Advanced Topics in Machine Learning (deep generative modelling)

Lecture 11: More continuous-time models

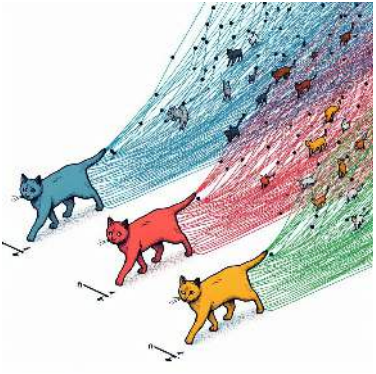

Nikolay Malkin

31 March 2026

Dynamics-based generative models in practice

'Edinburgh from Calton Hill, pointillist style'

Dynamics-based generative models in practice

‘Edinburgh from Calton Hill, pointillist style’

[Janner et al., ICML’22]

Dynamics-based generative models in practice

'Edinburgh from Calton Hill, pointillist style'

[Janner et al., ICML'22]

[Graikos et al., NeurIPS'22]

Dynamics-based generative models in practice

'Edinburgh from Calton Hill, pointillist style'

[Graikos et al., NeurIPS'22]

[Janner et al., ICML'22]

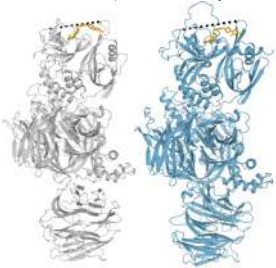

AlphaFold 3

Dynamics-based generative models in practice

[Zhang and Gienger, 2024]

Review of weeks 8-10

What is a diffusion model, how to understand it from different perspectives, and where to use it in practice:

Review of weeks 8-10

What is a diffusion model, how to understand it from different perspectives, and where to use it in practice:

- **Lecture 8:** Diffusion models are hierarchical latent variable models / deep VAEs

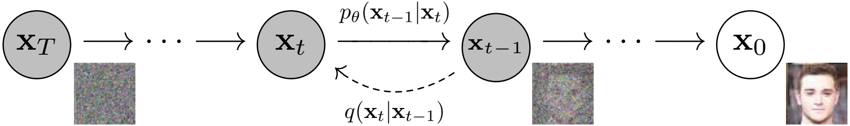

[Ho et al., 'Denoising diffusion probabilistic models']

Review of weeks 8-10

What is a diffusion model, how to understand it from different perspectives, and where to use it in practice:

- **Lecture 8:** Diffusion models are hierarchical latent variable models / deep VAEs
- **Lecture 9:** Diffusion models are score-based models

[Song et al., 'Generative modeling by estimating...' blog post]

Review of weeks 8-10

What is a diffusion model, how to understand it from different perspectives, and where to use it in practice:

- **Lecture 8:** Diffusion models are hierarchical latent variable models / deep VAEs
- **Lecture 9:** Diffusion models are score-based models
- **Lecture 10:** Diffusion models are continuous-time processes

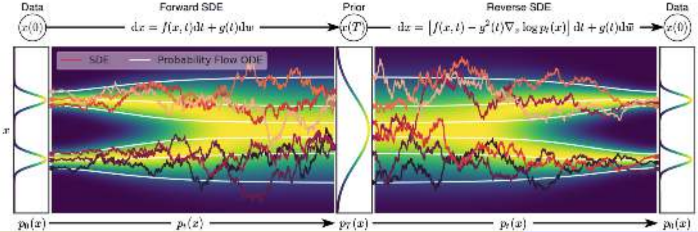

[Song et al., 'Score-based generative modeling. . .']

Outline of Lecture 11

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- Open problems in dynamics-based modelling
- Some announcements

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- Open problems in dynamics-based modelling
- Some announcements

From last time: Time reversal for Brownian motion

A simple noising SDE without drift:

$$dx_t = dw_t$$

From last time: Time reversal for Brownian motion

A simple noising SDE without drift:

$$dx_t = \color{red}{dw_t}$$

Initial conditions $p_0 \rightsquigarrow$ marginal densities p_t

- $x_t = x_0 + \mathcal{N}(0, tI_d)$ $([x_t \mid x_0] \sim \mathcal{N}(x_0, tI_d))$
- $p_t(x) = \frac{1}{\|\mathcal{D}\|} \sum_{x_0 \in \mathcal{D}} \mathcal{N}(x; x_0, tI_d)$

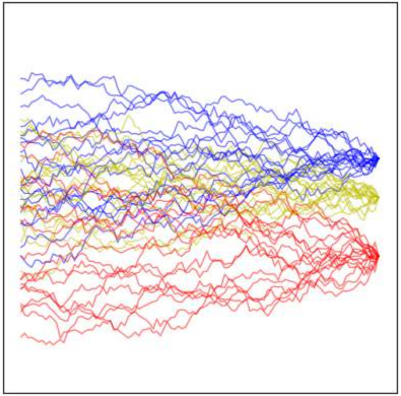

From last time: Time reversal for Brownian motion

A simple noising SDE without drift:

$$dx_t = dw_t$$

Initial conditions $p_0 \rightsquigarrow$ marginal densities p_t

- $x_t = x_0 + \mathcal{N}(0, tI_d)$ $([x_t \mid x_0] \sim \mathcal{N}(x_0, tI_d))$
- $p_t(x) = \frac{1}{\|\mathcal{D}\|} \sum_{x_0 \in \mathcal{D}} \mathcal{N}(x; x_0, tI_d)$

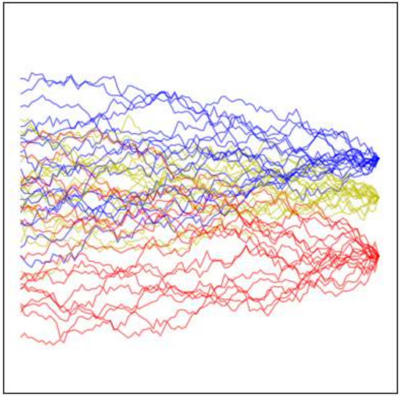

Nelson's identity: reverse SDE is

$$dx_t = -\nabla \log p_t(x_t)\, dt + \overline{dw}_t$$

From last time: Time reversal for Brownian motion

A simple noising SDE without drift:

$$dx_t = dw_t$$

Initial conditions $p_0 \rightsquigarrow$ marginal densities p_t

- $x_t = x_0 + \mathcal{N}(0, tI_d)$ $([x_t \mid x_0] \sim \mathcal{N}(x_0, tI_d))$
- $p_t(x) = \frac{1}{\|\mathcal{D}\|} \sum_{x_0 \in \mathcal{D}} \mathcal{N}(x; x_0, tI_d)$

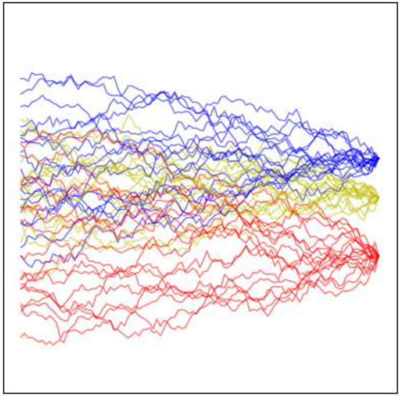

Nelson's identity: reverse SDE is

$$dx_t = -\nabla \log p_t(x_t)\, dt + \overline{dw}_t$$

- Learning to reverse noising SDE $\longleftrightarrow$ learning the score of the noised data

From last time: Time reversal for Brownian motion

SDE trajectories

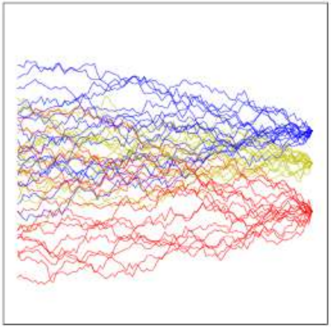

Reverse-time drift

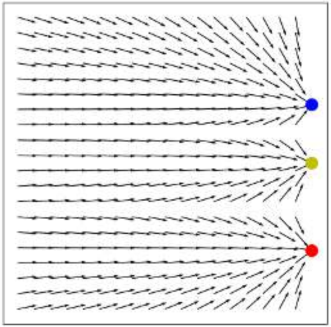

Forward-time drift

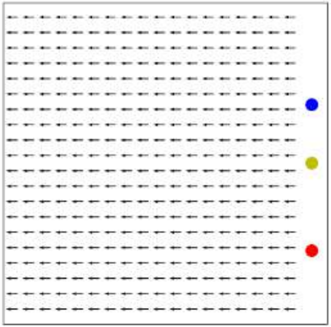

From last time: SDEs as generative models

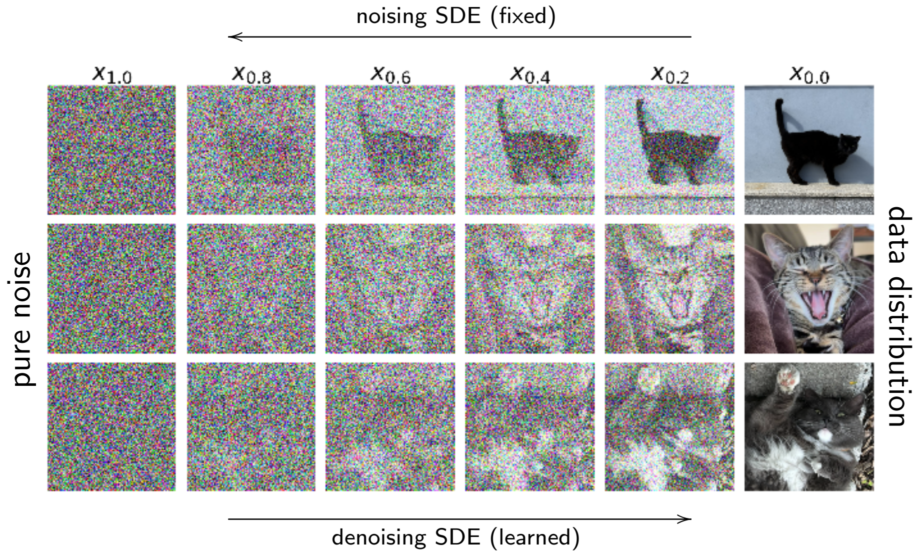

From last time: SDEs as generative models

- Learning score of noised data $\longleftrightarrow$ learning reverse SDE drift
- We can fit an approximation of the score for **all** values of t, then sample the reverse SDE with different numbers of steps
  - Instead of conditioning a denoiser / score model on $n \in \{1, \dots, N\}$, we condition it on $t \in [0, T]$
  - Each step is one approximate Langevin step (on a different target density)

[Song et al., 'Generative modeling by estimating...' blog post]

From last time: SDEs as generative models

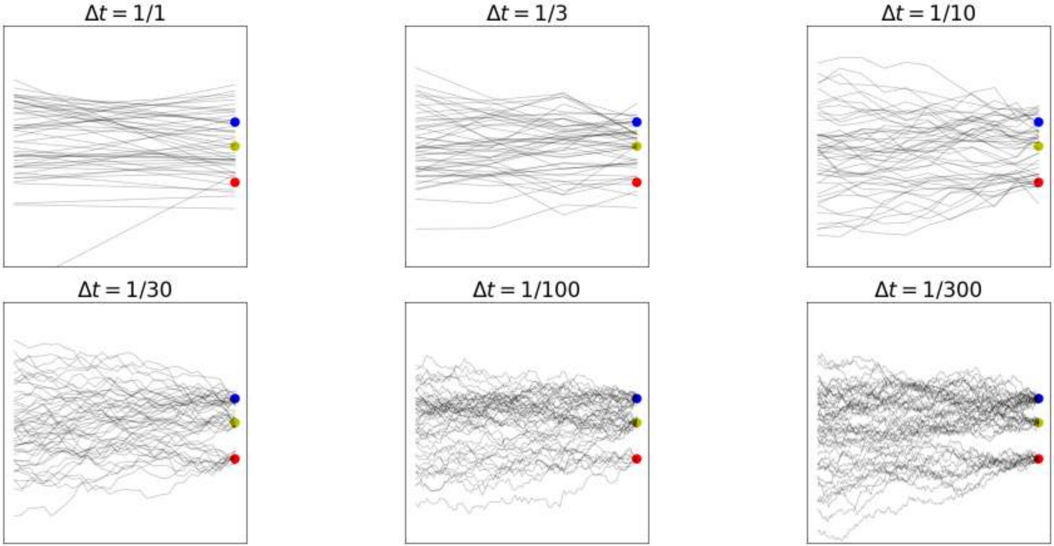

From last time: SDEs as generative models

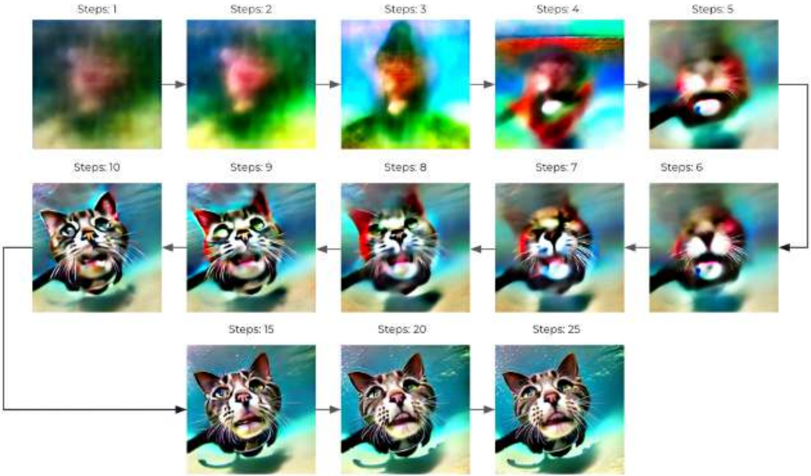

From last time: SDEs as generative models

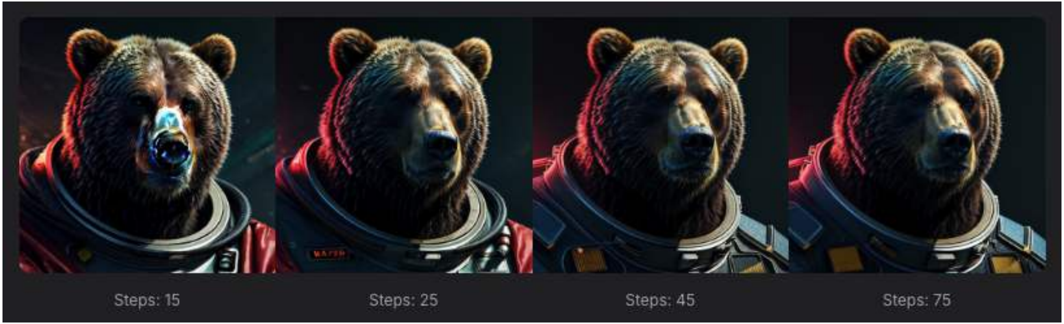

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- Open problems in dynamics-based modelling
- Some announcements

The probability flow ODE

- Forward SDE: $dx_t = f(x_t, t)\, dt + g(t)\, dw_t$
- Reverse SDE: $dx_t = (f(x_t, t) - g(t)^2 \nabla \log p_t(x_t))\, dt + g(t)\, \overline{dw}_t$

The probability flow ODE

- Forward SDE: $dx_t = f(x_t, t)\,dt + g(t)\,dw_t$
- Reverse SDE: $dx_t = (f(x_t, t) - g(t)^2 \nabla \log p_t(x_t))\,dt + g(t)\,\overline{dw}_t$
- Magic fact: this **probability flow ODE** has the same marginal densities p_t as the SDEs given the same initial conditions:

$$dx_t = \left(f(x_t, t) - \frac{1}{2} g(t)^2 \nabla \log p_t(x_t) \right) dt$$

The probability flow ODE

- Forward SDE: $dx_t = f(x_t, t)\, dt + g(t)\, dw_t$
- Reverse SDE: $dx_t = (f(x_t, t) - g(t)^2 \nabla \log p_t(x_t))\, dt + g(t)\, \overline{dw}_t$
- Magic fact: this **probability flow ODE** has the same marginal densities p_t as the SDEs given the same initial conditions:

$$dx_t = \left(f(x_t, t) - \frac{1}{2} g(t)^2 \nabla \log p_t(x_t) \right)\, dt$$

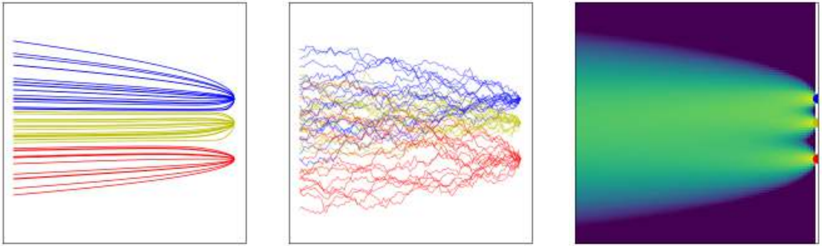

The probability flow ODE

- Forward SDE: $dx_t = f(x_t, t)\,dt + g(t)\,dw_t$
- Reverse SDE: $dx_t = (f(x_t, t) - g(t)^2 \nabla \log p_t(x_t))\,dt + g(t)\,\overline{dw}_t$
- Magic fact: this **probability flow ODE** has the same marginal densities p_t as the SDEs given the same initial conditions:

$$dx_t = \left(f(x_t, t) - \frac{1}{2} g(t)^2 \nabla \log p_t(x_t) \right) dt$$

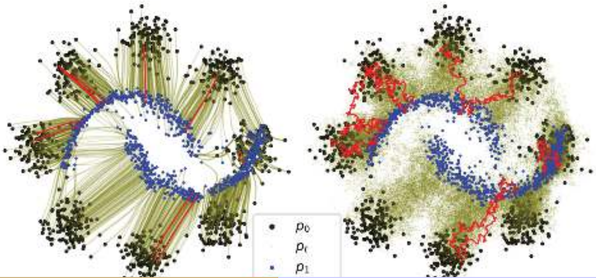

The probability flow ODE

- Forward SDE: $dx_t = f(x_t, t)\, dt + g(t)\, dw_t$
- Reverse SDE: $dx_t = (f(x_t, t) - g(t)^2 \nabla \log p_t(x_t))\, dt + g(t)\, \overline{dw}_t$
- Magic fact: this **probability flow ODE** has the same marginal densities p_t as the SDEs given the same initial conditions:

$$dx_t = \left(f(x_t, t) - \frac{1}{2} g(t)^2 \nabla \log p_t(x_t) \right) dt$$

  - Proof: F-P equation of SDE = continuity equation of ODE (tutorial exercise)
  - Allows to generate by sampling ODE instead of reverse SDE
  - Also allows density estimation (via Hutchinson trace estimator (tutorial exercise))

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- Open problems in dynamics-based modelling
- Some announcements

To noise or not to noise?

Advantages of stochastic modelling:

- Less sensitive to initial integration error, often better samples
- More expressive: non-Gaussian/non-Markovian noise, discrete spaces, ...
- Allows inference of trajectories, not just terminals
- Easier to adapt / fine-tune in certain ways (*e.g.*, using RL)

To noise or not to noise?

Advantages of stochastic modelling:

- Less sensitive to initial integration error, often better samples
- More expressive: non-Gaussian/non-Markovian noise, discrete spaces, ...
- Allows inference of trajectories, not just terminals
- Easier to adapt / fine-tune in certain ways (*e.g.*, using RL)

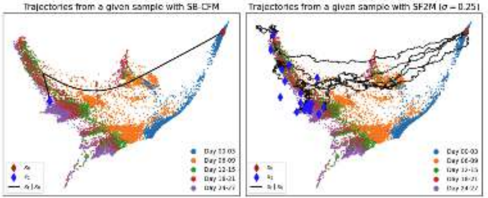

Figure 3: Simulation of trajectories from a given cell on 2D EB data. **Left:** Probability flow ODE trajectory, approximated by SB-CFM (Tong et al., 2024). **Right:** Five SDE trajectories from $[SF]^2M$; more target samples (20) in blue.

[Tong et al., 'Simulation-free Schrödinger bridges', AISTATS'24]

To noise or not to noise?

Advantages of stochastic modelling:

- Less sensitive to initial integration error, often better samples
- More expressive: non-Gaussian/non-Markovian noise, discrete spaces, ...
- Allows inference of trajectories, not just terminals
- Easier to adapt / fine-tune in certain ways (*e.g.*, using RL)

Advantages of deterministic (ODE) modelling:

- Integration in fewer steps (smooth paths, faster sampling, amortised integration)
- Deterministic mapping from x_0 to x_T (latent encoding!)
- Sometimes easier to generalise (flow matching – next slide)
- Easier to adapt in other ways (by modifying the noise distribution)

Flow matching

Algorithms for training the drift of an ODE

$$dx_t = f(x_t, t; \theta)\, dt$$

without assuming a underlying SDE...

Published as a conference paper at ICLR 2023

Flow Matching for Generative Modeling

Yaron Lipman[1,2] Ricky T. Q. Chen[1] Heli Ben-Hamu[2] Maximilian Nickel[1] Matt Le[1]
[1]Meta AI (FAIR) [2]Weizmann Institute of Science

Stochastic Interpolants: A Unifying Framework for Flows and Diffusions

Michael S. Albergo*[1], Nicholas M. Boffi*[2], and Eric Vanden-Eijnden[2]
[1]Center for Cosmology and Particle Physics, New York University
[2]Courant Institute of Mathematical Sciences, New York University

November 7, 2023

Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow

Xingchao Liu* University of Texas at Austin
Chengyue Gong* University of Texas at Austin
Qiang Liu University of Texas at Austin

Published in Transactions on Machine Learning Research

Improving and Generalizing Flow-Based Generative Models with Minibatch Optimal Transport

Alexander Tong*
Kilian Fatras*
Nikolay Malkin*

`https://pypi.org/project/torchcfm/`

Flow matching

Algorithms for training the drift of an ODE

$$dx_t = f(x_t, t; \theta)\, dt$$

without assuming a underlying SDE...

- Suitable for data-to-data modelling (bridges)
- Connections to optimal transport (straighter paths)

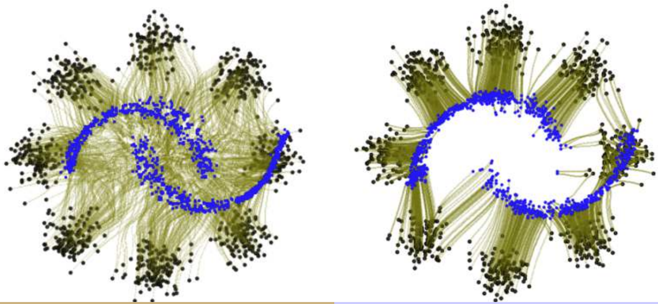

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- Open problems in dynamics-based modelling
- Some announcements

Diffusion on manifolds

- So far, we reversed Brownian motion on $\mathbb{R}^d$
- But Brownian motion defined on any Riemannian manifold
  - On some simple spaces (spheres, tori) we have ‘magic properties’

Diffusion on manifolds

- So far, we reversed Brownian motion on $\mathbb{R}^d$
- But Brownian motion defined on any Riemannian manifold
  - On some simple spaces (spheres, tori) we have ‘magic properties’
- Interesting open problems:
  - Diffusion and manifold learning

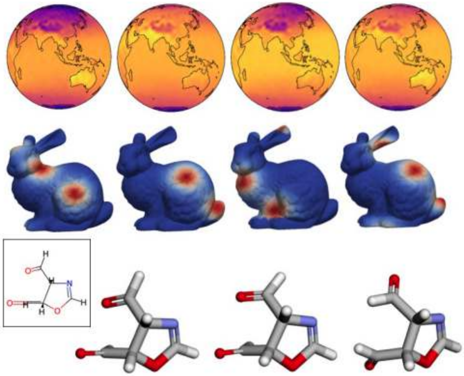

[Elhag et al., ‘Manifold diffusion fields’, ICLR’24]

Diffusion on manifolds

- So far, we reversed Brownian motion on $\mathbb{R}^d$
- But Brownian motion defined on any Riemannian manifold
  - On some simple spaces (spheres, tori) we have 'magic properties'
- Interesting open problems:
  - Diffusion and manifold learning
  - How to read data manifold properties from a diffusion model?

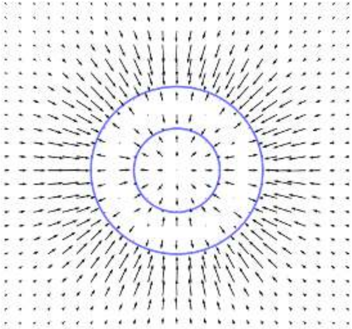

[Stanczuk et al., 'Diffusion models encode the intrinsic dimension of data manifolds', ICML'24]
see also [Kadkhodaie et al., 'Generalization in diffusion models arises from geometry-adaptive harmonic representations', ICLR'24], [Jones, 'Manifold diffusion geometry']

Interesting symmetric spaces

- What if we know the target distribution on $\mathbb{R}^d$ is invariant to some group action?
  - $\mathrm{O}(d)$: rotation and reflection invariance
  - $\mathrm{SE}(d) = \mathrm{SO}(d) \ltimes \mathbb{R}^d$: rotation and translation invariance (*e.g.*, molecule conformations, 3D objects)
  - S_n: permutation invariance (*e.g.*, point clouds)

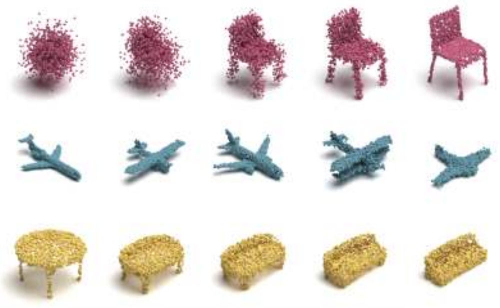

[Luo et al, 'Diffusion probabilistic models for 3D...', CVPR'21]

Interesting symmetric spaces

- What if we know the target distribution on $\mathbb{R}^d$ is invariant to some group action?
- Make the noising distribution group-invariant and the denoiser group-equivariant
  - Noising transformed data and denoising = transforming noised data and denoising = denoising and transforming the result

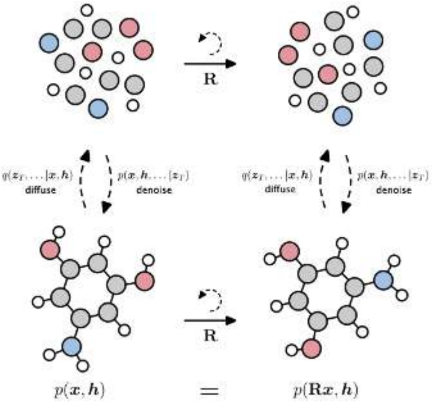

[Hoogeboom et al., 'Equivariant diffusion. . . ', ICML'22]

Interesting symmetric spaces

- What if we know the target distribution on $\mathbb{R}^d$ is invariant to some group action?
- Make the noising distribution group-invariant and the denoiser group-equivariant
  - Noising transformed data and denoising = transforming noised data and denoising = denoising and transforming the result
- But we can also learn invariances using data augmentation
  - AlphaFold's bitter lesson...

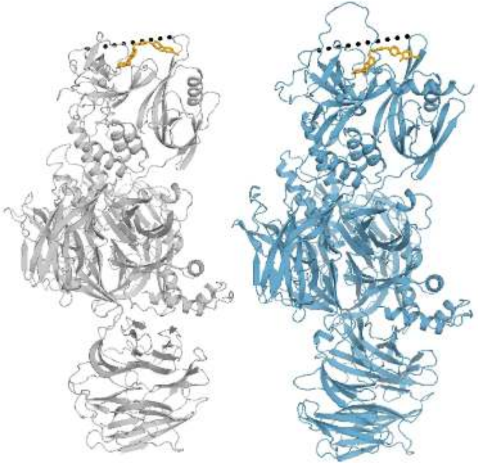

Diffusion for text I: Embedding spaces

Gaussian diffusion in an embedding space (latent diffusion model)

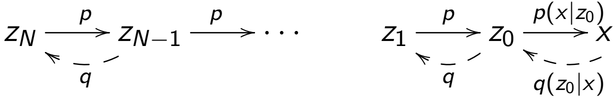

Diffusion for text I: Embedding spaces

Gaussian diffusion in an embedding space (latent diffusion model)

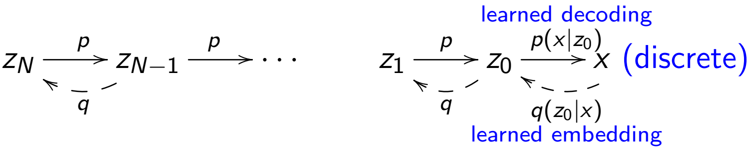

Diffusion for text I: Embedding spaces

Gaussian diffusion in an embedding space (latent diffusion model)

$$z_N \underset{q}{\overset{p}{\rightleftarrows}} z_{N-1} \xrightarrow{p} \cdots \qquad z_1 \underset{q}{\overset{p}{\rightleftarrows}} z_0 \underset{q(z_0|x)}{\overset{p(x|z_0)}{\rightleftarrows}} x \text{ (discrete)}$$

learned decoding $p(x|z_0)$

learned embedding $q(z_0|x)$

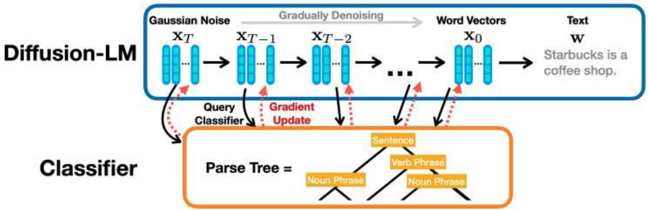

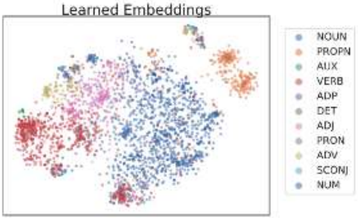

[Li et al, 'Diffusion-LM...', NeurIPS'22]

Diffusion for text II: Discrete-space diffusion

Replace Gaussian noise by random transitions (token replacement or **masking**)

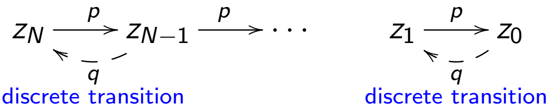

Diffusion for text II: Discrete-space diffusion

Replace Gaussian noise by random transitions (token replacement or **masking**)

$$z_N \underset{q}{\overset{p}{\rightleftarrows}} z_{N-1} \xrightarrow{p} \cdots \qquad z_1 \underset{q}{\overset{p}{\rightleftarrows}} z_0$$

discrete transition   discrete transition

- Have analogues of 'magic properties' enabling efficient training
- BERT is a special denoiser (masking noising process, 15% noise level)

Diffusion for text III: Simplices

A middle way: hidden states are full distributions over tokens

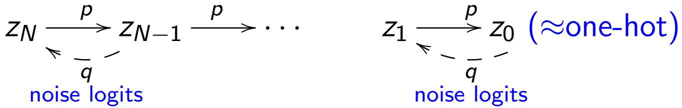

Diffusion for text III: Simplices

A middle way: hidden states are full distributions over tokens

$$z_N \underset{q}{\overset{p}{\rightleftarrows}} z_{N-1} \overset{p}{\longrightarrow} \cdots \quad z_1 \underset{q}{\overset{p}{\rightleftarrows}} z_0$$

(≈one-hot)

noise logits

t	$\operatorname{argmax} \boldsymbol{w}_{\text{logits},t}^{c:c+B}$	$\operatorname{argmax} \tilde{\boldsymbol{w}}_{t-1}^{c:c+B}$
2500	of the to the the the the the the the the the the the the the the the the the the the the the the	apeshifteriao41 fleeting frontman Nutdrop278temp Drama lime Employee cuc rival greatest kan snakes431 cav dreamedRange alloy originally Pact
1500	is the to be the, of the,,,\n the the the the the the the the the\n into the.	stunnedchildrenmetrywaveopensLayer Porn woman transcend242 Homs PluginNext Endsackle microbi spokesperson Brunswick awards":- Sharma Pinball Jr Rug wrapped
1300	of the mission it the as a, for the,,, as to as the the moons, and Cass Cassini is	178 whit promoters du basketballiche SchoolsPur Sack reward basketball corn////WeaponSpeaking squid Chains Caucasian McGivity Me SC rafthr jihadist
1100	was based on the in, 2014. Theini will be the the up is the the the the, Hubble but the the	battles swore starters test thanpadding ambiguityFri BADuitous Stuff depiction bankrupt >>> conversions240Genelvet aptLegweight Riy modesitanesday
900	of the Jarminiini Cass Gr, was supposed to be the most ambitious and most attempt to capture all most distant moons	Sim bag Ves serotonin._ Fab gameplay ransom Alisonorks Fargo expand Rhode pursuing most plagued formulateheter plainly troubled Professional Binary Creek geared
800	is all about Saturn. The Eini will, the closest that the instruments have reached will be to stop in on the Saturn	omial allcounter Saturn. The Directthank Ecuador two thelearning that the Animation have brothers will make tousands downtown governance the Further
700	will allow the Cass to finally see the planet's relatively small atmosphere and finally be able to procure an accurate way of understanding how	willPocket prelim Klux to finally see the planet intelligent relatively jumper atmosphere and halted Fly activityvirt00000 trem accurate way of Inferno what
600	will allow the scientists to better study the effects of Grand Impact, and also be able to get much more data and images of	will allowert scientists Damien better study the effects of Grand Impact, andasket bebery to get much more data and images of
500	will allow the scientists to better see the interior of its atmosphere, and also be able to get much more knowledge and understanding of	will allow the scientists to better see the interior of its atmosphere, and also be able to get much more knowledge and understanding of
1	will allow the scientists to better see the interior of its atmosphere, and also be able to get much more knowledge and observations of	will allow the scientists to better see the interior of its atmosphere, and also be able to get much more knowledge and observations of

[Han et al., ‘SSD-LM. . . ’, ACL’23]; *see also* [Han et al., ‘David helps Goliath. . . ’, NAACL’24]

Other unusual modalities

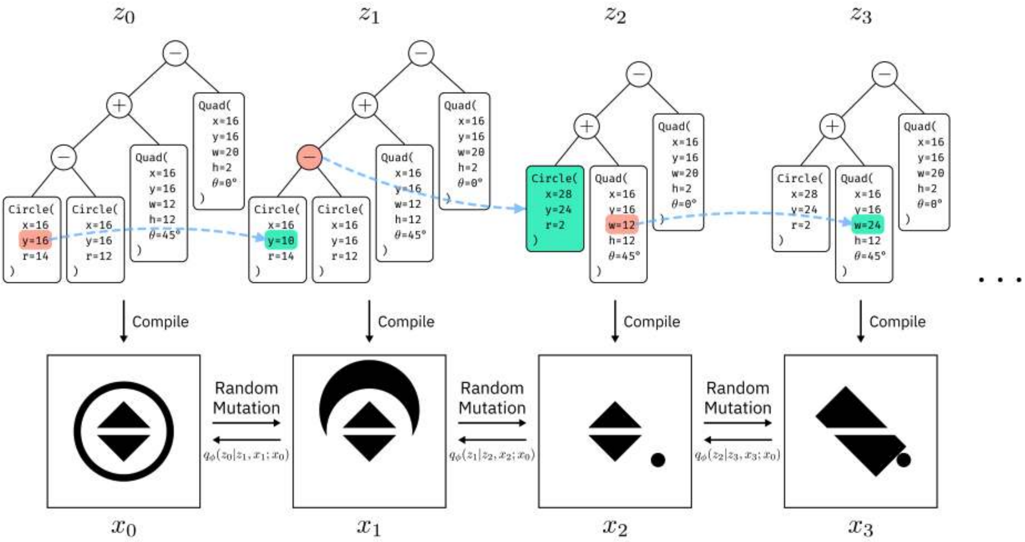

Other unusual modalities

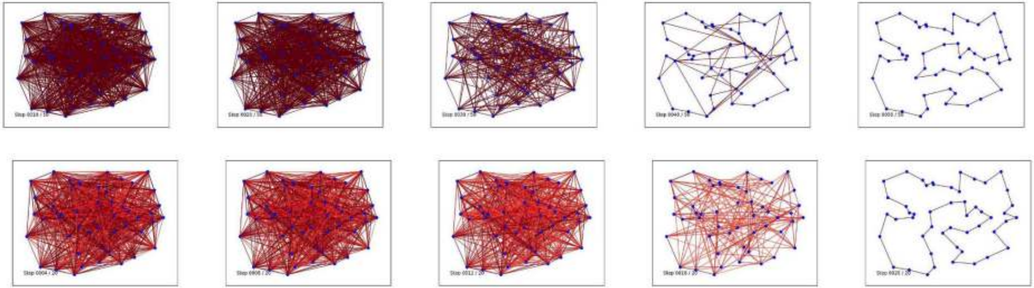

[Sun et al., ‘DIFUSCO. . . ’, NeurIPS’23]

Other unusual modalities

[Zhang and Gienger, χ:2409.01083]

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- Open problems in dynamics-based modelling
- Some announcements

Diffusion models without data?

- Diffusion models are trained from data to maximise ELBO, or minimise

$$\mathrm{KL}(\text{target distribution} \cdot \text{noising process} \,\|\, \text{denoising process})$$

Diffusion models without data?

- Diffusion models are trained from data to maximise ELBO, or minimise

$$\mathrm{KL}(\text{target distribution} \cdot \text{noising process} \,\|\, \text{denoising process})$$

- But what if we have only a target density / energy $R(\mathbf{x})$?
  - Thought of as unnormalised 'reward' (*e.g.*, a posterior $p(x \mid y) \propto p(x)p(y \mid x)$)
  - Related: product of diffusion prior $p(\mathbf{x})$ and constraint (recall the intractability of exact conditioning from Lecture 9)

Diffusion models without data?

- Diffusion models are trained from data to maximise ELBO, or minimise

$$\mathrm{KL}(\text{target distribution} \cdot \text{noising process} \,\|\, \text{denoising process})$$

- But what if we have only a target density / energy $R(\mathbf{x})$?
  - Thought of as unnormalised 'reward' (*e.g.*, a posterior $p(x \mid y) \propto p(x)p(y \mid x)$)
  - Related: product of diffusion prior $p(\mathbf{x})$ and constraint (recall the intractability of exact conditioning from Lecture 9)

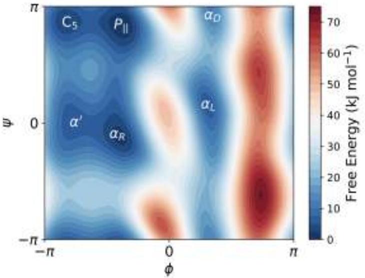

[Phillips et al., χ:2408.15905]

Diffusion models without data?

- Diffusion models are trained from data to maximise ELBO, or minimise

$$\mathrm{KL}(\text{target distribution} \cdot \text{noising process} \,\|\, \text{denoising process})$$

- But what if we have only a target density / energy $R(\mathbf{x})$?
  - Thought of as unnormalised 'reward' (*e.g.*, a posterior $p(x \mid y) \propto p(x)p(y \mid x)$)
  - Related: product of diffusion prior $p(\mathbf{x})$ and constraint (recall the intractability of exact conditioning from Lecture 9)

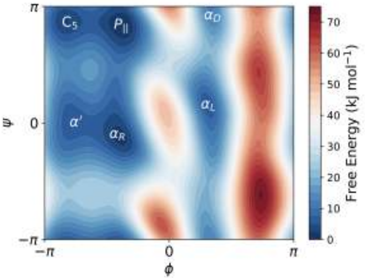

[Phillips et al., χ:2408.15905]

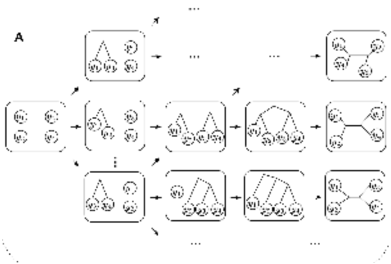

[Zhou et al., ICLR'24]

Diffusion models without data?

- Diffusion models are trained from data to maximise ELBO, or minimise

$$\mathrm{KL}(\text{target distribution} \cdot \text{noising process} \,\|\, \text{denoising process})$$

- But what if we have only a target density / energy $R(\mathbf{x})$?
  - Thought of as unnormalised 'reward' (*e.g.*, a posterior $p(x \mid y) \propto p(x)p(y \mid x)$)
  - Related: product of diffusion prior $p(\mathbf{x})$ and constraint (recall the intractability of exact conditioning from Lecture 9)
- One option: reverse KL

$$\mathrm{KL}(\text{denoising process} \,\|\, \text{target distribution} \cdot \text{noising process})$$

  - Works but has some issues (cost of differentiable simulation, mode-seeking, white-box reward)
    [Path integral sampler, ICLR'22]
    [Denoising diffusion sampler, ICLR'23]

Diffusion models as RL agents

- Denoising process $\leftrightarrow$ policy in a certain MDP

Diffusion models as RL agents

- Denoising process $\leftrightarrow$ policy in a certain MDP
- Target log-density $\log R(\cdot) \leftrightarrow$ reward

Diffusion models as RL agents

- Denoising process $\leftrightarrow$ policy in a certain MDP
- Target log-density $\log R(\cdot) \leftrightarrow$ reward
- A little different than typical RL
  - Sample proportionally to exp(reward), not maximise it

Diffusion models as RL agents

- Denoising process $\leftrightarrow$ policy in a certain MDP
- Target log-density $\log R(\cdot) \leftrightarrow$ reward
- A little different than typical RL
  - Sample proportionally to exp(reward), not maximise it
  - How to understand the 'reverse policy' (noising process)?

Diffusion models as RL agents

- Denoising process $\leftrightarrow$ policy in a certain MDP
- Target log-density $\log R(\cdot)$ $\leftrightarrow$ reward
- A little different than typical RL
  - Sample proportionally to exp(reward), not maximise it
  - How to understand the 'reverse policy' (noising process)?
  - Several RL / control theory / SDE formalisms for this [Vargas et al., ICLR'24], [Richter & Berner, ICLR'24] [Lahlou, . . . , Malkin, ICML'23], [Sendera, . . . , Malkin, NeurIPS'24] Unifying picture: [Berner, . . . , Malkin, TMLR, 2026]

Diffusion models as RL agents

- Denoising process $\leftrightarrow$ policy in a certain MDP
- Target log-density $\log R(\cdot) \leftrightarrow$ reward
- A little different than typical RL
  - Sample proportionally to exp(reward), not maximise it
  - How to understand the 'reverse policy' (noising process)?
  - Several RL / control theory / SDE formalisms for this [Vargas et al., ICLR'24], [Richter & Berner, ICLR'24] [Lahlou, . . . , Malkin, ICML'23], [Sendera, . . . , Malkin, NeurIPS'24] Unifying picture: [Berner, . . . , Malkin, TMLR, 2026]

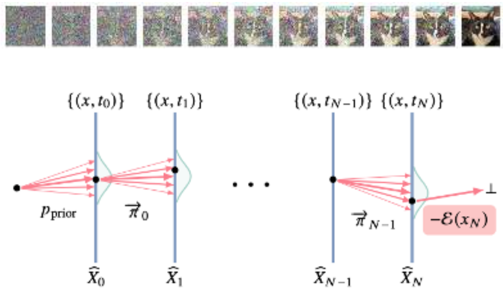

Amortising intractable posteriors under diffusion priors

What about sampling x_T from $p(x_T \mid y) \propto p(x_T)p(y \mid x_T)$, where $p(x_T)$ is a pretrained **diffusion prior** and $p(y \mid x_T)$ is a likelihood?

- Intractable in general (recall Lecture 9)

Amortising intractable posteriors under diffusion priors

What about sampling x_T from $p(x_T \mid y) \propto p(x_T)p(y \mid x_T)$, where $p(x_T)$ is a pretrained **diffusion prior** and $p(y \mid x_T)$ is a likelihood?

- Intractable in general (recall Lecture 9)

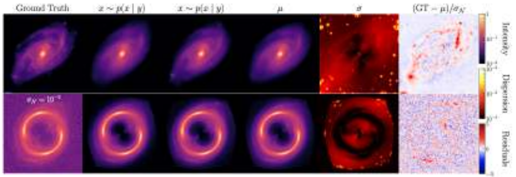

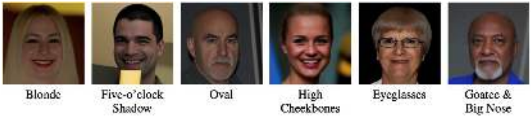

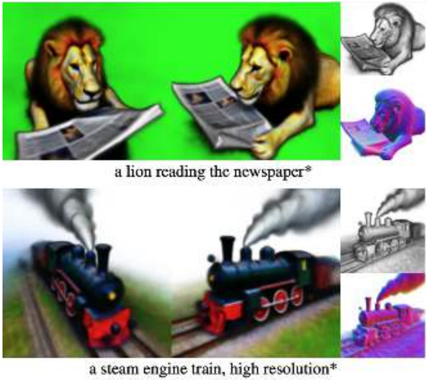

Amortising intractable posteriors under diffusion priors

What about sampling x_T from $p(x_T \mid y) \propto p(x_T)p(y \mid x_T)$, where $p(x_T)$ is a pretrained **diffusion prior** and $p(y \mid x_T)$ is a likelihood?

- Intractable in general (recall Lecture 9)

Amortising intractable posteriors under diffusion priors

Class-conditional image models from unconditional priors

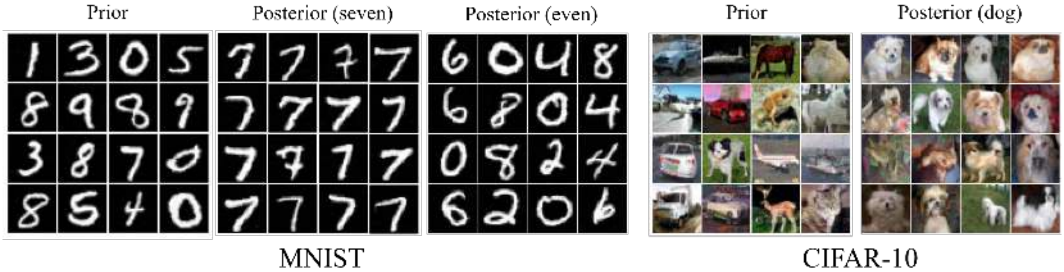

Amortising intractable posteriors under diffusion priors

(Non-!)linear inverse problems (with applications in inverse imaging)

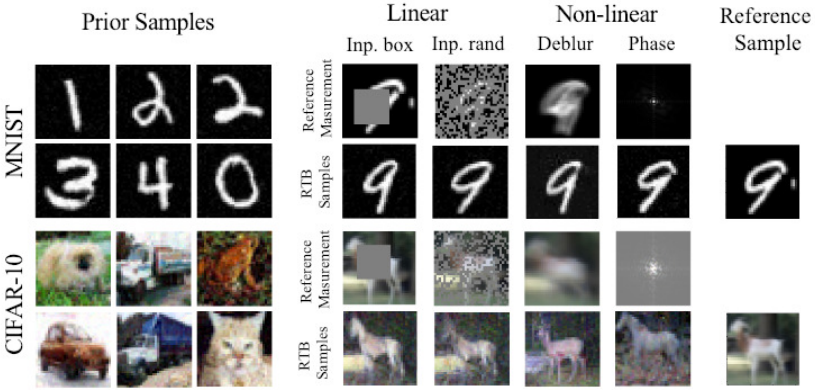

[Venkatraman et al., NeurIPS'24]

Amortising intractable posteriors under diffusion priors

(Non-!)linear inverse problems (with applications in inverse imaging)

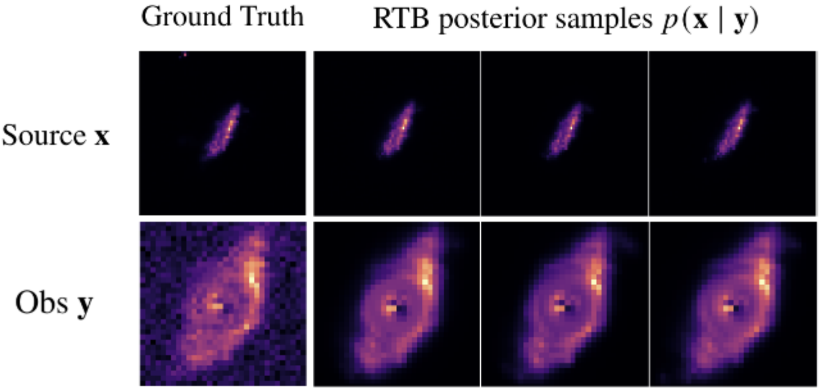

[Venkatraman et al., NeurIPS'24]

Amortising intractable posteriors under diffusion priors

Table 1: Sources of diffusion priors and constraints.

Domain	Prior $p(\mathbf{x})$	Constraint $r(\mathbf{x})$	Posterior
Conditional image generation (§4.1)	Image diffusion model $p(\mathbf{x})$	Classifier likelihood $p(c \mid \mathbf{x})$	Class-conditional distribution $p(\mathbf{x} \mid c)$
Text-to-image generation (§4.2)	Text-to-image foundation model	RLHF reward model	Aligned text-to-image model
Language infilling (§4.3)	Discrete diffusion model	Autoregressive completion likelihood	Infilling distribution
Offline RL policy extraction (§4.4)	Diffusion model as behavior policy	Boltzmann dist. of Q-function	Optimal KL-constrained policy

Other applications:

- Discrete-space diffusion (text)
- Offline RL policy extraction
- Black-box Bayesian optimisation [Yun et al., χ:2502.16824]

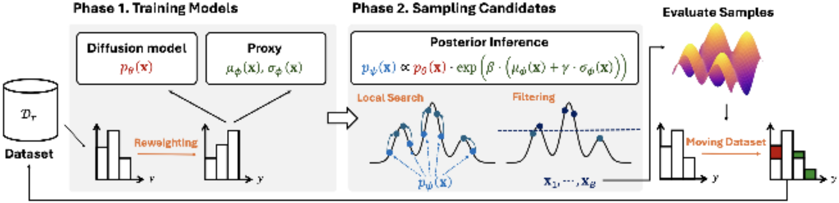

Inference in latent spaces of generative models

'Outsourced' diffusion sampling: sample posteriors in latent spaces of GANs, VAEs, etc., given a constraint on the output space

[Venkatraman et al., ICML'25]

Inference in latent spaces of generative models

'Outsourced' diffusion sampling: sample posteriors in latent spaces of GANs, VAEs, etc., given a constraint on the output space

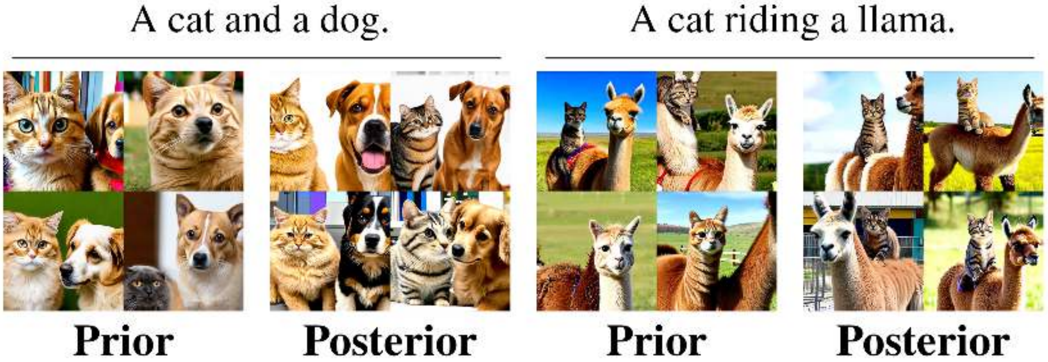

[Venkatraman et al., ICML'25]

Bridge problems

Translation $p_{\text{prior}} \leftrightarrow p_{\text{prior}} \cdot p(\text{class} \mid \cdot)$ in the latent space of a generative model

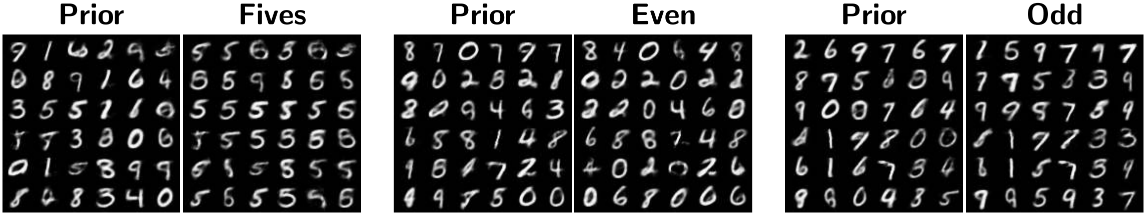

[Tamogashev & Malkin, ICLR'26]

Bridge problems

Translation $p_{\text{prior}} \leftrightarrow p_{\text{prior}} \cdot p(\text{class} \mid \cdot)$ in the latent space of a generative model

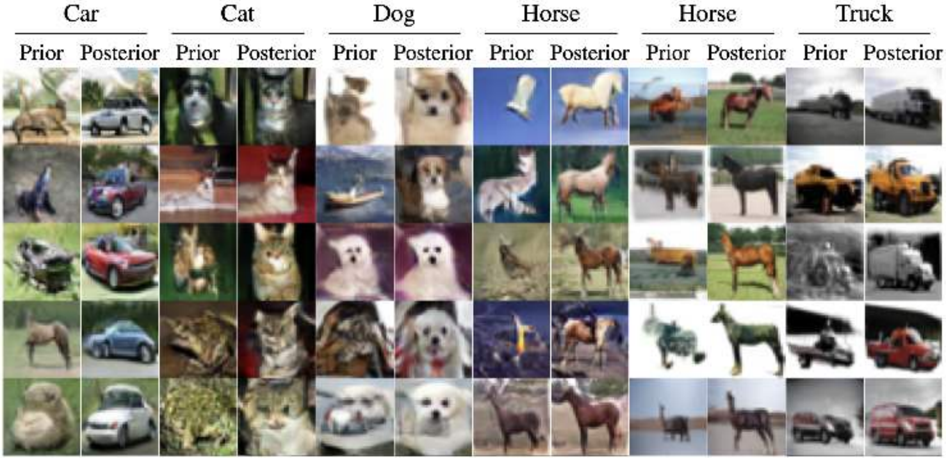

[Tamogashev & Malkin, ICLR'26]

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- **Open problems in dynamics-based modelling**
- Some announcements

Open problems I: Efficiency

How to reduce the number of neural net evaluations needed to generate samples?

- Distillation, amortised integrators

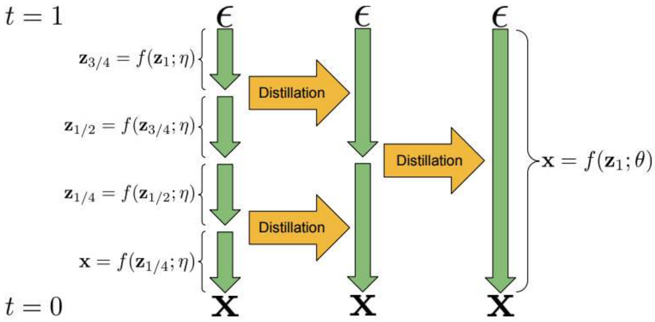

[Salimans & Ho, 'Progressive distillation...', ICLR'22]

Open problems I: Efficiency

How to reduce the number of neural net evaluations needed to generate samples?

- Distillation, amortised integrators
- Related theoretical questions

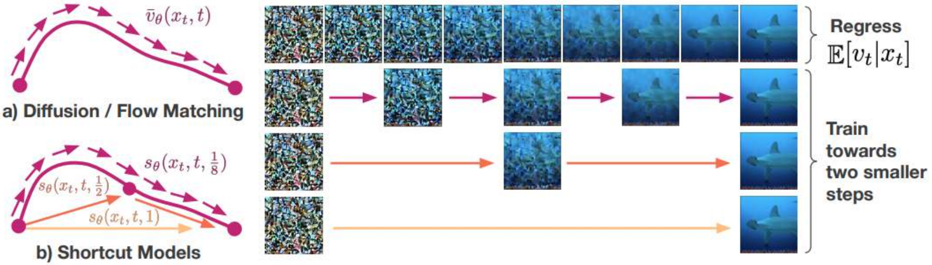

[Frans et al., 'One step diffusion via shortcut models', 2024]

see also: recent work on **consistency models**, **shortcut models**

Open problems II: Representation learning

Diffusion models have a simple denoising objective – what do they learn?

- What information is in hidden representations?

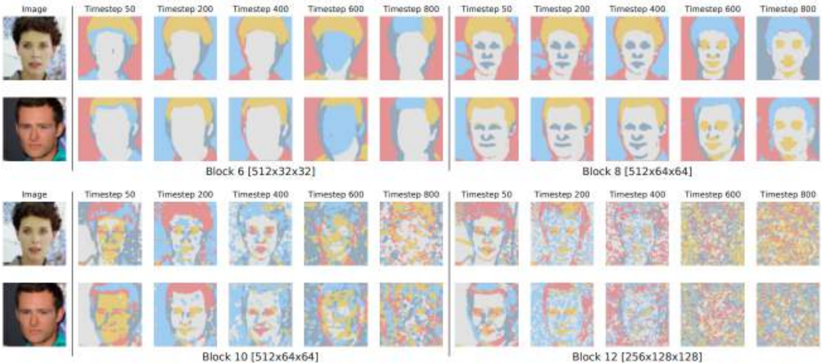

[Baranchuk et al., 'Label-efficient semantic segmentation with diffusion models, ICLR'22]

Open problems II: Representation learning

Diffusion models have a simple denoising objective – what do they learn?

- What information is in hidden representations?
- How to use diffusion models for representation learning or as encoders?

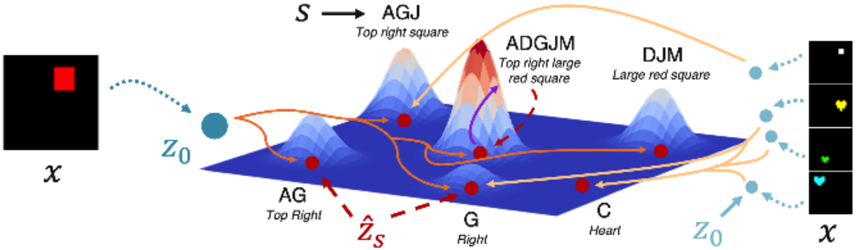

[Nam et al., 'Discrete, compositional, and symbolic representations through attractor dynamics', 2024]

Open problems III: Structure and generalisation

- Why do diffusion models generalise well compositionally despite not having explicit representation of structure?

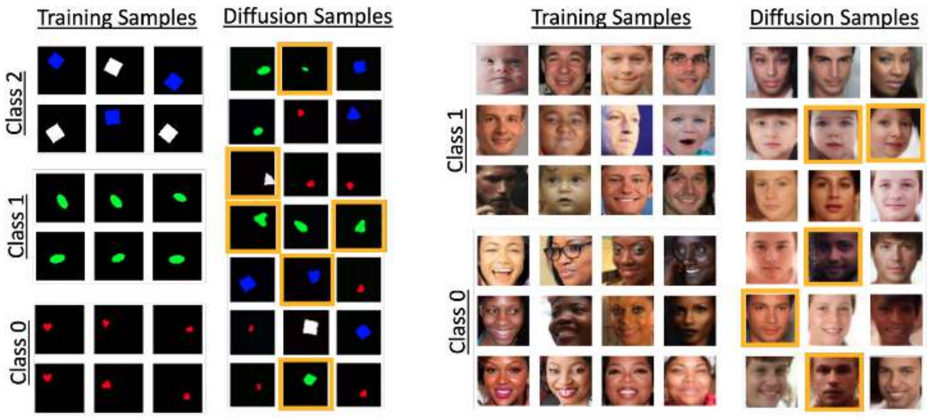

[Scimeca et al., 'Leveraging diffusion disentangled representations. . . ', 2023]

Open problems III: Structure and generalisation

- Why do diffusion models generalise well compositionally despite not having explicit representation of structure?
- How to impose (symbolic!) structure on the distribution or dynamics?

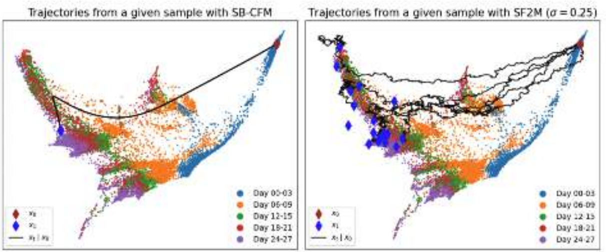

Figure 3: Simulation of trajectories from a given cell on 2D EB data. **Left:** Probability flow ODE trajectory, approximated by SB-CFM (Tong et al., 2024). **Right:** Five SDE trajectories from $[SF]^2M$; more target samples (20) in blue.

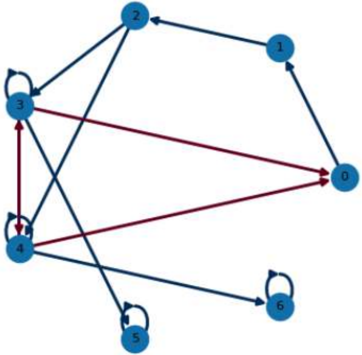

[Tong et al., 'Simulation-free Schrödinger bridges. . . ', AISTATS'24]

- Review of Lecture 10
- Finishing continuous-time models
  - SDEs vs. ODEs & generalisations
- Diffusion models for various data modalities
  - Constraints and invariances in diffusion models
  - Three approaches to discrete diffusion for text
  - Other unusual modalities
- Diffusion models without data
  - Diffusion models are RL agents
  - Applications to Bayesian inference, inverse problems, RLHF
- Open problems in dynamics-based modelling
- **Some announcements**

Some announcements

- Office hour: Monday, 20 April, 1-3pm, Appleton Tower 4.12.
- Ask questions by 23 April in Piazza.
- Exam: all the questions for this track are easy ;)

Some announcements

Thank you for your attention this term!
Go forth and generate!

- Office hour: Monday, 20 April, 1-3pm, Appleton Tower 4.12.
- Ask questions by 23 April in Piazza.
- Exam: all the questions for this track are easy ;)