.68	0.69	0.67	**0.77** *
Internal Diversity	0.47	0.41	0.45	0.44
SA score	3.04	**2.75** *	**2.82** *	3.08
Novelty	1.0	1.0	1.0	1.0
Validity	1.0	1.0	1.0	1.0
Uniqueness	1.0	1.0	1.0	1.0

QED score also significantly improved

Chemists 1 and 2 improved SA significantly

Polykovskiy, D., Zhebrak, A., et al. Molecular sets (MOSES): A benchmarking platform for molecular generation models. Frontiers in Pharmacology, 11 (**2020**).

Ertl, P., Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. Journal of cheminformatics, 22 (**2009**).

Wrapping up

- ★ AL can be employed during molecule generation to prevent the predictor hacking issue
- ★ AL results in higher true positive rate among predictor- high-scored molecules
- ★ Querying domain experts in AL proves to be successful for aligning generated molecules with real-world practical applications *(calibrated predictions, improved QED and SA scores)*
- ★ The question asked to experts must convey the goal of calibrating the predictor to avoid ambiguity
- ★ EPIG as query selection strategy is robust to noisy inputs

Acknowledgements

Thanks for all the support and valuable contributions!

Janosch Menke
Markus Heinonen
Mikhail Kabeshov
Jon Paul Janet
Eva Nittinger
Ola Engkvist
Samuel Kaski
Julien Martinelli
Iiris Sundin

Results: comparison with existing solutions