Statistical Mechanics of GFlowNets

Jacob Adamczyk, Rahul V. Kulkarni
Department of Physics,
UMass Boston

Generative Flow Networks

- Important problems across science share a common theme: <u>sampling</u>
  - Non-trivial when Z is intractable
- GFlowNets = new framework for sampling
- Compositional problem structure
  - Difficult sampling problem $\rightarrow$ easier pieces
- Deep learning tools

<u>Applications</u>

- Prob. inference
- Mat’l discovery
  - Batteries
  - Catalysts
- Drug discovery
- Spin Models
- $\mathcal{H}$ decomposition
- TSP
- LLMs

“GFlowNets for AI-driven scientific discovery” Jain, et al.

Introduction to GFlowNets

Stat. Mech. of MC’s

Bridge

Algorithms

Outlook

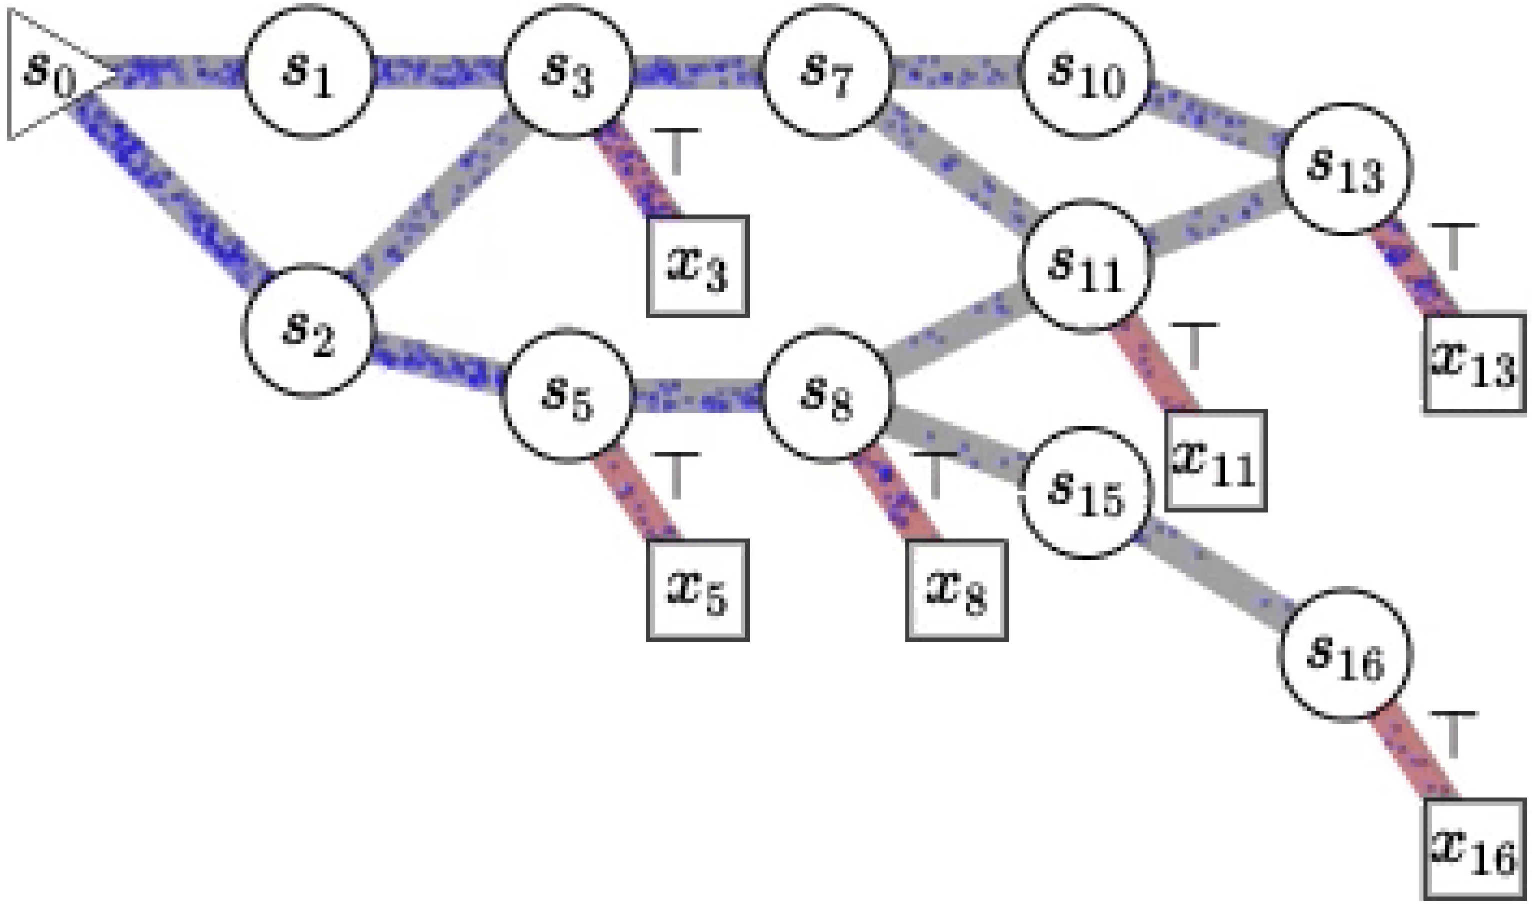

Example: “Small Molecule Generation”

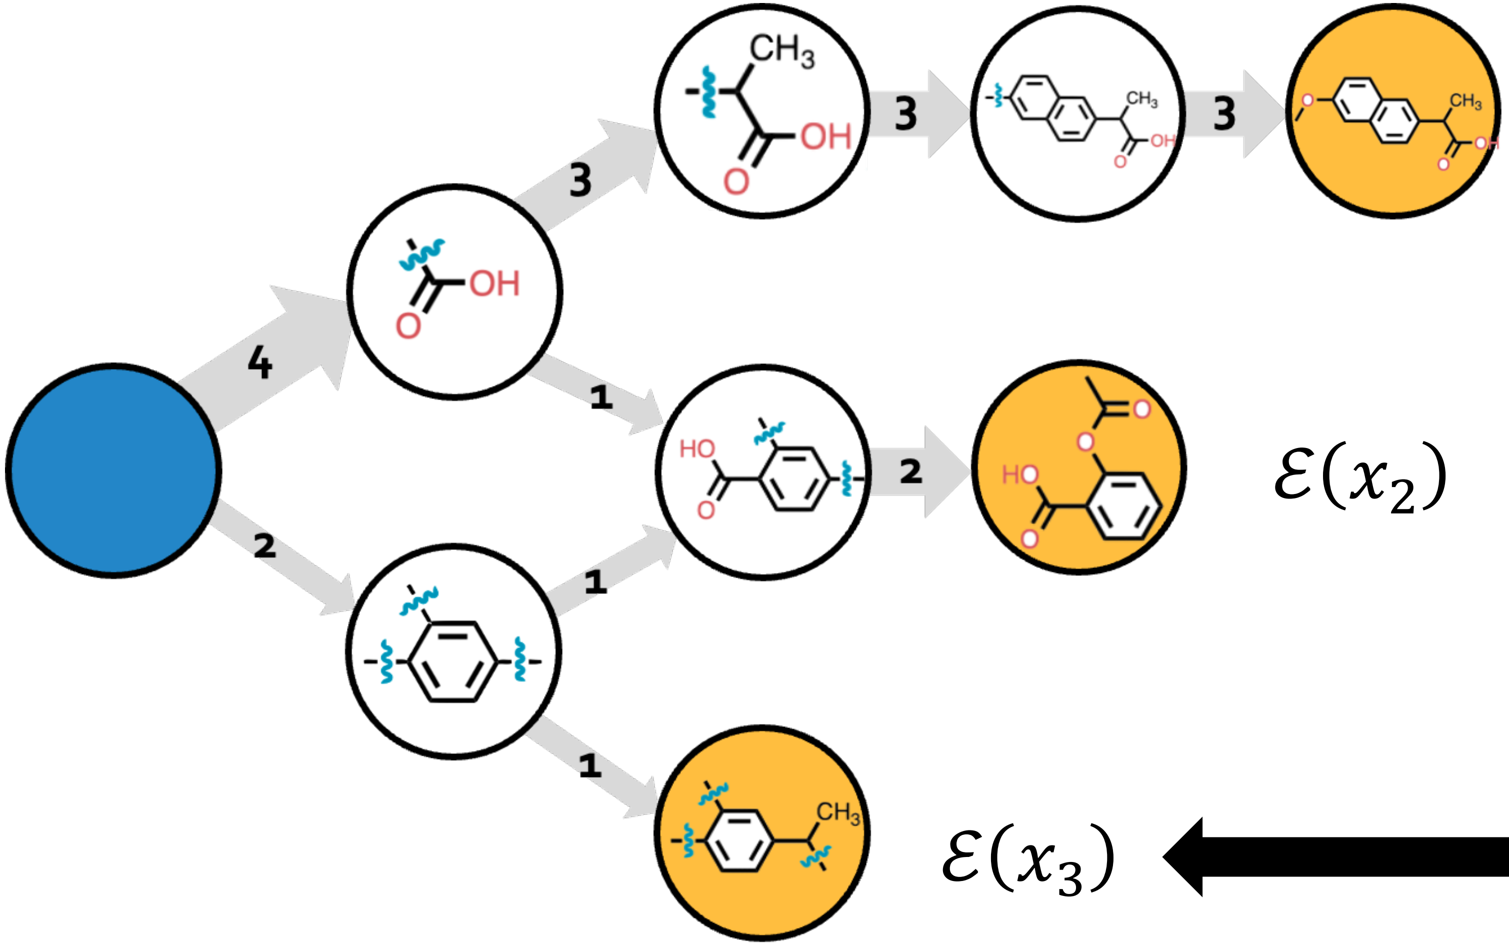

$\mathcal{E}(x_1)$

$\mathcal{E}(x_2)$

How to sample
“terminal states”
(complete products)
from Boltzmann?

$\mathcal{E}(x_3)$

Energy/Score/Reward for
each terminal state

GFlowNet Solution

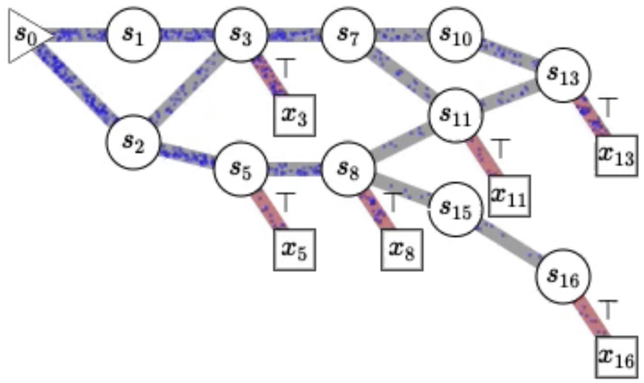

Find a flow over the DAG, $F(s \to s')$ satisfying:

$$\sum_{\bar{s}} F(\bar{s} \to s) = \sum_{s'} F(s \to s')$$

"Flow-matching condition"

B.C. at terminal states

$$F(x \to s_f) = e^{-\beta \mathcal{E}(x)}$$

Desired samples can then be constructed via:

$$P_F(s \to s') = \frac{F(s \to s')}{\sum_{s'} F(s \to s')}$$

$$\sum_{s'} F(s \to s') \doteq F(s)$$

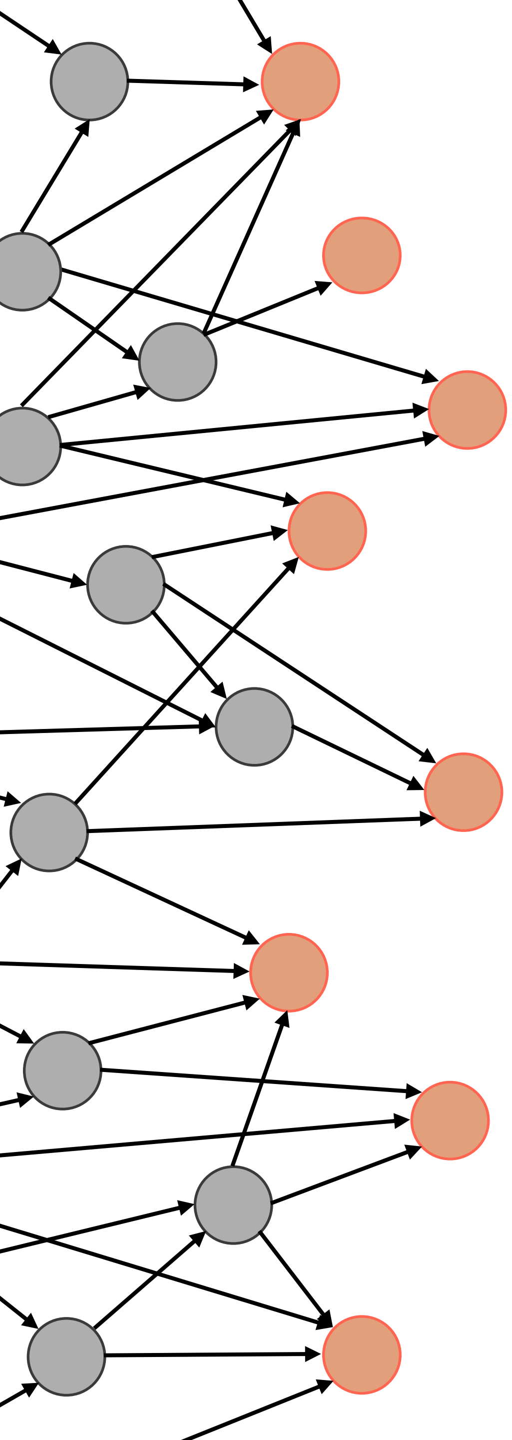

Central Problem:

$$\pi^*(x) = \frac{e^{-\beta\mathcal{E}(x)}}{Z}$$

Terminal states

Intractable!

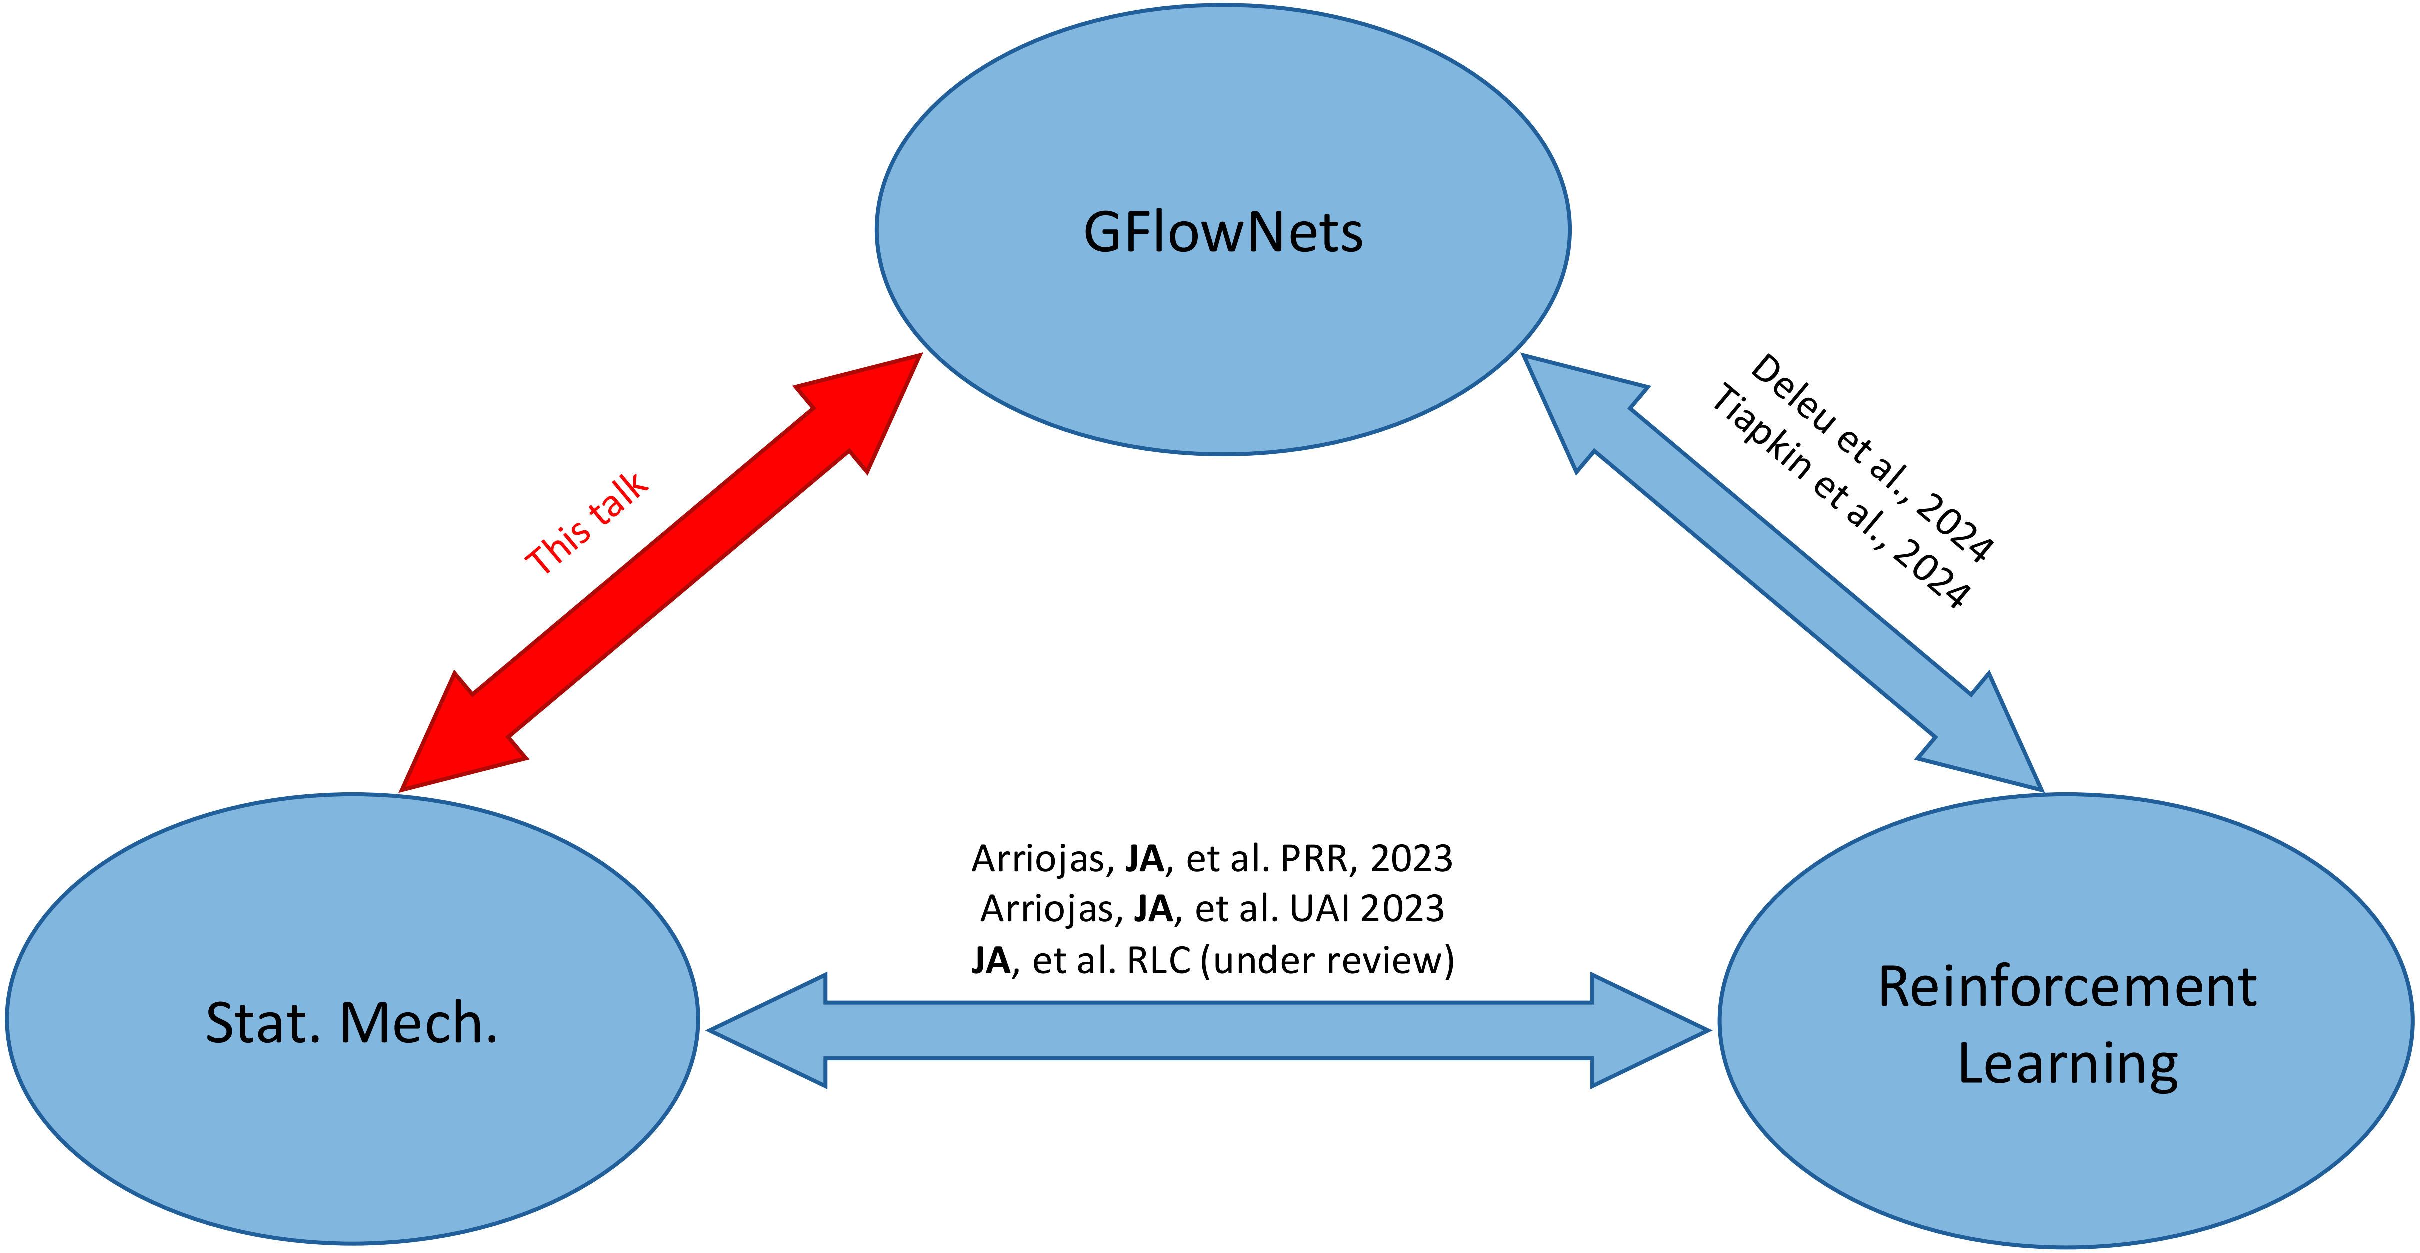
GFlowNets
This talk
Deleu et al., 2024
Tiapkin et al., 2024
Arriojas, **JA**, et al. PRR, 2023
Arriojas, **JA**, et al. UAI 2023
JA, et al. RLC (under review)
Stat. Mech.
Reinforcement Learning

Introduction to GFlowNets

Stat. Mech. of MC's

Bridge

Algorithms

Outlook

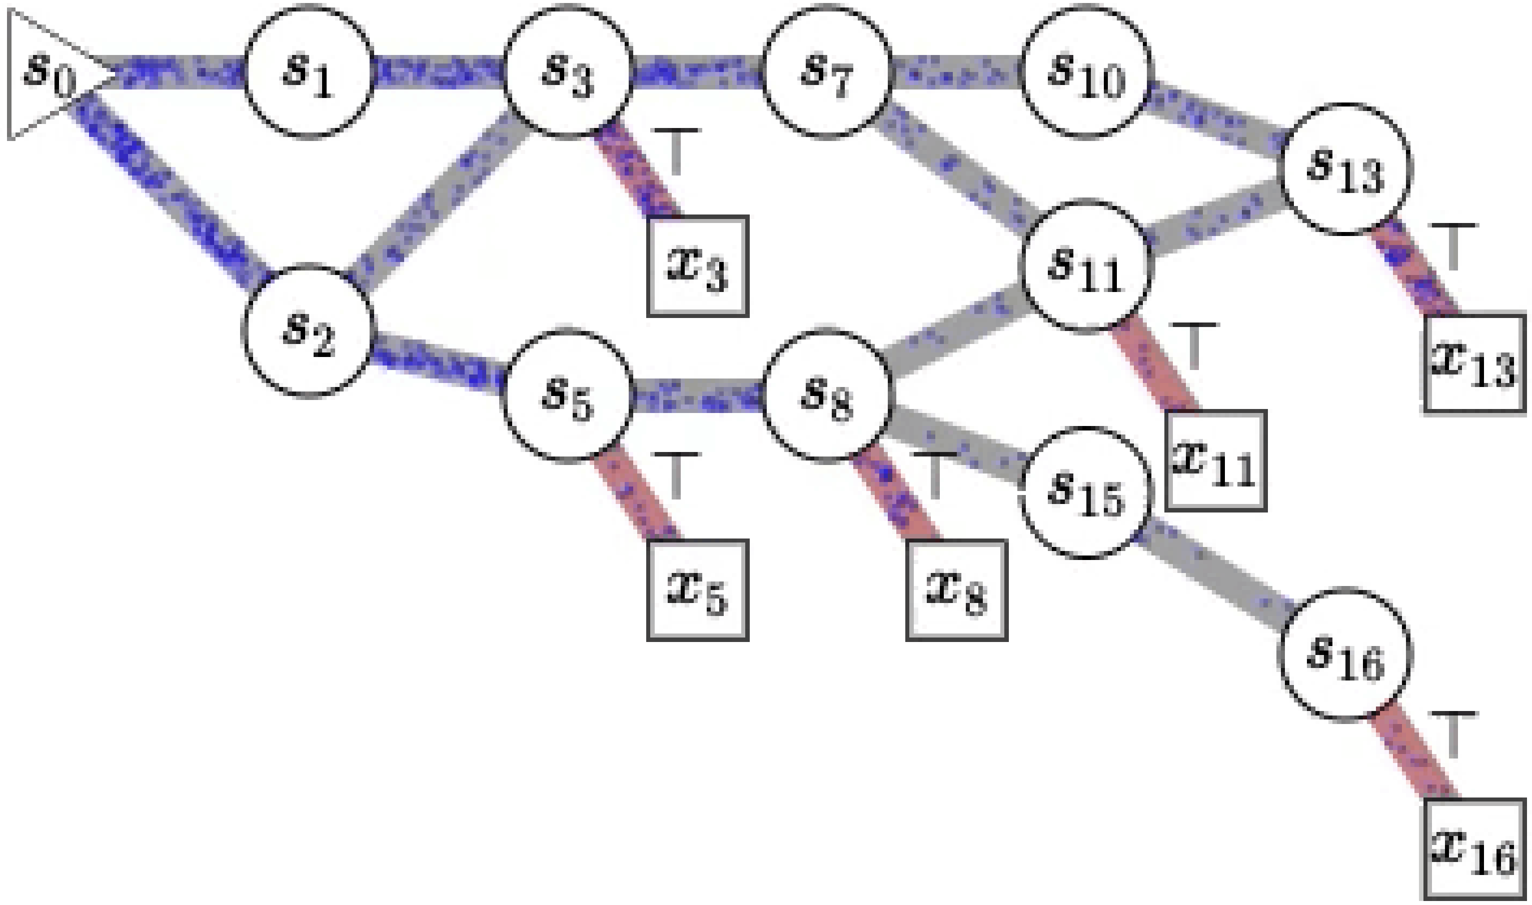

Stat. Mech. for Markov Chains

- Non-equilibrium dynamics
  - Steady-state has non-trivial currents (“flows”)
- Canonical ensemble → “s-ensemble”
  - Thermodynamics of trajectories

$$\tilde{P}(\tau) \propto e^{-\beta E(\tau)}$$

Useful Results from Stat Mech

$\tilde{P}$ generates the tilted dynamics, s-ensemble:

$$\tilde{P}(s \to s') = p_0(s \to s')e^{-\beta E(s,s')}$$

Conserves probability ❌

"Cloning" algorithm gives QSD:
$v(s), \rho$

P_F given by driven matrix (Doob's h transform):

$$P_F(s \to s') = \frac{u(s')\tilde{P}(s \to s')}{\rho u(s)}$$

Conserves probability ✅

Steady state of driven dynamics:
$u(s)v(s)$

Introduction to GFlowNets

Stat. Mech. of MC’s

Bridge

Algorithms

Outlook

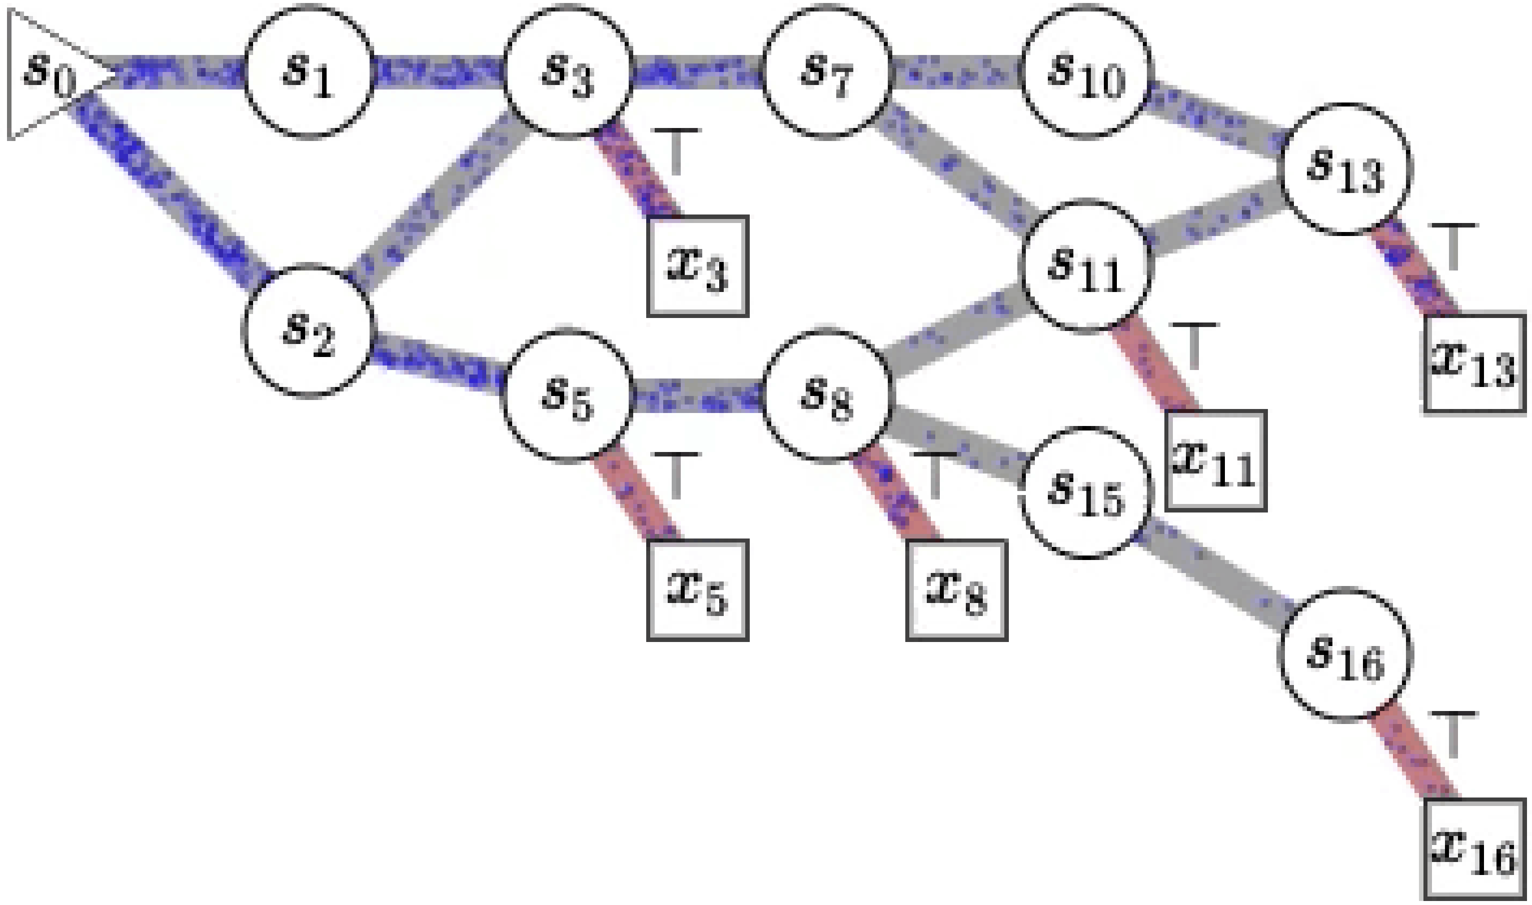

Flow-Matching Condition

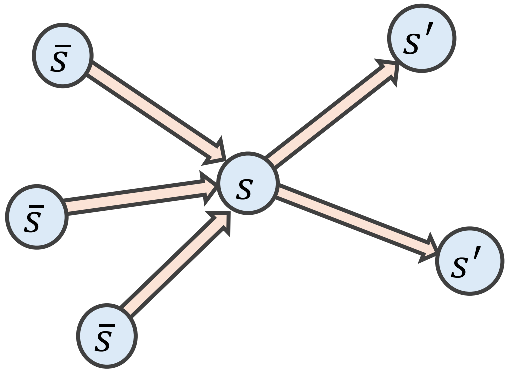

$$\sum_{\bar{s}} F(\bar{s} \to s) = \sum_{s'} F(s \to s')$$

Master Equation

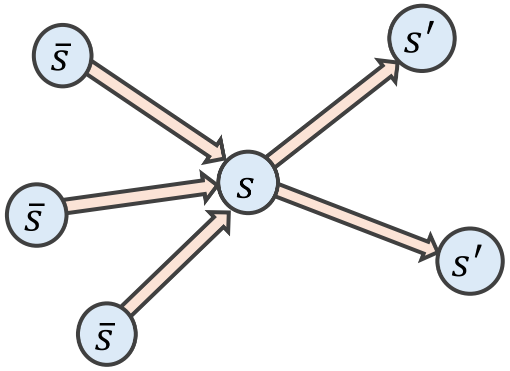

$$p_{t+1}(s) - p_t(s) = -\sum_{\bar{s}} p_t(\bar{s}) P_F(\bar{s} \to s) + \sum_{s'} p_t(s) P_F(s \to s')$$

In steady state:

$$\sum_{\bar{s}} p(\bar{s}) P_F(\bar{s} \to s) = \sum_{s'} p(s) P_F(s \to s')$$

$$p(s) = \frac{F(s)}{Z}$$

$$\sum_{\bar{s}} F(\bar{s}) P_F(\bar{s} \to s) = \sum_{s'} F(s)\, P_F(s \to s')$$

$$\sum_{\bar{s}} F(\bar{s} \to s) = \sum_{s'} F(s \to s')$$

Flow as Steady State

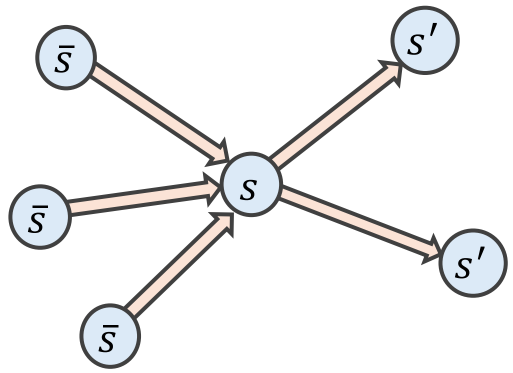

How to satisfy flow-matching?

- Steady-state of any driven dynamics → flow-matching

How to satisfy BC?

- Choose an appropriate tilted dynamics (internal energy function)

Does irreducible/steady-state match absorbing DAG setup?

- Steady-state of driven model = "One-shot" absorption probability

Flow as Steady State

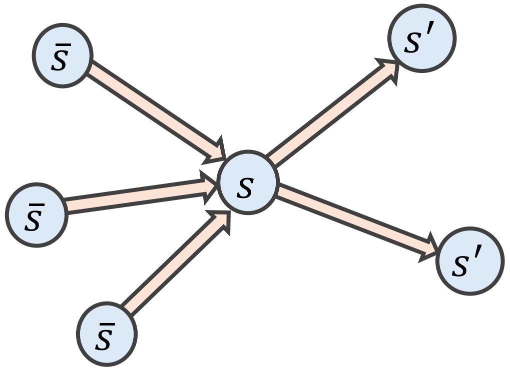

How to satisfy flow-matching?

- Steady-state of any driven dynamics → flow-matching

How to satisfy BC?

- Choose an appropriate tilted dynamics (internal energy function)

Does irreducible/steady-state match absorbing DAG setup?

- Steady-state of driven model = "One-shot" absorption probability

Steady-State Decomposition

Can recast the GFlowNet solution as finding interior energy function such that:

$$E(s, s') \to \tilde{P} \to u(x)v(x) \propto e^{-\beta \mathcal{E}(x)}$$

$$u(x) = e^{-\beta E(x, s_f)}$$

Many choices for u, v

Leads to new algorithms

Unifying framework for prior work

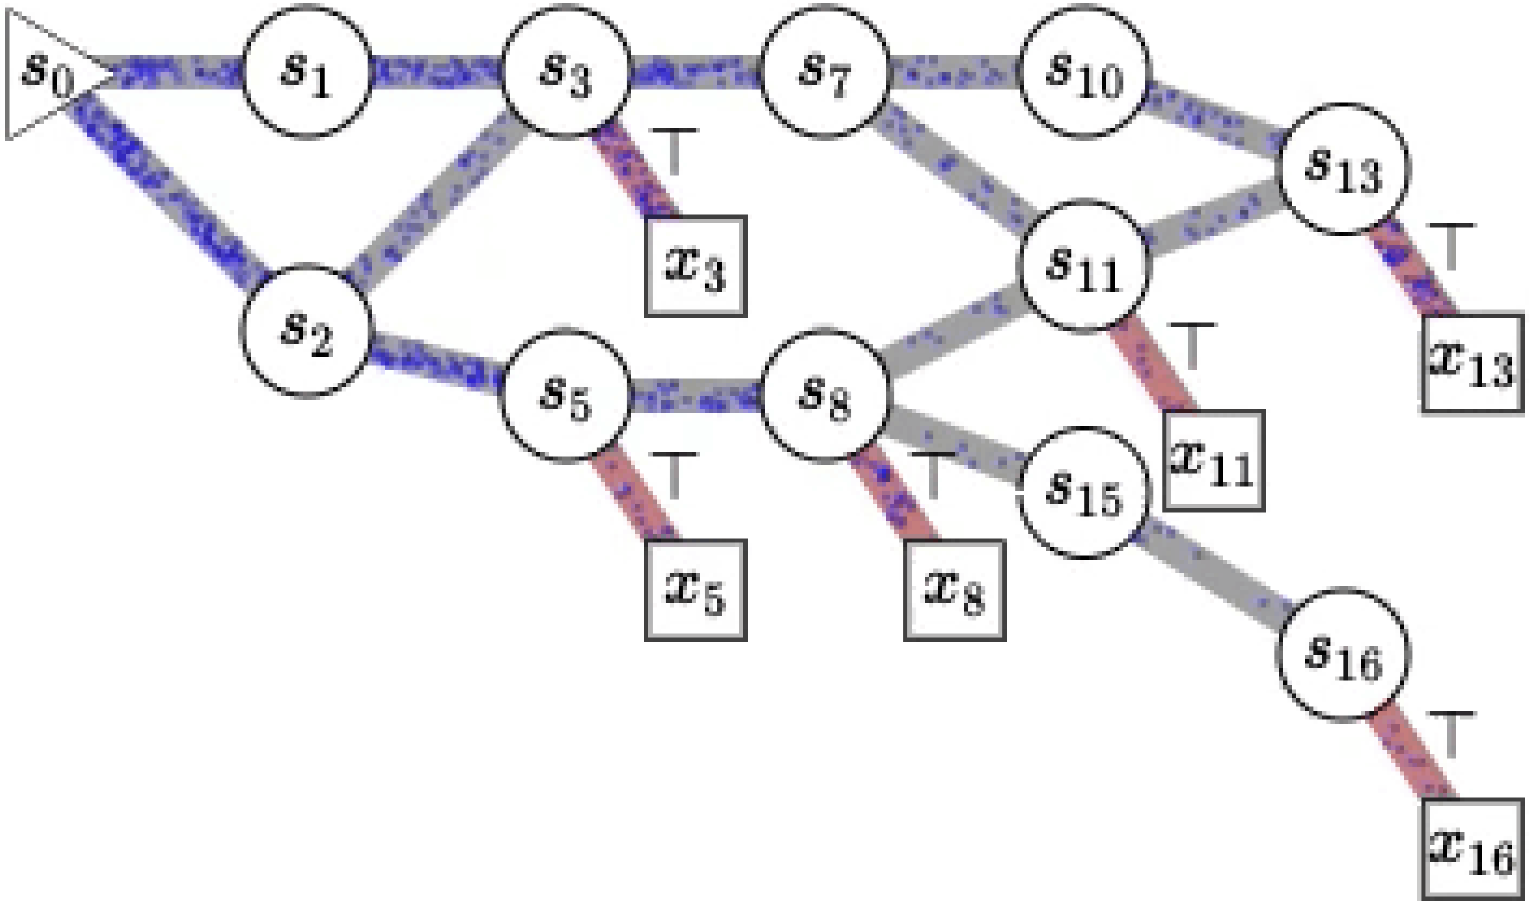
s_0
s_1
s_2
s_3
s_5
s_7
s_8
s_{10}
s_{11}
s_{13}
s_{15}
s_{16}
x_3
x_5
x_8
x_{11}
x_{13}
x_{16}
$\top$

Prior Work

$$p_{ss}(x) = u(x)v(x) \propto e^{-\beta\mathcal{E}(x)}$$

Tiapkin et al., AISTATS-24

$$u(x) = e^{-\beta\mathcal{E}(x)}$$

$$v(x) \propto 1$$

Energy $\rightarrow$ row-stochastic

RL problem

Mohammadpour et al., AISTATS-24

$$u(x) \propto e^{-\beta(\mathcal{E}(x)-\ell(x))}$$

$$v(x) = e^{-\beta\ell(x)}$$

Energy $= 0$

Dynamic Programming

Cloning

$$p_{ss}(x) = u(x)v(x) \propto e^{-\beta \mathcal{E}(x)}$$

Fix <u>any</u> interior energy function, solve for $v(x)$ through cloning:

Set u via energies

$$\frac{u(x_1)v(x_1)}{u(x_2)v(x_2)} \doteq e^{-\beta \Delta \mathcal{E}}$$

BC constraint

<u>Cloning from NESM</u> = <u>New GFlowNet algorithm!</u>

New Algorithms for GFlowNets

- Generalization to arbitrary prior dynamics
  - MaxEnt (all prior work) → EntReg
- Enables Posterior Policy Iteration (PPI)
  - Can solve ground state problems
- Energy function can be “shaped”
  - Extra degrees of freedom, PBRS (Ng et al ‘99, **JA** et al ‘23)
  - Dynamic shaping improves training (“BSRS” **JA**, et al. AAAI-25)

Introduction to GFlowNets

Stat. Mech. of MC’s

Bridge

Algorithms

Outlook

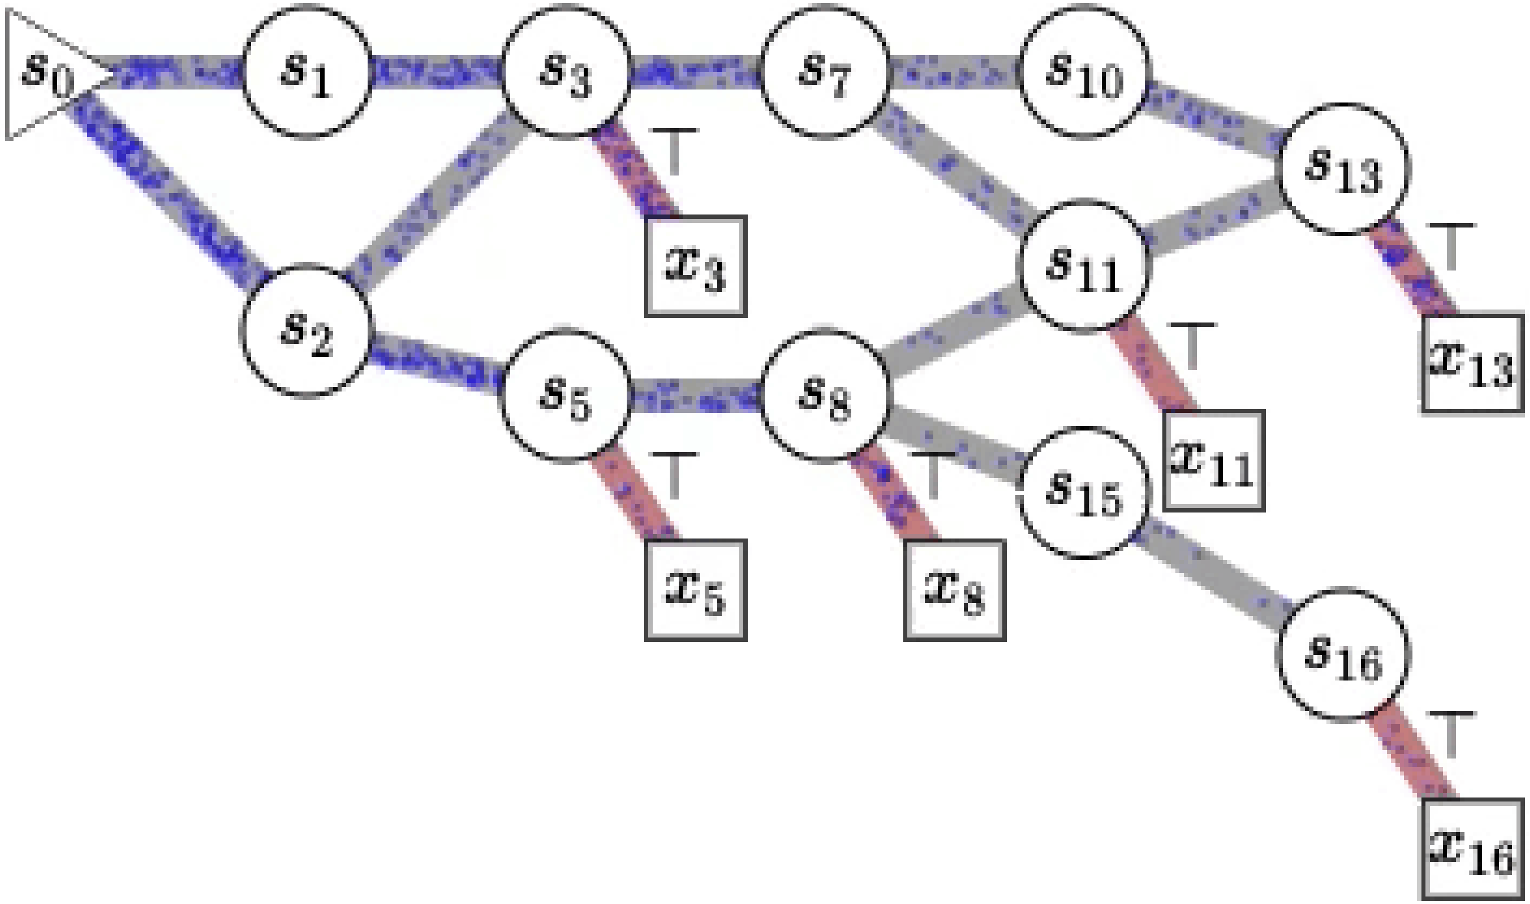

Bountiful Bridge

Stat Mech

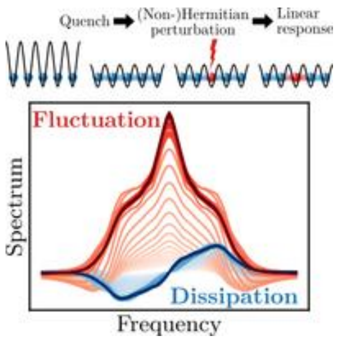

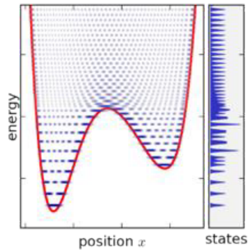

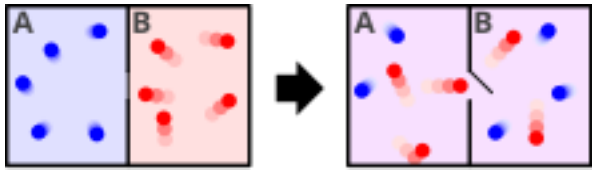

GFlowNets

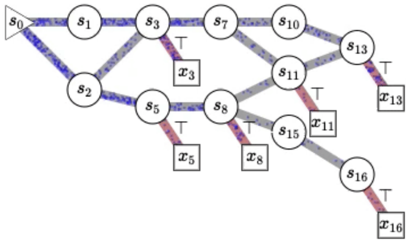

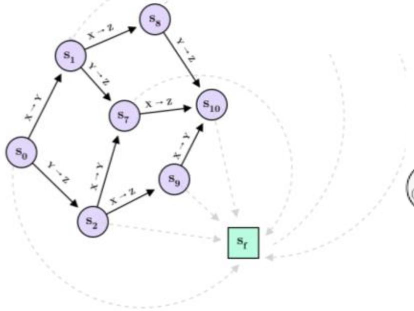

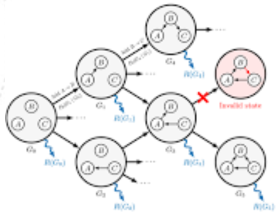

Outlook

Conclusions:

- Deeper understanding of GFlowNet and training objectives
- Unlocked internal energy functions
- New algorithms

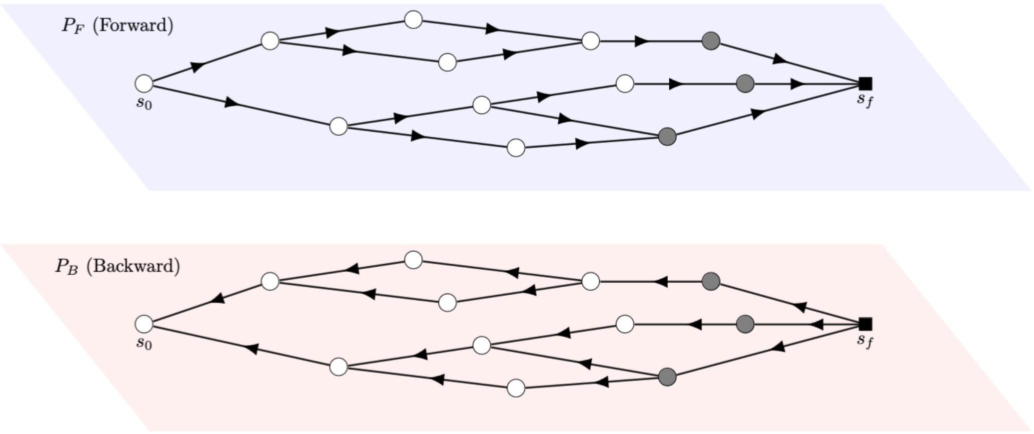

Future Work:

- Quantum
- Gauge Invariance on DAGs
- Non-equilibrium versions

APS Topical Group on Statistical & Nonlinear Physics

GSNP

Thank You!