A Molecular Bio-Assays Simulator to Unravel Predictors Hacking during Optimisation

Nicolas Drizard, Joseph-Andre Turk, Philippe Gendreau

- ❖ Generative models are being increasingly used in drug discovery campaigns, very often coupled with ADME or bio-assays QSAR predictive models to optimize a given set of properties which are easily hacked[1].
- ❖ We introduce a neural network based simulator for multi-target bio-assays to replicate real world prospective lead optimization scenario. We evaluate different approaches to mitigate the issue of *predictor hacking* using this oracle which lead to an increase in the proportion of molecules truly in the target product profile (TPP) among generated molecules.

Drug Discovery Workflow with AI

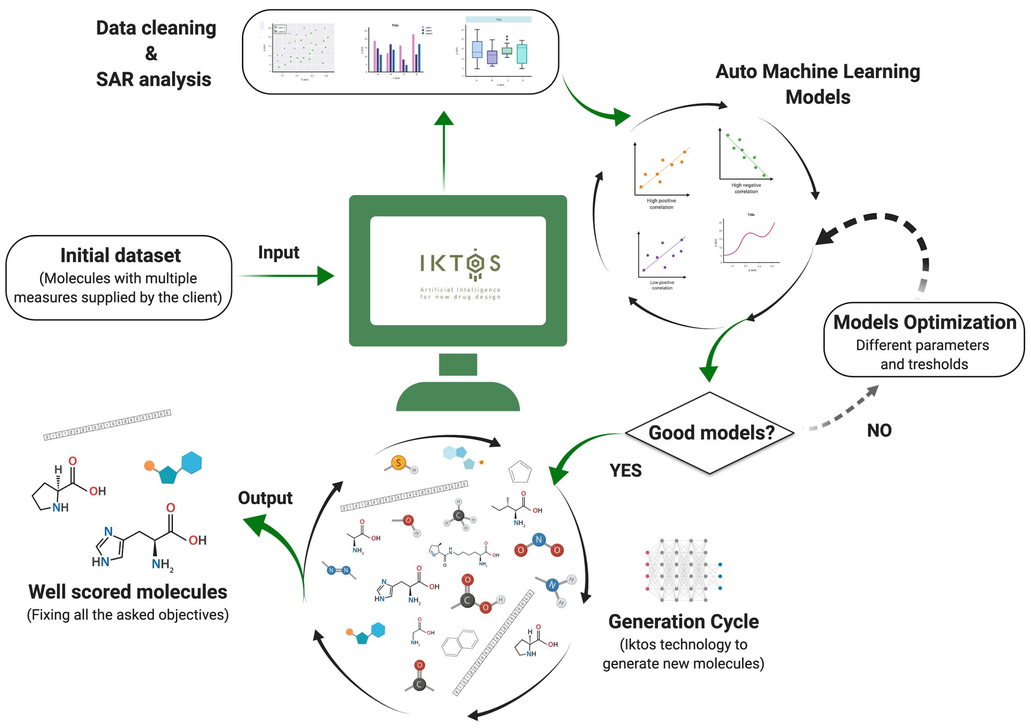

Presenting the Issue: Goodhart's Law

Since all models are false, the predictive models which are being optimized during the generation are *easily hacked*, as every flaw of the models which could give extra reward for the generator can be exploited. This is called "reward hacking" and is an evidence of the Goodhart's law.

> "When a measure becomes a target, it ceases to be a good measure"[2]

Bioassays Neural Simulator

The neural simulator returns continuous distributions used to replicate bio-assays behaviors. Its key features are:

- Inputs are MACCS fingerprints, which holds a differnet information than Morgan ECFP
- Initialization of the weights follow a power law $p(x) = kx^{\alpha}$

when α is high, close to a Gaussian distribution so easy to predict

when α is low, very non-linear so hard to predict

- Size constraint of the penultimate layer for moderate correlations between simulated targets

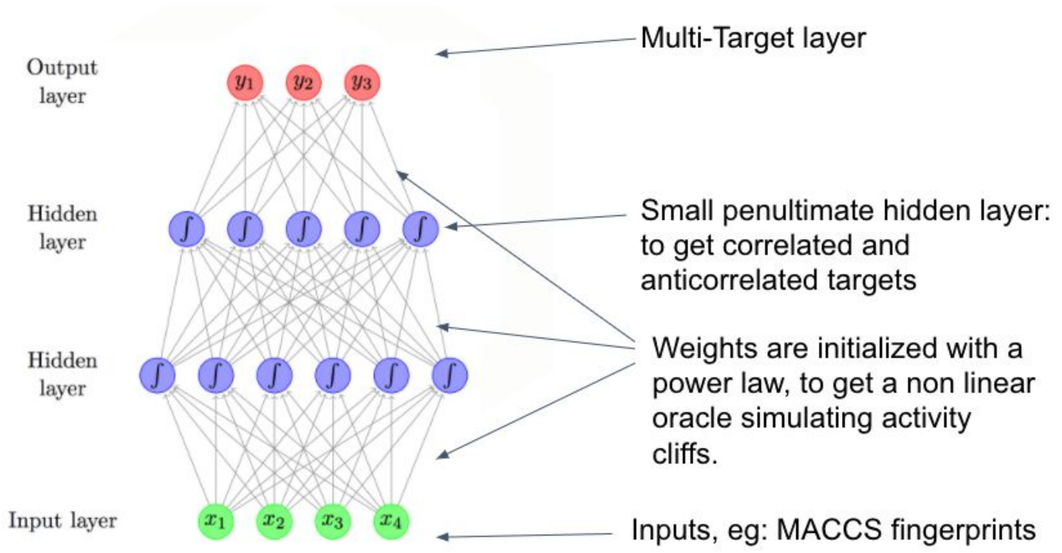

	Interpretable	Tunable difficulty/ activity cliffs	Multi-targets with correlation control	Deterministic	Cheap to compute
Molecular descriptor (e.g: QED)	✓	✗	✗	✓	✓
Docking	✓	✗	✗	✗	✗
Neural simulator	✗	✓	✓	✓	✓

Predictors Hacking In the Optimization

Definition

Predictor hacking stands for the over-overoptimization of the predicted score of a model which leads to a decrease or a stagnation of the real score during the optimisation while the predicted score increases.

Experiments

Using datasets representative of lead-optimization projects (~100s data points, limited number of chemical series), we generate molecules using a genetic algorithm from the Guacamole package[3] optimizing a random-forest based predictor built on the Morgan ECFPs - by batch of 100 at each step.

Observations

- Predicted score keeps increasing while the proportion of active decreases
- Controlling the *applicability domain* of the optimization, with both a similarity constraint and a start from an initial population within the chemical space (cs in graph) can mitigate the predictor hacking phenomenon. It increases as well the similarity within the generated batch

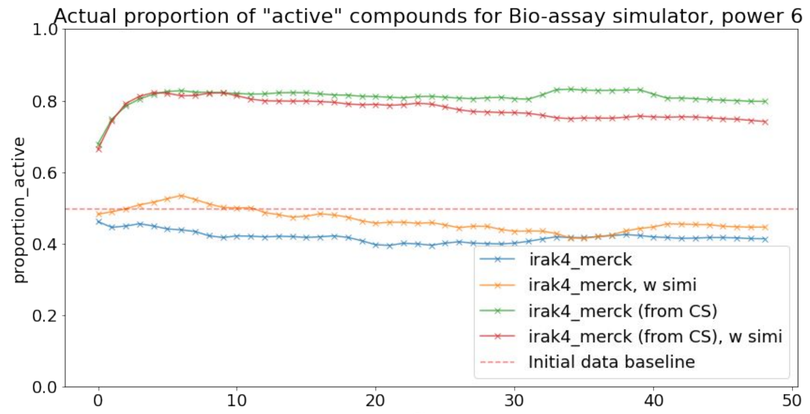

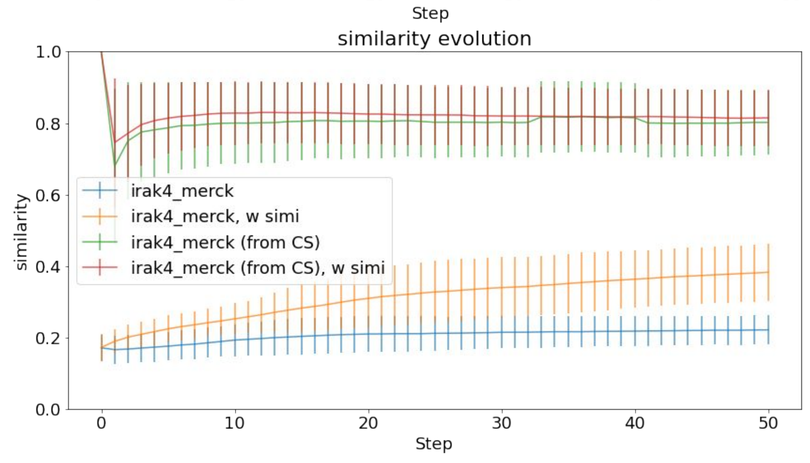

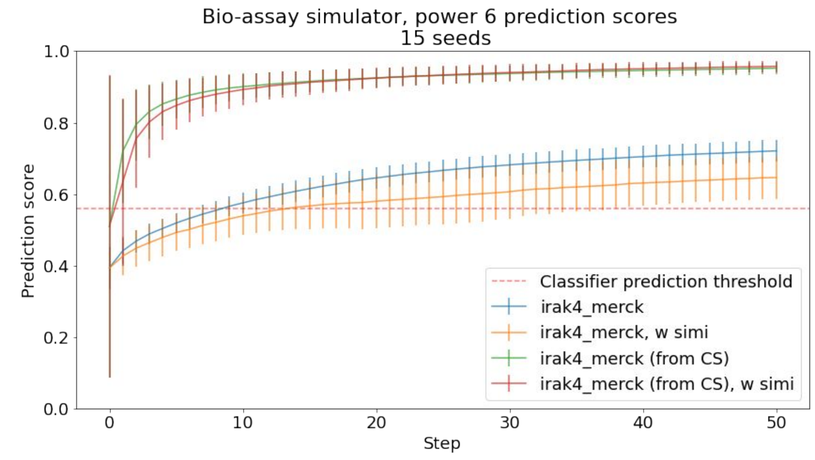

Next Steps

Evaluate how **rescoring** generation results with a different model (not used during optimisation) and **early stopping** the optimisation can mitigates the hacking;

References

[1]Renz, Philipp; Van Rompaey, Dries; Wegner, Jörg Kurt; Hochreiter, Sepp; Klambauer, Günter (2020): On Failure Modes of Molecule Generators and Optimizers. ChemRxiv. Preprint. https://doi.org/10.26434/chemrxiv.12213542.v1

[2] Marilyn Strathern

Improving ratings': audit in the British University system

https://doi.org/10.1002/(sici)1234-981x(199707)5:3<305::aid-euro184>3.0.co;2-4

[3]Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher.Guacamol: benchmarking models for de novo molecular design.Journal ofchemical information and modeling, 59(3):1096–1108, 2019

Contact

Nicolas Drizard: nicolas.drizard@iktos.com