Human-in-the-loop Active Learning for Goal-oriented Molecule Generation

Nahal Y., Menke J., Heinonen M., Kabeshov M., Nittinger Eva., Janet J.P., Engkvist O., Kaski S.

1. Motivation

Predictor hacking issue in goal-oriented molecule generation

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

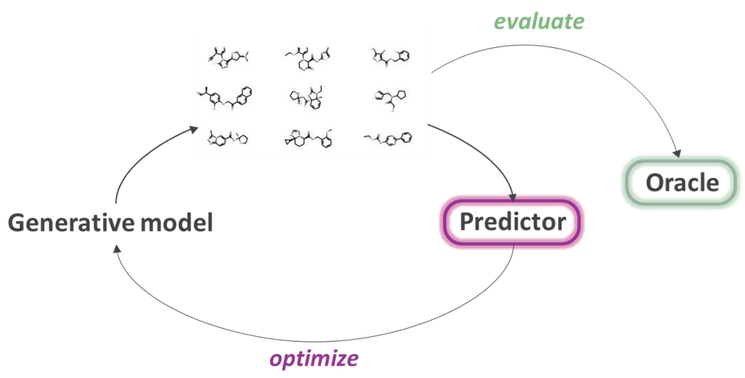

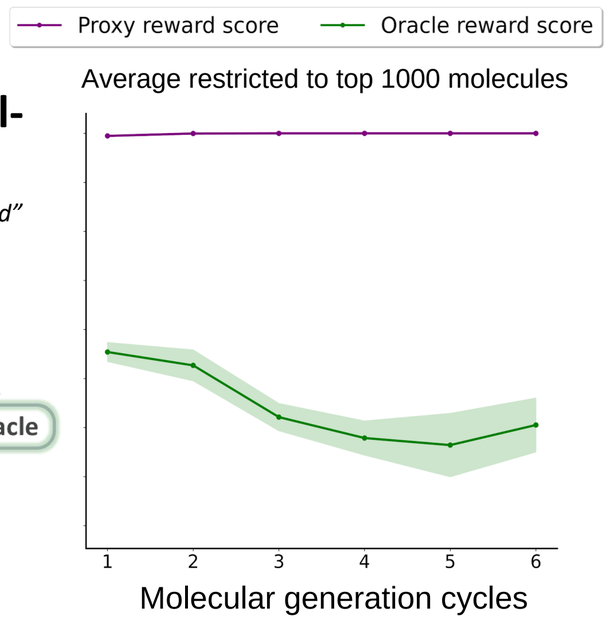

Related work

- Predictor hacking issue *(Renz et al. 2019, Langevin et al., 2022, Gendreau et al. 2023)*
- Constraining generative models w.r.t. training sets *(Griffiths et al. 2022)*
- Building more conservative predictors *(Fu et al. 2021)*

Fixing predictor hacking with Active Learning, but experiments are costly! *(Bilodeau et al. 2022)*

Active Learning + Domain Expertise : a potential way to mitigate predictor hacking with reduced cost

2. Fixing Predictor Hacking with Active Learning and Expert Knowledge

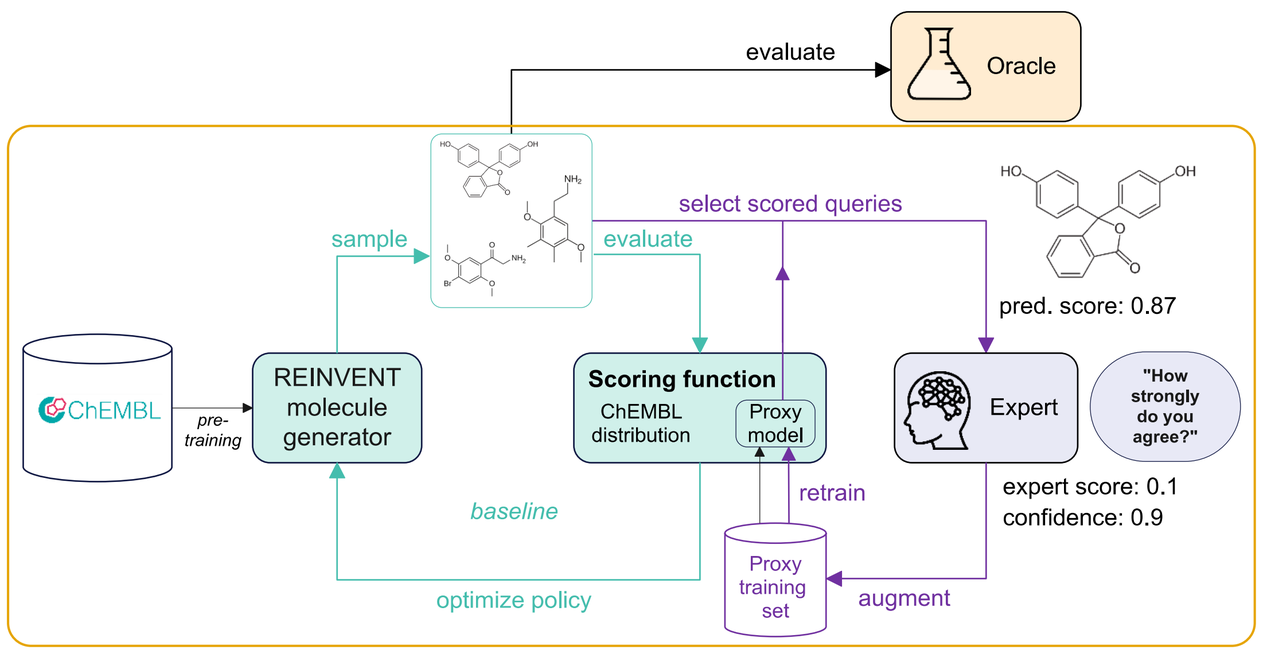

Oracle evaluation is only possible at the end of a generation cycle.

Before the generation: we train a predictive model using available *experimental* data.

During the generation: expert feedback (agreement score and confidence) is actively queried and used to update the predictor

3. How do we collect expert feedback?

1. **Simulated experts:**

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

controlling level of expertise

2. **Real experts:**

We developed Metis GUI based on rdEditor *(Bjerrum Esben, 2019)*.

- Allows to specify a selection strategy that will query a chemist.
- Automatically resumes REINVENT runs using the updated predictor.

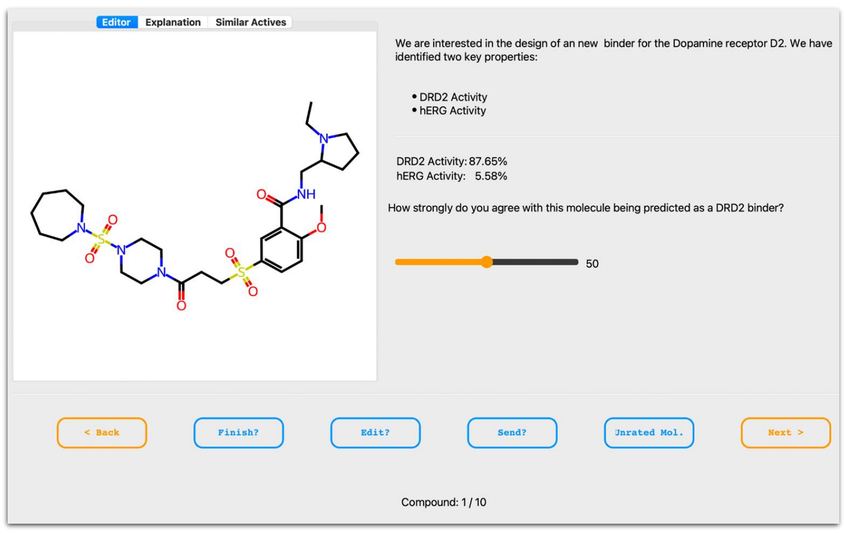

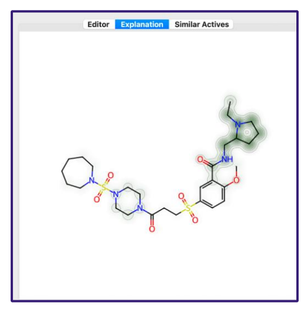

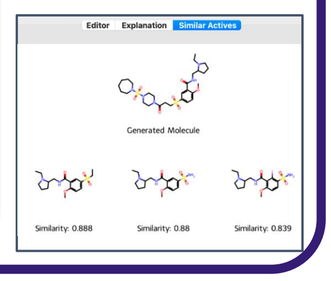

4. Experiments & Results

- **Use case: generate bioactive molecules against the Dopamine Receptor D2 (DRD2)**

Mono-objective, AL using simulated experts

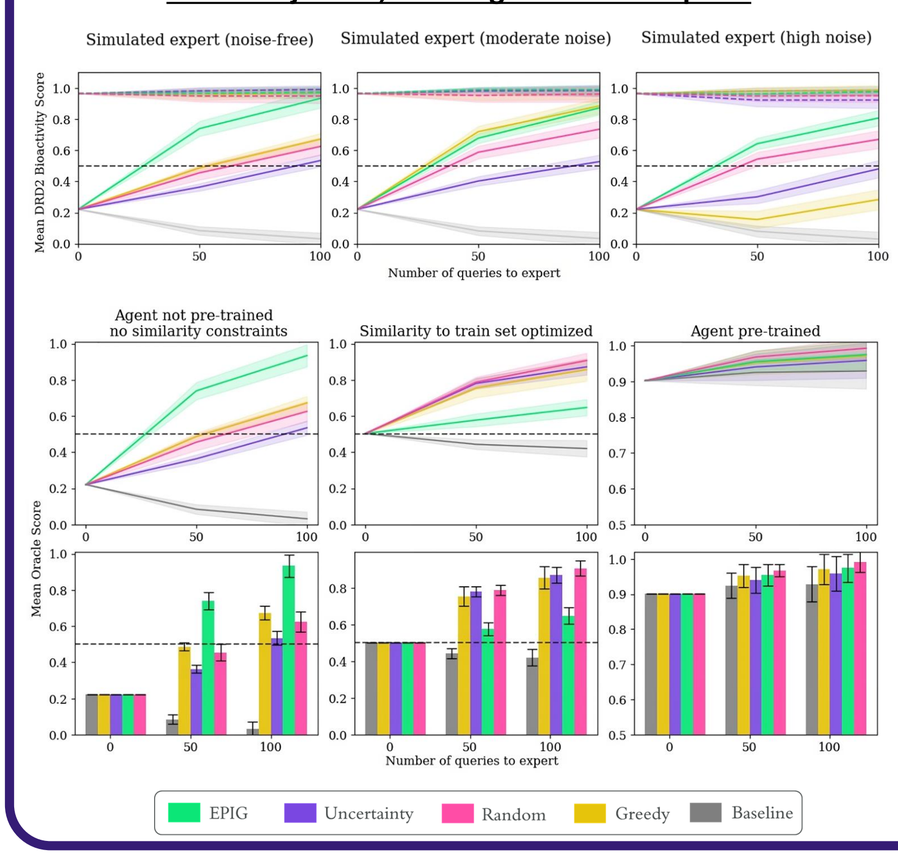

Multi-objective, AL using real experts

- Chemists 1 and 2 were asked "How much would you prioritize this molecule as a DRD2 active?"
- Chemist 3 was asked "How strongly do you agree with this molecule being predicted as DRD2 active?"

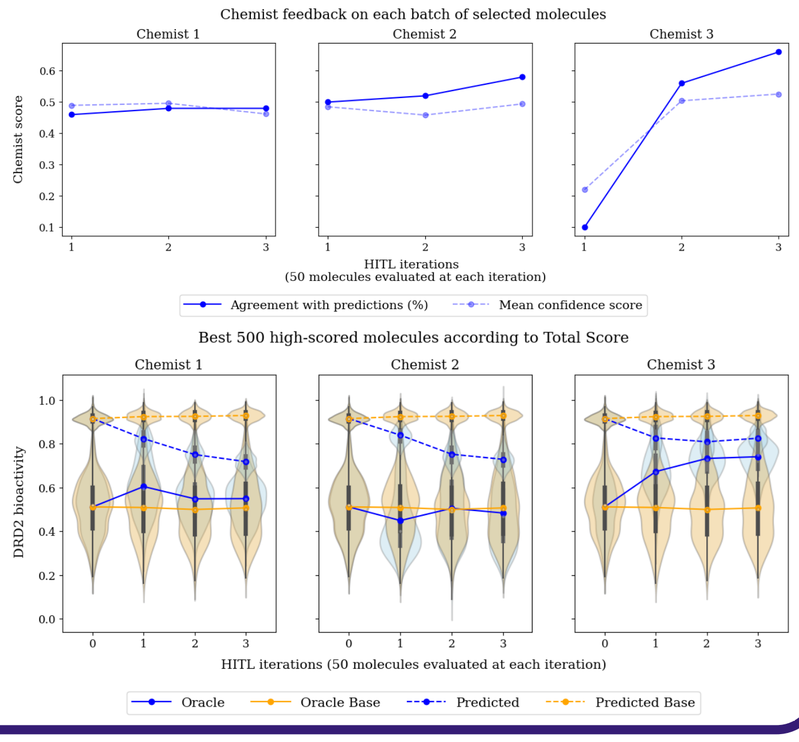

5. Conclusions

★ AL can be employed during molecule generation to prevent the predictor hacking issue

★ AL results in higher enrichment in true positives among predictor-top-scored molecules

★ Querying domain experts proves to be successful for aligning generated molecules with real-world practical applications *(better calibrated predictions, improved QED and SA scores)*

★ The question asked to experts must convey the goal of calibrating the predictor to avoid ambiguity

★ EPIG recommended as query selection strategy for its higher robustness to noisy inputs

This research is funded by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Innovative Training Network European Industrial Doctorate grant agreement No. 956832.

Preprint released soon!

yasmine.nahal@aalto.fi