

# Diffusion art or Digital Forgery?


Investigating data replication in diffusion models



Gowthami Somepalli



Vasu Singla



Jonas Geiping



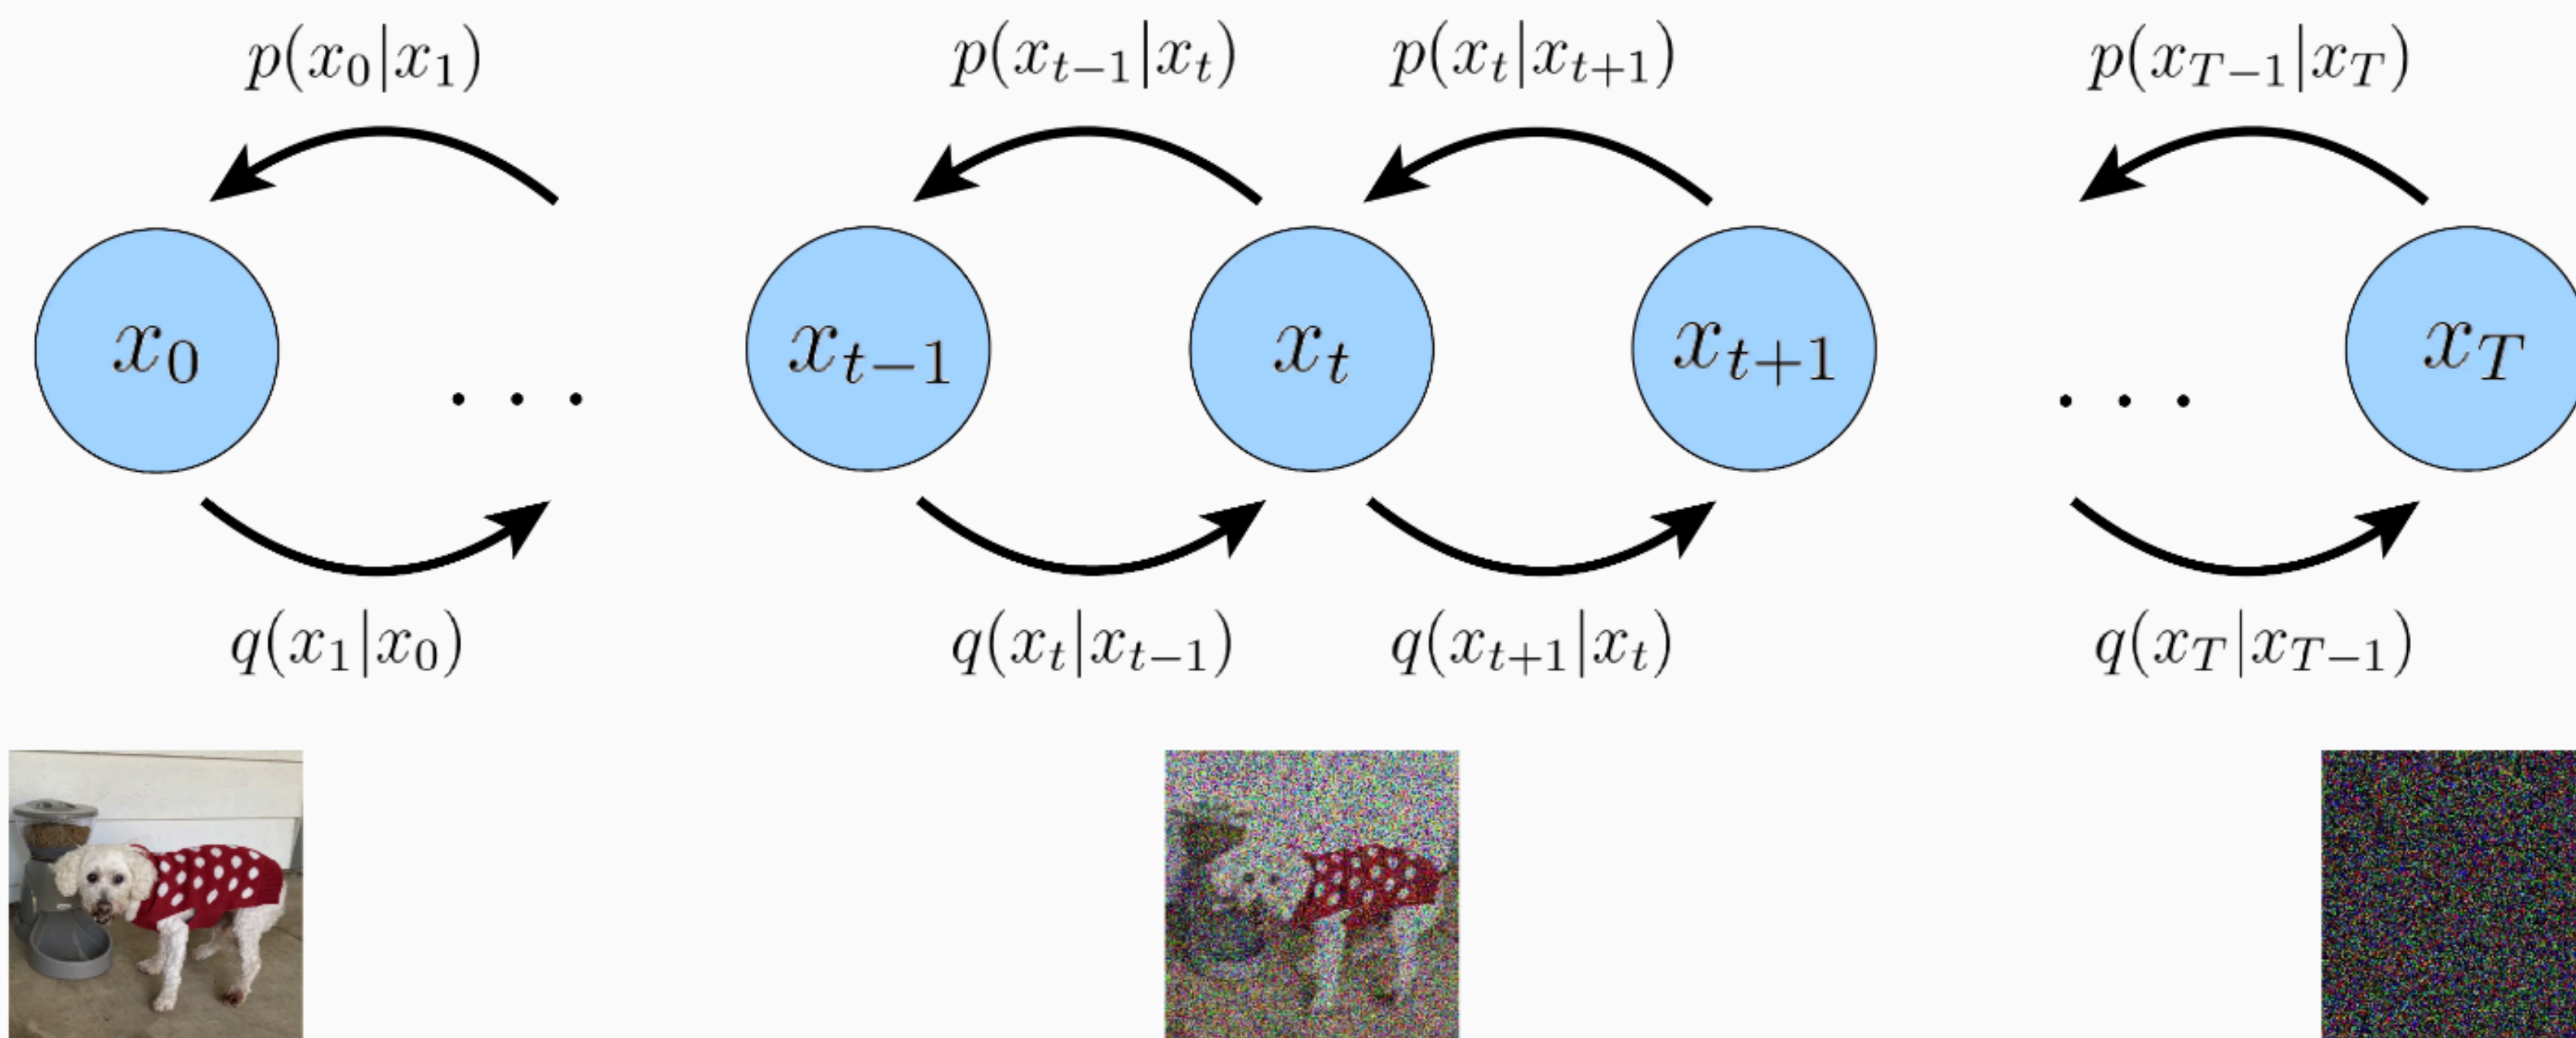
Micah Goldblum



Tom Goldstein




# Diffusion models



Visual representation of a Variational Diffusion Model. An input is steadily noised over time until it becomes identical to Gaussian noise; a diffusion model learns to reverse this process.

**Do you think diffusion models  
replicate training data?**

# Short answer: Yes.



# What counts as replication?

# What counts as replication?

Query image

Perceptual + Semantic

Semantic



# What counts as replication?

Query image



Perceptual + Semantic



Semantic




*“We say that a generated image has replicated content if it contains an object (either in the foreground or background) that appears identically in a training image, neglecting minor variations in appearance that could result from data augmentation.”*



**Can we detect replication with  
existing methods?**

# Current methods

- Instance retrieval / Copy detection
- Self-supervised/ Supervised learning methods




# Current datasets

- Oxford
- Paris
- INSTRE
- CUB-200
- GPR1200

Query image



Perceptual



Semantic




# No partial retrieval datasets!


Query image








Perceptual



Semantic



# Synthetic datasets



Ground truth

Segmix

Diagonal in-  
painting

Patch out-  
painting

# How good are these methods?

Type	Method	Average Rank ↓
CD/IR	Multigrain [6], ResNet-50	5.9
	SSCD [47], ResNet-50	<b>4.2</b>
PT	ViT [16] S/16, IN1k	5.5
	ViT-B/16, IN12k	9.2
	ViT-B/16, CLIP [51] on LAION [57]	5.4
	Swin Transformer [36], Base, IN1k	<b>4.1</b>
SSL	MoCo [15], ViT-B/16	5.1
	MoCo, ViT-B/16 + CutMix [68]	6.5
	VicRegL [5], ResNet-50	5.0
	DINO [12], ViT-B/16, split-product	<b>4.1</b>

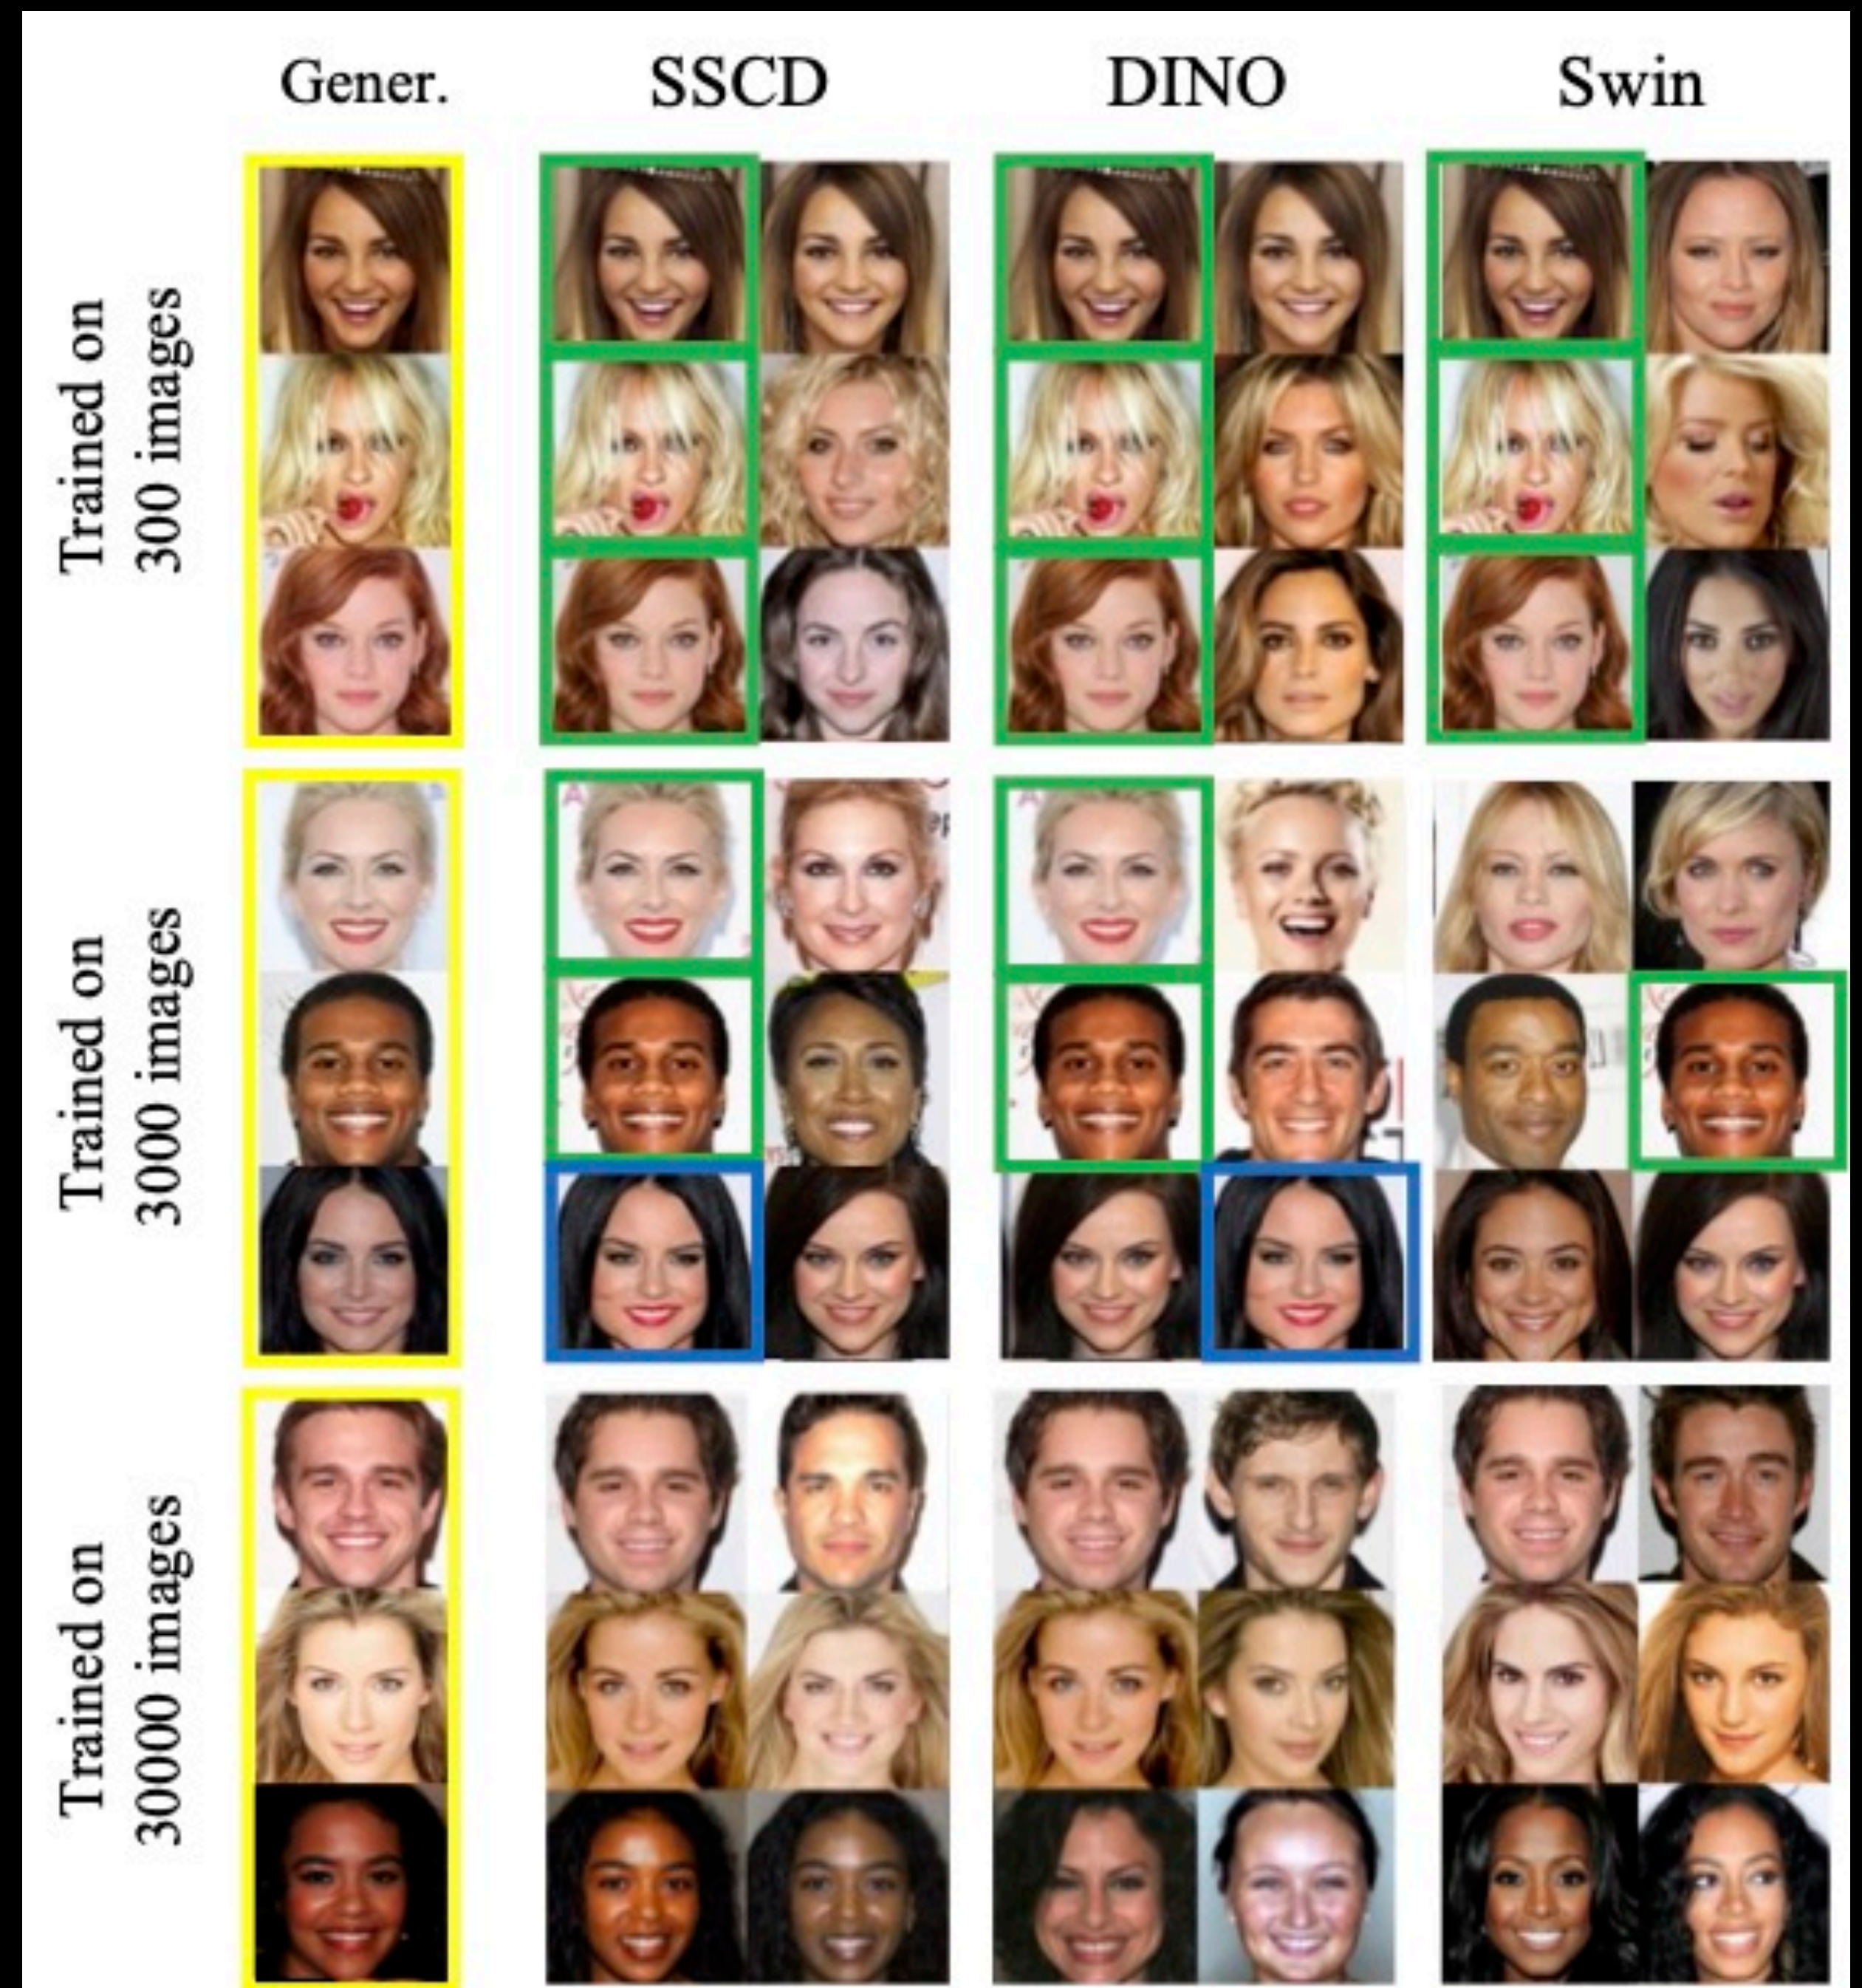
# Can diffusion models copy?

# Can diffusion models copy?

- Overfit a small dataset and see if copying exists.
- Setup: Train a diffusion model (DDPM) on Celeb-A HQ with varying amounts of data (300 | 3000 | 30,000)

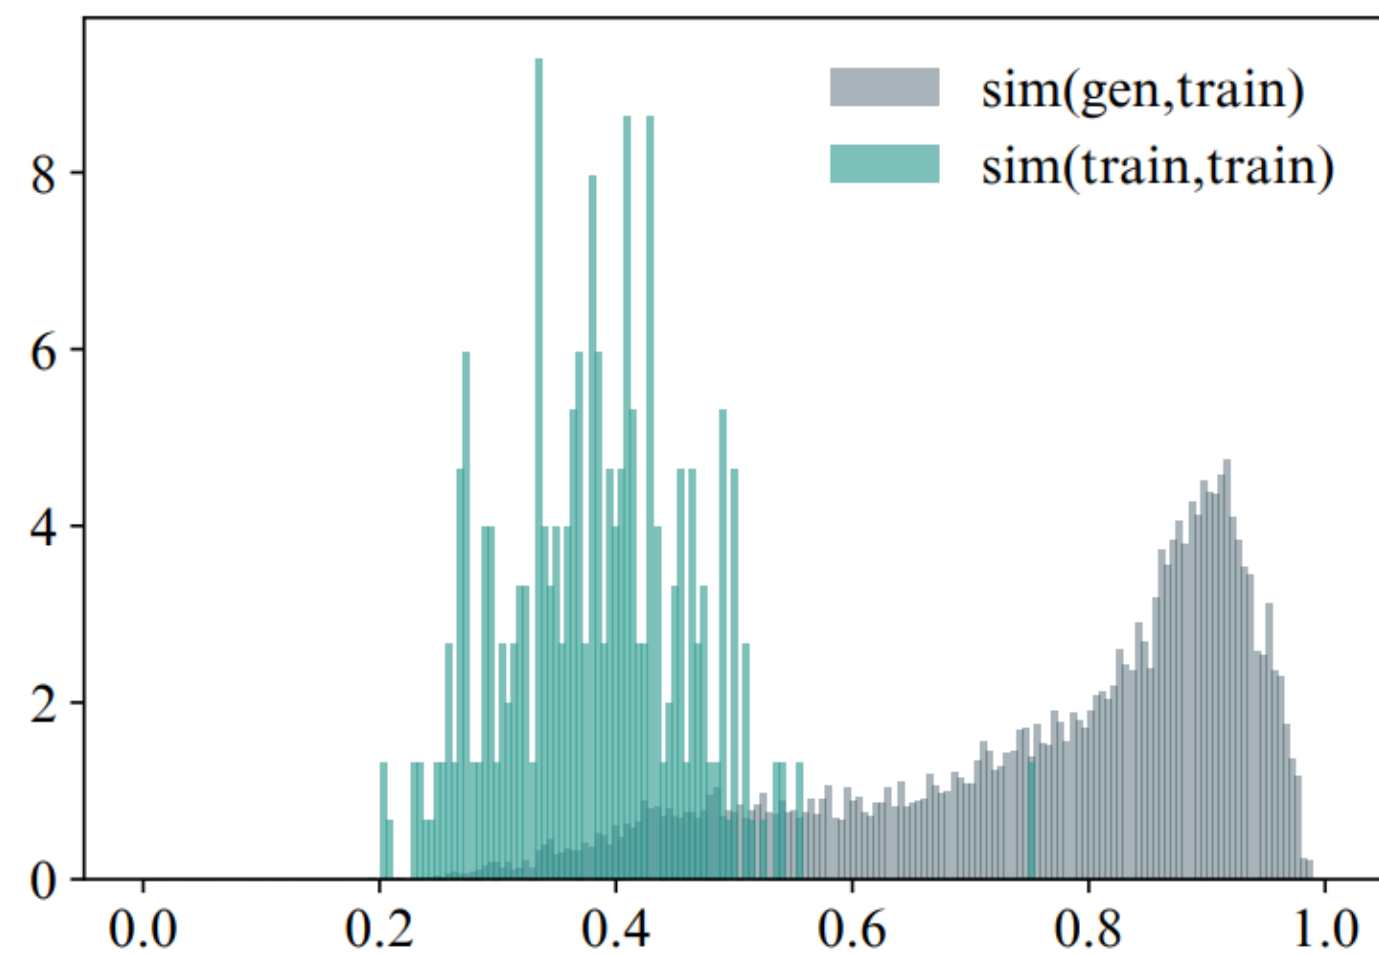


# Can diffusion models copy?

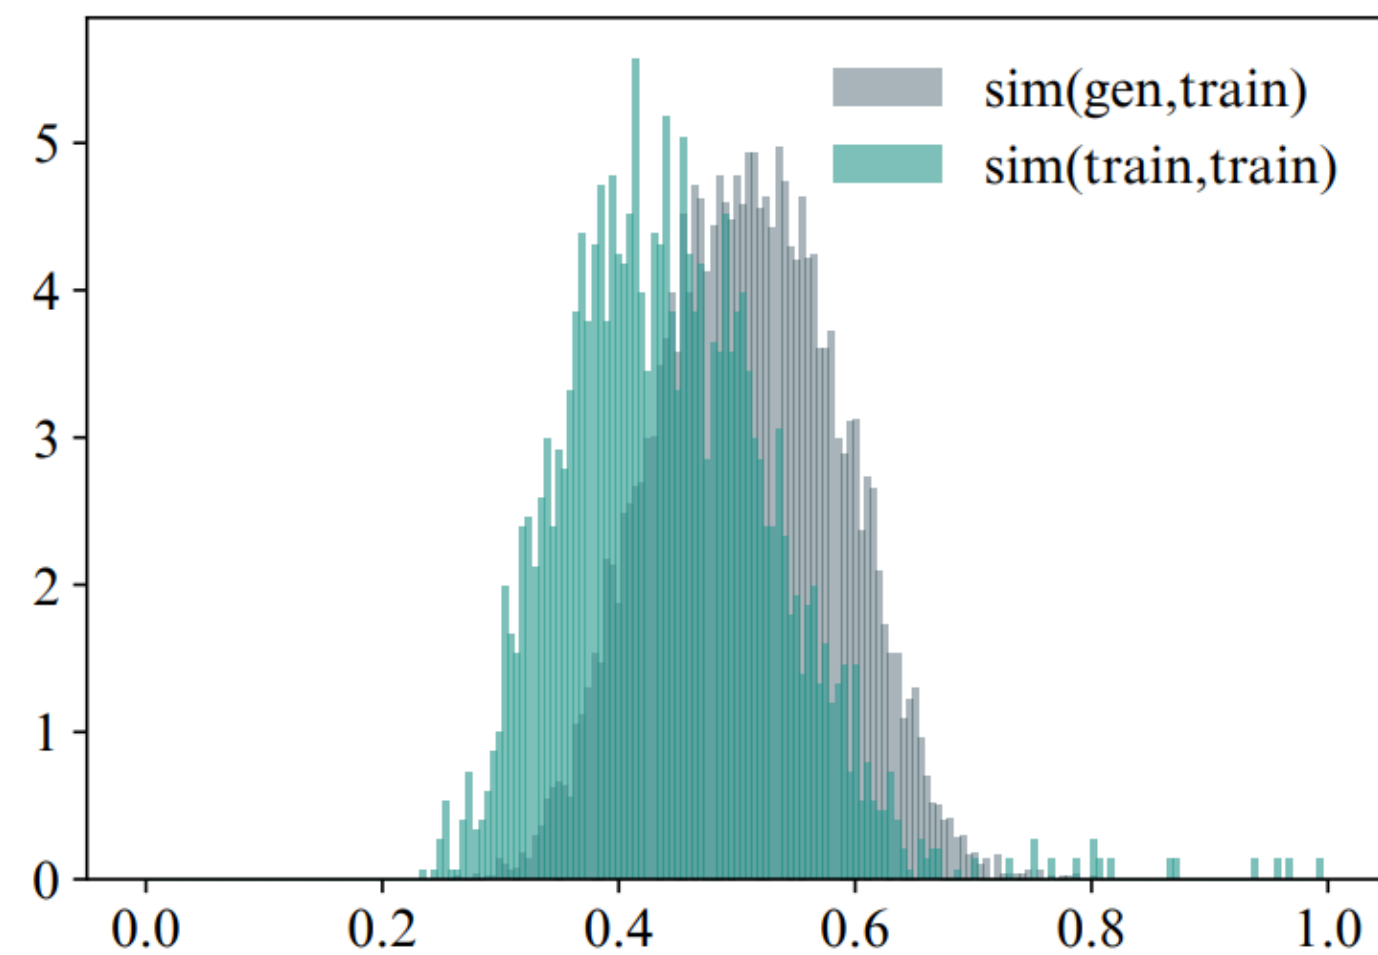


# Can diffusion models copy?


- Setup:
  - Train a diffusion model (DDPM) on Celeb-A HQ with varying amounts of data (300 | 3000 | 30,000)
  - Eval on 10,000 generations



(a) 300 points



(b) 3000 points



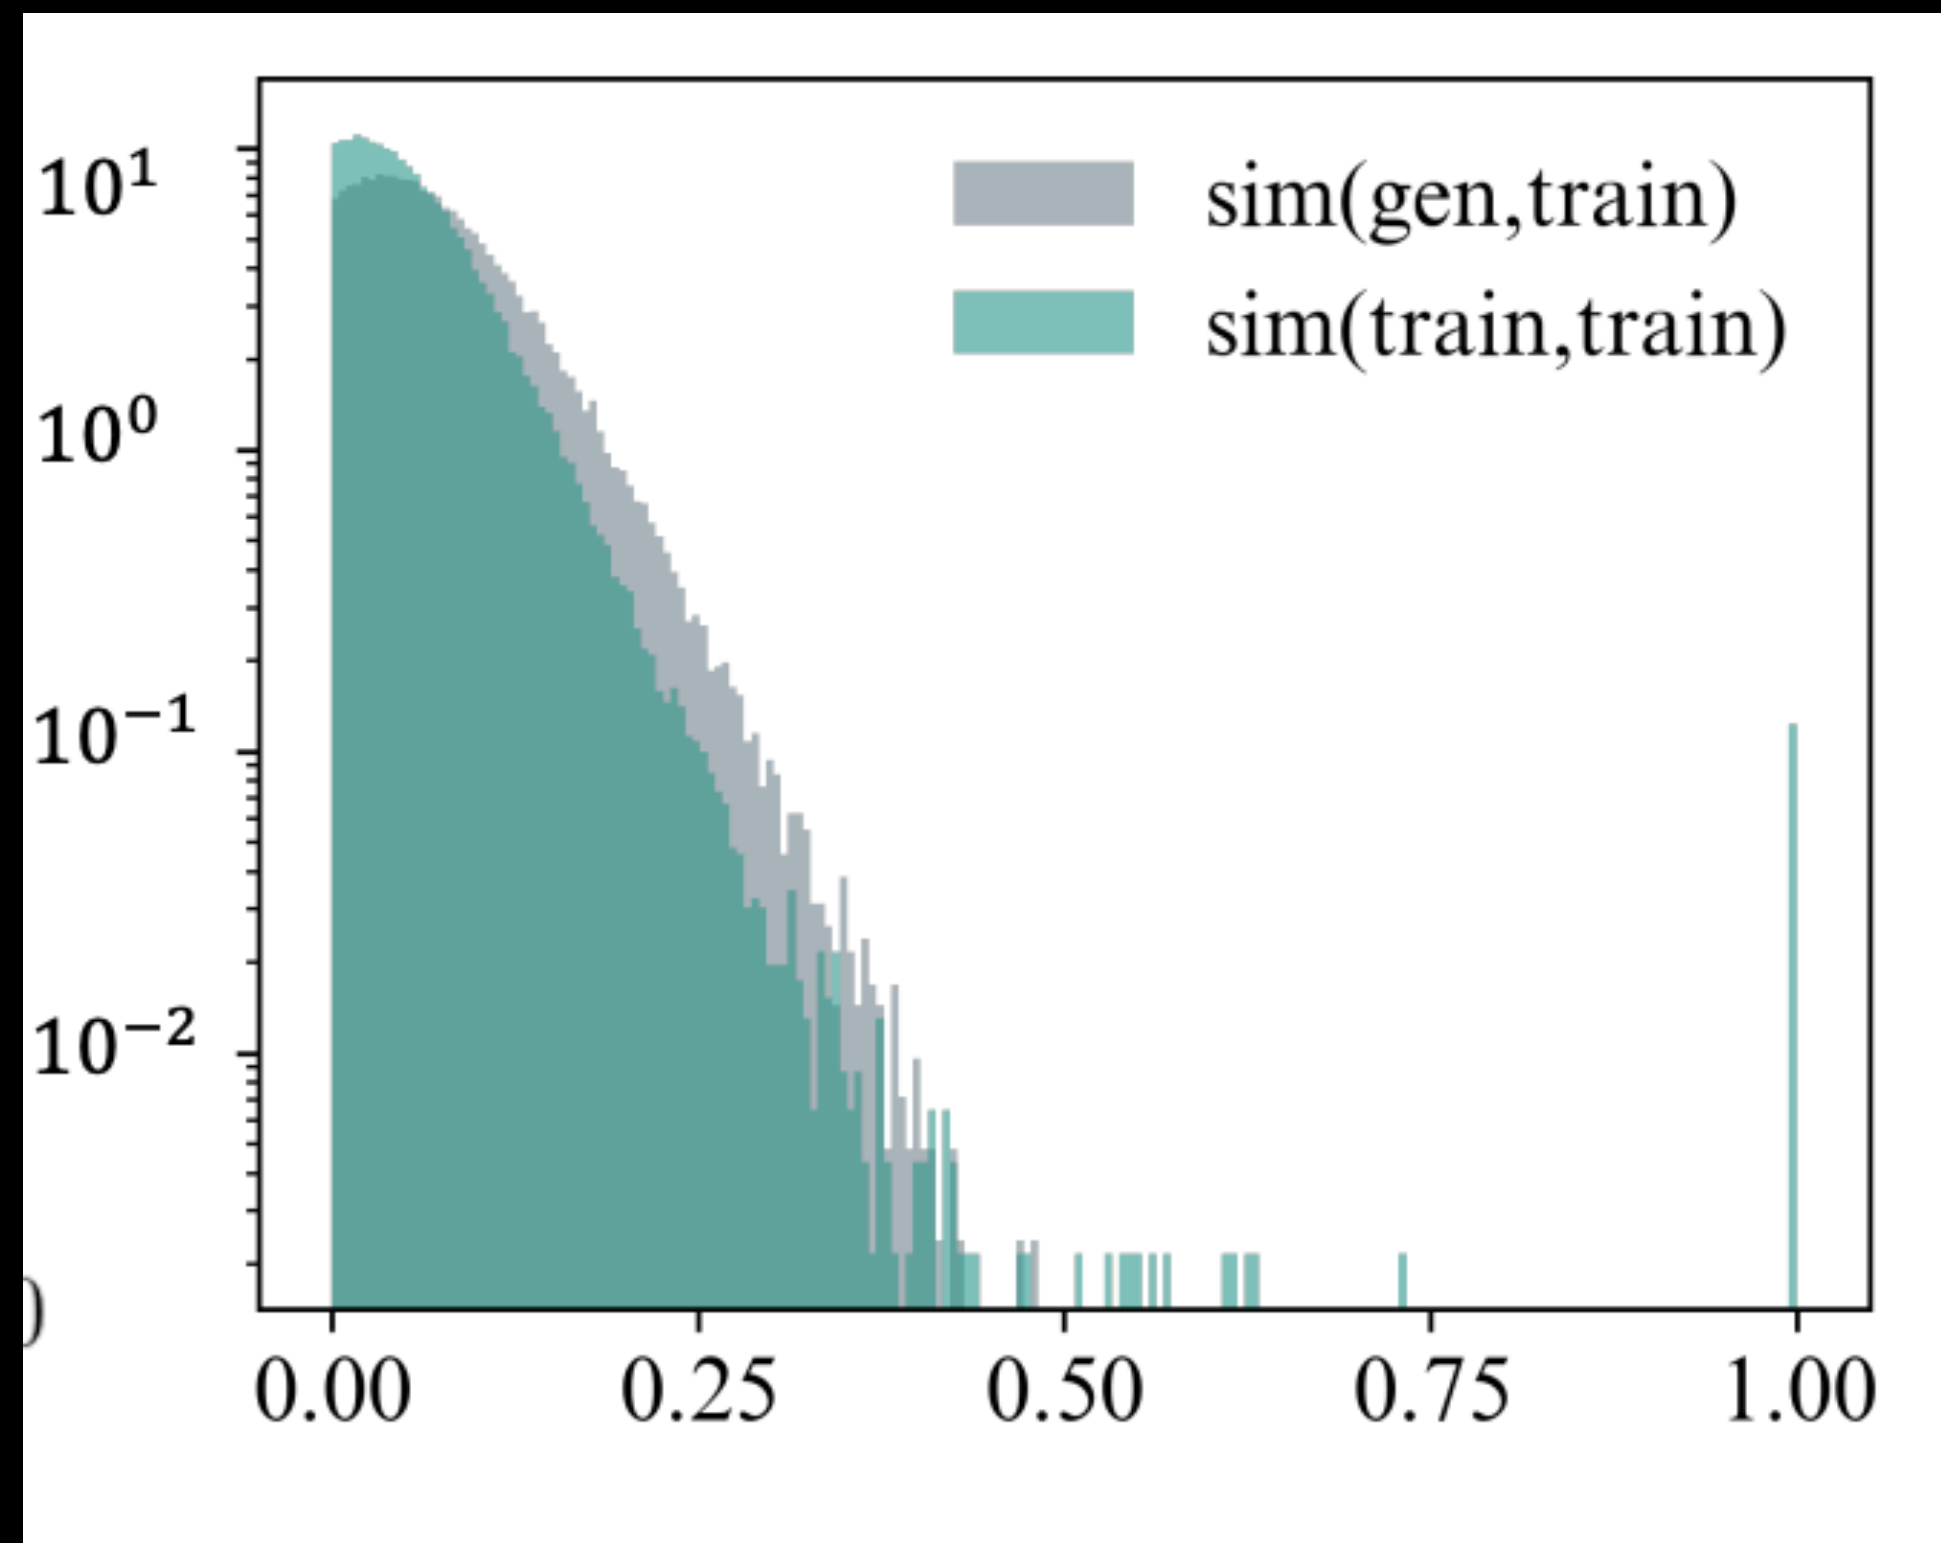
(c) All points

# Do larger diffusion models copy?

- Latent diffusion model (class-conditioned) trained on ImageNet.

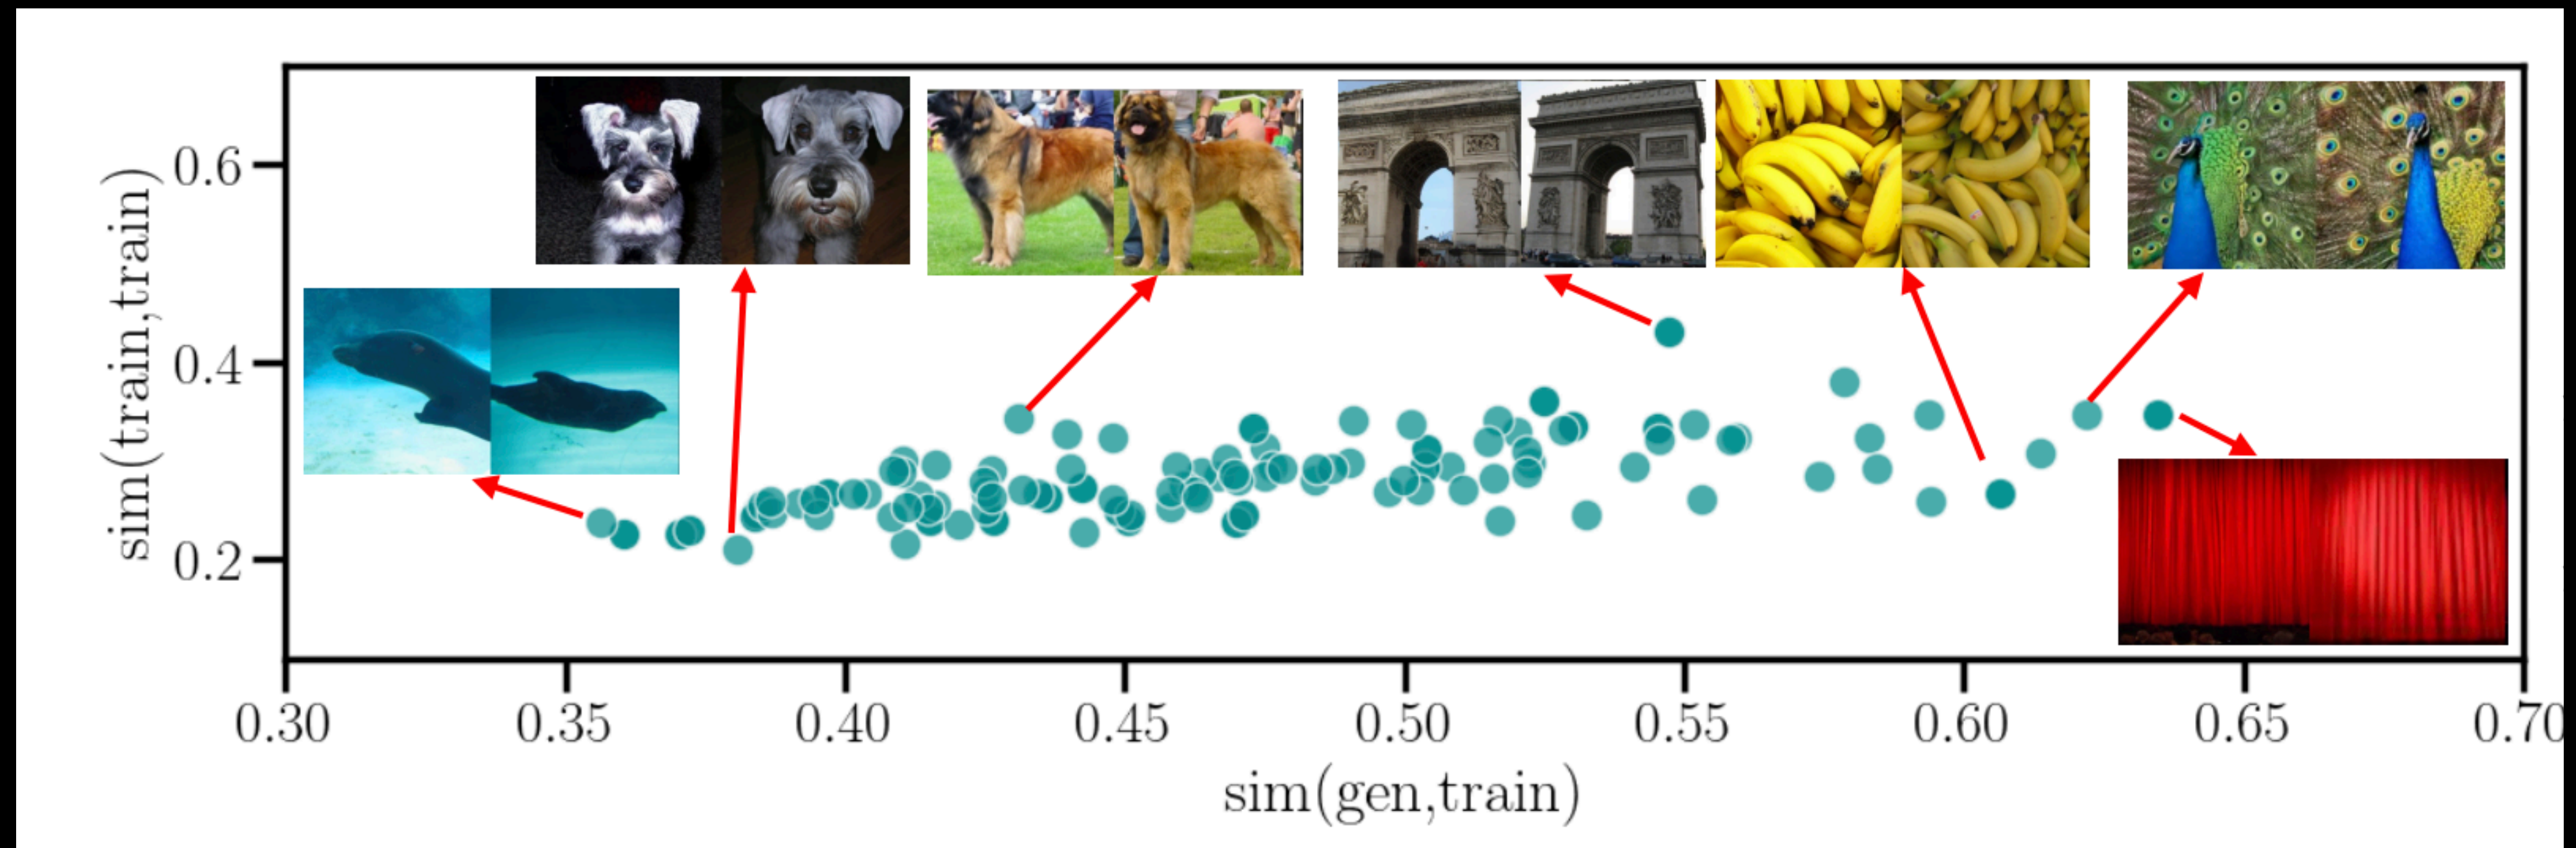
# Do larger diffusion models copy?

- Latent diffusion model (class-conditioned) trained on ImageNet.



# Do larger diffusion models copy?

- Latent diffusion model (class-conditioned) trained on ImageNet.



Diffusion models trained on larger datasets do not copy then? 🤔

# Stable diffusion copies from training data!



# Experimental setup

- SD v1.4 which is trained on LAION 2B dataset and finetuned on LAION-Aesthetics V5+ (~400M images)
- 10,000 generations with captions randomly sampled from LAION-A V6+ (~12M pairs)
- All generations are compared against LAION-A V6+ (a subset of whole training data)

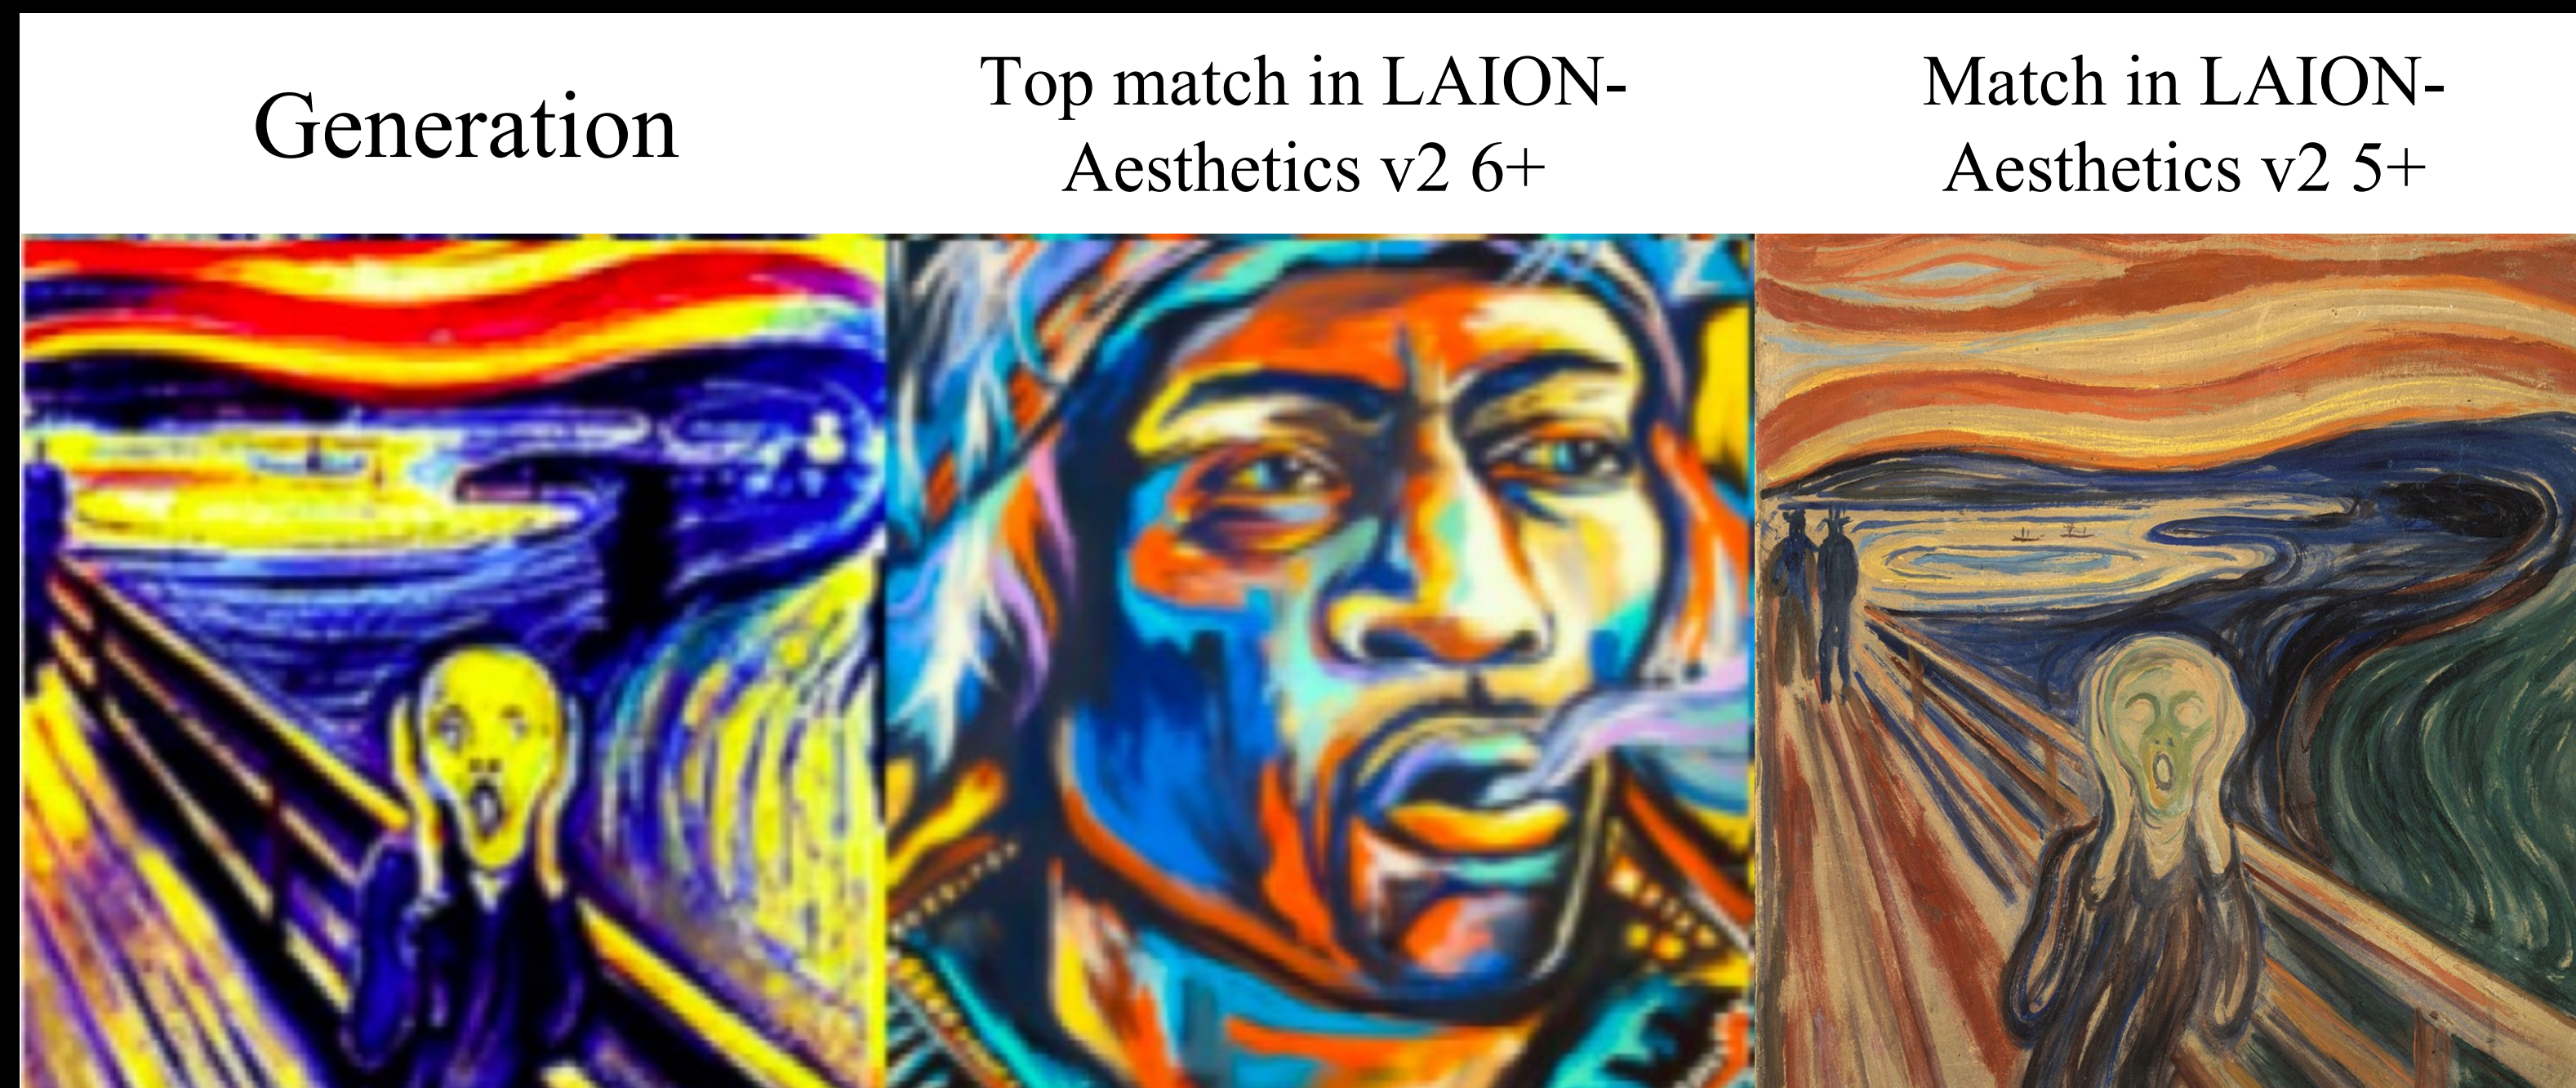


**~1.888%**

**Of generated data are potential copies.**

# How good is this number?

- A lower bound? Since search happened over a subset.



# How good is this number?

- ~~A lower bound? Since search happened over a subset.~~

We analysed only a few captions.

# How good is this number?


- ~~A lower bound? Since search happened over a subset.~~

We analysed only a few captions.

- An upper bound? Since captions are sampled from training data.


# Caption dependance - NONE

Generation



“Hill Country Castle by R Del Angel“

LAION-A match



“Ben Hogan Portrait Golf Legend” (2014) by GinetteCallaway”

# Caption dependance - NONE

Generation



“New Orleans House Galaxy Case”

LAION-A match



“iPhone Cover - Viking - EURO Trailer”

# Caption dependance - word/phrase match

Generation



“Sony Boss Confirms Bloodborne Expansion is Coming”

LAION-A match



“Bloodborne 32 Minute Speedrun”

# Caption dependance - word/phrase match



Prompt: <The description of the wall art> Canvas Wall Art Print



Prompt: A painting of the Great Wave off Kanagawa by Katsushika Hokusai.



# How good is this number?

- ~~A lower bound? Since search happened over a subset.~~

We only analysed only a few captions.

- ~~An upper bound? Since captions are sampled from training data.~~

Certain words and phrases can trigger copying, we do not need whole caption!

# Possible causes of replication

- Data duplication
- Text conditioning

# Replication spectrum


Content and style copied

Style copied

Generation



Top Match



# Thank you!

Project page: <https://somepage.github.io/diffrep.html>