Score Distillation via Reparametrized DDIM

Artem Lukoianov

Haitz S de Ocáriz Borde

Kristjan Greenewald

Vitor Guizilini

Timur Bagautdinov

Vincent Sitzmann

Justin Solomon

3D generation with the quality of 2D diffusion?

2D Diffusion (DDIM)

😀

Image Diffusion generates crisp, high-quality images

3D generation with the quality of 2D diffusion?

2D Diffusion (DDIM)

😀

Image Diffusion generates crisp, high-quality images

Score Distillation (SDS)

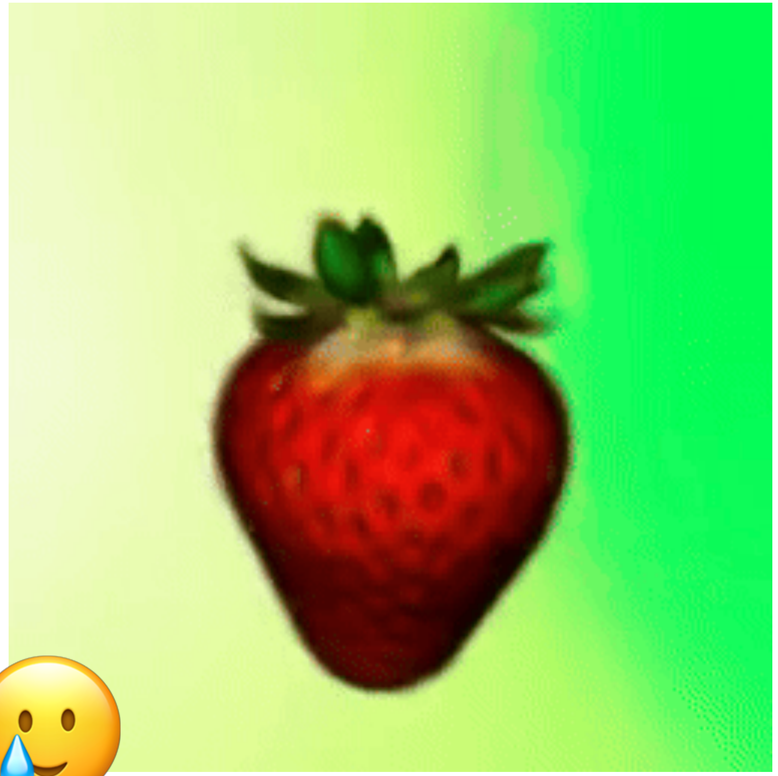

🥲

Score Distillation Sampling uses Image Diffusion, **but the results are blurry**

3D generation with the quality of 2D diffusion?

2D Diffusion (DDIM)

Image Diffusion generates crisp, high-quality images

Score Distillation (SDS)

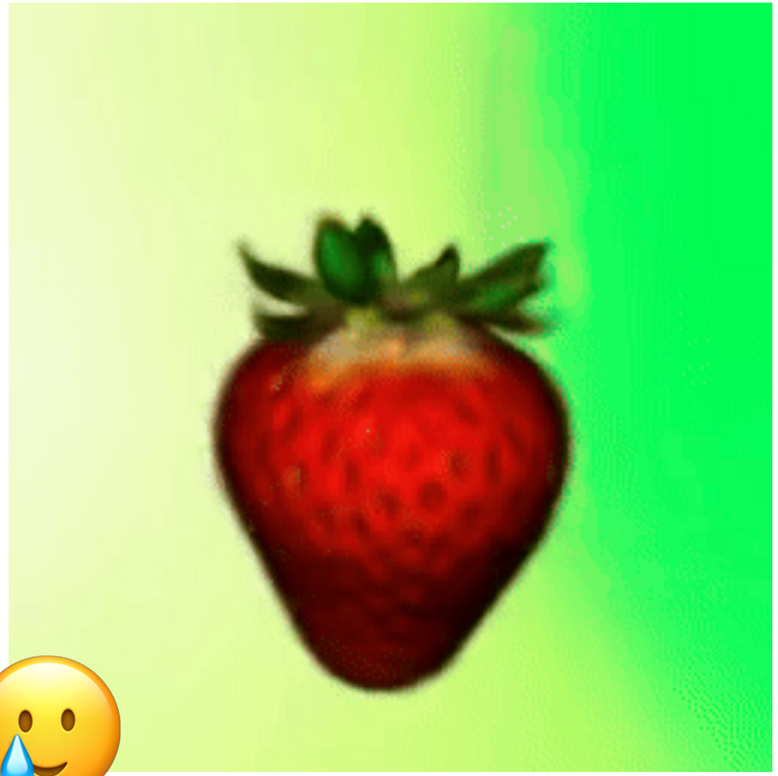

Score Distillation Sampling uses Image Diffusion, **but the results are blurry**

Ours (SDI)

We provide a theoretical analysis and **suggest a fix**

Score Distillation uses off-the-shelf 2D diffusion to generate 3D objects

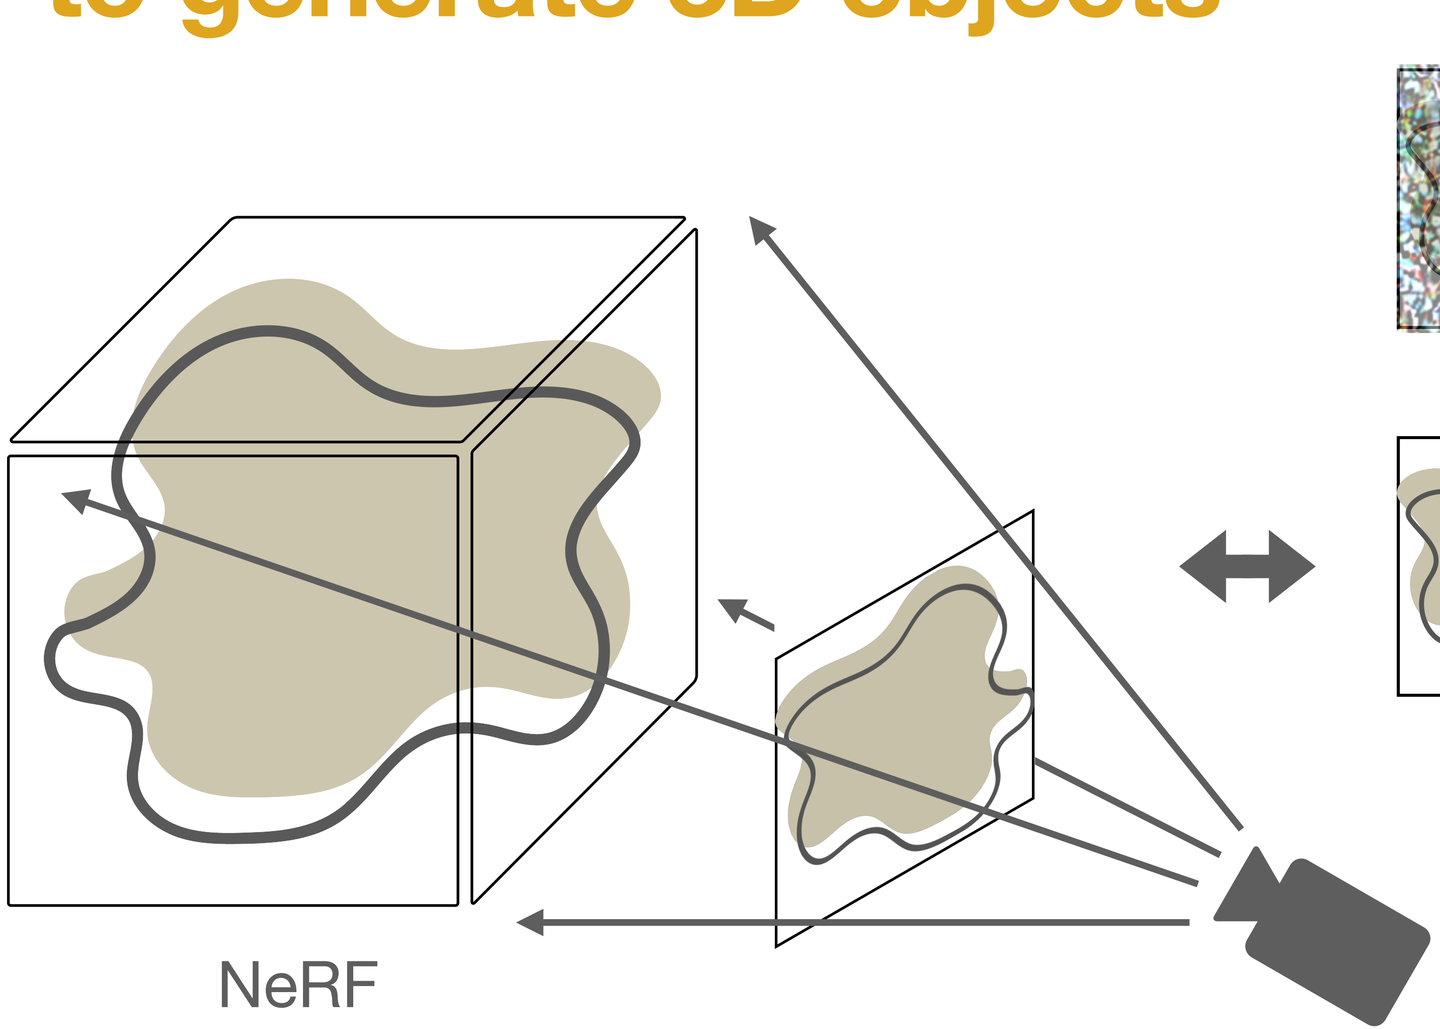

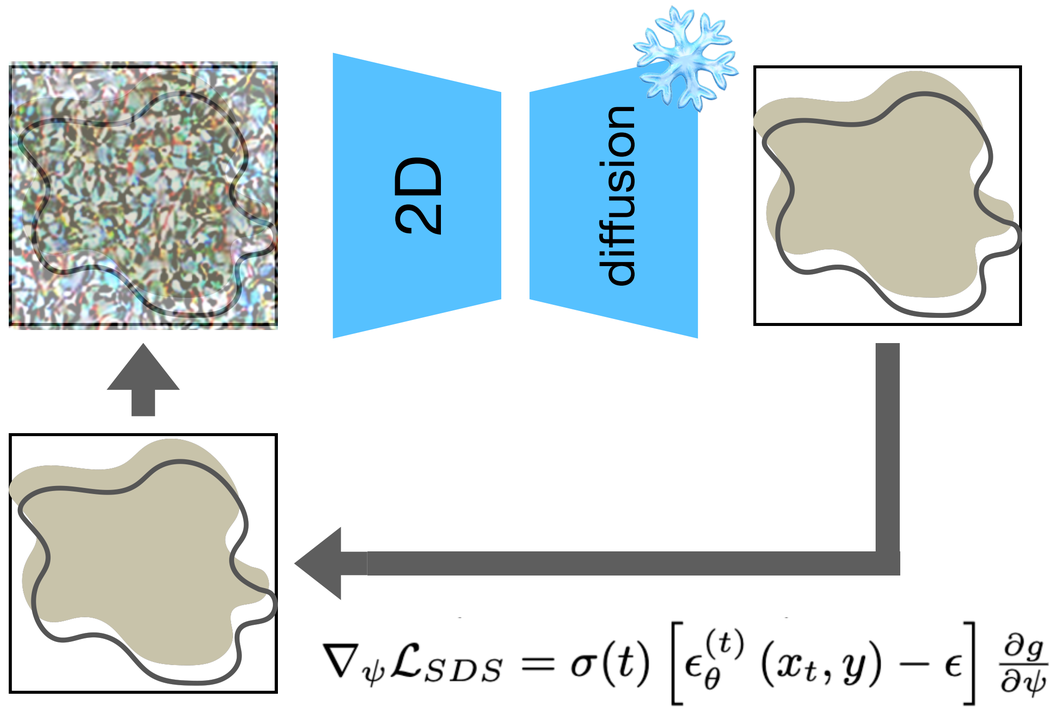

$$\nabla_\psi \mathcal{L}_{SDS} = \sigma(t)\left[\epsilon_\theta^{(t)}(x_t, y) - \epsilon\right]\frac{\partial g}{\partial \psi}$$

We show that SDS is a dual process of DDIM
The two trajectories match when using correct noise

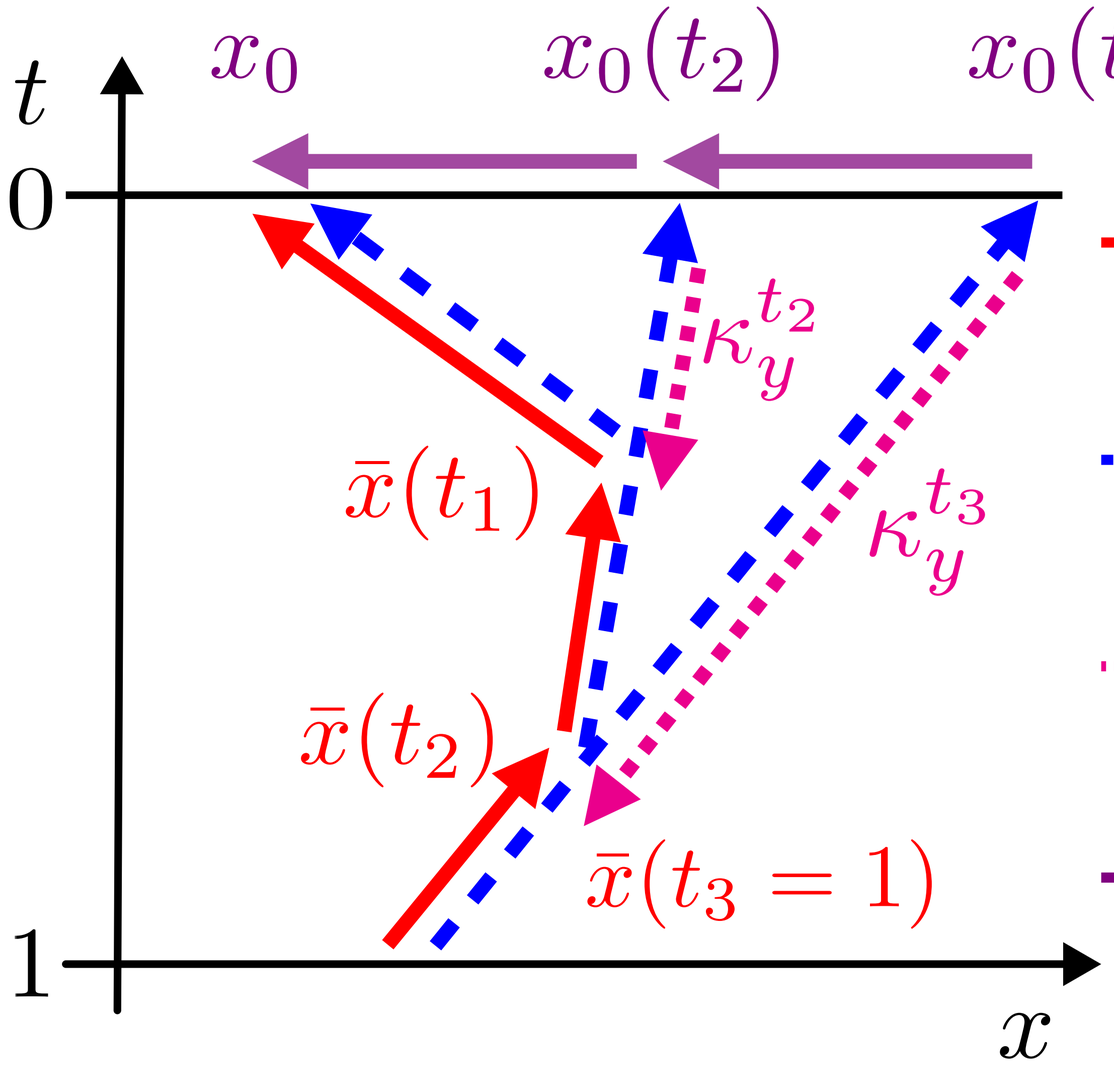

DDIM

Single-step denoising

Noising

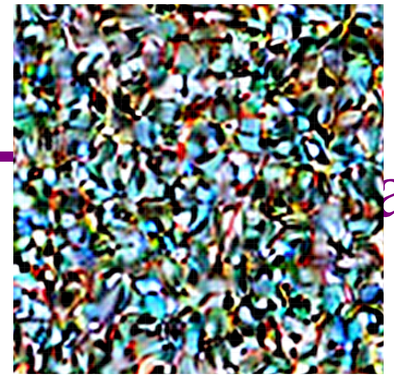

in ace

Reparametrizing DDIM

Formal derivation

$$x_0(t-\tau) = x_0(t) - \sigma(t-\tau)\Big[\underbrace{\epsilon_\theta^{t-\tau}\Big(\overbrace{\sqrt{\alpha(t-\tau)}x_0(t) + \sqrt{1-\alpha(t-\tau)}\kappa_y^t\big(x_0(t)\big)}^{x_0 \text{ noised with } \kappa_y^t \text{ to time } t-\tau}, y\Big)}_{\text{predicted noise}} - \underbrace{\kappa_y^t\big(x_0(t)\big)}_{\text{noise sample } \kappa_y^t}\Big].$$

$$\kappa_y^t = \epsilon_\theta^t\big(\sqrt{\alpha(t)}x_0(t) + \sqrt{1-\alpha(t)}\kappa_y^t, y\big)$$

$$\kappa_y^t \sim \mathcal{N}(0, I)$$

SDS as reparametrized DDIM

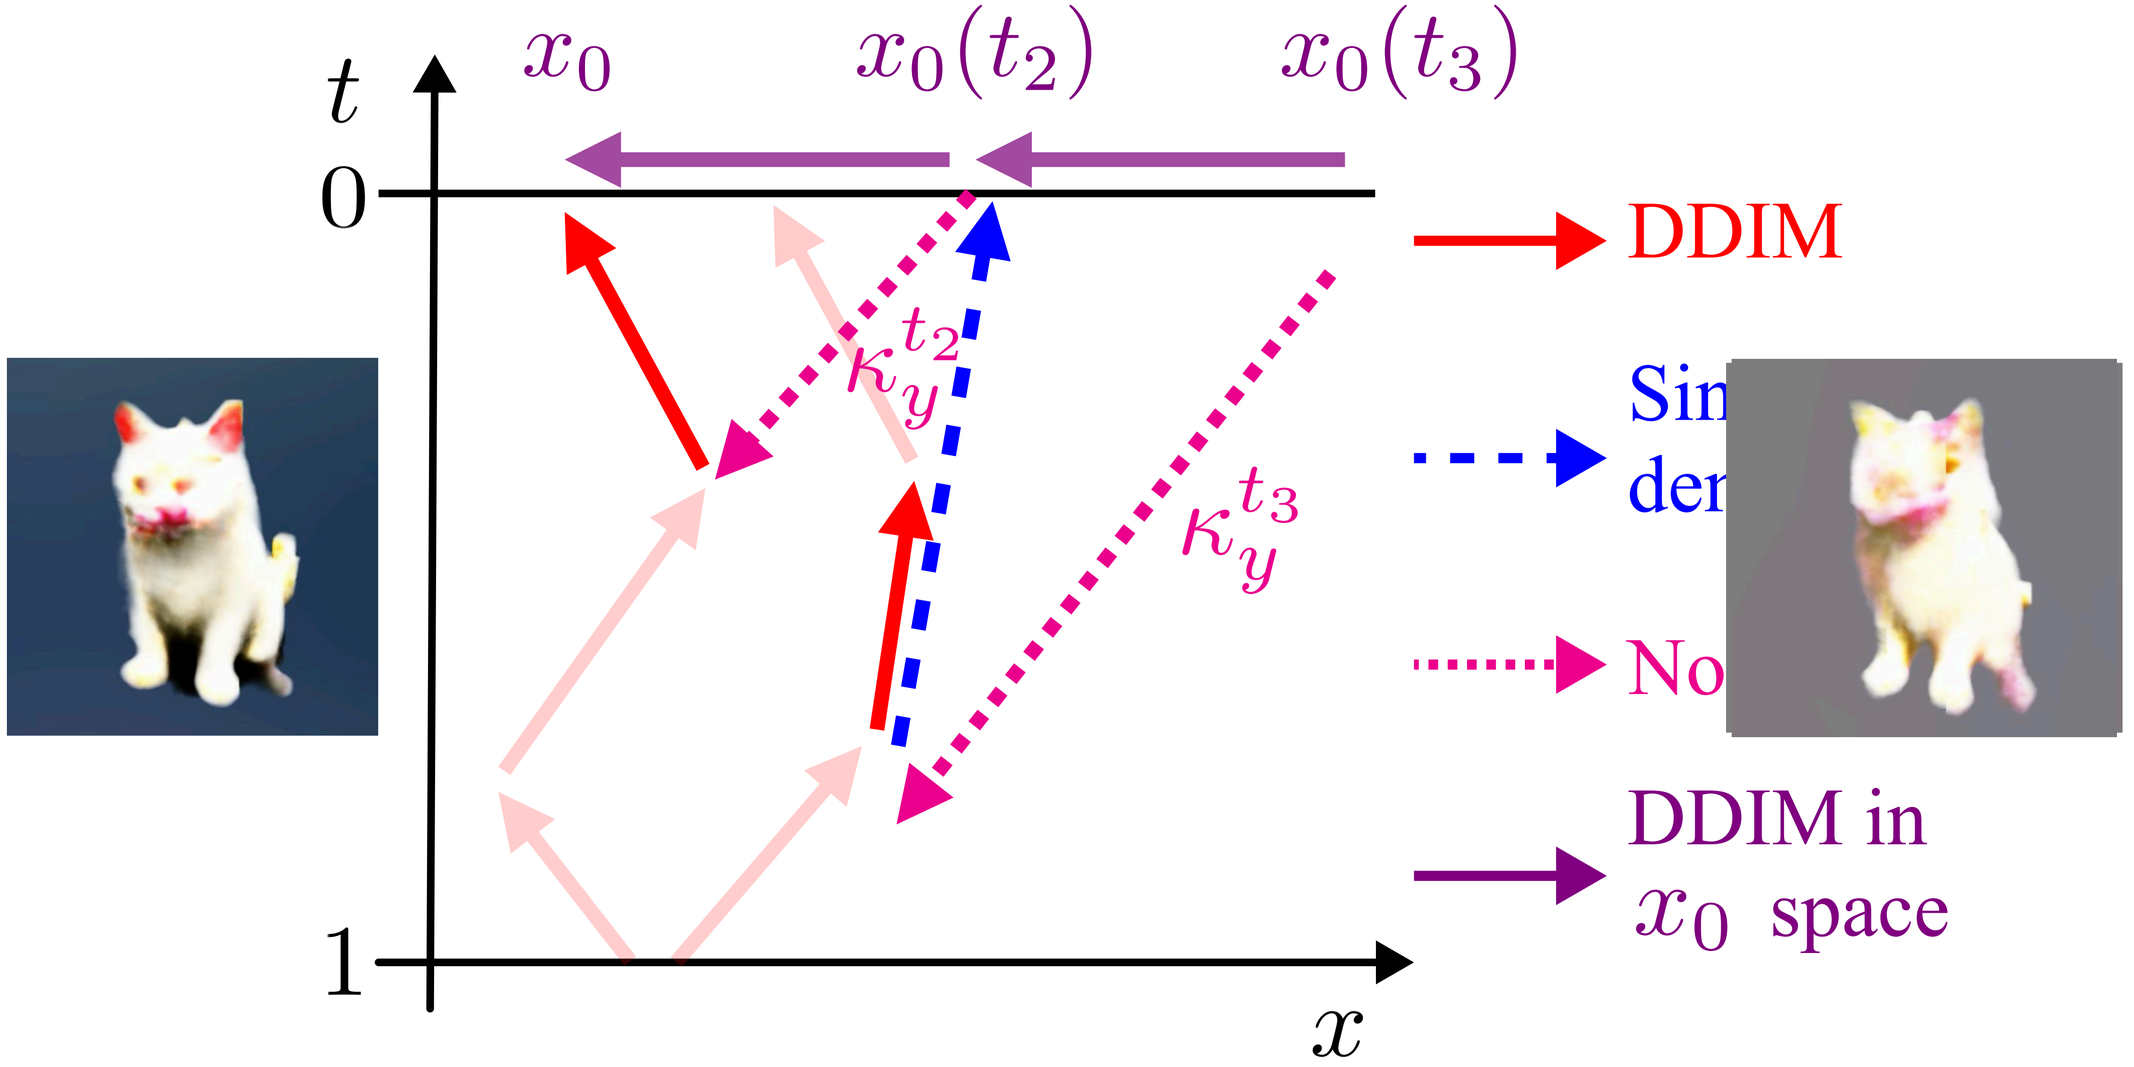

Noise is the problem!

$$\kappa_y^t = \epsilon_\theta^t\left(\sqrt{\alpha(t)}x_0(t) + \sqrt{1-\alpha(t)}\kappa_y^t, y\right)$$

Random, resampled

Random, fixed

SGD Optim.

DDIM Inversion

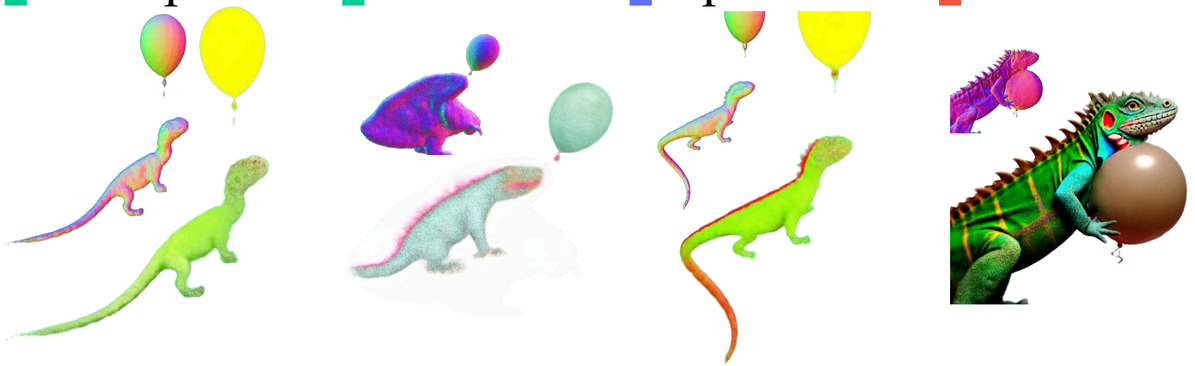

Prompt – “an iguana holding a balloon”

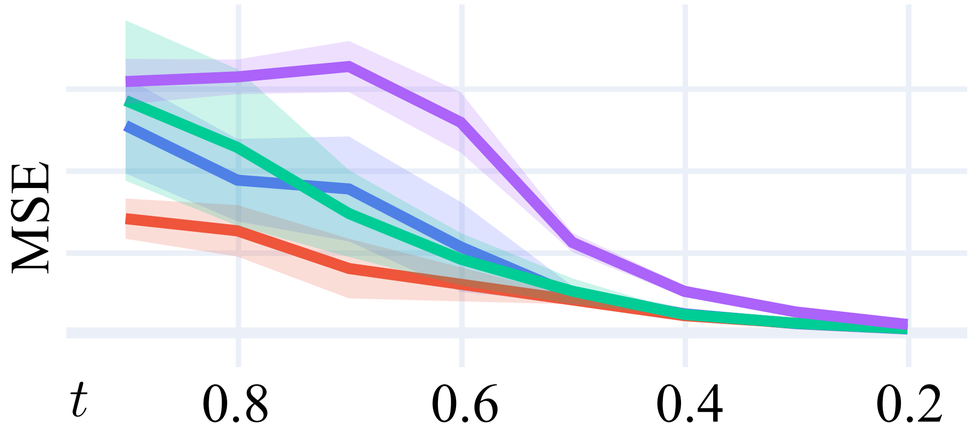

Fixed point iterations diverged

DDIM inversion

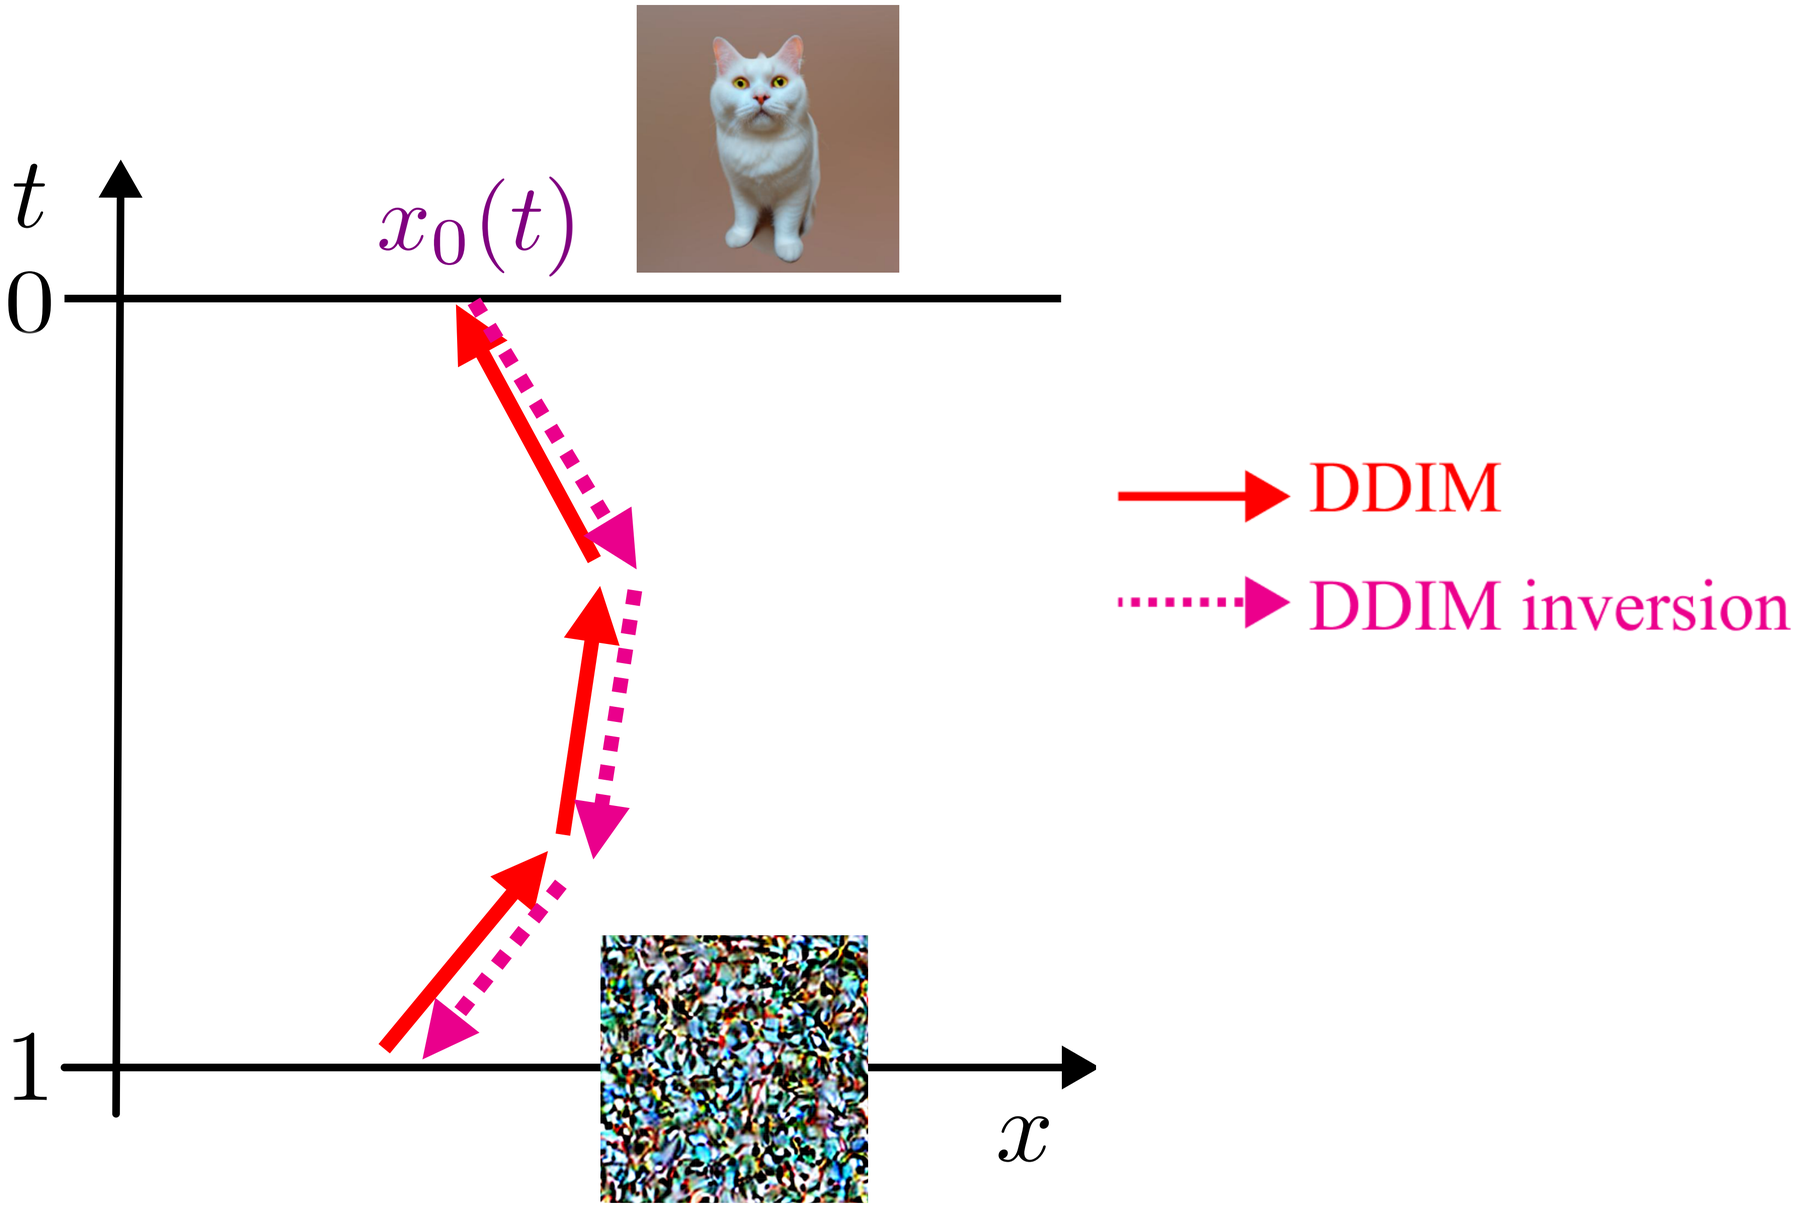

SDI as reparametrized DDIM

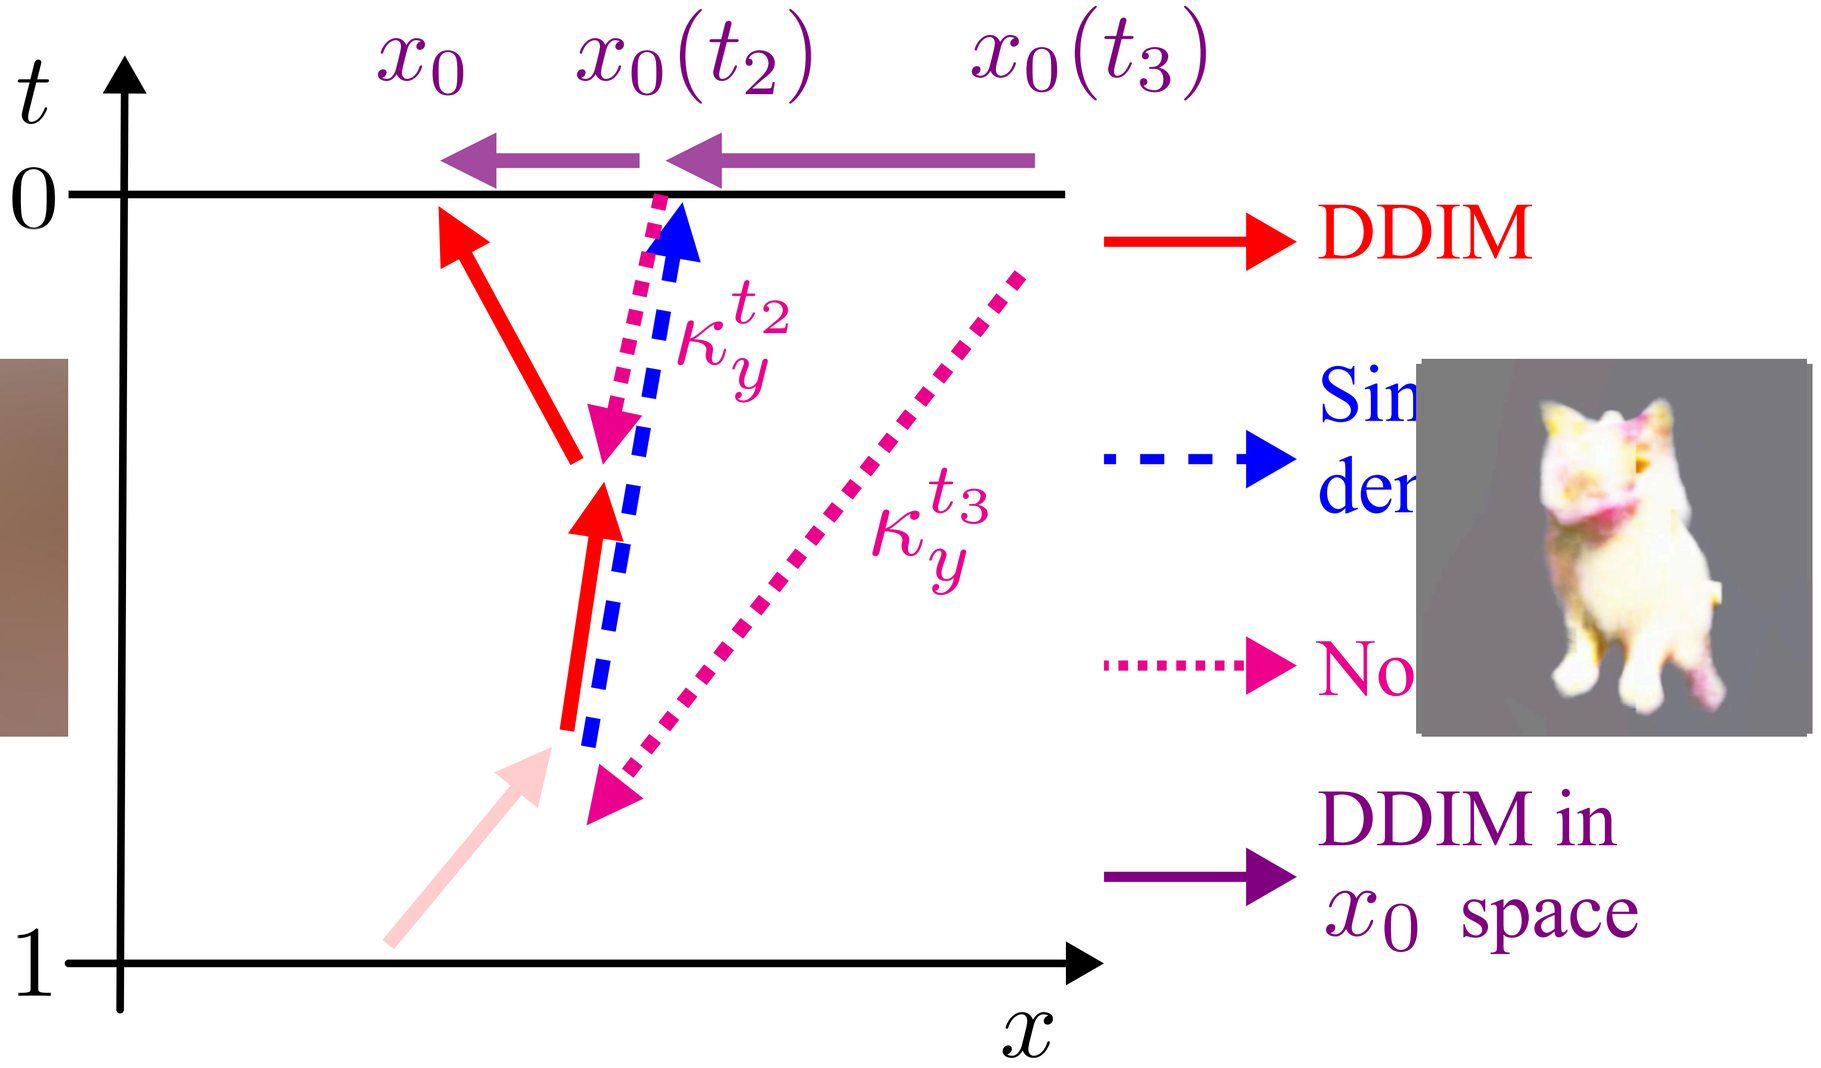

By fixing the noise term we are bringing the generation quality in 3D much closer to the 2D models

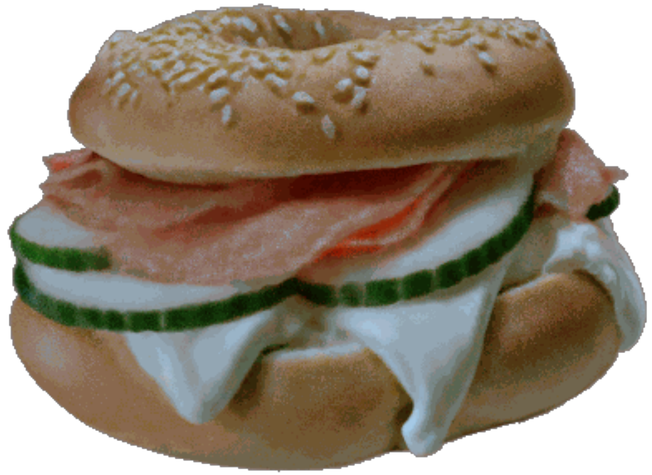

MORE ABOUT THE WORK

CODE AND THE FULL PAPER ARE AVAILABLE

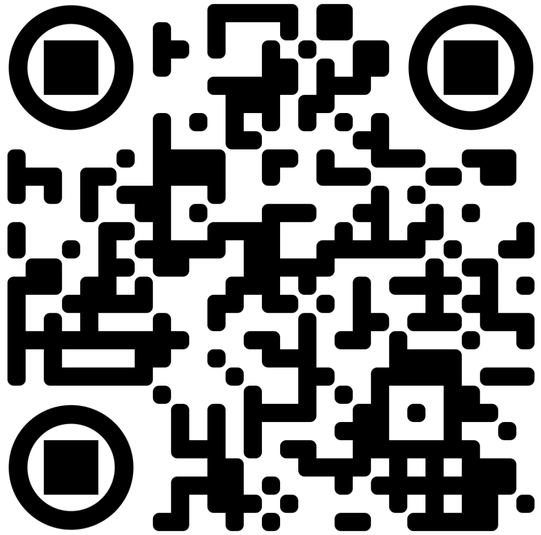

lukoianov.com/sdi