Learning GFlowNets From Partial Episodes For Improved Convergence And Stability

Kanika Madan, Jarrid Rector-Brooks, Maksym Korablyov, Emmanuel Bengio, Moksh Jain, Andrei Cristian Nica, Tom Bosc, Yoshua Bengio, Nikolay Malkin

Generative Flow Networks (GFlowNets)

Generative Flow Networks (GFlowNets)

Flow of unnormalized probabilities

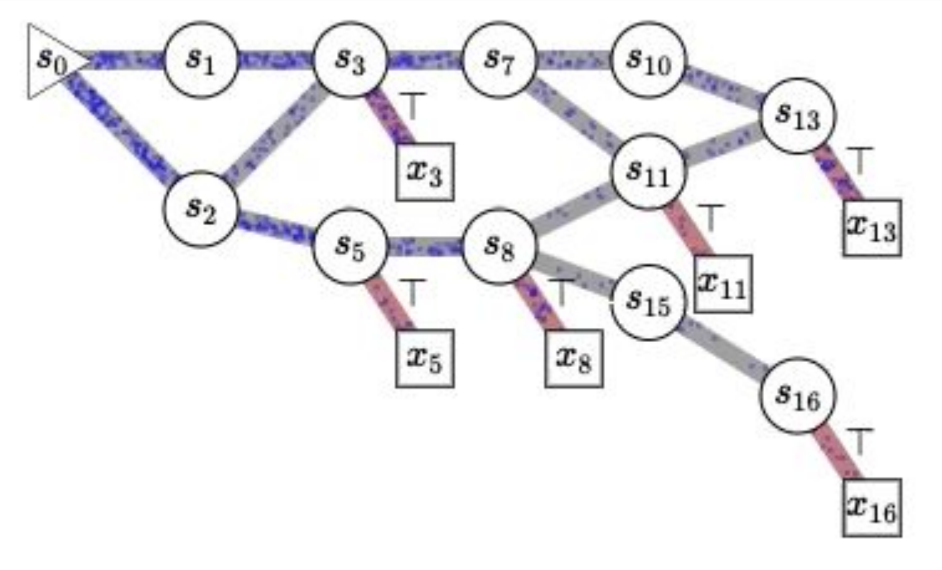

Generative Flow Networks (GFlowNets)

Flow of unnormalized probabilities

Flow Based Network

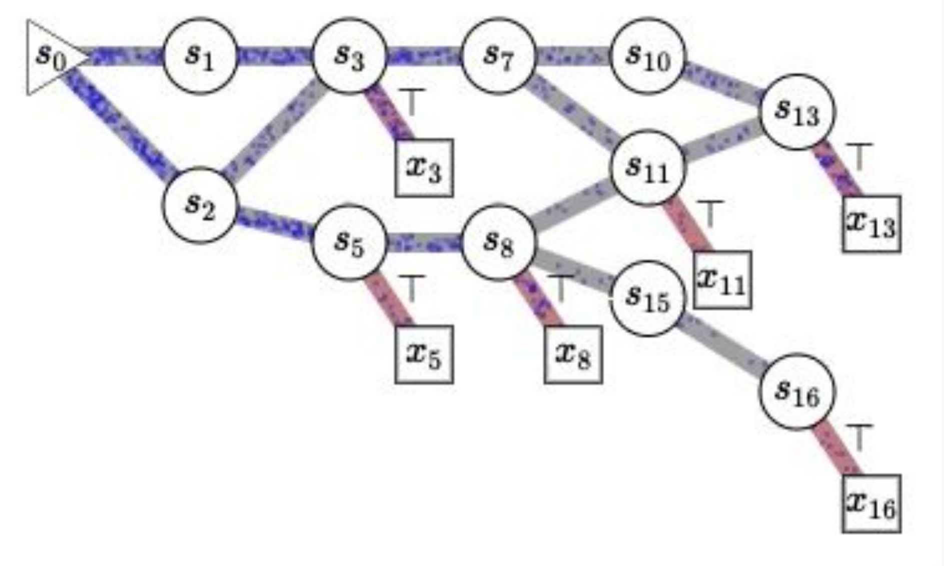

Generative Flow Networks (GFlowNets)

Flow of unnormalized probabilities

Flow Based Network

- Analogy:

Water flowing from source to sink

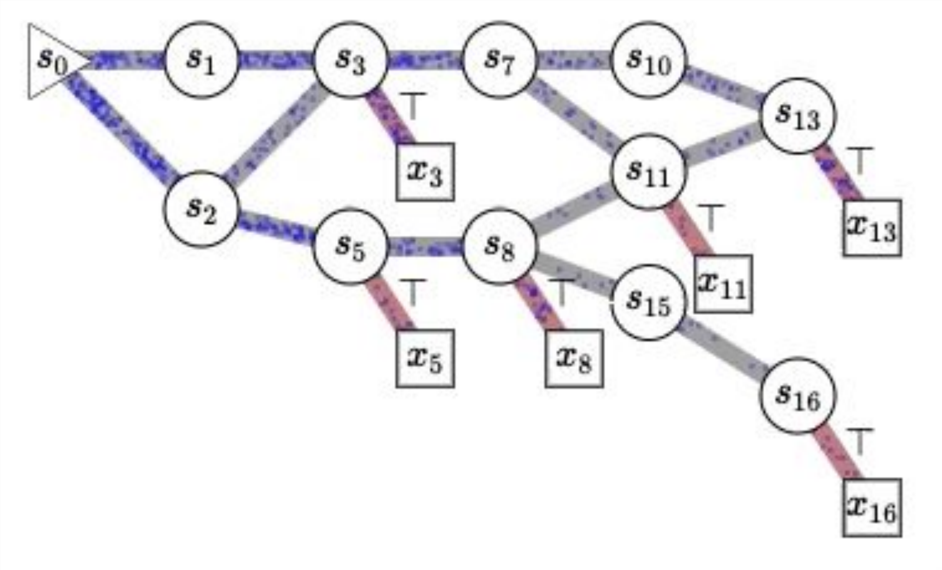

Generative Flow Networks (GFlowNets)

- Flow Based Network
- Generative Model

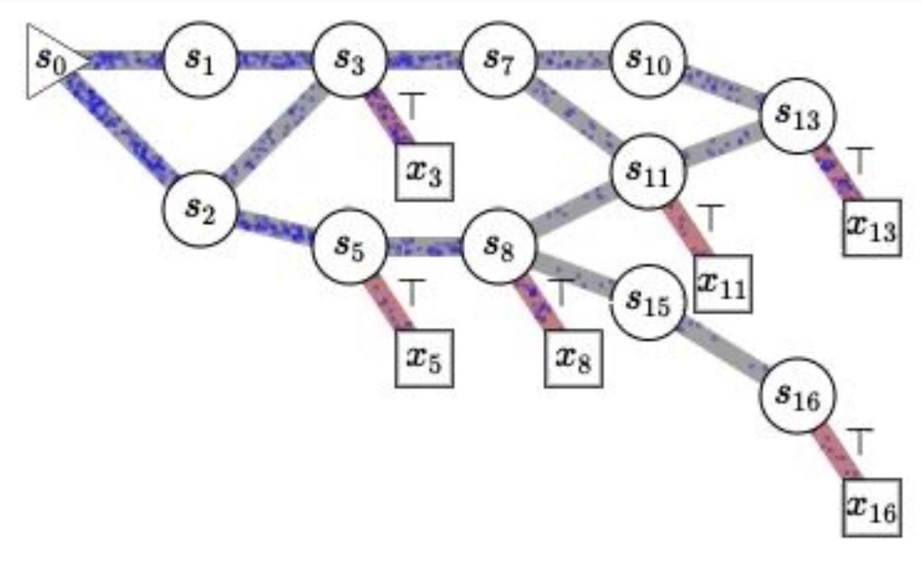

Generative Flow Networks (GFlowNets)

- Flow Based Network
- Generative Model
  - Stochastic Policy

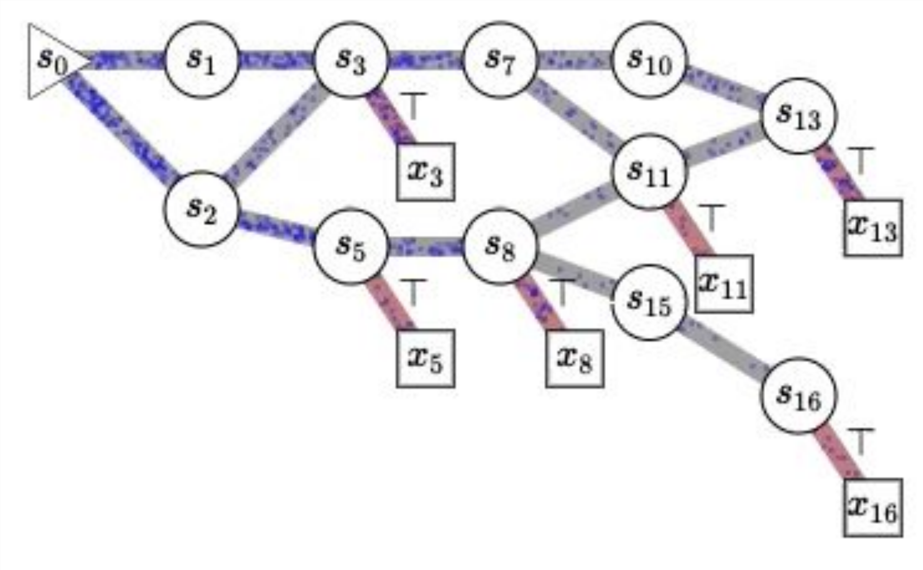

Generative Flow Networks (GFlowNets)

- Flow Based Network
- Generative Model
  - Stochastic Policy
- Generates objects sequentially

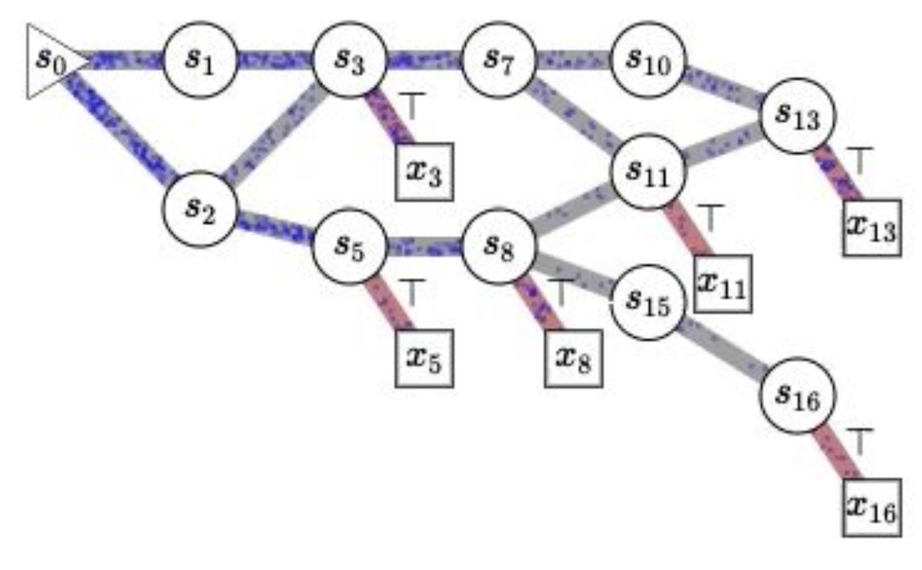

Generative Flow Networks (GFlowNets)

- Flow Based Network
- Generative Model
  - Stochastic Policy
- Generates objects sequentially
- Directed Acyclic Graph

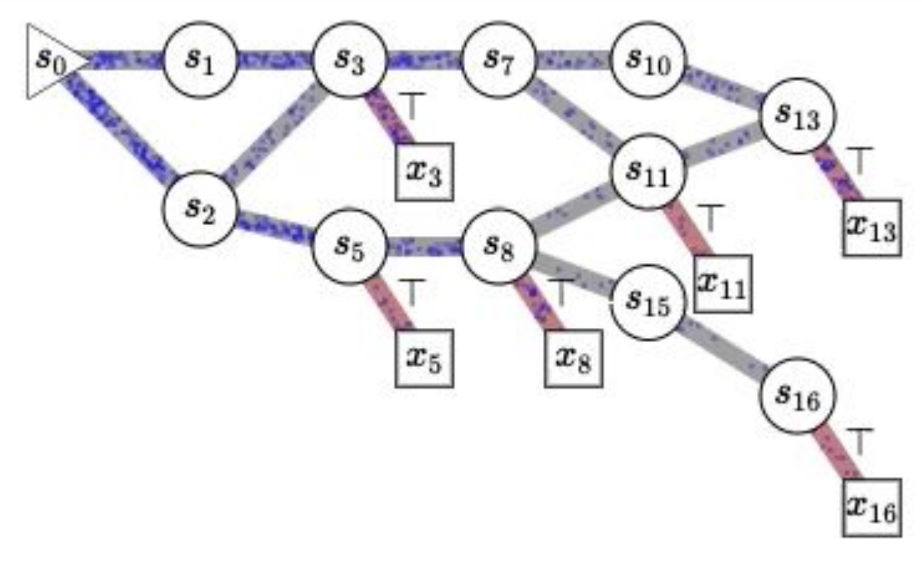

→ Generate Objects Sequentially

Generate Objects Sequentially

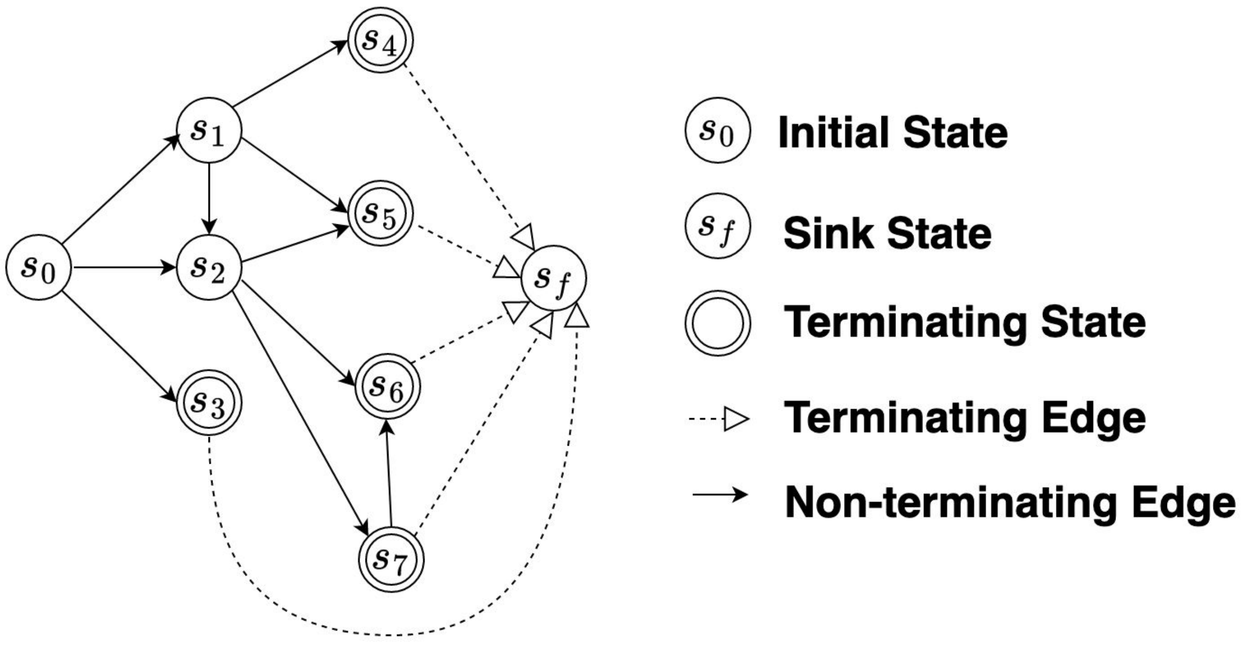

Generate Objects Sequentially - Trajectory 1

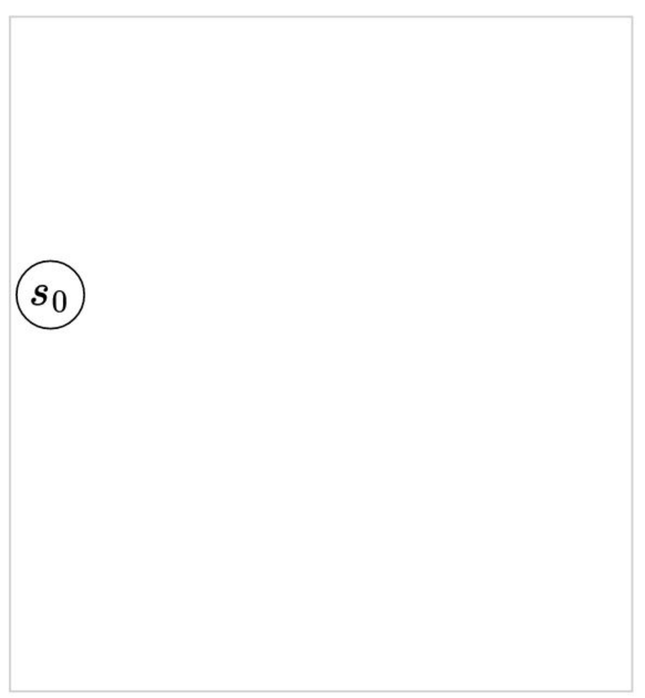

Generate Objects Sequentially - Trajectory 1

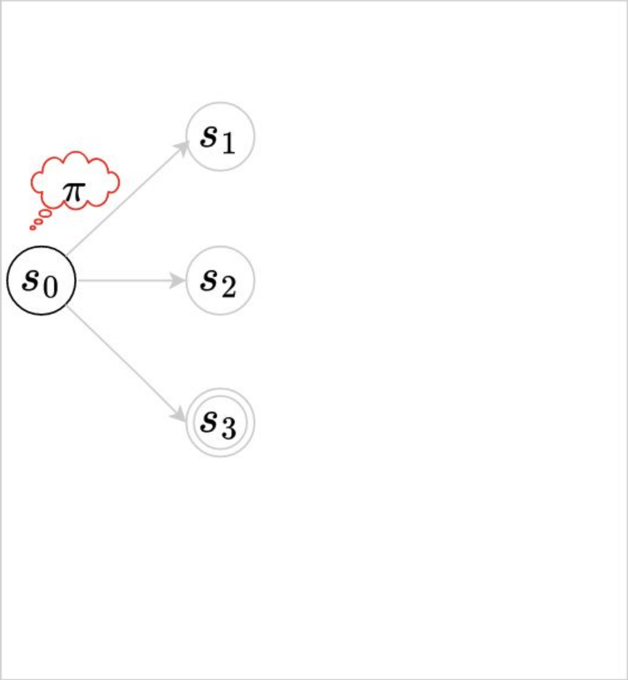

Generate Objects Sequentially - Trajectory 1

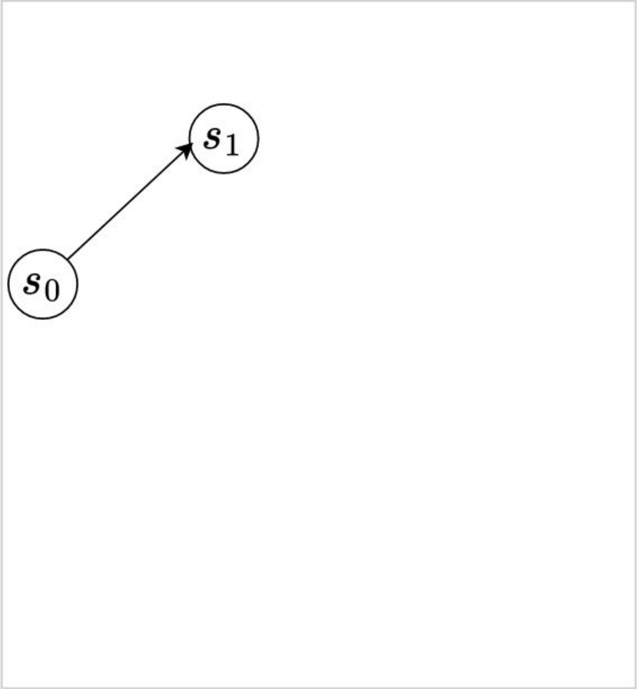

Generate Objects Sequentially - Trajectory 1

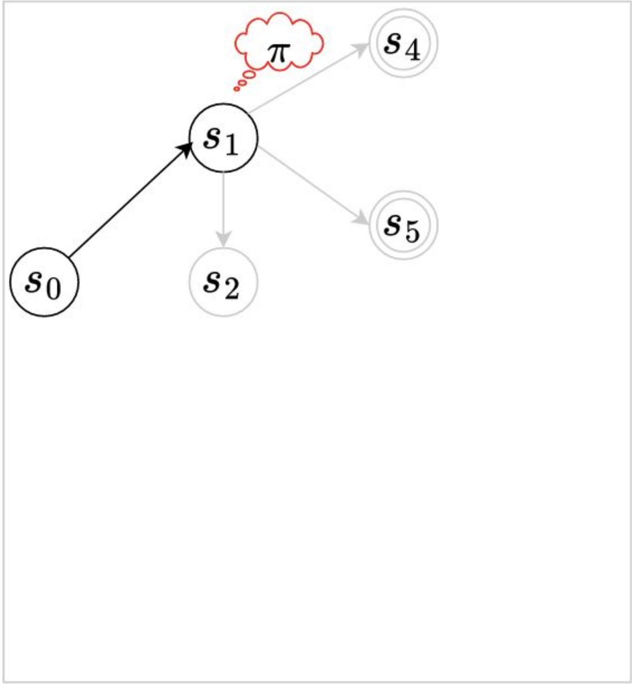

Generate Objects Sequentially - Trajectory 1

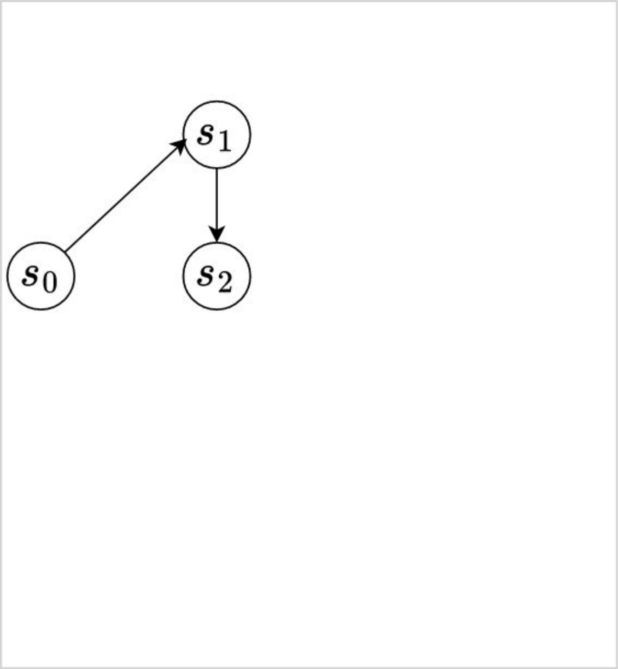

Generate Objects Sequentially - Trajectory 1

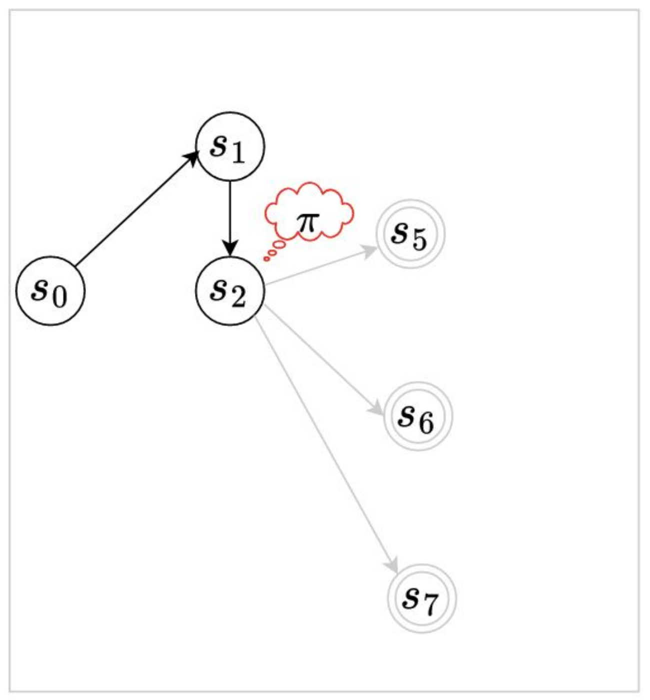

Generate Objects Sequentially - Trajectory 1

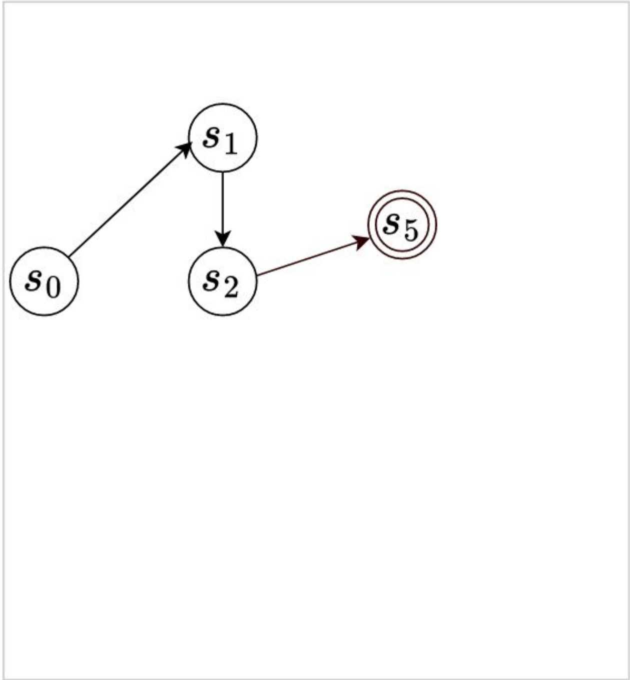

Generate Objects Sequentially - Trajectory 1

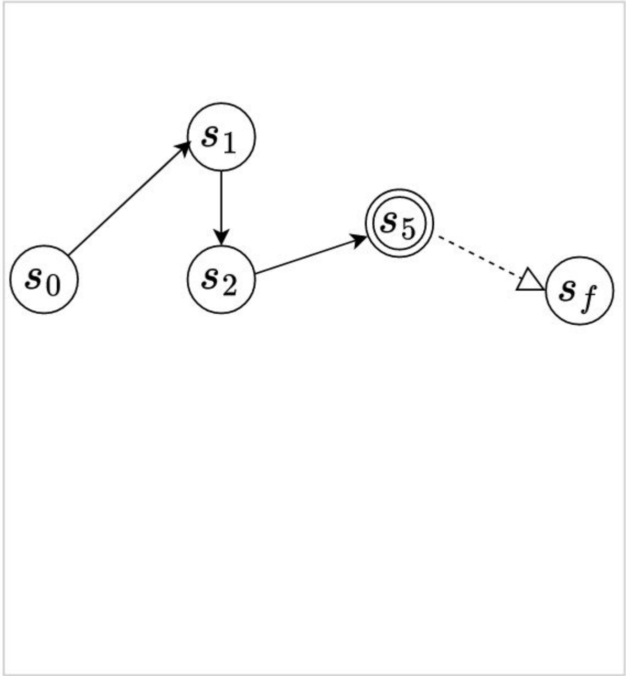

Generate Objects Sequentially

$\rightarrow$ Stochastic Policy

Generate Objects Sequentially - Trajectory 2

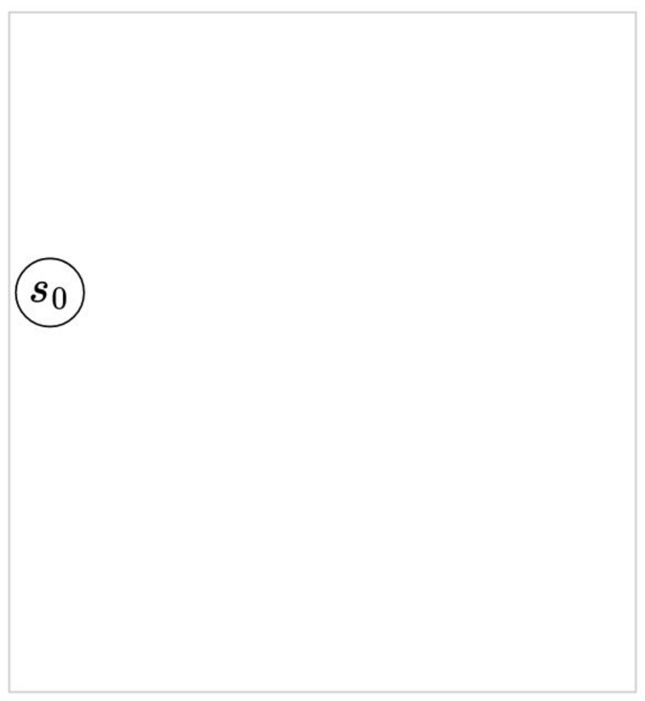

Generate Objects Sequentially - Trajectory 2

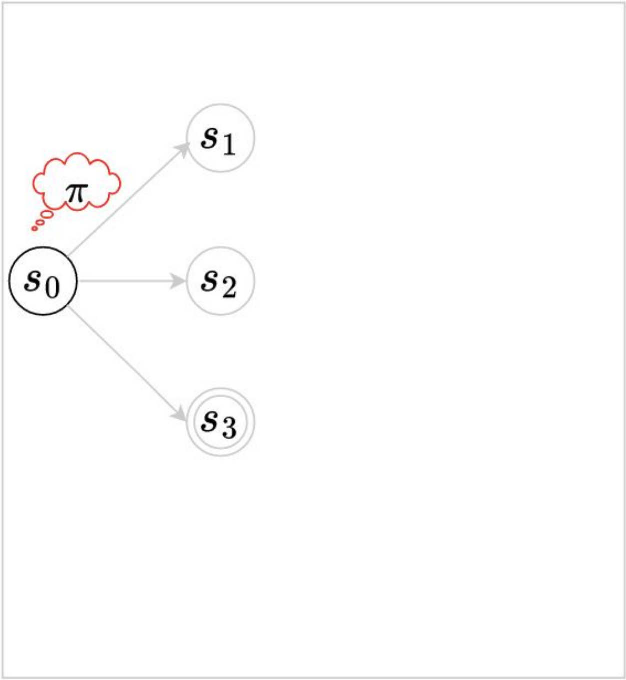

Generate Objects Sequentially - Trajectory 2

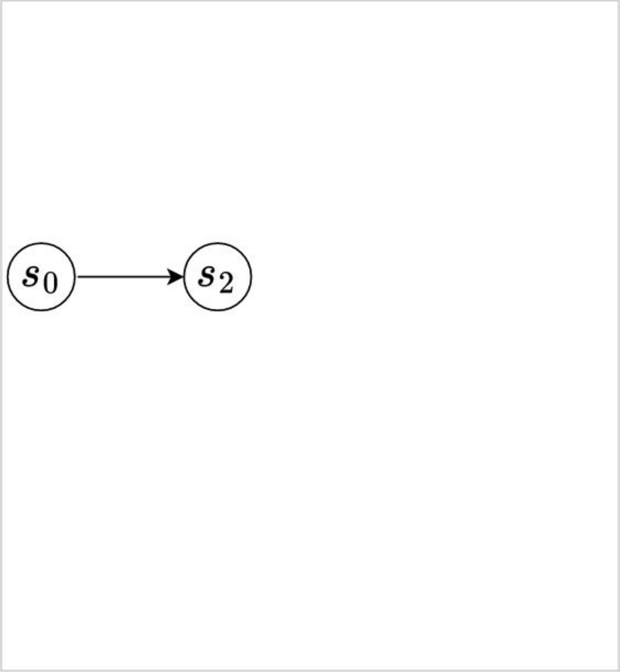

Generate Objects Sequentially - Trajectory 2

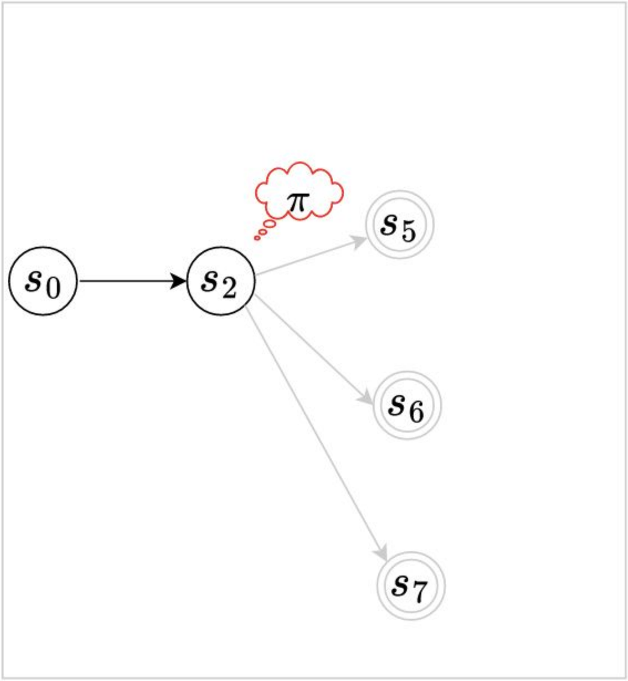

Generate Objects Sequentially - Trajectory 2

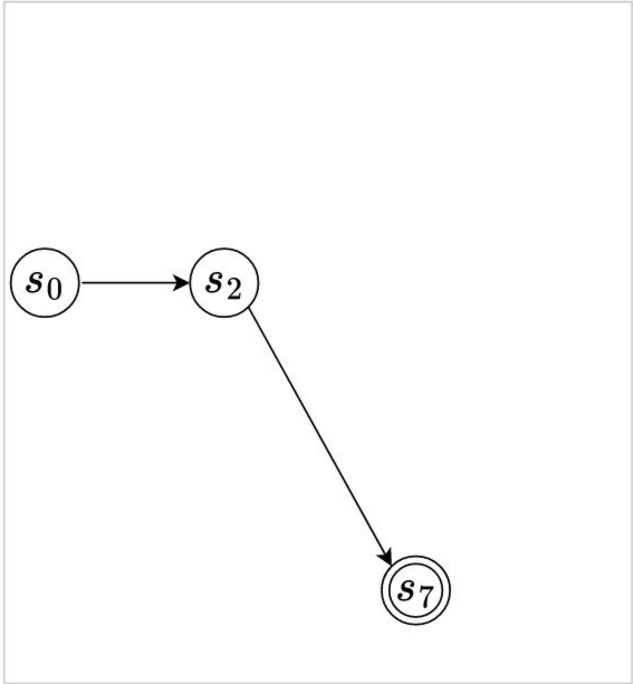

Generate Objects Sequentially - Trajectory 2

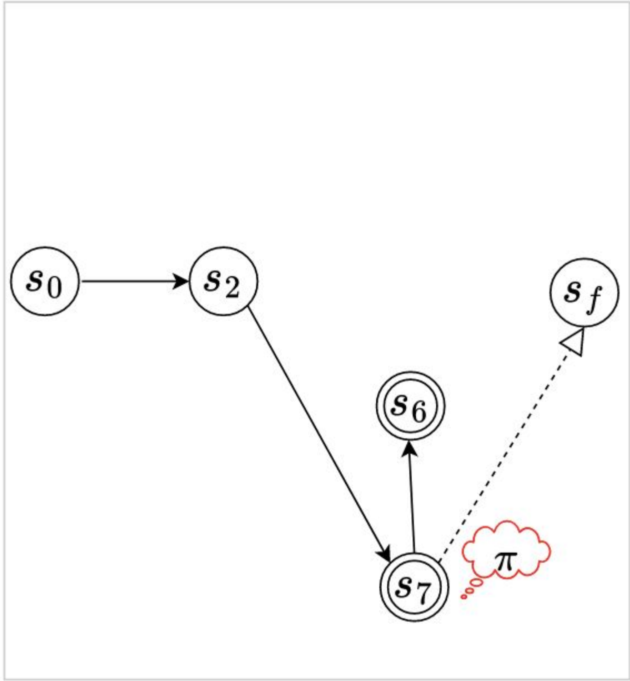

Generate Objects Sequentially - Trajectory 2

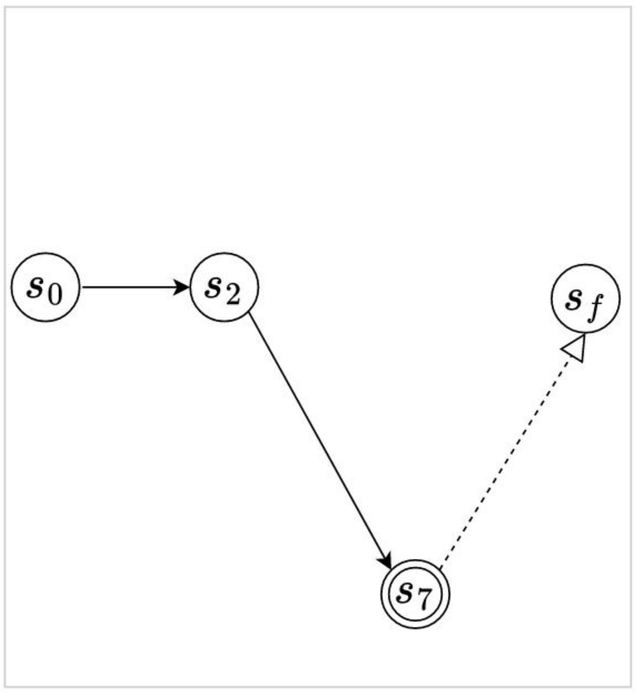

Generative Model

- Objects Generated Proportional to Reward

Generate Objects Sequentially - Trajectories

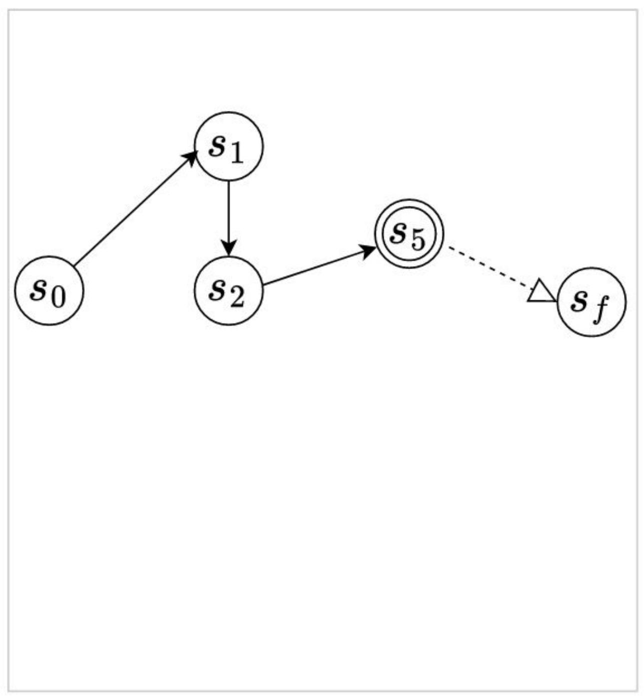

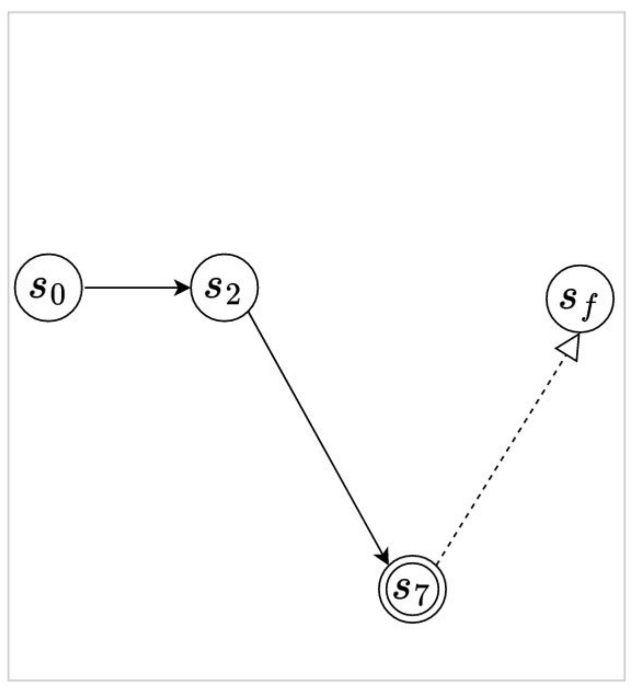

Generate Objects Sequentially - Trajectories

$$p(x_1) = p(s_0 s_1 s_2 s_5) \propto R(x_1)$$

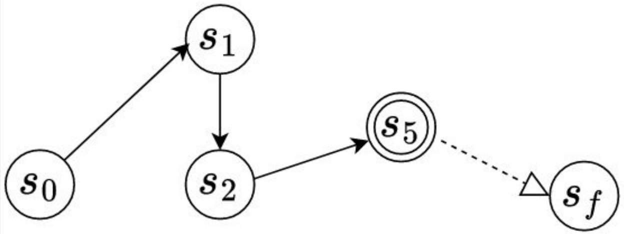

Generate Objects Sequentially - Trajectories

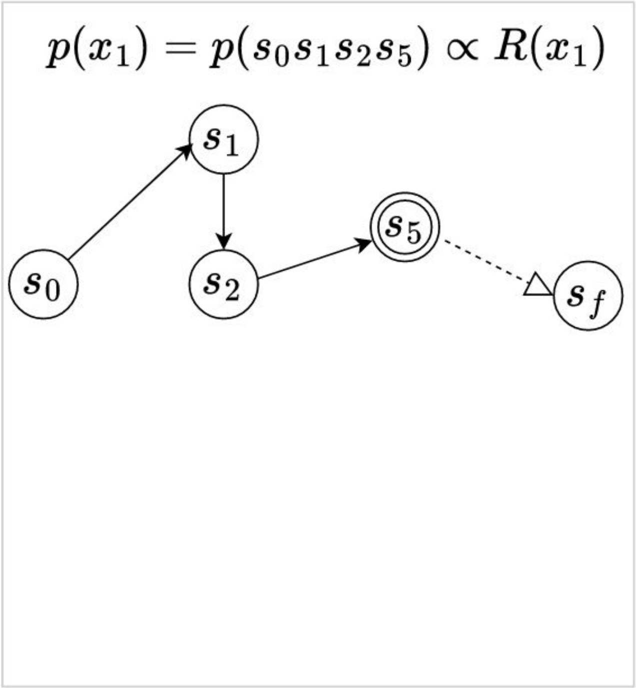

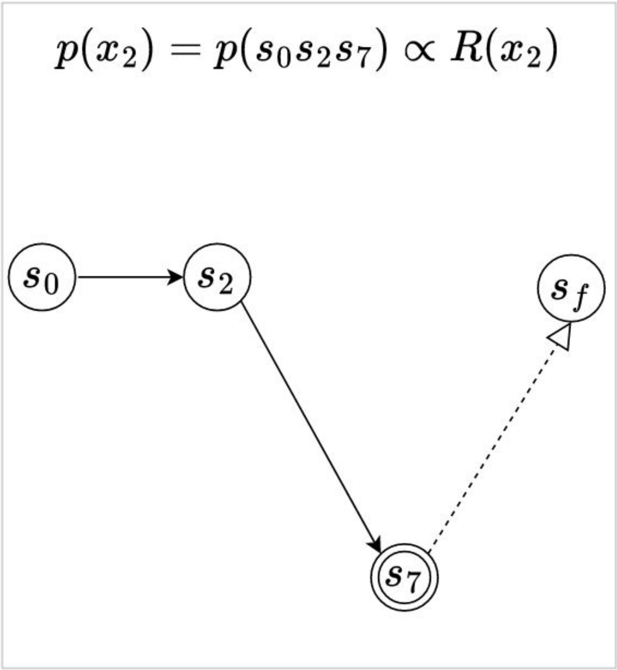

Training GFlowNets

GFlowNets - Training

- Flow Consistency Equations

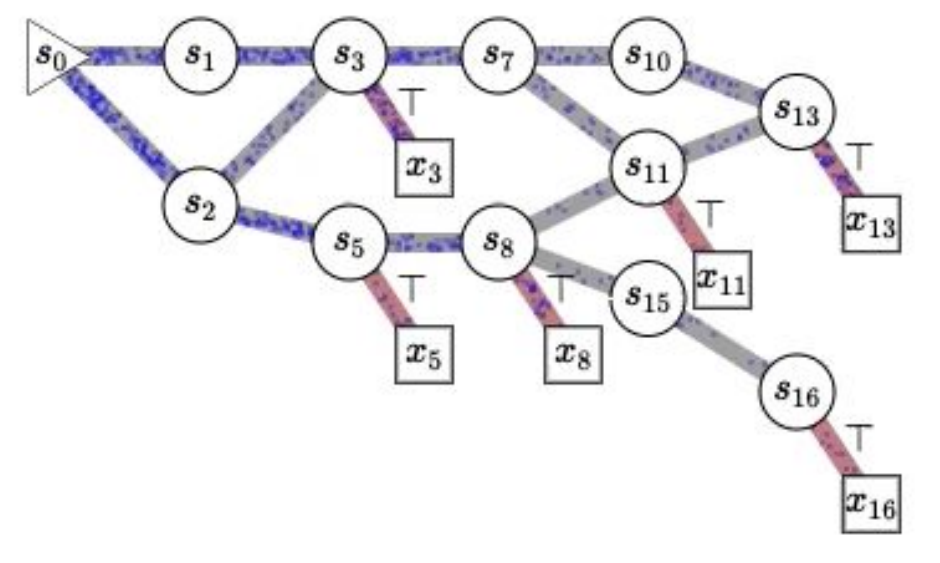

GFlowNets - Training

- Flow Consistency Equations
- Forward Flow = Backward Flow

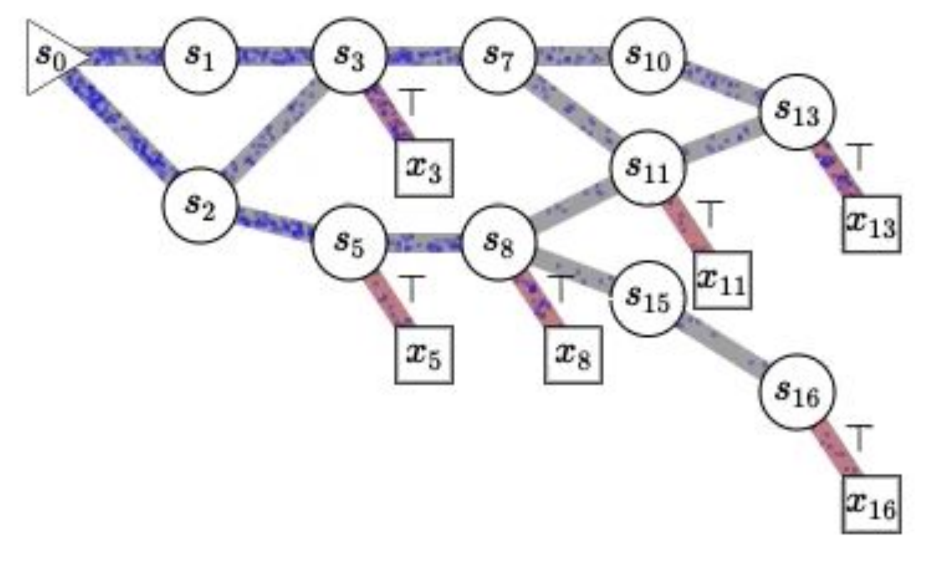

GFlowNets Training Objectives

Flow Matching Objective

- Flow Consistency Equations

- Flow Matching

$$\mathcal{L}_{\text{FM}}(s) = \left(\log \frac{\sum_{s:(s\to t)\in\mathcal{A}} F(s\to t;\theta) + \epsilon}{\sum_{u:(t\to u)\in\mathcal{A}} F(t\to u;\theta) + \epsilon}\right)^2$$

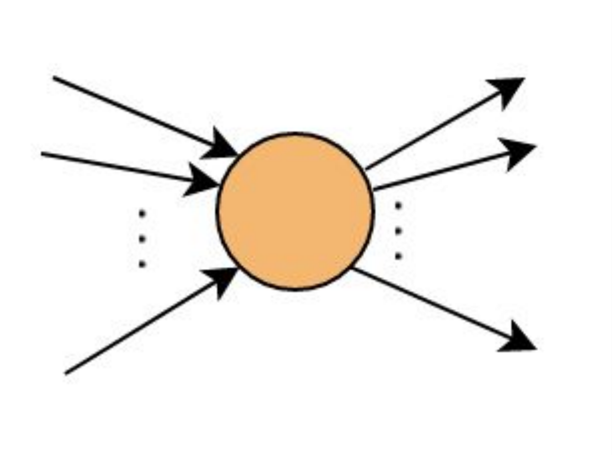

Flow Matching Objective

- Flow Consistency Equations
- Flow Matching
  - State level flow matching

$$\mathcal{L}_{\mathrm{FM}}(s) = \left(\log \frac{\sum_{s:(s \to t) \in \mathcal{A}} F(s \to t; \theta) + \epsilon}{\sum_{u:(t \to u) \in \mathcal{A}} F(t \to u; \theta) + \epsilon} \right)^2$$

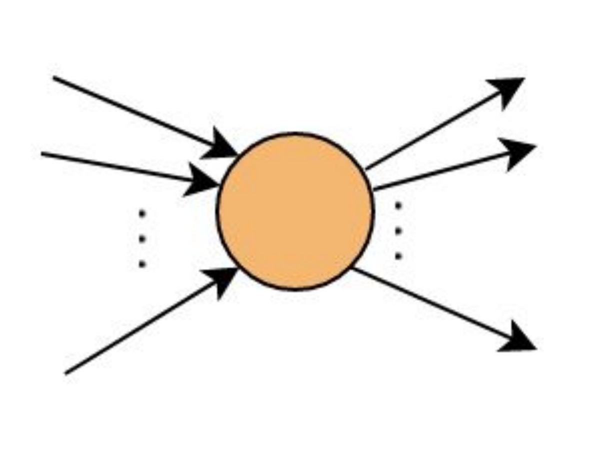

Detailed Balance Objective

- Flow Consistency Equations

- Detailed Balance:

$$\mathcal{L}_{\mathrm{DB}}(s, s') = \left(\log \frac{F_\theta(s) P_F(s' \mid s; \theta)}{F_\theta(s') P_B(s \mid s'; \theta)} \right)^2$$

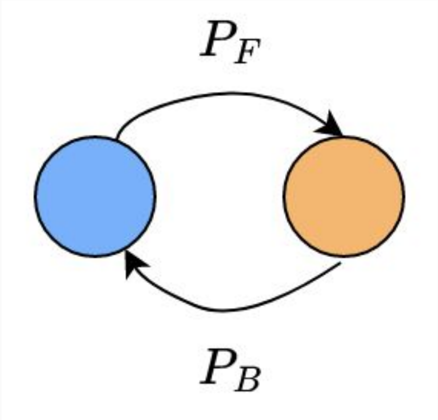

Detailed Balance Objective

- Flow Consistency Equations
- Detailed Balance:
  - Edge level flow matching

$$\mathcal{L}_{\mathrm{DB}}(s, s') = \left(\log \frac{F_\theta(s) P_F(s' \mid s; \theta)}{F_\theta(s') P_B(s \mid s'; \theta)}\right)^2$$

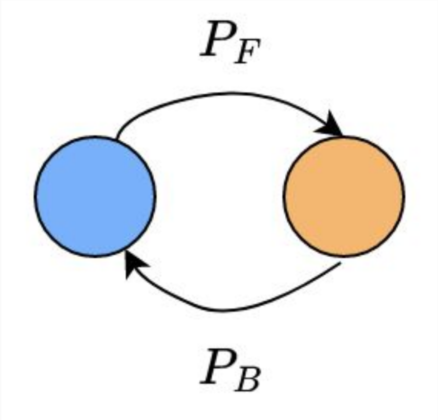

Trajectory Balance Objective

- Flow Consistency Equations

- Trajectory Balance:

$$\mathcal{L}_{\mathrm{TB}}(\tau) = \left(\log \frac{Z_\theta P_F(\tau;\theta)}{R(x_\tau) P_B(\tau \mid x_\tau;\theta)} \right)^2$$

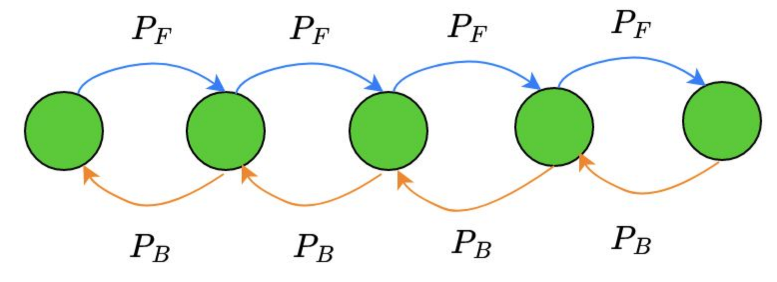

Trajectory Balance Objective

- Flow Consistency Equations

- Trajectory Balance:
  - Trajectory level flow matching

$$\mathcal{L}_{\mathrm{TB}}(\tau) = \left(\log \frac{Z_\theta P_F(\tau;\theta)}{R(x_\tau) P_B(\tau \mid x_\tau;\theta)} \right)^2$$

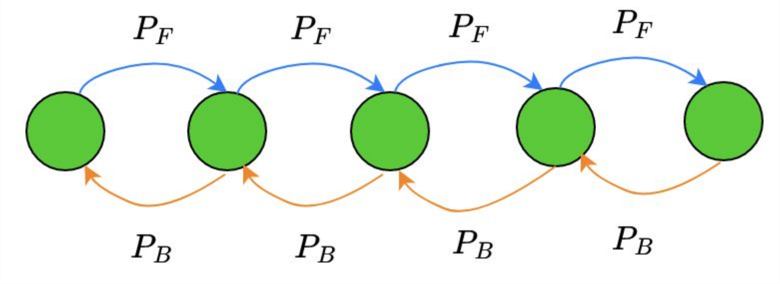

SubTrajectory Balance:
GFlowNet Objectives Unified

GFlowNets Training Objectives

Objective	Parametrization	Locality
FM	edge flow $F(s{\to}t;\theta)$	state s
DB	state flow $F(s;\theta)$, policies $P_F(-\|-;\theta)$, $P_B(-\|-;\theta)$	action $s{\to}t$
TB	initial state flow Z_θ, policies $P_F(-\|-;\theta)$, $P_B(-\|-;\theta)$	complete trajectory τ

GFlowNets Training Objectives

Objective	**Parametrization**	**Locality**
FM	edge flow $F(s \rightarrow t; \theta)$	state s
DB	state flow $F(s; \theta)$, policies $P_F(-\|-; \theta)$, $P_B(-\|-; \theta)$	action $s \rightarrow t$
TB	initial state flow Z_θ, policies $P_F(-\|-; \theta)$, $P_B(-\|-; \theta)$	complete trajectory τ
SubTB(λ)	state flow $F(s; \theta)$, policies $P_F(-\|-; \theta)$, $P_B(-\|-; \theta)$	(partial) trajectory τ

SubTrajectory Balance

Subtrajectory balance	(partial) trajectory τ

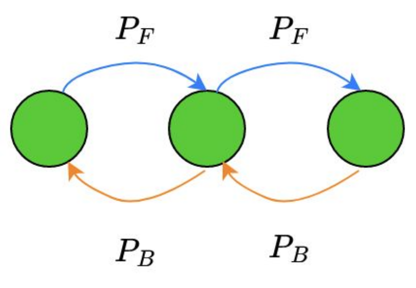

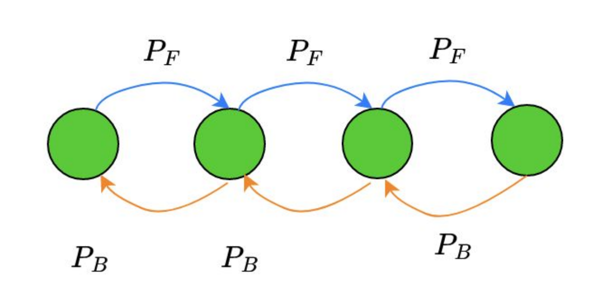

SubTrajectory Balance

Subtrajectory balance	(partial) trajectory τ

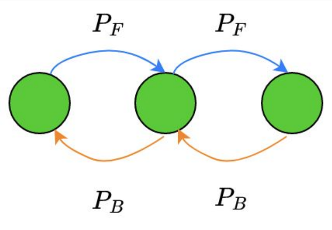

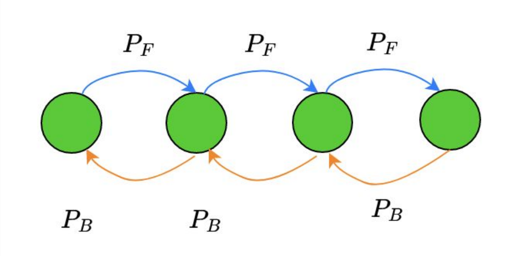

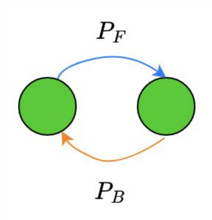

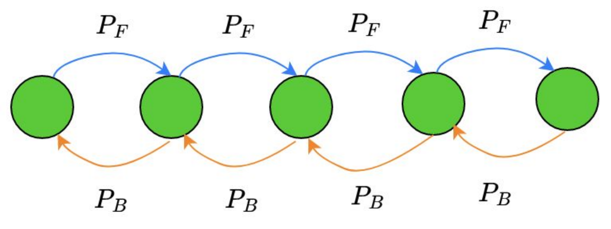

SubTrajectory Balance

$$F(s_m;\theta)\prod_{i=m}^{n-1} P_F(s_{i+1}|s_i;\theta) = F(s_n;\theta)\prod_{i=m}^{n-1} P_B(s_i|s_{i+1};\theta),$$

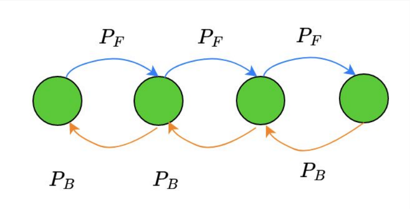

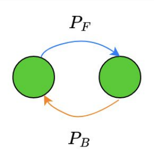

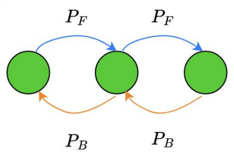

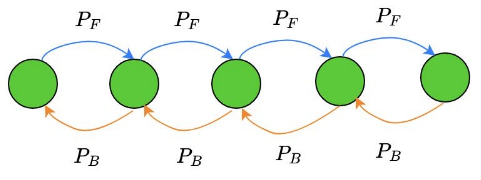

SubTrajectory Balance

$$F(s_m;\theta)\prod_{i=m}^{n-1} P_F(s_{i+1}|s_i;\theta) = F(s_n;\theta)\prod_{i=m}^{n-1} P_B(s_i|s_{i+1};\theta),$$

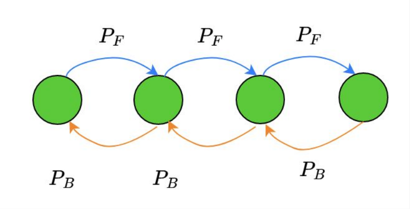

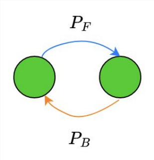

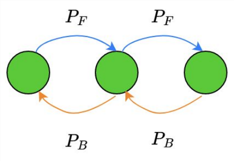

$$\mathcal{L}_{\text{SubTB}}(\tau_{m:n}) = \left(\log \frac{F(s_m;\theta)\prod_{i=m}^{n-1} P_F(s_{i+1}|s_i;\theta)}{F(s_n;\theta)\prod_{i=m}^{n-1} P_B(s_i|s_{i+1};\theta)}\right)^2$$

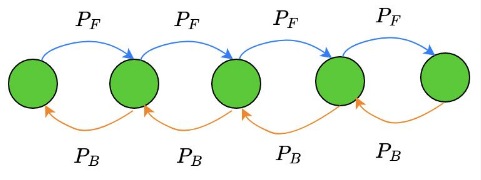

SubTrajectory(λ) or SubTB(λ):

GFlowNet Objectives Unified & Extended

SubTrajectory (λ) or SubTB (λ)

$$\mathcal{L}_{\mathrm{SubTB}(\lambda)}(\tau) = \frac{\sum_{0\le m<n\le N} \lambda^{n-m} \mathcal{L}_{\mathrm{SubTB}}(\tau_{m:n})}{\sum_{0\le m<n\le N} \lambda^{n-m}}.$$

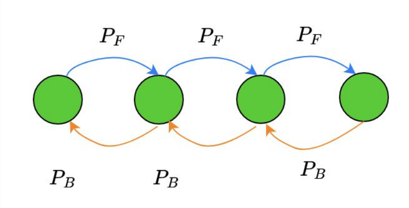

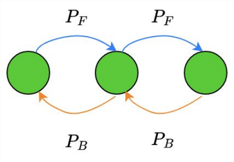

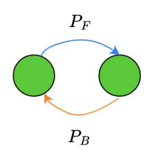

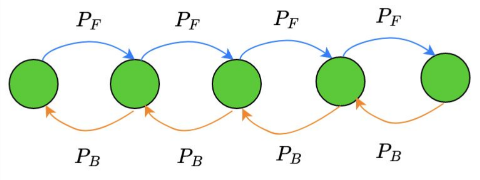

SubTrajectory (λ) or SubTB (λ)

- SubTrajectory (λ)

$$\mathcal{L}_{\mathrm{SubTB}(\lambda)}(\tau) = \frac{\sum_{0 \le m < n \le N} \lambda^{n-m} \mathcal{L}_{\mathrm{SubTB}}(\tau_{m:n})}{\sum_{0 \le m < n \le N} \lambda^{n-m}}.$$

- Unifies Detailed Balance and Trajectory Balance

SubTrajectory (λ) or SubTB (λ)

- SubTrajectory (λ)

$$\mathcal{L}_{\mathrm{SubTB}(\lambda)}(\tau) = \frac{\sum_{0 \le m < n \le N} \lambda^{n-m} \mathcal{L}_{\mathrm{SubTB}}(\tau_{m:n})}{\sum_{0 \le m < n \le N} \lambda^{n-m}}.$$

- Unifies Detailed Balance and Trajectory Balance

$$\lambda \to 0^+ : \sum_i \mathcal{L}_{\mathrm{DB}}(s_i, s_{i+1})$$

SubTrajectory (λ) or SubTB (λ)

- SubTrajectory (λ)

$$\mathcal{L}_{\mathrm{SubTB}(\lambda)}(\tau) = \frac{\sum_{0 \le m < n \le N} \lambda^{n-m} \mathcal{L}_{\mathrm{SubTB}}(\tau_{m:n})}{\sum_{0 \le m < n \le N} \lambda^{n-m}}.$$

- Unifies Detailed Balance and Trajectory Balance

$$\lambda \to 0^+ : \sum_i \mathcal{L}_{\mathrm{DB}}(s_i, s_{i+1})$$

$$\lambda \to +\infty : \mathcal{L}_{\mathrm{TB}}(\tau)$$

SubTrajectory (λ) or SubTB (λ)

- SubTrajectory (λ)

$$\mathcal{L}_{\text{SubTB}(\lambda)}(\tau) = \frac{\sum_{0 \leq m < n \leq N} \lambda^{n-m} \mathcal{L}_{\text{SubTB}}(\tau_{m:n})}{\sum_{0 \leq m < n \leq N} \lambda^{n-m}}.$$

- Unifies Previous Objectives
  - Detailed Balance: $\lambda \to 0^+$
  - Trajectory Balance: $\lambda \to +\infty$
- Lower gradient variance

SubTrajectory (λ) or SubTB (λ)

- SubTrajectory (λ)

$$\mathcal{L}_{\mathrm{SubTB}(\lambda)}(\tau) = \frac{\sum_{0 \le m < n \le N} \lambda^{n-m} \mathcal{L}_{\mathrm{SubTB}}(\tau_{m:n})}{\sum_{0 \le m < n \le N} \lambda^{n-m}}.$$

- Unifies Previous Objectives
  - Detailed Balance: $\lambda \to 0^+$
  - Trajectory Balance: $\lambda \to +\infty$
- Lower gradient variance
- Better stability and Faster convergence

SubTrajectory (λ) or SubTB (λ)

- SubTrajectory (λ)

$$\mathcal{L}_{\text{SubTB}(\lambda)}(\tau) = \frac{\sum_{0 \le m < n \le N} \lambda^{n-m} \mathcal{L}_{\text{SubTB}}(\tau_{m:n})}{\sum_{0 \le m < n \le N} \lambda^{n-m}}.$$

- Unifies Previous Objectives
  - Detailed Balance: $\lambda \to 0^+$
  - Trajectory Balance: $\lambda \to +\infty$
- Lower gradient variance
- Better stability and Faster convergence
- Wider set of applications

SubTB(λ): Experiments & Results

Experiments: SubTB (λ)

- 6 domains:

1. Hypergrid: Multi-dimensional grid
2. Small Molecule Synthesis: sequential generation of molecules from fixed graphs
3. Bit Sequence Generation: sequences of bits with fixed length
4. AMP: Antimicrobial Peptide sequence generation
5. GFP: Fluorescent Protein Generation - long sequences
6. Inverse protein folding: Non-autoregressive sequence generation

Experiments: Hypergrid

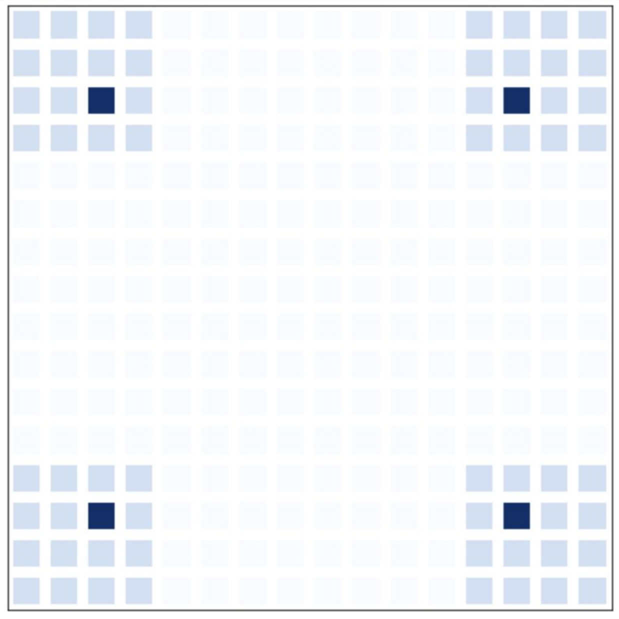

Experiments: Hypergrid

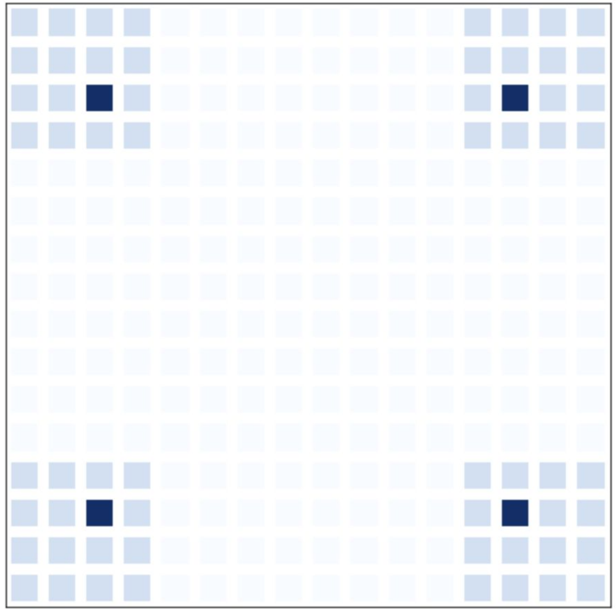

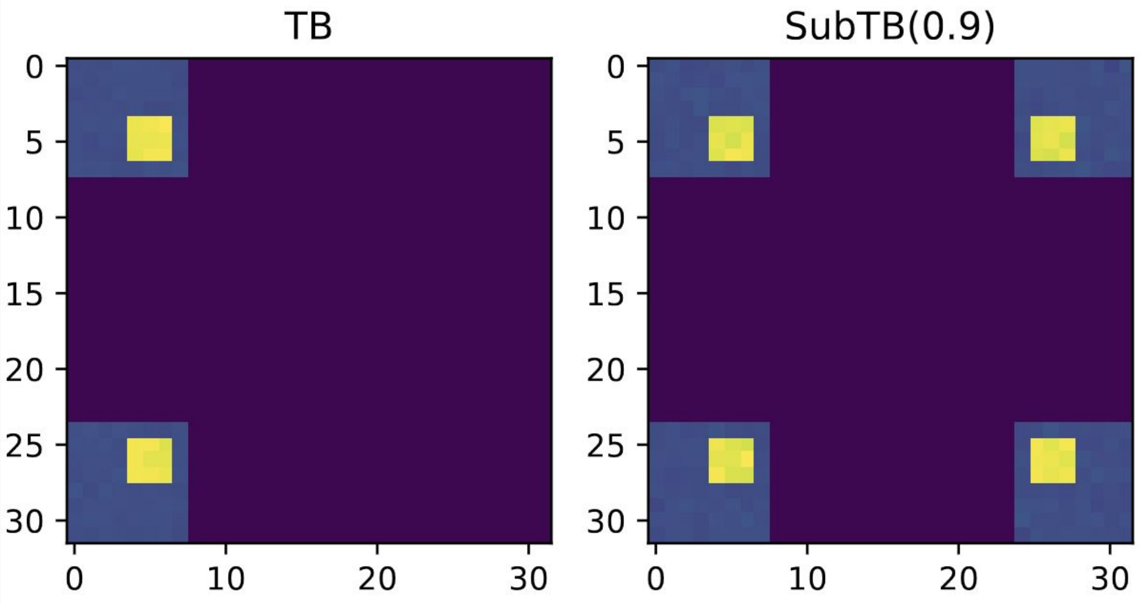

Experiments: Hypergrid

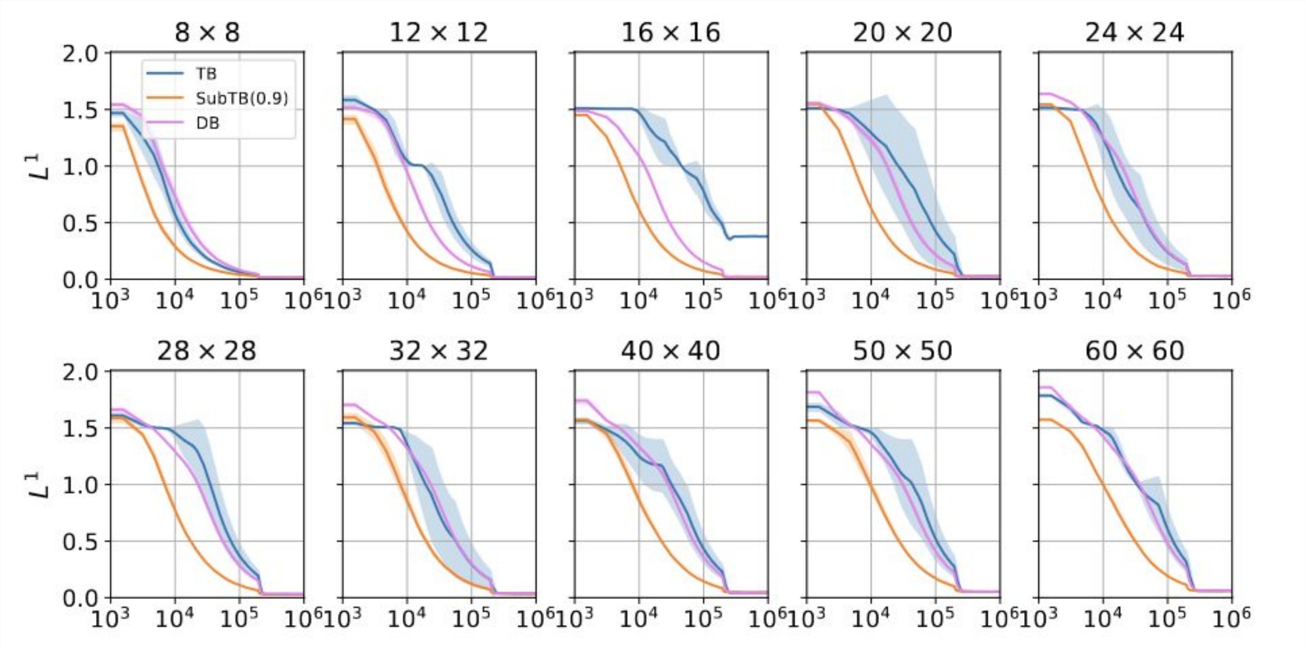

Experiments: Hypergrid

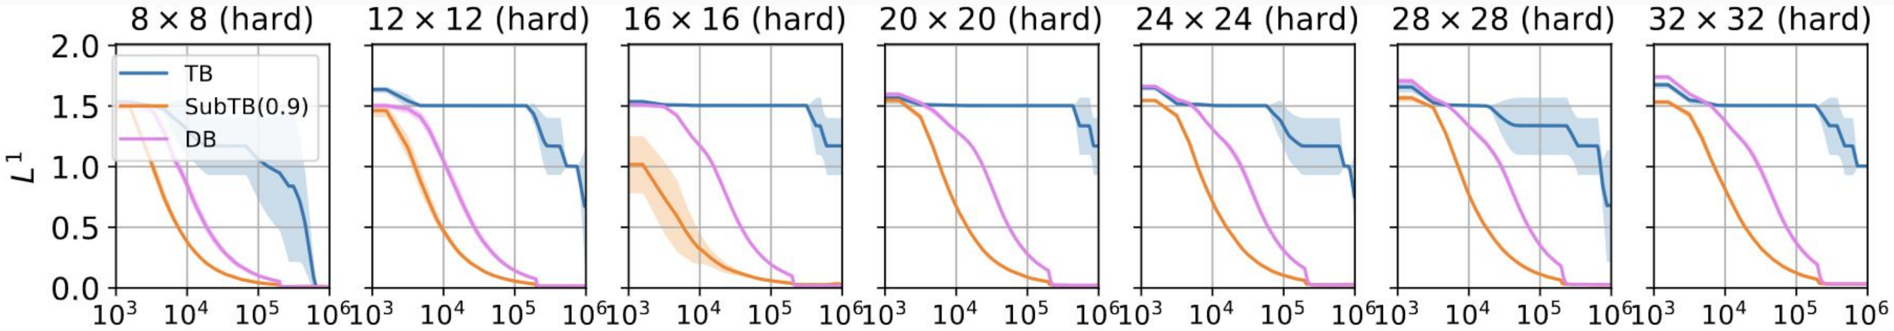

Experiments: Bit Sequence

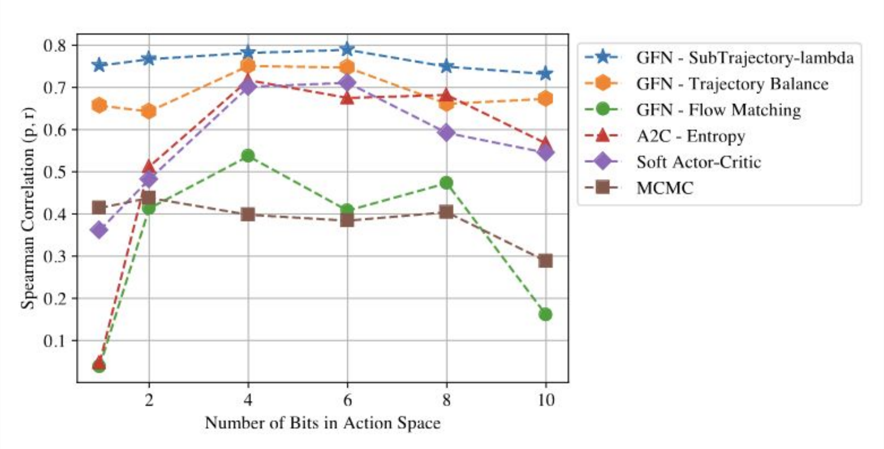

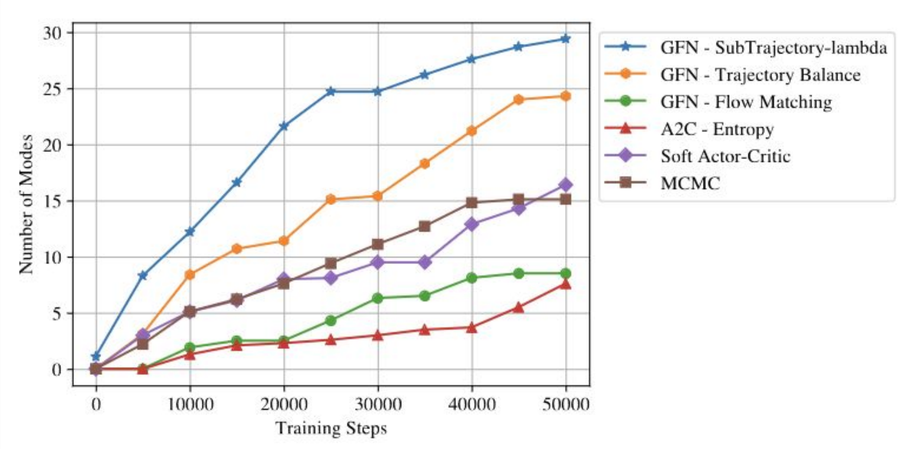

Experiments: Small Molecule

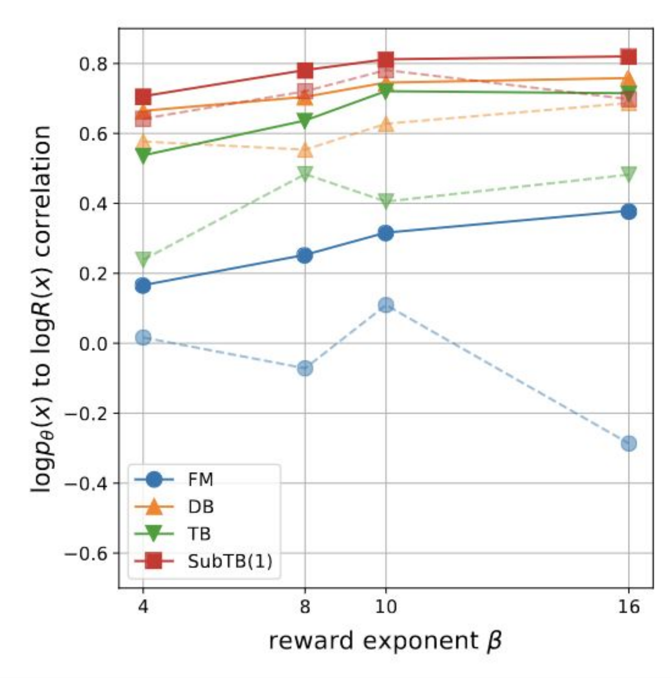

Experiments: AMP and GFP sequence

Algorithm	Top-100 Reward	Top-100 Diversity
GFN-$\mathcal{L}_{\mathrm{SubTB}(\lambda)}$	0.96 ± 0.02	42.23 ± 3.4
GFN-$\mathcal{L}_{\mathrm{TB}}$	0.90 ± 0.03	31.42 ± 2.9
GFN-$\mathcal{L}_{\mathrm{FM}}/\mathcal{L}_{\mathrm{DB}}$	0.78 ± 0.05	12.61 ± 1.32
SAC	0.80 ± 0.01	8.36 ± 1.44
AAC-ER	0.79 ± 0.02	7.32 ± 0.76
MCMC	0.75 ± 0.02	12.56 ± 1.45
GFN-$\mathcal{L}_{\mathrm{SubTB}(\lambda)}$	1.18 ± 0.10	204.44 ± 0.45
GFN-$\mathcal{L}_{\mathrm{TB}}$	0.76 ± 0.19	204.31 ± 0.44
GFN-$\mathcal{L}_{\mathrm{FM}}/\mathcal{L}_{\mathrm{DB}}$	0.30 ± 0.08	190.21 ± 6.78
SAC	0.23 ± 0.03	120.32 ± 15.57
AAC-ER	0.22 ± 0.02	113.65 ± 21.31
MCMC	0.28 ± 0.01	169.17 ± 12.44

SubTB(λ): Gradient Analysis

Gradient Analysis: SubTB(λ)

- **Small-batch** SubTB(λ) gradient is a good estimator of large-batch TB gradient.

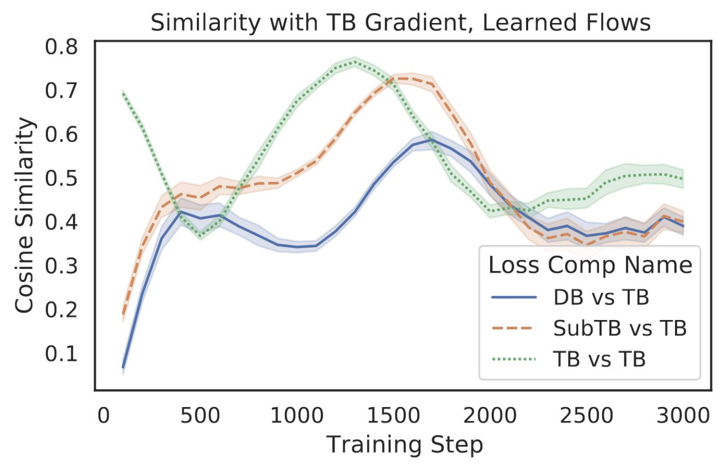

Gradient Analysis: SubTB(λ)

- **Small-batch** SubTB(λ) gradient is a good estimator of large-batch TB gradient.
- Despite its bias, the **small-batch** SubTB(λ) gradient estimates the **full-batch** TB gradient better than small-batch TB gradient.

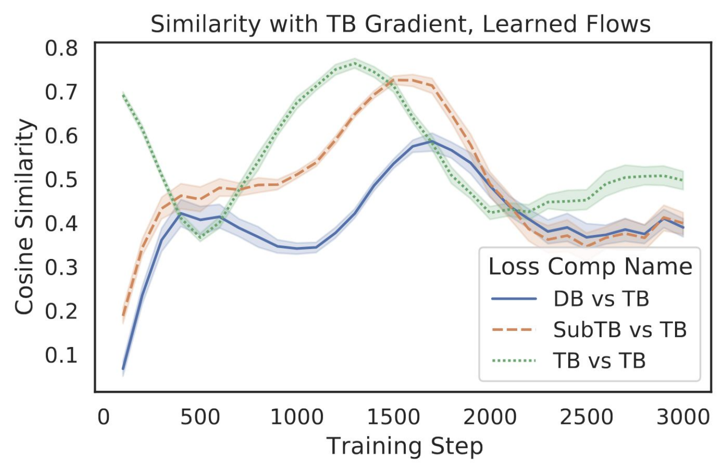

Gradient Analysis: SubTB(λ)

- **Small-batch** SubTB(λ) gradient is a good estimator of large-batch TB gradient.
- Despite its bias, the **small-batch** SubTB(λ) gradient estimates the **full-batch** TB gradient better than small-batch TB gradient.
- SubTB(λ) interpolates between the unbiased gradient estimates of TB and the biased gradient estimates of DB.

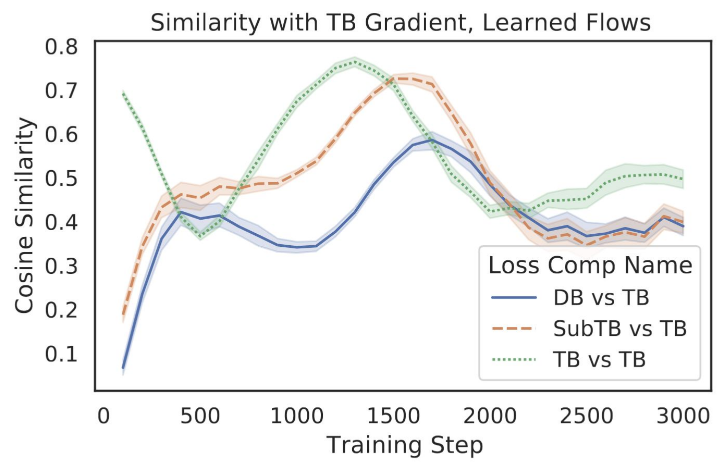

Thanks!

Poster -
26 Jul @ 2 p.m
Oral A6 Ballroom C

Paper