Introduction to flow matching - MIT Class 6.S184

Base distribution $\quad p_0 = \mathcal{N}(\vec{0}, I_d)$

Target dist $\quad \vec{z}^{(1)}, \vec{z}^{(2)}, \ldots, \vec{z}^{(N)} \sim p_{data}$

Probability path $\quad p_t$

Vector field $\quad \vec{u}_t(\vec{x}) = \vec{u}_t(\vec{x}, t) \in \mathbb{R}^d$

Interpolated sample $\quad \vec{x}_t = \vec{x}_0 + \int_0^t u_s(\vec{x}_s)\, ds$

Goal: Find some vector field $\vec{u}_t^{target}$ s.t.

$$\vec{x}_0 \sim p_0$$
$$\Rightarrow \vec{x}_1 \sim p_{data}$$
$$\vec{x}_t \sim p_t$$

Q1: What should be the form of $p_t, \vec{u}_t$?

Q2: How to train a NN u_t^θ to approximate $\vec{u}_t^{target}$?

Q3: How to extend to stochastic dynamics?

Conditional probability path

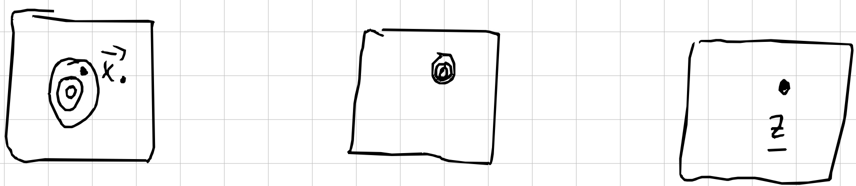

$t=0$ $\qquad$ $t=0.5$ $\qquad$ $t=1$

$p_0(\cdot\,|\,\vec{z}) = p_0$ $\qquad$ $p_t(\cdot\,|\,\vec{z})$?? $\qquad$ $p_1(\cdot\,|\,\vec{z}) = \delta_z$

Think about 1 sample : $\vec{X}_t = \alpha_t \vec{z} + \beta_t \vec{X}_0 \sim p_t(\cdot\,|\,z)$

$$\alpha_0 = 0, \quad \alpha_1 = 1 \quad \alpha_t = t$$

$$\beta_0 = 1, \quad \beta_1 = 0 \quad \beta_t = 1-t$$

$X_{t=0} = X_0 \sim p_0$ $\qquad$ $\mathbb{E}[X_t] = \alpha_t z$

$X_{t=1} = z \sim \delta_z$ $\qquad$ $Var[X_t] = \beta_t^2 I_d$

$$p_t(\cdot\,|\,z) = \mathcal{N}(\alpha_t z, \beta_t^2 I_d)$$

$$p_0(\cdot\,|\,z) = \mathcal{N}(\underline{0}, I_d) \checkmark$$

$$p_1(\cdot\,|\,z) = \mathcal{N}(z, 0) = \delta_z \checkmark$$

Marginal probability path

We can sample from $p_t(x)$

We can evaluate $p_t(x)$:

1 $\boxed{p_t(x) = \int p_t(x|z)\, p_{data}(z)\, dz}$

Intractable $\Rightarrow$ we don't know $p_{data}(z)$

Sampling from $p_t(x)$

$p_t(x) \Rightarrow x_0, z$ $\qquad \mathbb{E}_{x_t \sim p_t}[\dots]$

$p_t(\cdot|z)$ $\qquad = \mathbb{E}_{x_t \sim p_t(\cdot|z),\, t \sim \text{unif},\, z \sim p_{data}}[\dots]$

$\vec{x}_0, \vec{z} \Rightarrow \vec{x}_t = t \cdot \vec{z} + (1-t)\vec{x}_0$

Condirional vector field

$X_t = \alpha_t \vec{z} + \beta_t \vec{X}_0$ $\qquad X_t \equiv X(t)$

$\dot{X}_t = \dot{\alpha}_t \vec{z} + \dot{\beta}_t \vec{X}_0 \equiv u_t^{target}(X_t | z)$

$\vec{u}_t(\vec{x}|\vec{z}) = \vec{z} - \vec{x}_0$

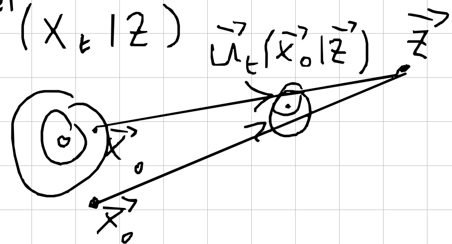

Marginal vector field

$$\text{(2)} \quad \boxed{\vec{u}_t^{target}(x) = \int \underbrace{\vec{u}_t^{target}(\vec{x}|\vec{z})}_{\text{condirional v.f}} \underbrace{\frac{p_t(\vec{x}|\vec{z})\, p_{data}(\vec{z})}{p_t(\vec{x})}}_{\text{Posterior distribution over } z \;\; p_t(\vec{z}|\vec{x})} d\vec{z}}$$

But how do we know this is the correct field

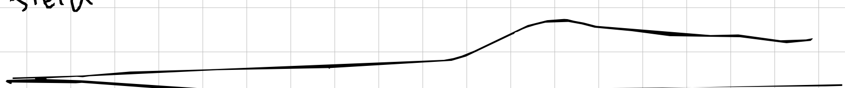

=> Tractable : condirional Prob. Path + vec. field

$p_t(x|z)$, $u_t^{target}(x|z)$ => choose form

=> Intractable : Marginal " + "

$p_t(x)$, $u_t^{target}(x)$ <=> implies form

So far p_t, u_t^{target}

Theorem 1. Q1

If $\frac{d}{dt} X_t = u_t(X_t)$ define according to 2

$\Rightarrow X_t \sim p_t$ samples flow along marginal prob. path defined in 1

Proof:

$$\partial_t p_t(x) = -\nabla \cdot (u_t(x) p_t(x)) \quad (3)$$

inflow prob. mass

$$\text{LHS: } \partial_t p_t(x) \overset{(1)}{=} \partial_t \int p_t(x|z) p_{data}(z) dz$$

$$= \int \partial_t p_t(x|z) p_{data}(z) dz$$

$$\overset{(3)}{=} - \int \nabla \cdot (p_t(x|z) u_t(x|z)) p_{data}(z) dz$$

$$= - \nabla \cdot \left(\int p_t(x|z) u_t(x|z) p_{data}(z) dz \right)$$

$$= - \nabla \cdot \left(p_t(x) \int u_t(x|z) \frac{p_t(x|z) p_{data}(z)}{p_t(x)} dz \right)$$

$$\overset{(2)}{=} - \nabla \cdot \left(p_t(x) u_t(x) \right)$$

Q2: How to train u_t^{θ} to approximate u_t^{target}?

$$\underset{t=0}{X_0 \sim p_0} \xrightarrow[X_t \sim p_t]{u_t^{target}(x)} \underset{t=1}{X_1 \sim p_{data}}$$

<u>Flow matching loss</u>

$$\mathcal{L}_{FM}(\theta) = \mathbb{E}_{x \sim p_t,\, t \sim Unif}\left[\| u_t^{\theta}(x) - u_t^{target}(x) \|^2\right]$$

$\Updownarrow$ equivalent

Conditional FM loss

$$\mathcal{L}_{CFM}(\theta) = \mathbb{E}_{t \sim Unif,\, z \sim p_{data},\, x \sim p_t(\cdot | z)}\left[\| u_t^{\theta}(x) - u_t^{target}(x | z) \|^2\right]$$

Theorem 2.

$$\mathcal{L}_{FM}(\theta) = \mathcal{L}_{CFM}(\theta) + C$$

$$\nabla_\theta \mathcal{L}_{FM}(\theta) = \nabla_\theta \mathcal{L}_{CFM}(\theta)$$

at The minimizer θ^* of $\mathcal{L}_{CFM}(\theta)$, $u_t^{\theta^*} = u_t^{target}$

Proof Thm. 2

$$\mathcal{L}_{FM}(\theta) = \mathbb{E}_{t\sim \text{Unif},\, x\sim p_t}\left[\|u_t^\theta(\vec{x}) - u_t^{target}(\vec{x})\|^2\right]$$

$$= \mathbb{E}_{t,z,x}\left[\|u_t^\theta(x)\|^2 - 2\,\underset{(1)}{u_t^\theta(x)^T u_t^{target}(x)} + \underbrace{\|u_t^{target}(x)\|^2}_{C_1}\right]$$

(1)

$$\mathbb{E}_{t,\,x\sim p_t}\left[u_t^\theta(x)^T u_t^{target}(x)\right]$$

$$= \int_0^1\int u_t^\theta(x)^T u_t^{target}(x)\, p_t(x)\, dt\, dx$$

(2)

$$= \int_0^1\int \cancel{p_t(x)}\, u_t^\theta(x)^T \left[\int u_t^{target}(x|z)\, \frac{p_t(x|z)\, p_{data}(z)}{\cancel{p_t(x)}}\, dz\right] dt\, dx$$

$$= \int_0^1\int\int u_t^\theta(x)^T u_t^{target}(x|z)\, p_t(x|z)\, p_{data}(z)\, dt\, dx\, dz$$

$$= \mathbb{E}_{t\sim\text{Unif},\, z\sim p_{data},\, x\sim p_t(\cdot|z)}\left[u_t^\theta(x)^T u_t^{target}(x|z)\right]$$

$$\mathcal{L}_{FM}(\theta) = \mathbb{E}_{t\sim \text{Unif}, z\sim p_{data}, x\sim p_t(\cdot|z)}\Big[\|u_t^\theta(x)\|^2$$

$$- 2\,\vec{u}_t^{\,\theta}(\vec{x})^T\vec{u}_t^{\,target}(\vec{x}|\underline{z})\Big] + C_1$$

$$= \mathbb{E}_{t, z\sim p_{data}, x\sim p_t(\cdot|z)}\Big[\|u_t^\theta(\underline{x})\|^2 - 2\underline{u}_t^\theta(\underline{x})^T u_t^{target}(\underline{x}|\underline{z})$$

$$+ \|u_t^{target}(\underline{x}|\underline{z})\|^2 - \|u_t^{target}(\underline{x}|\underline{z})\|^2\Big] + C_1$$

$$= \mathbb{E}_{t, z\sim p_{data}, x\sim p_t(\cdot|z)}\Big[\|u_t^\theta(\underline{x}) - u_t^{target}(\underline{x}|\underline{z})\|^2\Big] + C_1 + C_2$$

$$= \mathcal{L}_{CFM}(\theta) + C$$

Summary

① Construct conditional Prob. path $p_t(x|z)$ that satisfies $p_0(x|z) = p_0(x)$, $p_1(x|z) = \delta_z$

② construct conditional v.f

$$\frac{d}{dt}X_t = u_t(x|z) \Rightarrow \begin{matrix}\text{Linear path} \\ \alpha_t = t, \beta_t = 1-t\end{matrix} \quad u_t(x|z) = z - x_0$$

③ $$u_t(x) = \int u_t(x|z)\frac{p_t(x|z)\,p_{data}(z)}{p_t(x)}\,dz$$

satisfies $X_t \sim p_t$ Thm 1.

④ Regress marginal v.f using Thm 2.

$$\mathcal{L}_{FM}(\theta) = \mathbb{E}_{t\sim U[0,1],\, z\sim p_{data},\, x\sim p_t(\cdot|z)}\left[\|u_t^\theta(x) - u_t(x|z)\|^2\right]$$

$$= \mathbb{E}_{t,z,x}\left[\|u_t^\theta(x) - (z - x_0)\|^2\right]$$

Extension to SDEs

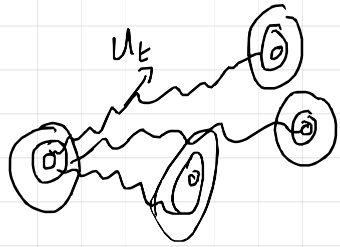

FM: $\frac{d}{dt} X_t = u_t(X_t)$

Score-matching: $dX_t = \tilde{u}_t(X_t)\,dt + \sigma_t\,dW_t$

?

$\Rightarrow$ We want $X_t \sim p_t$

$$\boxed{\begin{aligned}&\text{Fokker-Planck eqn.}\\ &\partial_t p_t(x_t) = -\nabla \cdot \left(p_t(x_t)\tilde{u}_t(x_t)\right) + \frac{\sigma_t^2}{2}\Delta p_t(x_t)\end{aligned}}$$

$\tilde{u}_t = u_t + s_t$ LHS:

from continuity: $\partial_t p_t(x_t) =$

$$-\nabla \cdot \left(p_t(x)u_t(x)\right) = -\nabla \cdot \left(p_t(x)\tilde{u}_t(x)\right) + \frac{\sigma_t^2}{2}\Delta p_t(x)$$

$$= -\nabla \cdot \left(p_t(x)u_t(x)\right) - \nabla \cdot \left(p_t(x)s_t(x)\right) + \frac{\sigma_t^2}{2}\Delta p_t(x)$$

$$\Rightarrow \nabla \cdot \left(p_t(x)s_t(x)\right) = \frac{\sigma_t^2}{2}\nabla \cdot \left(\nabla p_t(x)\right)$$

$$= \frac{\sigma_t^2}{2}\nabla \cdot \left(\frac{\nabla p_t(x)}{p_t(x)}p_t(x)\right)$$

$$= \frac{\sigma_t^2}{2}\nabla \cdot \left(p_t(x)\nabla \log p_t(x)\right)$$

$$\Rightarrow s_t(x) = \frac{\sigma_t^2}{2}\nabla \log p_t(x)$$

Hence SDE:

$$dX_t = u_t(X_t)dt + \underbrace{\frac{\sigma_t^2}{2}\nabla \log p_t(x)}_{\text{Score}}$$

We want to learn the score

$$\nabla \log p_t(\underline{x}) = \int \nabla \log p_t(\underline{x}|\underline{z}) \frac{p_t(\underline{x}|\underline{z})\, p_{data}(\underline{z})}{p_t(\underline{x})}\, d\underline{z}$$

Proof

$$\nabla \log p_t(x) = \frac{\nabla p_t(x)}{p_t(x)} = \frac{\nabla \int p_t(x|z)\, p_{data}(z)\, dz}{p_t(\underline{x})}$$

$$= \int \frac{\nabla p_t(x|z)\, p_{data}(z)\, dz}{p_t(x)}$$

$$= \int \nabla \log p_t(x|z) \frac{p_t(x|z)\, p_{data}(z)}{p_t(x)}\, dz$$

Gaussian prob path

$$\nabla \log p_t(x|z) = \nabla \log \mathcal{N}(x; \alpha_t z, \beta_t^2 I_d)$$

$$= -\frac{x - \alpha_t z}{\beta_t^2}$$

To obtain density/Likelihood

$$\log P_1(x_1) = \log P_0(x_0) - \int_0^1 \mathrm{Tr}\left(\frac{\partial u_t^\theta(x_t)}{\partial x_t}\right) dt$$