Conditioning in SDEs

9 February 2024

R255 Lecture

Shreyas Padhy

So far - Reversal of SDEs

- We have a forward SDE $dX_t = f^+(X_t, t)\, dt + \sigma(X_t, t)\, dW_t$
- We can reverse the SDE $dX_t = f^-(X_t, t)\, dt + \sigma(X_t, t)\, d\overline{W}_t$

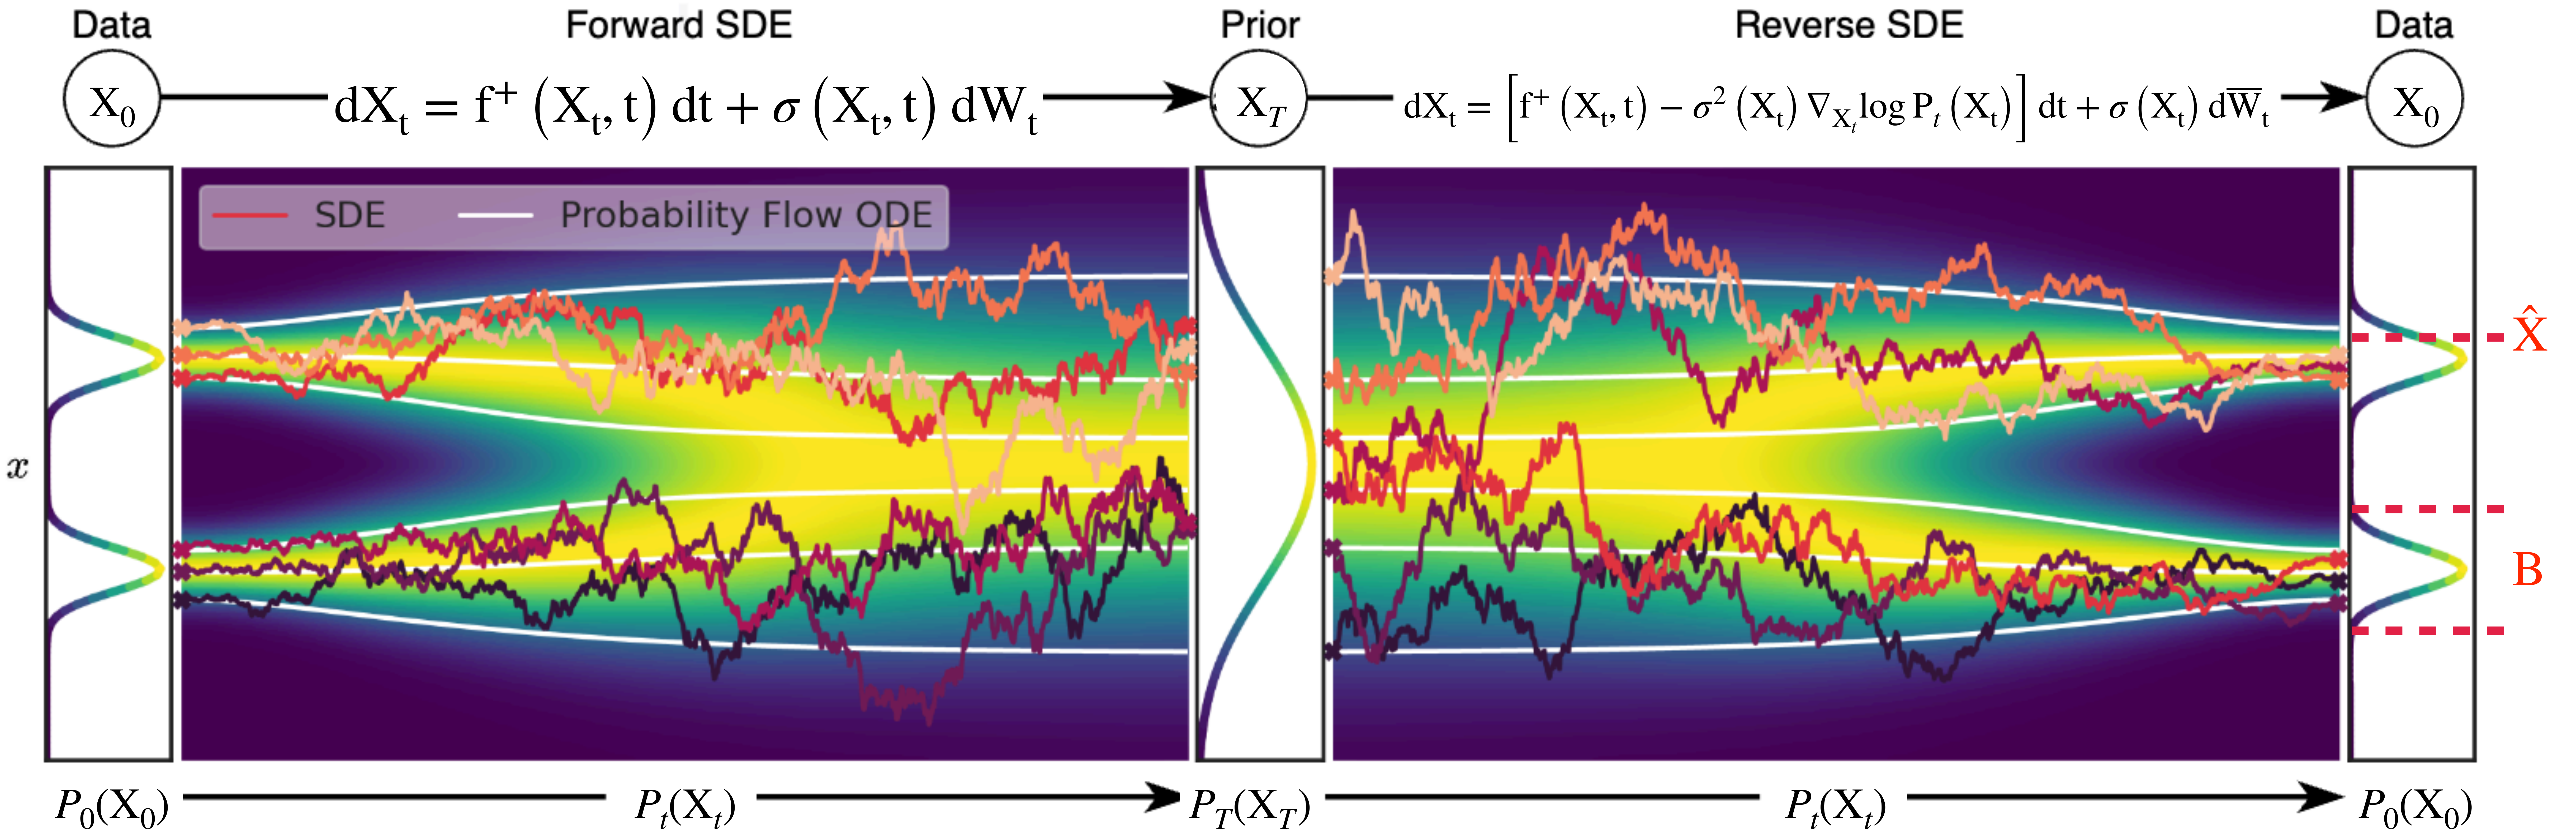

Conditioning of SDEs

- We want the reverse SDE to hit a certain constraint set $X_0 \in B$ at time $T = 0$

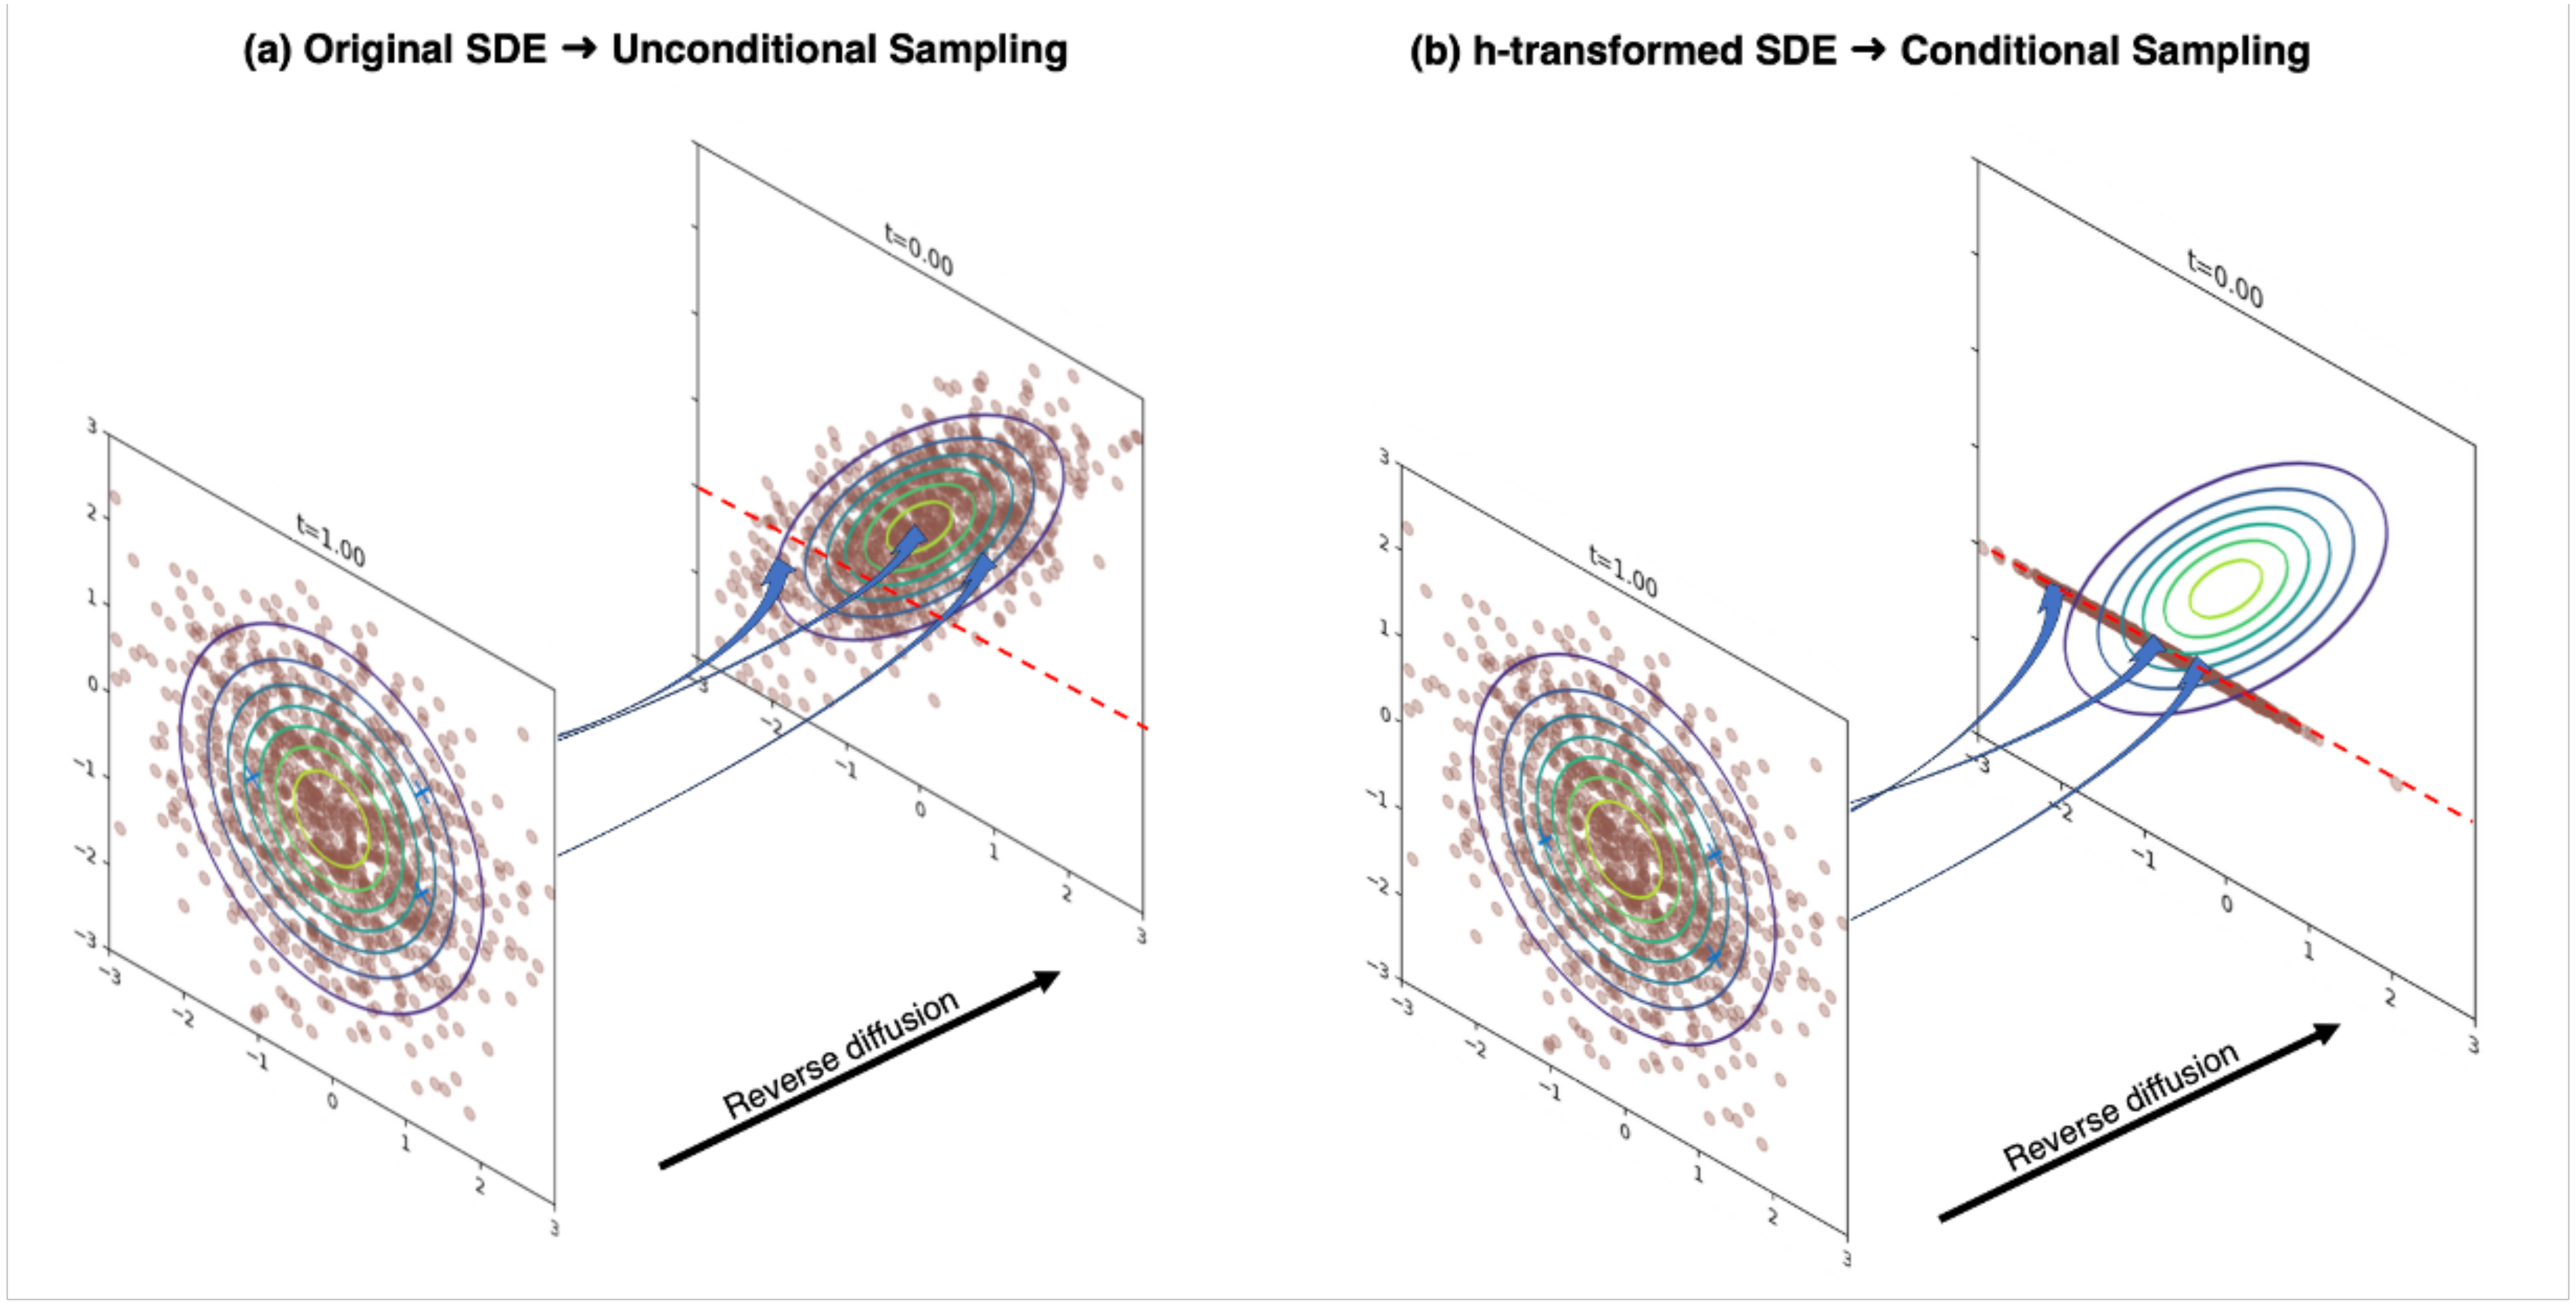

Conditioning of SDEs

- We want the reverse SDE to hit a certain constraint set $X_0 \in B$ at time $T = 0$

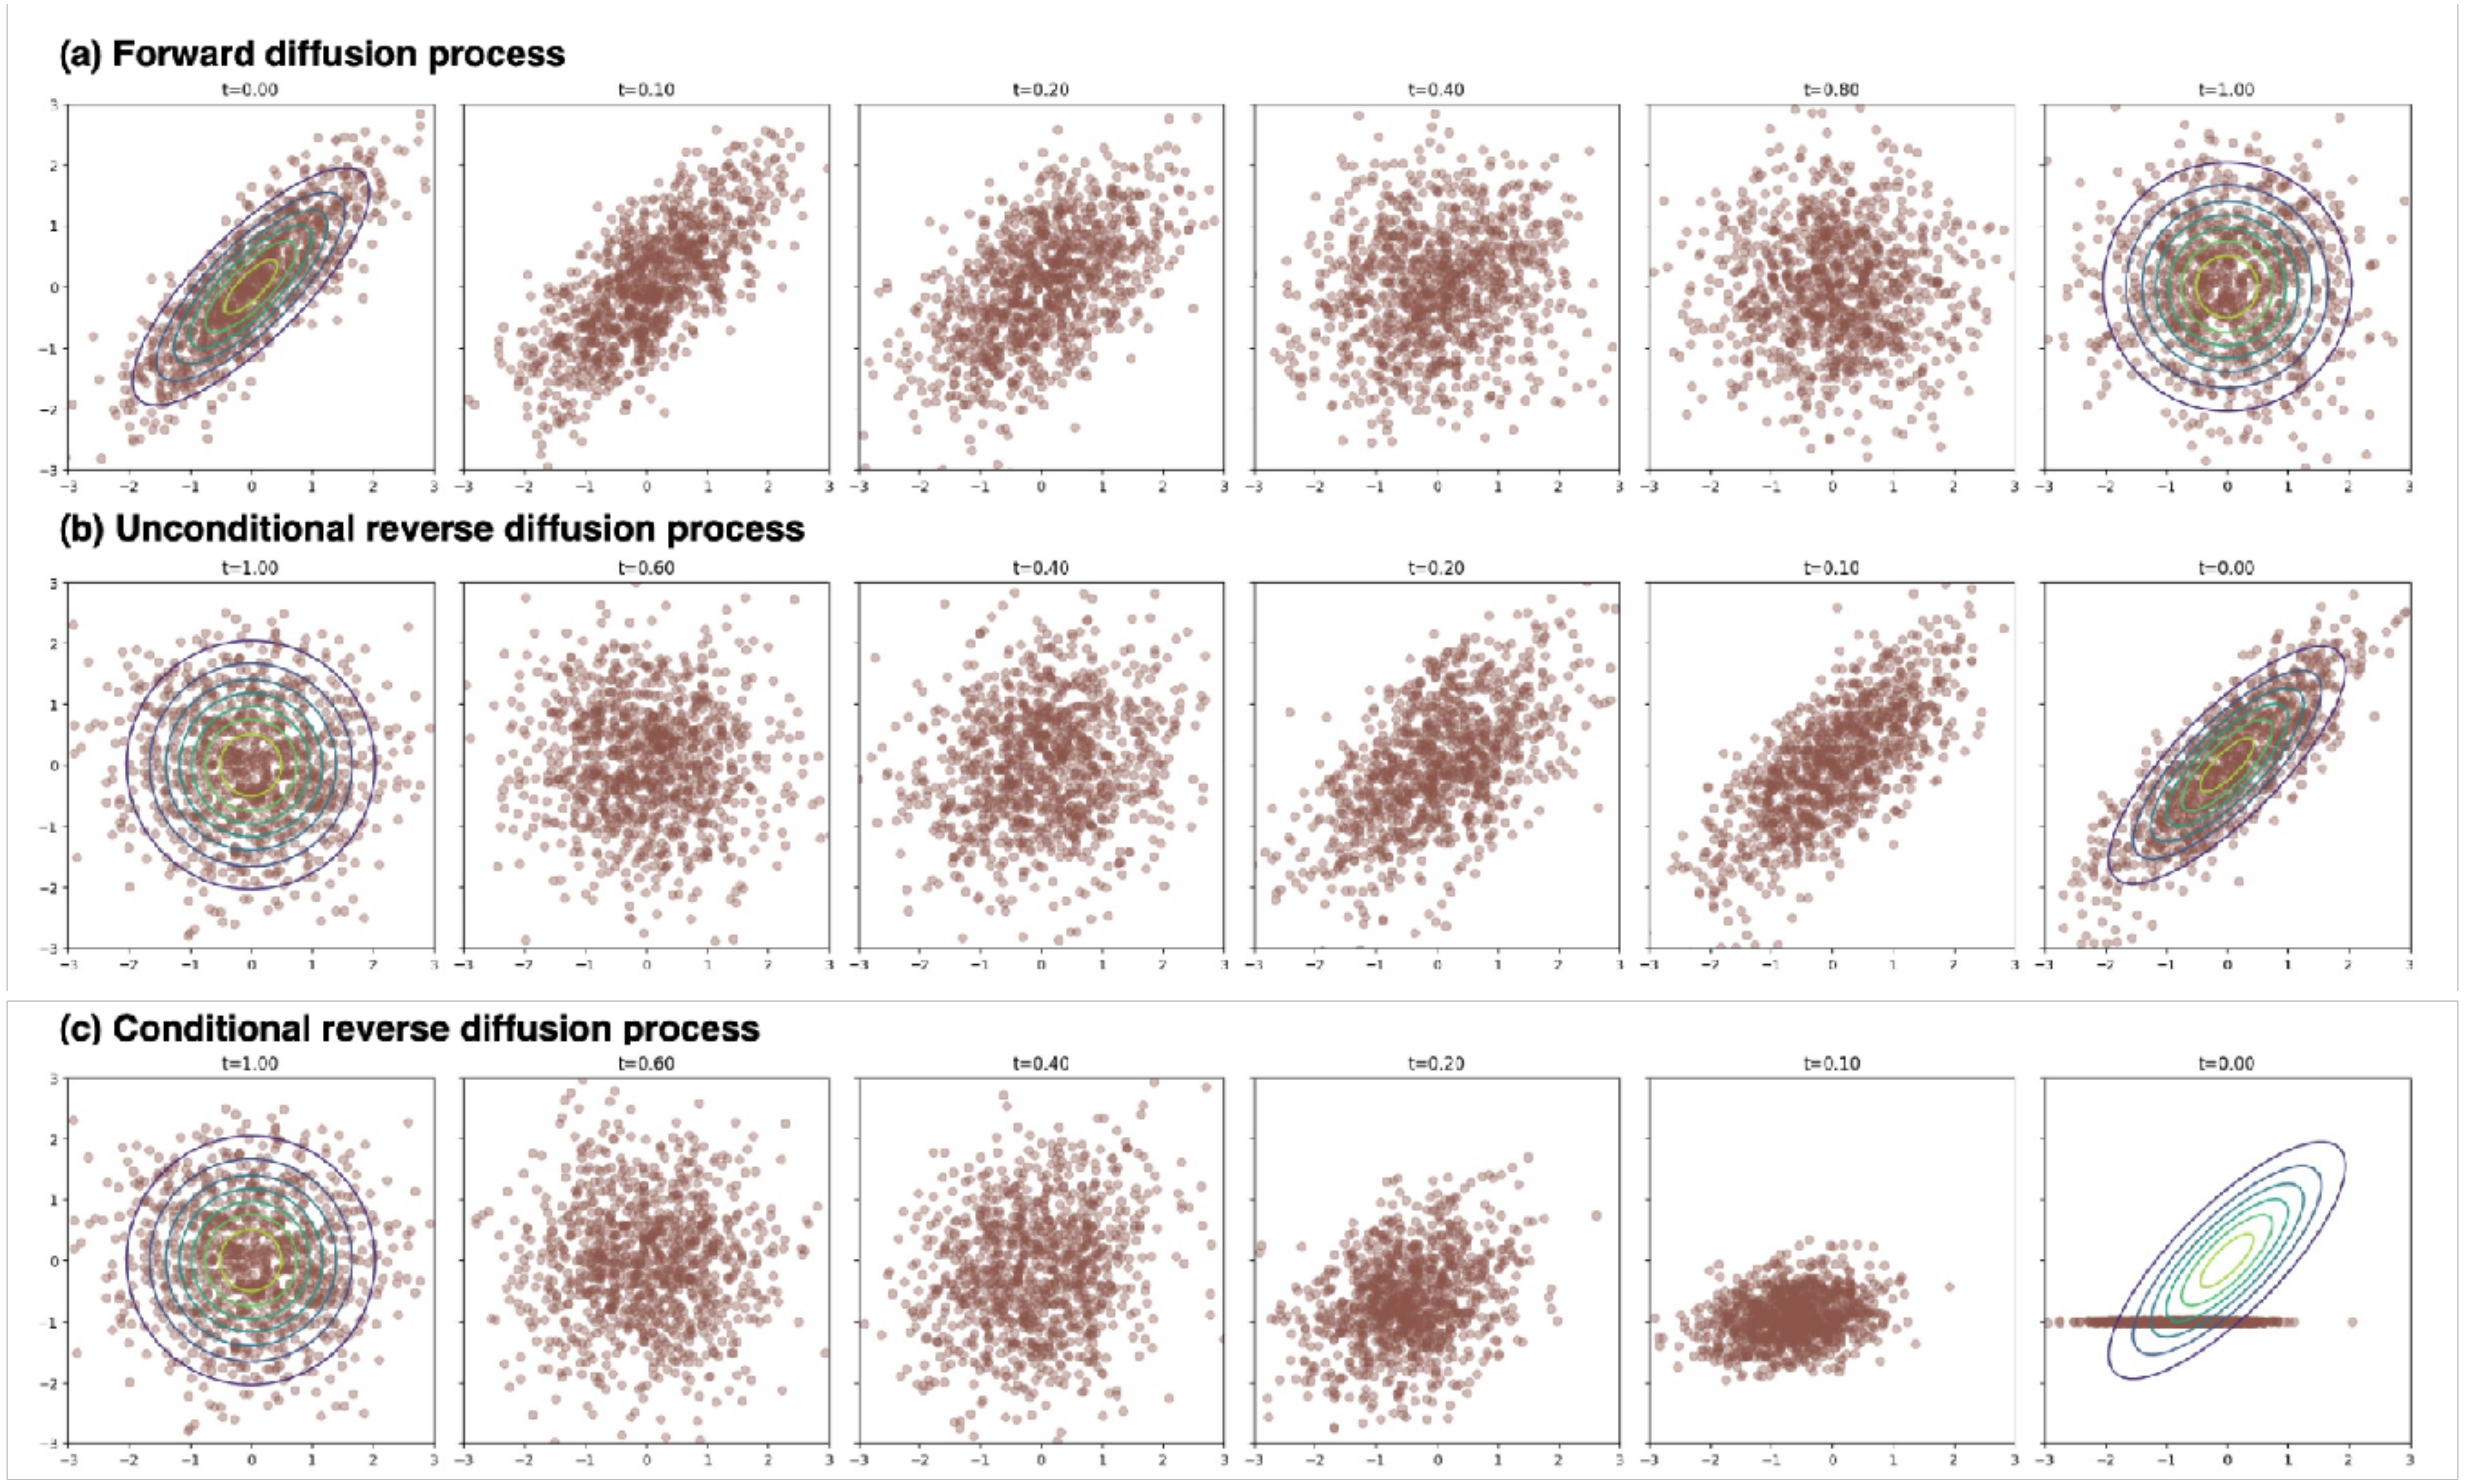

Conditioning of SDEs

- We want the reverse SDE to hit a certain constraint set $X_0 \in B$ at time $T = 0$
- Applications -
  - Class-conditional sampling

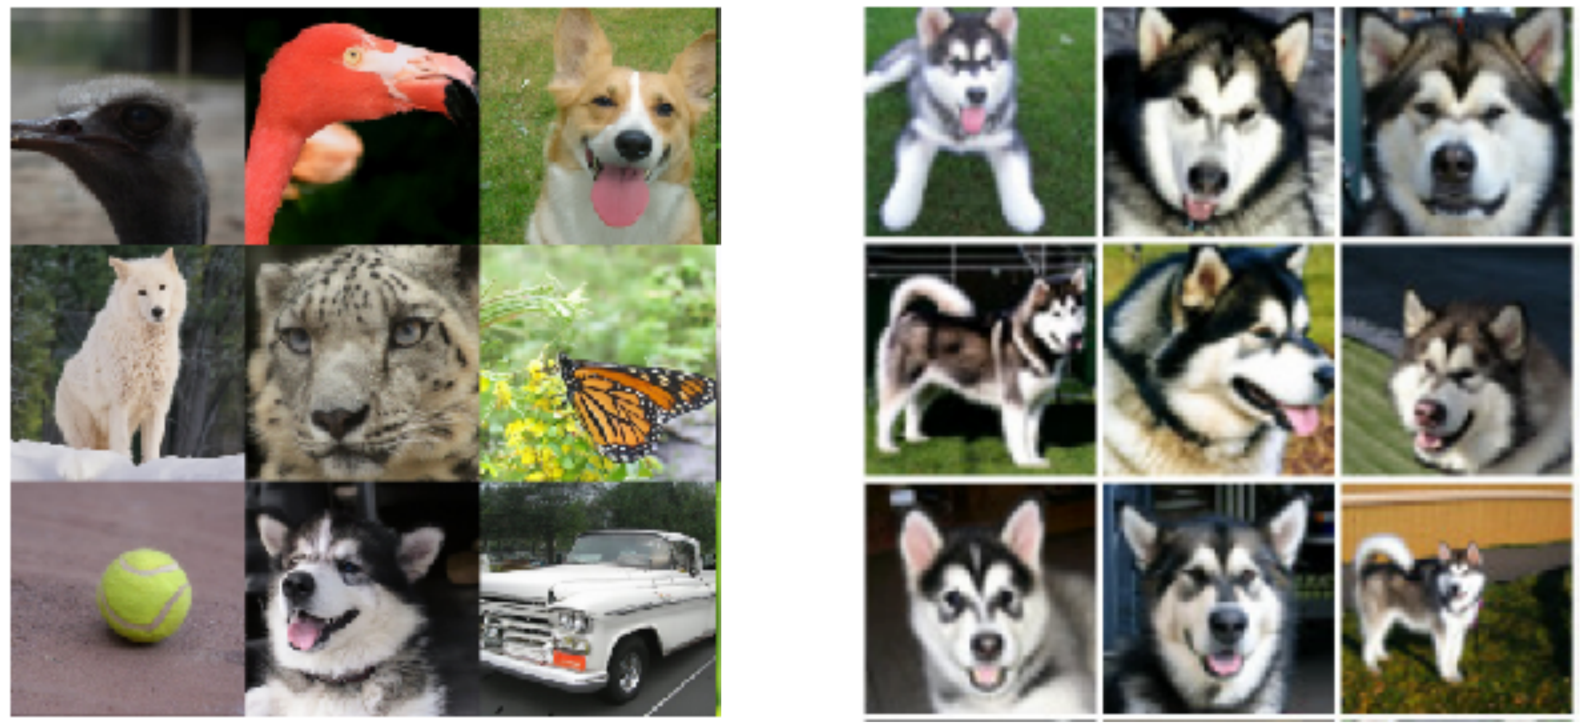

  - Inverse modeling, where we rewrite $X_0 \in B$ to equivalently $\mathscr{A}(X_0) = y$

In-painting

Deblurring

Colorisation

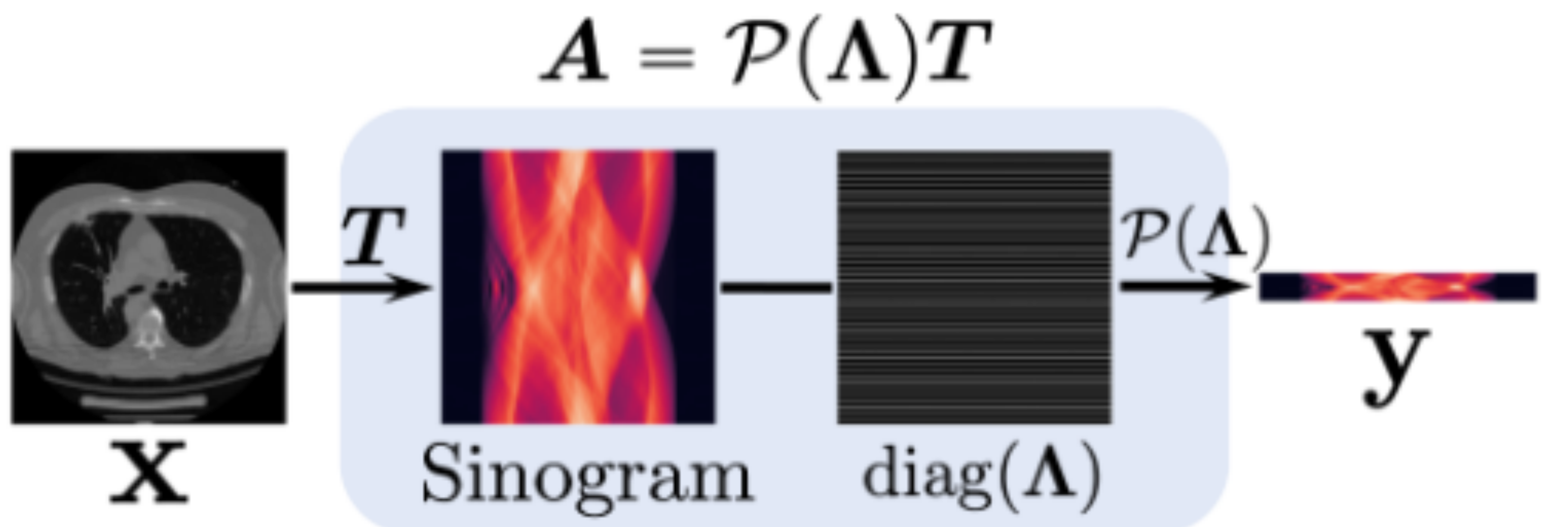

Medical Imaging

Conditioning of SDEs - Doob's *h*-transform

[Rogers and Williams 2000] Consider the forward SDE:

$$\mathrm{dX_t} = \mathrm{f}^{+}\left(\mathrm{X_t}, \mathrm{t}\right)\mathrm{dt} + \sigma_t \mathrm{dW_t}$$

where time flows forwards and with transition densities $p_{t|s}$. It then follows that the conditioned process $\mathrm{X}_t \mid \mathrm{X}_T \in B$ is a solution of

$$\mathrm{dH}_t = \left(f_t^{+}\left(\mathrm{H}_t, t\right) + \sigma_t^2 \nabla_{\mathrm{H}_t} \ln P_{T|t}\left(\mathrm{X}_T \in B \mid \mathrm{H}_t\right)\right)\mathrm{d}t + \sigma_t \mathrm{dW}_t,$$

such that $\mathrm{Law}\left(\mathrm{H}_t \mid \mathrm{H}_s\right) = \vec{p}_{t|s,T}\left(\mathrm{h}_t \mid \mathrm{h}_s, \mathrm{x}_T \in B\right)$ and $\mathbb{P}\left(\mathrm{H}_T \in B\right) = 1$.

- By conditioning a diffusion process to hit a particular event $\mathrm{X}_T \in B$, we get a resulting conditional process that is also an SDE with an additional drift term.
- $h\left(t, \mathrm{H}_t\right) \triangleq P_{T|t}\left(\mathrm{X}_T \in B \mid \mathrm{H}_t\right)$ is known as the ***h*-transform**

Aside - FPK and Infinitesimal Generators

- The FPK, a PDE that describes the evolution of the probability density of an SDE
  - $\mathrm{dX_t} = \mu\left(\mathrm{X_t}, \mathrm{t}\right)\mathrm{dt} + \sigma\left(\mathrm{X_t}, \mathrm{t}\right)\mathrm{dW_t}$
  - $\frac{\partial}{\partial t}p(x, t) = -\frac{\partial}{\partial x}[\mu(x, t)p(x, t)] + \frac{\partial^2}{\partial x^2}[D(x, t)p(x, t)]$
- Let us define an operator called the (infinitesimal) generator of a stochastic process
  - $\mathscr{A}\phi(x_t, t) = \lim_{s\to 0}\frac{\mathbb{E}\left[\phi(x_{t+s}, t+s)\right] - \phi(x_t, t)}{s}$
- Why is it called the generator?
  - $\mathbb{E}\left[\phi(x_{t+s}, t+s)\right] \simeq \phi(x_t, t) + s\mathscr{A}\phi(x_t, t)$

Aside - FPK and Infinitesimal Generators

$$\mathscr{A}\phi(x_t, t) = \lim_{s\to 0} \frac{\mathbb{E}\left[\phi(x_{t+s}, t+s)\right] - \phi(x_t, t)}{s}$$

- For the Ito process $\mathrm{dX_t} = \mu\left(\mathrm{X_t, t}\right)\mathrm{dt} + \sigma\left(\mathrm{X_t, t}\right)\mathrm{dW_t}$, the generator is given by

$$\mathscr{A}(\,\bullet\,) = \sum_i \frac{\partial(\,\bullet\,)}{\partial x_i} f_i(x_t, t) + \frac{1}{2}\sum_{i,j}\left(\frac{\partial^2(\,\bullet\,)}{\partial x_i \partial x_j}\right)\left[\sigma^2(X_t, t)\right]_{ij}$$

- We can show that the transition densities of the SDE satisfy the backwards Kolmogorov equation defined by the generator acting on the transition density

$$-\frac{\partial p(x_{t+s}, t+s \mid x_t, t)}{\partial t} = \mathscr{A}p(x_{t+s}, t+s \mid x_t, t)$$

Doob's h-transform (Quick Proof)

- Given the SDE with transition density $p_{t|s}(x_t \mid x_s)$, where $t > s$,
- We wish to obtain the process after conditioning the SDE to hit a deterministic end point z at time T,

$$p_{t|s,T}\left(x_t \mid x_s, x_T = z\right) = \frac{p_{t|s,T}\left(x_T = z \mid x_t\right) p_{t|s}\left(x_t \mid x_s\right)}{p\left(x_T = z \mid x_s\right)}$$

- Let us try to find the SDE that would have this transition density

Doob's h-transform (Quick Proof)

- Apply Bayes' rule

$$p_{t|s,T}\left(x_t \mid x_s, x_T = z\right) = \frac{p_{T|s,t}\left(x_T = z \mid x_s, x_t\right) p_{t|s}\left(x_t \mid x_s\right)}{p_{T|s}\left(x_T = z \mid x_s\right)}$$

- Now, apply the Markov property $p_{T|s,t}\left(x_T = z \mid x_s, x_t\right) \to p_{T|s,t}\left(x_T = z \mid x_t\right)$,

$$p_{t|s,T}\left(x_t \mid x_s, x_T = z\right) = \frac{p_{T|t}\left(x_T = z \mid x_t\right) p_{t|s}\left(x_t \mid x_s\right)}{p_{T|s}\left(x_T = z \mid x_s\right)}$$

- Now, define $h(x_t, t) = p_{T|t}(x_T = z \mid x_t)$, and replace above to get

$$p_{t|s,T}\left(x_t \mid x_s, x_T = z\right) = \frac{h(x_t, t) p_{t|s}\left(x_t \mid x_s\right)}{h(x_s, s)}$$

Doob's *h*-transform (Quick Proof)

- In order for this to be a valid Markov kernel, we require

$$\int_{x_t} p_{t|s,T}\left(x_t \mid x_s, x_T = z\right) \mathrm{d}x_t = 1$$

$$\int_{x_t} \frac{h(x_t, t) p_{t|s}\left(x_t \mid x_s\right)}{h(x_s, s)} \mathrm{d}x_t = 1$$

- Taking $h(x_s, s)$ to the RHS (independent of x_t), we get the following property that $h(x_s, s)$ satisfies
  - $h(x_s, s) = \int_{x_t} h(x_t, t) p_{t|s}\left(x_t \mid x_s\right) = \mathbb{E}\left[h(x_t, t)\right]$
- From the definition for infinitesimal generator $\mathscr{A}$, we have (assuming $(t, t+s)$ instead of (s, t))
  - $\mathscr{A}\phi(\mathbf{x}, t) = \lim_{s \downarrow 0} \frac{\mathrm{E}[\phi(\mathbf{x}(t+s), t+s)] - \phi(\mathbf{x}(t), t)}{s}$
- Therefore, $\mathscr{A}h(x_t, t) = 0$

Doob's h-transform (Quick Proof)

- $\mathscr{A}h(x_t, t) = 0$, and the new transition probability $p^h = p\dfrac{h(x_{t+s}, t+s)}{h(x_t, t)}$
- Writing down the infinitesimal generator for some arbitrary ϕ, $\mathbb{E}$ w.r.t p^h, we get

$$\mathscr{A}\phi = \lim_{s \to 0} \frac{\mathbb{E}^h[\phi(x_{t+s}, t+s)] - \phi(x_t, t)}{s}$$

- Noting that $\mathbb{E}^h[\phi] = \mathbb{E}\left[\phi\dfrac{h(x_{t+s}, t+s)}{h(x_t, t)}\right]$, we have

$$\mathscr{A}^h\phi = \lim_{s \downarrow 0} \frac{\mathrm{E}[\phi(x_{t+s})h(x_{t+s}, t+s)] - \phi(x_t)h(x_t, t)}{sh(x_t, t)}$$

$$= \frac{1}{h(x_t, t)}\mathscr{A}\{h(x_t, t))\phi(x_t)\}$$

- Now, remember the form of $\mathscr{A}$ below, and apply the product rule to $h(\,.\,)\phi(\,.\,)$

$$\mathscr{A}(\bullet) = \sum_i \frac{\partial(\bullet)}{\partial x_i} f_i(x_t, t) + \frac{1}{2}\sum_{i,j}\left(\frac{\partial^2(\bullet)}{\partial x_i \partial x_j}\right)\left[\sigma^2(X_t, t)\right]_{ij}$$

Doob's h-transform (Quick Proof)

- Now, remember the form of $\mathscr{A}$ below, and apply the product rule to $h(\,.\,)\phi(\,.\,)$

$$\mathscr{A}(\bullet) = \sum_i \frac{\partial(\bullet)}{\partial x_i} f_i(x_t, t) + \frac{1}{2}\sum_{i,j}\left(\frac{\partial^2(\bullet)}{\partial x_i \partial x_j}\right)\left[\sigma^2(X_t, t)\right]_{ij}$$

- We get this behemoth

$$\mathscr{A}^h\phi = \frac{1}{h(t,x_t)}\left\{\frac{\partial h(t,x_t)}{\partial t}\phi + \sum_i\left[\frac{\partial h(t,x_t)}{\partial x_i}\phi(x_t) + h(t,x_t)\frac{\partial \phi(x_t)}{\partial x_i}\right] f_i(x_t,t) + \frac{1}{2}\sum_{i,j}\frac{\partial^2[h(t,x_t)\phi(x_t)]}{\partial x_i \partial x_j}\left[\sigma^2(x_t,t)\right]_{ij}\right\}$$

- This gives us an even bigger behemoth

$\mathscr{A}h = 0$

$$= \frac{1}{h(t,x_t)}\left\{\left[\frac{\partial h(t,x_t)}{\partial t} + \sum_i \frac{\partial h(t,x_t)}{\partial x_i} f_i(x_t,t) + \frac{1}{2}\sum_{i,j}\frac{\partial^2 h(t,x_t)}{\partial x_i \partial x_j}\left[\sigma^2(x_t,t)\right]_{ij}\right]\phi(x_t) + \sum_i h(t,x_t)\frac{\partial \phi(x_t)}{\partial x_i} f_i(x_t,t) + \right.$$

- $$\left. \frac{1}{2}\sum_{i,j}\left[\frac{\partial h(t,x_t)}{\partial x_j}\frac{\partial \phi(x_t)}{\partial x_i} + \frac{\partial h(t,x_t)}{\partial x_i}\frac{\partial \phi(x_t)}{\partial x_j} + h(t,x_t)\frac{\partial^2 \phi(x_t)}{\partial x_i \partial x_j}\right]\left[\sigma^2(x_t,t)\right]_{ij}\right\}$$

Doob's h-transform (Quick Proof)

- Finally, arranging terms we get

$$= \sum_i \left[f_i(x_t, t) + \sigma^2(x_t, t) \frac{\nabla h(t, x_t)}{h(t, x_t)} \right] \frac{\partial \phi(x_t)}{\partial x_i} + \frac{1}{2} \sum_{i,j} \frac{\partial^2 \phi(x_t)}{\partial x_i \partial x_j} \left[\sigma^2(x_t, t)\right]_{ij}$$

- Comparing to $\mathscr{A}(\bullet) = \sum_i \frac{\partial(\bullet)}{\partial x_i} f_i(x_t, t) + \frac{1}{2} \sum_{i,j} \left(\frac{\partial^2(\bullet)}{\partial x_i \partial x_j} \right) \left[\sigma^2(X_t, t)\right]_{ij}$

- The new drift is $f_i(x_t, t) + \sigma^2(x_t, t) \frac{\nabla h(t, x_t)}{h(t, x_t)} = f_i(x_t, t) + \sigma^2(x_i, t) \nabla \log h(x_t, t)$

Example: Pinned Brownian Motion

- Consider a Brownian motion starting from arbitrary $X_0 \sim p_{\text{data}}$, and $\mathrm{d}X_t = \sigma \mathrm{d}W_t$
- Let us condition this Brownian motion to hit $X_T = 0$ at time T,
  - Therefore $h(x_t, t) = P(X_T = 0 \mid X_t) = \mathcal{N}\left(X_t, \sigma(T-t)\right)$
  - Which gives us $\nabla_{X_t} \log h(X_t, t) = \nabla_{X_t}\left(C - \dfrac{(X_t - 0)^2}{2\sigma^2(T-t)} \right) = -\dfrac{X_t}{\sigma^2(T-t)}$
- Plugging into the conditioned process formula, we get
  - $\mathrm{dH}_t = \left(f_t^+\left(\mathrm{H}_t, t\right) + \sigma_t^2 \nabla_{\mathrm{H}_t} \ln P_{T|t}\left(\mathrm{X}_T \in B \mid \mathrm{H}_t\right)\right) \mathrm{d}t + \sigma_t \mathrm{dW}_t,$
  - $\mathrm{dH}_t = -\dfrac{X_t}{T-t} + \sigma \mathrm{dW}_t$

Doob's *h*-transform on the reverse SDE

[Rogers and Williams 2000] Consider the reverse SDE:

$$\mathrm{d}\mathrm{X}_\mathrm{t} = \mathrm{f}^{-}\left(\mathrm{X}_\mathrm{t}, \mathrm{t}\right)\mathrm{d}\mathrm{t} + \sigma_t \mathrm{d}\overline{\mathrm{W}}_\mathrm{t}$$

where time flows backwards and with transition densities $\bar{p}_{t|s}$. It then follows that the conditioned process $\mathrm{X}_t \mid \mathrm{X}_0 \in B$ is a solution of

$$\mathrm{d}\mathrm{H}_t = \left(f_t^{-}\left(\mathrm{H}_t, t\right) - \sigma_t^2 \nabla_{\mathrm{H}_t} \ln \bar{P}_{0|t}\left(\mathrm{X}_0 \in B \mid \mathrm{H}_t\right)\right)\mathrm{d}t + \sigma_t \bar{\mathrm{d}}_t \bar{\mathrm{W}}_t,$$

such that $\mathrm{Law}\left(\mathrm{H}_s \mid \mathrm{H}_t\right) = \vec{p}_{s|t,0}\left(\mathrm{h}_s \mid \mathrm{h}_t, \mathrm{x}_0 \in B\right)$ and $\mathbb{P}\left(\mathrm{H}_0 \in B\right) = 1$.

Specifying hard constraints

- We consider events of the form $X_0 \in B$ that can be described by
  - An equality constraint $\mathscr{A}(X_0) = y$
- Consider Doob's h-transform with hard constraints
  - $\mathrm{dH}_t = \left(f_t^-\left(\mathrm{H}_t, t\right) - \sigma_t^2 \nabla_{\mathrm{H}_t} \ln \bar{P}_{0|t}\left(\mathscr{A}(X_0) = y \mid \mathrm{H}_t\right)\right) \mathrm{d}t + \sigma_t \bar{\mathrm{d}}_t \bar{\mathrm{W}}_t,$
- Replacing $f_t^-(H_t, t) = f^+(H_t, t) - \sigma^2 \nabla_{H_t} \ln P_t(H_t)$, we get
  - $\mathrm{dH}_t = f_t^+\left(\mathrm{H}_t, t\right) - \sigma_t^2 \left(\nabla_{H_t} \ln P_t(H_t) + \nabla_{\mathrm{H}_t} \ln \bar{P}_{0|t}\left(\mathscr{A}(X_0) = y \mid \mathrm{H}_t\right)\right) \mathrm{d}t + \sigma_t \bar{\mathrm{d}}_t \bar{\mathrm{W}}_t,$
- Sampling from this SDE gives us $x \sim p_{\text{data}}$ that satisfy $\mathscr{A}(x) = y$

Known unconditional score

Unknown conditional score

$\mathscr{A} = \{0,1\}^{H \times W \times C}$

In-painting

Specifying soft constraints

- We consider events of the form $X_0 \in B$ that can be described by
    - Noisy observations $y = \mathscr{A}(X_0) + \eta$, and a density $p(y \mid X_0)$
        - We wish to recover the posterior $p(X_0 \mid y)$

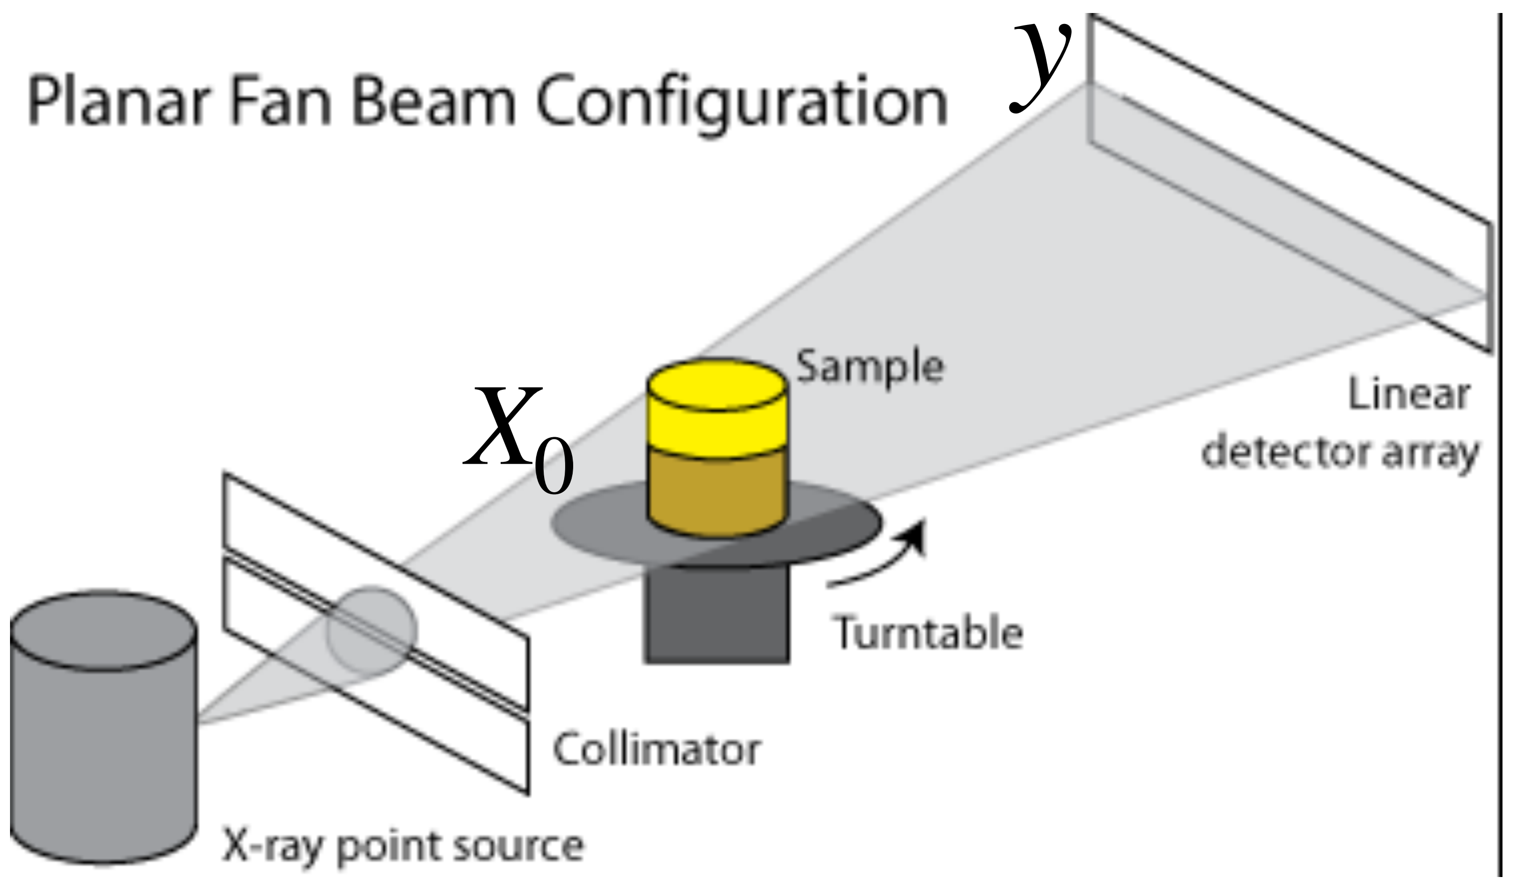

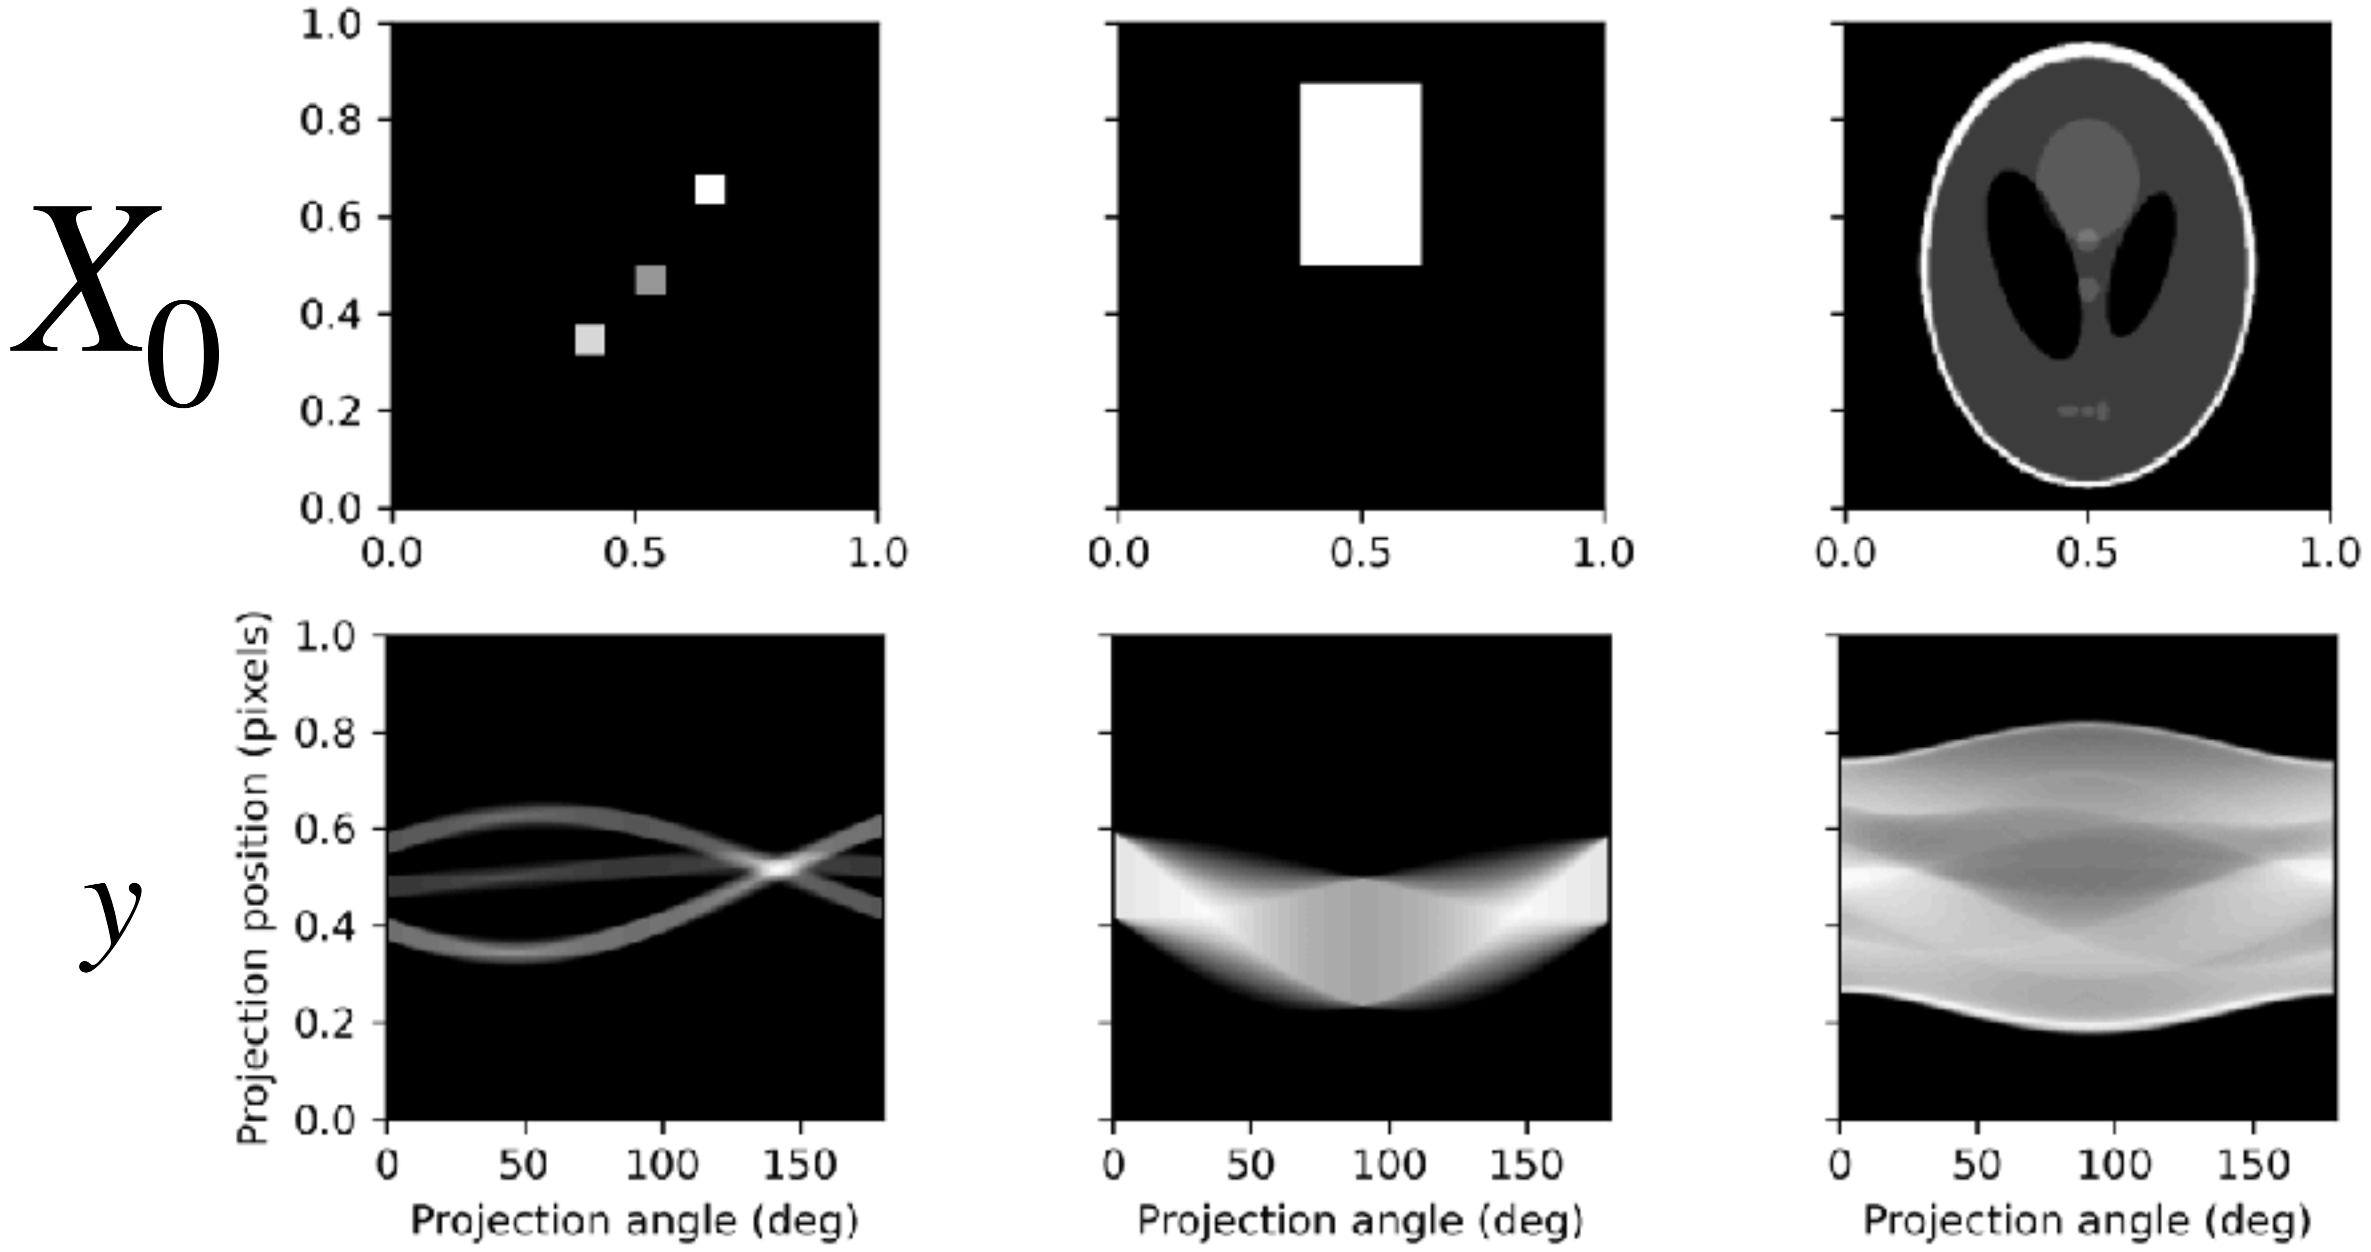

Specifying soft constraints

- We consider events of the form $X_0 \in B$ that can be described by
    - Noisy observations $y = \mathscr{A}(X_0) + \eta$, and a density $p(y \mid X_0)$
        - We wish to recover the posterior $p(X_0 \mid y)$

- Consider Doob's h-transform with soft constraints
    - $\mathrm{dH}_t = \left(f_t^-\left(\mathrm{H}_t, t\right) - \sigma_t^2 \nabla_{\mathrm{H}_t} \ln \bar{P}_{y|t}\left(Y = y \mid \mathrm{H}_t\right)\right)\mathrm{d}t + \sigma_t \bar{\mathrm{d}}_t \bar{\mathrm{W}}_t,$
- Again, replacing $f_t^-(H_t, t) = f_t^+(H_t, t) - \sigma^2 \nabla_{H_t} \ln P_t(H_t)$, we get
    - $\mathrm{dH}_t = f_t^+\left(\mathrm{H}_t, t\right) - \sigma_t^2\left(\nabla_{H_t} \ln P_t(H_t) + \nabla_{\mathrm{H}_t} \ln \bar{P}_{y|t}\left(Y = y \mid \mathrm{H}_t\right)\right)\mathrm{d}t + \sigma_t \bar{\mathrm{d}}_t \bar{\mathrm{W}}_t,$
- Sampling from this SDE gives us $x \sim p(x_0 \mid Y = y) = \dfrac{p(y \mid x_0) p_{\text{data}}(x_0)}{p(y)}$

Different approaches for conditioning

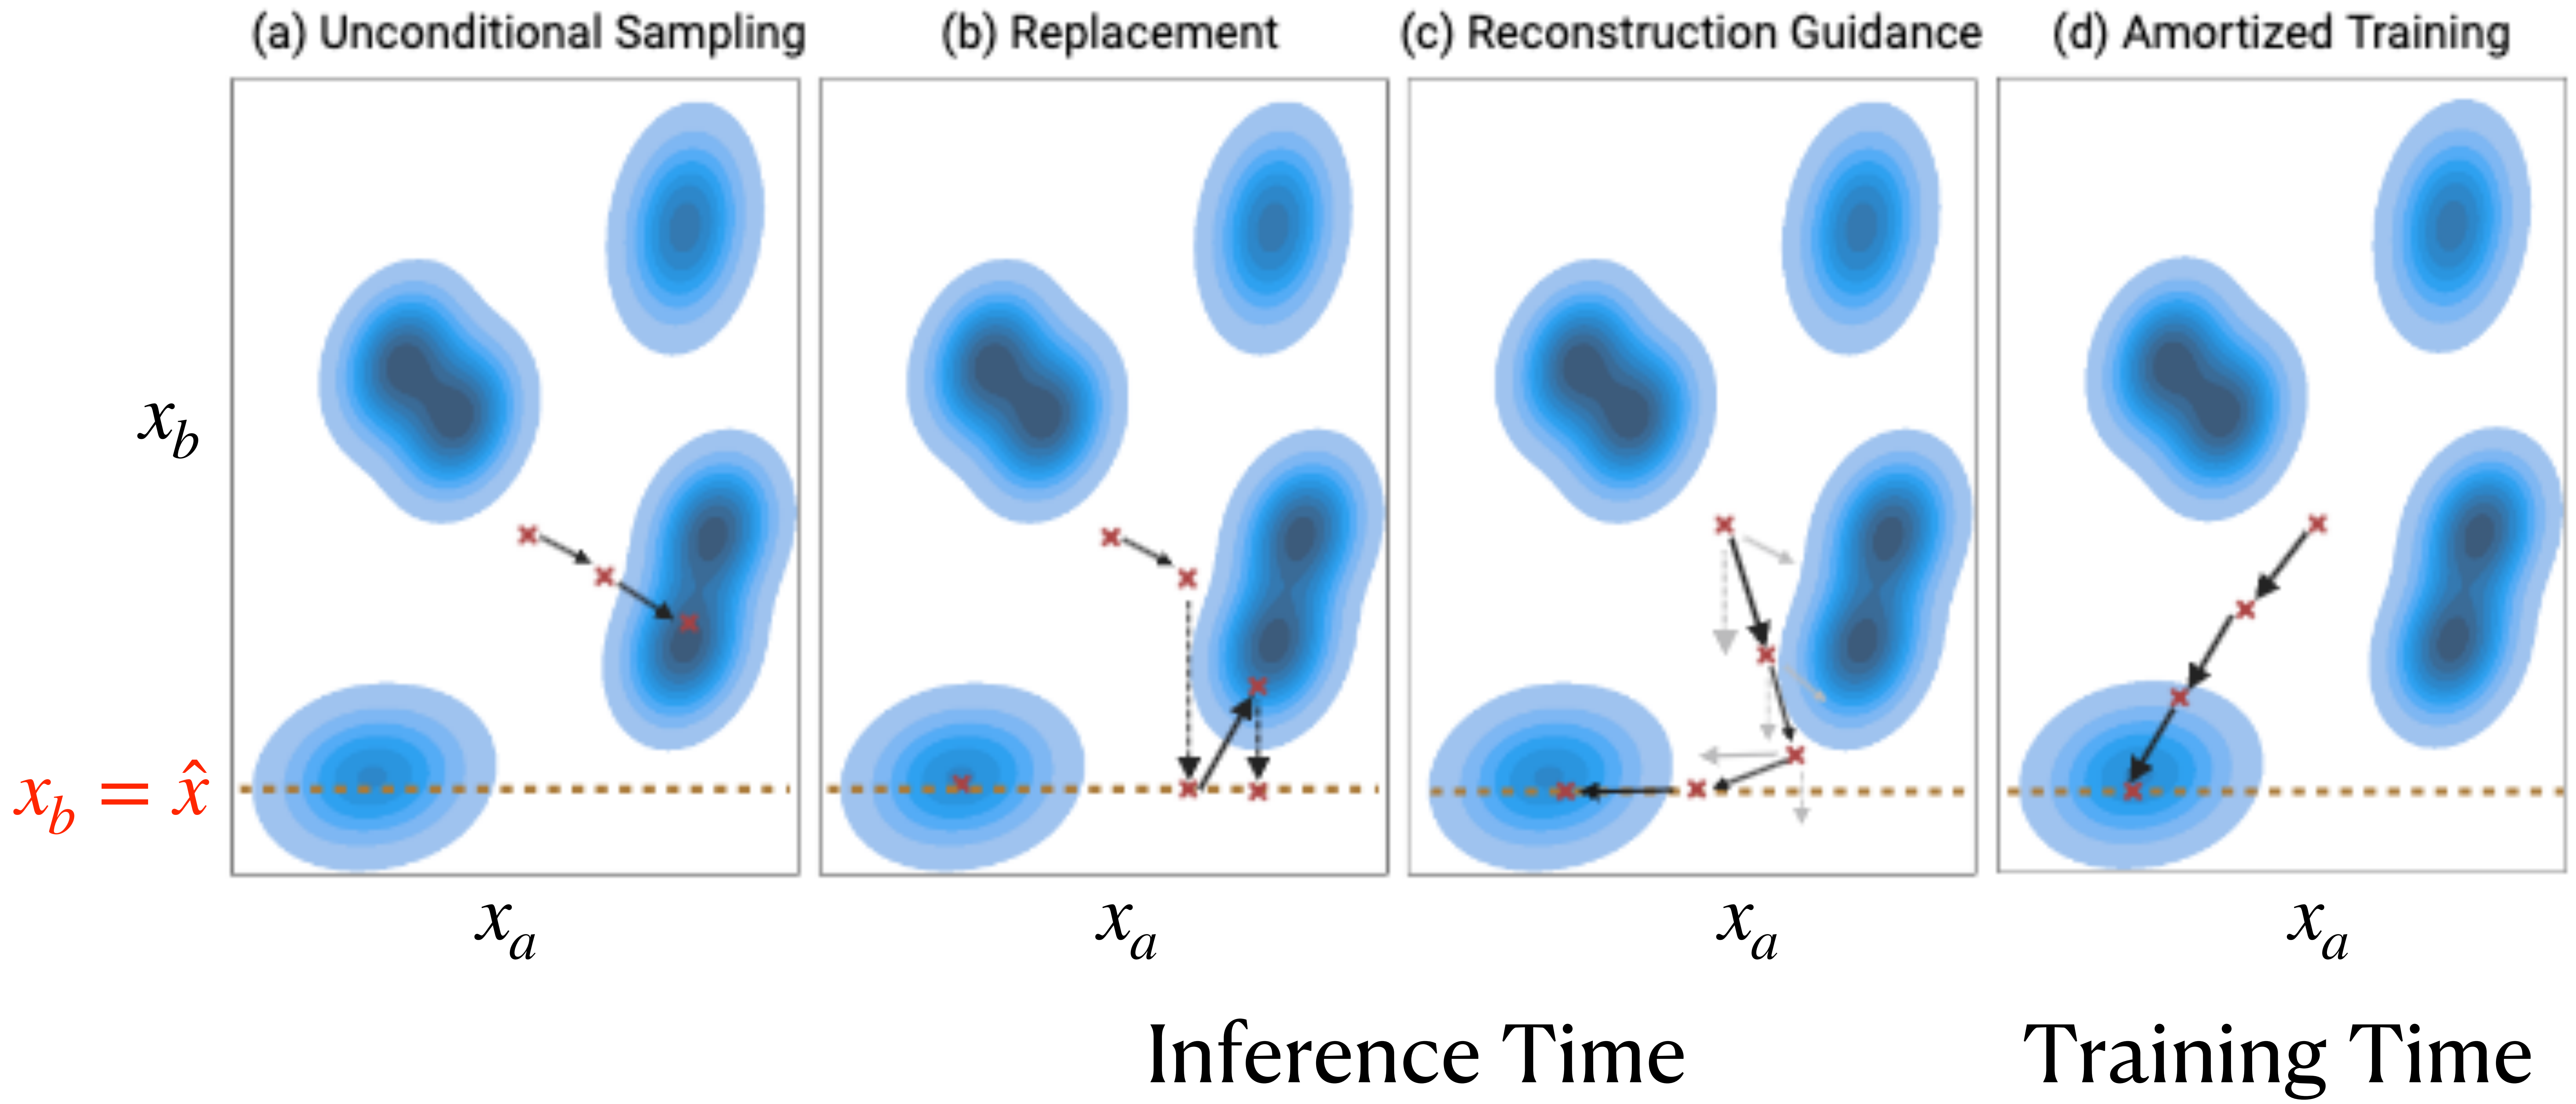

How do we sample from conditioned SDE?

- We have
$$\mathrm{dH}_t = f_t^+\left(\mathrm{H}_t, t\right) - \sigma_t^2\left(\nabla_{H_t} \ln P_t(H_t) + \nabla_{\mathrm{H}_t} \ln \bar{P}_{y|t}\left(Y = y \mid \mathrm{H}_t\right)\right)\mathrm{d}t + \sigma_t \bar{\mathrm{d}}_t \bar{\mathrm{W}}_t,$$

- The challenge is $\bar{P}_{y|t}\left(Y = y \mid \mathrm{H}_t\right) = \int p(y \mid x_0)\bar{p}_{0|t}\left(x_0 \mid \mathrm{H}_t\right)\mathrm{d}x_0$

- We can sample from the reverse SDE to obtain samples from $p(x_0 \mid H_t)$
- We cannot estimate the integral without *many* samples, or evaluate at fixed x_0
- Can we approximate $p(x_0 \mid H_t)$?

Reconstruction Guidance

- Approximate $p(x_0 \mid x_t)$ with a Gaussian approximation using **Tweedie's formula**

[**Tweedie's formula**]. Let $p(a \mid b) = p_0(a)\exp(b^\top T(a) - \psi(b))$. Then the posterior mean $\hat{b} = \mathbb{E}[b \mid a]$ should satisfy

$$\left(\nabla_a T(a)\right)^\top \hat{b} = \nabla_a \log p(a) - \nabla_a \log p_0(a)$$

- **Diffusion Posterior Sampling** [Chung et al 2023]: Forward process is VP-SDE or DDPM sampling, approximate posterior mean using Tweedie's formula:

$$\hat{x}_0 = \mathbb{E}[x_0 \mid x_t] = \frac{1}{\sqrt{\bar{\alpha}(t)}}\left(x_t + (1 - \bar{\alpha}(t))\nabla_{x_t} \log p_t(x_t)\right)$$

- Then approximate

$$\bar{P}_{y|t}\left(Y = y \mid \mathrm{x}_t\right) = \int p(y \mid x_0)\bar{p}_{0|t}\left(x_0 \mid \mathrm{x}_t\right) \mathrm{d}x_0 \simeq p(y \mid \hat{x}_0)$$

- Simulate the following reverse SDE
  - $\mathrm{dH}_t = f_t^+\left(\mathrm{H}_t, t\right) - \sigma_t^2\left(\nabla_{H_t} \ln P_t(H_t) + \nabla_{\mathrm{H}_t} \ln p\left(y \mid \hat{x}_0\right)\right)\mathrm{d}t + \sigma_t \mathrm{d}\bar{\mathrm{W}}_t,$

Known unconditional score

Extra backprop through score network

Classifier Guidance

- When we have external labels to guide diffusion
- Consider $\mathrm{dX}_t = f_t^+(\mathrm{X}_t, t) - \sigma_t^2 \left(\nabla_{X_t} \ln P_t(X_t) + \nabla_{X_t} \ln P\left(Y = y \mid \mathrm{X}_t\right) \right) \mathrm{d}t + \sigma_t \mathrm{d}\bar{\mathrm{W}}_t,$
  - We can train a time-dependent classifier by creating training data
    - Sample (x_0, y) from dataset, then sample $x_t \sim$ SDE to obtain (x_t, y)
    - Using trained $p(y \mid x_t)$ classifier, we backprop to get $\nabla_{x_t} \ln p(y \mid x_t)$
    - $\nabla_{x_t} \ln p(x_t)$ is unconditional score model
    - We can overweight "guidance", $\nabla_{X_t} \ln P_t(X_t) + \gamma \nabla_{X_t} \ln P\left(y \mid \mathrm{X}_t\right)$

Classifier Free Guidance

- What if we don't want to train a separate classifier model?
- Consider $\mathrm{dX}_t = f_t^+\left(\mathrm{X}_t, t\right) - \sigma_t^2\left(\nabla_{X_t}\ln P_t(X_t) + \nabla_{X_t}\ln P\left(Y = y \mid \mathrm{X}_t\right)\right)\mathrm{d}t + \sigma_t \mathrm{d}\bar{\mathrm{W}}_t,$
  - From Bayes' rule, $\nabla_{X_t}\ln p(X_t \mid y) = \nabla_{X_t}\ln p(X_t) + \nabla_{X_t}\ln p(y \mid X_t)$
  - Train a single model to approximate both $p(X_t \mid y)$ and $p(X_t)$ on (X_0, y) inputs
  - Occasionally, 10-20% of the time, drop y to learn $p(X_t)$, otherwise $p(X_t \mid y)$
  - At inference time, $\nabla_x \log p_\gamma(x \mid y) = \nabla_x \log p(x) + \gamma\left(\nabla_x \log p(x \mid y) - \nabla_x \log p(x)\right)$

Amortised training of h-transform

- Can we learn a score model for the h-transform directly?
- Consider $\mathrm{dX}_t = f_t^+\left(\mathrm{X}_t, t\right) - \sigma_t^2\left(\nabla_{X_t}\ln P(X_t \mid y)\right)\mathrm{d}t + \sigma_t \mathrm{d}\bar{\mathrm{W}}_t,$
- Let us learn a score model $f(t, X_t, y, \mathscr{A})$ to approximate $\nabla_{X_t}\ln p(X_t \mid y)$

$$f^* = \arg\min_f \mathbb{E}_{Y\sim p_{|\mathscr{A},X_0},\mathscr{A}\sim p, X_0\sim p_{\text{data}}}\left[\int_0^T \left\| f\left(t, X_t, y, \mathscr{A}\right) - \nabla_{X_t}\ln \vec{p}_{t|0}\left(X_t \mid X_0\right)\right\|^2 \mathrm{d}t\right]$$

- Sample $X_0 \sim p_{\text{data}}$, sample an operator $\mathscr{A}$, sample $y \sim p(y \mid X_0, \mathscr{A})$
- We “amortise” the score model over y and $\mathscr{A}$
- The minimiser is given by the conditional score

$$\mathrm{f}^*\left(\mathrm{t}, \mathrm{X}_\mathrm{t}, \mathrm{y}, \mathrm{A}\right) = \nabla_{\mathrm{x}_\mathrm{t}}\ln \mathrm{p}_{\mathrm{t}|0}\left(\mathrm{X}_\mathrm{t} \mid \mathrm{Y} = \mathrm{y}, \mathscr{A} = \mathrm{A}\right)$$

Proof of minimiser of loss function

$$f\left(x_t, y, A\right) = \mathbb{E}_{X_0|X_t=x_t, Y=y, \mathscr{A}=A}\left[\nabla_{X_t} \ln p_{t|0}\left(X_t \mid X_0\right)\right]$$

$$= \int \nabla_{x_t} \ln p_{t|0}\left(x_t \mid X_0\right) p_{0|t}\left(X_0 \mid x_t, y, A\right) dX_0$$ (Write expectation as integral)

$$= \int \nabla_{x_t} \ln p_{t|0}\left(x_t \mid X_0\right) \frac{p_{t|0}\left(x_t \mid X_0\right) p_0\left(X_0\right)}{p_t\left(x_t \mid y, A\right)} dX_0$$ (Bayes' rule)

$$= \int \frac{\nabla_{x_t} p_{t|0}\left(x_t \mid X_0\right)}{p_{t|0}\left(x_t \mid X_0\right)} \frac{p_{t|0}\left(x_t \mid X_0\right) p_0\left(X_0\right)}{p_t\left(x_t \mid y, A\right)} dX_0$$ (Expand the gradient of log)

$$= \frac{1}{p_t\left(x_t \mid y, A\right)} \int \frac{\nabla_{x_t} p_{t|0}\left(x_t \mid X_0\right)}{p_{t|0}\left(x_t \mid X_0\right)} p_{t|0}\left(x_t \mid X_0\right) p_0\left(X_0\right) dX_0$$ (Cancel terms)

$$= \frac{1}{p_t\left(x_t \mid y, A\right)} \int \nabla_{x_t}\left[p_{t|0}\left(x_t \mid X_0\right)\right] p_t\left(x_t \mid y, A\right) dX_0$$ (Swap integral derivative - Dominated Convergence Theorem)

$$= \frac{1}{p_t\left(x_t \mid y, A\right)} \nabla_{x_t} \int \left[p_{t|0}\left(x_t \mid x_0\right)\right] p_t\left(x_t \mid y, A\right) dX_0$$ (Marginalise over X_0)

$$= \frac{1}{p_t\left(x_t \mid y, A\right)} \nabla_{x_t} p_t\left(x_t \mid y, A\right)$$

$$= \nabla_{x_t} \ln p_t\left(x_t \mid y, A\right)$$