Neural Networks Designing New Drugs

Mariya Popova, Olexandr Isayev, Alexandr Tropsha

THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL

Drug Discovery Timeline

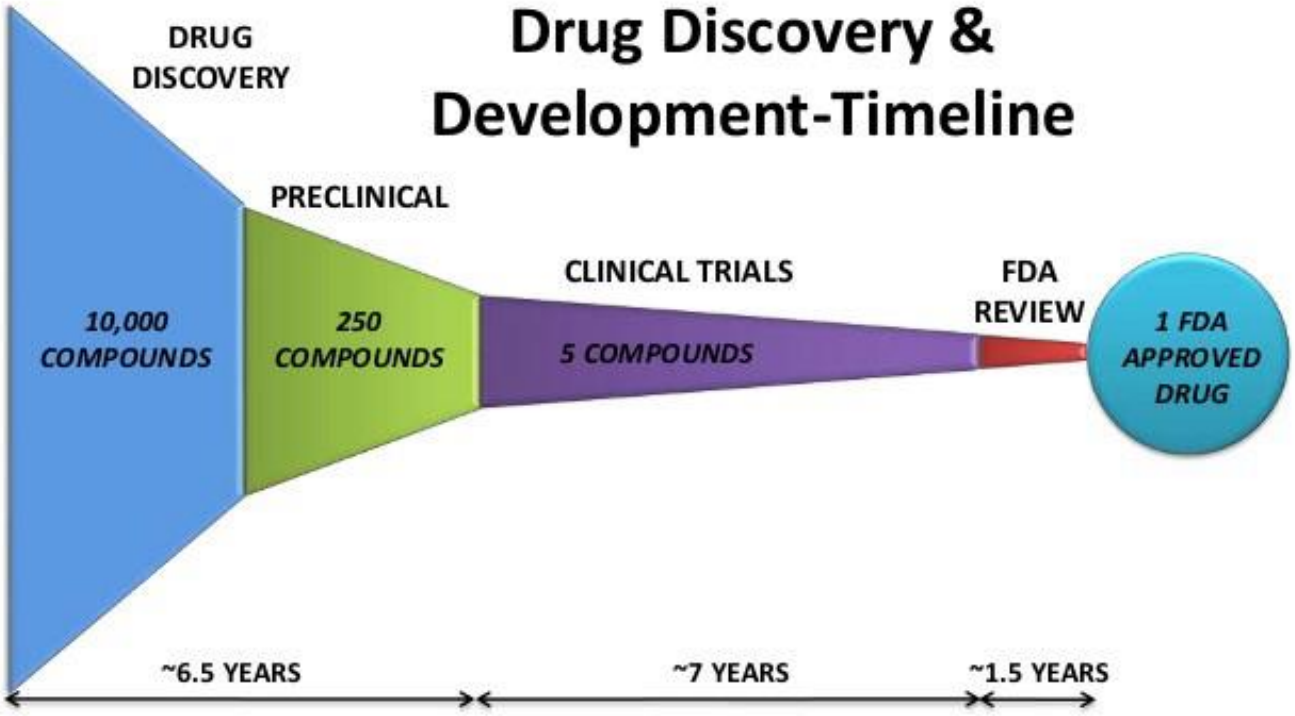

Conventional Virtual Screening Pipeline

CHEMICAL STRUCTURES → CHEMICAL DESCRIPTORS → PREDICTIVE MODELS ← PROPERTY/ ACTIVITY

Chemical database

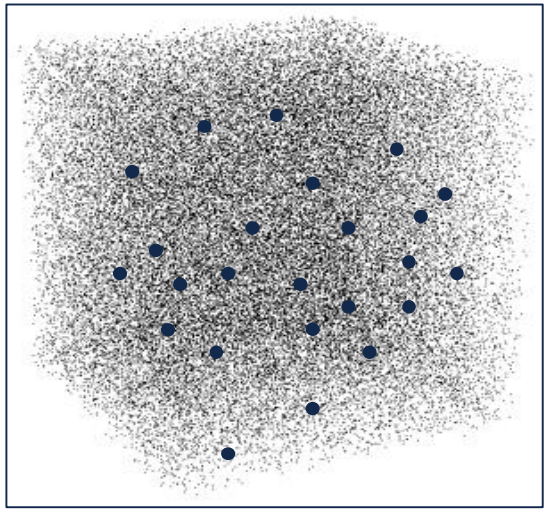

~10^6 – 10^9 molecules

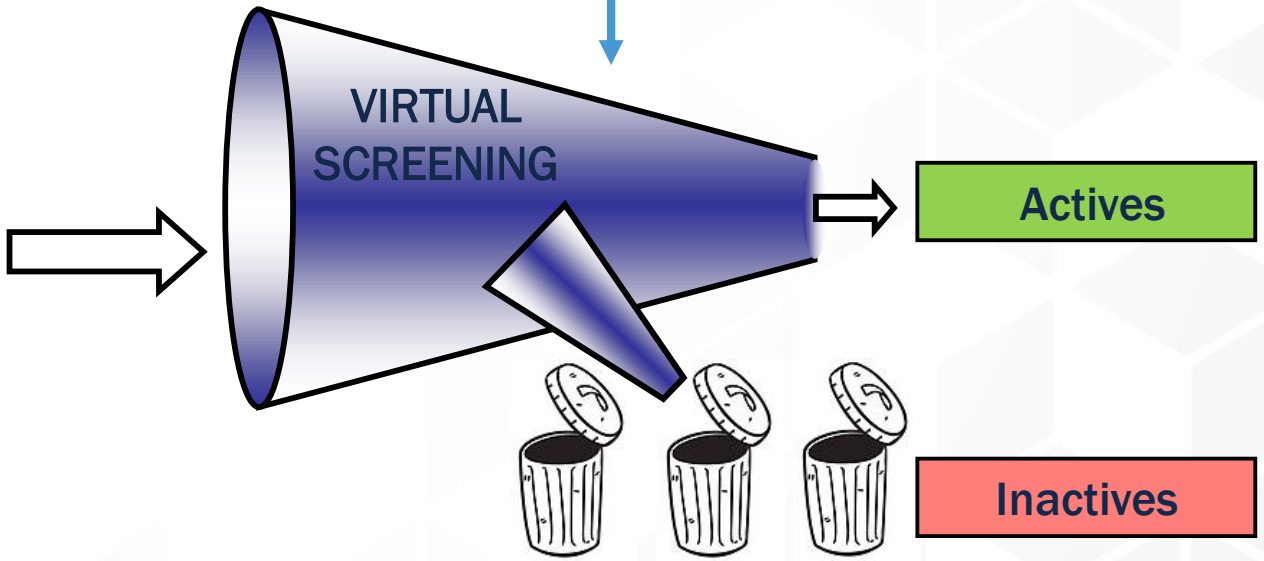

Why Do We Need Generative Models?

- Biggest database of molecules has ~10^9 compounds
- Estimates for the size of chemical space – up to 10^{60}
- Searching for new drug candidates in existing databases – observation bias

Generative Models Overview

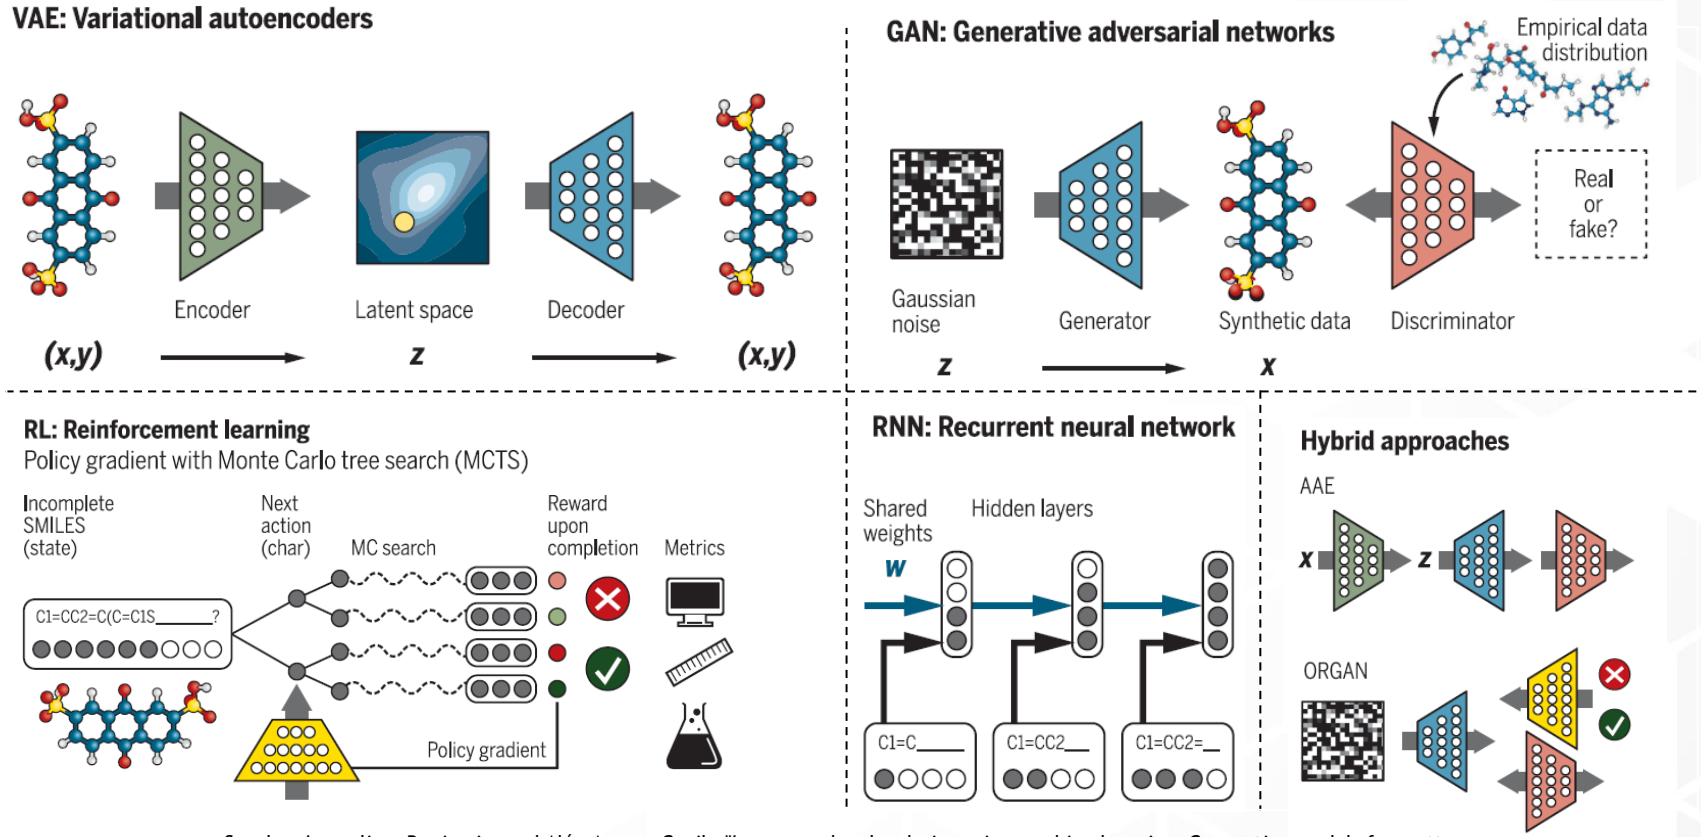

Sanchez-Lengeling, Benjamin, and Alán Aspuru-Guzik. "Inverse molecular design using machine learning: Generative models for matter engineering." *Science* 361.6400 (2018): 360-365.

Our Approach

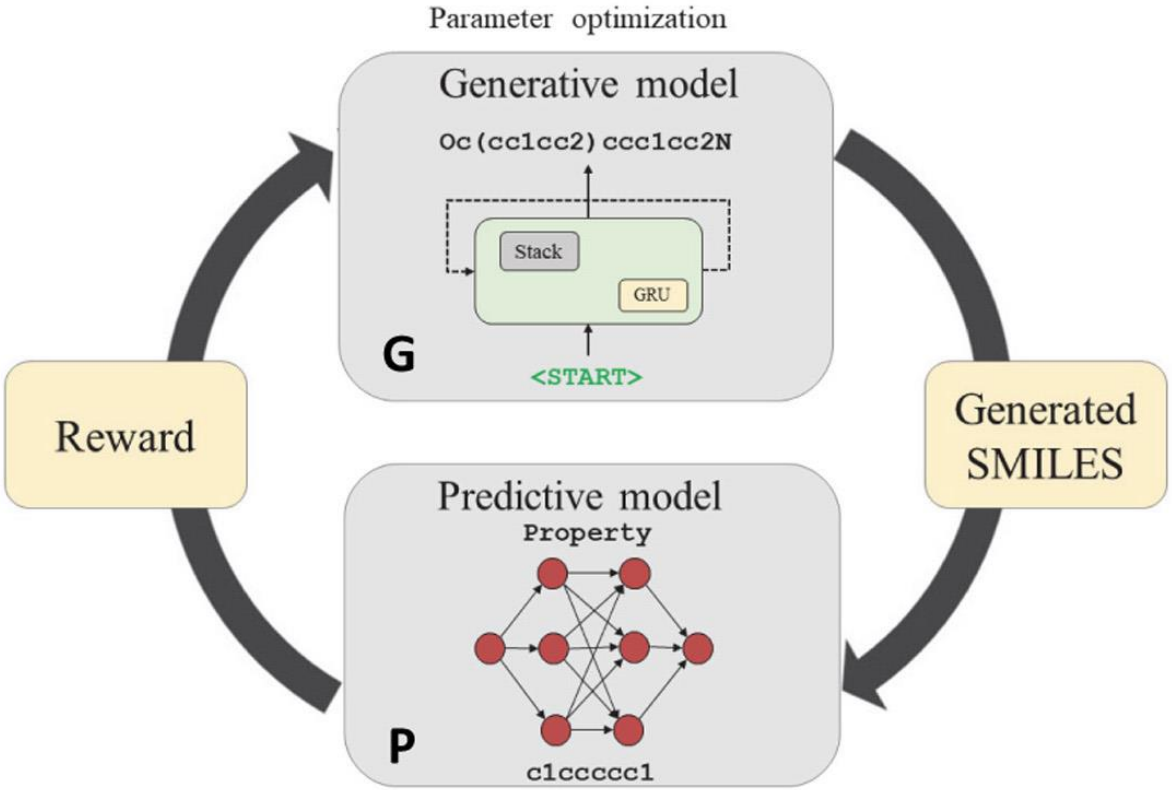

- Generative model for SMILES G
- Predictive model for the desired property P
- G and P combined with RL in one pipeline to bias the property of generated molecules.

Popova et. al. "Deep reinforcement learning for de novo drug design." *Science advances* 4.7 (2018): eaap7885.

SMILES-based Generative Model

- SMILES (**simplified molecular-input line-entry system**) is a sequence of characters then encodes the molecular graph
- One sequence = one molecule
- Has alphabet

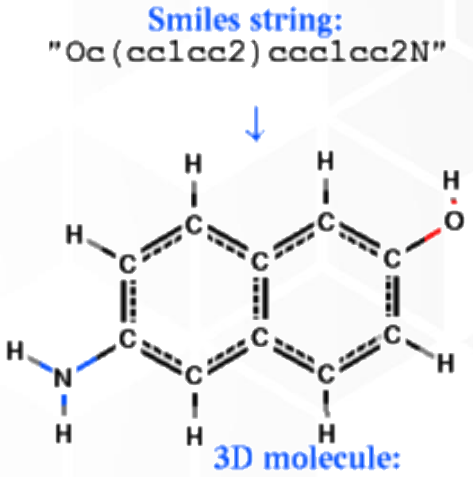

Use language model for producing novel SMILES strings

$$p(s_t|s_1 \dots s_{t-1};\ \theta) = f(s_1 \dots s_{t-1}|\theta)$$

Generative Model: training mode

- Trained on 1.5 million of drug-like compounds from ChEMBL in a supervised manner

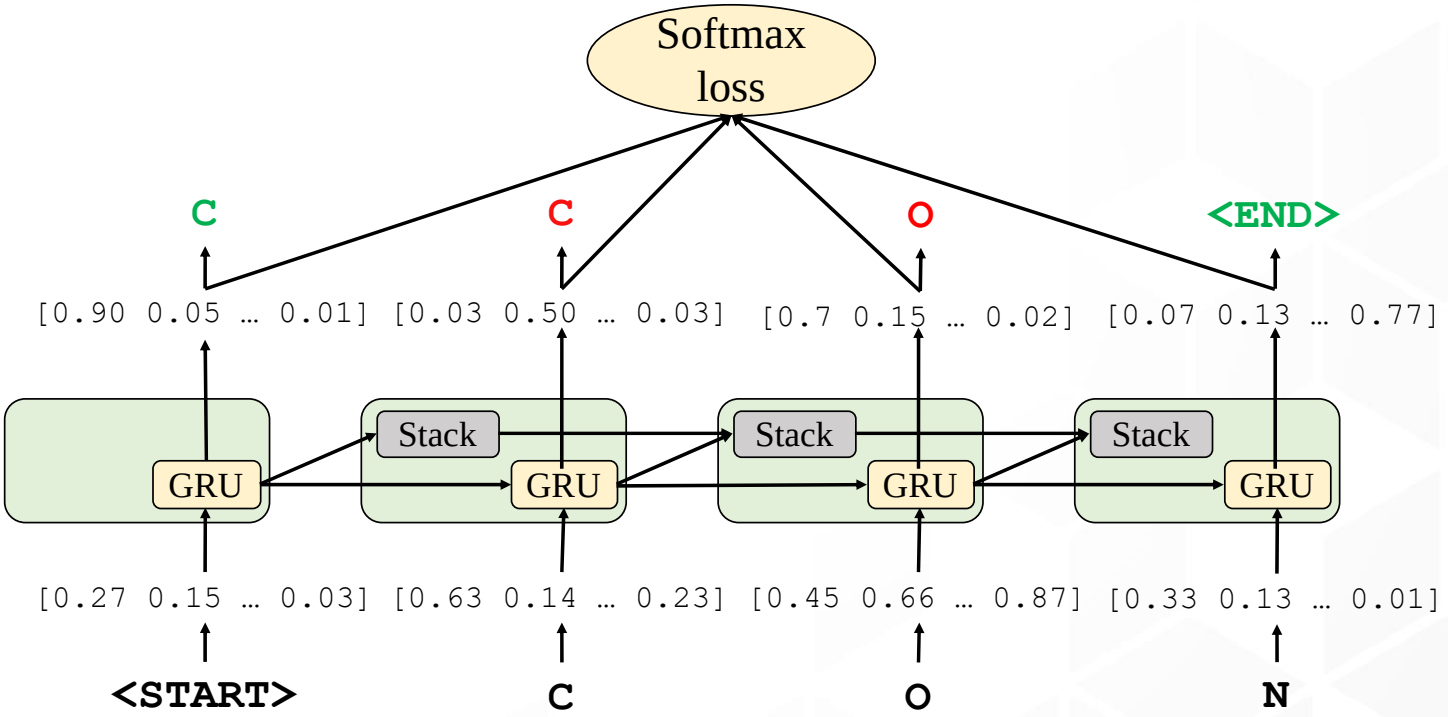

Generative Model: inference mode

Model takes its own predictions as next input character: $p(s_t|s_1 \dots s_{t-1};\ \theta) = f(s_1 \dots s_{t-1}|\theta)$

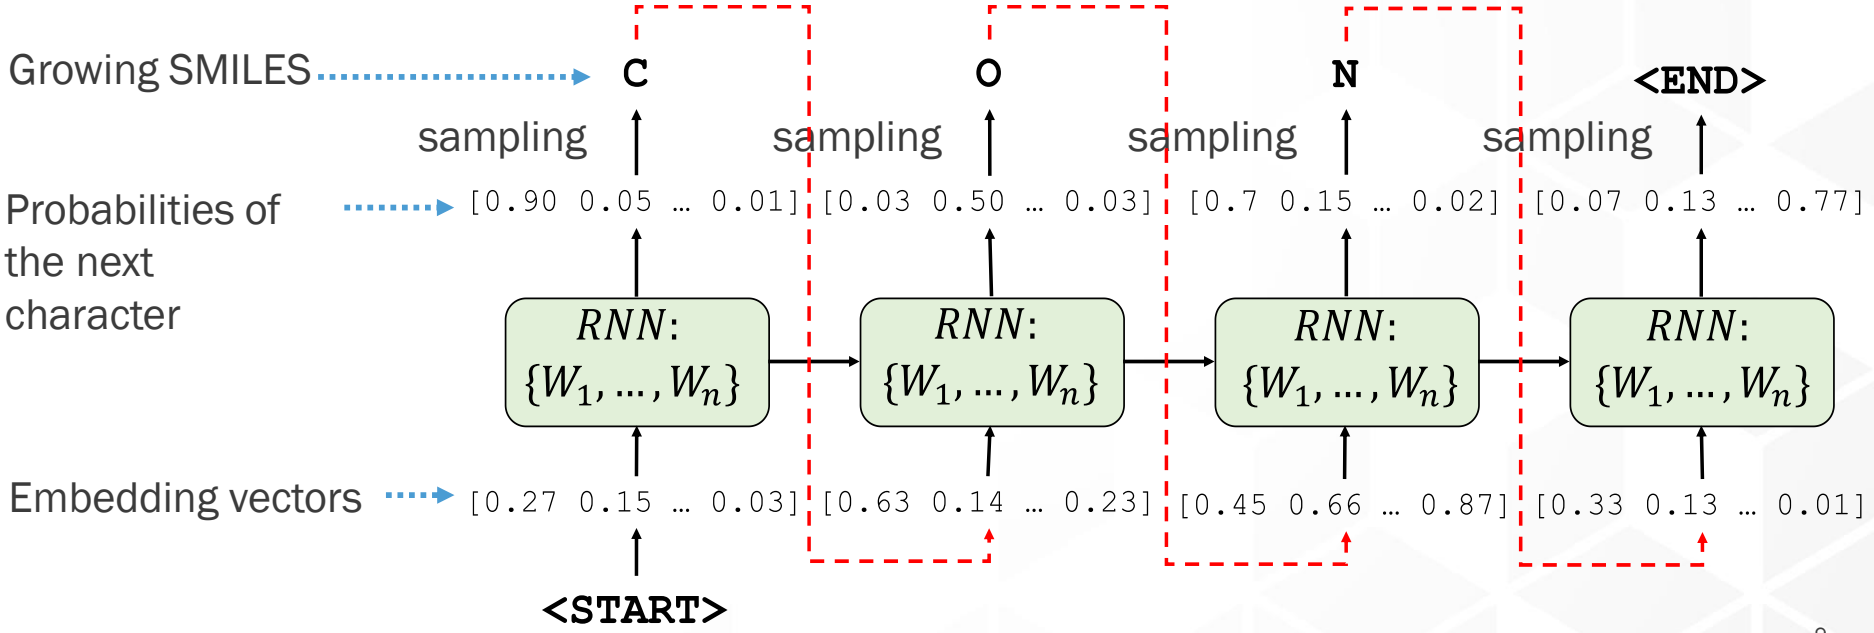

RL formulation for SMILES generation

- Action – generate symbol s
- Set of actions – SMILES alphabet A
- State – generated prefix $s_1 s_2 \dots s_{t-1}$
- Set of states – set of all possible strings in SMILES alphabet A with lengths from 0 to T -- $\mathbb{A} = \{A^t, t = 0 \dots T\}$
- Environment – set of states $\mathbb{A}$, set of actions A and transition probabilities $p(s_t = a | s_1 \dots s_{t-1};\ \theta),\ a\ \in A$
- Reward function – $R(S_t)$
- Objective – maximize the expected reward:

$$p(s_t | s_1 \dots s_{t-1};\ \theta) = f(s_1 \dots s_{t-1} | \theta)$$

$$\mathbb{E}[R(S_t)|\theta] = \sum_{S \in \mathbb{A}} p(S|\theta)\, R(S) \rightarrow max_\theta$$

RL Pipeline For Molecule Generation

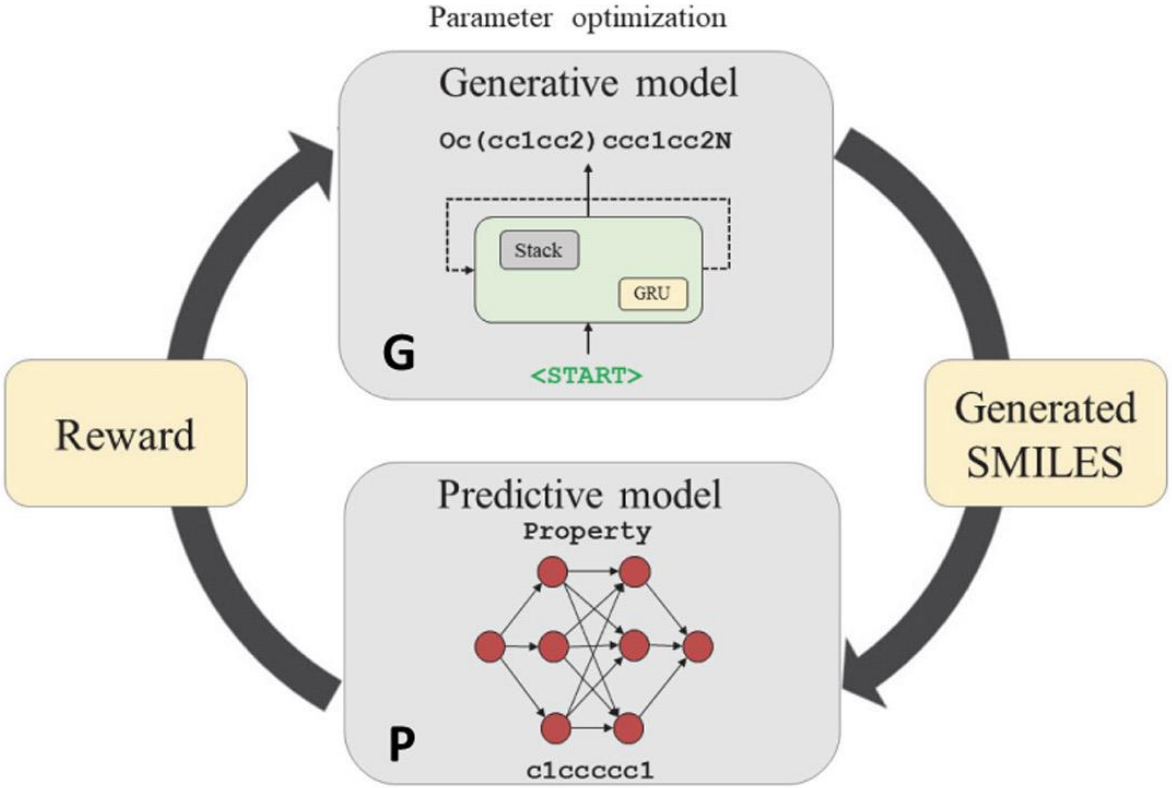

- Generative model is a policy network
- Predictive model is a simulator of the real-world
- Reward is assigned based on the property prediction and researcher's objective

Results: optimizing lipophilicity

- Lipophilicity is possibly the lost important physicochemical property of a potential drug
- It plays a role in solubility, absorption, membrane penetration, etc
- Log P is quantitative measure of lipophilicity, is the ratio of concentrations of a compound in a mixture of two immiscible phases at equilibrium
- Log P is a component of Lipinski's Rule of 5 a rule of thumb to predict drug-likeness
- According to Lipinski's rule must be in a range between 0 and 5 for drug-like molecules

Predictive Model for log P

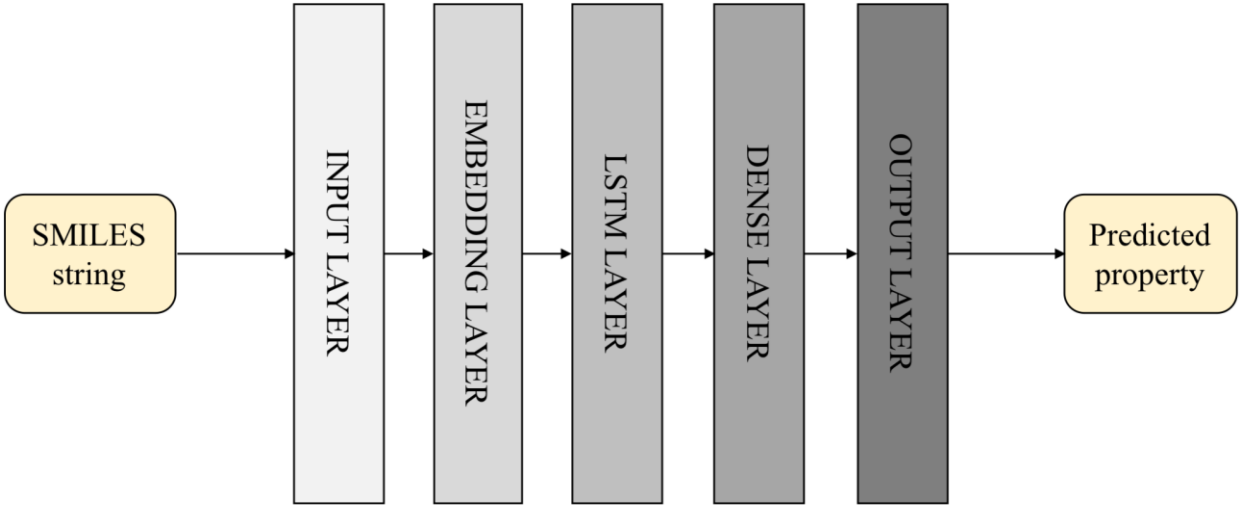

- SMILES-based RNN
- Dataset of 14k compounds with logP measurements
- 5 fold cross-validation
- RMSE = 0.57
- $R^2 = 0.90$

Results: optimizing lipophilicity

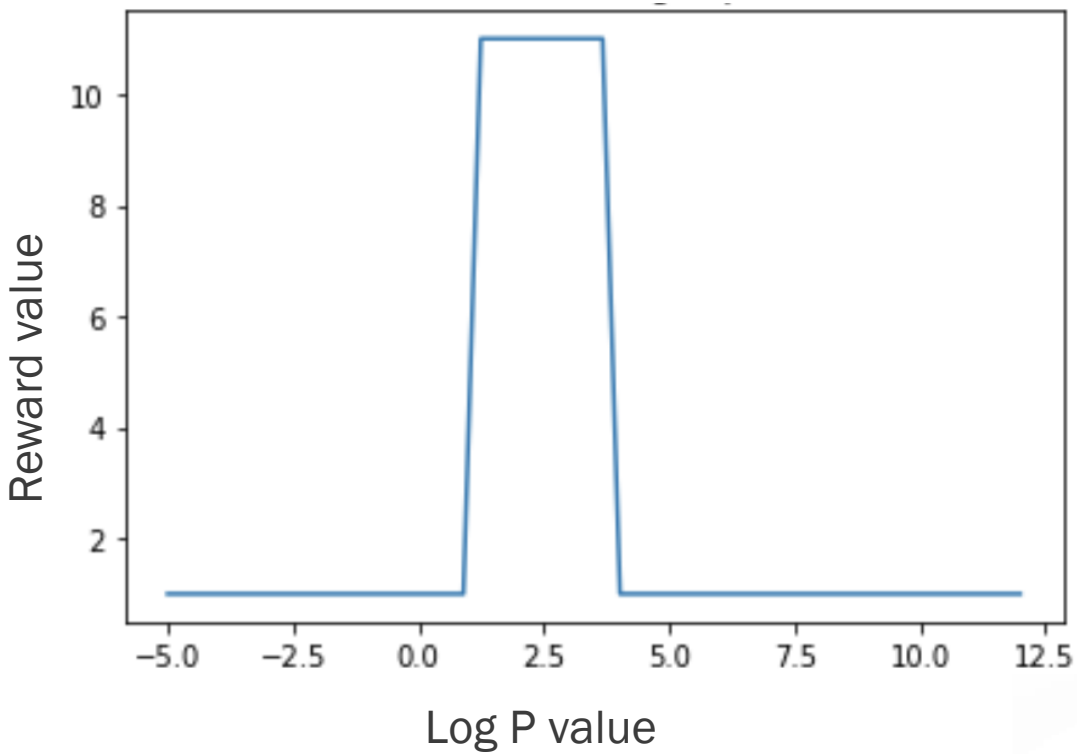

$$R(S) = \begin{cases} 11, if\ logP(s) \in [0.5; 4.5] \\ 0, otherwise \end{cases}$$

Results: optimizing lipophilicity

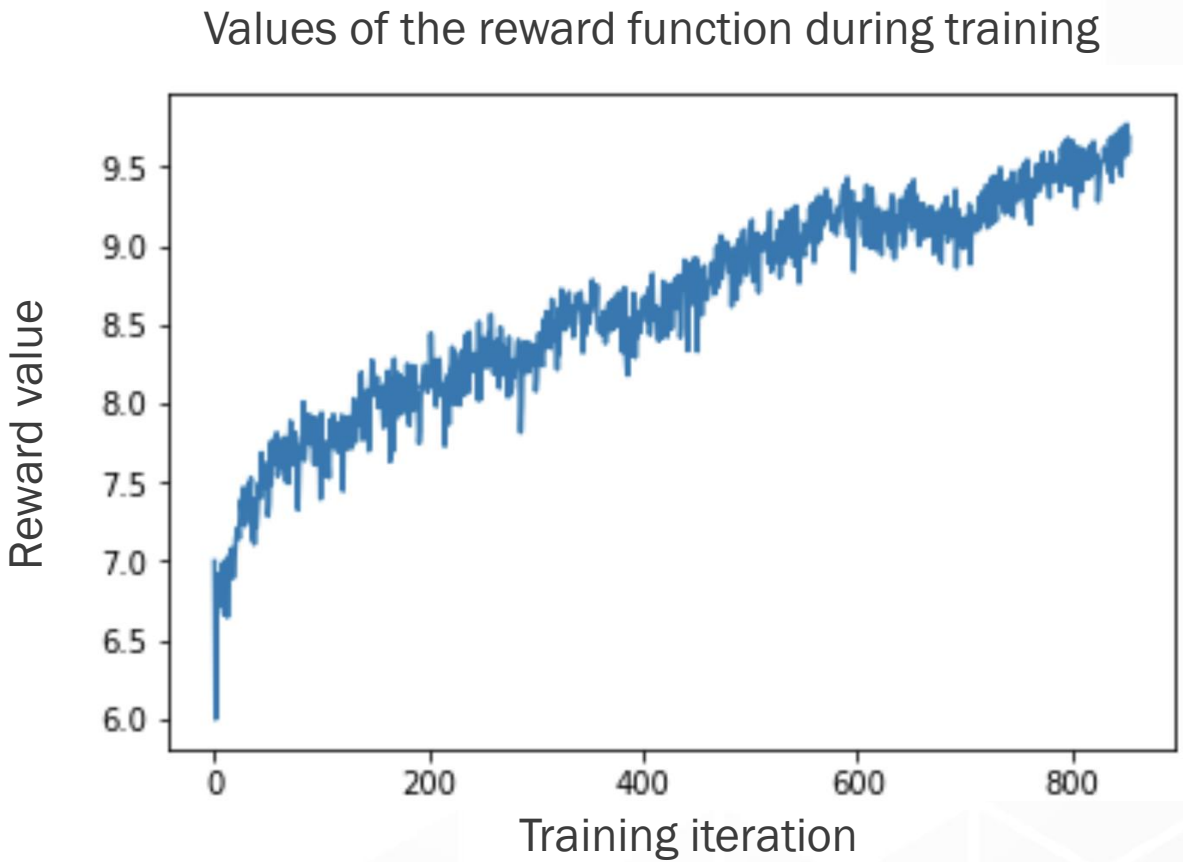

Results: optimizing lipophilicity

Distribution of unbiased and optimized log P values

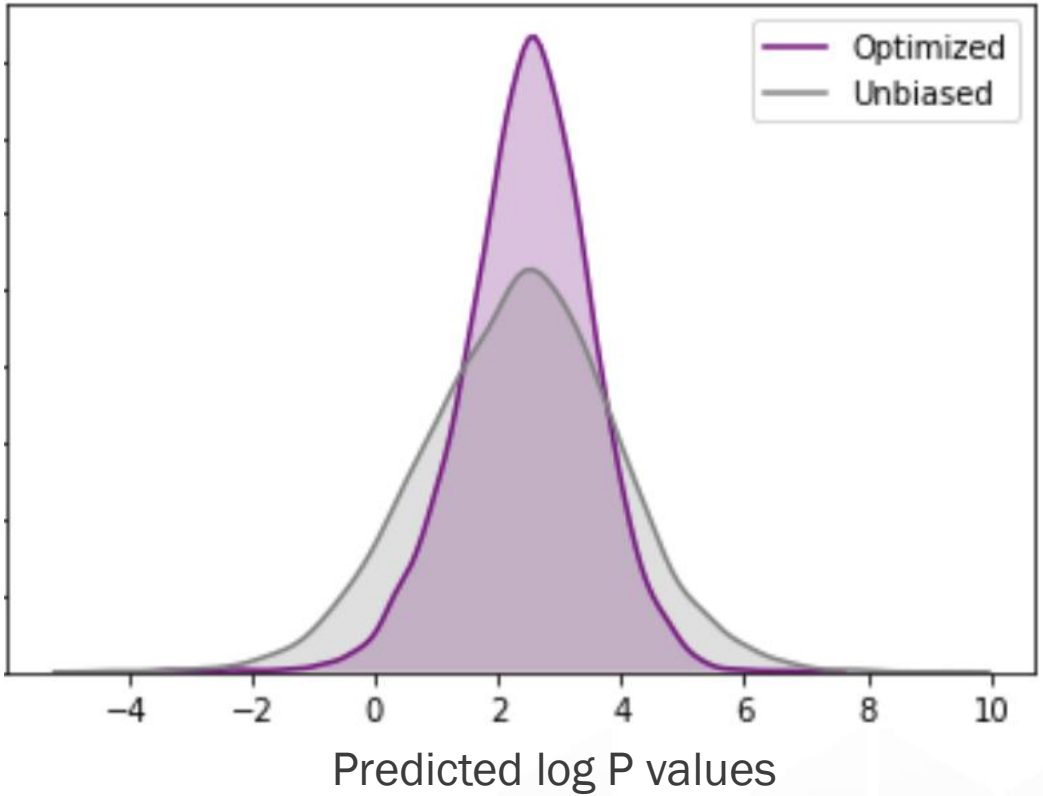

- Statistics are calculated from 10000 randomly generated SMILES
- 100% of optimized SMILES were predicted to have log P within drug-like region

Limitations

Worked well for a relatively simple physical property

What if a molecule with a high reward is a rear event?

It could take very long until the model receives a high or non-zero reward

Tricks

- Flexible reward
  - First give high reward for worse molecules, then gradually increase threshold
- Fine-tuning on a dataset of “good” molecules in a supervised manner
  - Fine-tune on generated molecules with high rewards
  - Fine-tune on experimental ground truth data
  - High exploitation, low exploration
- Using experience replay for policy gradient optimization
  - Remember generated molecules with high rewards and replay on them
  - Replay on experimental ground truth data

More results: EGFR

Epidermal growth factor receptor (EGFR)

- Associated with cancer and inflammatory disease
- Has ~10k experimental measurements for molecules

More results: EGFR

- Built a binary classification (active/inactive) predictive model for EGFR (F-1 score 0.9)
- Took pretrained on ChEMBL generative network
- Generated 10k random molecules and predicted probability of class “active”

More results: EGFR

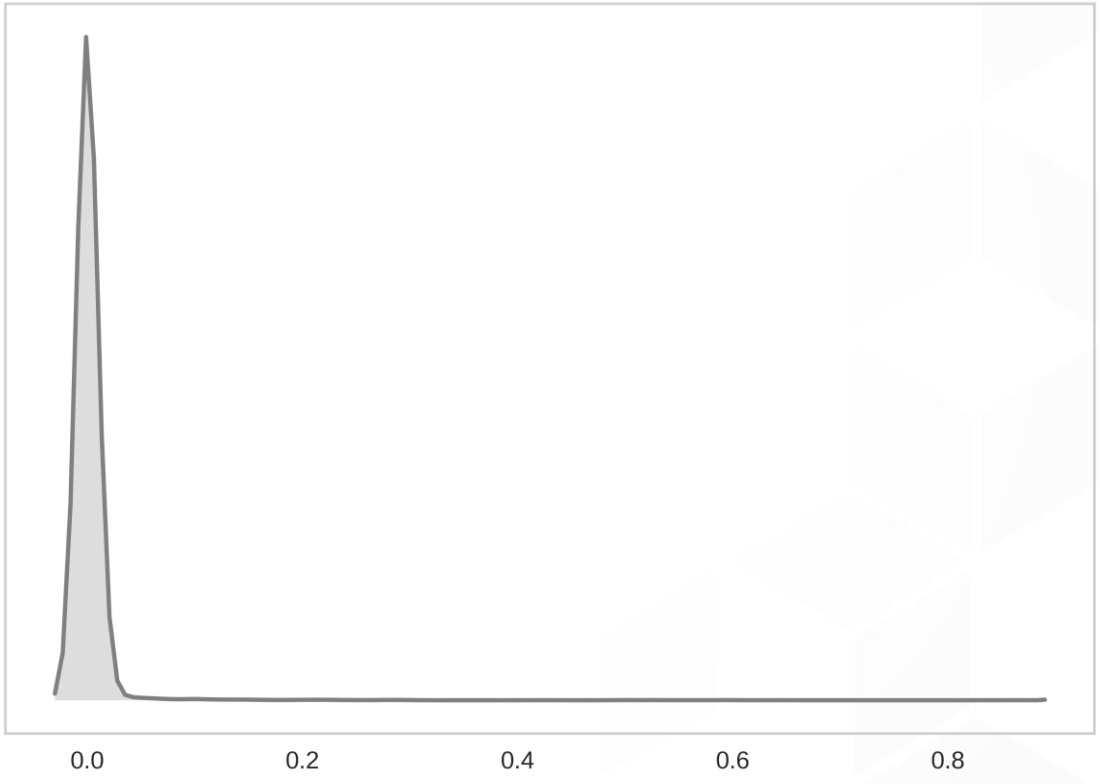

More results: EGFR

- Flexible reward:

$$R(S) = \begin{cases} 10, if\ P(S) > threshold \\ 0, otherwise \end{cases}$$

- Initial threshold = 0.05
- After every update we generate 10k compound
- If 15% of them predicted to have property > threshold, we increase threshold by 0.05
- Fine-tuning on generated molecules with high rewards
- Experience replay on experimental measurements and on generated molecules with high rewards

More results: EGFR

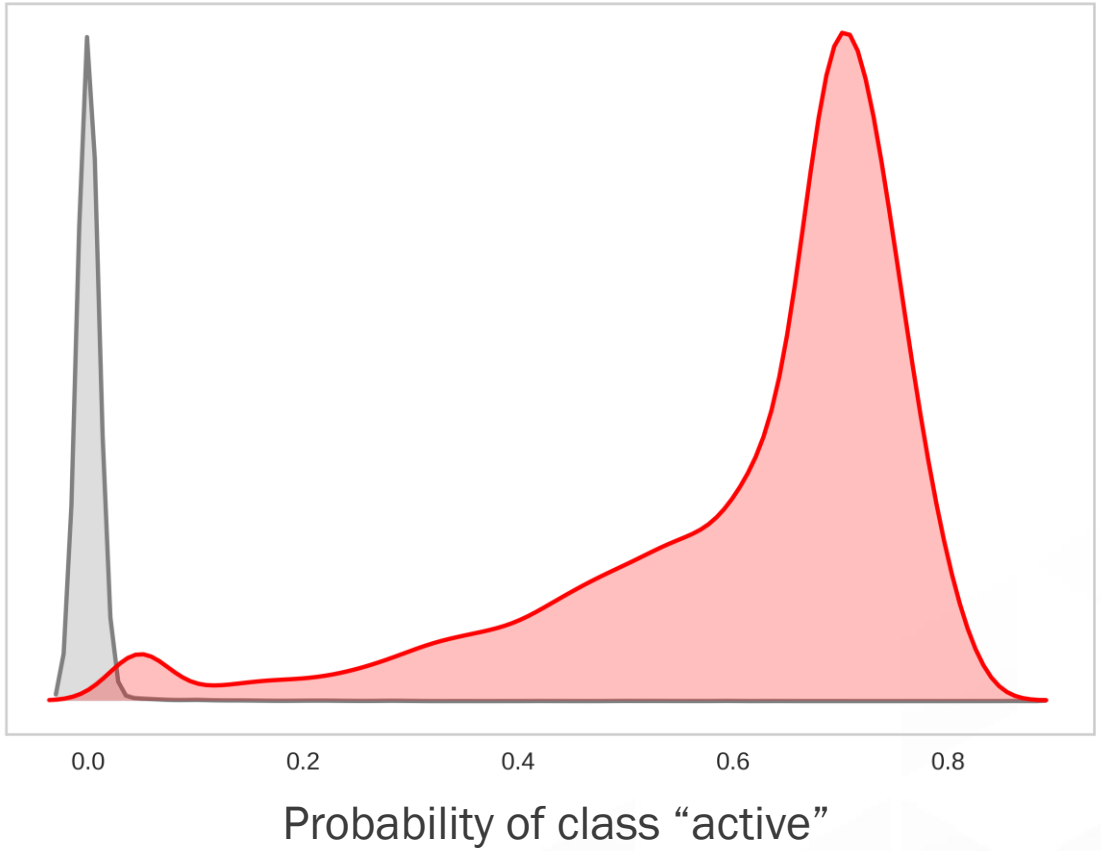

More results: EGFR

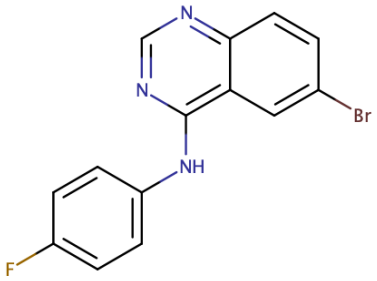

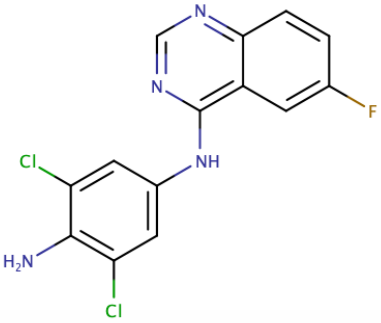

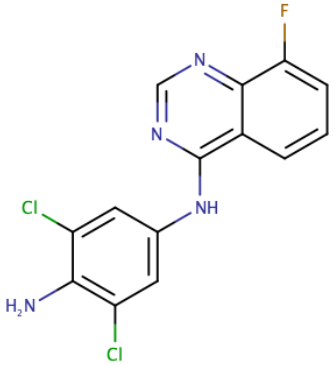

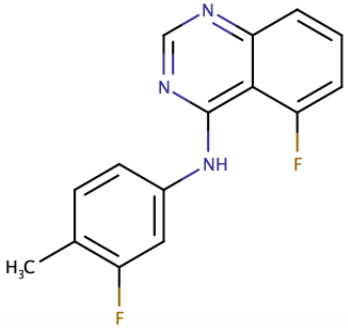

Experimental validation:

- Selected several commercially available and validated our results experimentally
- Found 4 active compounds

Future work

- Develop graph-based generative models:
  - SMILES-based models generate some amount of invalid molecules
- Develop lead optimization methods:
  - Start from a given scaffold/structure
  - Impossible to do with SMILES
- Develop models for predicting route for synthesis:
  - To be able to perform custom synthesis

Code Links

RL for de novo drug design

https://github.com/isayev/ReLeaSE

Acknowledgements

University of North Carolina at Chapel Hill:

Olexandr Isayev

Alexandr Tropsha