Data-to-Energy Stochastic Dynamics

Kirill Tamogashev
University of Edinburgh

Nikolay Malkin
University of Edinburgh, CIFAR Fellow

Abstract

The Schrödinger bridge problem is concerned with finding the optimal transportation dynamics between two distributions. Existing algorithms allow to infer such dynamics only for cases where samples from both distributions are available. We propose a novel method that allows to model Schrödinger bridges when one (or both) distributions are given by their unnormalised densities. We apply the newly developed method to the problem of sampling posterior distributions in latent spaces of generative models, creating a scalable data-free image-to-image translation method.

1 Introduction

At the heart of generative modelling lies the problem of bridging two distributions by a stochastic or deterministic dynamical system that transforms one distribution into the other. In recent years there has been a surge in development of dynamics-based generative models, including diffusion models Sohl-Dickstein et al. [2015], Ho et al. [2020], Song et al. [2021] and flow matching models Liu et al. [2023], Albergo et al. [2023], Lipman et al. [2023]. However, the former approach typically deals with the case where one of the distributions is Gaussian, whereas the latter approach produces deterministic dynamics. The generalised problem of inferring stochastic dynamics between two arbitrary distributions is called the Schrödinger bridge problem, which was initially proposed in Schrödinger [1931, 1932] and recently applied to machine learning problems in De Bortoli et al. [2021], Vargas et al. [2021], Chen et al. [2021a], Shi et al. [2023].

Let p_0 and p_1 be two given marginal distributions over the space $\mathbb{R}^d$, and let $\mathbb{Q}_t$ be a reference process on the time interval $[0, 1]$ taking values in $\mathbb{R}^d$ (usually an Ornstein-Uhlenbeck process, such as the Wiener process). Given this notation, the Schrödinger bridge problem can be formalised in the following way:

$$\mathbb{P}_t^* = \arg\min_{\mathbb{P}_t} \left\{ \mathrm{KL}\left(\mathbb{P}_t \mid \mathbb{Q}_t\right) \text{ s.t. } (\pi_0)_\# \mathbb{P}_t = p_0, (\pi_1)_\# \mathbb{P}_t = p_1 \right\}, \tag{1}$$

where the minimisation is over all processes $\mathbb{P}_t$ whose marginals $(\pi_0)_\#\mathbb{P}_t$ and $(\pi_1)_\#\mathbb{P}_t$ at times $t = 0$ and $t = 1$ equal p_0 and p_1, respectively. The solution $\mathbb{P}_t^*$ is a bridge – that is, a stochastic dynamical system that transports p_0 to p_1 – that is the closest to $\mathbb{Q}_t$ in KL divergence, which amounts to minimising a certain transport cost.

Computationally, the Schrödinger bridge problem can be solved using the Iterative Proportional Fitting (IPF) algorithm Sinkhorn [1964], Vargas et al. [2021], De Bortoli et al. [2021]. IPF defines a recursion initialised at $\overleftarrow{\mathbb{P}}_t^0 = \mathbb{Q}_t$:

$$\overrightarrow{\mathbb{P}}_t^{n+1} = \arg\min_{\mathbb{P}_t} \left\{ \mathrm{KL}\left(\mathbb{P}_t \mid \overleftarrow{\mathbb{P}}_t^n\right) \text{ s.t. } (\pi_1)_\#\mathbb{P}_t = p_1 \right\} \tag{2a}$$

$$\overleftarrow{\mathbb{P}}_t^{n+1} = \arg\min_{\mathbb{P}_t} \left\{ \mathrm{KL}\left(\mathbb{P}_t \mid \overrightarrow{\mathbb{P}}_t^{n+1}\right) \text{ s.t. } (\pi_0)_\#\mathbb{P}_t = p_0 \right\}. \tag{2b}$$

Under mild conditions, these processes can be represented as Itô SDEs in forward and reverse time. It can be shown that upon convergence $\overrightarrow{\mathbb{P}}_t$ and $\overleftarrow{\mathbb{P}}_t$ become time reversals of each other and solve the prob-

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: 'Frontiers in Probabilistic Inference: Sampling Meets Learning'.

lem in (1). Until now, the ways to turn (2) into learnable objectives had been studied only for cases where samples from both p_0 and p_1 are available Vargas et al. [2021], De Bortoli et al. [2021], Shi et al. [2023].

We propose to extend the IPF algorithm (2) to the case where one (or both) distributions are given by unnormalised densities, but samples are unavailable: $p(x) = e^{-\mathcal{E}(x)}/Z,\ Z = \int e^{-\mathcal{E}(x)}\,\mathrm{d}x$, where $\mathcal{E}$ can be queried, but Z is unknown. We call this the data-to-energy (or energy-to-energy) setting. To solve the optimisation problems in each IPF step, we make use of losses popular in the diffusion sampling literature Richter et al. [2020], Richter and Berner [2024] as well as stabilisation techniques from off-policy reinforcement learning. We apply this newly developed algorithm to the outsourced sampling problem Venkatraman et al. [2025], proposing a way to do data-free stochastic transport in latent space.

2 Methods

2.1 Data-to-data Schrödinger bridge

As background to our main contributions, we summarise the IPF algorithm in the data-to-data case as proposed by Vargas et al. [2021] and our modifications to it.

IPF in a time discretisation. We assume that the processes $\overrightarrow{\mathbb{P}}^n_t,\ \overleftarrow{\mathbb{P}}^n_t$ from (2) are given as Itô SDEs in forward and reverse time initialised at the source and target distributions, respectively:

$$\mathrm{d}X_t = \overrightarrow{F}_\theta(X_t, t)\,\mathrm{d}t + \sigma_t\,\mathrm{d}\overrightarrow{W}_t,\ X_0 \sim p_0, \qquad \mathrm{d}X_t = \overleftarrow{F}_\varphi(X_t, t)\,\mathrm{d}t + \sigma_t\,\mathrm{d}\overleftarrow{W}_t,\ X_1 \sim p_1, \tag{3}$$

where σ_t should coincide with the diffusion coefficient of the reference process. We discretise $\overrightarrow{\mathbb{P}}_t$ and $\overleftarrow{\mathbb{P}}_t$ over K steps using the Euler-Maruyama scheme with $\Delta t = \frac{1}{K}$, giving a joint distribution over discrete-time trajectories $\tau = (x_0, \ldots, x_K)$:

$$\overrightarrow{p}_\theta(\tau) = p_0(x_0)\overbrace{\prod_{k=0}^{K-1}\overrightarrow{p}_\theta(x_{k+1} \mid x_k)}^{\overrightarrow{p}_\theta(\tau|x_0)},\ \overrightarrow{p}_\theta(x_{k+1} \mid x_k) = \mathcal{N}\left(x_k + \overrightarrow{F}_\theta(x_k, k\Delta t)\Delta t, \sigma^2_{k\Delta t}\Delta t\right) \tag{4a}$$

$$\overleftarrow{p}_\varphi(\tau) = p_1(x_1)\underbrace{\prod_{k=1}^{K}\overleftarrow{p}_\varphi(x_{k-1} \mid x_k)}_{\overleftarrow{p}_\varphi(\tau|x_1)},\ \overleftarrow{p}_\varphi(x_{k-1} \mid x_k) = \mathcal{N}\left(x_k + \overleftarrow{F}_\varphi(x_k, k\Delta t)\Delta t, \sigma^2_{k\Delta t}\Delta t\right) \tag{4b}$$

For each n, we represent the processes $\overrightarrow{\mathbb{P}}^n_t,\ \overleftarrow{\mathbb{P}}^n_t$ as SDEs with parameters θ_n and φ_n. The two IPF optimisation problems (2a) and (2b) are approximately solved by maximum likelihood:

$$\theta_{n+1} = \arg\max_\theta \mathbb{E}_{x_1\sim p_1, \tau\sim\overleftarrow{p}_{\varphi_n}(\tau|x_1)} \sum_{k=0}^{K-1} \log \overrightarrow{p}_\theta(x_{k+1} \mid x_k) \tag{5a}$$

$$\varphi_{n+1} = \arg\max_\varphi \mathbb{E}_{x_0\sim p_0, \tau\sim\overrightarrow{p}_{\theta_{n+1}}(\tau|x_0)} \sum_{k=1}^{K} \log \overleftarrow{p}_\varphi(x_{k-1} \mid x_k). \tag{5b}$$

We call this iterative algorithm the *log-likelihood method*, contrasting with slightly different discretisation methods proposed in De Bortoli et al. [2021].

Learned variances. The majority of existing work Chen et al. [2021a], Vargas et al. [2021], De Bortoli et al. [2021], Shi et al. [2023] trains only the drift functions $\overrightarrow{F}_\theta, \overleftarrow{F}_\varphi$ or related objects using objectives similar to (5). In contrast to that, **we propose to train not only the drift, but also diffusion coefficients of both processes**, by replacing the variances $\sigma^2_{k\Delta t}$ in (4) by learned functions $\overrightarrow{\sigma}^2_\theta(x_k, k\Delta t)$ and $\overleftarrow{\sigma}^2_\varphi(x_k, k\Delta t)$. We expect this to correct for the effect of time discretisation error, inspired by results for diffusion samplers in Gritsaev et al. [2025].

The optimisation problems in (5) are solved with respect to the parameters of both the drift and diffusion coefficients. We compare our parametrisation to others in Table 1.

2.2 Data-to-energy Schrödinger bridge

IPF for data-to-energy setting For the case of data-to-energy Schrödinger bridge we assume that the marginal distribution at time $t = 1$ has the form $p_1(x) = e^{-\mathcal{E}_1(x)}/Z,\ Z = \int e^{-\mathcal{E}_1(x)}\mathrm{d}x$. Since true unbiased samples from p_1 are now unavailable, optimising (5a) is now impossible. Instead, for every x_0 we want to enforce the following relation: $\overrightarrow{p}_\theta(\tau \mid x_0) \propto \overleftarrow{p}_\varphi(\tau \mid x_1) p_1(x_1)$.

To enforce this proportionality, we use a variant of the VarGrad loss Richter et al. [2020] used for training diffusion samplers. We sample x_0 from some training distribution p_0^{train} and minimise the following variance over trajectories τ that begin at x_0:

$$\mathcal{L}(\theta) = \text{Var}\left(\log \frac{\overrightarrow{p}_\theta(\tau \mid x_0)}{p_1(x_1)\overleftarrow{p}_\varphi(\tau \mid x_1)}\right) = \text{Var}\left(\sum_{k=1}^{K} \log \frac{\overrightarrow{p}_\theta(x_k \mid x_{k-1})}{\overleftarrow{p}_\varphi(x_{k-1} \mid x_k)} + \mathcal{E}_1(x_1)\right), \quad \tau \sim p^{\text{train}}(\tau \mid x_0), \tag{6}$$

where $p^{\text{train}}(\tau \mid x_0)$ is some training distribution over trajectories.

Off-policy training methods. The above objective leaves room for the choice of training distributions $p_0^{\text{train}}(x_0)$ and $p^{\text{train}}(\tau \mid x_0)$. A naïve choice would take $p_0^{\text{train}}(x_0) = p_0(x)$ and $p^{\text{train}}(\tau \mid x_0) = \overrightarrow{p}_\theta(\tau \mid x_0)$ (*on-policy* training). However, to improve exploration, we use a few tricks from off-policy reinforcement learning inspired by approaches used for diffusion samplers Sendera et al. [2024]:

- We keep a replay buffer Schaul et al. [2016] of samples x_K from the trained forward process $\overrightarrow{p}_\theta(\tau \mid x_0)$ and use them during training by sampling x_K from buffer, then sampling in reverse $\tilde{x}_0 \sim \overleftarrow{p}_\phi(\cdot \mid x_1)$ to obtain x_0 for training.
- As trajectories τ starting at x_0, we use both this reverse trajectory and on-policy samples from $\overrightarrow{p}_\theta(\tau \mid x_0)$.
- We periodically update the buffer using a few steps of unadjusted Langevin on the density p_1.

Details are described in Appendix B, and these methods are ablated in Fig. 1.

2.3 Outsourced sampling with Schrödinger bridge

We apply the newly developed method to the problem of Bayesian posterior sampling under a pretrained generative model prior $p(x)$. The posterior has the form $p(x \mid y) \propto p(x)r(x, y)$, where $r(x, y)$ is a constraint function that encodes the conditional information about the sample x, e.g., class or text prompt. If the pretrained generative model is expressed as a deterministic function f of a random noise variable $z \sim p(z)$, Venkatraman et al. [2025] proposed to sample the posterior pulled back to the noise space, with density $p(z \mid y) \propto p(z)r(f(z), y)$, using a diffusion sampler. If z is distributed with this density, then samples $f(z)$ follow the desired posterior distribution $p(x \mid y)$ in data space.

Instead of using a diffusion sampler, we propose to model a Schrödinger bridge between distributions $p(z)$ and $p(z \mid y)$. Since neither the normalizing constant nor samples from the latter distribution are available, we use the data-to-energy algorithm described in §2.2.

3 Experiments

3.1 Benefits of trainable variance

Table 1: Comparison of data-to-data IPF methods.

Distributions →	Gauss ↔ GMM		Gauss ↔ S-curve	
Algorithm ↓ Metric →	$\mathcal{W}_2^2$	Path Energy	$\mathcal{W}_2^2$	Path Energy
DSBM-IMF Shi et al. [2023]	0.038±0.008	**0.524±0.022**	0.053±0.012	**0.299±0.047**
DSBM-IMF+ Shi et al. [2023]	0.033±0.005	0.623±0.162	0.057±0.016	0.316±0.051
SF^{2}M Tong et al. [2024]	0.044±0.025	0.827±0.388	0.052±0.037	0.706±0.082
IPF-based				
DSB mean De Bortoli et al. [2021]	0.037±0.006	0.685±0.039	0.135±0.064	0.638±0.013
DSB score De Bortoli et al. [2021]	0.040±0.022	0.685±0.070	0.351±0.240	**0.279±0.356**
SDE Chen et al. [2021a]	0.019±0.005	0.718±0.176	0.023±0.004	0.466±0.018
LL fixed var. ≈ Vargas et al. [2021]	0.017±0.000	0.765±0.033	0.022±0.004	0.593±0.127
LL trainable var. (ours)	**0.015±0.004**	**0.551±0.097**	**0.012±0.004**	0.406±0.418

In Table 1 we present a comparisons of existing SB methods with our proposed parameterisation with both learned and fixed variance. To do that we use three 2-dimensional distributions: Gauss ↔ GMM, Gauss ↔ S-curve. We note that the same distributions were used in Shi et al. [2023]. Our method is compared to [Shi et al., 2023, DSBM and DSBM++], [De Bortoli et al., 2021, DSB] that also use IPS for training, Chen et al. [2021a] that uses continuous version of IPF, and [Tong et al., 2024, SF^{2}M].

We use $\mathrm{d}X_t = \sqrt{2}\,\mathrm{d}W_t$ as a reference process, training is done using 4000 steps for both backward and forward processes and 20 IPF steps. SF^{2}M is trained using 160'000 training steps. Path energy is computed as KL divergence between reference process and the learned process and $\mathcal{W}_2$ distance is computed between samples from $p_1(x)$ and samples obtained from $\overrightarrow{p}_\theta(\tau)$.

3.2 Data-to-Energy Schrödinger bridge

To prove viability of data-to-energy Schrödinger bridge training we compare the bridge between Gaussian and GMM distributions trained using data-to-data and data-to-energy IPF versions. Fig. 1 shows the resulting IPF trajectories and the difference in $\mathcal{W}_2^2$ metric and Path Energy.

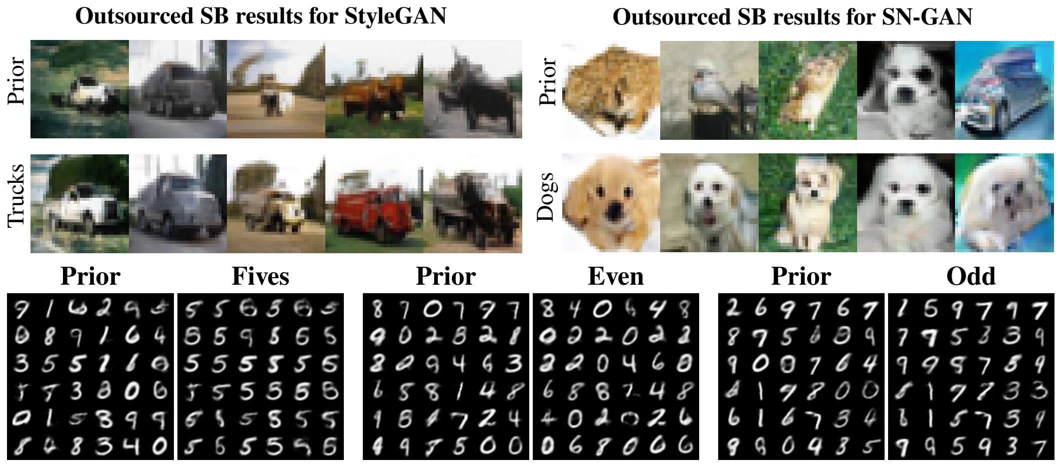

Figure 2: Outsourced Schrödinger bridge on MNIST and CIFAR-10.

3.3 Outsourced Schrödinger bridge

Finally, we show the scalability of our method by running data-to-energy Schrödinger bridge algorithm in the latent space of generative model. We use generators of StyleGAN Karras et al. [2020, 2021] and SN-GAN Miyato et al. [2018] generator trained on CIFAR-10 Krizhevsky [2009] and VAE Kingma and Welling [2014], Rezende et al. [2014] generator trained on MNIST LeCun et al. [1998]. We train the bridge model between the latent space prior $p(z)$, which in our case is always a Gaussian distribution, and reward-reweighted prior of the form $r(f(z), y) \cdot p(z)$. The reward function $r(x, y)$ is a classifier that returns the probability of the object x belonging to the class y (i.e., the probability that x is, for example, boat for CIFAR-10 or the digit 5 for MNIST).

MNIST experiments are conducted in two setups: (a) reward function returns the probability that x is even or odd (b) reward function returns the probability that $x = 5$. We present the results in the lower panel of Fig. 2.

For CIFAR-10 experiments we use StyleGAN Karras et al. [2020, 2021] generator with latent dimension 512 and SN-GAN Miyato et al. [2018] generator with latent dimension 128. We use a pretrained classifier model as a reward function. We chose the target class to be *cars* for StyleGAN and *dogs* for SN-GAN. The results are shown in the upper panel of Fig. 2.

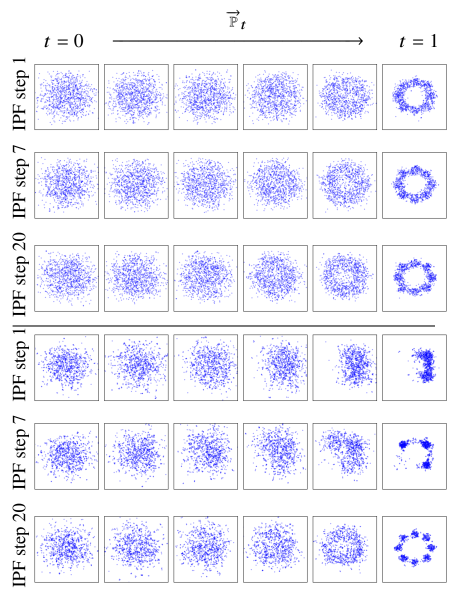

Algorithm ↓ Metric →	Gauss ↔ GMM $\mathcal{W}_2^2$
data-to-data SB	$0.015_{\pm 0.004}$
data-to-energy SB	$0.317_{\pm 0.267}$
on-policy	$0.476_{\pm 0.369}$
no Langevin	$0.533_{\pm 0.234}$

Figure 1: Comparison of learnt processes at various IPF iterations for data-to-data (*top*) or data-to-energy (*middle*) settings. The table (*bottom*) shows that our data-to-energy algorithm works well despite having no access to target data samples.

We would like to note a benefit of training a Schrödinger bridge model, as opposed to a diffusion sampler or an arbitrary stochastic mapping. It can be seen in Fig. 2 that images already belonging to the target class change little, while those belonging to other classes maintain features that are unrelated to the target class: the background and global structure are preserved. This suggests that style transfer for higher-dimensional images can be a promising application of our method.

4 Conclusion

This paper shows the potential of training data-to-energy Schrödinger bridges in a time discretisation with learnable drift and variance for the forward and backward processes. We showed that, despite its complexity, our method can be successfully scaled. Future work should explore applications to high-dimensional data, generalisation to discrete data, style transfer applications, as well as compatibility with a broad range of pretrained generative models.

References

Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

Arip Asadulaev, Rostislav Korst, Vitalii Shutov, Alexander Korotin, Yaroslav Grebnyak, Vahe Egiazarian, and Evgeny Burnaev. Optimal transport maps are good voice converters. *arXiv preprint arXiv:2411.02402*, 2024.

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024.

Julius Berner, Lorenz Richter, Marcin Sendera, Jarrid Rector-Brooks, and Nikolay Malkin. From discrete-time policies to continuous-time diffusion samplers: Asymptotic equivalences and faster training. *arXiv preprint arXiv:2501.06148*, 2025.

Denis Blessing, Julius Berner, Lorenz Richter, and Gerhard Neumann. Underdamped diffusion bridges with applications to sampling. *International Conference on Learning Representations (ICLR)*, 2025.

Tianrong Chen, Guan-Horng Liu, and Evangelos Theodorou. Likelihood training of Schrödinger bridge using forward-backward SDEs theory. *International Conference on Learning Representations (ICLR)*, 2021a.

Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal transport in systems and control. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):89–113, 2021b.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

W Edwards Deming and Frederick F Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11(4): 427–444, 1940.

Milena Gazdieva, Petr Mokrov, Litu Rout, Alexander Korotin, Andrey Kravchenko, Alexander Filippov, and Evgeny Burnaev. An optimal transport perspective on unpaired image super-resolution. *Journal of Optimization Theory and Applications*, 207(2):40, 2025.

Timofei Gritsaev, Nikita Morozov, Kirill Tamogashev, Daniil Tiapkin, Sergey Samsonov, Alexey Naumov, Dmitry Vetrov, and Nikolay Malkin. Adaptive destruction processes for diffusion samplers. *arXiv preprint arXiv:2506.01541*, 2025.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Leonid V. Kantorovich. Mathematical methods of organizing and planning production. *Management science*, 6(4):366–422, 1960.

Leonid V. Kantorovich and Gennady S. Rubinshtein. On a space of totally additive functions. *Vestnik of the St. Petersburg University: Mathematics*, 13(7):52–59, 1958.

Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. *Computer Vision and Pattern Recognition (CVPR)*, 2020.

Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representations (ICLR)*, 2014.

Alexander Korotin, Daniil Selikhanovych, and Evgeny Burnaev. Neural optimal transport. *International Conference on Learning Representations (ICLR)*, 2022.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical Report TR-2009, University of Toronto, 2009.

Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghuai Zhang, Alexandra Volokhova, Alex Hernández-García, Léna Néhale Ezzine, Yoshua Bengio, and Nikolay Malkin. A theory of continuous generative flow networks. *International Conference on Machine Learning (ICML)*, 2023.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *International Conference on Learning Representations (ICLR)*, 2023.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *International Conference on Learning Representations (ICLR)*, 2023.

Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems*, 34(4):1533–1574, 2014.

Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *International Conference on Learning Representations (ICLR)*, 2018.

Gaspard Monge. *Mémoire sur la Théorie des Déblais et des Remblais*. Imprimerie royale, 1781.

Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *International Conference on Machine Learning (ICML)*, 2014.

Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. *International Conference on Learning Representations (ICLR)*, 2024.

Lorenz Richter, Ayman Boustati, Nikolas Nüsken, Francisco Ruiz, and Omer Deniz Akyildiz. VarGrad: a low-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *International Conference on Learning Representations (ICLR)*, 2016.

Erwin Schrödinger. Über die Umkehrung der Naturgesetze. *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse*, pages 144–153, 1931.

Erwin Schrödinger. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. In *Annales de l'institut Henri Poincaré*, volume 2, pages 269–310, 1932.

Marcin Sendera, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks, Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. Improved off-policy training of diffusion samplers. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger bridge matching. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics*, 35(2):876–879, 1964.

Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *International Conference on Machine Learning (ICML)*, 2015.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations (ICLR)*, 2021.

Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. *Artificial Intelligence and Statistics (AISTATS)*, 2024.

Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving Schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, August 2021.

Francisco Vargas, Will Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. *International Conference on Learning Representations (ICLR)*, 2023.

Siddarth Venkatraman, Mohsin Hasan, Minsu Kim, Luca Scimeca, Marcin Sendera, Yoshua Bengio, Glen Berseth, and Nikolay Malkin. Outsourced diffusion sampling: Efficient posterior inference in latent spaces of generative models. *International Conference on Machine Learning (ICML)*, 2025.

Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2008.

Qinsheng Zhang and Yongxin Chen. Path Integral Sampler: A stochastic control approach for sampling. *International Conference on Learning Representations (ICLR)*, 2022.

A Related works

In this section we establish links between our method and other research direction in the literature.

Optimal transport Optimal transport is a well elaborated area of research with solid theoretical background and scalable applied algorithm. The problem is concerned with finding the optimal transportation map which minimises a given transportation cost. Originally, the problem was proposed by Monge in Monge [1781]. In 20th century this problem was reformulated and generalised by a Kantorovich in a serier of works including Kantorovich and Rubinshtein [1958], Kantorovich [1960]. Since then the problem has been rigourously studied, refer to Peyré et al. [2019], Villani [2008] for a detailed presentation of theory. In a discrete case optimal transport can be solved using Sinkhorn's Sinkhorn [1964], Peyré et al. [2019] algorithm. Recent works have applied optimal transport to a series of machine learning problems, including image-to-image translation Korotin et al. [2022], voice conversion Asadulaev et al. [2024], and super-resolution Gazdieva et al. [2025].

Schrödinger bridges Schrödinger bridge problem is concerned with finding stochastic optimal transport dynamics between two distributions. The problem was originally proposed by Schrödinger in Schrödinger [1931, 1932]. The Schrödinger bridge problem can be seen as a regularised version of dynamic optimal transport Léonard [2014] and has interesting connections to optimal control theory Chen et al. [2021b]. Computationally, the problem can be solved using Iterative Proportional Fitting (IPF) algorithm Deming and Stephan [1940], Sinkhorn [1964]. De Bortoli et al. [2021], Vargas et al. [2021] proposed the scalable formulation of this method that allows to compute Schrödinger bridge between a pair of distributions given by unbiased samples. Chen et al. [2021a] proposes a continuous-time variant of IPF. Methods distinct from IPF have also been proposed Shi et al. [2023], Tong et al. [2024]; all of them assume access to samples from the target distribution for an unbiased objective.

Diffusion samplers Data-to-energy Schrödinger bridge is related to the problem of sampling from an unnormalised density. Diffusion samplers Zhang and Chen [2022], Vargas et al. [2023], Richter and Berner [2024], Berner et al. [2024], Blessing et al. [2025] represent one of the approaches that solve this problem. Some methods use off-policy reinforcement learning techniques Lahlou et al. [2023], Sendera et al. [2024] to amortise sampling from intractable density. The theoretical connection among various objectives was established in Berner et al. [2025].

Outsourced sampling The concept of outsourced diffusion sampling – modelling continuous-time dynamics in latent space for posterior inference under pretrained priors – was proposed in Venkatraman et al. [2025]. The work shows that sampling can be efficiently conducted in the latent space of a generator, where the density landscape is smoother.

B Experiment setup

B.1 Data-to-data experiments

All data-to-data experiments are conducted under the unified setup. For neural network we use MLP with 3 hidden layers and 64 neurons in each layer, each layer is followed by a LayerNormBa et al. [2016] and SiLU activation function. All neural networks are trained using AdamW optimiser with learning rate 0.0008. Sampling is done in 20 steps with $t_{\max} = 0.2$ and $\mathrm{d}t = 0.01$. We train forward and backward model for 4000 steps at each IPF iteration and we train each model for 20 IPF iterations. $\mathrm{SF}^2\mathrm{M}$ is trained using 160'000 steps.

B.2 2D data-to-energy experiments

Data-to-energy experiments we use the same neural networks. We use 20 steps for sampling with $t_{\max} = 0.8$ and $\mathrm{d}t = 0.04$. Neural nets are optimised with AdamW optimiser with learning rate 0.0005. When Langevin update is used, we update buffer samples every 500 during training of the forward process. Langevin is used with the step size 0.01 and we do 50 updates each time. We use 2 trajectories from each x_0 for computations of VarGrad loss.

B.3 Outsourced Schrödinger bridge experiments

Experiments on MNIST For MNIST experiments we use our own VAE with 3 layers in both decoder and encoder. Moreover, we also use custom MNIST classifier. We train forward and backward networks for 5000 steps during each IPF iteration, we have 20 IPF iterations in total. All the networks are trained with AdamW optimiser using learning rate 0.0008. We do not use Langevin updates for this experiment, only buffer. We use the sample MLPs as in 2D data-to-energy experiment.

Experiments on CIFAR-10 with SN-GAN and StyleGAN . For CIFAR-10 experiments we use MLP with 3 hidden layers, it has 256 hidden units for SN-GAN experiments and 512 hidden units for StyleGAN experiments. We train forward network for 500 steps and backward network for 100 steps during each IPF iteration, we have 300 IPF iterations in total. All the networks are trained with AdamW optimiser using learning rate 0.0005. Langevin updates are made every 500 iterations during the training of the forward network. We run Langevin for 500 steps with initial step size 0.01, we anneal step size to the 0.001 during the updates.

We use 20 steps for sampling with $\mathrm{d}t = 0.04$ for StyleGAN experiments and $\mathrm{d}t = 0.005$ for SN-GAN experiments. We use Wiener process $\mathrm{d}X_t = \sqrt{2}\mathrm{d}W_t$ as a reference process.