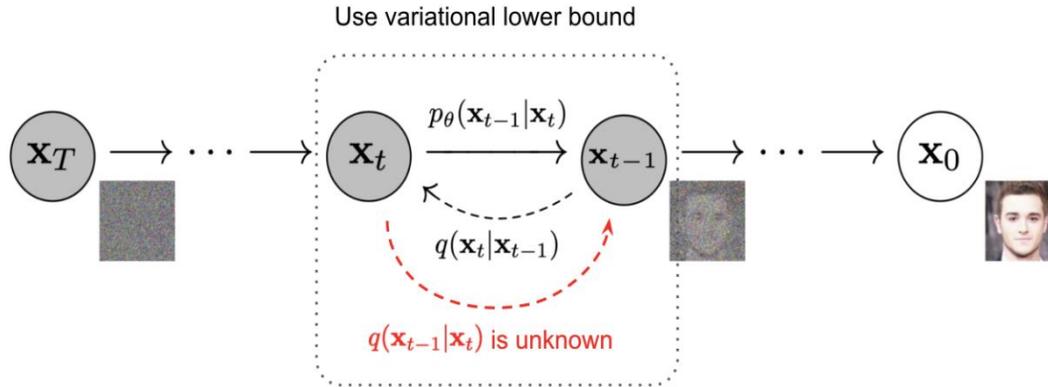


# Towards the Efficient and Human Preference Aligned Diffusion Model-based Generation

Fu-Yun Wang

fywang0126@gmail.com

# Diffusion Models: Markovian Perspective



## Assumption:

$$p(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$$

## Forward Process:

- $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}_{t-1}, (1 - \alpha_t)\mathbf{I})$
- $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$

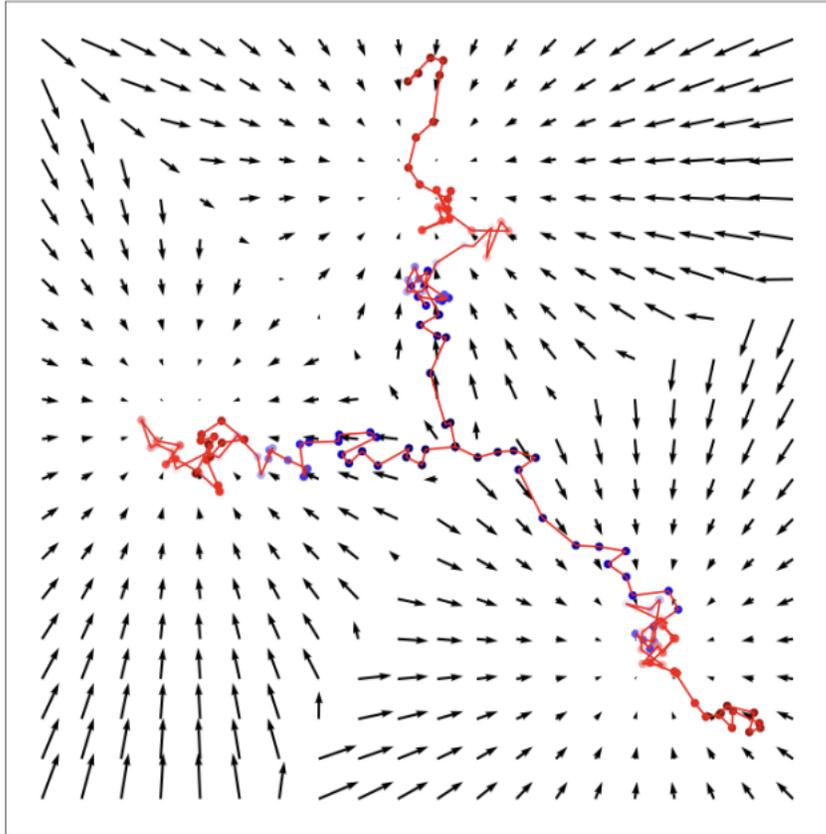
## Reverse Process:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0)q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)}$$

- Maximum Likelihood Estimation (MLE) is Equivalent to

$$\arg \min_{\theta} \mathbb{E}_{t \sim U\{2, T\}} [\mathbb{E}_{q(\mathbf{x}_t|\mathbf{x}_0)} [D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \parallel p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))]]$$

# Vanilla Score Matching



- Energy-based model

- $p_{\theta}(\mathbf{x}) = \frac{1}{Z_{\theta}} e^{-f_{\theta}(\mathbf{x})}$

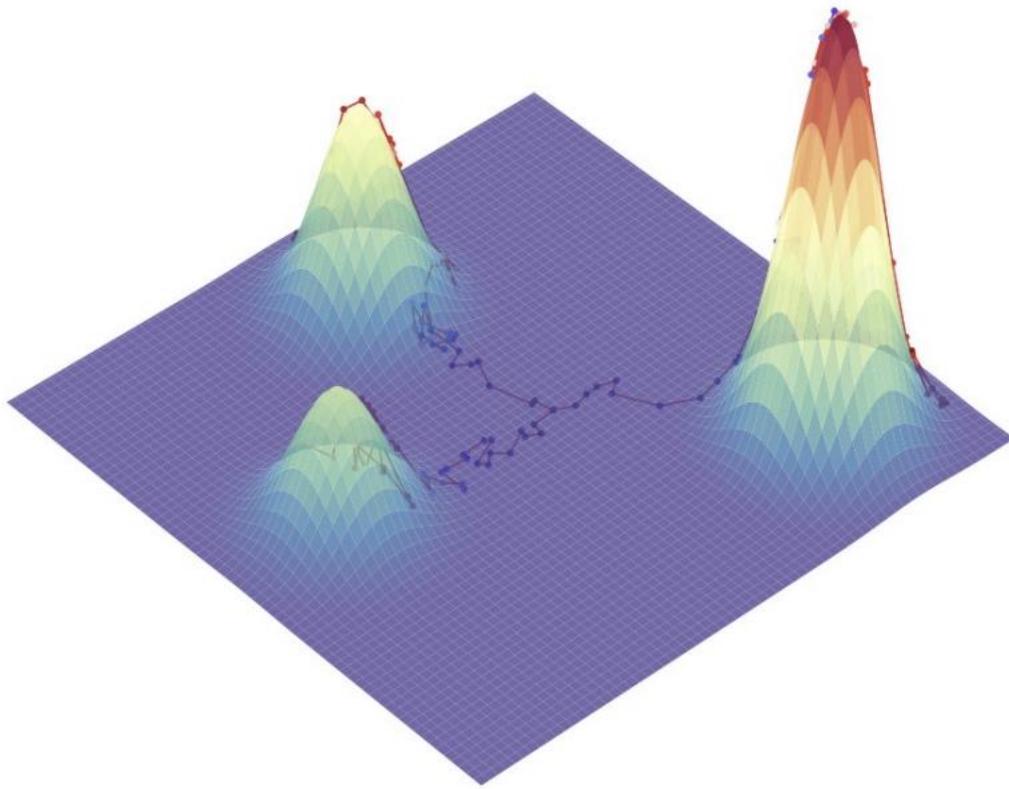
- Score-based model

- $$\begin{aligned} \nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}) &= \nabla_{\mathbf{x}} \log \left( \frac{1}{Z_{\theta}} e^{-f_{\theta}(\mathbf{x})} \right) \\ &= \nabla_{\mathbf{x}} \log \frac{1}{Z_{\theta}} + \nabla_{\mathbf{x}} \log e^{-f_{\theta}(\mathbf{x})} \\ &= -\nabla_{\mathbf{x}} f_{\theta}(\mathbf{x}) \\ &\approx \mathbf{s}_{\theta}(\mathbf{x}) \end{aligned}$$

- Score Matching

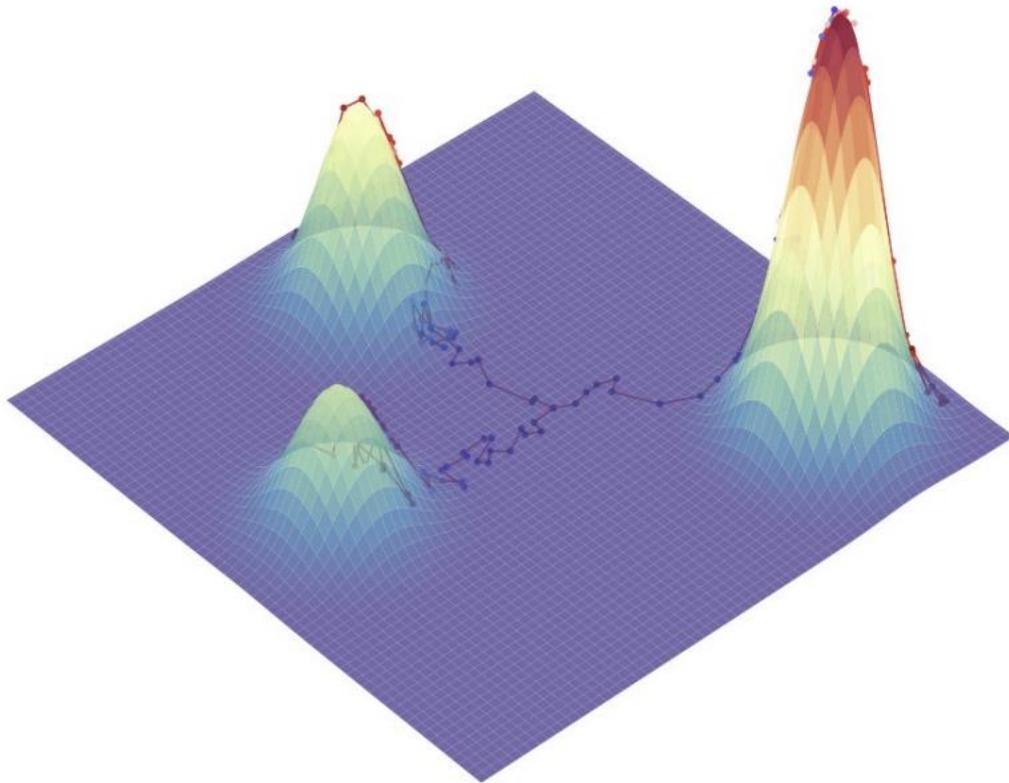
$$\mathbb{E}_{p(\mathbf{x})} \left[ \|\mathbf{s}_{\theta}(\mathbf{x}) - \nabla \log p(\mathbf{x})\|_2^2 \right]$$

## Real Data Distribution



- Sparse
- Disconnected
- Non-overlapping modes

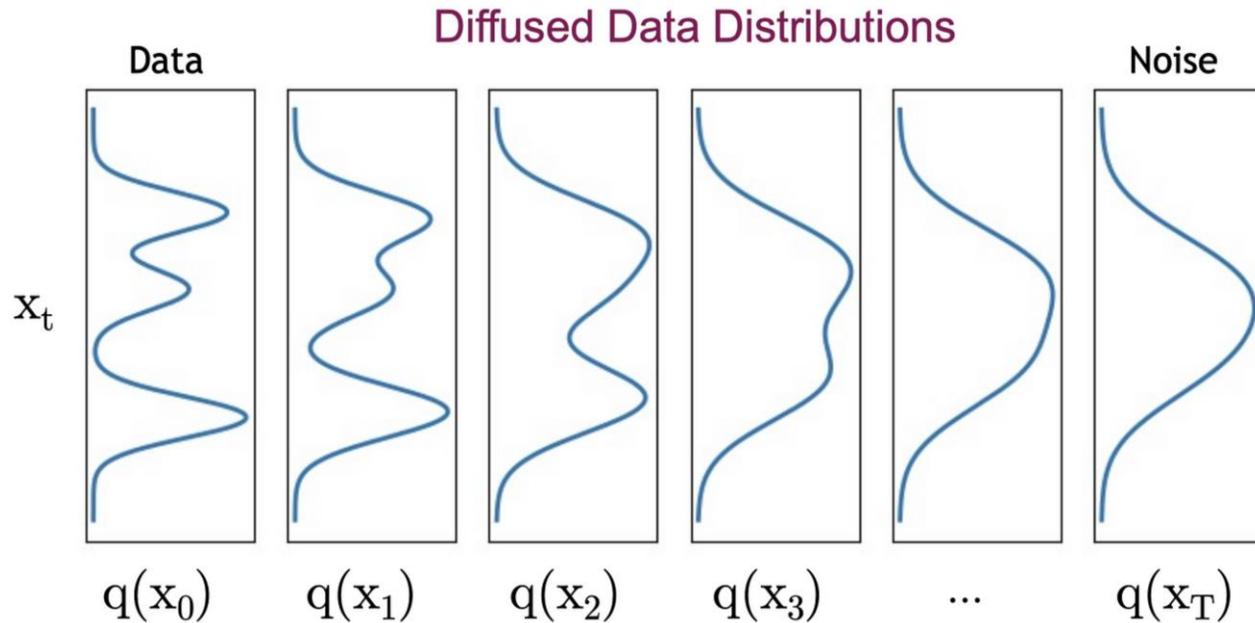
# Vanilla Score Matching



- Limitations:

- Poorly defined for real-world data
- Inaccurate score estimation for low-density region
- Poor sampling with large low-density region

# Score-Based Generative Model



- Extending the distribution

- $$p_{\sigma_t}(\mathbf{x}_t) = \int p(\mathbf{x}) \mathcal{N}(\mathbf{x}_t; \mathbf{x}, \sigma_t^2 \mathbf{I}) d\mathbf{x}$$

- Score Matching for all noise levels

- $$\arg \min_{\theta} \sum_{t=1}^T \lambda(t) \mathbb{E}_{p_{\sigma_t}(\mathbf{x}_t)} \left[ \|\mathbf{s}_{\theta}(\mathbf{x}, t) - \nabla \log p_{\sigma_t}(\mathbf{x}_t)\|_2^2 \right]$$

# Score-Based Generative Model

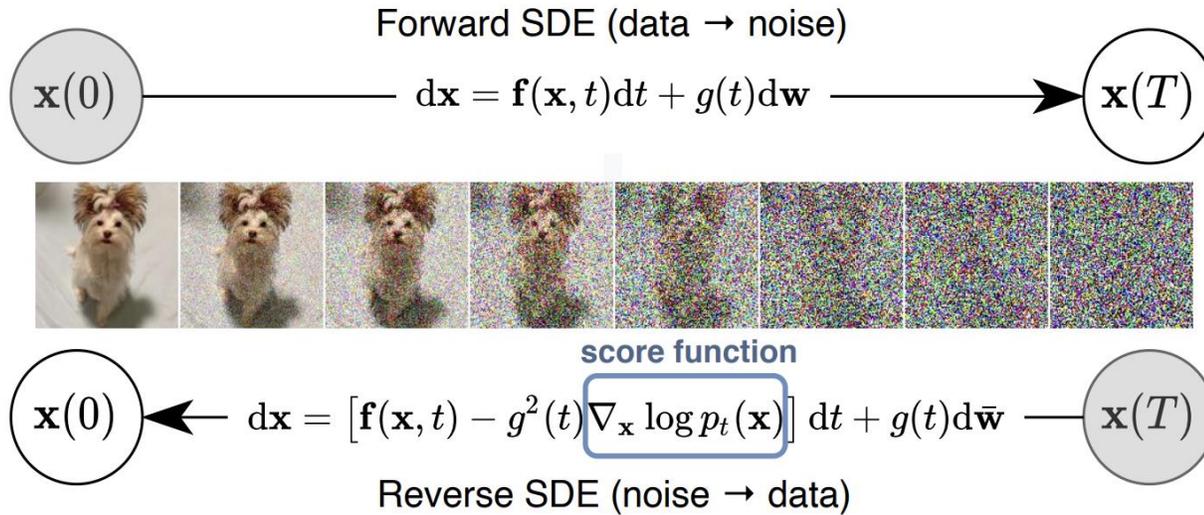
## ■ Limitations of vanilla score matching:

- Poorly defined for real-world data
- Inaccurate score estimation for low-density region
- Poor sampling with large low-density region

## ■ Advantages of score-based generative model

- The support of a Gaussian noise distribution is the entire space.
- Increase the area of each mode by adding noise.
- Different modes are connected by adding noise.

# Diffusion Models: Stochastic Differential Equation Perspective



Probability Flow ODE:

A deterministic reverse process

$$d\mathbf{x} = \left[ \mathbf{f}(\mathbf{x}, t) - \frac{1}{2}g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt$$

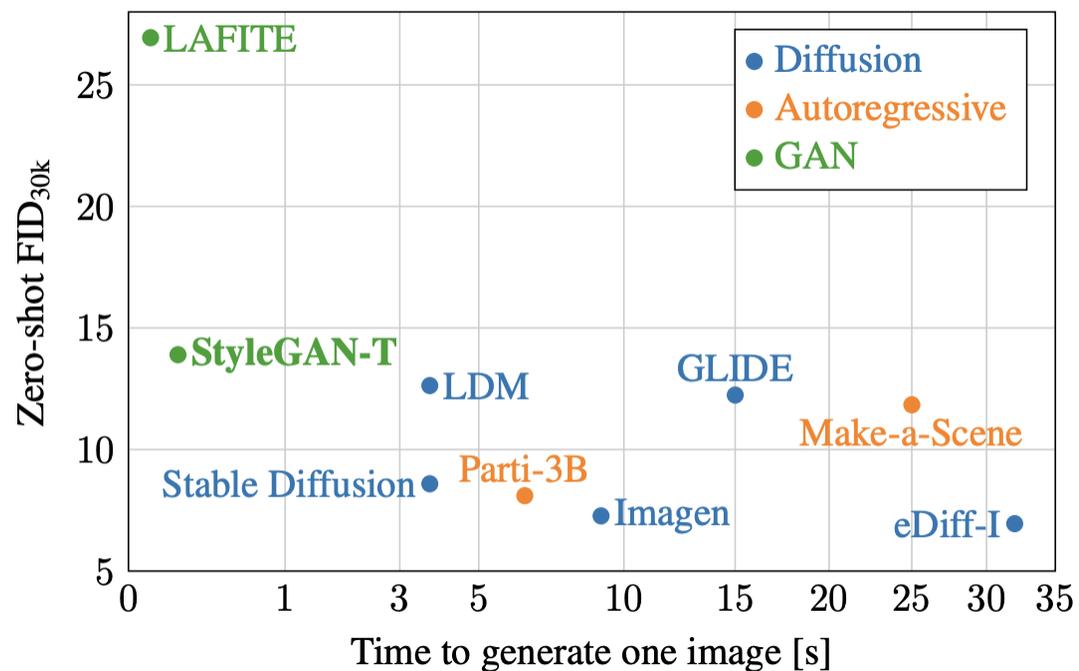
Exact Solution form of PF-ODE

$$\mathbf{x}_t = \frac{\alpha_t}{\alpha_s} \mathbf{x}_s - \alpha_t \int_{\lambda_s}^{\lambda_t} e^{-\lambda} \hat{\boldsymbol{\epsilon}}_{\theta}(\hat{\mathbf{x}}_{\lambda}, \lambda) d\lambda.$$

The only unknown term is the **score function**.

Train a neural network through score matching!

## Diffusion Models: Slow Inference Speed



How to speed up the diffusion generation?

- Reducing the number of function evaluation (NFE).
- Better Solvers.
- Adversarial post-training.
- Parallel Sampling.
- Distillation.
  - Naïve distillation.
  - Guided distillation.
  - Score distillation.
  - Consistency distillation.
  - Rectification.

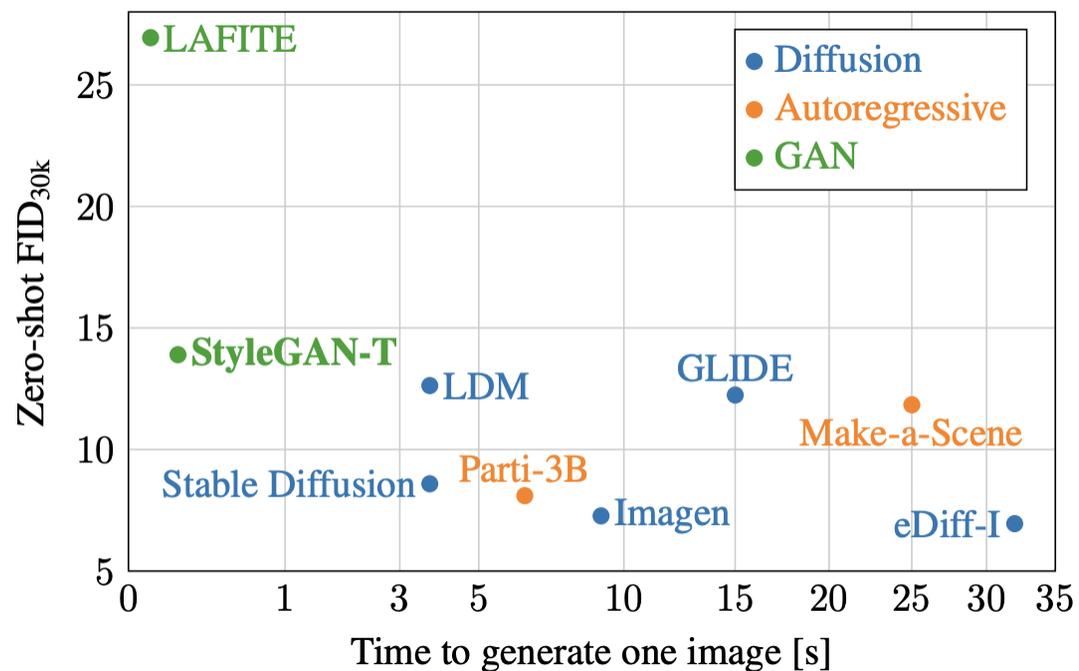
## DPM-Solver

$$\mathbf{x}_{t_{i-1} \rightarrow t_i} = \frac{\alpha_{t_i}}{\alpha_{t_{i-1}}} \tilde{\mathbf{x}}_{t_{i-1}} - \alpha_{t_i} \int_{\lambda_{t_{i-1}}}^{\lambda_{t_i}} e^{-\lambda} \hat{\mathbf{e}}_{\theta}(\hat{\mathbf{x}}_{\lambda}, \lambda) d\lambda.$$

$$\hat{\mathbf{e}}_{\theta}(\hat{\mathbf{x}}_{\lambda}, \lambda) = \sum_{n=0}^{k-1} \frac{(\lambda - \lambda_{t_{i-1}})^n}{n!} \hat{\mathbf{e}}_{\theta}^{(n)}(\hat{\mathbf{x}}_{\lambda_{t_{i-1}}}, \lambda_{t_{i-1}}) + \mathcal{O}((\lambda - \lambda_{t_{i-1}})^k),$$

$$\mathbf{x}_{t_{i-1} \rightarrow t_i} = \frac{\alpha_{t_i}}{\alpha_{t_{i-1}}} \tilde{\mathbf{x}}_{t_{i-1}} - \alpha_{t_i} \sum_{n=0}^{k-1} \hat{\mathbf{e}}_{\theta}^{(n)}(\hat{\mathbf{x}}_{\lambda_{t_{i-1}}}, \lambda_{t_{i-1}}) \int_{\lambda_{t_{i-1}}}^{\lambda_{t_i}} e^{-\lambda} \frac{(\lambda - \lambda_{t_{i-1}})^n}{n!} d\lambda + \mathcal{O}(h_i^{k+1}),$$

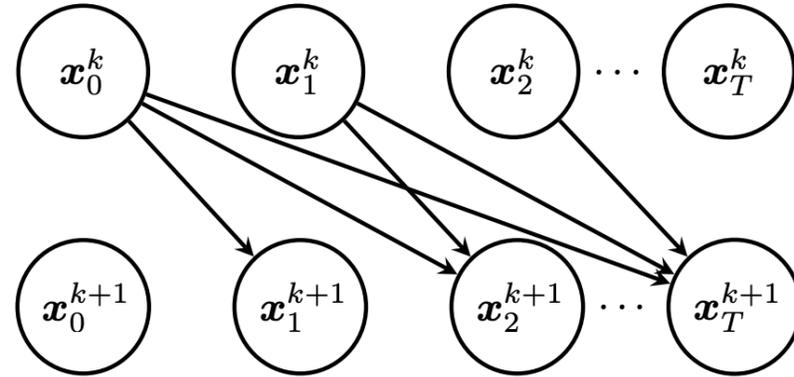
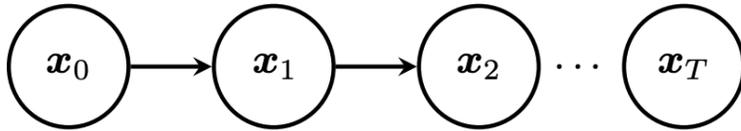
## Diffusion Models: Slow Inference Speed



How to speed up the diffusion generation?

- Reducing the number of function evaluation (NFE).
- Better Solvers.
- Adversarial post-training.
- **Parallel Sampling.**
- Distillation.
  - Naïve distillation.
  - Guided distillation.
  - Score distillation.
  - Consistency distillation.
  - Rectification.

# Picard Iteration



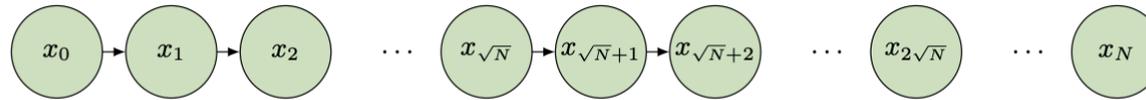
$$\mathbf{x}_t = \mathbf{x}_0 + \int_0^t s(\mathbf{x}_u, u) du.$$

$$\mathbf{x}_t^{k+1} = \mathbf{x}_0^k + \int_0^t s(\mathbf{x}_u^k, u) du.$$

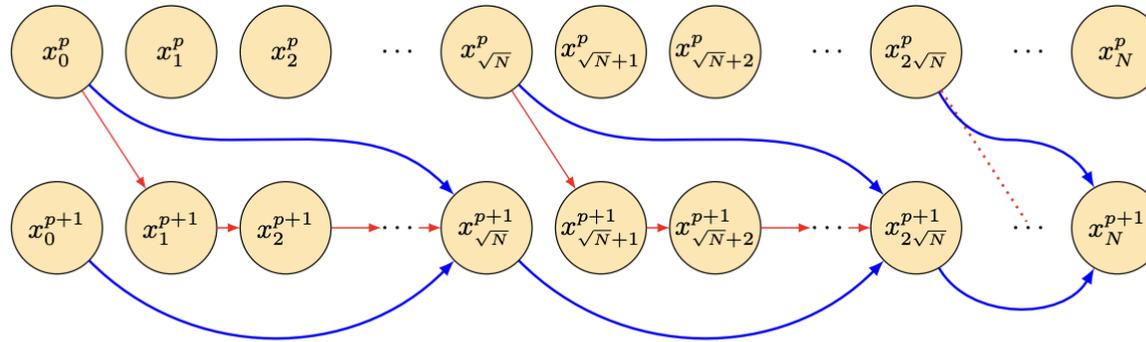
## Lower Bound of Picard Iteration = Sequential Denoising

$$\begin{aligned}\mathbf{x}_{t+1}^{k+1} &= \mathbf{x}_0^k + \frac{1}{T} \sum_{i=0}^t s(\mathbf{x}_i^k, \frac{i}{T}) \\ &= \left( \mathbf{x}_0^k + \frac{1}{T} \sum_{i=0}^{t-1} s(\mathbf{x}_i^k, \frac{i}{T}) \right) + \frac{1}{T} s(\mathbf{x}_t^k, \frac{t}{T}) \\ &= \mathbf{x}_t^{k+1} + \frac{1}{T} s(\mathbf{x}_t^k, \frac{t}{T}) \\ &= \mathbf{x}_t^{k+1} + \frac{1}{T} s(h_{t-1}(\dots h_2(h_1(\mathbf{x}_0))), \frac{t}{T}) \\ &= \mathbf{x}_t^* + \frac{1}{T} s(\mathbf{x}_t^*, \frac{t}{T}) = \mathbf{x}_{t+1}^*.\end{aligned}$$

# Parareal Algorithm



(a) Sequential Sampling



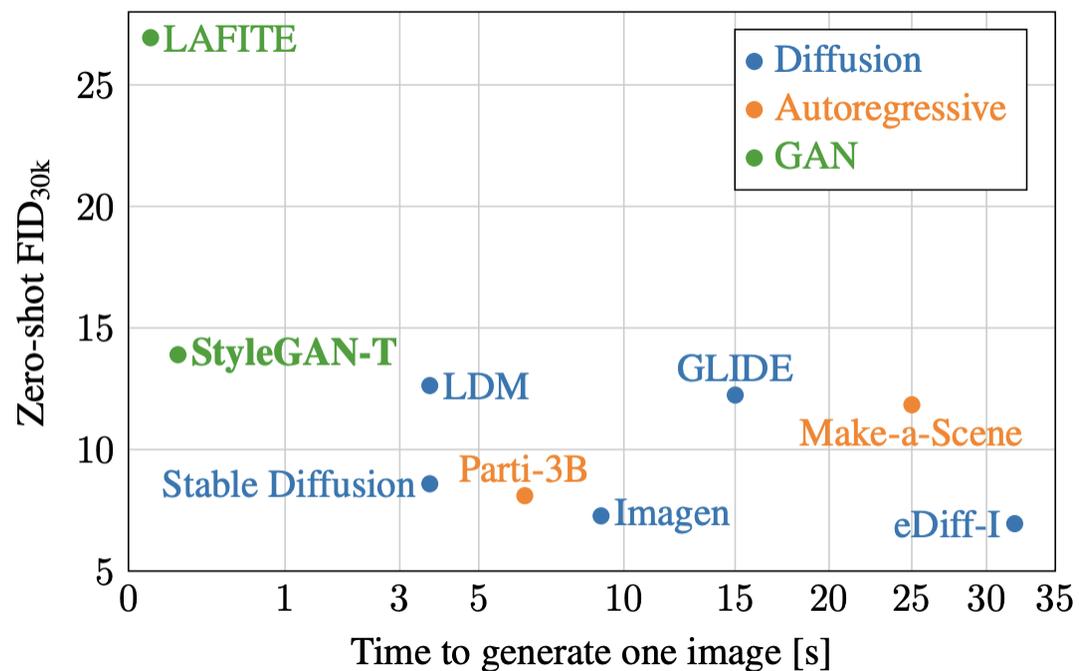
(b) SRDS

$$x_{i+1}^{p+1} = \mathcal{F}(x_i^p, t_i, t_{i+1}) + \left( \mathcal{G}(x_i^{p+1}, t_i, t_{i+1}) - \mathcal{G}(x_i^p, t_i, t_{i+1}) \right)$$

Fine Solver (Parallel)

Coarse Solver (Sequential)

## Diffusion Models: Slow Inference Speed



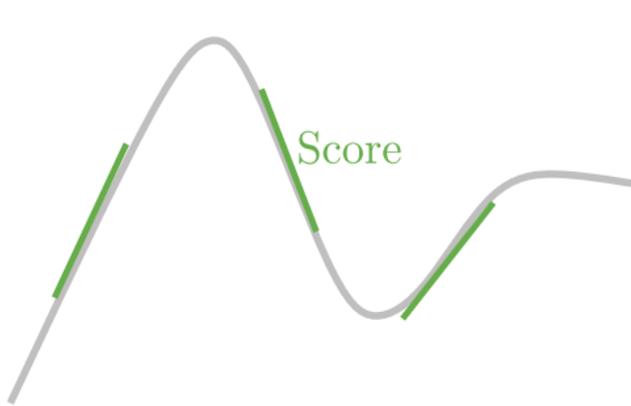
How to speed up the diffusion generation?

- Reducing the number of function evaluation (NFE).
  - Better Solvers.
  - Adversarial post-training.
  - Parallel Sampling.
- Distillation.
    - Naïve distillation.
    - Guided distillation.
    - Score distillation.
    - Consistency distillation.
    - Rectification.

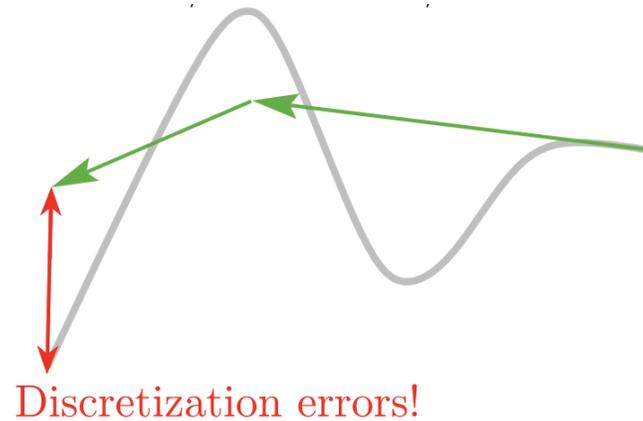
# Understanding Diffusion Models from the PF-ODE path

We know the derivative w.r.t. time  $t$ .

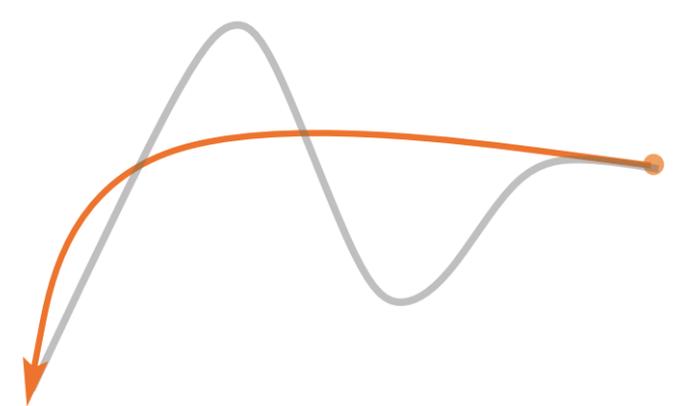
$$d\mathbf{x} = \left[ \mathbf{f}(\mathbf{x}, t) - \frac{1}{2}g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt$$



PF-ODE

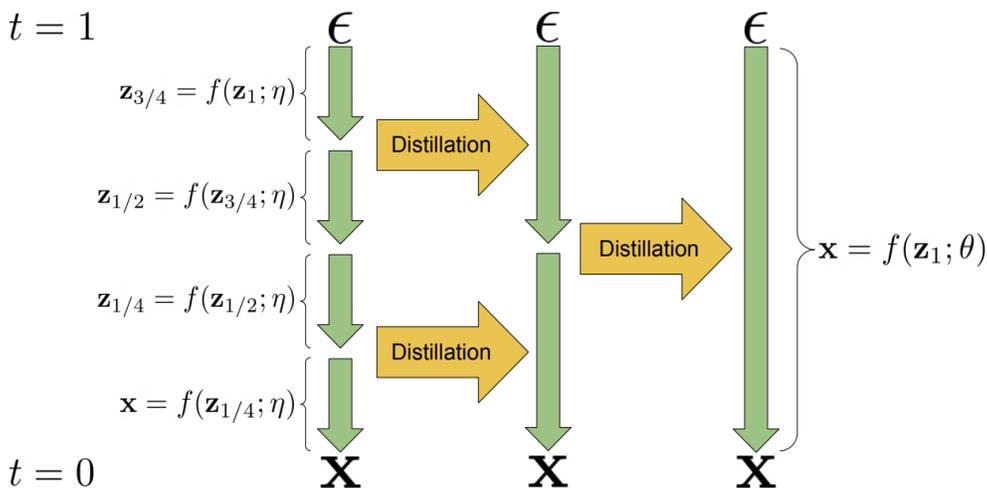


Discretized numerical solving.



Naïve distillation.

# Distillation Techniques: Progressive Distillation




---

## Algorithm 2 Progressive distillation

---

**Require:** Trained teacher model  $\hat{\mathbf{x}}_\eta(\mathbf{z}_t)$

**Require:** Data set  $\mathcal{D}$

**Require:** Loss weight function  $w()$

**Require:** Student sampling steps  $N$

**for**  $K$  iterations **do**

$\theta \leftarrow \eta$        $\triangleright$  Init student from teacher

**while** not converged **do**

$\mathbf{x} \sim \mathcal{D}$

$t = i/N, i \sim \text{Cat}[1, 2, \dots, N]$

$\epsilon \sim N(0, I)$

$\mathbf{z}_t = \alpha_t \mathbf{x} + \sigma_t \epsilon$

**# 2 steps of DDIM with teacher**

$t' = t - 0.5/N, t'' = t - 1/N$

$\mathbf{z}_{t'} = \alpha_{t'} \hat{\mathbf{x}}_\eta(\mathbf{z}_t) + \frac{\sigma_{t'}}{\sigma_t} (\mathbf{z}_t - \alpha_t \hat{\mathbf{x}}_\eta(\mathbf{z}_t))$

$\mathbf{z}_{t''} = \alpha_{t''} \hat{\mathbf{x}}_\eta(\mathbf{z}_{t'}) + \frac{\sigma_{t''}}{\sigma_{t'}} (\mathbf{z}_{t'} - \alpha_{t'} \hat{\mathbf{x}}_\eta(\mathbf{z}_{t'}))$

$\tilde{\mathbf{x}} = \frac{\mathbf{z}_{t''} - (\sigma_{t''}/\sigma_t) \mathbf{z}_t}{\alpha_{t''} - (\sigma_{t''}/\sigma_t) \alpha_t}$        $\triangleright$  Teacher  $\hat{\mathbf{x}}$  target

$\lambda_t = \log[\alpha_t^2 / \sigma_t^2]$

$L_\theta = w(\lambda_t) \|\tilde{\mathbf{x}} - \hat{\mathbf{x}}_\theta(\mathbf{z}_t)\|_2^2$

$\theta \leftarrow \theta - \gamma \nabla_\theta L_\theta$

**end while**

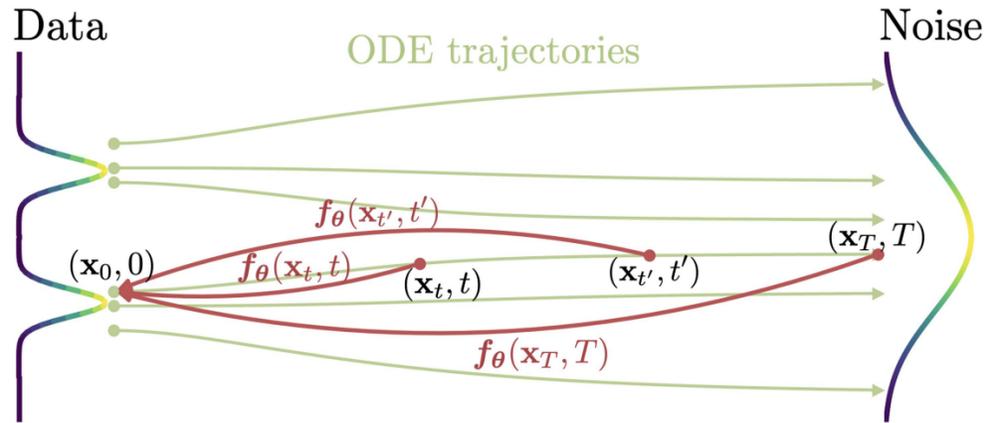
$\eta \leftarrow \theta$        $\triangleright$  Student becomes next teacher

$N \leftarrow N/2$        $\triangleright$  Halve number of sampling steps

**end for**

---

# Distillation Techniques: Consistency Distillation




---

## Algorithm 2 Consistency Distillation (CD)

---

**Input:** dataset  $\mathcal{D}$ , initial model parameter  $\theta$ , learning rate  $\eta$ , ODE solver  $\Phi(\cdot, \cdot; \phi)$ ,  $d(\cdot, \cdot)$ ,  $\lambda(\cdot)$ , and  $\mu$

$\theta^- \leftarrow \theta$

**repeat**

  Sample  $\mathbf{x} \sim \mathcal{D}$  and  $n \sim \mathcal{U}[[1, N - 1]]$

  Sample  $\mathbf{x}_{t_{n+1}} \sim \mathcal{N}(\mathbf{x}; t_{n+1}^2 \mathbf{I})$

$\hat{\mathbf{x}}_{t_n}^{\phi} \leftarrow \mathbf{x}_{t_{n+1}} + (t_n - t_{n+1})\Phi(\mathbf{x}_{t_{n+1}}, t_{n+1}; \phi)$

$\mathcal{L}(\theta, \theta^-; \phi) \leftarrow$

$\lambda(t_n)d(\mathbf{f}_{\theta}(\mathbf{x}_{t_{n+1}}, t_{n+1}), \mathbf{f}_{\theta^-}(\hat{\mathbf{x}}_{t_n}^{\phi}, t_n))$

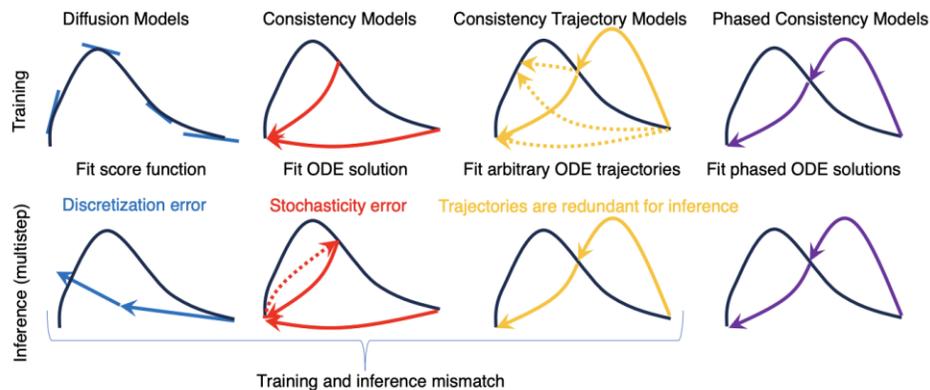
$\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}(\theta, \theta^-; \phi)$

$\theta^- \leftarrow \text{stopgrad}(\mu \theta^- + (1 - \mu)\theta)$

**until** convergence

---

# Distillation Techniques: Phased Consistency Distillation




---

## Algorithm 1 Phased Consistency Distillation with CFG-augmented ODE solver (PCD)

---

**Input:** dataset  $\mathcal{D}$ , initial model parameter  $\theta$ , learning rate  $\eta$ , ODE solver  $\Psi(\cdot, \cdot, \cdot, \cdot)$ , distance metric  $d(\cdot, \cdot)$ , EMA rate  $\mu$ , noise schedule  $\alpha_t, \sigma_t$ , guidance scale  $[w_{\min}, w_{\max}]$ , number of ODE step  $k$ , discretized timesteps  $t_0 = \epsilon < t_1 < t_2 < \dots < t_N = T$ , edge timesteps  $s_0 = t_0 < s_1 < s_2 < \dots < s_M = t_N \in \{t_i\}_{i=0}^N$  to split the ODE trajectory into  $M$  sub-trajectories.

Training data :  $\mathcal{D}_x = \{(\mathbf{x}, \mathbf{c})\}$

$\theta^- \leftarrow \theta$

**repeat**

  Sample  $(\mathbf{z}, \mathbf{c}) \sim \mathcal{D}_z, n \sim \mathcal{U}[0, N - k]$  and  $\omega \sim [\omega_{\min}, \omega_{\max}]$

  Sample  $\mathbf{x}_{t_{n+k}} \sim \mathcal{N}(\alpha_{t_{n+k}} \mathbf{z}; \sigma_{t_{n+k}}^2 \mathbf{I})$

  Determine  $[s_m, s_{m+1}]$  given  $n$

$\mathbf{x}_{t_n}^\phi \leftarrow (1 + \omega) \Psi(\mathbf{x}_{t_{n+k}}, t_{n+k}, t_n, \mathbf{c}) - \omega \Psi(\mathbf{x}_{t_{n+k}}, t_{n+k}, t_n, \emptyset)$

$\tilde{\mathbf{x}}_{s_m} = \mathbf{f}_{\theta}^m(\mathbf{x}_{t_{n+k}}, t_{n+k}, \mathbf{c})$  and  $\hat{\mathbf{x}}_{s_m} = \mathbf{f}_{\theta^-}(\mathbf{x}_{t_n}^\phi, t_n, \mathbf{c})$

  Obtain  $\tilde{\mathbf{x}}_s$  and  $\hat{\mathbf{x}}_s$  through adding noise to  $\tilde{\mathbf{x}}_{s_m}$  and  $\hat{\mathbf{x}}_{s_m}$

$\mathcal{L}(\theta, \theta^-) = d(\tilde{\mathbf{x}}_{s_m}, \hat{\mathbf{x}}_{s_m}) + \lambda(\text{ReLU}(1 + \tilde{\mathbf{x}}_s) + \text{ReLU}(1 - \hat{\mathbf{x}}_s))$

$\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}(\theta, \theta^-)$

$\theta^- \leftarrow \text{stopgrad}(\mu \theta^- + (1 - \mu) \theta)$

**until** convergence

---

# Application: AnimateLCM

- [AnimateLCM](#) support
  - NOTE: You will need to use `autoselect` or `lcm` or `lcm[100_ots]` `beta_schedule`. To use fully with LCM, be sure to use appropriate LCM lora, use the `lcm` `sampler_name` in KSampler nodes, and lower `cfg` to somewhere around 1.0 to 2.0. Don't forget to decrease steps (minimum = ~4 steps), since LCM converges faster (less steps). Increase step count to increase detail as desired.
- [AnimateLCM-I2V](#) support, big thanks to [Fu-Yun Wang](#) for providing me the original diffusers code he created during his work on the paper
  - NOTE: Requires same settings as described for AnimateLCM above. Requires `Apply AnimateLCM-I2V Model Gen2` node usage so that `ref_latent` can be provided; use `Scale Ref Image` and `VAE Encode` node to preprocess input images. While this was intended as an `img2video` model, I found it works best for `vid2vid` purposes with `ref_drift=0.0`, and to use it for only at least 1 step before switching over to other models via chaining with `toher Apply AnimateDiff Model (Adv.)` nodes. The `apply_ref_when_disabled` can be set to `True` to allow the `img_encoder` to do its thing even when the `end_percent` is reached. AnimateLCM-I2V is also extremely useful for maintaining coherence at higher resolutions (with ControlNet and SD LoRAs active, I could easily upscale from 512x512 source to 1024x1024 in a single pass). TODO: add examples

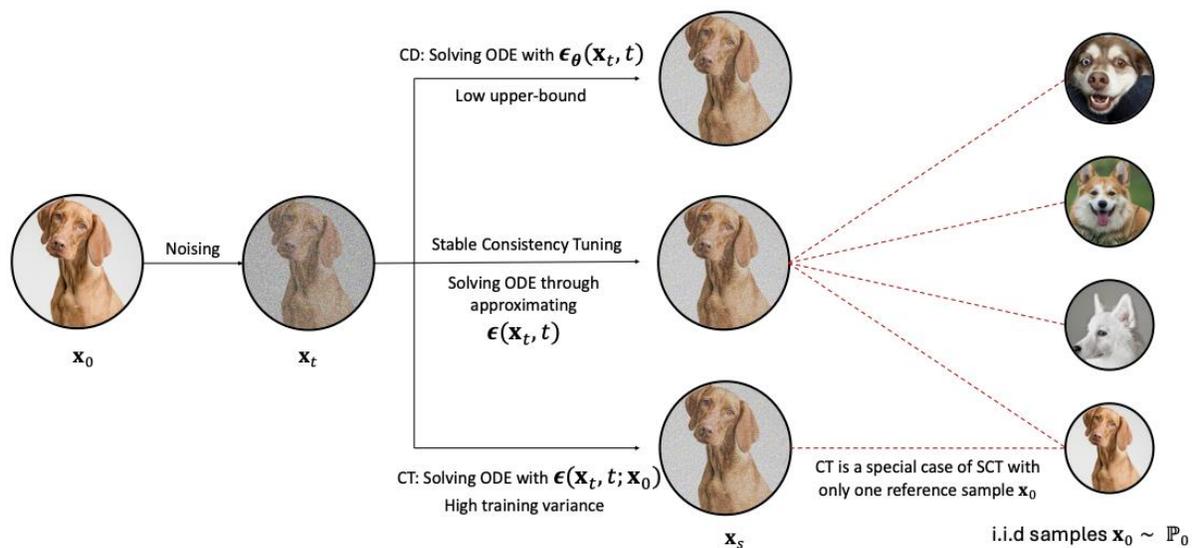
Downloads last month

**73,760**





# Consistency Training



$$h_{\theta}(\mathbf{x}_t, t) \xleftarrow{\text{fit}} \mathbf{r} + h_{\theta-}(\mathbf{x}_r, r)$$

Bootstrapping

$$\mathbf{r} \approx \epsilon_{\phi}(\mathbf{x}_t, t) \int_{\lambda_t}^{\lambda_r} e^{-\lambda} d\lambda + \mathcal{O}((\lambda_r - \lambda_t)^2)$$

Consistency Distillation

$$\mathbf{r} \approx \epsilon(\mathbf{x}_t, t; \mathbf{x}_0) \int_{\lambda_t}^{\lambda_r} e^{-\lambda} d\lambda + \mathcal{O}((\lambda_r - \lambda_t)^2)$$

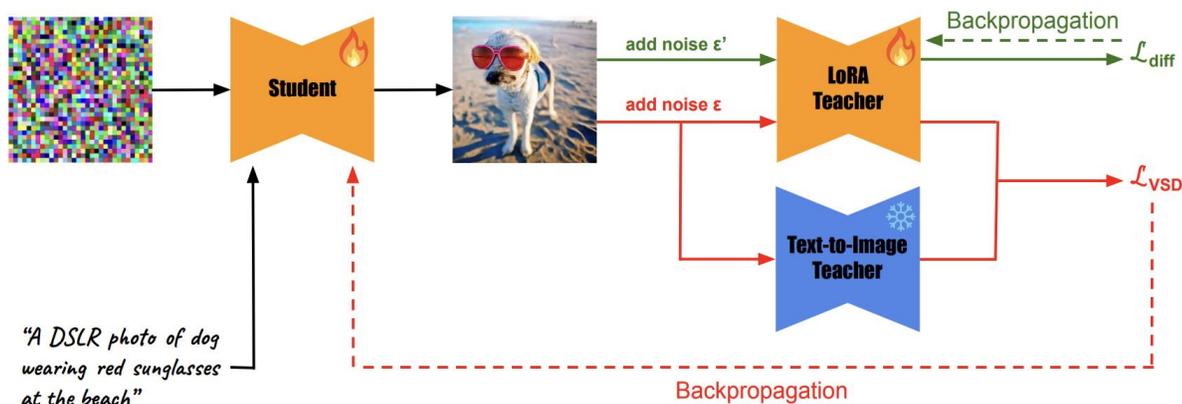
Consistency Training

## Ground Truth of Score Estimation: Stable Consistency Tuning

$$\begin{aligned}
 \nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{c}) &= \mathbb{E}_{\mathbb{P}(\mathbf{x}_0 \mid \mathbf{x}_t, \mathbf{c})} [\nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{x}_0, \mathbf{c})] \\
 &= \mathbb{E}_{\mathbb{P}(\mathbf{x}_0 \mid \mathbf{c})} \left[ \frac{\mathbb{P}(\mathbf{x}_0 \mid \mathbf{x}_t, \mathbf{c})}{\mathbb{P}(\mathbf{x}_0 \mid \mathbf{c})} \nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{x}_0, \mathbf{c}) \right] \\
 &= \mathbb{E}_{\mathbb{P}(\mathbf{x}_0 \mid \mathbf{c})} \left[ \frac{\mathbb{P}(\mathbf{x}_t \mid \mathbf{x}_0, \mathbf{c})}{\mathbb{P}(\mathbf{x}_t \mid \mathbf{c})} \nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{x}_0, \mathbf{c}) \right] \\
 &= \mathbb{E}_{\mathbb{P}(\mathbf{x}_0 \mid \mathbf{c})} \left[ \frac{\mathbb{P}(\mathbf{x}_t \mid \mathbf{x}_0)}{\mathbb{P}(\mathbf{x}_t \mid \mathbf{c})} \nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{x}_0) \right] \\
 &\approx \frac{1}{n} \sum_{\substack{i=0, \dots, n-1 \\ \{\mathbf{x}_0^{(i)}\} \sim \mathbb{P}(\mathbf{x}_0 \mid \mathbf{c})}} \frac{\mathbb{P}(\mathbf{x}_t \mid \mathbf{x}_0^{(i)})}{\mathbb{P}(\mathbf{x}_t \mid \mathbf{c})} \nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{x}_0^{(i)}) \\
 &\approx \frac{1}{n} \sum_{\substack{i=0, \dots, n-1 \\ \{\mathbf{x}_0^{(i)}\} \sim \mathbb{P}(\mathbf{x}_0 \mid \mathbf{c})}} \frac{\mathbb{P}(\mathbf{x}_t \mid \mathbf{x}_0^{(i)})}{\sum_{\mathbf{x}_0^{(j)} \in \{\mathbf{x}_0^{(i)}\}} \mathbb{P}(\mathbf{x}_t \mid \mathbf{x}_0^{(j)}, \mathbf{c})} \nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{x}_0^{(i)}) \\
 &= \frac{1}{n} \sum_{\substack{i=0, \dots, n-1 \\ \{\mathbf{x}_0^{(i)}\} \sim \mathbb{P}(\mathbf{x}_0 \mid \mathbf{c})}} \frac{\mathbb{P}(\mathbf{x}_t \mid \mathbf{x}_0^{(i)})}{\sum_{\mathbf{x}_0^{(j)} \in \{\mathbf{x}_0^{(i)}\}} \mathbb{P}(\mathbf{x}_t \mid \mathbf{x}_0^{(j)})} \nabla_{\mathbf{x}_t} \log \mathbb{P}_t(\mathbf{x}_t \mid \mathbf{x}_0^{(i)})
 \end{aligned}$$

# Distillation Techniques: Score Distillation

$$\nabla_{\theta} \mathcal{L}_{\text{VSD}}(\theta) \triangleq \mathbb{E}_{t, \epsilon, c} \left[ \omega(t) (\epsilon_{\text{pretrain}}(\mathbf{x}_t, t, y^c) - \epsilon_{\phi}(\mathbf{x}_t, t, c, y)) \frac{\partial g(\theta, c)}{\partial \theta} \right]$$




---

## Algorithm 1 SwiftBrush Distillation

---

- 1: **Require:** a pretrained text-to-image teacher  $\epsilon_{\psi}$ , a LoRA teacher  $\epsilon_{\phi}$ , a student model  $f_{\theta}$ , two learning rates  $\eta_1$  and  $\eta_2$ , a weighting function  $\omega$ , a prompts dataset  $Y$ , the maximum number of time steps  $T$  and the noise schedule  $\{(\alpha_t, \sigma_t)\}_{t=1}^T$  of the teacher model
  - 2: **Initialize:**  $\phi \leftarrow \psi, \theta \leftarrow \psi$
  - 3: **while not converged do**
  - 4:   Sample input noise  $z \sim \mathcal{N}(0, I)$
  - 5:   Sample text caption input  $y \sim Y$
  - 6:   Compute student output  $\hat{x}_0 = f_{\theta}(z, y)$
  - 7:   Sample timestep  $t \sim \mathcal{U}(0.02T, 0.98T)$
  - 8:   Sample added noise  $\epsilon \sim \mathcal{N}(0, I)$
  - 9:   Compute noisy sample  $\hat{x}_t = \alpha_t \hat{x}_0 + \sigma_t \epsilon$
  - 10:    $\theta \leftarrow \theta - \eta_1 \left[ \omega(t) (\epsilon_{\psi}(\hat{x}_t, t, y) - \epsilon_{\phi}(\hat{x}_t, t, y)) \frac{\partial \hat{x}_0}{\partial \theta} \right]$
  - 11:   Sample timestep  $t' \sim \mathcal{U}(0, T)$
  - 12:   Sample added noise  $\epsilon' \sim \mathcal{N}(0, I)$
  - 13:   Compute noisy sample  $\hat{x}_{t'} = \alpha_{t'} \hat{x}_0 + \sigma_{t'} \epsilon'$
  - 14:    $\phi \leftarrow \phi - \eta_2 \nabla_{\phi} \|\epsilon_{\phi}(\hat{x}_{t'}, t', y) - \epsilon'\|^2$
  - 15: **end while**
  - 16: **return** trained student model  $f_{\theta}$
-

## Distillation Techniques: Score Distillation

$$\mathcal{L}(\theta) = \mathcal{D}^{[0,T]}(p_\theta, q) = \int_{t=0}^T w(t) \mathbb{E}_{\mathbf{x}_t \sim \pi_t} [\mathbf{d}(\mathbf{s}_{p_{\theta,t}}(\mathbf{x}_t) - \mathbf{s}_{q_t}(\mathbf{x}_t))] dt,$$

Generator  $g_\theta : p_z \rightarrow p_\theta$ ,  $p_{\theta,t} = p_\theta * \mathcal{N}(0, I)$ ,  $\mathbf{s}_{\theta,t}(\mathbf{x}_t) = \nabla_{\mathbf{x}_t} \log p_{\theta,t}(\mathbf{x}_t)$

impossible to compute  $\frac{d}{d\theta} \mathbf{s}_{\theta,t}(\mathbf{x}_t)$

## Score Divergence Gradient Theorem

$$\mathcal{L}(\theta) = \mathcal{D}^{[0,T]}(p_\theta, q) = \int_{t=0}^T w(t) \mathbb{E}_{\mathbf{x}_t \sim \pi_t} [\mathbf{d}(\mathbf{s}_{p_{\theta,t}}(\mathbf{x}_t) - \mathbf{s}_{q_t}(\mathbf{x}_t))] dt,$$

$$\mathbb{E}_{\mathbf{x}_t \sim p_{\text{sg}[\theta],t}} \left[ \mathbf{d}'(\mathbf{s}_{p_{\theta,t}}(\mathbf{x}_t) - \mathbf{s}_{q_t}(\mathbf{x}_t)) \frac{\partial}{\partial \theta} \mathbf{s}_{p_{\theta,t}}(\mathbf{x}_t) \right] \quad (3.6)$$

$$= -\frac{\partial}{\partial \theta} \mathbb{E}_{\substack{\mathbf{x}_0 \sim p_{\theta,0}, \\ \mathbf{x}_t | \mathbf{x}_0 \sim q_t(\mathbf{x}_t | \mathbf{x}_0)}} \left[ \left\{ \mathbf{d}'(\mathbf{s}_{p_{\text{sg}[\theta],t}}(\mathbf{x}_t) - \mathbf{s}_{q_t}(\mathbf{x}_t)) \right\}^T \left\{ \mathbf{s}_{p_{\text{sg}[\theta],t}}(\mathbf{x}_t) - \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | \mathbf{x}_0) \right\} \right].$$

Simplify

$$\mathbf{d}_\psi(\mathbf{x}_t, t) - \mathbf{x}_0$$

Ignore

# Application: Casual Vid

## Bidirectional Teacher

Preparing...

Progress: 0/1

```
16] -1/-1/-1->0->-1 [17] -1/-1/-1->0->-1 [18] -1/-1/-1->0->-1 [19] -1/-1/-1->0->-1 [20] -1/-1/-1->0->-1 [21] -1/-1/-1->0->-1 [22] -1/-1/-1->0->-1 [23] -1/-1/-1->0->-1 [24] -1/-1/-1->0->-1 [25]
-1/-1/-1->0->-1 [26] -1/-1/-1->0->-1 [27] -1/-1/-1->0->-1 [28] -1/-1/-1->0->-1 [29] -1/-1/-1->0->-1 [30] -1/-1/-1->0->-1 [31] -1/-1/-1->0->-1
xuhuang-0464291930-0-0:504194:506058 [0] NCCL INFO P2P Chunksize set to 131072
xuhuang-0464291930-0-0:504194:506058 [0] NCCL INFO Connected all rings
xuhuang-0464291930-0-0:504194:506058 [0] NCCL INFO Connected all trees
xuhuang-0464291930-0-0:504194:506058 [0] NCCL INFO 32 coll channels, 32 collnet channels, 0 nvls
channels, 32 p2p channels, 32 p2p channels per peer
xuhuang-0464291930-0-0:504194:506058 [0] NCCL INFO TUNER/Plugin: Failed to find ncclTunerPlugin_
v2, using internal tuner instead.
xuhuang-0464291930-0-0:504194:506058 [0] NCCL INFO ncclCommInitRank comm 0x55a809dbc940 rank 0 n
ranks 1 cudaDev 0 nmlDev 0 busId 53000 commId 0x54a120ee0c121fbc - Init COMPLETE
```

## CausVid (Ours)

Preparing...

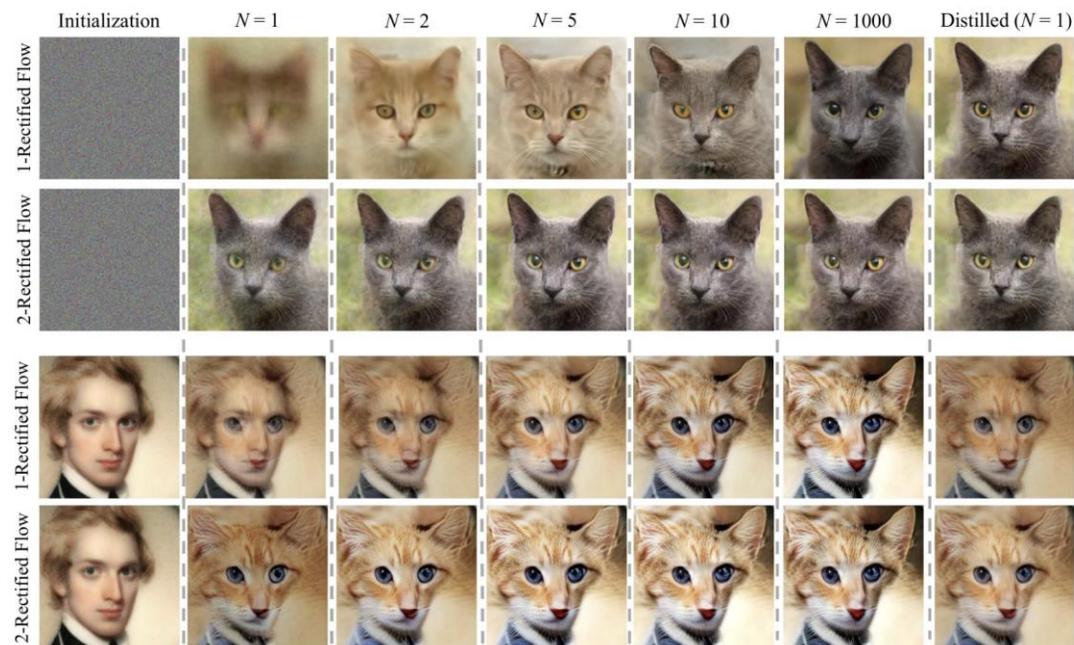
Progress: 0/15

```
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 5 device #3 0000:a4:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 6 device #0 0000:b8:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 6 device #1 0000:b7:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 6 device #2 0000:b6:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 6 device #3 0000:b5:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 7 device #0 0000:c9:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 7 device #1 0000:c8:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 7 device #2 0000:c7:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI NIC group 7 device #3 0000:c6:00.0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO NET/OFI Libfabric provider associates MRs w
ith domains
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO Using non-device net plugin version 0
xuhuang-0464291930-0-0:472752:473617 [0] NCCL INFO Using network AWS Libfabric
```



00:00.

# Distillation Techniques: Rectified Flow



## Advantages:

- High-quality few-step generation.
- Flexibility on inference steps.
- Simple forms.

# Distillation Techniques: Rectified Flow

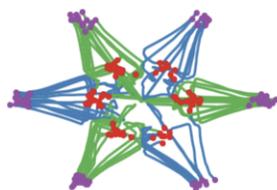
- Linear interpolation.

$$X_t = tX_1 + (1 - t)X_0$$

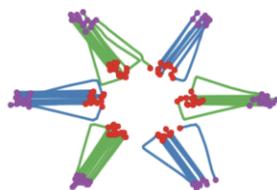
- $v$ -prediction.

$$dX_t = (X_1 - X_0)dt$$

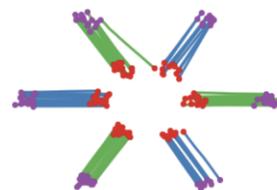
- Rectification (Reflow).



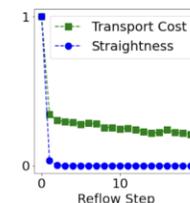
(a) The 1st rectified flow  $Z^1$   
 $Z^1 = \text{RectFlow}((X_0, X_1))$



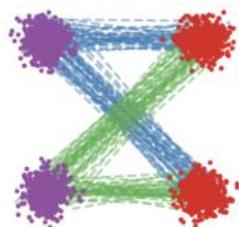
(b) Reflow  $Z^2$   
 $Z^2 = \text{RectFlow}((Z_0^1, Z_1^1))$



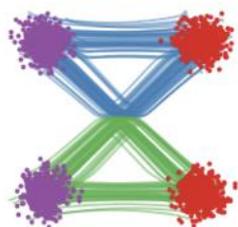
(c) Reflow  $Z^3$   
 $Z^3 = \text{RectFlow}((Z_0^2, Z_1^2))$



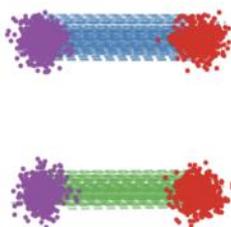
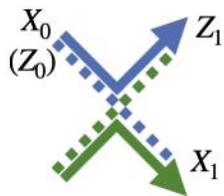
(d) Transport cost, Straightness



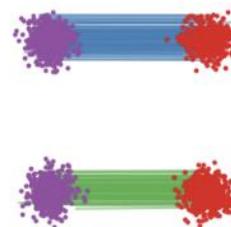
(a) Linear interpolation  
 $X_t = tX_1 + (1 - t)X_0$



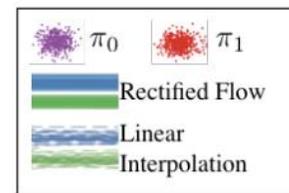
(b) Rectified flow  $Z_t$   
induced by  $(X_0, X_1)$



(c) Linear interpolation  
 $Z_t = tZ_1 + (1 - t)Z_0$

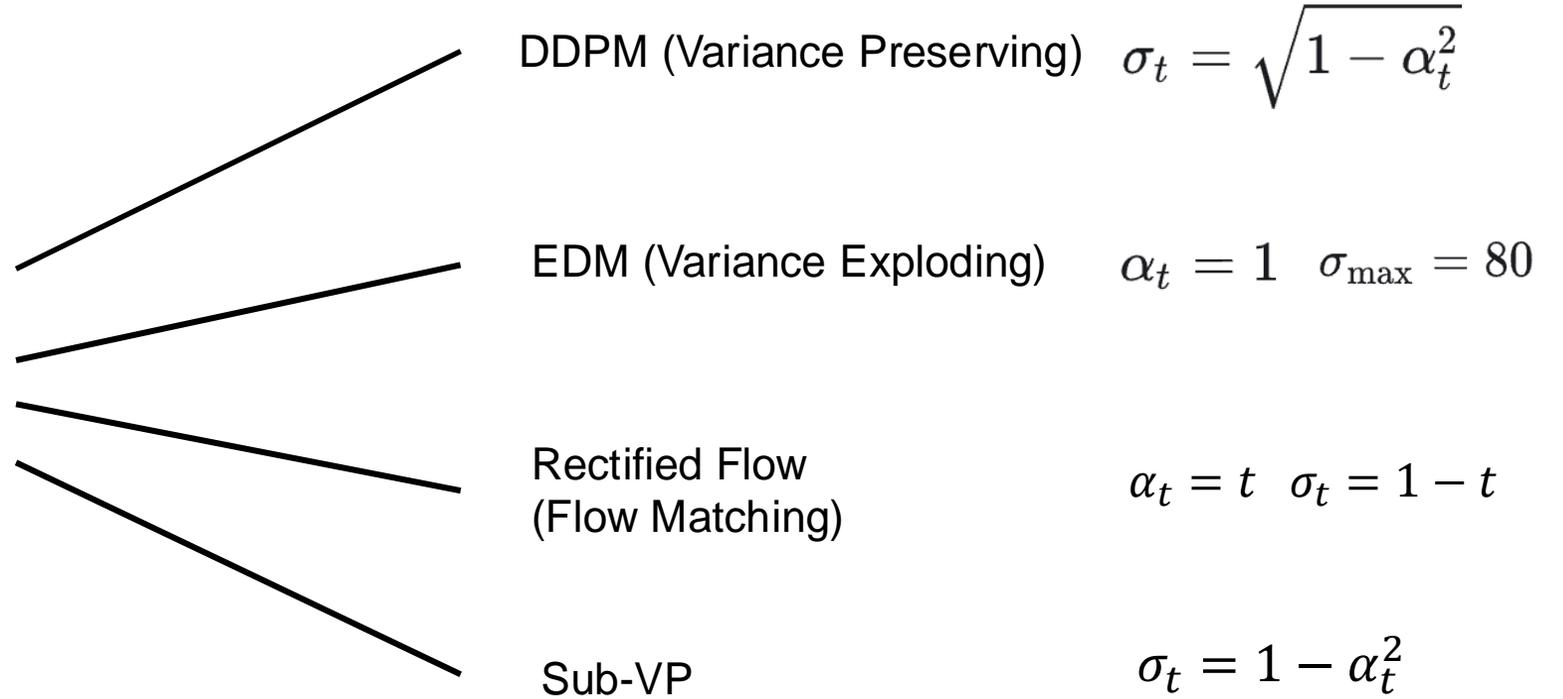


(d) Rectified flow  $Z'_t$   
induced by  $(Z_0, Z_1)$



# Diffusion Models: A (relative) Unified Perspective

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}$$



# The Magic of Rectified Flow: Retraining with Matched Noise-Sample Pairs

---

## Algorithm 1 Flow Matching $v$ -Prediction

---

**Input:**

Sample  $\mathbf{x}_0$  from the data distribution

Sample time  $t$  from a predefined schedule or uniformly from  $[0, 1]$

Sample noise  $\epsilon$  from normal distribution

Compute  $\mathbf{x}_t$ :  $\mathbf{x}_t = (1 - t) \cdot \mathbf{x}_0 + t \cdot \epsilon$

Predict velocity  $\hat{v}$  using the model:  $\hat{v} = \text{Model}(\mathbf{x}_t, t)$

Compute loss:  $\mathcal{L} = \|\hat{v} - (\mathbf{x}_0 - \epsilon)\|_2^2$

Backpropagate and update parameters

---

---

## Algorithm 3 Rectified Flow $v$ -Prediction

---

**Input: noise-data pair  $(\epsilon, \hat{\mathbf{x}}_0)$**

~~Sample  $\mathbf{x}_0$  from the data distribution~~

Sample time  $t$  from a predefined schedule or uniformly from  $[0, 1]$

~~Sample noise  $\epsilon$  from normal distribution~~

Compute  $\mathbf{x}_t$ :  $\mathbf{x}_t = (1 - t) \cdot \hat{\mathbf{x}}_0 + t \cdot \epsilon$

Predict velocity  $\hat{v}$  using the model:  $\hat{v} = \text{Model}(\mathbf{x}_t, t)$

Compute loss:  $\mathcal{L} = \|\hat{v} - (\hat{\mathbf{x}}_0 - \epsilon)\|_2^2$

Backpropagate and update parameters

---

# Rectified Flow Training Is a Subset of Diffusion Training

---

## Algorithm 1 Flow Matching $v$ -Prediction

---

**Input:**

Sample  $\mathbf{x}_0$  from the data distribution

Sample time  $t$  from a predefined schedule or uniformly from  $[0, 1]$

Sample noise  $\epsilon$  from normal distribution

Compute  $\mathbf{x}_t$ :  $\mathbf{x}_t = (1 - t) \cdot \mathbf{x}_0 + t \cdot \epsilon$

Predict velocity  $\hat{v}$  using the model:  $\hat{v} = \text{Model}(\mathbf{x}_t, t)$

Compute loss:  $\mathcal{L} = \|\hat{v} - (\mathbf{x}_0 - \epsilon)\|_2^2$

Backpropagate and update parameters

---

---

## Algorithm 2 Diffusion Training $\epsilon$ -Prediction

---

**Input:**  $\alpha_t, \sigma_t$

Sample  $\mathbf{x}_0$  from the data distribution

Sample time  $t$  from a predefined schedule or uniformly from  $[0, 1]$

Sample noise  $\epsilon$  from normal distribution

Compute  $\mathbf{x}_t$ :  $\mathbf{x}_t = \alpha_t \cdot \mathbf{x}_0 + \sigma_t \cdot \epsilon$

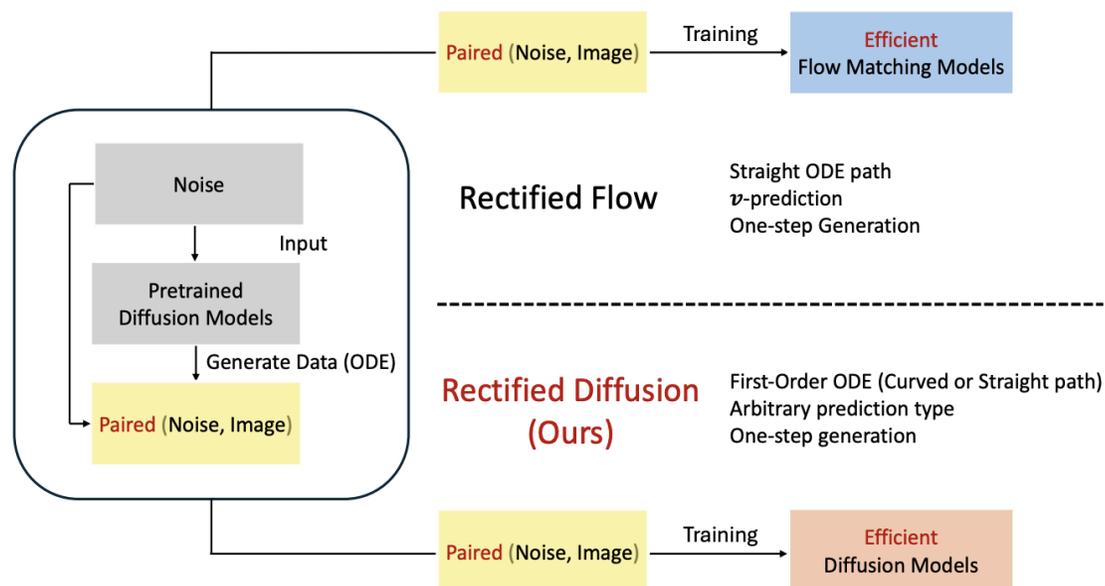
Predict noise  $\hat{\epsilon}$  using the model:  $\hat{\epsilon} = \text{Model}(\mathbf{x}_t, t)$

Compute loss:  $\mathcal{L} = \|\hat{\epsilon} - \epsilon\|_2^2$

Backpropagate and update parameters

---

# Rectified Diffusion: Extending Rectified Flow to General Diffusion Models




---

## Algorithm 4 Rectified Diffusion $\epsilon$ -Prediction

---

**Input:** noise-data pair  $(\epsilon, \hat{\mathbf{x}}_0)$ ,  $\alpha_t, \sigma_t$

Sample  $\mathbf{x}_0$  from the data distribution

Sample time  $t$  from a predefined schedule or uniformly from  $[0, 1]$

Sample noise  $\epsilon$  from normal distribution

Compute  $\mathbf{x}_t$ :  $\mathbf{x}_t = \alpha_t \cdot \hat{\mathbf{x}}_0 + \sigma_t \cdot \epsilon$

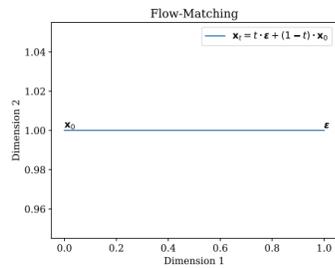
Predict noise  $\hat{\epsilon}$  using the model:  $\hat{\epsilon} = \text{Model}(\mathbf{x}_t, t)$

Compute loss:  $\mathcal{L} = \|\hat{\epsilon} - \epsilon\|_2^2$

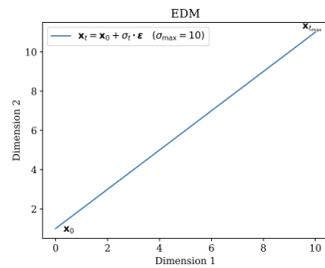
Backpropagate and update parameters

---

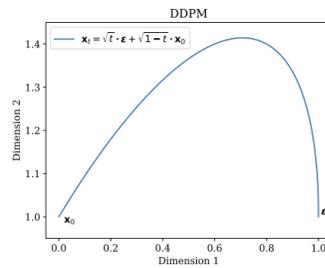
# Rectified Diffusion: the Essential Training Target Is First-Order Approximated ODE



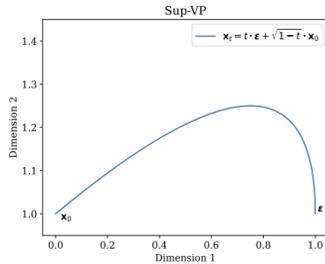
(a) Flow Matching



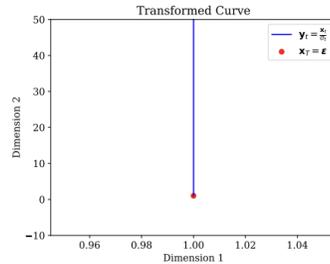
(b) EDM



(c) DDPM



(d) Sub-VP



(e) Transformed

Important points of first-order approximated ODE:

- It has the same form of predefined diffusion forms.

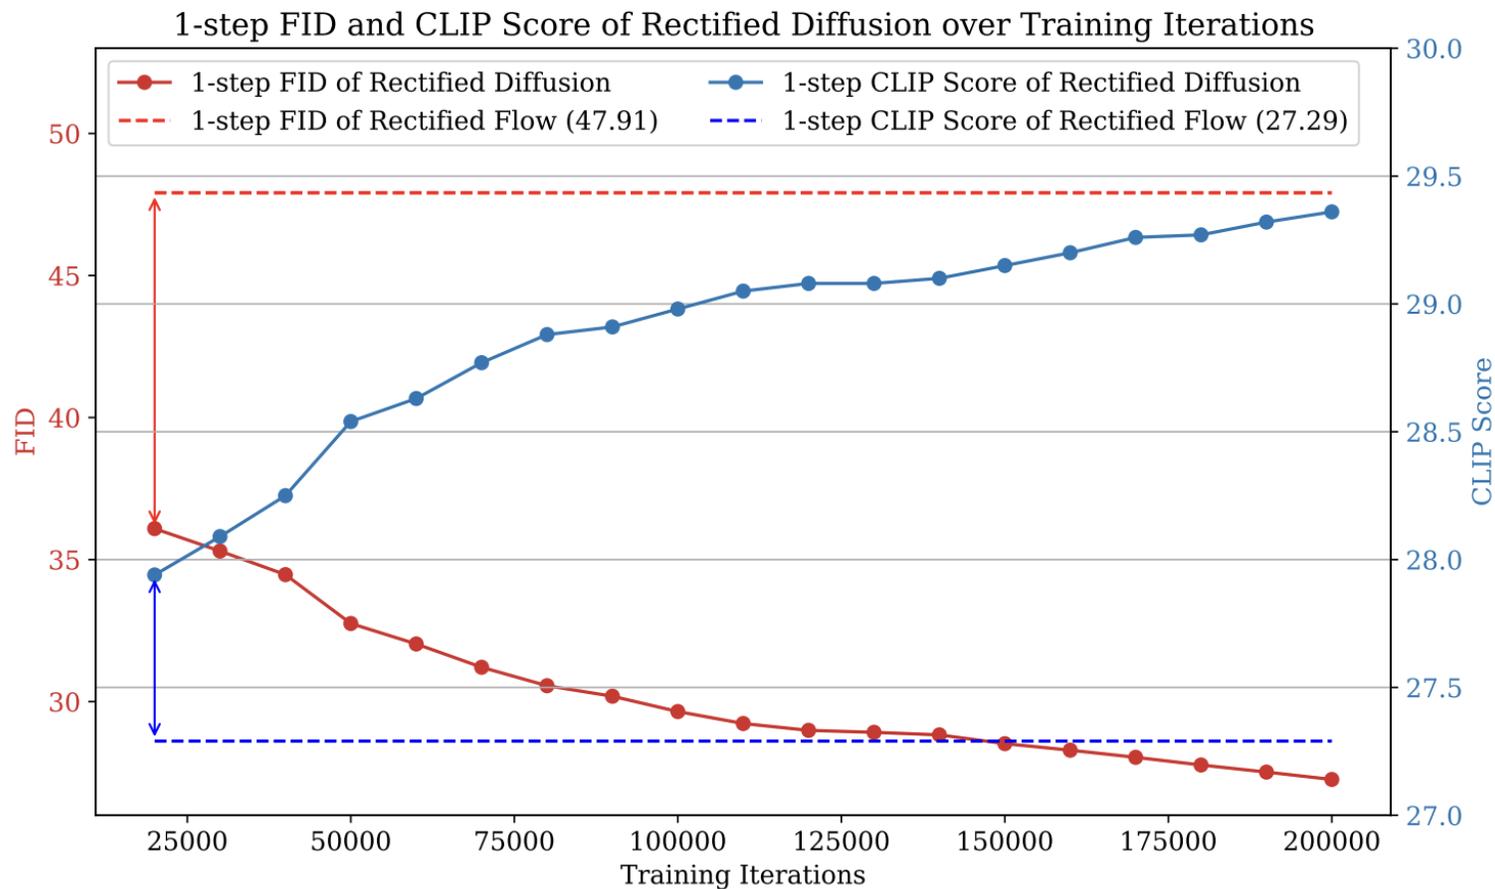
$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \epsilon$$

- It can be inherently curved.

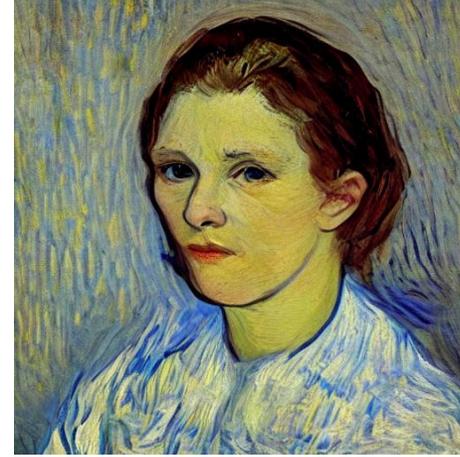
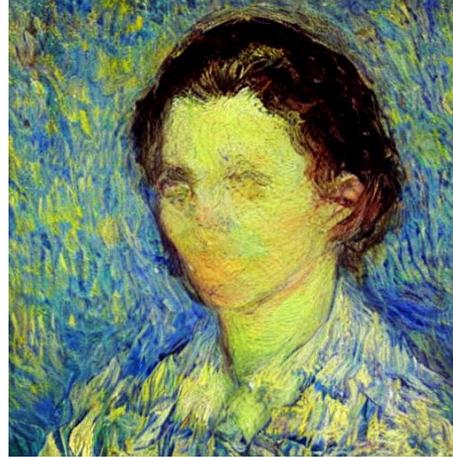
- It can be transformed into straight lines with timestep dependent scaling.

$$\mathbf{y}_t = \frac{\alpha_t}{\sigma_t} \mathbf{x}_0 + \epsilon$$

# Rectified Diffusion Vs Rectified Flow



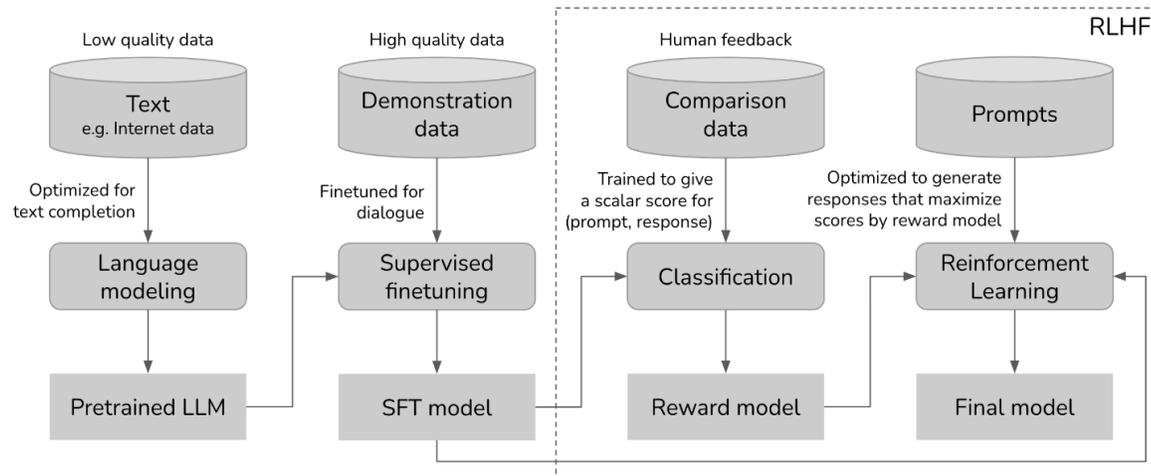
# Rectified Diffusion Vs Rectified Flow



Rectified-Flow

Rectified-Diffusion

# Human Preference Learning



Scale  
May '23

>1 trillion  
tokens

10K - 100K  
(prompt, response)

100K - 1M comparisons  
(prompt, winning\_response, losing\_response)

10K - 100K  
prompts

Examples  
Bolded: open  
sourced

GPT-x, Gopher, **Falcon**,  
LLaMa, **Pythia**, **Bloom**,  
**StableLM**

**Dolly-v2**, **Falcon-Instruct**

InstructGPT, ChatGPT,  
Claude, **StableVicuna**

Three ways for Preference Optimization:

- Differential Reward
- Reinforcement Learning
- Direct Preference Optimization

# Reinforcement Learning

The generation process of generative models can be seen as Markov decision process (MDP)

- Large language models.
  - Token-by-token prediction.
  - Each token sampling can be seen as an action following the implicitly defined policy.
  - All the generated tokens can be seen as state.
  - Reward Models: LLMs.
  
- Diffusion models.
  - Step-by-step prediction.
  - Each step can be seen as an action following the implicitly defined policy.
  - Last denoised results can be seen as state.
  - Reward models: VLMs or CLIP.

## Direct Preference Optimization

$$\max_{\pi} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi} [r(x, y)] - \beta \mathbb{D}_{\text{KL}}[\pi(y|x) || \pi_{\text{ref}}(y|x)]$$

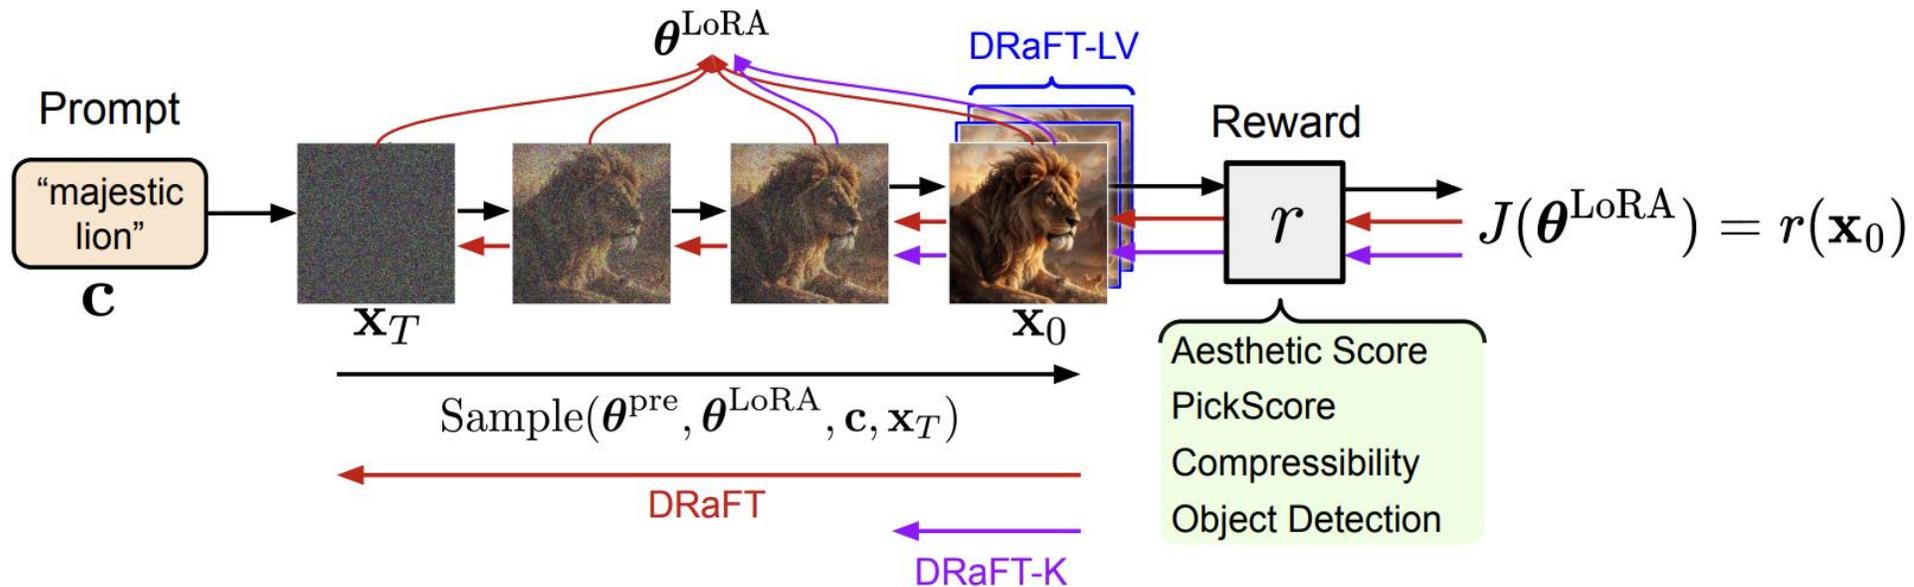
$$\pi(y|x) = \pi^*(y|x) = \frac{1}{Z(x)} \pi_{\text{ref}}(y|x) \exp\left(\frac{1}{\beta} r(x, y)\right)$$

$$r^*(x, y) = \beta \log \frac{\pi^*(y|x)}{\pi_{\text{ref}}(y|x)} + \beta \log Z(x)$$

## Direct Preference Optimization

$$\begin{aligned} p^*(y_1 \succ y_2 | x) &= \frac{\exp\left(\beta \log \frac{\pi^*(y_1|x)}{\pi_{\text{ref}}(y_1|x)} + \beta \log Z(x)\right)}{\exp\left(\beta \log \frac{\pi^*(y_1|x)}{\pi_{\text{ref}}(y_1|x)} + \beta \log Z(x)\right) + \exp\left(\beta \log \frac{\pi^*(y_2|x)}{\pi_{\text{ref}}(y_2|x)} + \beta \log Z(x)\right)} \\ &= \frac{1}{1 + \exp\left(\beta \log \frac{\pi^*(y_2|x)}{\pi_{\text{ref}}(y_2|x)} - \beta \log \frac{\pi^*(y_1|x)}{\pi_{\text{ref}}(y_1|x)}\right)} \\ &= \sigma\left(\beta \log \frac{\pi^*(y_1|x)}{\pi_{\text{ref}}(y_1|x)} - \beta \log \frac{\pi^*(y_2|x)}{\pi_{\text{ref}}(y_2|x)}\right). \end{aligned}$$

# Differential Reward



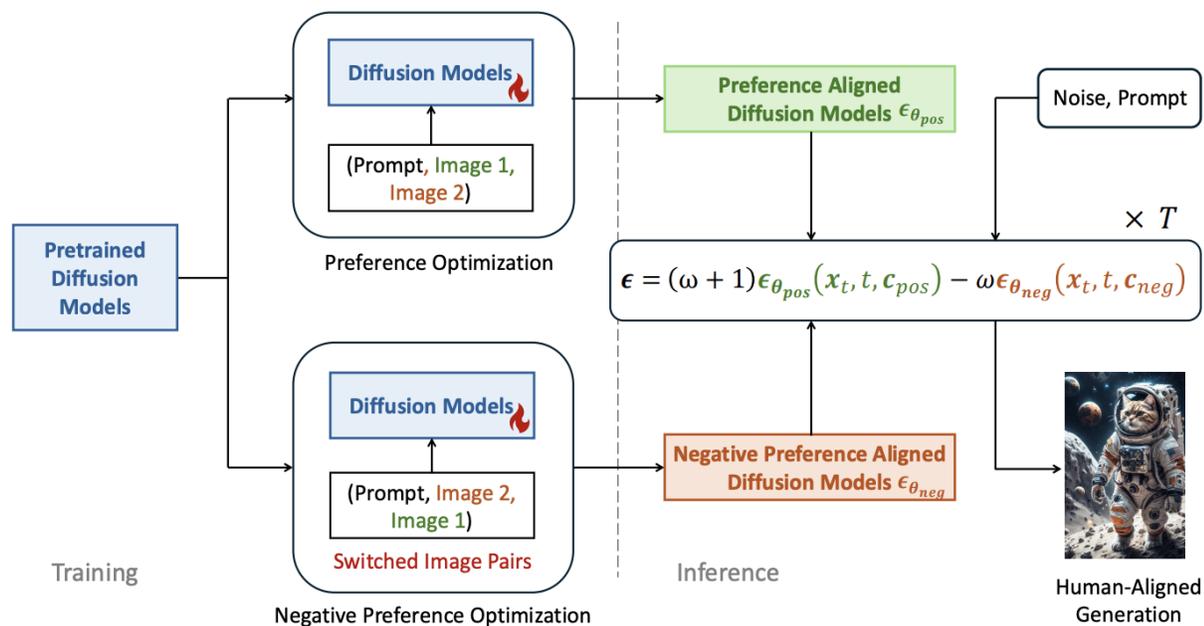
## Classifier-free guidance

$$\nabla_{\mathbf{x}_t} \log \mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{c}; t) + \nabla_{\mathbf{x}_t} \log \left[ \frac{\mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{c}; t)}{\mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{c}'; t)} \right]^{\omega}$$

$$\epsilon_{\theta}^{\omega} = (\omega + 1)\epsilon_{\theta}(\mathbf{x}_t, \mathbf{c}, t) - \omega\epsilon_{\theta}(\mathbf{x}_t, \mathbf{c}', t)$$

$$\epsilon_{\theta}^{\omega} = (\omega + 1)\epsilon_{\theta_{pos}}(\mathbf{x}_t, \mathbf{c}, t) - \omega\epsilon_{\theta_{neg}}(\mathbf{x}_t, \mathbf{c}', t)$$

# Diffusion Negative Preference Optimization



$$\nabla_{\mathbf{x}_t} \log \mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{c}; t) + \nabla_{\mathbf{x}_t} \log \left[ \frac{\mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{c}; t)}{\mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{c}'; t)} \right]^{\omega}$$

$$\epsilon_{\theta}^{\omega} = (\omega + 1)\epsilon_{\theta_{pos}}(\mathbf{x}_t, \mathbf{c}, t) - \omega\epsilon_{\theta_{neg}}(\mathbf{x}_t, \mathbf{c}', t)$$

To train Diffusion-NPO, we only need one line code.

# Diffusion Negative Preference Optimization

We only need one line code for Negative Preference Optimization

- Reinforcement learning or differential reward
  - Negating the output of reward model

$$R_{\text{NPO}}(\mathbf{x}, \mathbf{c}) = 1 - R(\mathbf{x}, \mathbf{c})$$

- Direct preference optimization
  - Switch the order of preference annotations

$$r = (\mathbf{x}_0, \mathbf{x}_1, \mathbf{c})$$

$$r_{\text{NPO}} = (\mathbf{x}_1, \mathbf{x}_0, \mathbf{c})$$

## Diffusion Negative Preference Optimization



w/o NPO

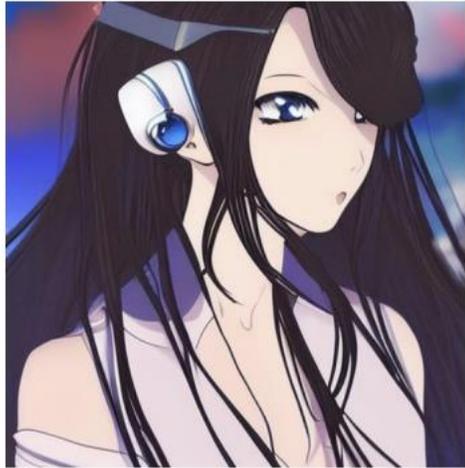
w/ NPO

w/o NPO

w/ NPO

# Diffusion Negative Preference Optimization

Prompt: "An anime woman"



w/o NPO



w/ NPO

Prompt: "Preteen girls with no underwear neither other clothes in a sofa with a childish faces ... (over 30 words)"



w/o NPO



w/ NPO

## References

- [1] Denoising Diffusion Probabilistic Models.
- [2] Denoising Diffusion Implicit Models.
- [3] Score-Based Generative Modeling through Stochastic Differential Equations.
- [4] Flow Matching for Generative Modeling.
- [5] Elucidating the Design Space of Diffusion-Based Generative Models.
- [6] DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps.
- [7] Discrete Flow Matching.
- [8] Consistency Models.
- [9] Consistency Models Made Easy.
- [10] Latent Consistency Models: Synthesizing High-Resolution Images with Few-step Inference.
- [11] Phased Consistency Model.
- [12] Multistep Consistency Models.
- [13] PeRFlow: Piecewise Rectified Flow as Universal Plug-and-Play Accelerator.
- [14] Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow.
- [15] InstaFlow: One Step is Enough for High-Quality Diffusion-Based Text-to-Image Generation.
- [16] StyleGAN-T: Unlocking the Power of GANs for Fast Large-Scale Text-to-Image Synthesis.
- [17] Stable Consistency Tuning: Understanding and Improving Consistency Models.
- [18] SwiftBrush : One-Step Text-to-Image Diffusion Model with Variational Score Distillation.
- [19] One-step Diffusion with Distribution Matching Distillation.
- [20] Progressive Distillation for Fast Sampling of Diffusion Models

## Our Works

[1] Stable Consistency Tuning: Understanding and Improving Consistency Models.

[2] Rectified Diffusion: Straightness Is Not Your Need in Rectified Flow.

[3] Phased Consistency Model.

[4] AnimateLCM: Computation-Efficient Personalized Style Video Generation without Personalized Video Data.

[5] Diffusion-NPO: Negative Preference Optimization for Better Preference Aligned Generation of Diffusion Models

# Thank you!

Fu-Yun Wang

[fywang@link.cuhk.edu.hk](mailto:fywang@link.cuhk.edu.hk)