Flow Matching II

Discussion #11

Written by:
Brent, CMK

1 Flow Matching Implementation

In last week's discussion, we explored geometric intuition for conditional and marginal flows. This week, we'll connect this idea to the actual implementation.

Consider a toy 1D flow matching problem. Our "dataset" has just two points: $A = +0.5$ and $B = -0.5$, each sampled with probability 0.5. We run a simple flow matching training loop:

```
L1: Sample x_1 ~ data (A or B with prob 0.5 each)
L2: Sample ε ~ N(0, 1)
L3: Sample t ~ Uniform(0, 1)
L4: Compute x_t = (1 − t)ε + t · x_1
L5: Compute target velocity:  u = x_1 − ε
L6: Update θ to minimize ||u_θ(x_t, t) − u||^2
```

For a given (x_t, t), the training loop sometimes supervises toward A and sometimes toward B.

1.1 Consider the case where $t = 0$ is sampled at L3, and $x_t = 0.5$ is computed at L4.

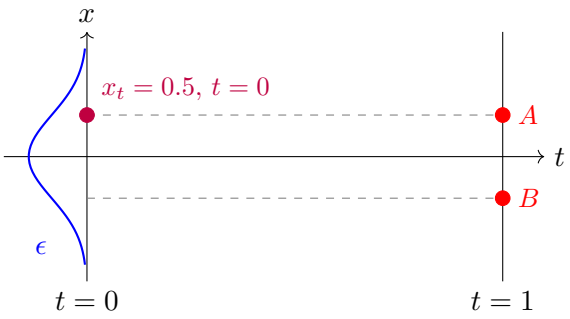

(a) If $x_1 = A$ was sampled at L1, what ϵ was sampled at L2? What if $x_1 = B$?

(b) Using L5, compute the target velocity u for each case. ($x_1 = A$ and $x_1 = B$).

(c) How often does the training loop produce $(x_t{=}0.5, t{=}0, x_1{=}A)$ vs. $(x_t{=}0.5, t{=}0, x_1{=}B)$?

(d) What does $u_\theta(0.5, 0)$ converge to after training? Which x_1 does this velocity vector point at?

1.2 Consider the case where $t = 0.5$ is sampled at L3, and $x_t = 0$ is computed at L4.

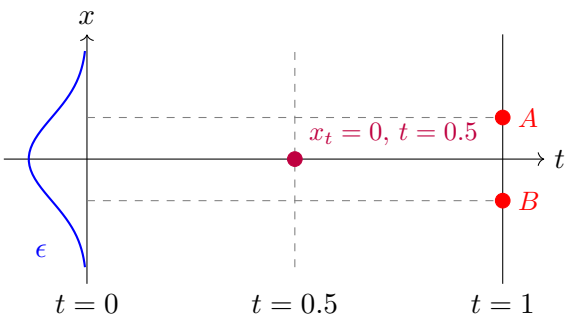

(a) If $x_1 = A$ was sampled at L1, what ϵ was sampled at L2? What if $x_1 = B$?

(b) Using L5, compute the target velocity u for each case.

(c) What does $u_\theta(0, 0.5)$ converge to?

1.3 Consider the case where $t = 0.9$ is sampled at L3, and $x_t = 0.5$ is computed at L4.

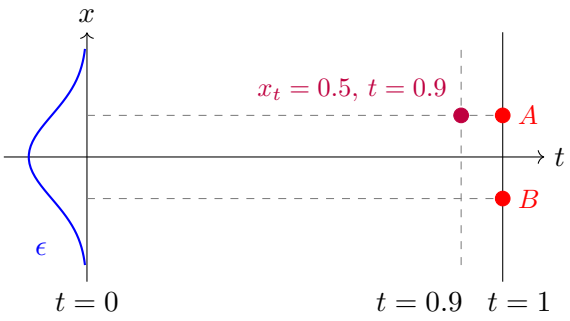

(a) If $x_1 = A$ was sampled at L1, what ϵ was sampled at L2? What if $x_1 = B$?

(b) Using L5, compute the target velocity u for each case.

(c) How often does the training loop produce supervision toward A vs. toward B at this (x_t, t)?

(d) What does $u_\theta(0.5, 0.9)$ converge to?

1.4 Both 1.1 and 1.3 have $x_t = 0.5$. Why does u_θ converge to -0.5 at $t = 0$ but ≈ 0 at $t = 0.9$?

2 Mixing and Matching Models and Samplers

Suppose you have a pre-trained **ϵ-prediction model** $\epsilon_\theta(x_t, t)$ that predicts noise (like DDPM), and a sampler that implements **Euler integration** for flow matching:

$$x_{t+\Delta t} = x_t + \Delta t \cdot u_\theta(x_t, t)$$

The sampler expects a velocity u, but your model outputs ϵ.

2.1 Using the flow matching interpolation $x_t = (1-t)\epsilon + t \cdot x_{\text{clean}}$ and velocity definition $u = x_{\text{clean}} - \epsilon$, derive a formula for u in terms of x_t, ϵ, and t.

2.2 Rewrite the Euler update step using ϵ_θ instead of u_θ.