

Generative flow models learn a (possibly stochastic) mapping between source and target distributions. Common paradigms include diffusion models, score matching models, and continuous normalizing flows.

In this talk I will first present methods for improved training of flow models using flow matching objectives using ideas from optimal transport. I will then show how these improved methods can be applied to the tasks of (1) generative protein design, with applications to biologic drug discovery. (2) amortized sampling from an unnormalized probability distribution without data.