Distributional GFlowNets with Quantile Flows

Dinghuai Zhang

Basics of Generative Flow Networks

- Goal: sample proportional to a given reward function $R(x), x \in \mathscr{X}$
- Approach: match sum of all flows into x to be equal to reward values
  - Flow matching: in-flow = out-flow (which incl. reward)
  - $$\sum_{s:(s \to s') \in \mathscr{A}} F(s \to s') = \sum_{s'':(s' \to s'') \in \mathscr{A}} F(s' \to s'')$$
  - Other methods: trajectory balance, …

Limitation of Current GFlowNets

- Stochastic reward setting
  - Given sufficiently large capacity and computation, the obtained GFlowNet would sample with probability proportional to $\exp\left(\mathbb{E}[\log R(x)]\right)$
  - Cannot capture uncertainty / stochasticity

Distributional Modeling

- Distributional modeling of GFlowNet edge flows
- Distributional flow matching
  - $$Z(s') = \sum_{(s \to s') \in \mathcal{A}} Z(s \to s') = \sum_{(s' \to s'') \in \mathcal{A}} Z(s' \to s'')$$
  - Z denotes random variable
  - $=$ denotes equation in distribution

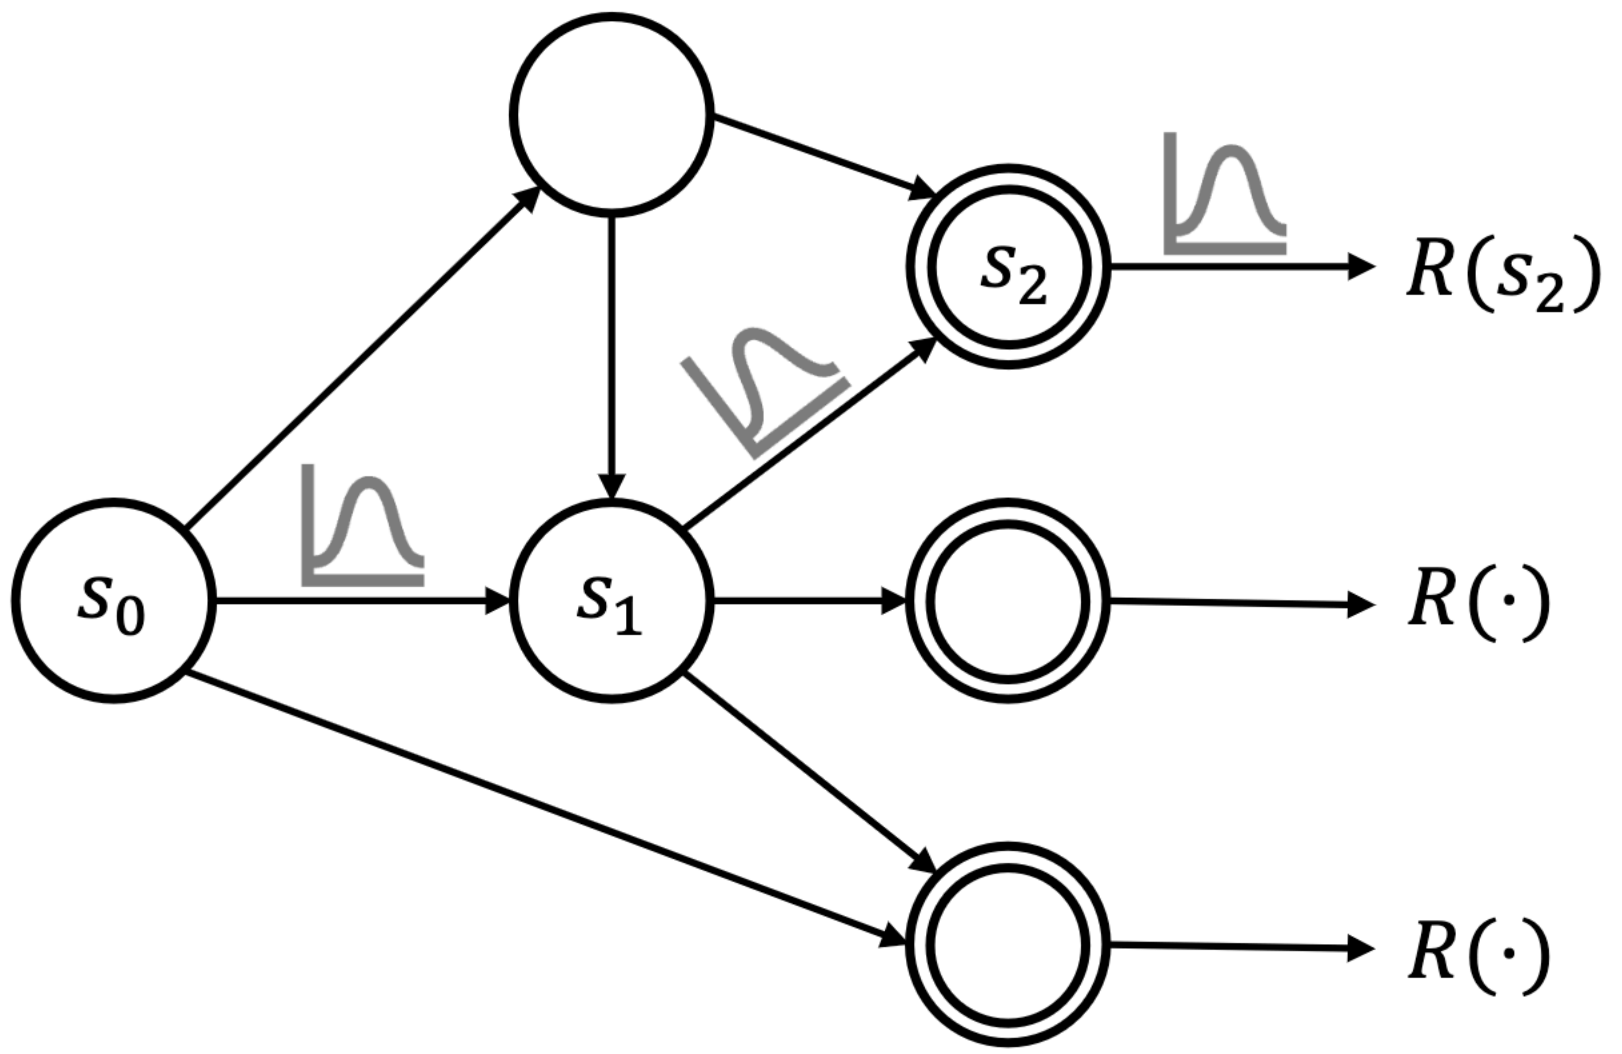

Quantile Flows

- Model the β-quantile function of distribution of each edge flow
    - $Z_\beta(s \to s'; \theta), \quad \beta \in [0,1]$
- **Quantile matching** algorithm

$$\delta^{\beta,\tilde{\beta}}(s';\theta) = \log \sum_{s' \to s''} \exp Z_{\tilde{\beta}}^{\log}(s' \to s''; \theta)$$

- $$-\log \sum_{s \to s'} \exp Z_{\beta}^{\log}(s \to s'; \theta),$$
- minimize δ with quantile regression

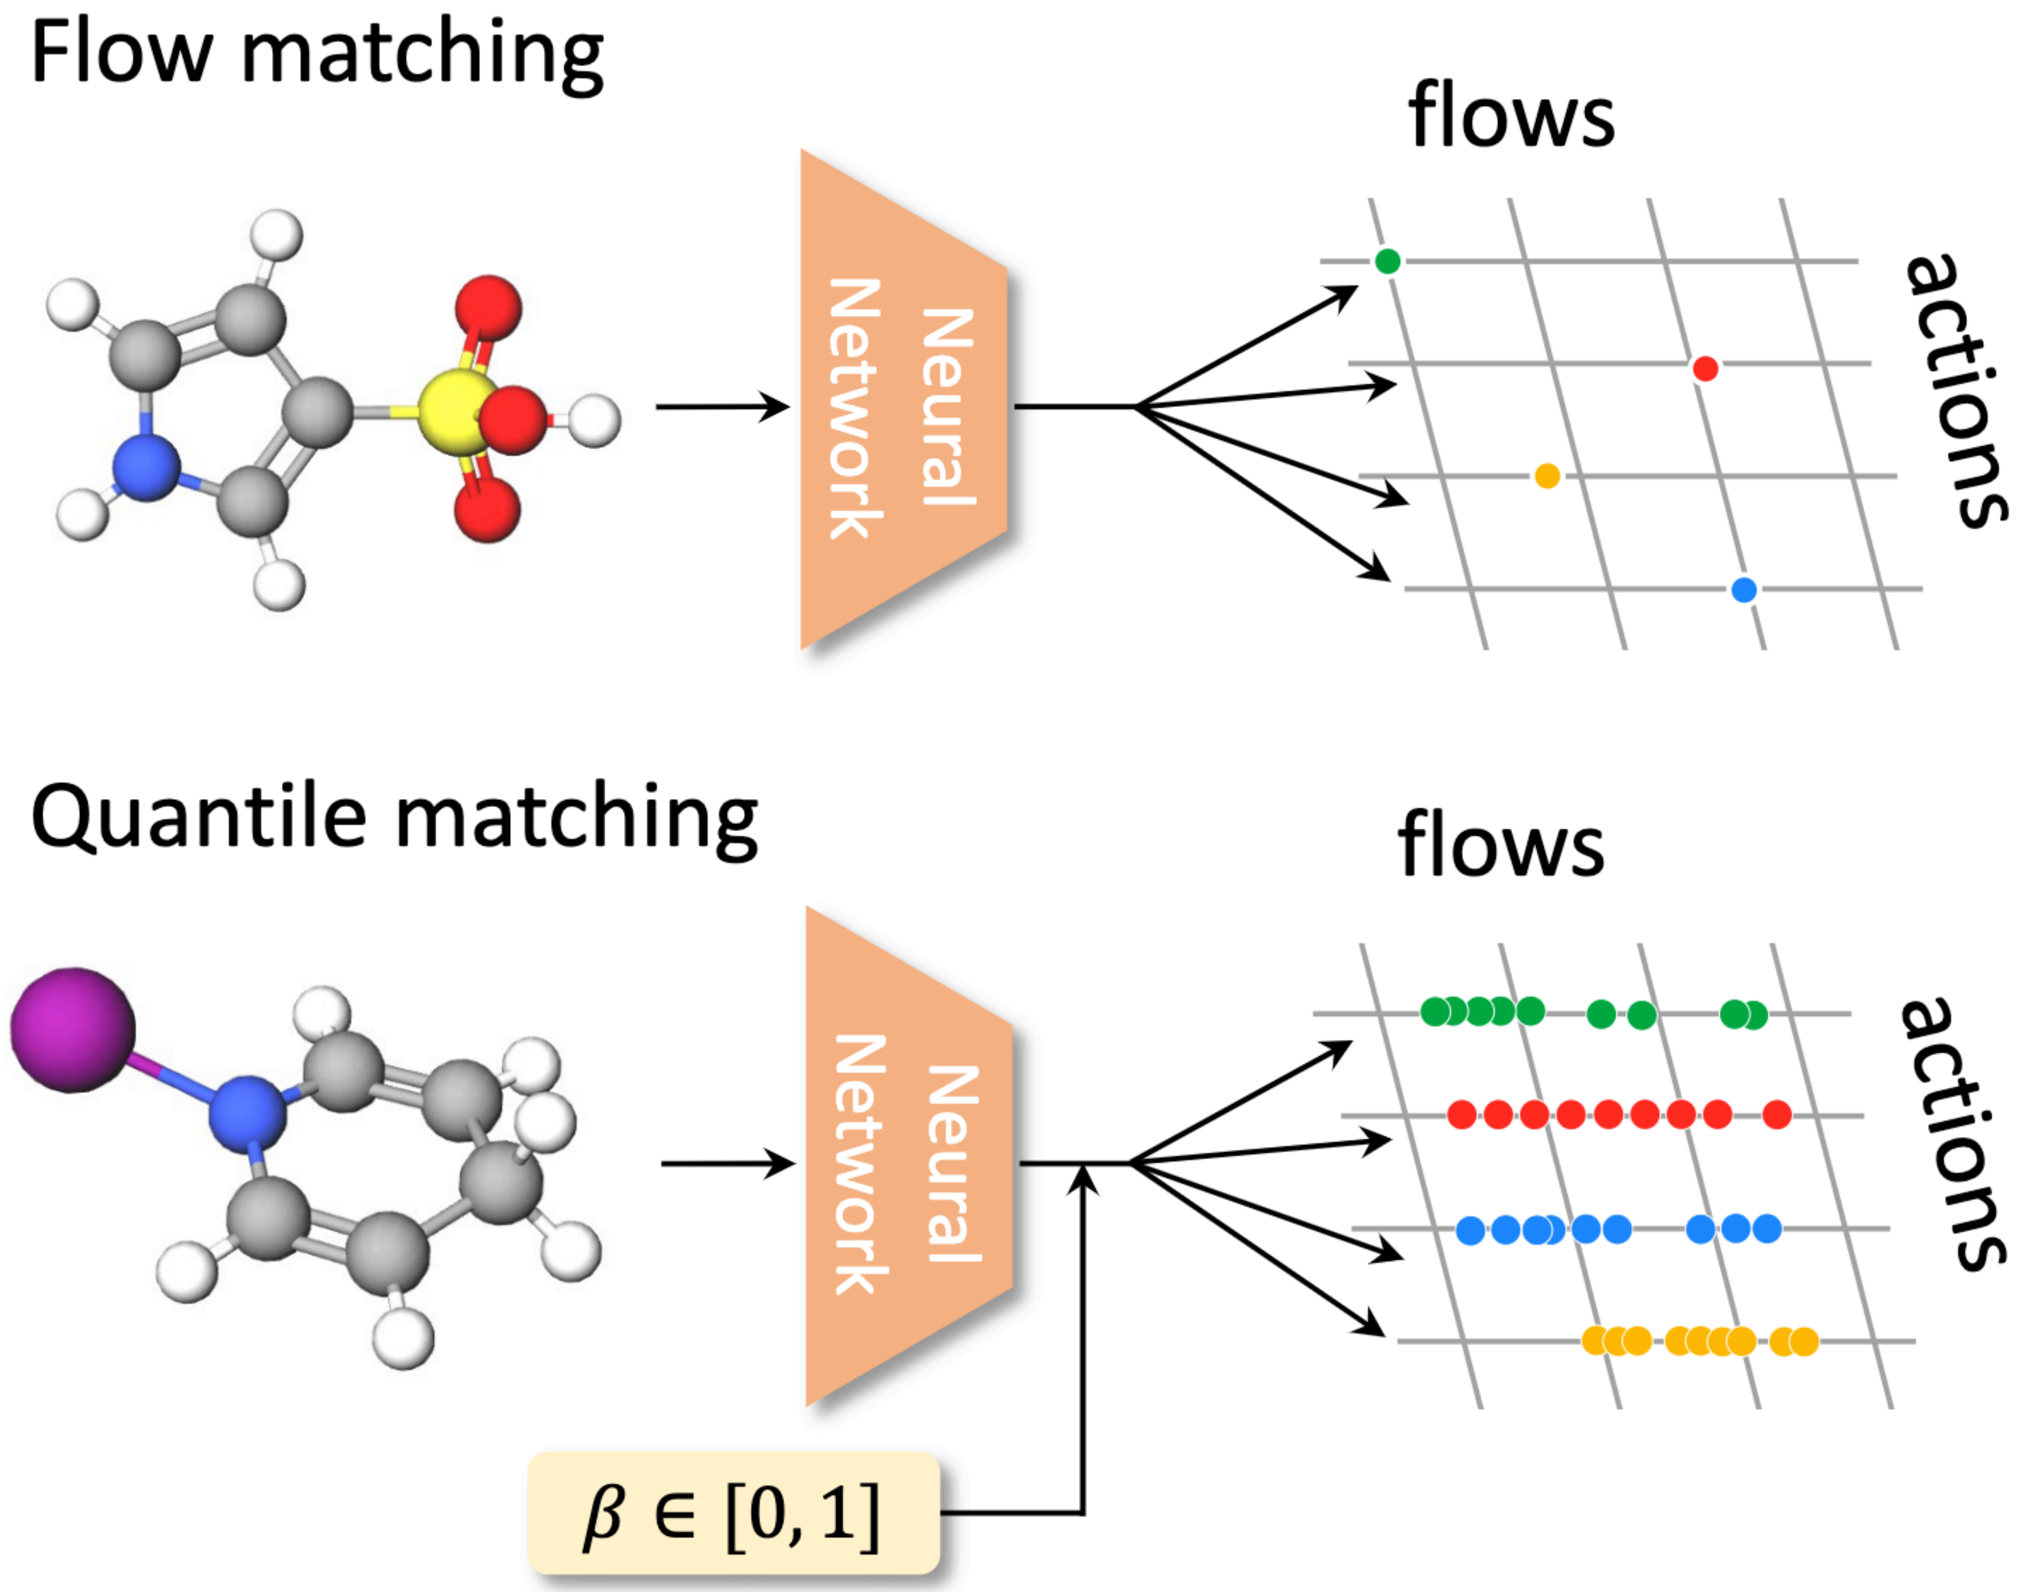

Risk-sensitive Flows

- Standard risk measure (mean): $\mathbb{E}\,[Z] = \int_0^1 Q_Z(\beta)d\beta$
- Distortion risk measure: $\mathbb{E}^g\,[Z] = \int_0^1 Q_Z(g(\beta))d\beta, \quad g : [0,1] \to [0,1]$
- Risk-averse modeling example:
  - $g(\beta) = 0.1 * \beta$ => only estimate the mean of the lowest 10% data
    - conditional value-at-risk (CVaR)
  - Using risk-averse distortion functions $g(\cdot)$ leads to conservative behaviors

Risk-averse GFlowNets

- Risky hypergrid environment
- Distributional GFlowNets with rise-averse distortion function step less into risky regions (i.e., lower violation rate in figure)
    - rise-averse: CVaR(0.1), Wang(-0.75)
    - risk-neutral: CPW

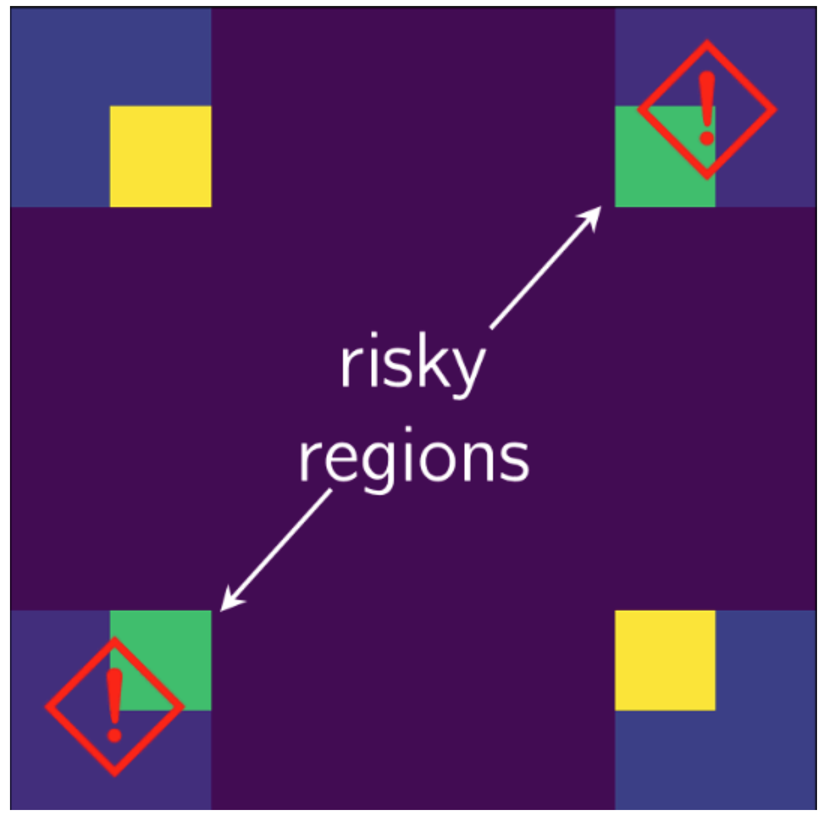

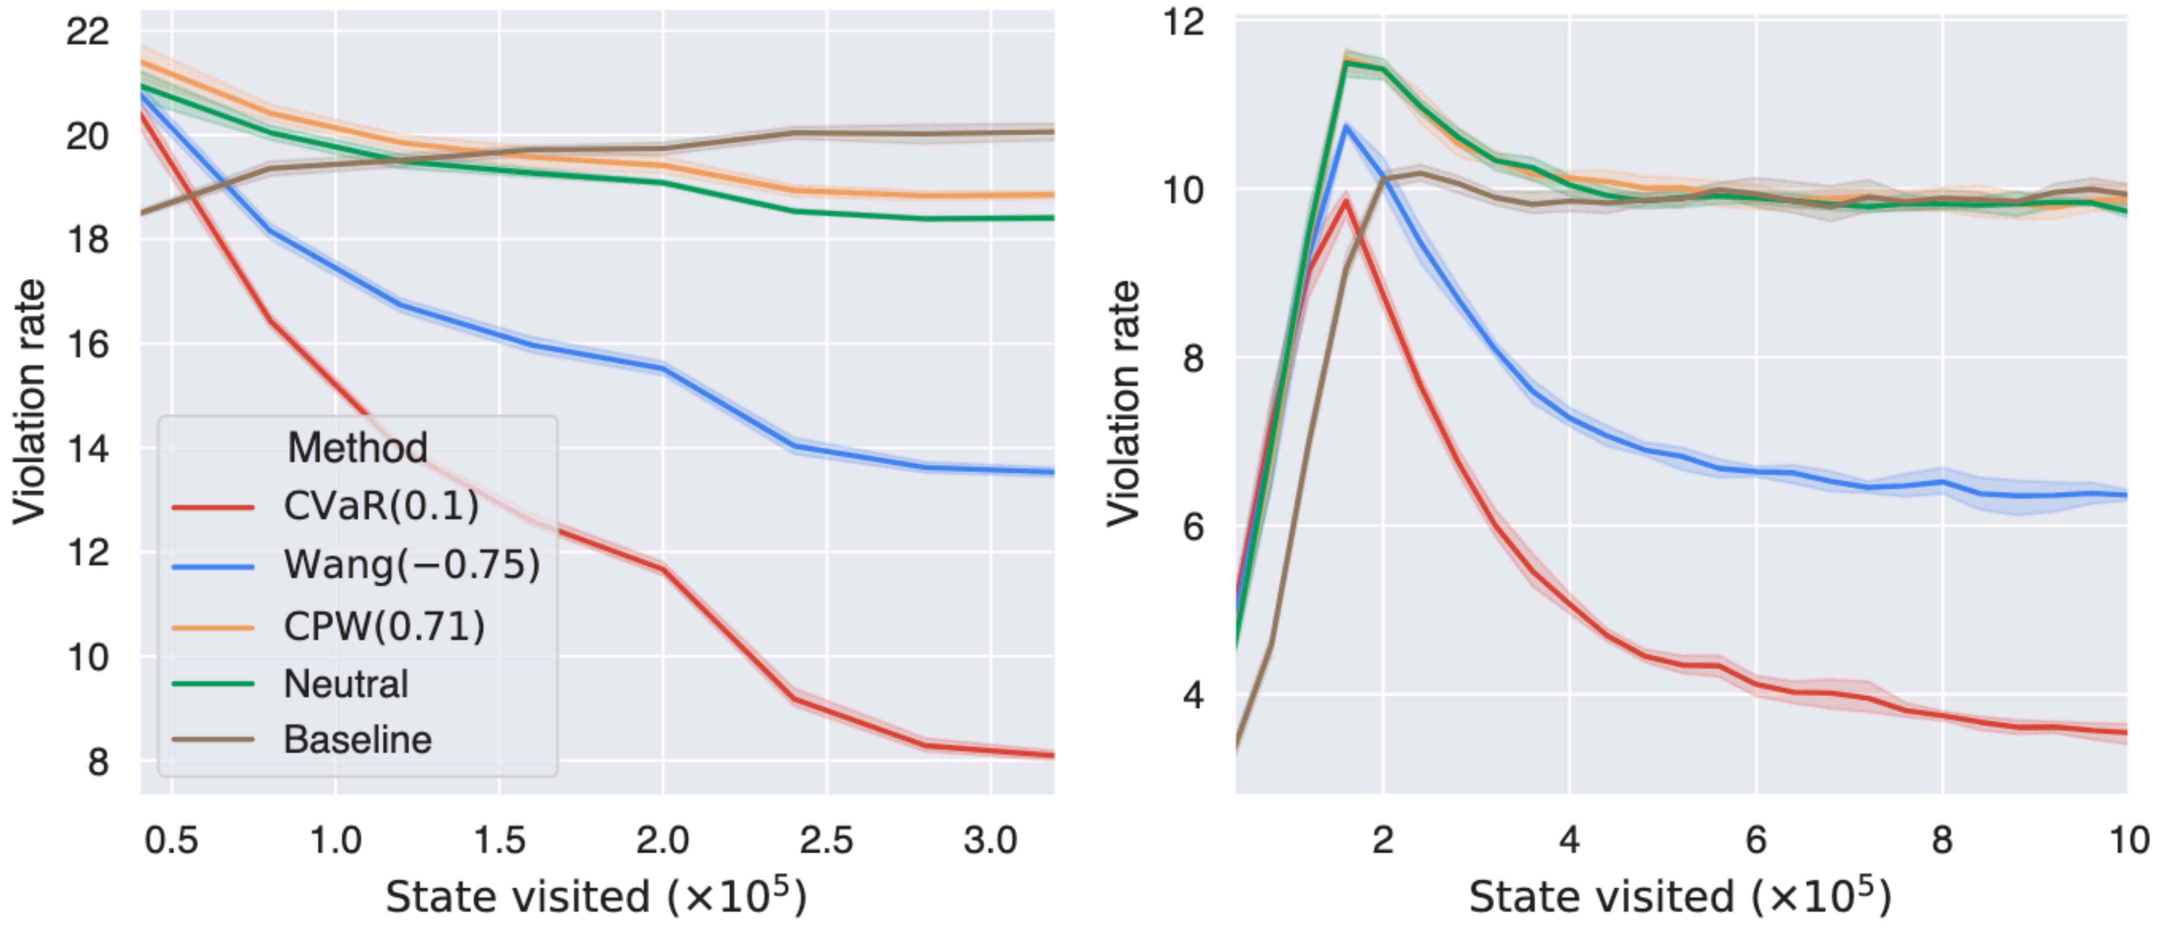

(a) small

(b) large

Benchmarking Experiments

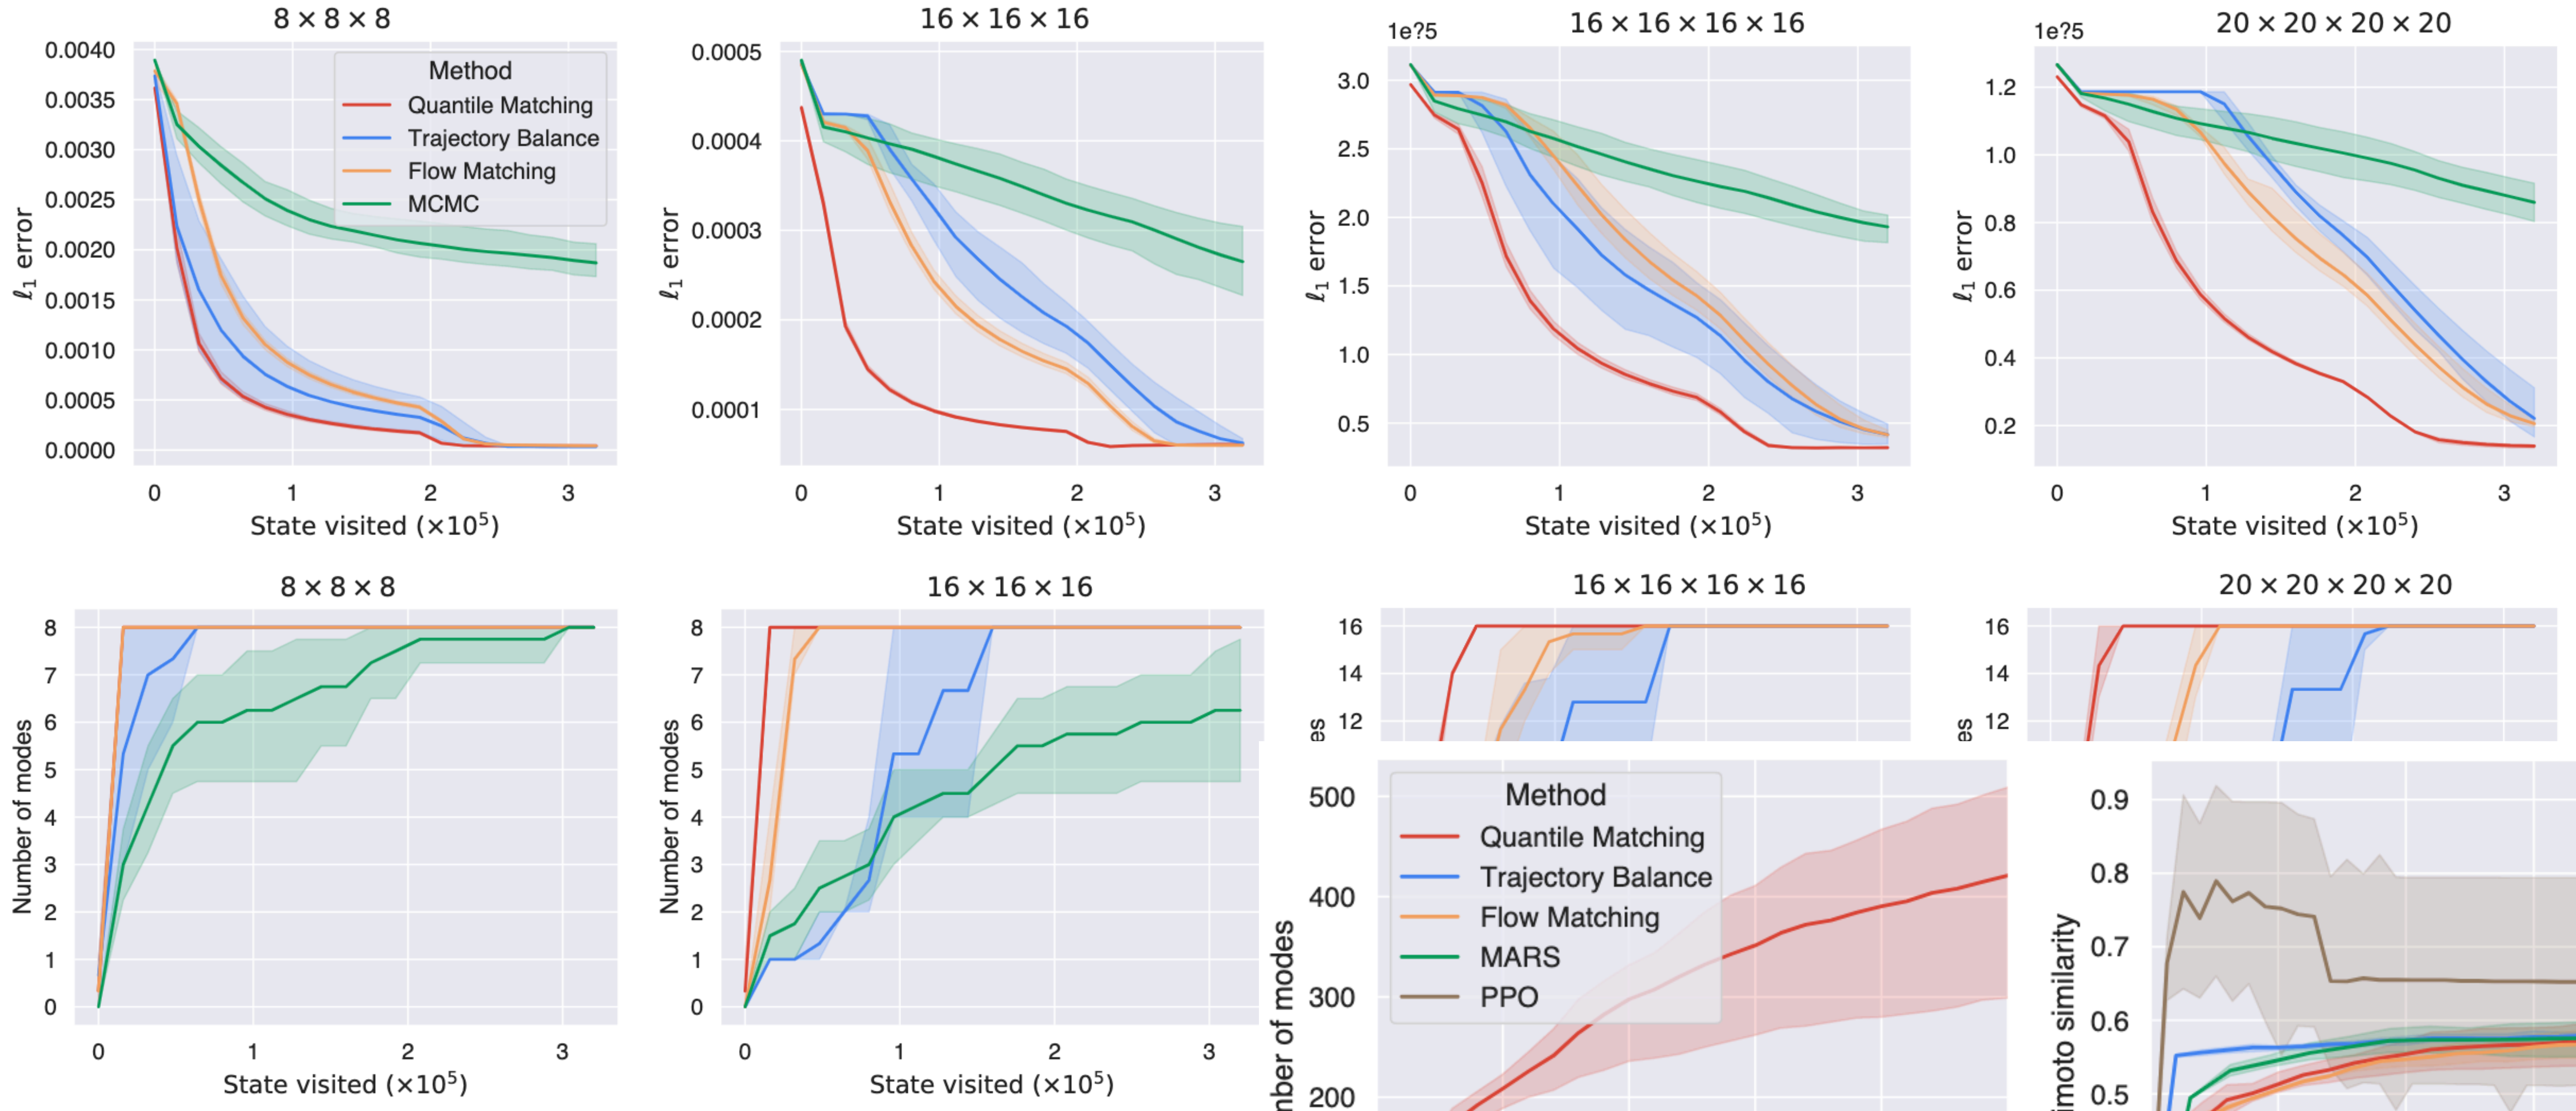

HyperGrid

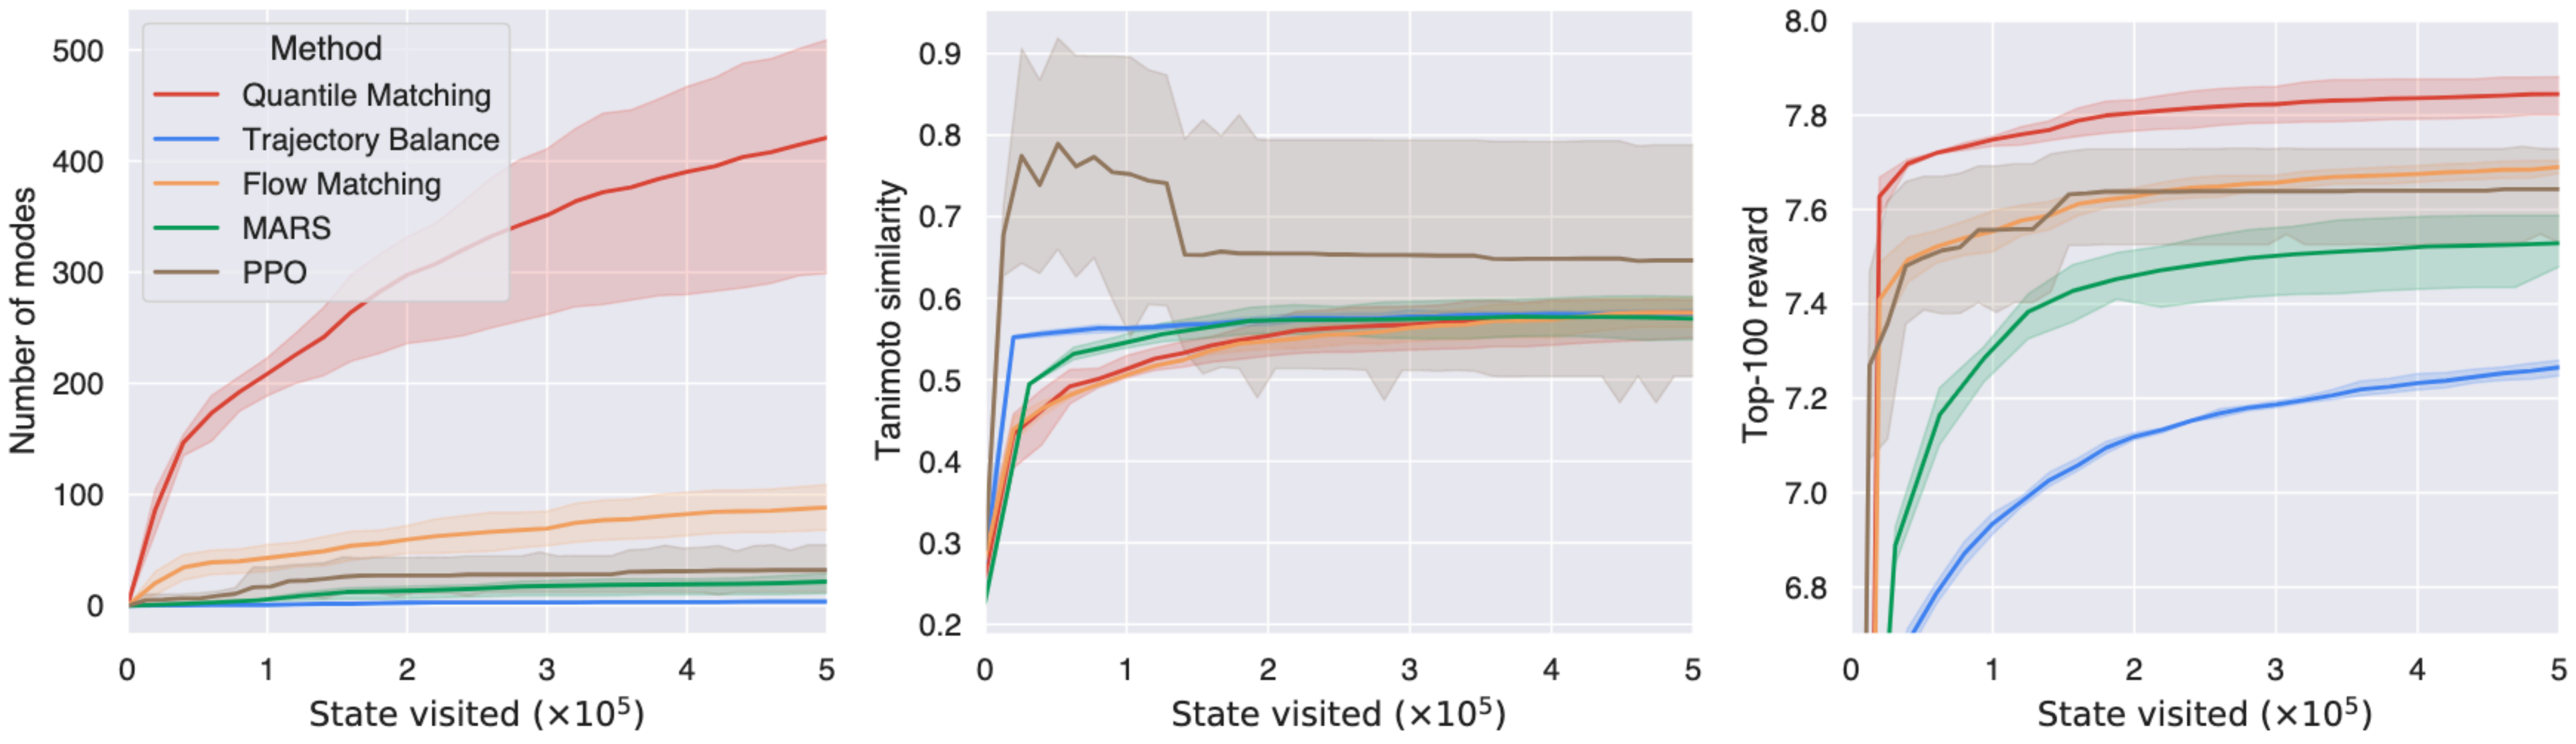

Molecule synthesis experiments

Thank you very much!