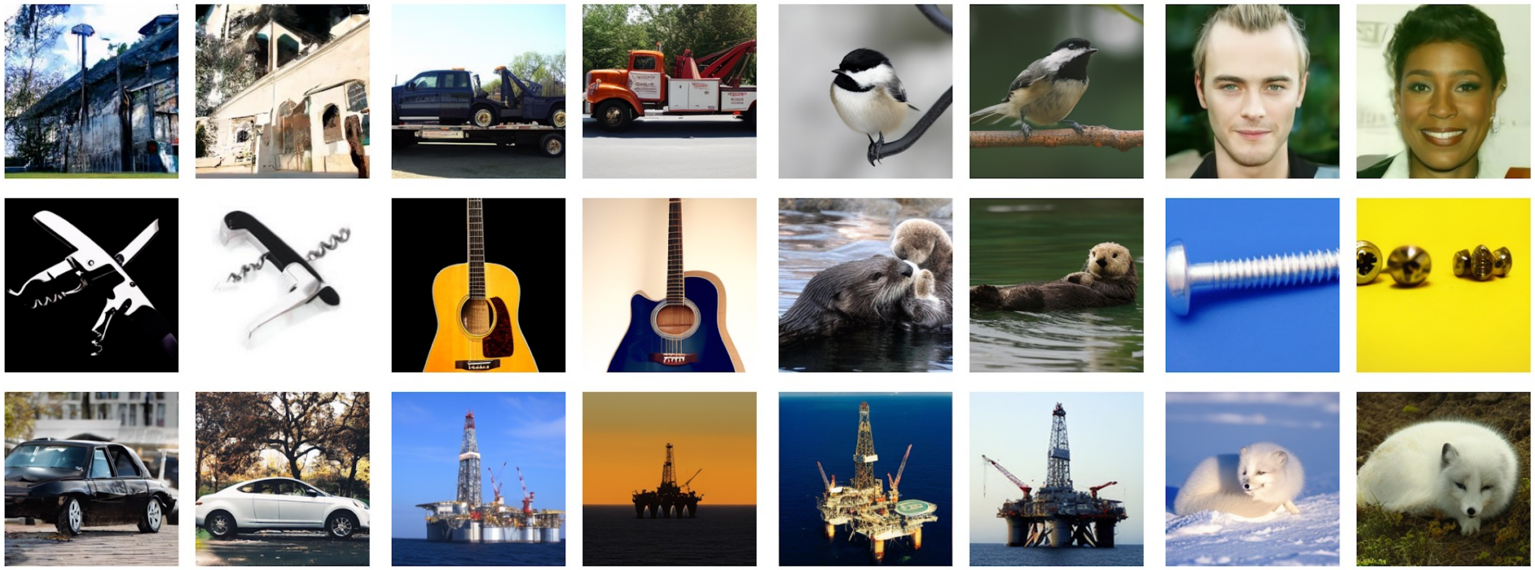

Antithetic Noise in Diffusion Models

Jing Jia[1], Sifan Liu[2], Bowen Song[3], Wei Yuan[4], Liyue Shen[3], Guanyang Wang[4]

1. Department of Computer Science, Rutgers University

2. Duke University

3. Department of EECS, University of Michigan

4. Department of Statistics, Rutgers University

Scan for project

RUTGERS UNIVERSITY

MICHIGAN ENGINEERING
UNIVERSITY OF MICHIGAN

Duke
UNIVERSITY

Agenda

- Overview
- Antithetic Noise
- Application

Diffusion Model

- A class of generative models (like GANs, VAEs, autoregressive models)
- Achieve **state-of-the-art** quality in image generation
- **Real world product:**
  - Stable Diffusion
  - MidJourney
  - DALL·E 3 (OpenAI)
  - Imagen (Google DeepMind)

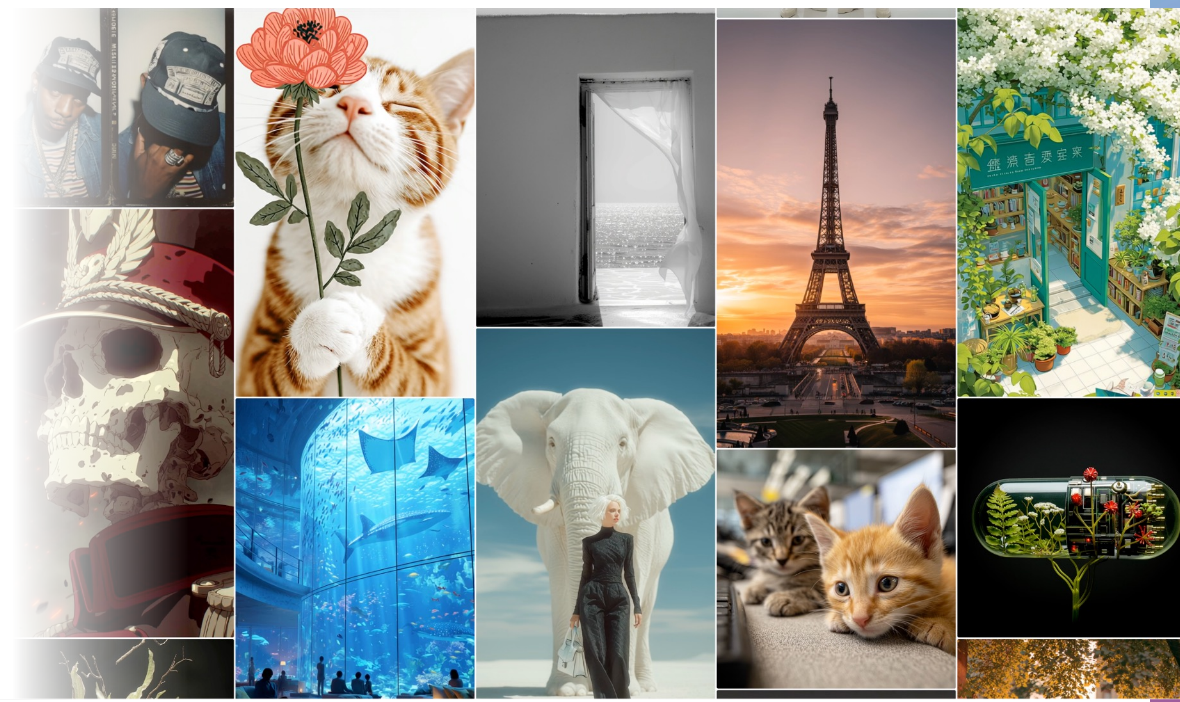

Image from MidJourney gallery

RUTGERS

Diffusion Model

- **Training: image → noise**
  - Corrupted with **additive Gaussian noise**
  - Asymptotically transformed into pure Gaussian noise
- **Goal**: learn the reverse process
- **Generation: noise → image**
  - Model reverses the process with **predicted noise**
  - **Denoises** from pure Gaussian noise to image samples

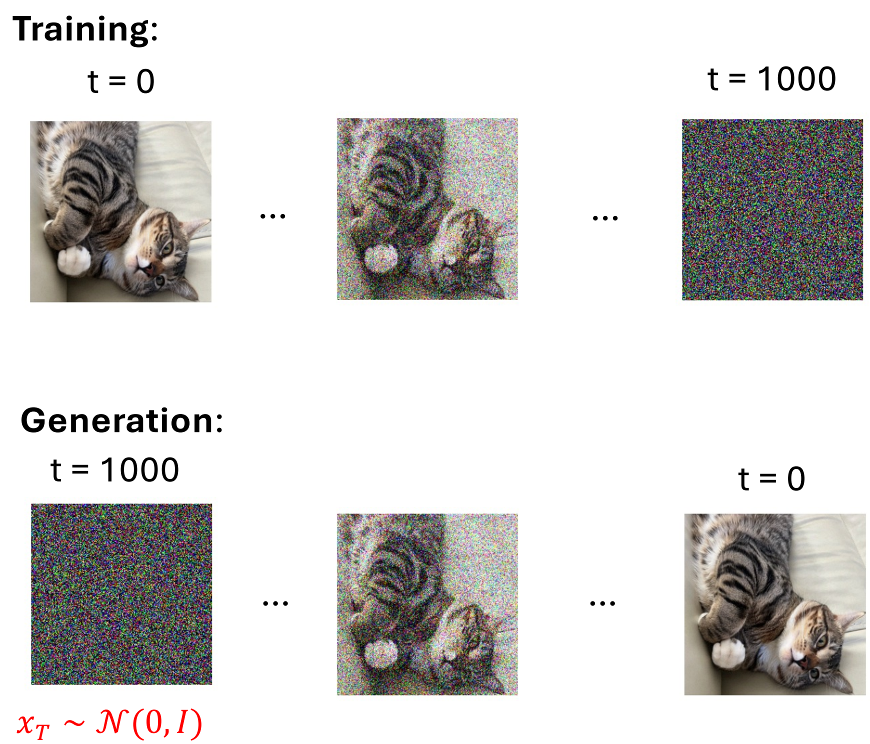

Diffusion Model

- **Network weights**
  - Architectures: U-Net, Transformers, etc.
- **Denoising schedule**
  - Defines how noise is removed step by step
  - DDPM – **stochastic** reverse process
  - DDIM – **deterministic, fast** reverse process
- **Initial Gaussian noise**
  - $x_T \sim \mathcal{N}(0, I)$

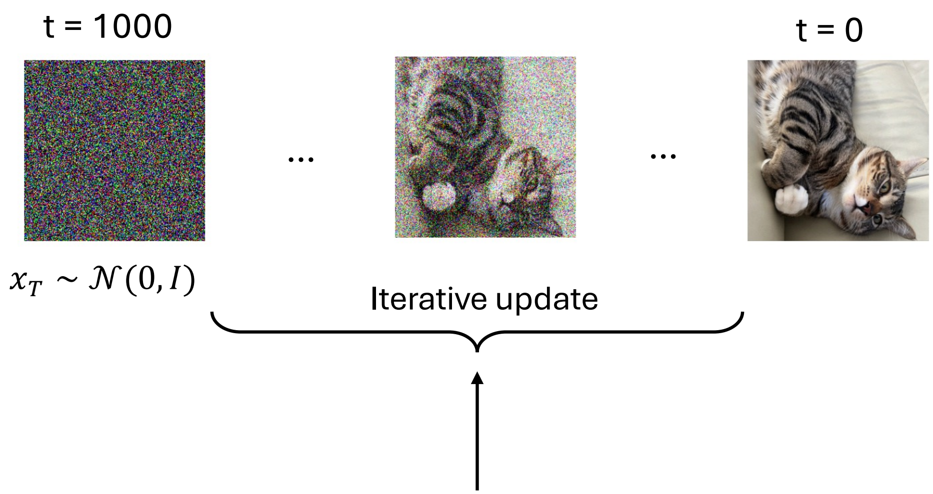

One step DDIM update formula

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1-\alpha_t}\, \epsilon_\theta^{(t)}(\mathbf{x}_t)}{\sqrt{\alpha_t}} \right) + \sqrt{1-\alpha_{t-1}}\, \epsilon_\theta^{(t)}(\mathbf{x}_t)$$

α_t coefficients determined by the noise schedule

$\epsilon_\theta^{(t)}(\mathbf{x})$ noise predicted by the denoising network with weight θ

Diffusion Model

- **Network weights**
  - Architectures: U-Net, Transformers, etc.
- **Denoising schedule**
  - Defines how noise is removed step by step
  - DDPM – **stochastic** reverse process
  - DDIM – **deterministic, fast** reverse process
- **Initial Gaussian noise**
  - $x_T \sim \mathcal{N}(0, I)$

Question:

How would initial noise change outcome?

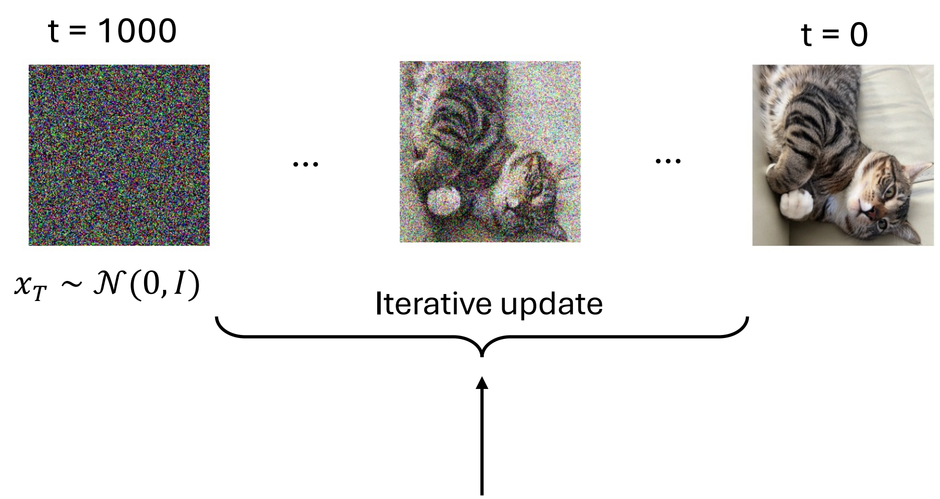

One step DDIM update formula

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1-\alpha_t}\, \epsilon_\theta^{(t)}(\mathbf{x}_t)}{\sqrt{\alpha_t}} \right) + \sqrt{1-\alpha_{t-1}}\, \epsilon_\theta^{(t)}(\mathbf{x}_t)$$

α_t coefficients determined by the noise schedule

$\epsilon_\theta^{(t)}(\mathbf{x})$ noise predicted by the denoising network with weight θ

Agenda

- Overview
- Antithetic Noise
- Application

Antithetic Initial Noise

- **Random/Random (RR):** Two independent Gaussian noise vectors z_1 and z_2 are sampled.
- **Positive/Negative (PN):** A single Gaussian noise vector z is paired with its negation $-z$.

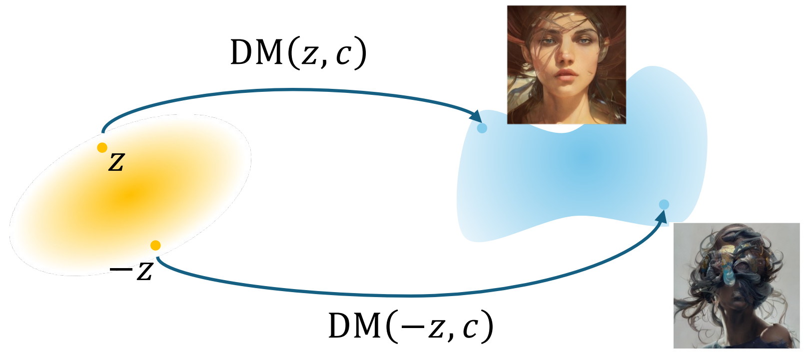

$\boldsymbol{z} \sim \mathcal{N}(0, I)$

c, condition if available (text prompt, class label)

RUTGERS

Antithetic Initial Noise

- **Well-studied technique**
  - Classic variance reduction method in Monte Carlo simulation
  - Natural to adapt
- **Preserve quality**
  - Since $z \sim \mathcal{N}(0,1) \Rightarrow -z \sim \mathcal{N}(0,1)$
  - Generated samples $\mathrm{DM}(z, c)$ and $\mathrm{DM}(-z, c)$ share the same distribution → no loss in quality
- **Maximal separation in noise space**
  - z and $-z$ lie on opposite poles → largest possible perturbation

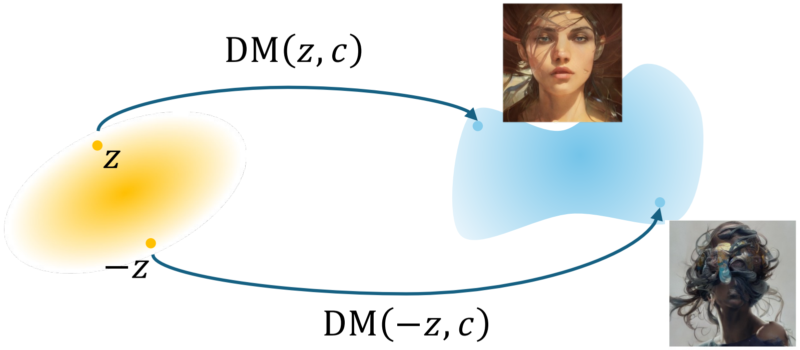

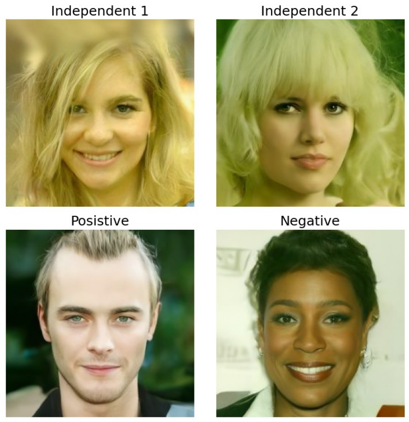

Discovery

- **Apply Antithetic Noise:**
  - Datasets: CIFAR10, CelebA, Church
  - Architecture : U-Net (SD1.5) and transformer (DiT)

- **Pixel-wise Correlation**
  - Standard Pearson correlation
  - Centralized Pearson correlation
    - Correlation with mean subtracted
    - To correct image structure-level bias

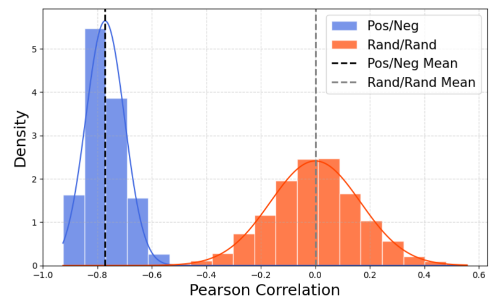

LSUN-Church

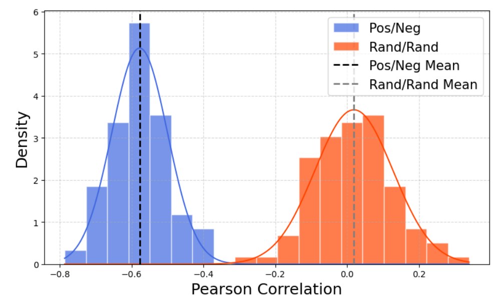

Stable Diffusion

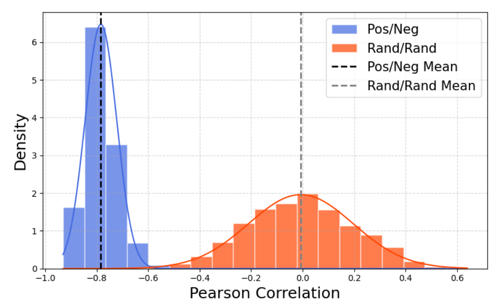

CelebA-HQ

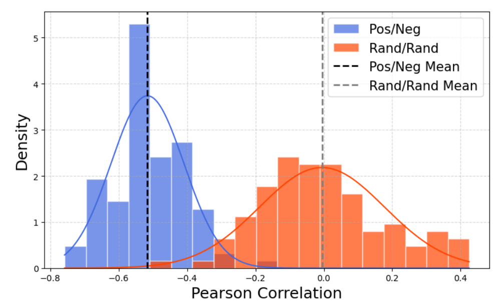

DiT

Discovery

- **Main Finding:**
  - Pairing noise z with its negation −z produces strongly negatively correlated samples
- **Universal Phenomenon:**
  - Works across all unconditional, conditioned latent-diffusion, and diffusion-posterior samplers
  - Works on generative models beyond diffusion model: VAE, normalizing flow

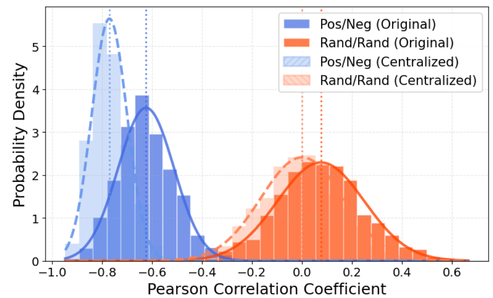

LSUN-Church

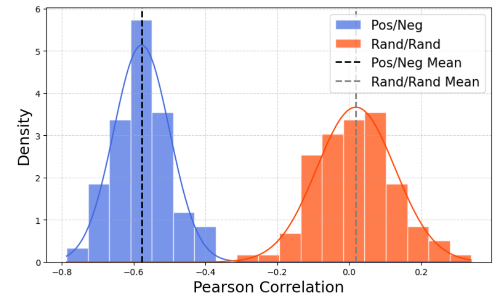

Stable Diffusion

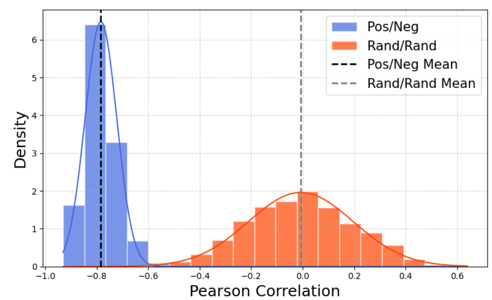

CelebA-HQ

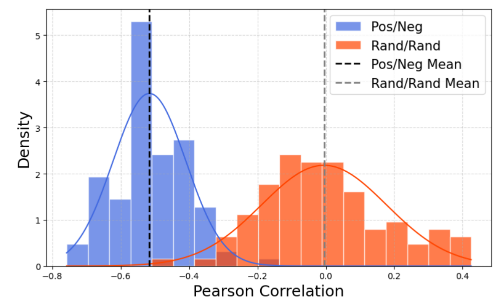

DiT

RUTGERS

Agenda

- Overview
- Antithetic Noise
- Application

Diversity Enhancement

A young woman standing on a vintage train platform with a suitcase, golden hour light casting long shadows, wind blowing her scarf. Old-school train in the distance, surrounded by mountains and fall trees - nostalgic

Art illustration of a moonlit jungle clearing where luminous giant mushrooms tower over a lone wanderer in a crimson cloak, fireflies tracing golden spirals through misty air; illustrated as a rich digital painting with painterly brush-work soft rim-lighting, and deep teal-violet shadows for a dreamy, storybook fantasy mood

RUTGERS

Diversity Enhancement

- **Problem:** Standard sampling produces similar images
- **Solution:** Antithetic pairs explore opposite regions of data manifold
- **Benefits:**
  - No quality degradation
  - Zero computational overhead
- **Metric:** higher diversity
  - SSIM (lower better)
  - LPIPS (higher better)

Metric	Unconditional Diffusion			Conditional Diffusion	
	CIFAR-10	Church	CelebA-HQ	SD1.5	DiT
SSIM (%)	88.78	45.69	36.78	28.32	23.99
LPIPS (%)	6.69	3.54	15.14	5.78	10.62

Table 1: Average percentage improvement of PN over RR pairs.

RUTGERS

Diversity Visualization

PN Pairs

Pos | Neg | Pos | Neg

RR Pairs

Rand1 | Rand2 | Rand1 | Rand2

✅ **More diverse !**

RUTGERS

Diversity Visualization

PN Pairs

Pos Neg Pos Neg

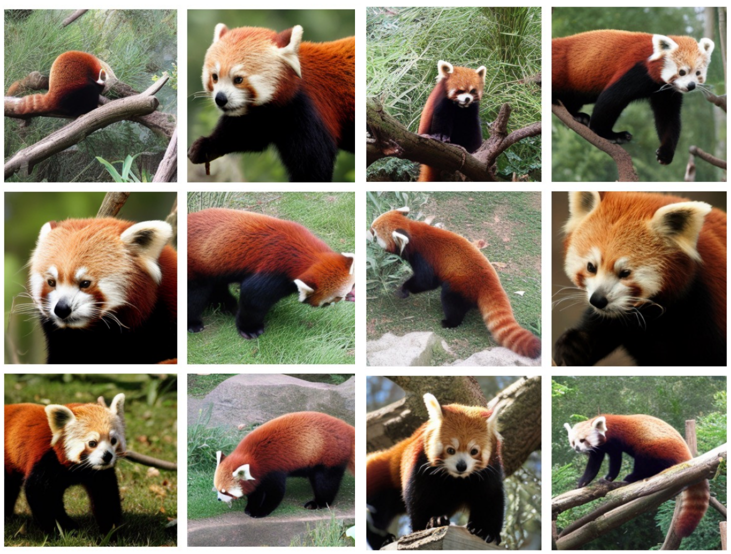

RR Pairs

Rand1 Rand2 Rand1 Rand2

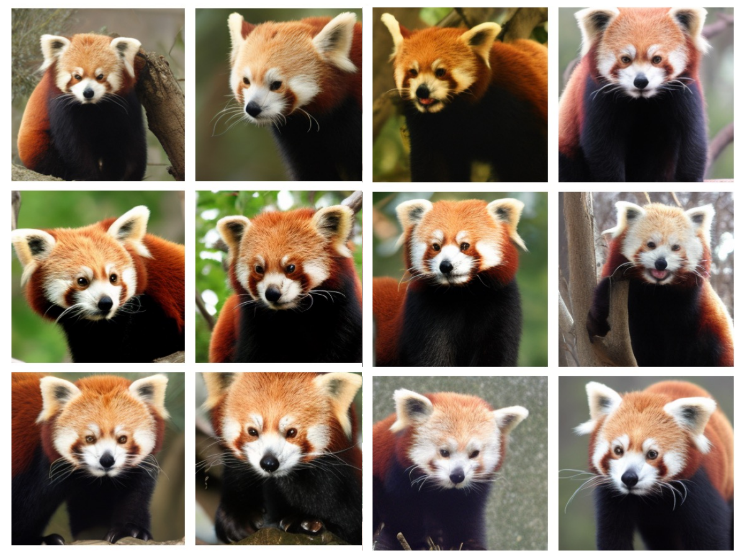

✅ **More diverse !**

Diversity Visualization

PN Pairs

Pos Neg Pos Neg

RR Pairs

Rand1 Rand2 Rand1 Rand2

✅ **More diverse !**

RUTGERS

Diversity Visualization

PN Pairs

Pos Neg Pos Neg

RR Pairs

Rand1 Rand2 Rand1 Rand2

✅ **More diverse !**

RUTGERS

Pixel Statistics Uncertainty

- **Idea:** Use negative correlated samples to reduce variance.
- **Metric**:
  - Confidence interval
  - Relative efficiency w.r.t. MC $\left(\frac{CI_{MC}}{CI}\right)^2$
- **Estimators:**
  - **Standard Monte Carlo:**
    - draw N independ $z_1, \ldots, z_N \sim \mathcal{N}(0, I)$
  - **Antithetic Monte Carlo (AMC):**
    - Generate $2K$ pairs negatively correlated sample

Dataset		Brightness	Pixel Mean	Contrast
Stable Diff.	MC	5.45	5.59	2.18
	AMC	**1.26 (18.83)**	**1.41 (15.84)**	**0.89 (6.25)**
CIFAR-10	MC	2.00	2.04	1.08
	AMC	**0.35 (32.66)**	**0.39 (27.12)**	0.23 (22.05)
Church	MC	1.67	1.67	1.02
	AMC	**0.14 (136.37)**	**0.16 (106.79)**	**0.20 (27.16)**
DiT	MC	7.52	7.74	3.32
	AMC	**3.13 (5.79)**	**3.12 (6.16)**	**1.72 (3.74)**

Table 2: CI lengths and efficiency $\left(\frac{CI_{MC}}{CI}\right)^2$

RUTGERS

Summary

- ✓ **Discovery:** Universal negative correlation from antithetic initial noise
- ✓ **Theory:** Symmetry conjecture explaining the phenomenon
- ✓ **Applications:** Enhanced diversity + uncertainty quantification, and more…
- ✓ **Extensions:** Quasi-Monte Carlo integration
- ✓ **Impact:** Training-free, model-agnostic, zero overhead

Thank you!