Georgia Tech

ICLR 2021

Molecule Optimization by Explainable Evolution

Binghong Chen[1*], Tianzhe Wang[1,5*], Chengtao Li[2], Hanjun Dai[3], Le Song[4]

[1]Georgia Institute of Technology, [2]Galixir, [3]Google Research, Brain Team,
[4]Mohamed bin Zayed University of AI, [5]Shanghai Qi Zhi Institute

Molecule Optimization

- Design new molecules with desired properties:
  - Property scoring function f (potent, non-toxic, easy to synthesize, …)
  - Challenges: searching over the vast space of $> 10^{60}$ molecules.

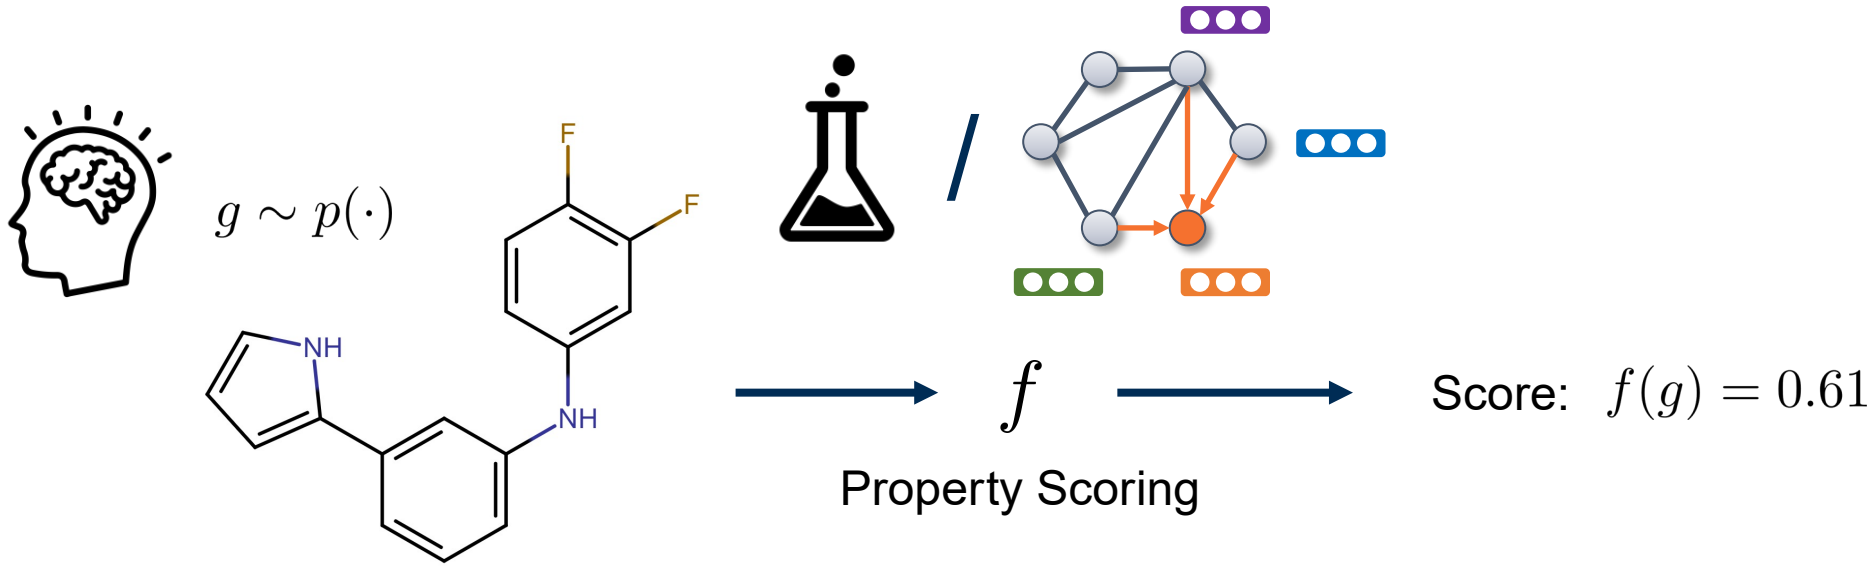

- Task: Learn a molecule generative model $p(\cdot)$ to maximize $\max_{p(\cdot)} \mathbb{E}_{g \sim p(\cdot)}[f(g)]$

RL for Generative Design

- Generation policy p: decide a new atom (and bonds) to add to the current partial molecule.

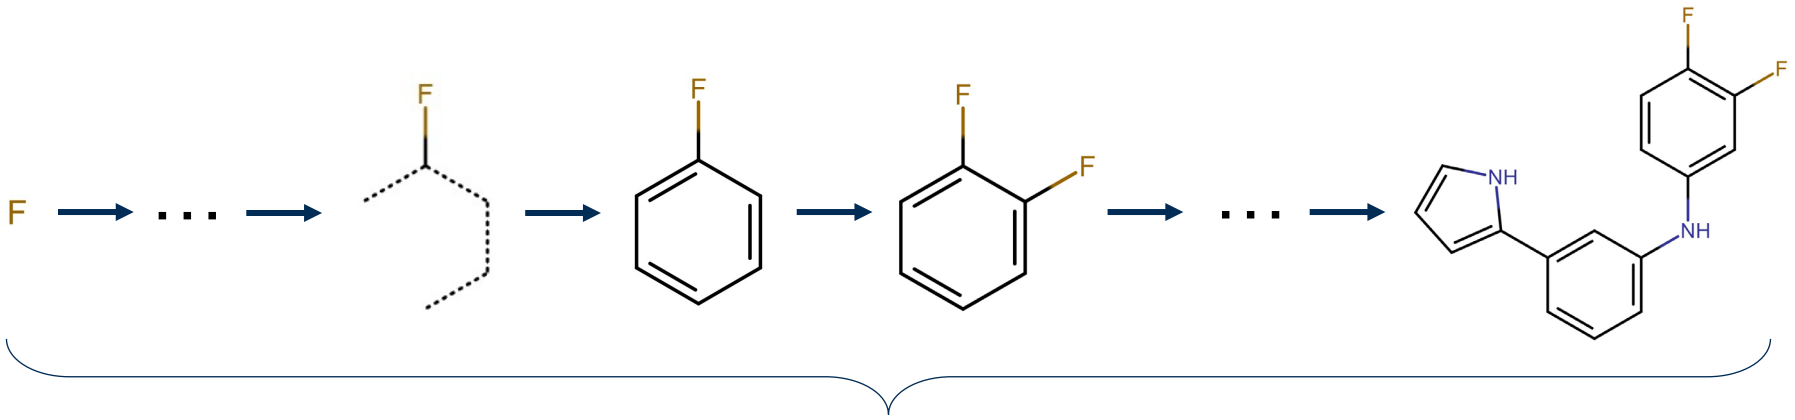

Autoregressive generative process

- Use RL to optimize p:
  - Reward $r = f(g)$ only obtained at the end.
  - Sparse reward, long horizon ⟶ hard to optimize.

Conditioning on Substructures

- Rationales – substructures that most contributes to the desired molecular properties.
- Conditioning generation policy p on rationales.

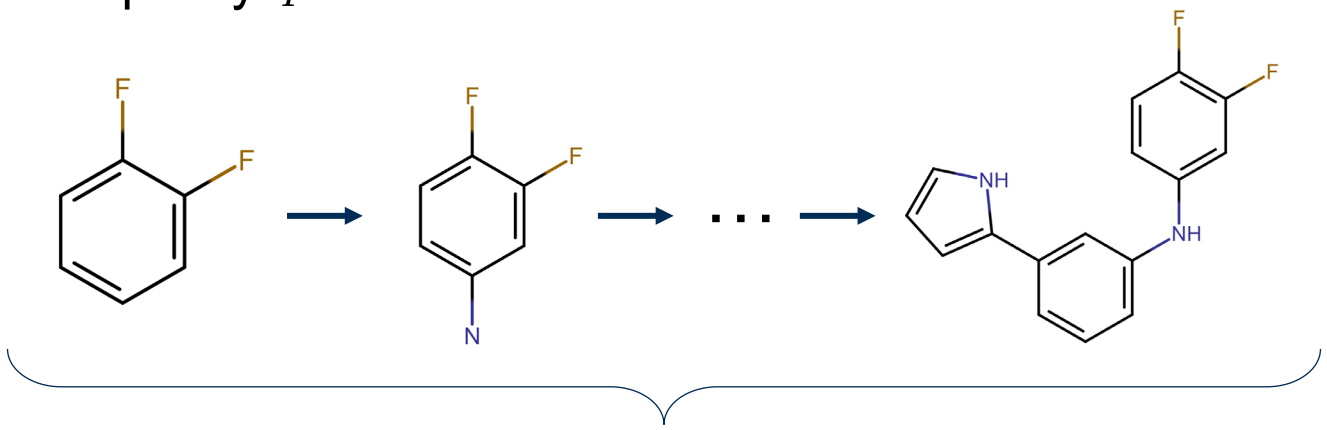

Conditioning autoregressive generative process

- Use RL to optimize p:
  - Shorter horizon ⟶ easier to optimize.
  - Obtaining rationales is hard
    - Designed manually: require human effort.
    - MCTS (Jin et al., 2020): unable to optimize rationales jointly with p.

Our Approach: MolEvol

- Hierarchical Generative Model

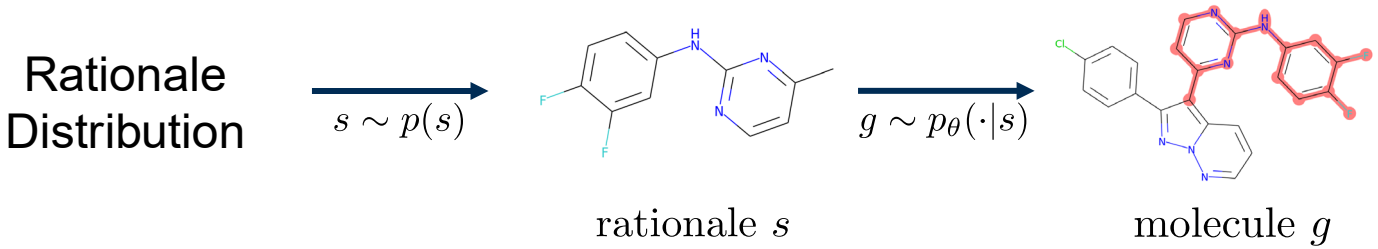

- Alternating Optimization (EM-style)

$$J(\theta, p(s)) = \mathbb{E}_{g \sim p_\theta(\cdot)}[f(g)] + \lambda \cdot \mathbb{H}[p(s)]$$

  - E-step
    - Fix $p_\theta(g|s)$, update $p(s)$.
  - M-step
    - Fix $p(s)$, update $p_\theta(g|s)$.

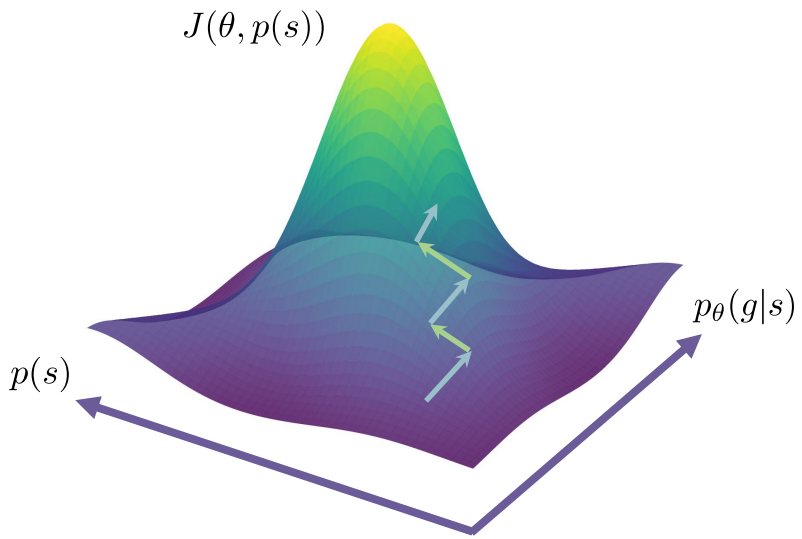

MolEvol: Algorithm Overview

- Init
  - A set of seed molecules are given.
  - Parameter θ^0.
- E-step
  - Produce a set of rationales with explainable graph model.
  - Optimize $p(s)$ (closed form).
- M-step
  - Produce a set of seed molecules.
  - Optimize $p_\theta(g|s)$ (RL).

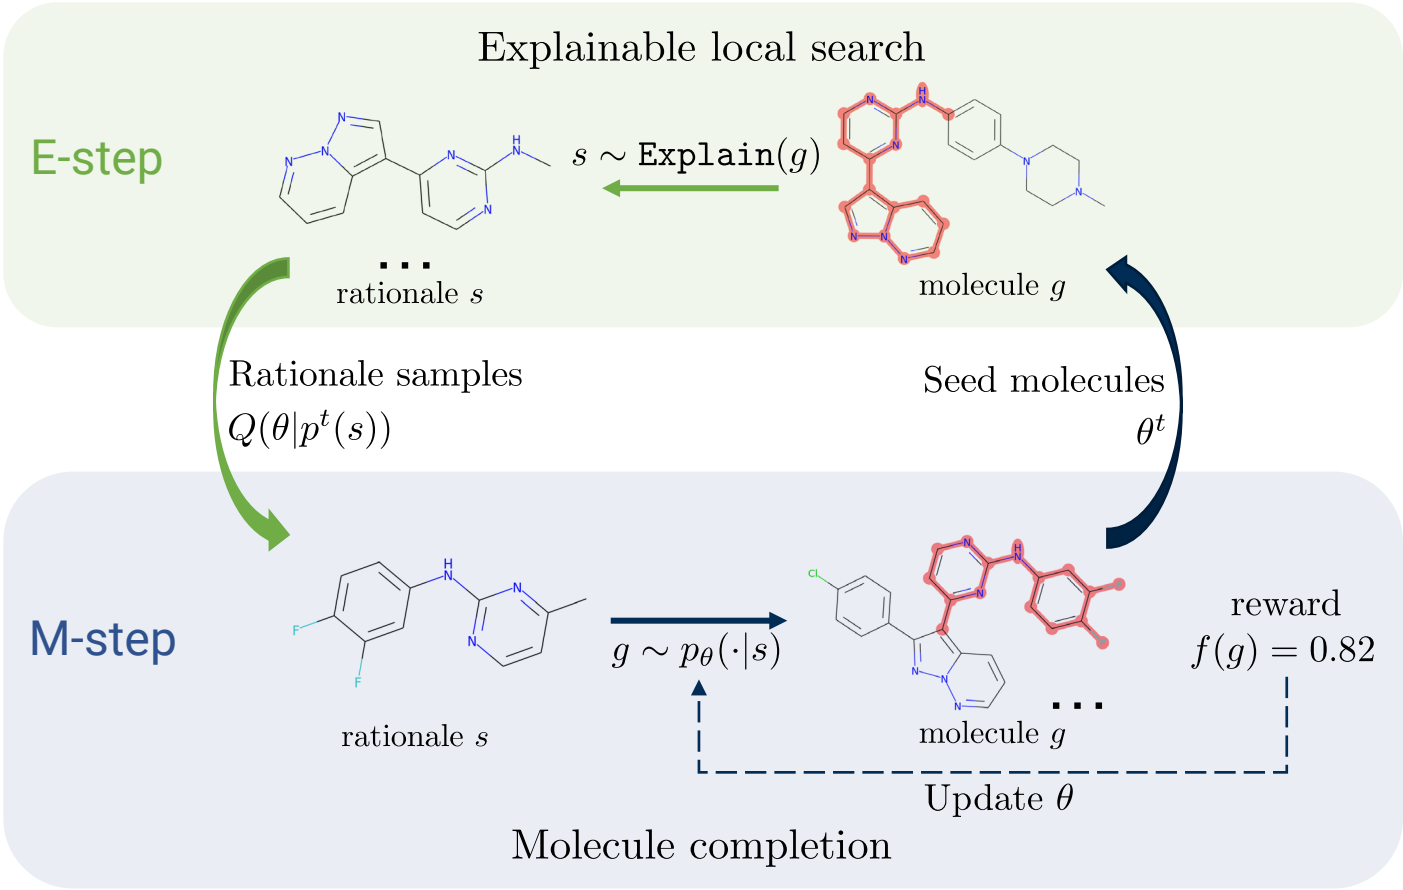

MolEvol: E-step

- We use weighted particles to represent $p^t(s)$.
- Particles are obtained from an explainable graph model.
  - Extract key subgraphs (rationales) s from seed molecules $g \in \mathcal{G}^{t-1}$.
  - s explains $f(g)$.
- Weights can be computed using θ^{t-1}.

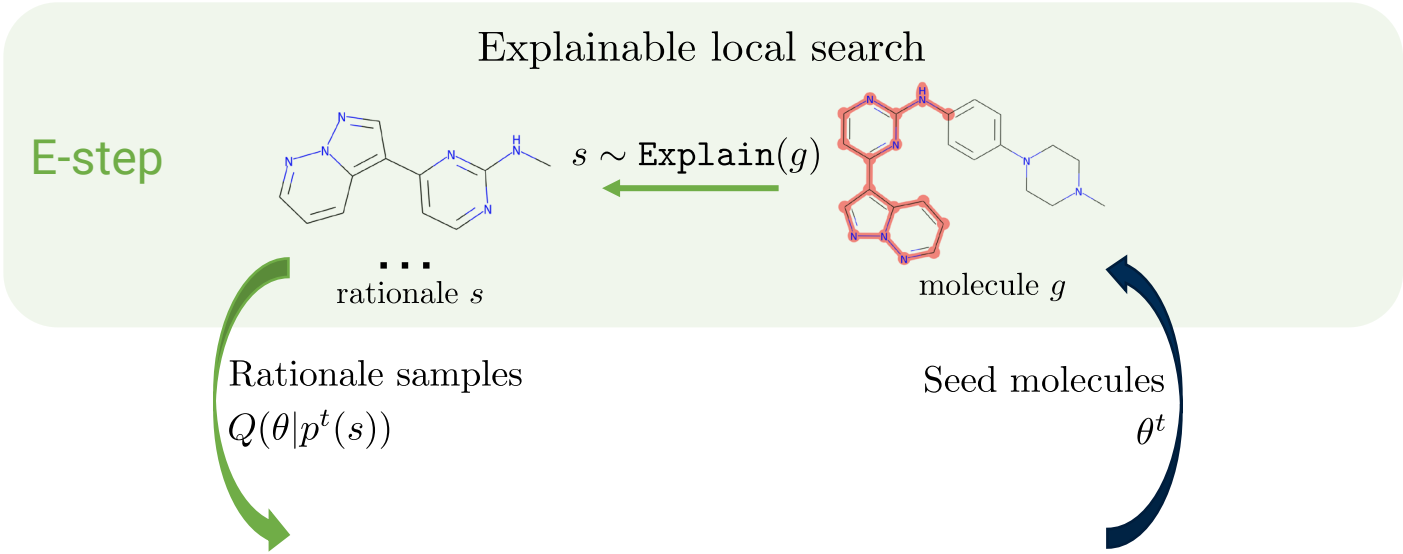

MolEvol: Explainable Graph Model

- To explain $\mathbb{P}(Y=1|g) \triangleq f(g)$, we maximize the mutual information between Y and rationale s.

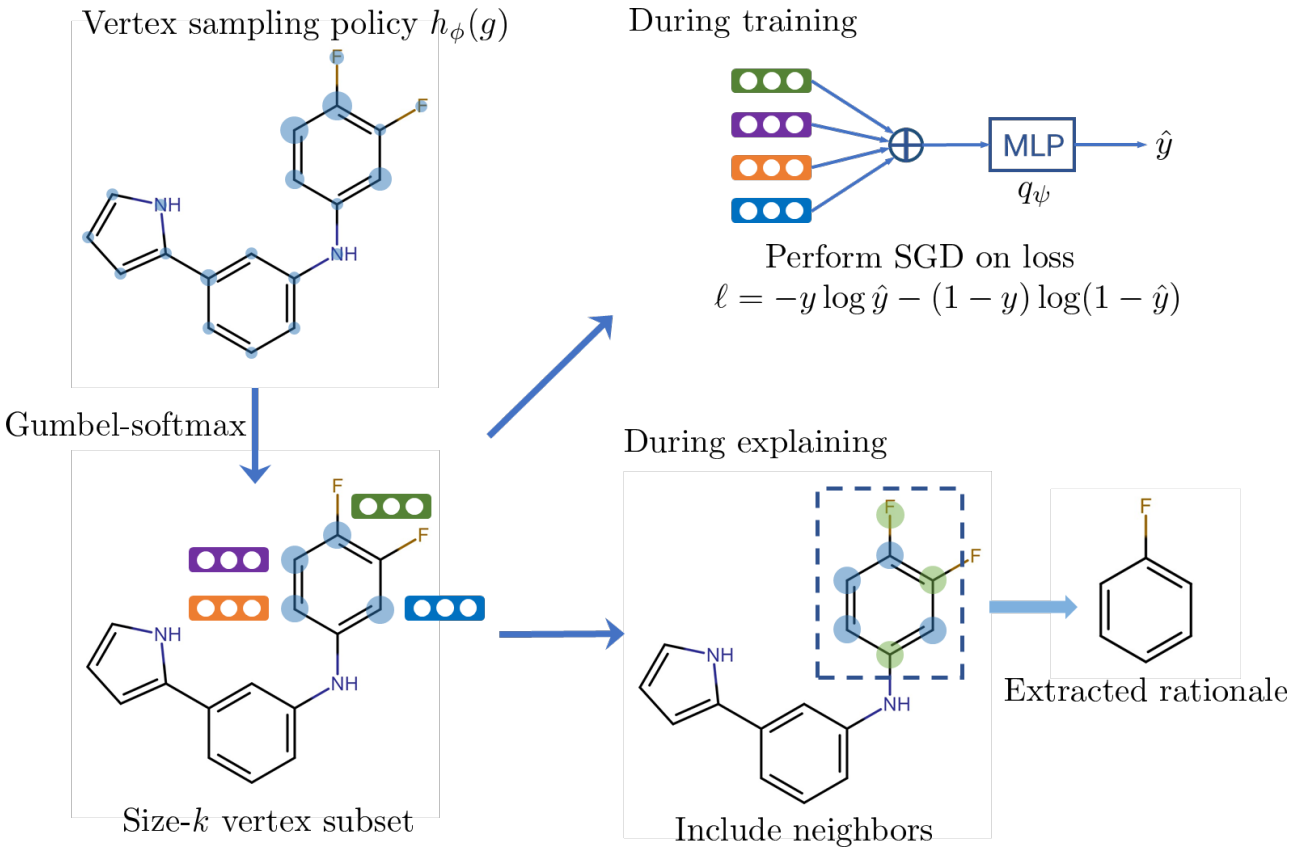

MolEvol: M-step

- We update θ^t from θ^{t-1} using RL,
  - Init state $s \sim p^t(s)$,
  - Reward $r = f(g)$.

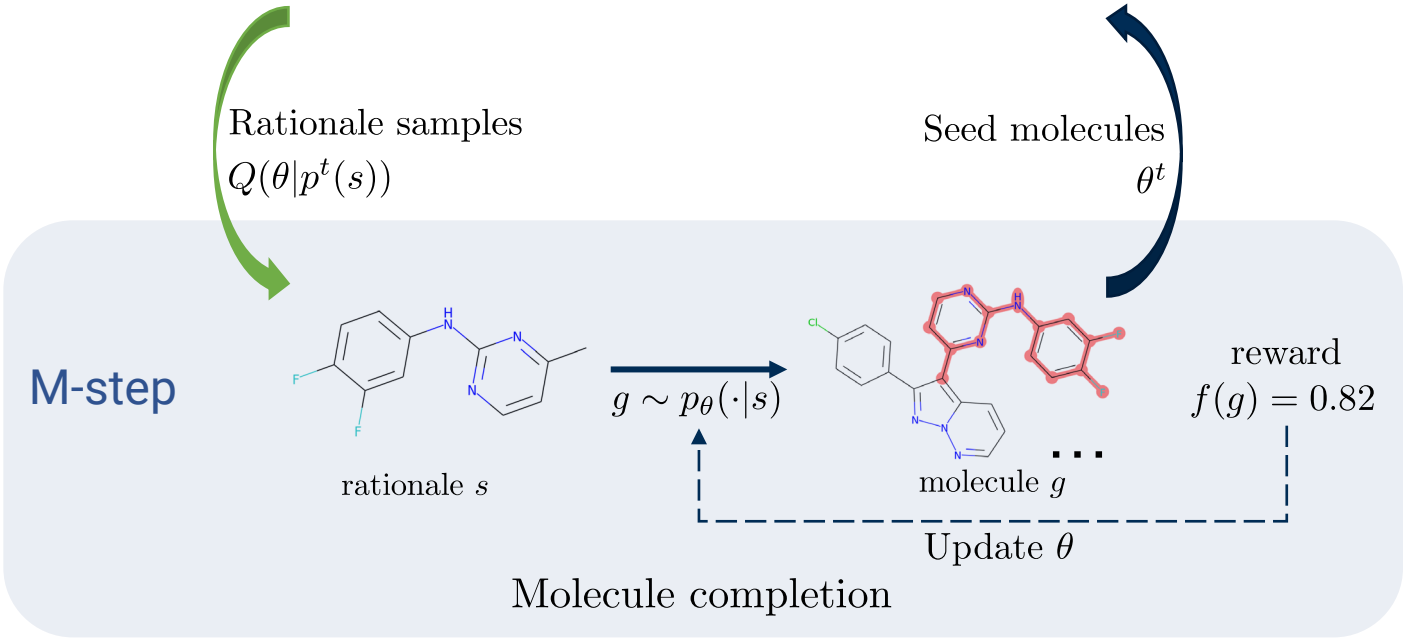

Property Score Distributions

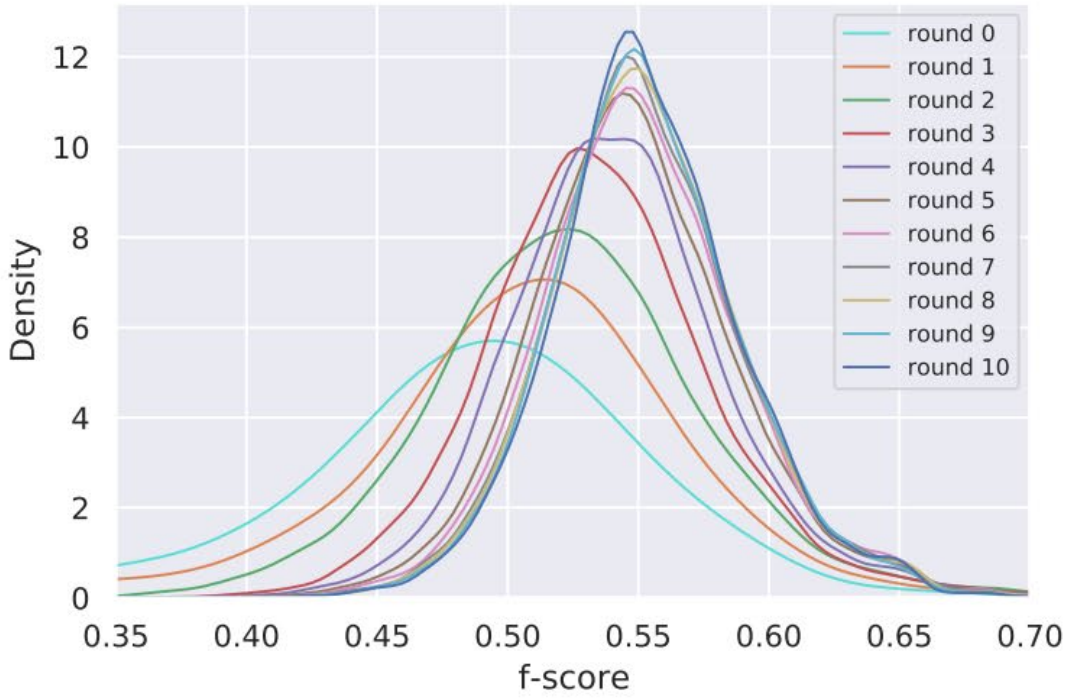

Distribution of $f(g)$ from each iteration.

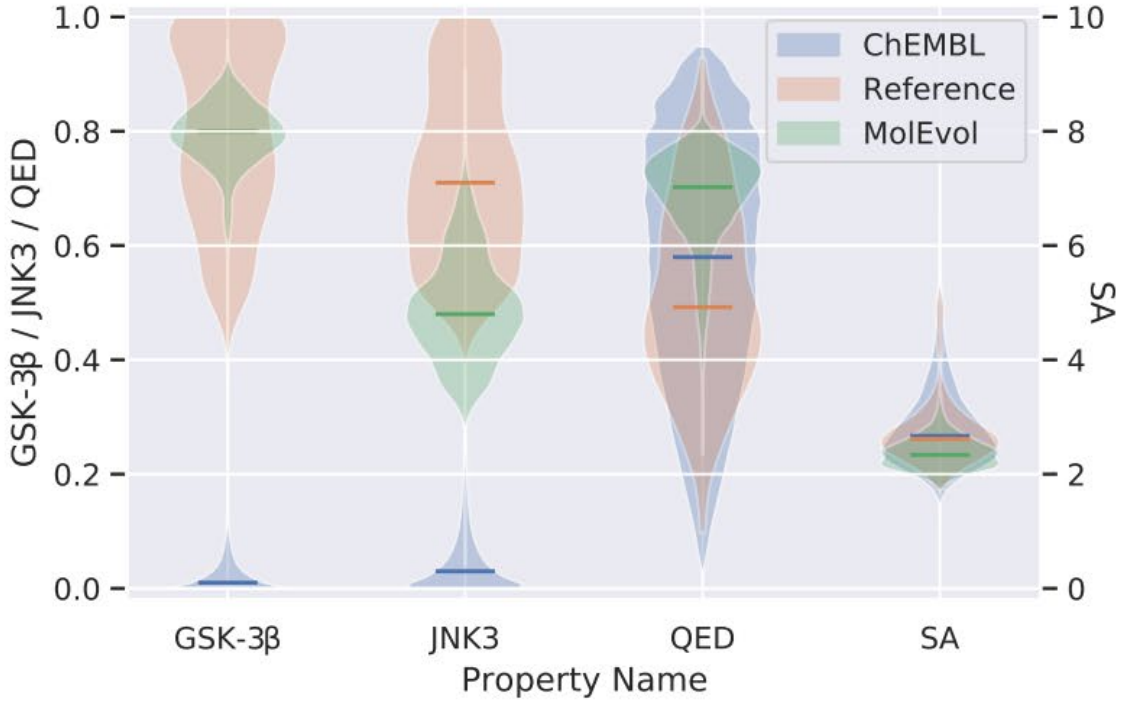

Distribution of the property scores.

Comparing to Baselines

Algorithm	MolEvol	[MCTS]	[FixM]	[FixR]	RationaleRL	REINVENT	MSO	GA-D(t)
Success rate	**93.0%**	77.7%	67.3%	66.3%	61.1%	46.6%	57.7%	62.0%
Novelty	**75.7%**	72.5%	67.4%	54.6%	57.4%	66.4%	28.6%	19.4%
Diversity	0.681	0.707	0.723	0.727	**0.749**	0.666	-	-

Ablation studies

RL baselines

EA baselines

Thanks for listening!

For more details, please refer to our paper/full slides/poster/repo:

Paper

Full Slides

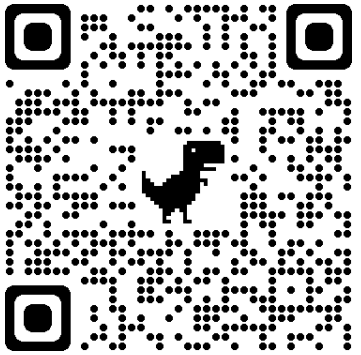

Poster

GitHub