

# Flow Matching from dummies

Sometimes you gotta go with the flow...




**!!! Disclaimer: This presentation contains memes. OMG!**

# Why do Flow Matching 101...



# Why do Flow Matching 102...





# Join the Flow Matching Cult



# Flow Matching

1



2



# Flow Matching




3 numbers



# Flow Matching



# Flow Matching






# Flow Matching




# Flow Matching



Here are the 3 numbers

# Flow Matching




How do you get them?

# Flow Matching




# Flow Matching



# Flow Matching




# Flow Matching






# Flow Matching




$$\pi_t = t \cdot \pi_1 + (1 - t) \cdot \pi_0$$

# Flow Matching



# Flow Matching


One more way!



Interpolation!


# Flow Matching

One more way!



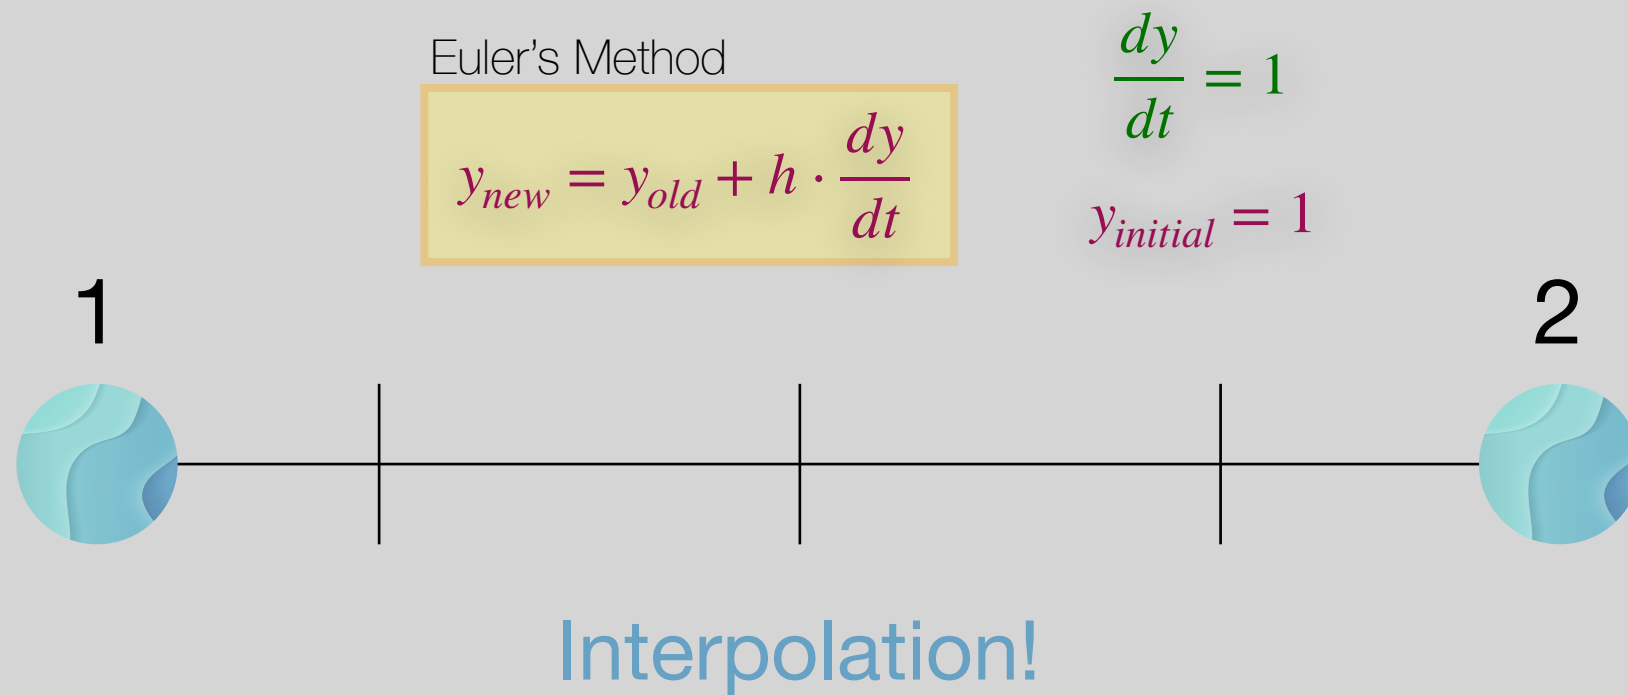
# Flow Matching

One more way!



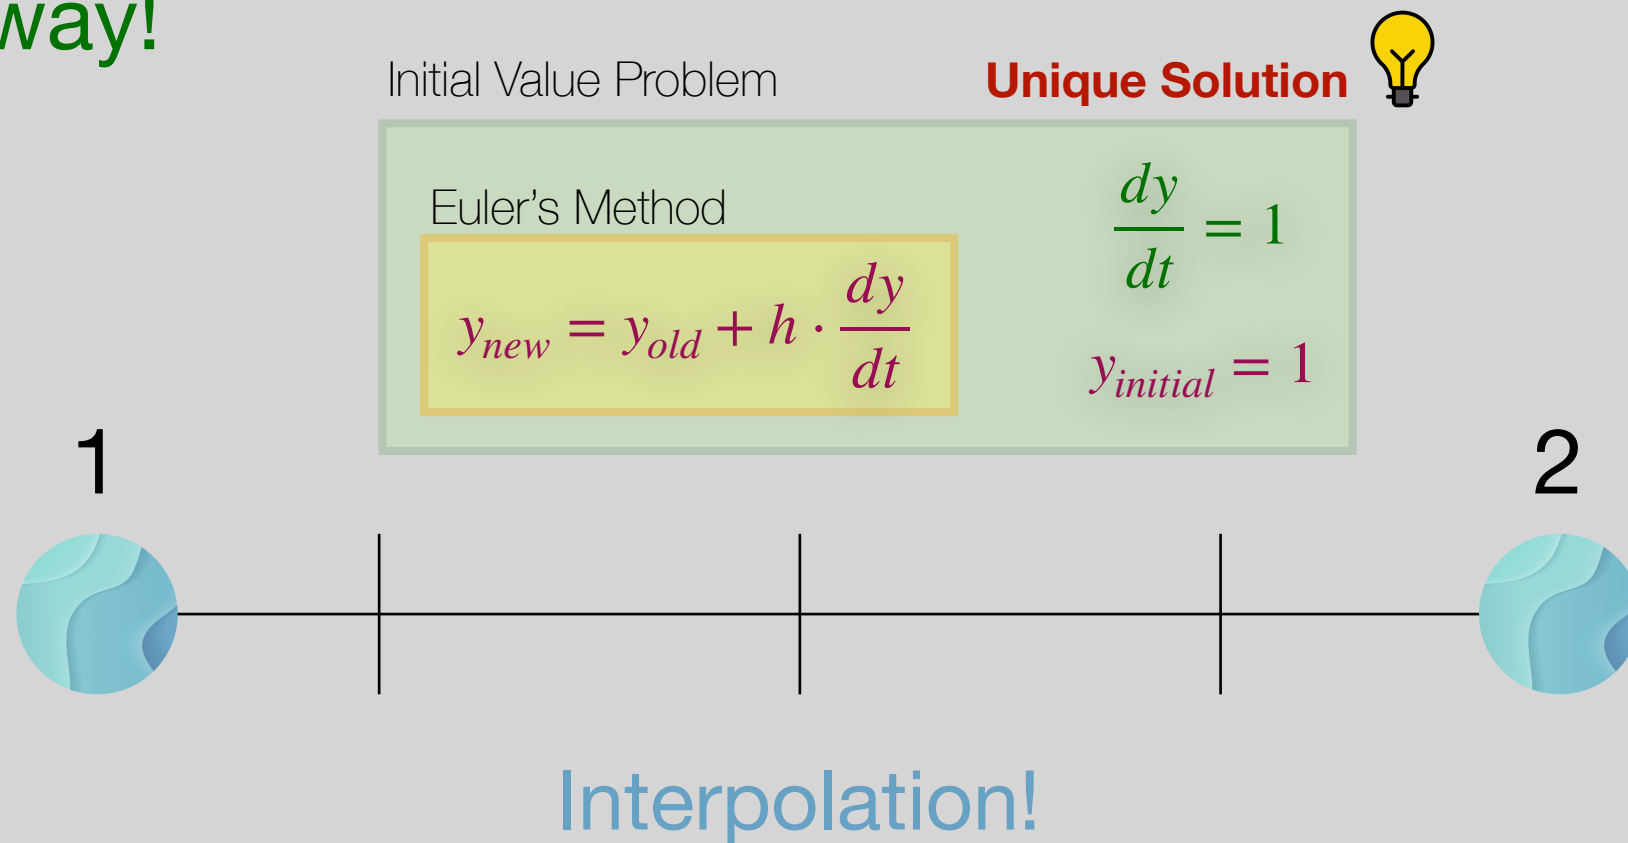
# Flow Matching

One more way!



# Flow Matching

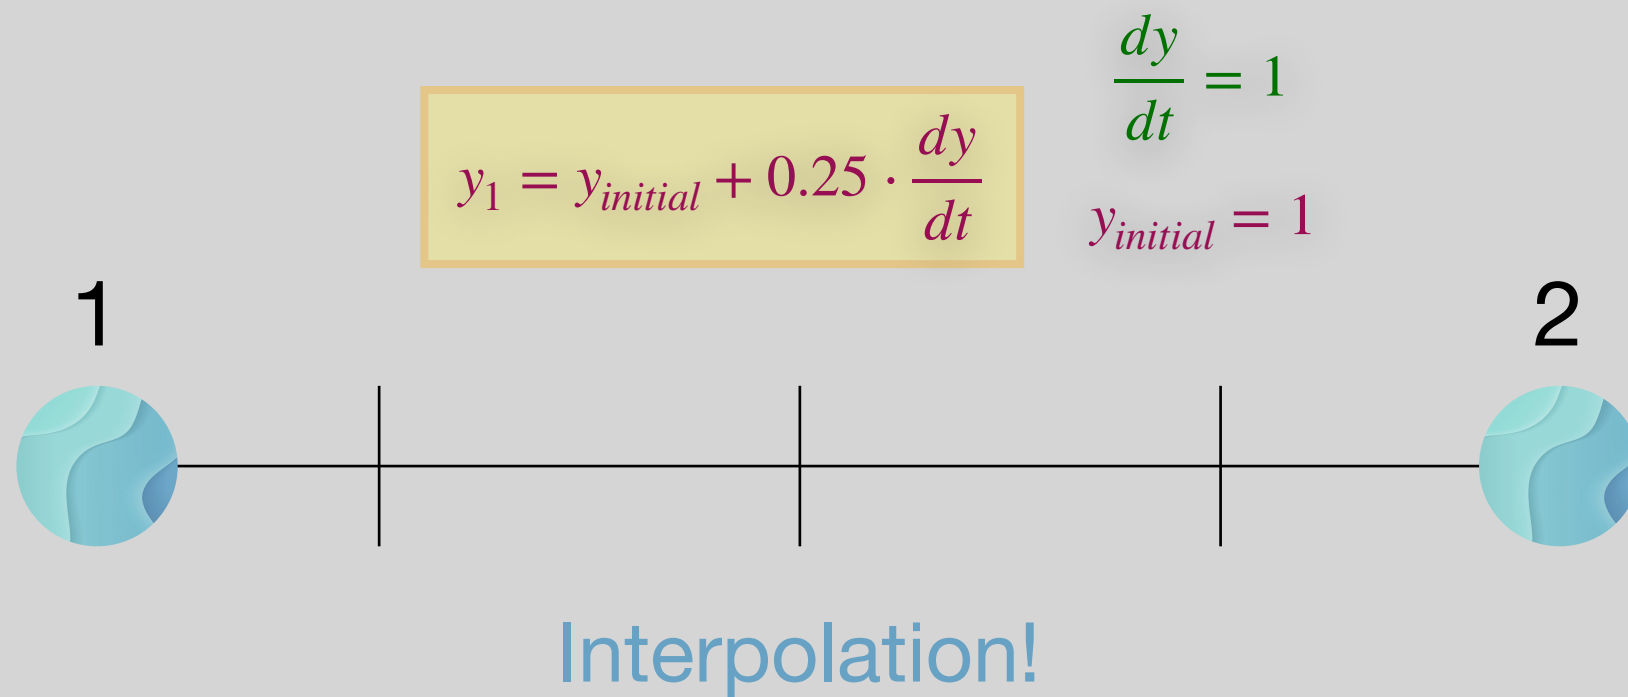
One more way!





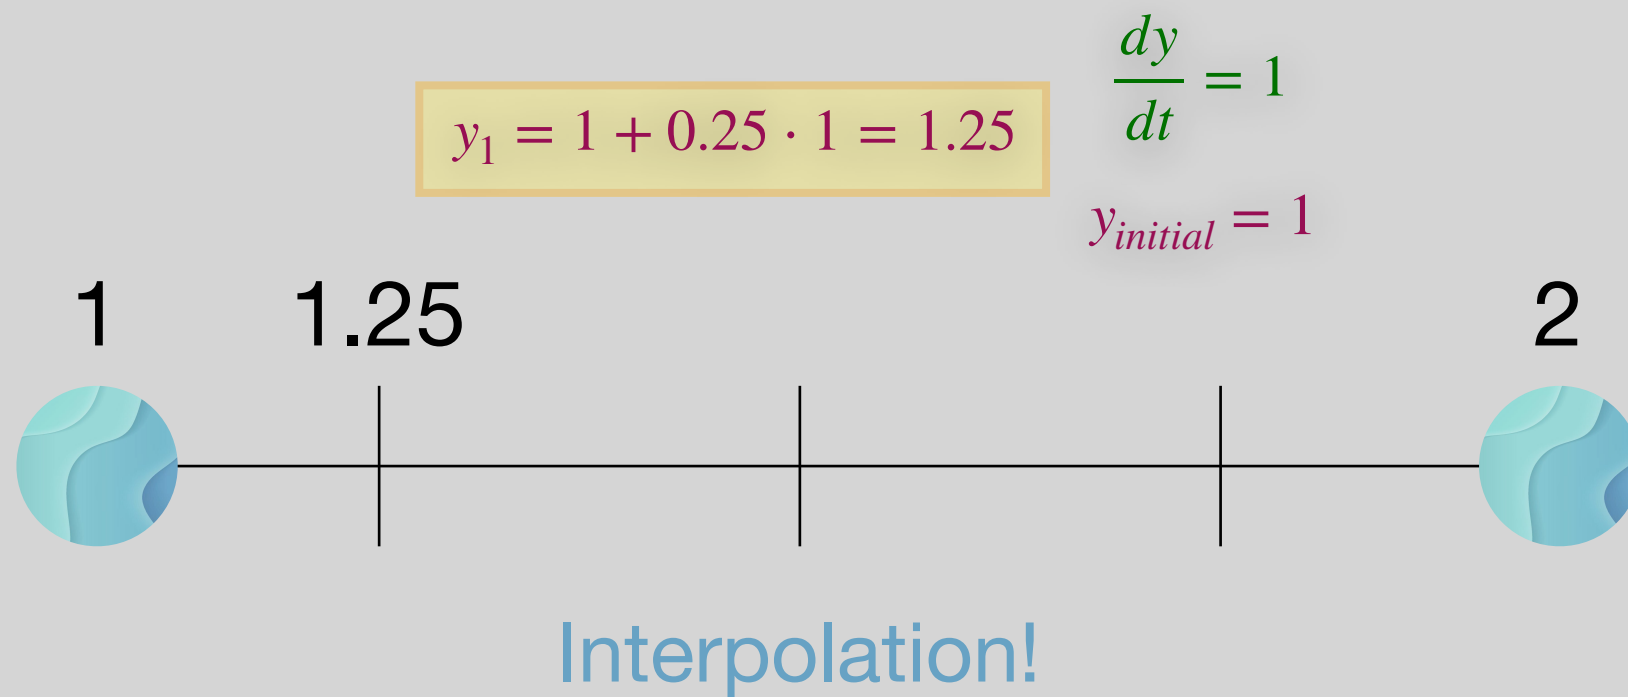
# Flow Matching

One more way!



# Flow Matching


One more way!



# Flow Matching


One more way!

$$y_2 = 1.25 + 0.25 \cdot 1 = 1.5$$



# Flow Matching


One more way!



# Flow Matching is literally **Interpolation**




# Flow Matching is literally Interpolation



Source Distribution


$$\pi_0$$



Target Distribution

$$\pi_1$$

# Flow Matching is literally Interpolation



Source Distribution


$\pi_0$

Target Distribution

$\pi_1$



# Flow Matching is literally Interpolation



Source Distribution


$\pi_0$

Target Distribution



$\pi_1$

$$\pi_t = t \cdot \pi_1 + (1 - t) \cdot \pi_0$$

# Flow Matching is literally Interpolation



Source Distribution




Target Distribution

$$\pi_t = \frac{1}{4} \cdot \pi_1 + \frac{3}{4} \cdot \pi_0$$



*Lipman, Yaron, et al. "Flow Matching for Generative Modeling." ICLR 2023.*

*Liu, Xingchao, et al. "Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow." ICLR 2023.*

# Flow Matching is literally Interpolation



Source Distribution



Target Distribution


$$\pi_{\frac{1}{4}} = \frac{1}{4} \cdot \pi_1 + \frac{3}{4} \cdot \pi_0$$

$$\pi_{\frac{1}{4}} = \pi_0 + \frac{1}{4} \cdot \nabla \pi_t$$

*Lipman, Yaron, et al. "Flow Matching for Generative Modeling." ICLR 2023.*

*Liu, Xingchao, et al. "Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow." ICLR 2023.*

# Flow Matching is literally Interpolation




$$\pi_{\frac{1}{4}} = \frac{1}{2} \cdot \pi_1 + \frac{1}{2} \cdot \pi_0$$

$$\pi_{\frac{1}{2}} = \pi_{\frac{1}{4}} + \frac{1}{4} \cdot \nabla \pi_t$$

*Lipman, Yaron, et al. "Flow Matching for Generative Modeling." ICLR 2023.*

*Liu, Xingchao, et al. "Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow." ICLR 2023.*

# Flow Matching is literally Interpolation



$$\pi_{\frac{3}{4}} = \frac{3}{4} \cdot \pi_1 + \frac{1}{4} \cdot \pi_0$$

$$\pi_{\frac{3}{4}} = \pi_{\frac{1}{2}} + \frac{1}{4} \cdot \nabla \pi_t$$


Lipman, Yaron, et al. "Flow Matching for Generative Modeling." ICLR 2023.

Liu, Xingchao, et al. "Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow." ICLR 2023.


# Flow Matching is **So Simple!**



# However...




Source Distribution




Target Distribution???



# However...




Source Distribution







Target samples

# Can we get the velocity to get to intermediates?




Source Distribution




Target samples

# Solution 101 : Random coupling




Source Samples




Target samples

# Solution 101 : Random coupling




Source Samples



Target samples

# Solution 101 : Intermediate gradient/ velocity

What are the velocities to make the intermediates?




Source Samples

Target samples

# !!! Velocity depends on Interpolation paths

Which one to pick?



Source Samples

Target samples

# Solution 101 : Go linear




Source Samples

Target samples

# Solution 101 : Gradients

## Model prediction




Source Samples

$X_0$

Target samples

$X_1$




Lipman, Yaron, et al. "Flow Matching for Generative Modeling." ICLR 2023.

Liu, Xingchao, et al. "Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow." ICLR 2023.



# Solution 101 : Gradients

## Model prediction




Source Samples

$X_0$

Target samples

$X_1$



Lipman, Yaron, et al. "Flow Matching for Generative Modeling." ICLR 2023.

Liu, Xingchao, et al. "Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow." ICLR 2023.


# Overall Training

*Lipman, Yaron, et al. "Flow Matching for Generative Modeling." ICLR 2023.*

*Liu, Xingchao, et al. "Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow." ICLR 2023.*


# Overall Training

## (1) Random Coupling



Source Samples

$X_0$




Target samples

$X_1$


# Overall Training

(2) Randomly sample  $t$  in  $[0,1)$



Source Samples

$X_0$




Target samples

$X_1$

# Overall Training

## (3) Linear interpolation (model input)



Source Samples

$X_0$


Target samples

$X_1$

$$X_t = t \cdot X_1 + (1 - t) \cdot X_0$$

# Overall Training

## (4) Predict Velocity (model output)



Source Samples

$X_0$


Target samples

$X_1$



# Overall Training

## (5) Optimize



$$\mathbb{E}[||u_\theta(t) - (X_1 - X_0)||^2]$$

# Sampling

## Euler's Method

$$x_{next} = x_{old} + \frac{1}{num\_steps} \cdot u_{\theta}(t)$$

$$x_{initial} \sim \mathcal{N}(0, I)$$



# Code 101...



```
1  ## Setting up the flow matching pipeline ##
2  class FlowMatchingPipeline:
3      """Implements the Flow Matching pipeline"""
4
5      def __init__(self):
6          """Constructor"""
7          self.model = model()
8          self.optim = torch.optim.AdamW(self.model.parameters(), lr=1e-4)
9          self.loss = nn.MSELoss()
```

# Code 101...



```
1  def interpolate(self, x1, t):  
2      """Interpolate x1 data from x0 data."""  
3  
4      x0 = torch.randn(*x1.shape, device=x1.device, dtype=x1.dtype)  
5  
6      xt = t * x1 + (1 - t) * x0  
7  
8      vel = x1 - x0  
9  
10     return xt, vel
```

# Code 101...



```
1  def training_step(self, data):
2      """Implements one single training step for the model"""
3
4      data = data.to(self.device)
5
6      t = torch.rand(
7          data.shape[0],
8          dtype=data.dtype,
9          device=self.device
10     ).view(-1,1)
11
12     xt, true_vt = self.interpolate(data, t)
13
14     self.optim.zero_grad()
15
16     vt = self.model(xt, t)
17
18     loss = self.loss(true_vt, vt)
19     loss_val = loss.clone().detach().cpu().item()
20
21     loss.backward()
22
23     self.optim.step()
24
25     return loss_val
```

# Code 101...




```
1  @torch.no_grad()
2  def sample(self, num_steps, init_val):
3      """Samples using euler steps using the model"""
4
5      step_size = 1 / num_steps
6      xt = init_val
7
8      for step in range(1, num_steps):
9          t = torch.tensor(
10             [step / num_steps] * init_val.shape[0],
11             dtype=xt.dtype,
12             device=self.device
13             ).view(-1,1)
14          vt = self.model(xt, t)
15          xt = xt + step_size * vt
16
17      return xt
18
```

# Gimme the Theory...




# Theory of Flow Matching 101...

## Recap: Random Coupling




Source Samples  
 $X_0$



Target samples  
 $X_1$


# Theory of Flow Matching 101...

## Recap: Random Coupling



# Theory of Flow Matching 101...


## Conditional Probability Path






# Theory of Flow Matching 101...

## Conditional Probability Path




# Theory of Flow Matching 101...

## Conditional Probability Path




# Theory of Flow Matching 101...

## Conditional Probability Path




# Theory of Flow Matching 101...

## Conditional Probability Path




# Theory of Flow Matching 101...

## Conditional Probability Path




# Theory of Flow Matching 101...

## Conditional Probability Path



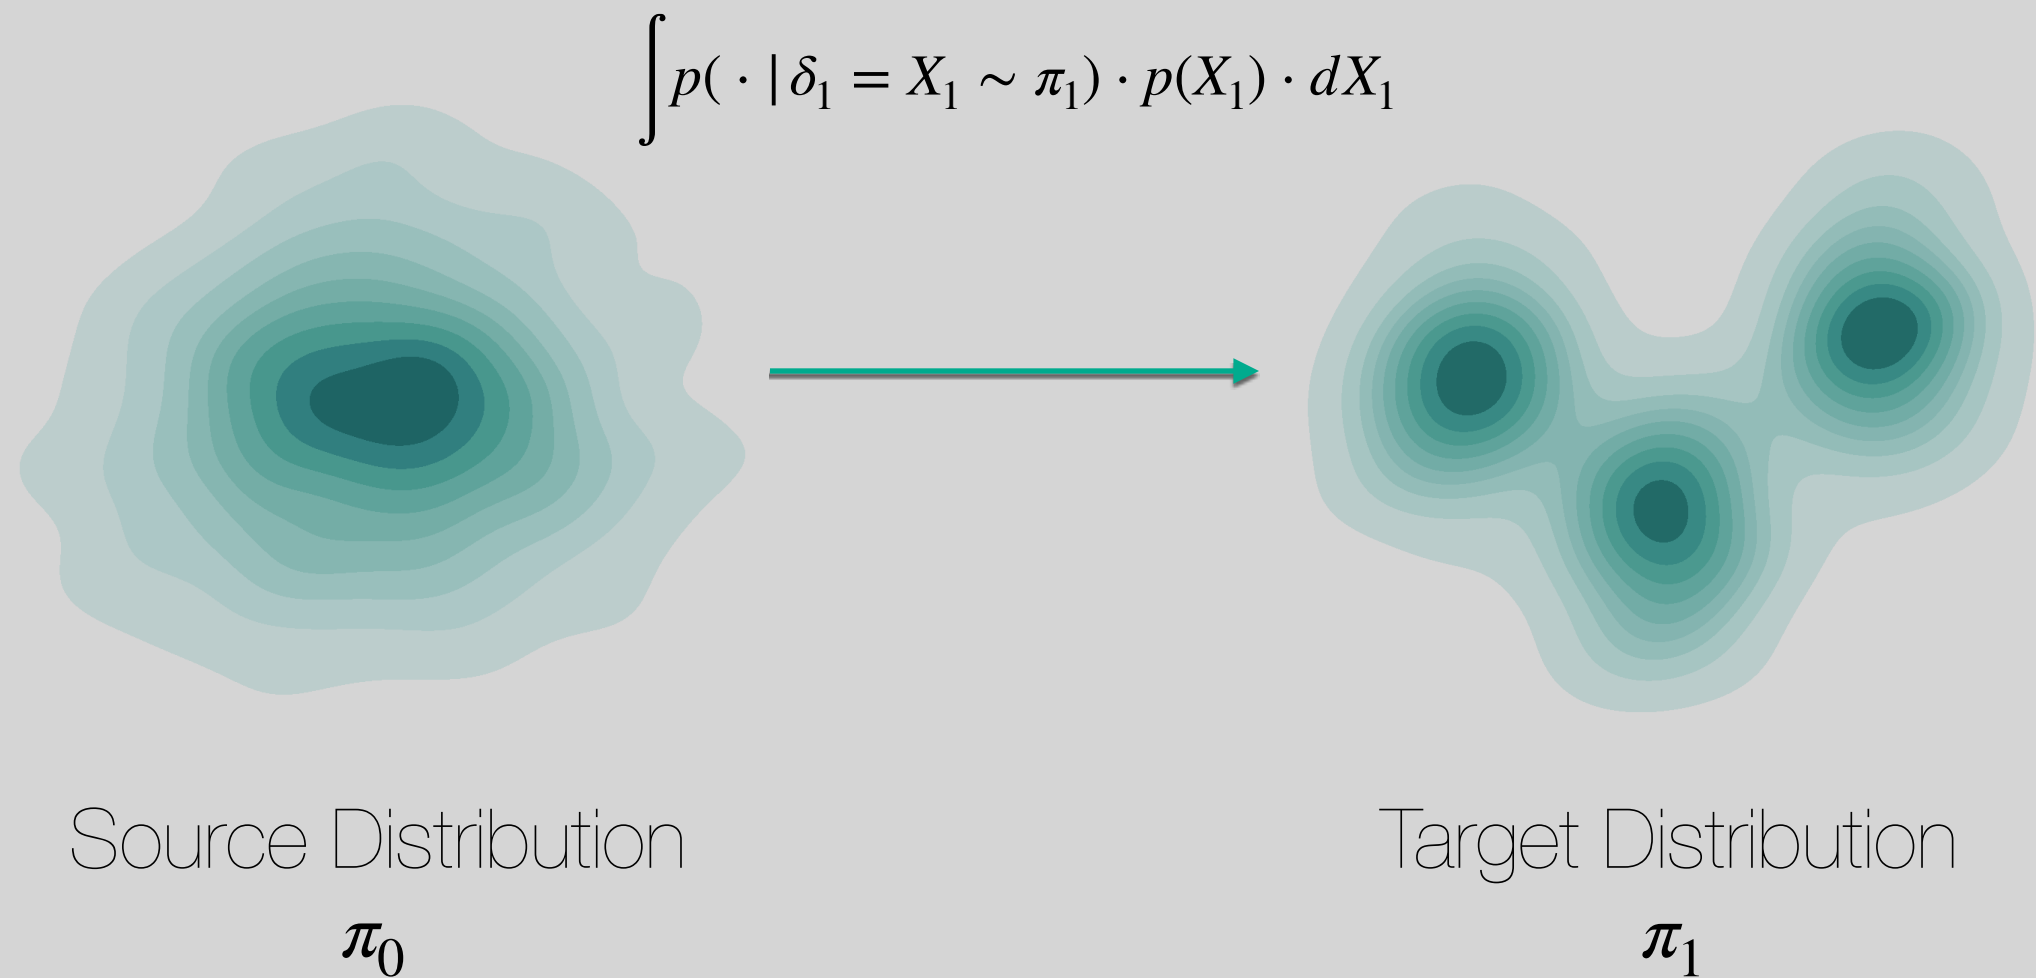
# Theory of Flow Matching 101...

## Conditional Probability Path



# Theory of Flow Matching 101...


## Marginal Probability Path






# Theory of Flow Matching 101...

## Marginal Velocity Field




# Theory of Flow Matching 101...

Marginal Velocity Field : Follows marginal probability path



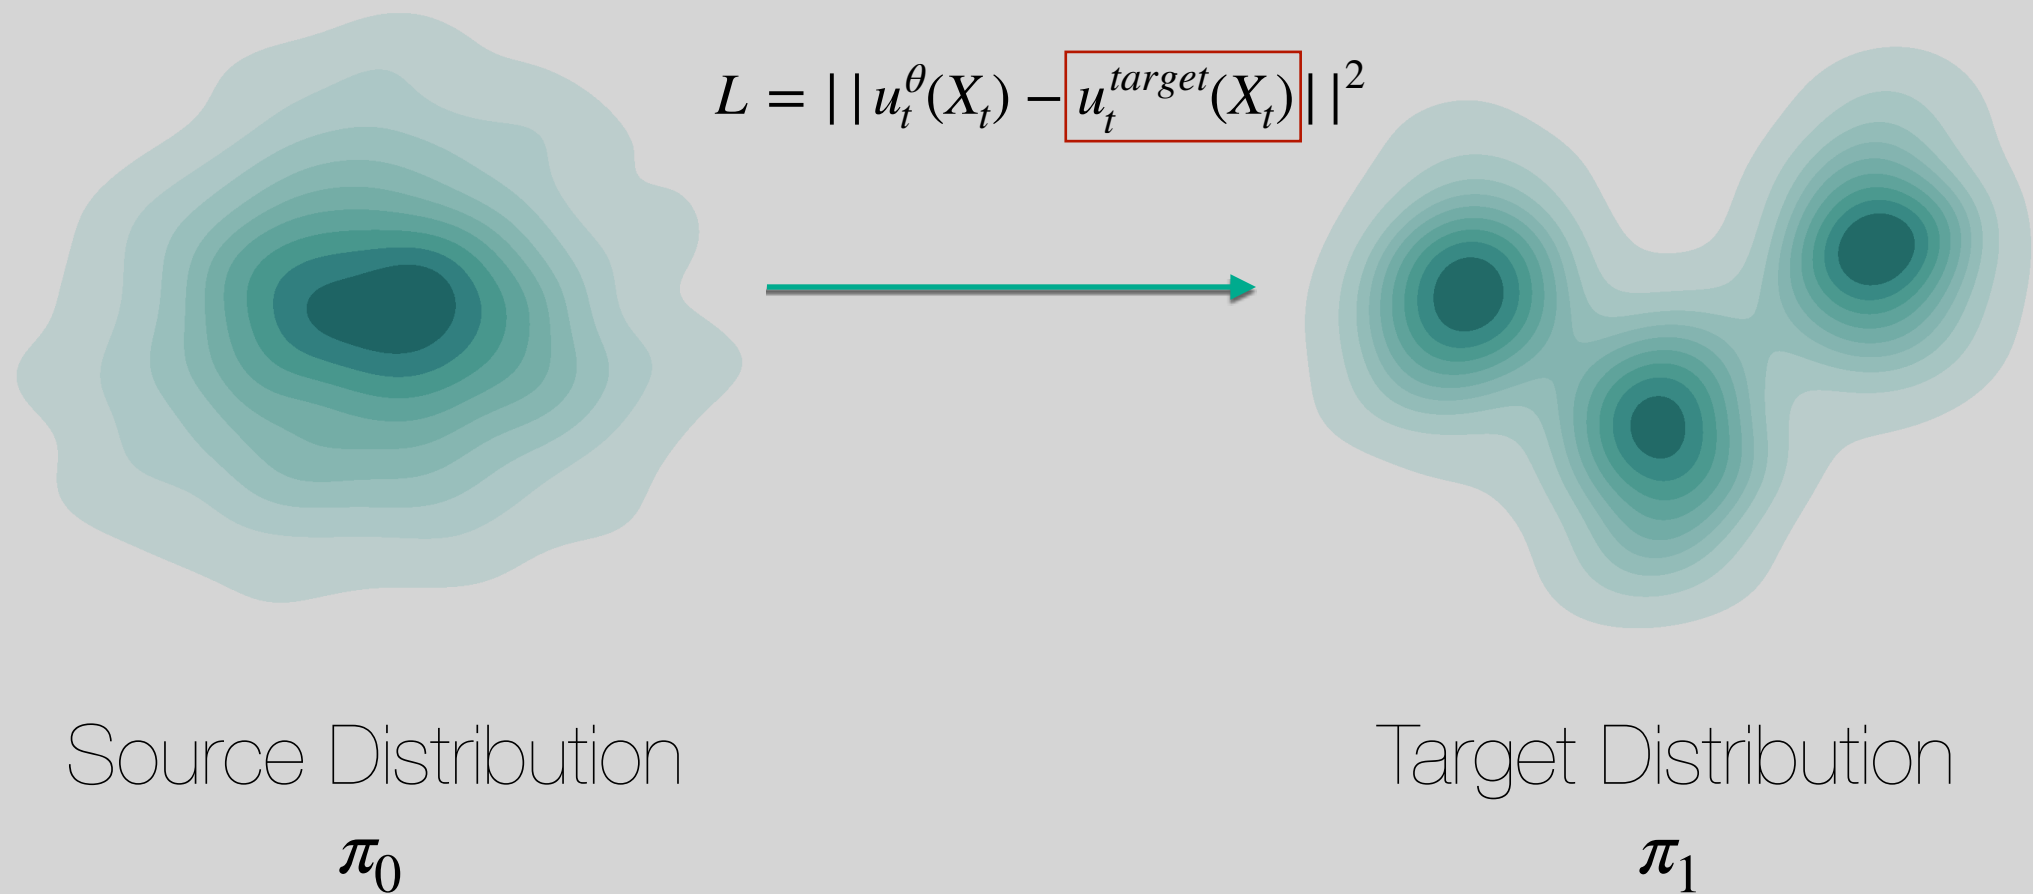
# Theory of Flow Matching 101...

Marginal Velocity Field : **Minimize**



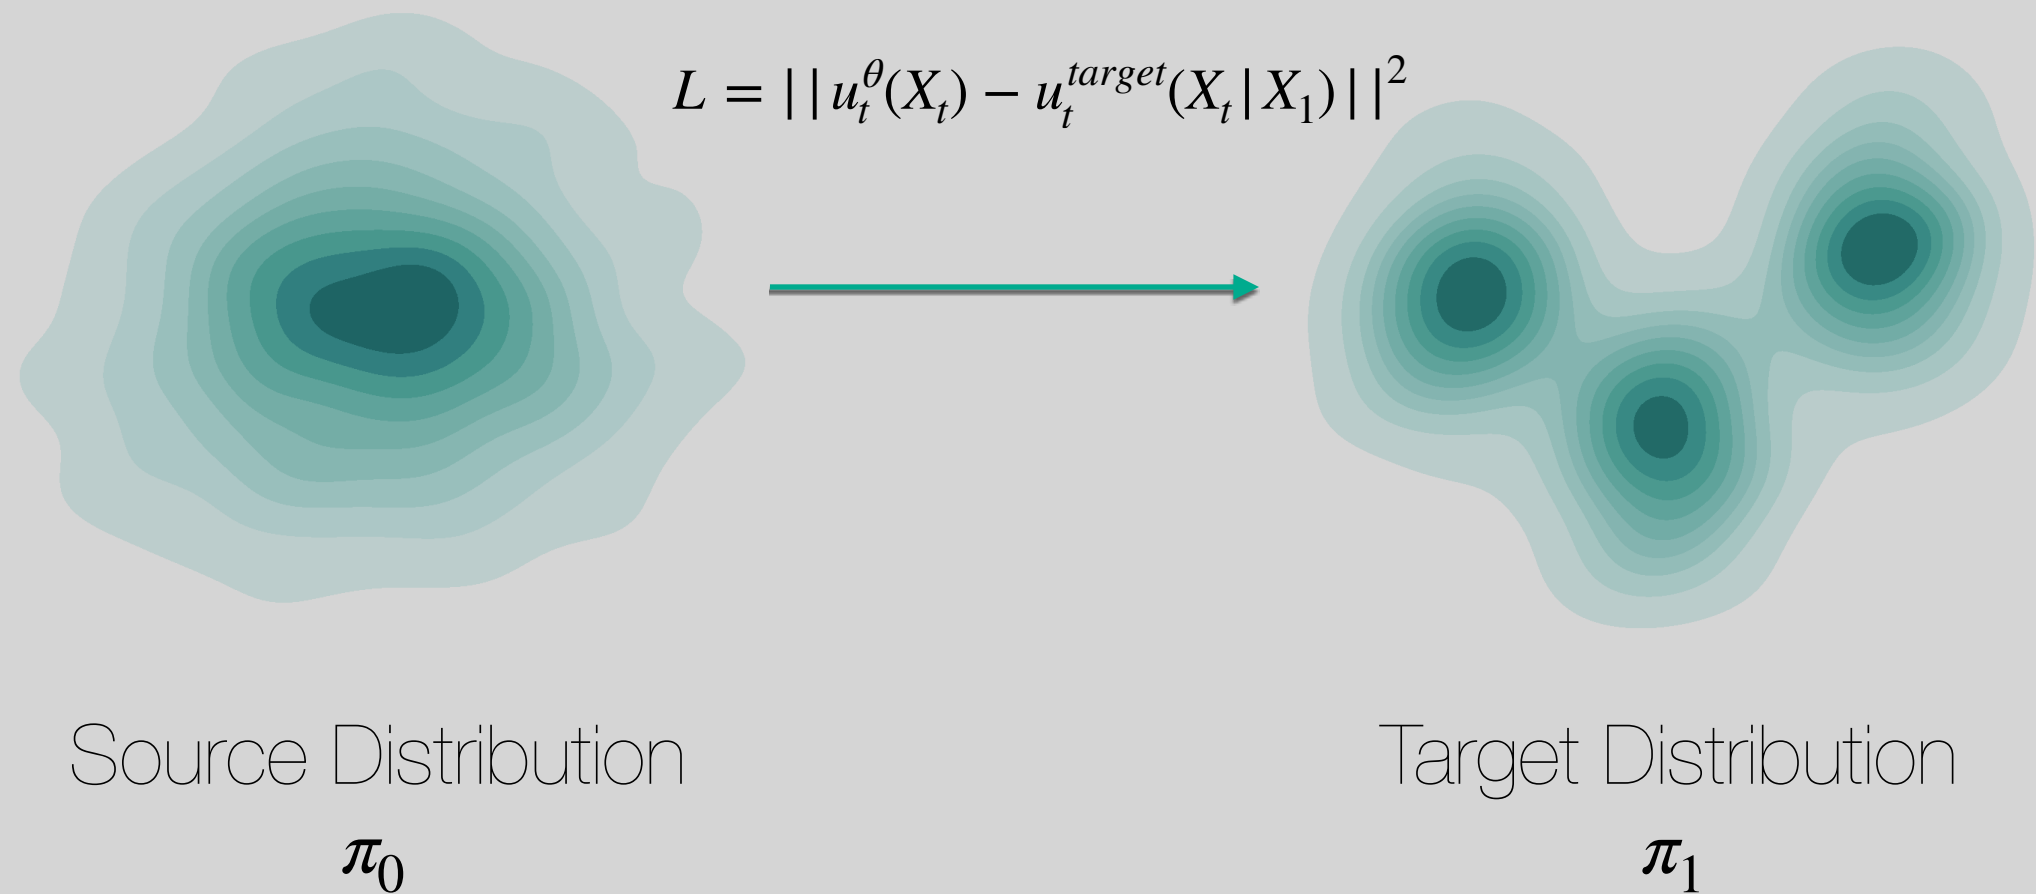
# Theory of Flow Matching 101...

## Marginal Velocity Field



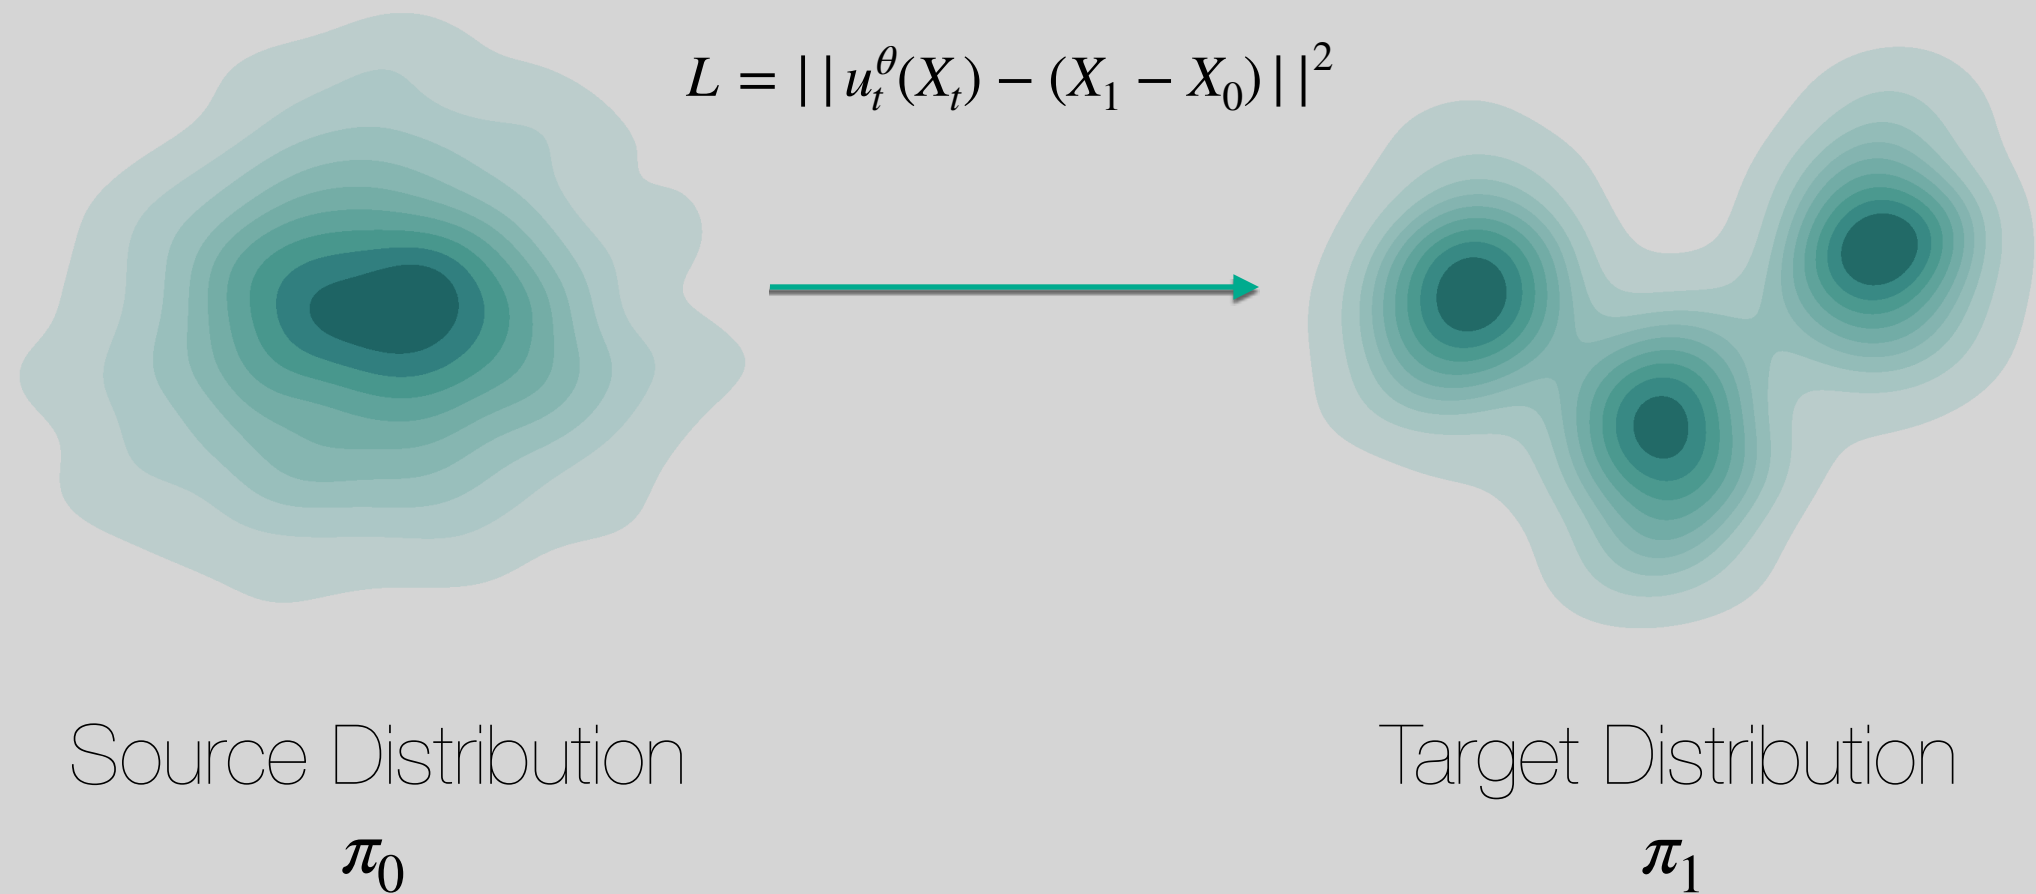
# Theory of Flow Matching 101...

## Conditional Velocity Field



# Theory of Flow Matching 101...


## Conditional Velocity Field for optimal transport path



# Update 101: ReFlow

# Update 101: ReFlow


## Recap: Random Coupling





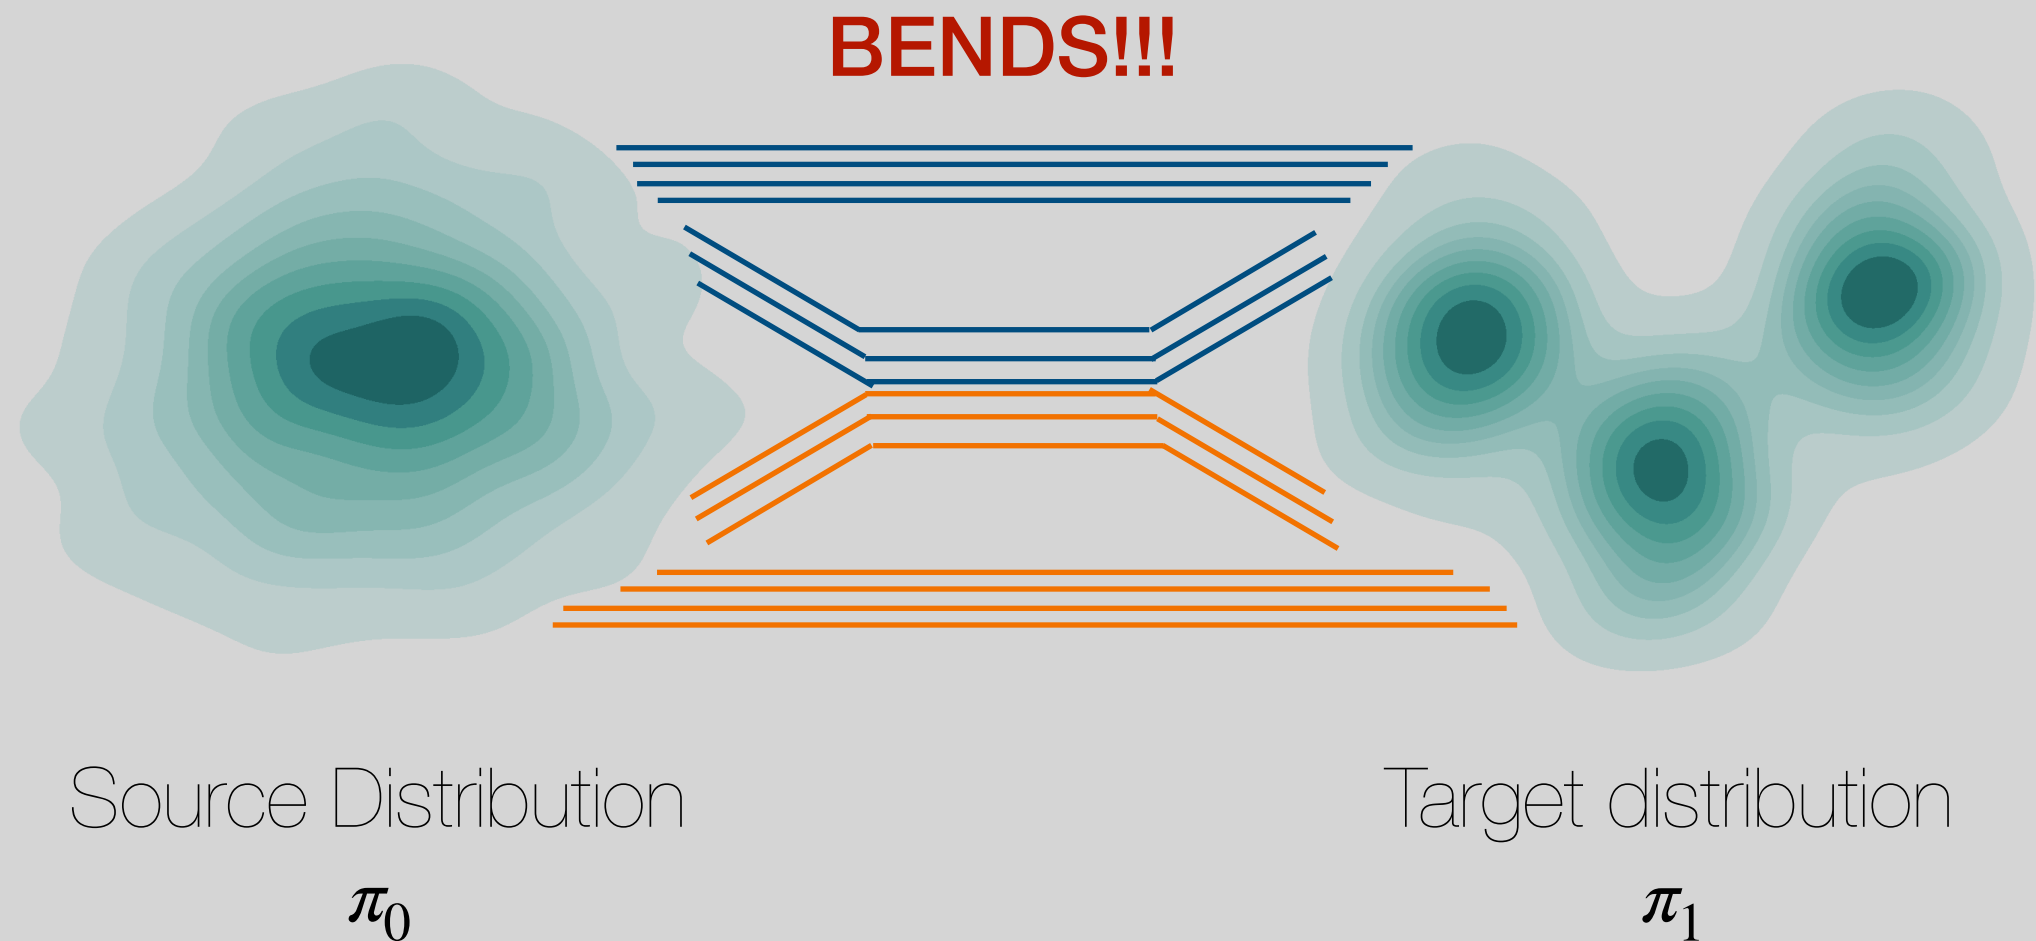
# Update 101: ReFlow

## Learnt Vector Fields



# Update 101: ReFlow


## Learnt Vector Fields



# Update 101: ReFlow


## Learnt Vector Fields

Requires more Sampling steps




# Update 101: ReFlow

Distill with straight pairs




# Update 101: ReFlow

## Learnt Vector Fields



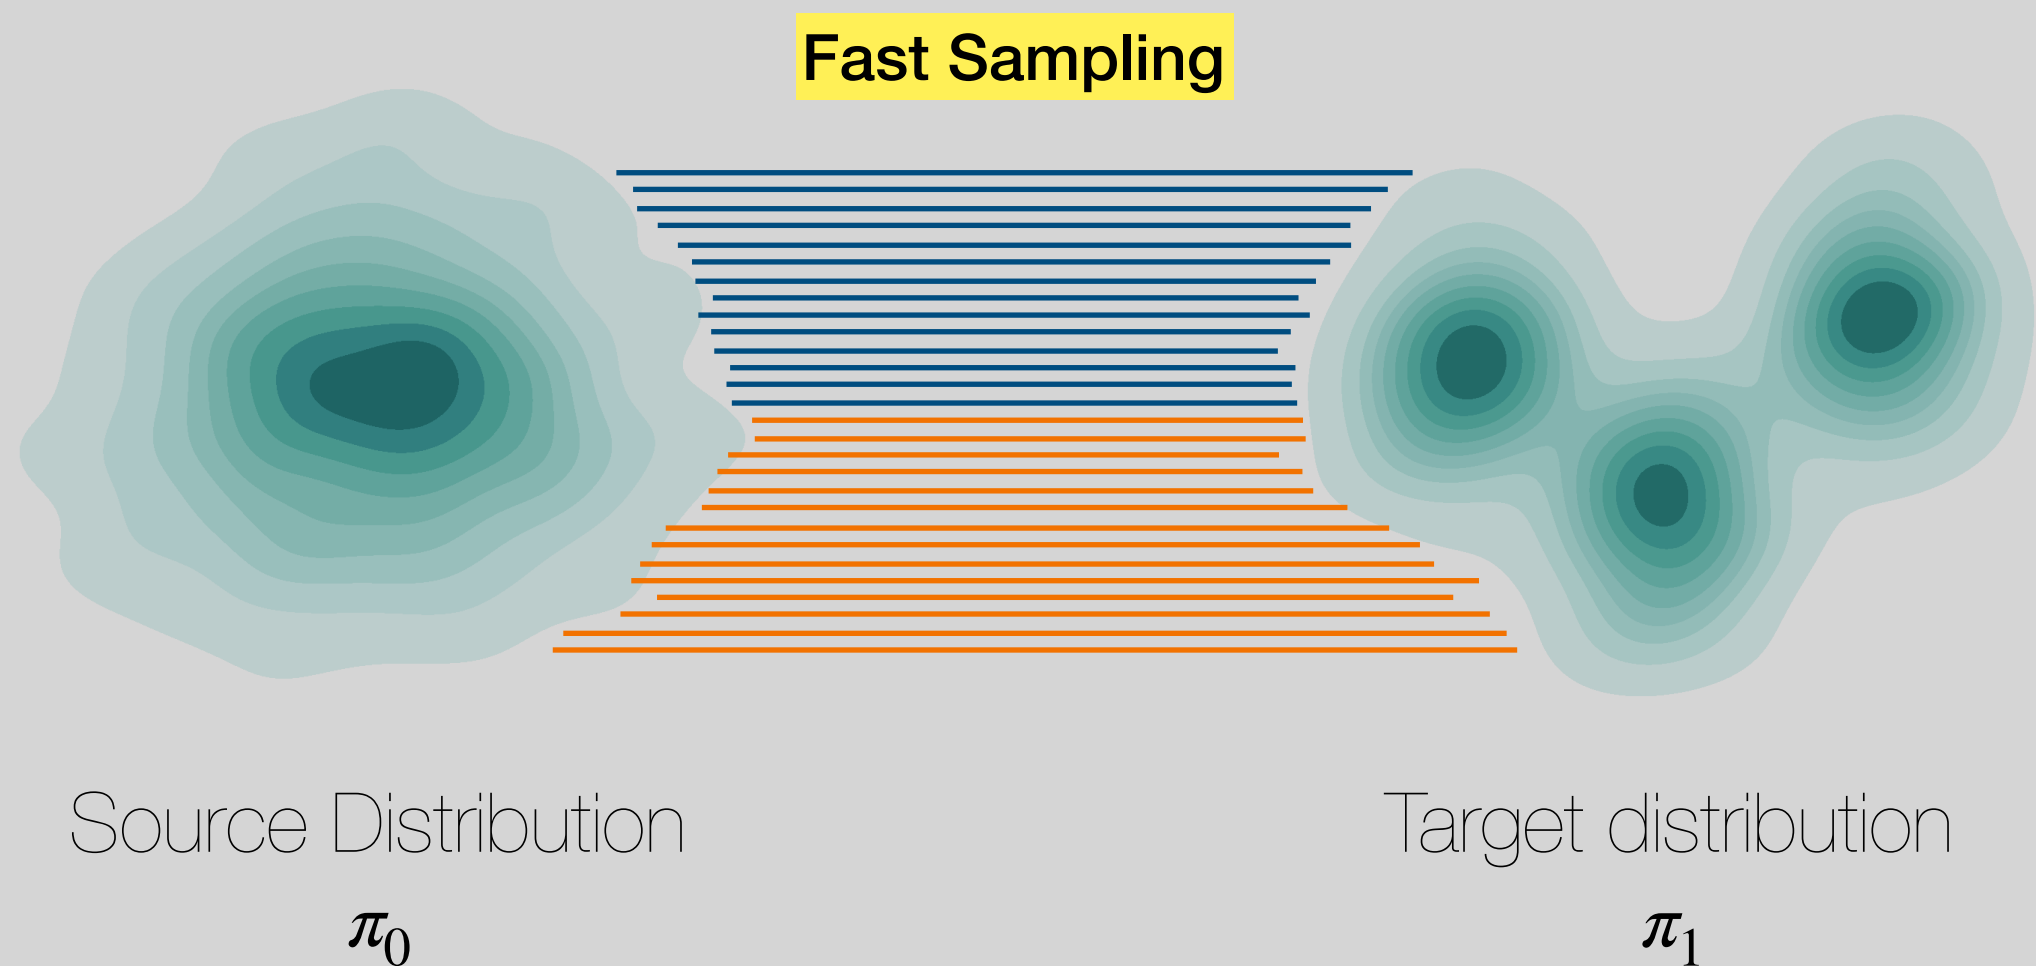
# Update 101: ReFlow

## Learnt Vector Fields



# Update 101: ReFlow

## Learnt Vector Fields



# Update 102: Better ODE Solvers

## Recap: Sampling

### Euler's Method

$$x_{next} = x_{old} + \frac{1}{num\_steps} \cdot u_{\theta}(t)$$

$$x_{initial} \sim \mathcal{N}(0, I)$$



# Update 102: Better ODE Solvers


Why not use better ones?

**Heun Method**

**RK Method**


# Update 103: Non-uniform timestep sampling

Recap: Randomly sample  $t$  in  $[0,1)$



Source Samples

$X_0$




Target samples

$X_1$


# Update 103: Non-uniform timestep sampling

## Recap: Learnt Vector Fields



# Update 103: Non-uniform timestep sampling

Change in single path mostly in the middle



# Update 103: Non-uniform timestep sampling

Logit-Normal Sampling


$$t = \frac{1}{1 + \exp(u)}$$

$$u \sim \mathcal{N}(0,1)$$

# Connection to Diffusion Models


# Connection to Diffusion Models

Flow Matching: Learn a vector field



# Connection to Diffusion Models

## Flow Matching Sampling : Solving an ODE






# Connection to Diffusion Models

## Flow Matching Sampling : Solving an ODE


$$dx = v_t \cdot dt$$

$$x_{t+h} = x_t + h \cdot v_t$$



Source Distribution

$\pi_0$



Target distribution

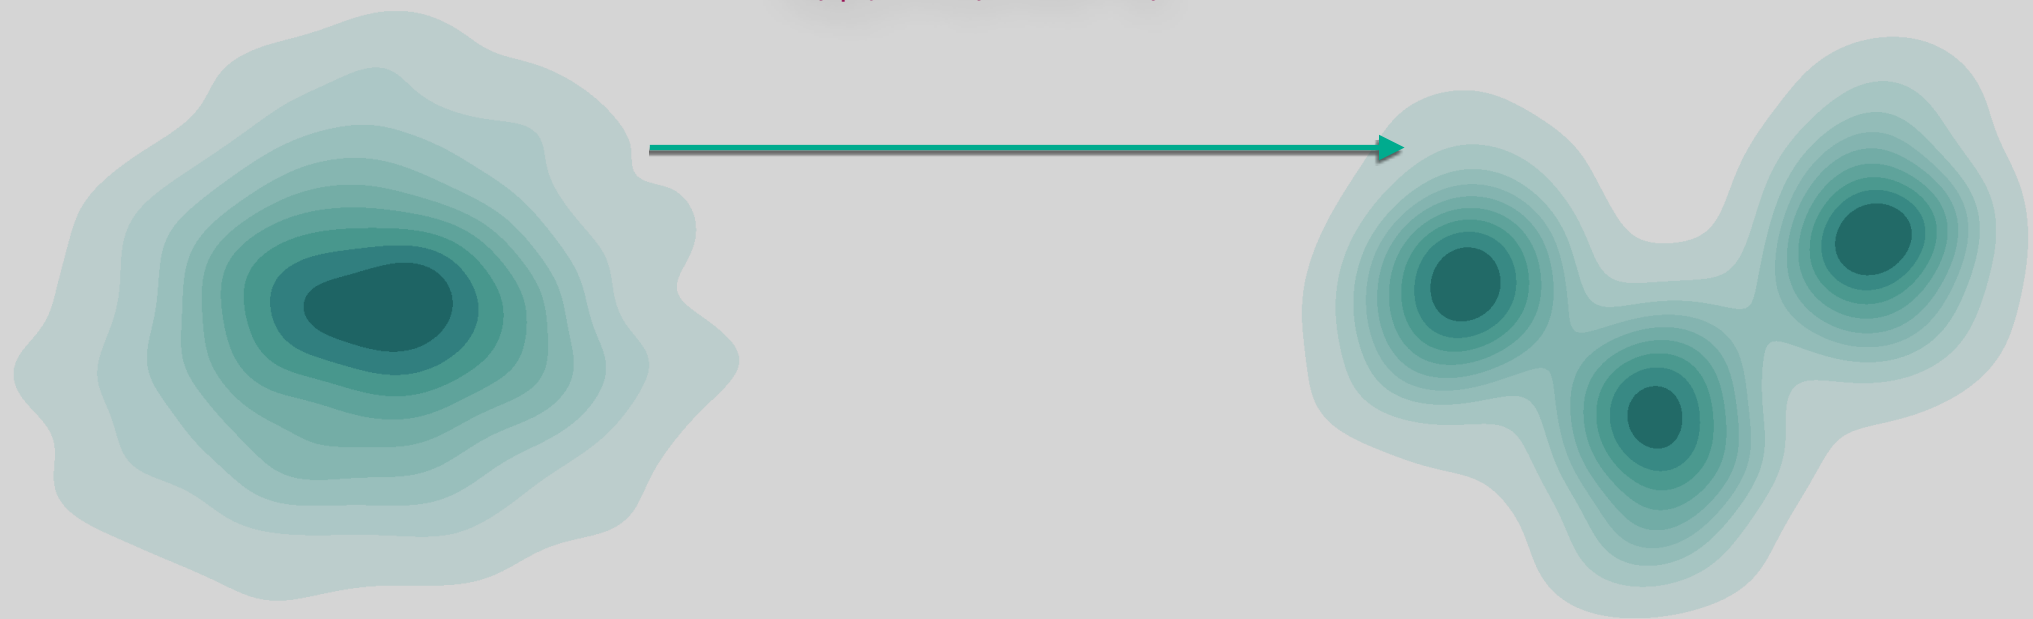
$\pi_1$

# Connection to Diffusion Models

## Flow Matching Sampling : Solving an ODE

$$dx = v_t \cdot dt$$

$$x_{t+h} = x_t + h \cdot v_t$$



Source Distribution

$\pi_0$


Target distribution

$\pi_1$

# Connection to Diffusion Models


## SDE and Diffusion

$$dx = [v_t + \lambda(w_t) \cdot s_t] \cdot dt + w_t \cdot dW$$



Source Distribution

$\pi_0$



Target distribution

$\pi_1$

# Connection to Diffusion Models

## SDE and Diffusion

$$dx = [\boxed{v_t} + \lambda(w_t) \cdot s_t] \cdot dt + w_t \cdot dW$$



Source Distribution

$\pi_0$

Target distribution

$\pi_1$

# Connection to Diffusion Models

## SDE and Diffusion

$$dx = [v_t + \lambda(w_t) \cdot s_t] \cdot dt + w_t \cdot dW$$

Noise



Source Distribution

$\pi_0$

Target distribution


$\pi_1$

# Connection to Diffusion Models

## SDE and Diffusion

$$dx = [v_t + \lambda(w_t) \cdot s_t] \cdot dt + w_t \cdot dW$$

Noise



Source Distribution

$\pi_0$


Target distribution

$\pi_1$

# Connection to Diffusion Models

## SDE and Diffusion

$$dx = \underbrace{[v_t + \lambda(w_t) \cdot s_t]}_{\text{Correction}} \cdot dt + \underbrace{w_t \cdot dW}_{\text{Noise}}$$



Source Distribution

$\pi_0$

Target distribution

$\pi_1$

# Connection to Diffusion Models

## Diffusion Models

Predict Score

$$dx = [v_t + \lambda(w_t) \cdot \overset{\downarrow}{s_t}] \cdot dt + w_t \cdot dW$$



# Connection to Diffusion Models

## Diffusion Models

Predict Score

$$dx = [v_t + \lambda(w_t) \cdot \overset{\downarrow}{s_t}] \cdot dt + w_t \cdot dW$$

Fun Fact: Using **Tweedie's formula**

$$v_t \leftrightarrow s_t$$

# Connection to Diffusion Models

## Diffusion Models

$$x_{t+h} = x_t +$$

# Connection to Diffusion Models

## Diffusion Models

$$x_{t+h} = x_t + h \cdot (s_t$$

# Connection to Diffusion Models

## Diffusion Models

$$x_{t+h} = x_t + h \cdot \left( s_t + \frac{1}{\lambda(w_t)} \cdot v_t \right)$$

# Connection to Diffusion Models

## Diffusion Models

DDIM Sampler

$$x_{t+h} = x_t + h \cdot \left( s_t + \frac{1}{\lambda(w_t)} \cdot v_t \right)$$

# Connection to Diffusion Models

## Diffusion Models

$$x_{t+h} = x_t + h \cdot \left( s_t + \frac{1}{\lambda(w_t)} \cdot v_t \right) + w_t \cdot \epsilon$$

# Connection to Diffusion Models

## Diffusion Models

DDPM Sampler

$$x_{t+h} = x_t + h \cdot \left( s_t + \frac{1}{\lambda(w_t)} \cdot v_t \right) + w_t \cdot \epsilon$$

# Thank you!

