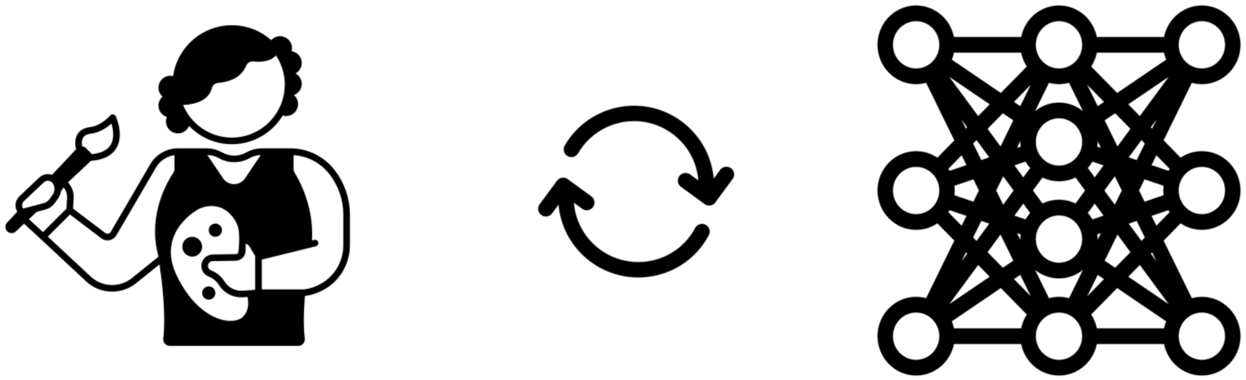

Ensuring Data Ownership in Generative Visual Models

Jun-Yan Zhu

Generative Intelligence Lab

Carnegie Mellon University

Carnegie Mellon University

Generative Models (2016)

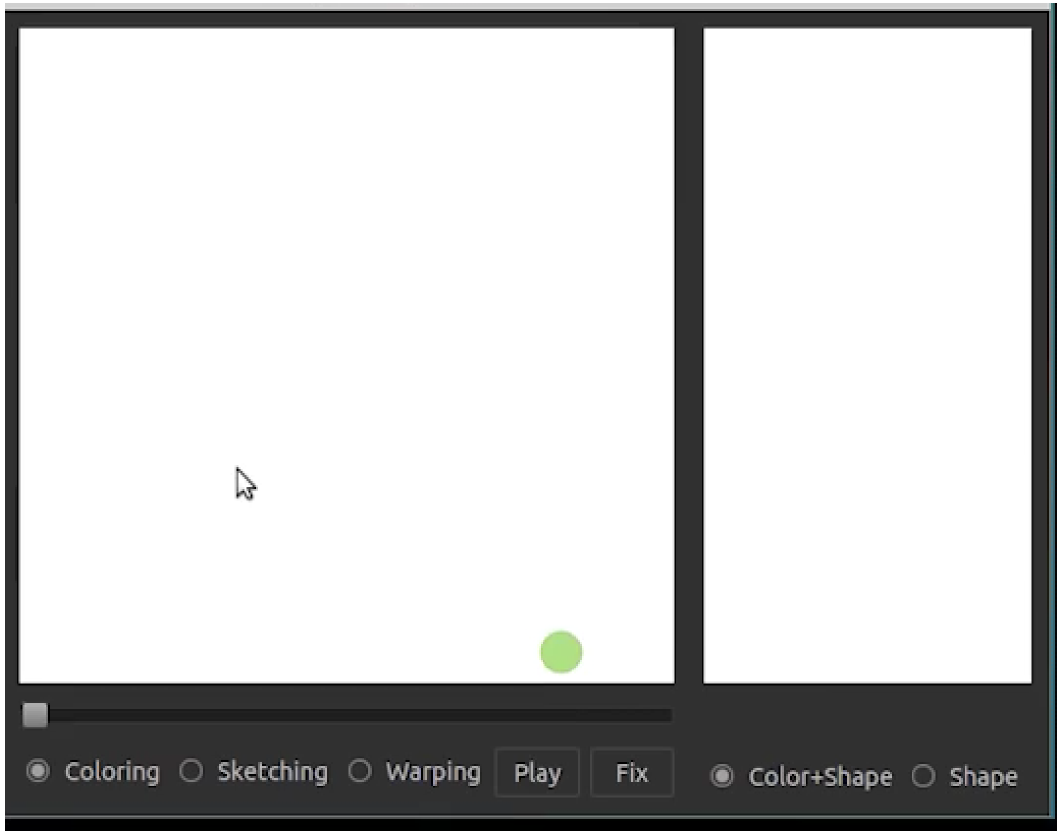

GAN Inversion [Zhu et al., ECCV 2016]

A toilet seat sits open in the grass field.

text2image [Mansimov et al., ICLR 2016] from Ruslan Salakhutdinov's group

#edges2cats with pix2pix (2017)

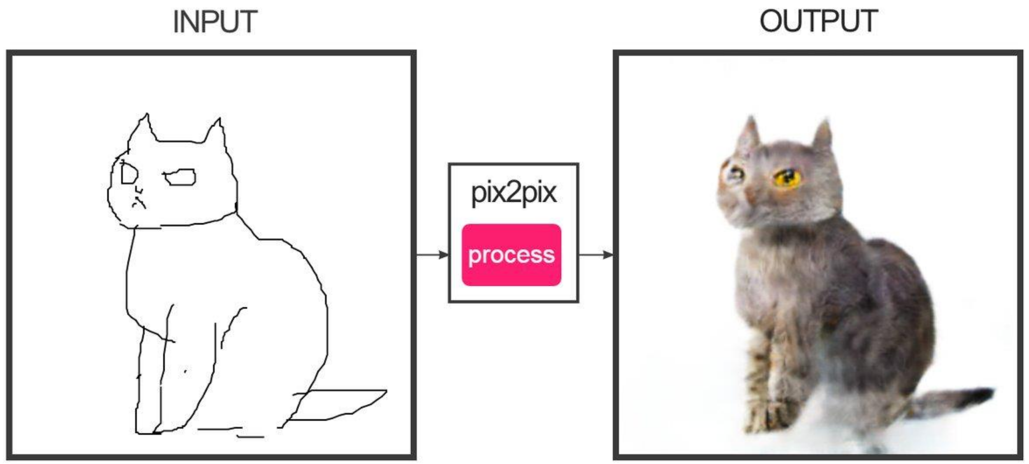

@gods_tail

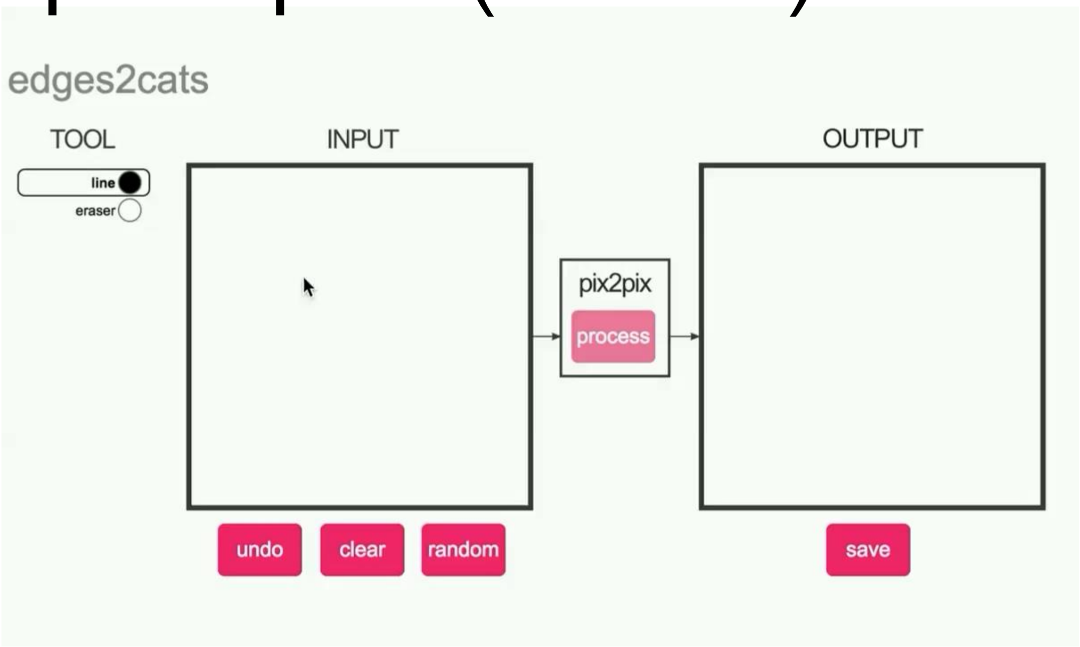

@matthematician

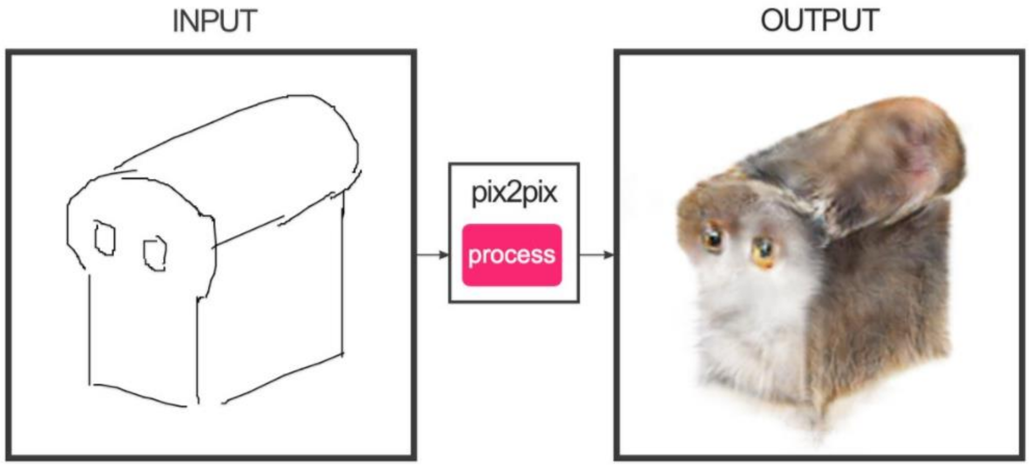

Ivy Tasi @ivymyt

Vitaly Vidmirov @vvid

[Isola, Zhu, Zhou, Efros. CVPR 2017]

By Christopher Hesse https://affinelayer.com/pixsrv/

GauGAN System (2019)

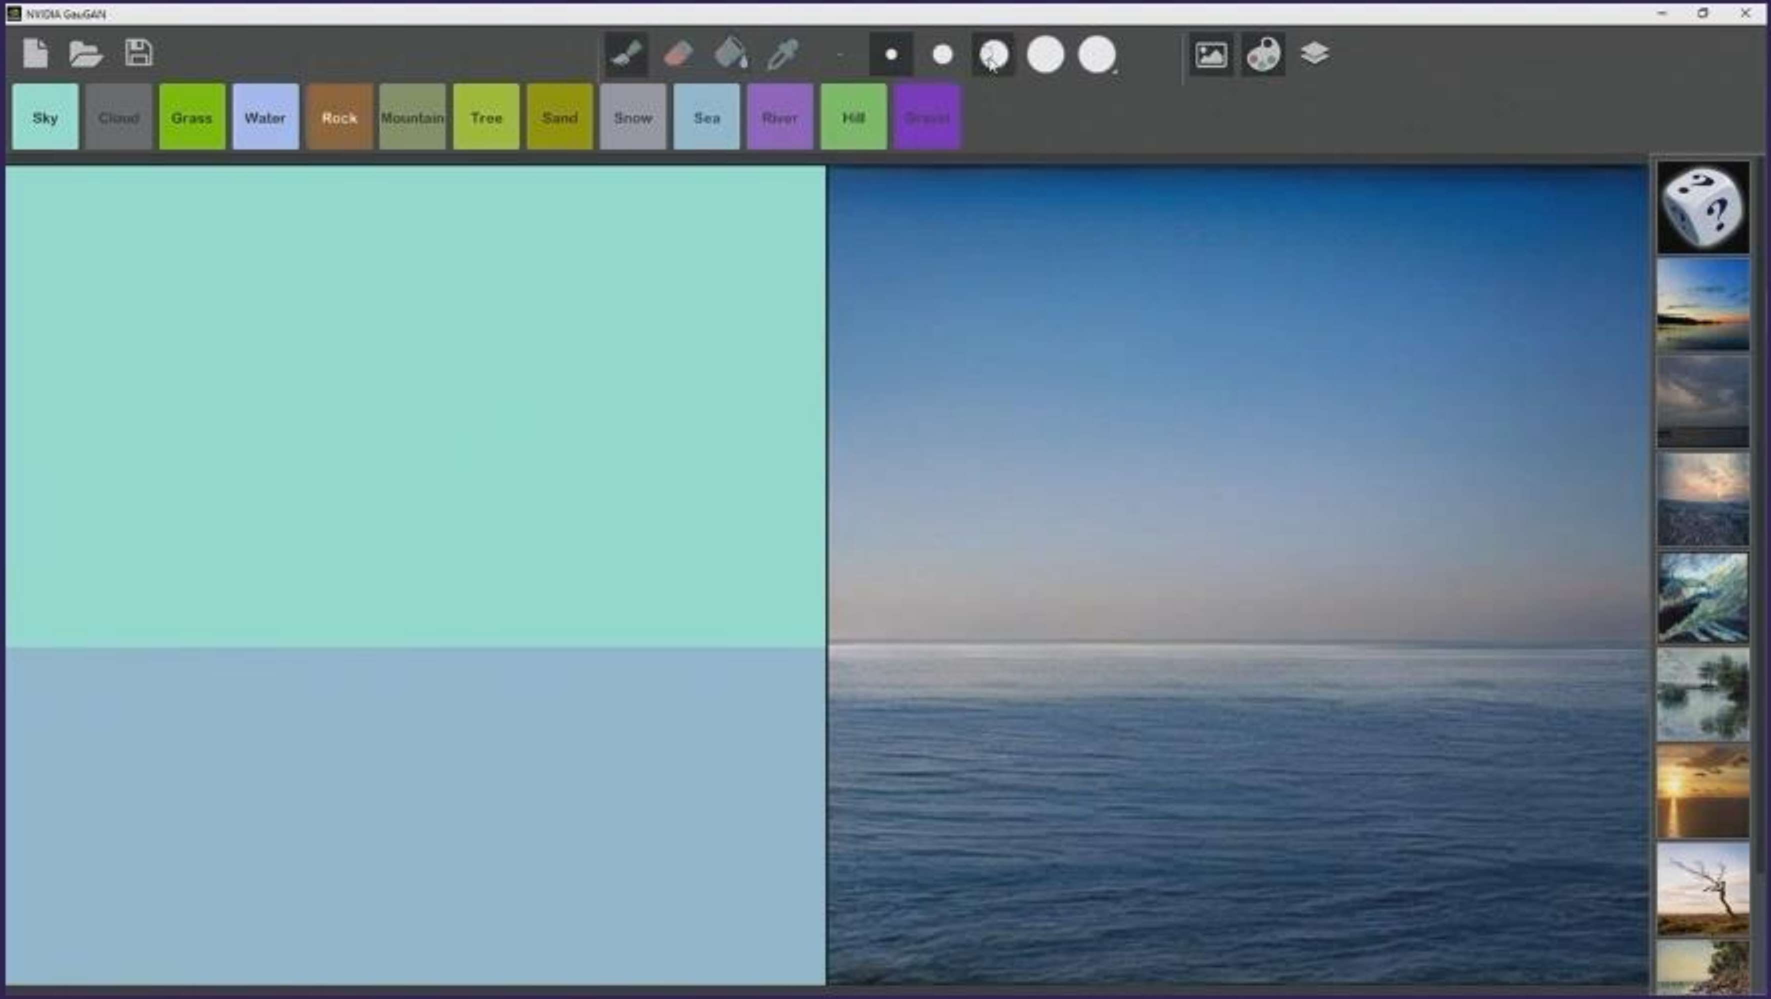

thrive
SIGGRAPH2019
LOS ANGELES 28 JULY - 1 AUGUST
© 2019 SIGGRAPH. ALL RIGHTS RESERVED.

[Park, Liu, Wang, Zhu. CVPR 2019]

SIGGRAPH 2019 Real-time Live! “Best of Show Award" and "Audience Choice Award"

SDEdit: Guided Image Synthesis with Diffusion

Input User Drawing

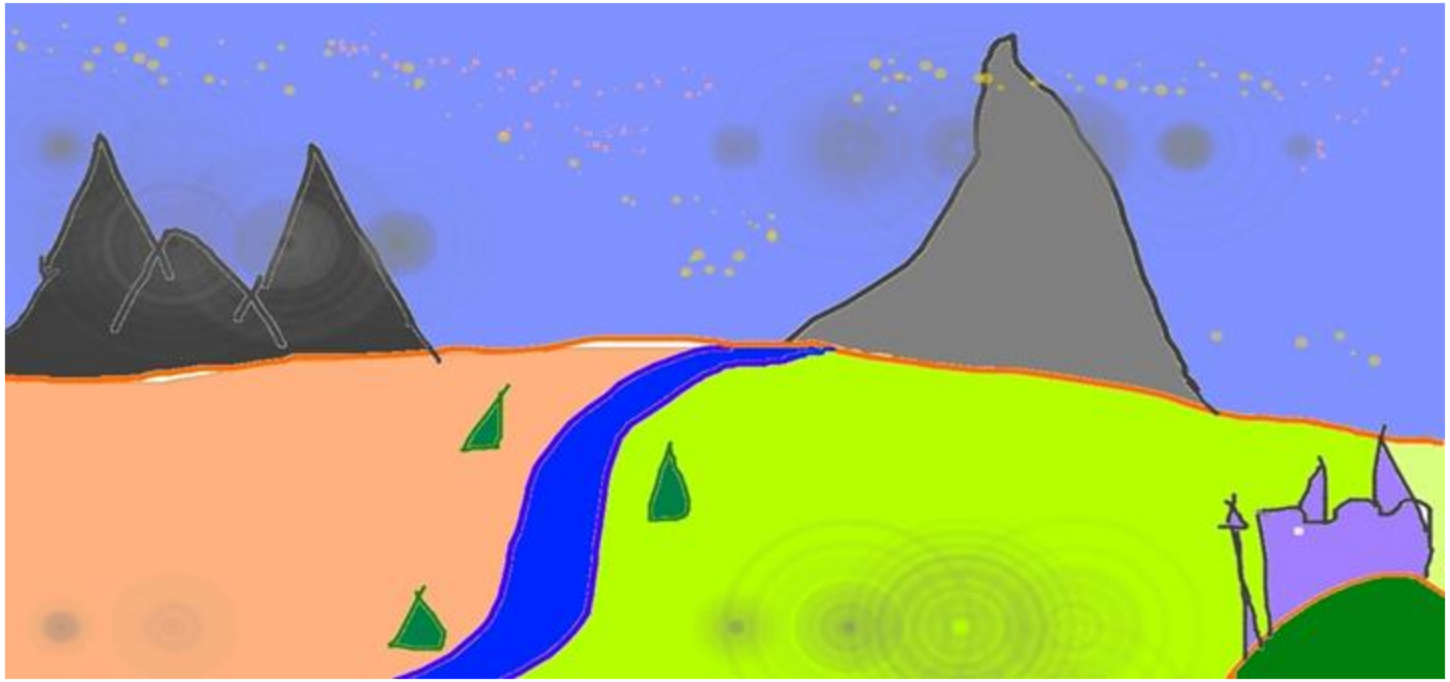

Used in Stable Diffusion Image-to-Image (“img2img’)

[Meng et al., ICLR 2022]

SDEdit: Guided Image Synthesis with Diffusion

Text prompt: “A fantasy landscape, trending on artstation”

Used in Stable Diffusion Image-to-Image (“img2img’)

[Meng et al., ICLR 2022]

SDEdit: Guided Image Synthesis with Diffusion

https://www.reddit.com/r/StableDiffusion/comments/wyq04v/using_img2img_to_upgrade_my_sons_artwork/

<u>Concurrent work with SDEdit:</u> ILVR [Choi et al., 2021]

<u>See more recent works</u>: prompt-to-prompt, Imagic, pix2pix-zero, Edict, Plug & Play, Instruct-pix2pix, ControlNet, etc.

[Meng et al., ICLR 2022]

pix2pix-turbo (2024)

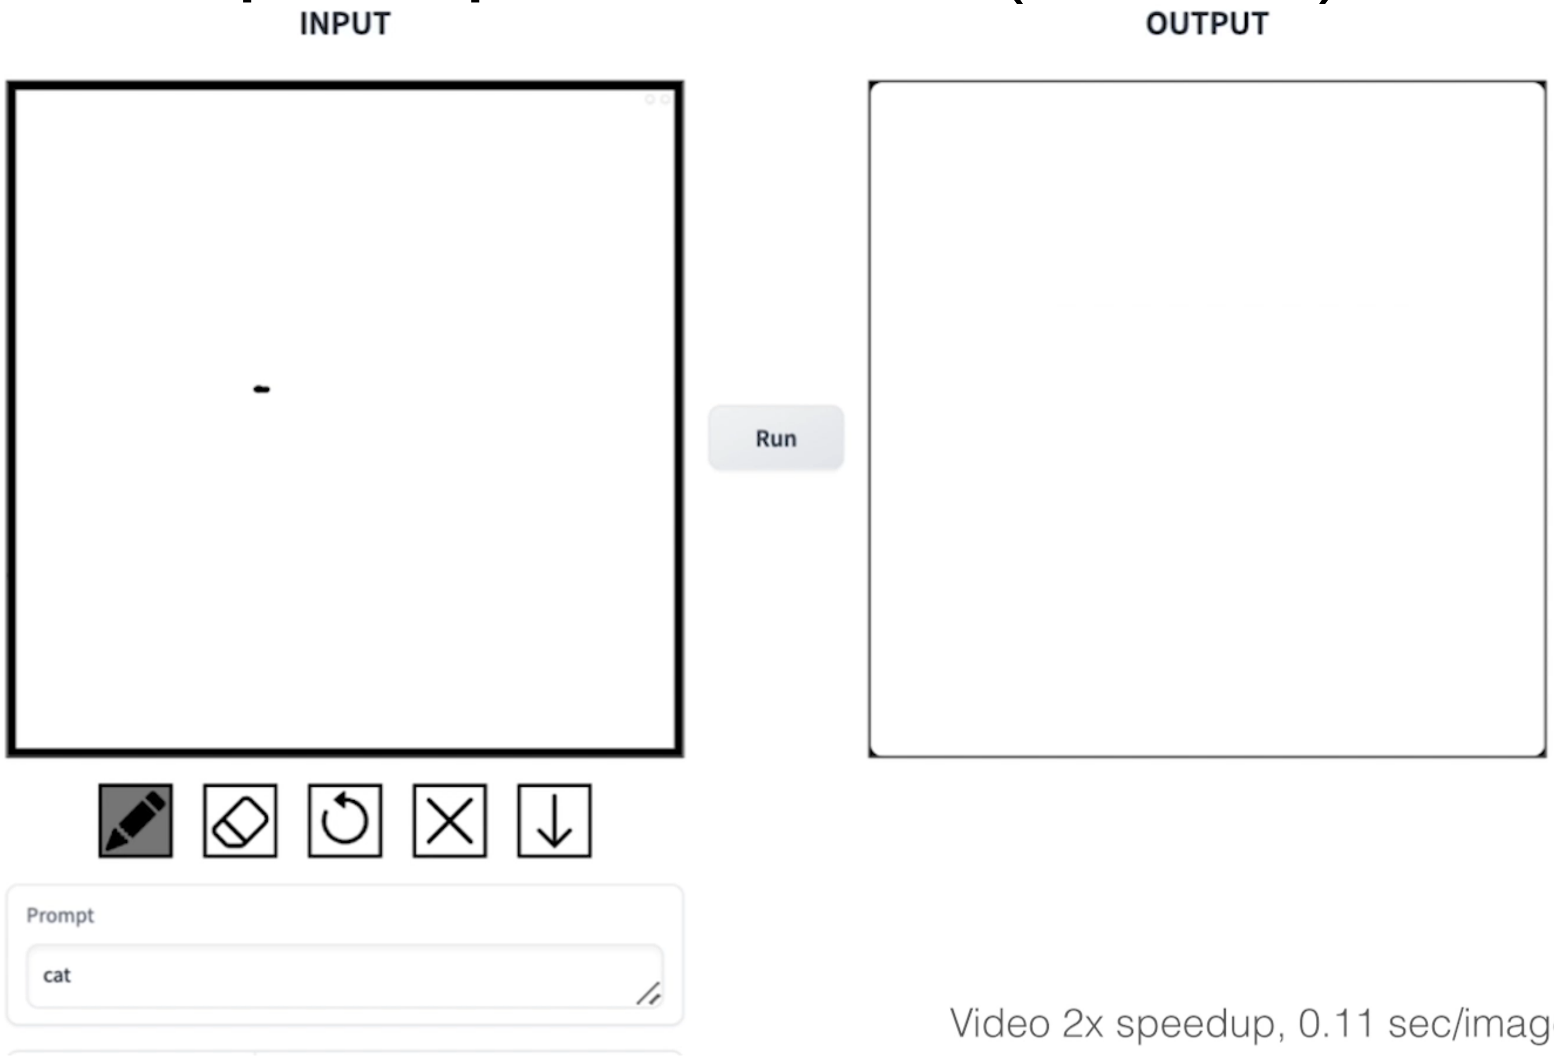

Video 2x speedup, 0.11 sec/image on A100

[Parmer et al., 2024]

Generative Models (2024)

Diffusion models
(DALL-E 2, Imagen, SD)

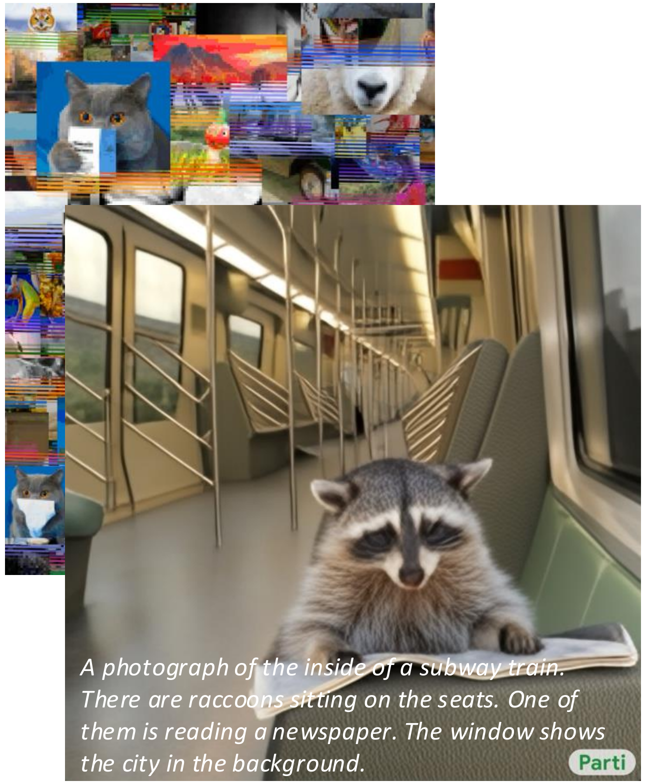

Autoregressive models
(Image GPT, Parti)

GANs, Masked GIT
(GigaGAN, MUSE)

Generative Models (2024)

By DALL·E 3

Generative ~~Models~~ AI (2024)

Stability AI, the startup behind Stable Diffusion, raises $101M
By Kyle Wiggers, Tech Crunch. (Image credits: Bryce Durbin)

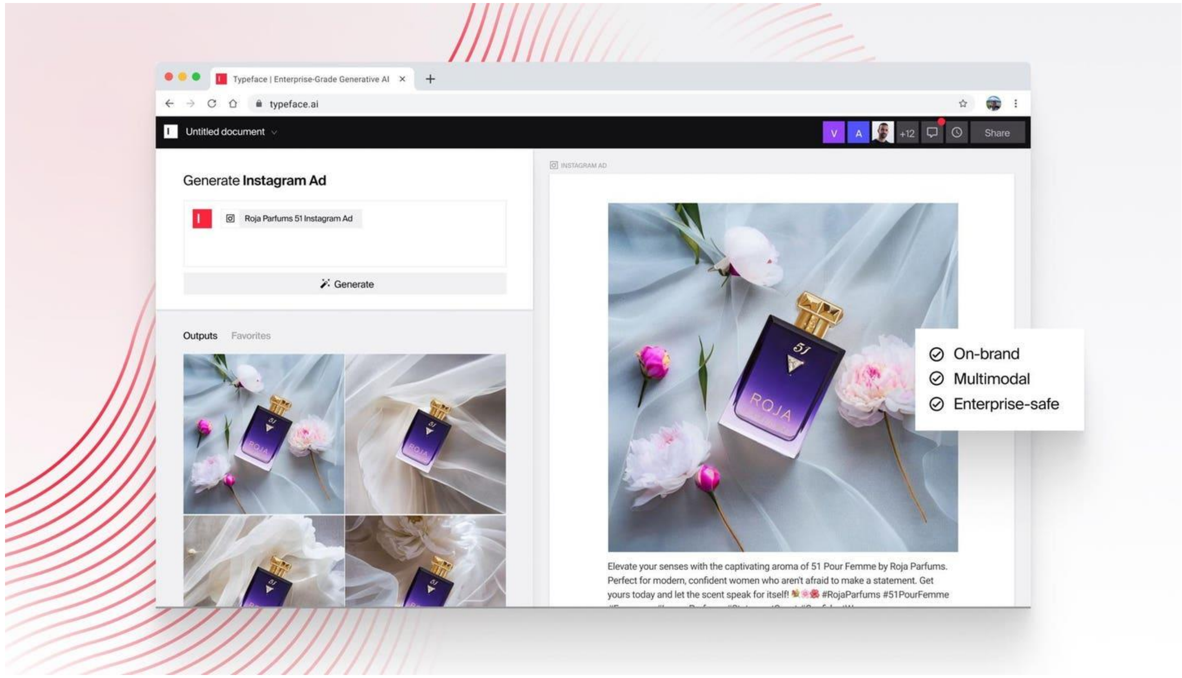

Typeface Raises $100 Million To Set Up AI ‘Content Factories’ For Enterprises. By Rashi Shrivastava, Forbes
(Image credits: Typeface)

Machine Learning Pipeline

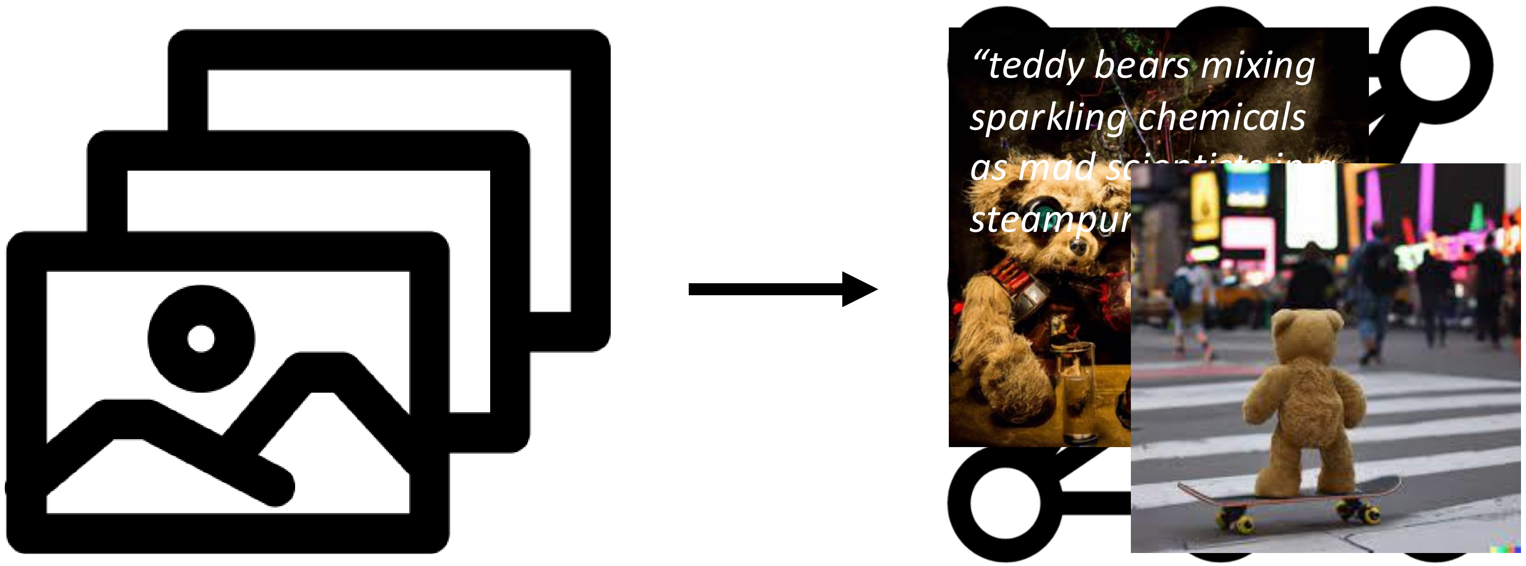

Training images

Model

Data Comes from People!

So researchers & founders are excited, but…

Ongoing Legal Battles

ARTIFICIAL INTELLIGENCE / TECH / LAW

Getty Images sues AI art generator Stable Diffusion in the US for copyright infringement

An illustration from Getty Images' lawsuit, showing an original photograph and a similar image (complete with Getty Images watermark) generated by Stable Diffusion. Image: Getty Images

/ Getty Images has filed a case against Stability AI, alleging that the company copied 12 million images to train its AI model 'without permission ... or compensation.'

By JAMES VINCENT
Feb 6, 2023, 11:56 AM EST | 16 Comments / 16 New

Getty Images has filed a lawsuit in the US against Stability AI, creators of open-source AI art generator Stable Diffusion, escalating its legal battle against the firm.

ARTIFICIAL INTELLIGENCE / TECH / CREATORS

AI art tools Stable Diffusion and Midjourney targeted with copyright lawsuit

A collage of AI-generated images created using Stable Diffusion. Image: The Verge via Lexica

/ The suit claims generative AI art tools violate copyright law by scraping artists' work from the web without their consent.

By JAMES VINCENT
Jan 16, 2023, 6:28 AM EST | 28 Comments / 28 New

A trio of artists have launched a lawsuit against Stability AI and Midjourney, creators of AI art generators Stable Diffusion and Midjourney, and artist portfolio platform DeviantArt, which recently created its own AI art generator, DreamUp.

Source: The Verge

Ongoing Legal Battles

REUTERS

World Business Markets Legal Breakingviews Technology Investigations Sports

Copyright | Technology | Intellectual Property | Litigation | Data Privacy

2 minute read · February 22, 2023 8:41 PM EST · Last Updated 2 months ago

AI-created images lose U.S. copyrights in test for new technology

By Blake Brittain

REUTERS/Andrew Kelly

Feb 22 (Reuters) - Images in a graphic novel that were created using the artificial-intelligence system Midjourney should not have been granted copyright protection, the U.S. Copyright Office said in a letter seen by Reuters.

I'm not so sure. As we've seen, a key assumption for a "non-expressive use" defense is that Stable Diffusion only learns uncopyrightable facts—not creative expression—from its training images. That's *mostly* true. But it's not entirely true. And the exceptions could greatly complicate Stability AI's legal defense.

Stable Diffusion's copying problem

Here's one of the most awkward examples for Stability AI:

Enlarge

https://arstechnica.com/

Hollywood Strikes against AI

In Hollywood writers' battle against AI, humans win (for now). By JAKE COYLE, AP News

If artificial intelligence uses your work, it should pay you
By Joseph Gordon-Levitt, The Washington Post

Source: AP news

Digital Artists are Pushing Back

@loisvb's Instagram Post

[https://hyperallergic.com/806026/digital-artists-are-pushing-back-against-ai]

Digital Artists are Pushing Back

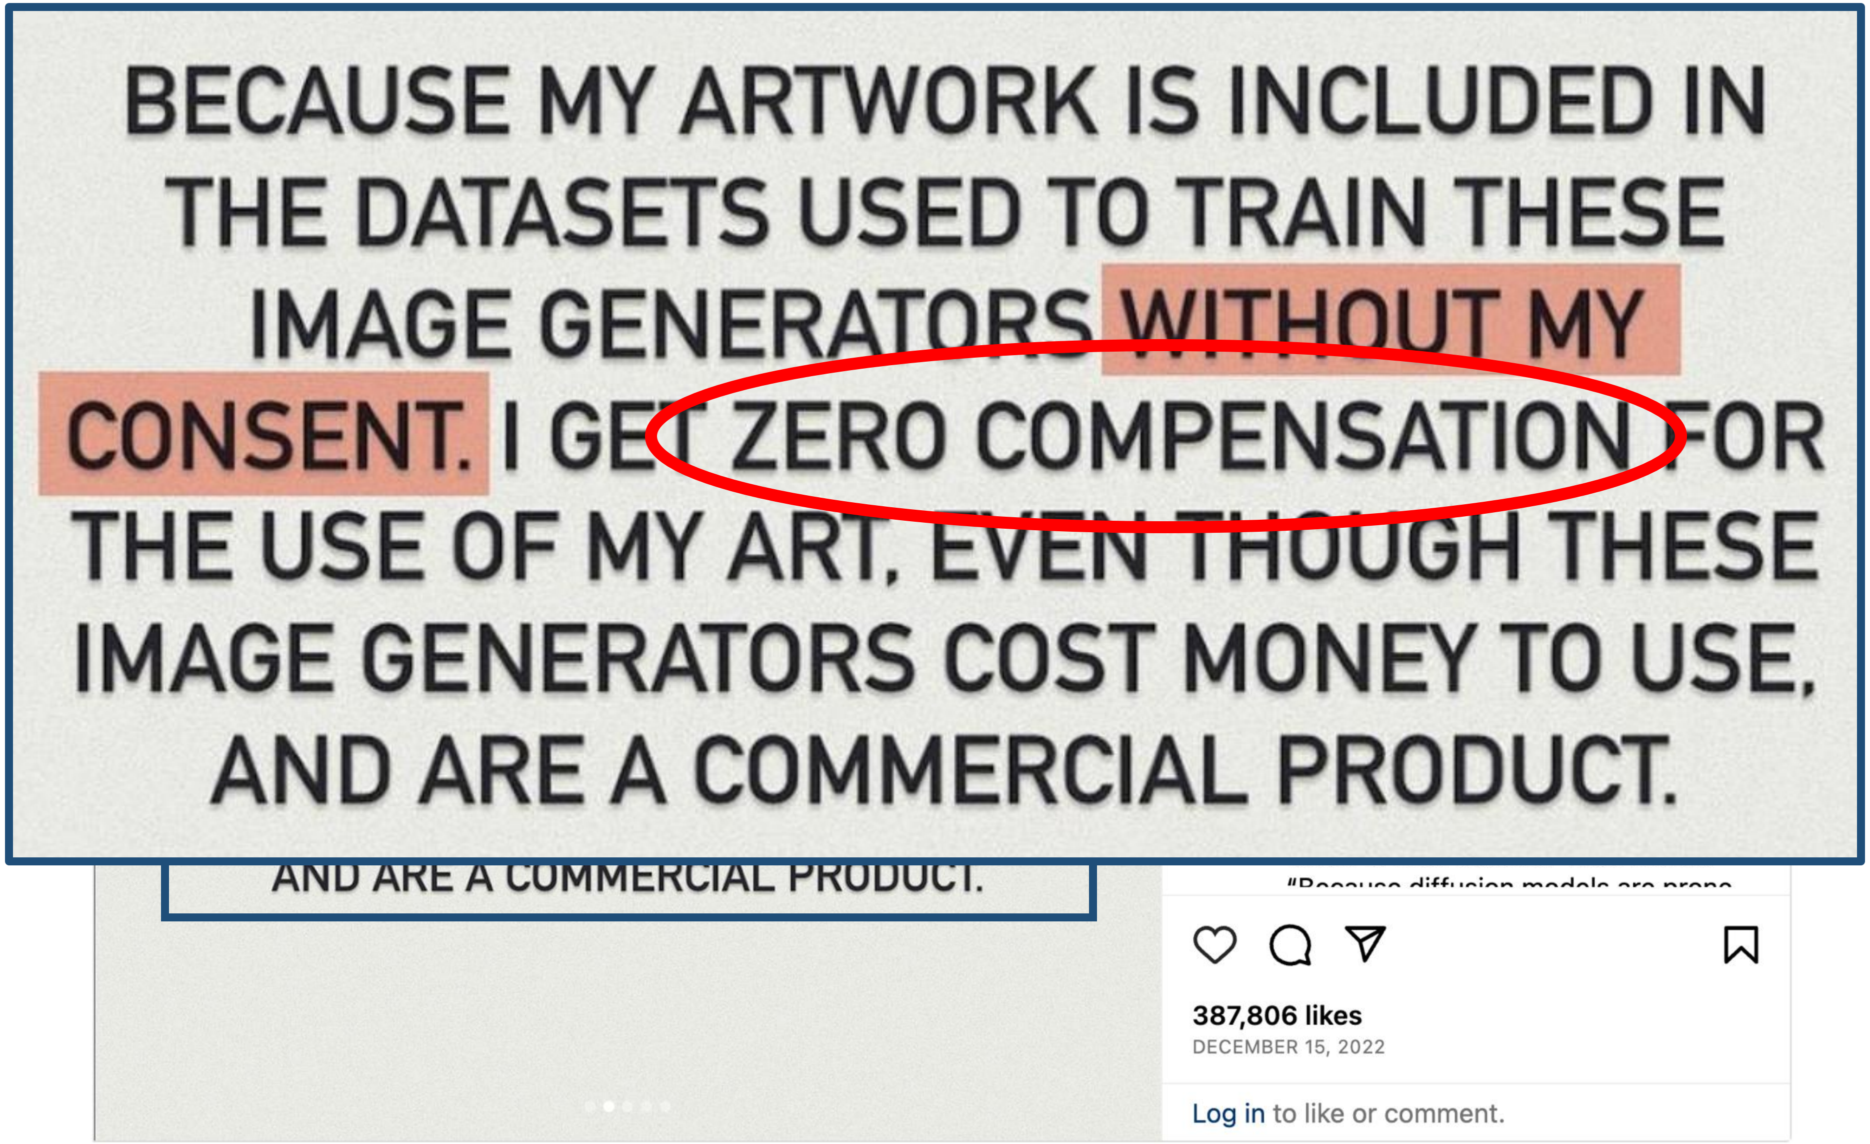

@loisvb's Instagram Post

[https://hyperallergic.com/806026/digital-artists-are-pushing-back-against-ai]

Generative models use training data of artists, photographers, and creators

without Consent

without Compensation

Copyright Issues

- Copyrighted images.
- Company IPs / logos.
- Artist styles of living artists.

Getty Images

Greg Rutkowski

Memorized Style

Stable Diffusion

Greg Rutkowski

A painting of a boat on the water in the style of Greg Rutkowski

*image taken from https://rutkowski.artstation.com

Memorized Instances

THE TWO-WAY

Grumpy Cat Awarded $710,000 In Copyright Infringement Suit

January 25, 2018 · 8:45 AM ET

By Scott Neuman

EU GDPR: Right to erasure (right to be forgotten)

Concept Ablation: remove <u>**copyrighted training data**</u>!

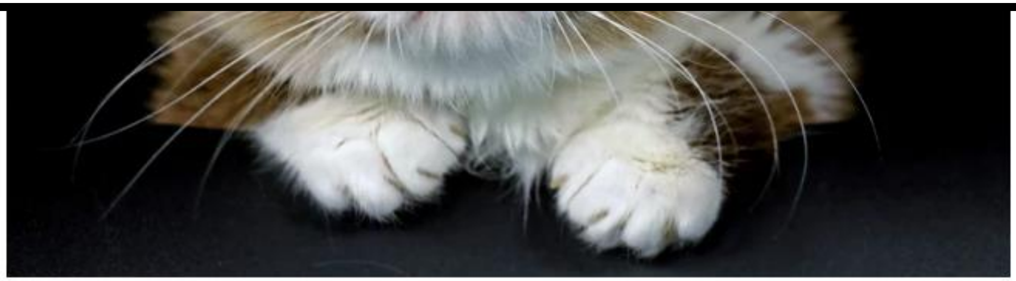

Grumpy Cat appears unimpressed posing for a photo during an interview at The Associated Press bureau in Los Angeles in December 2015.

Richard Vogel/AP

Diffusion Model Overview

Diffusion Model Overview

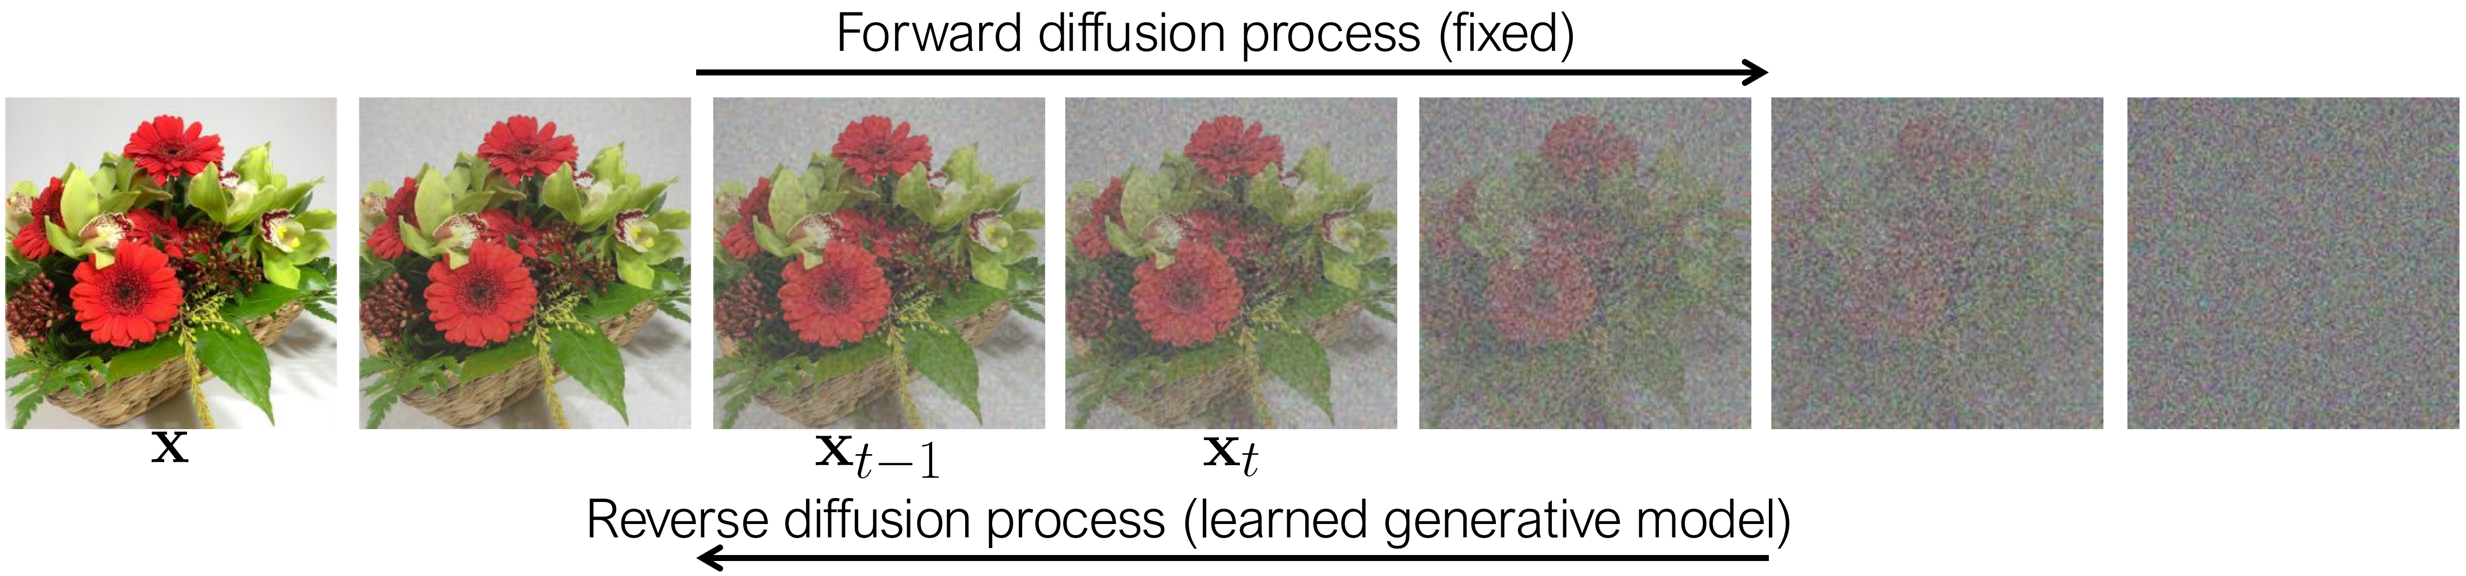

*slides motivated from https://cvpr2022-tutorial-diffusion-models.github.io

Diffusion Model Training

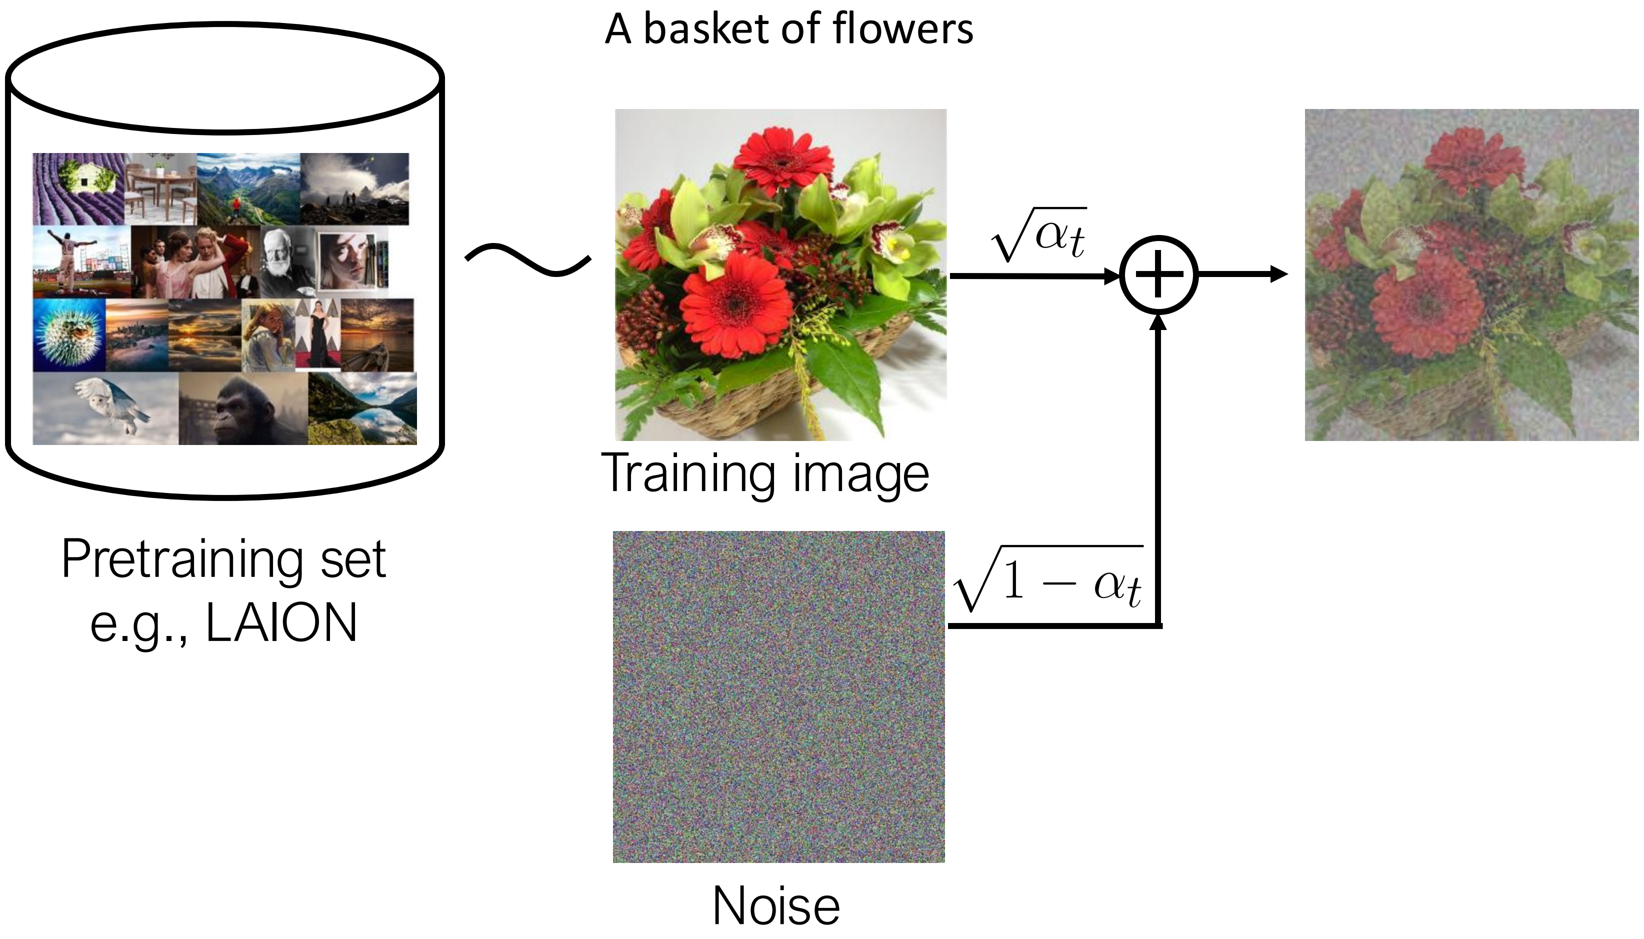

Diffusion Model Training

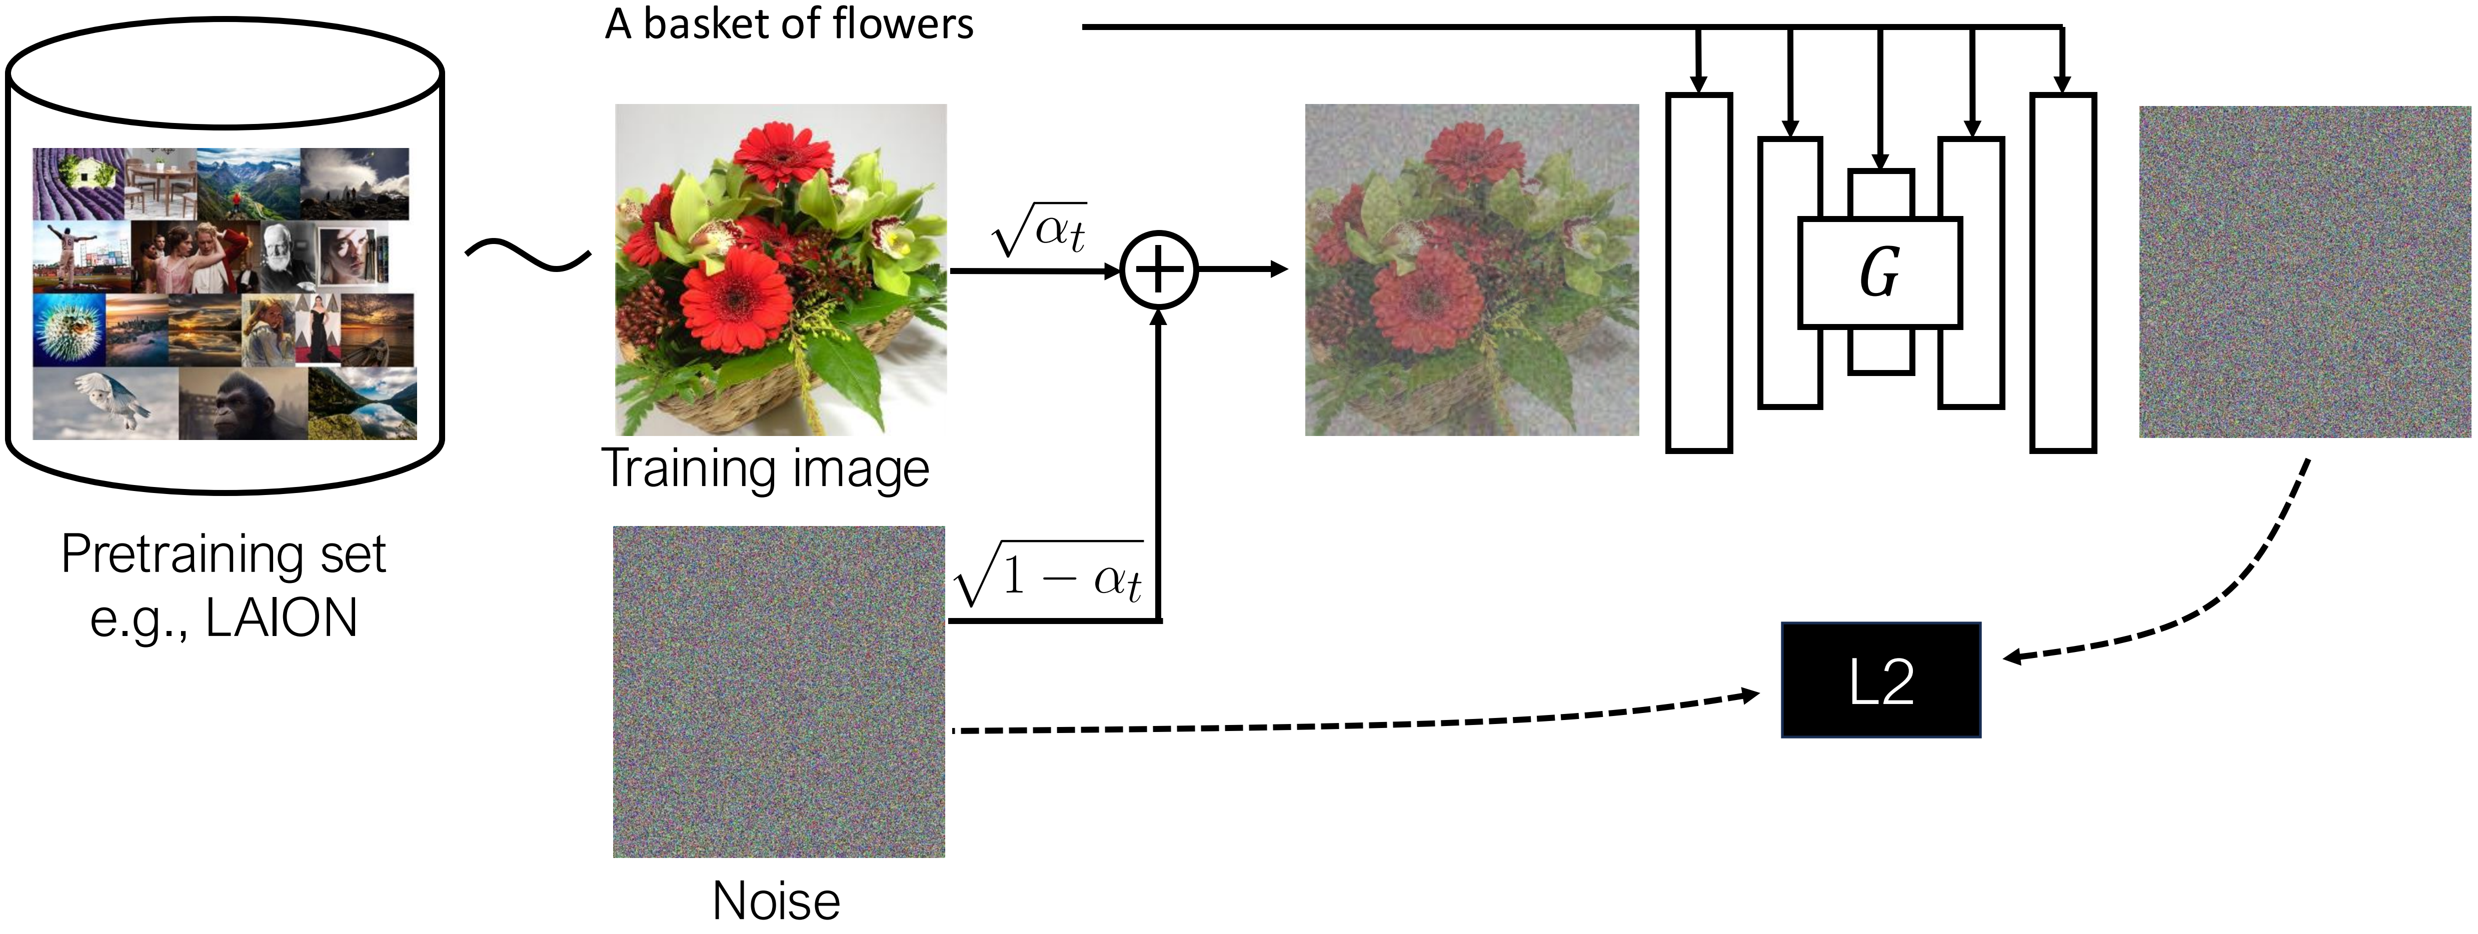

*slides motivated from https://cvpr2022-tutorial-diffusion-models.github.io

Solution I: Remove + Retraining

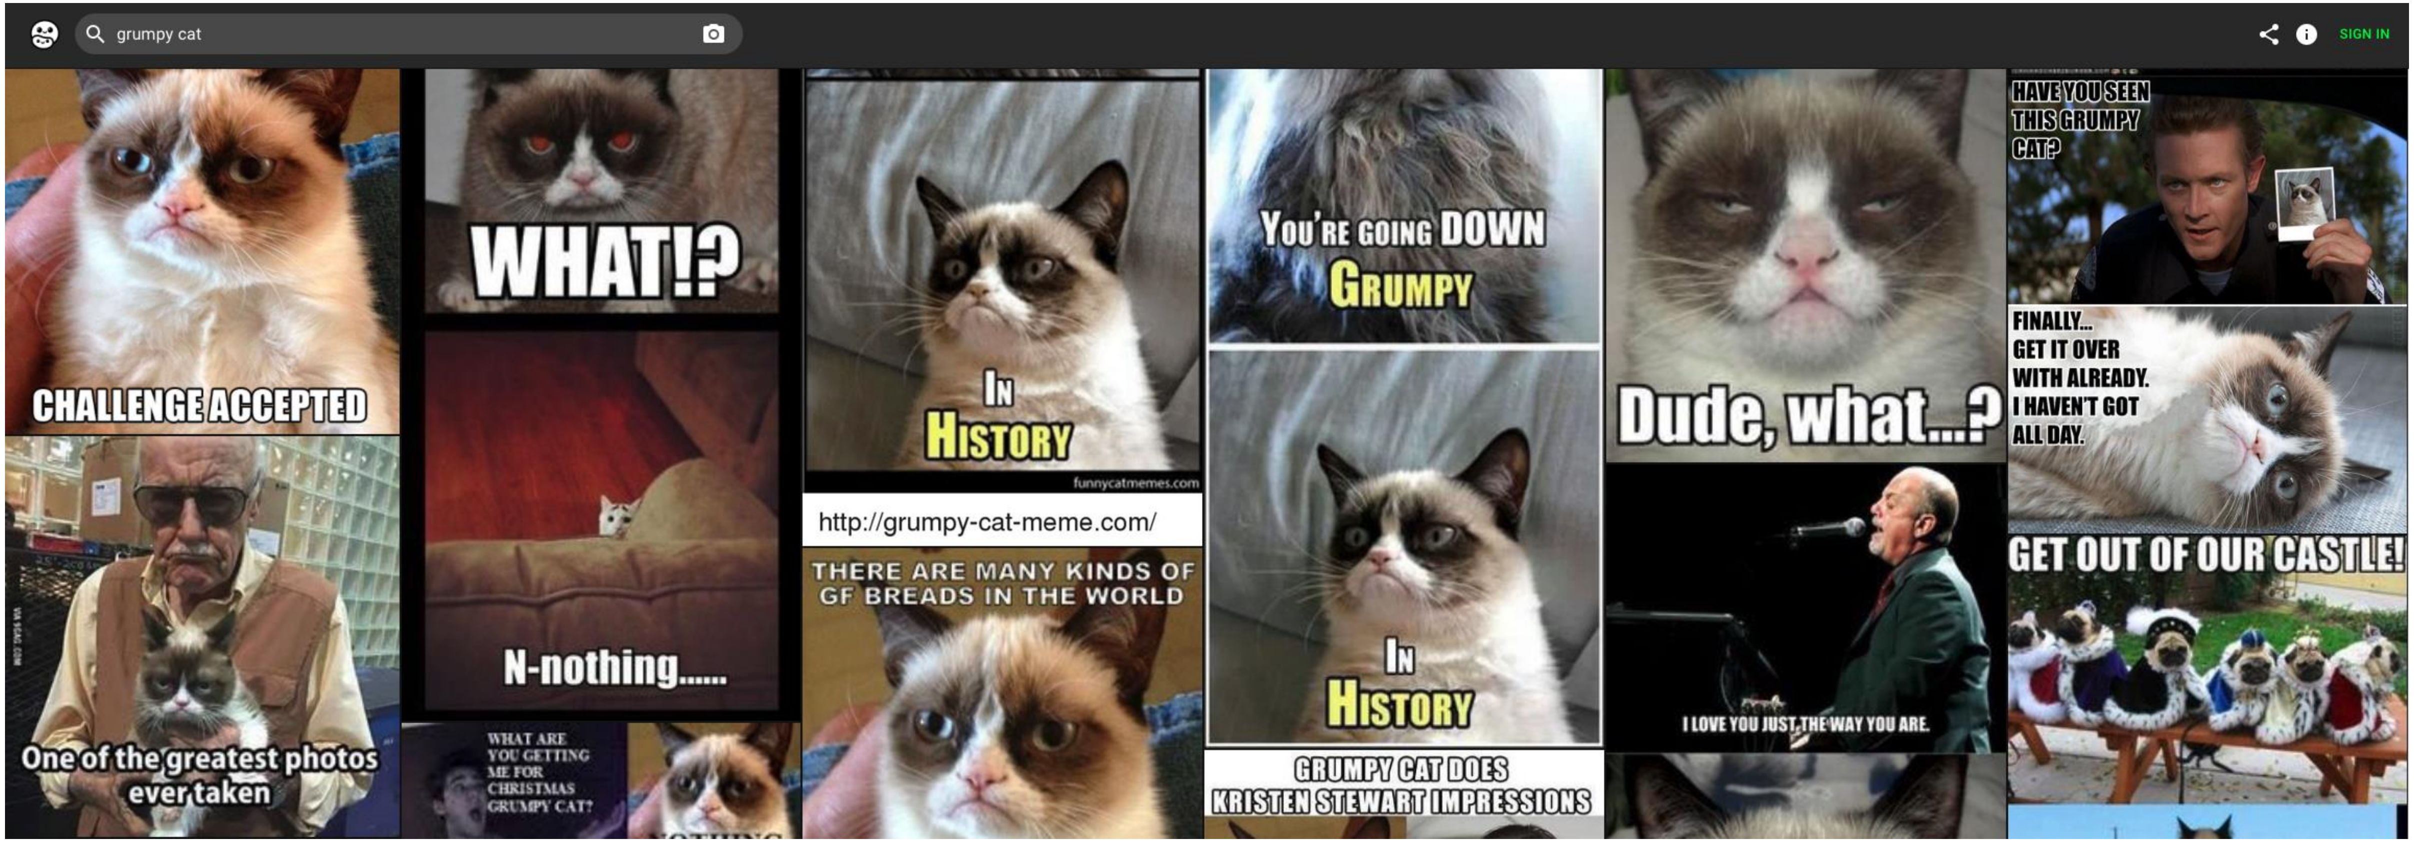

Time-consuming and Computationally-expensive

https://haveibeentrained.com/?search_text=grumpy%20cat

Solution II: Maximize Loss

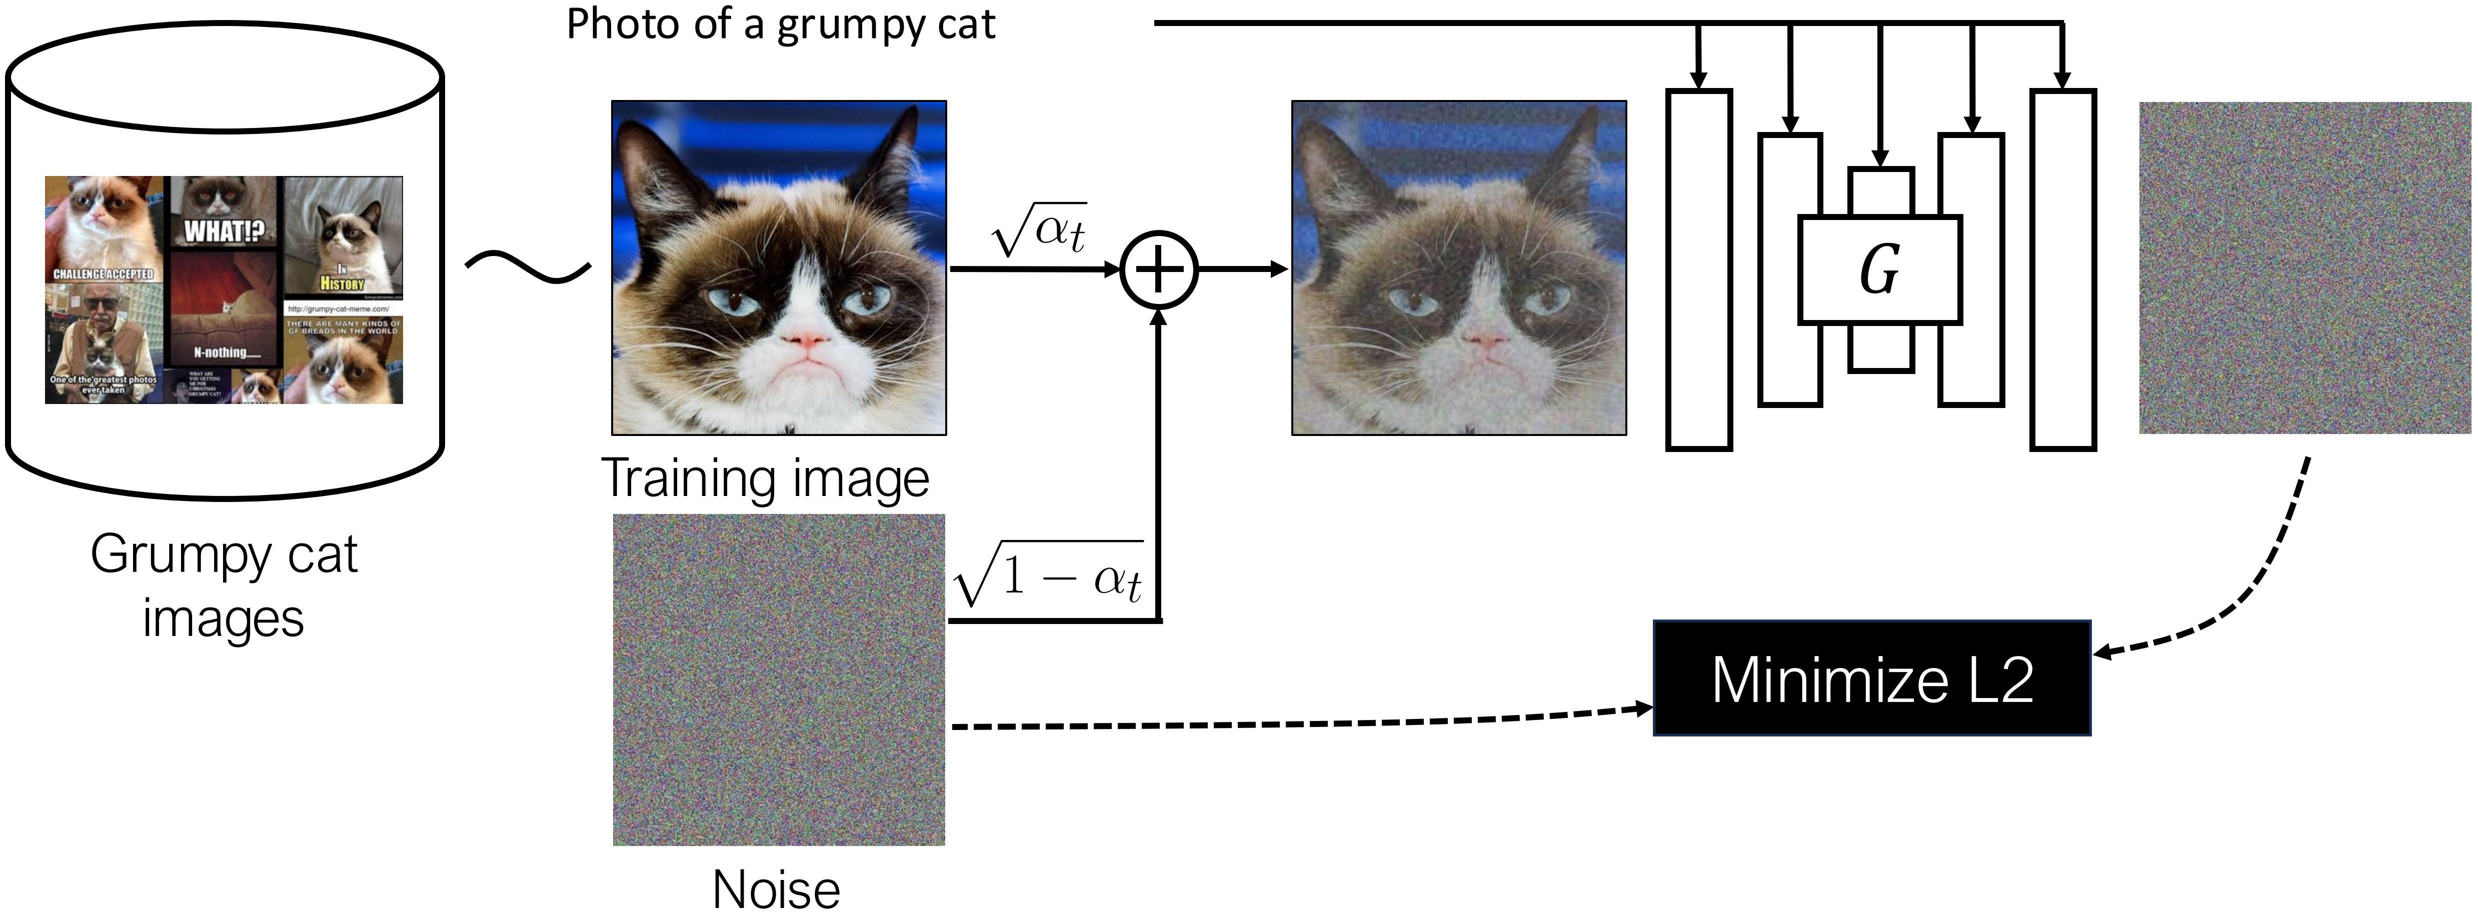

c.f. Bourtoule et al., Machine Unlearning. 2019

Photo of a grumpy cat

Target concept

Photo of a british shorthair cat

Nearby concept

Challenges

- Data opt-out and compensation are standard practices for content creation platforms.

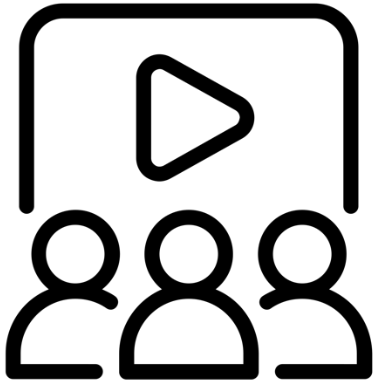

Challenges

- Data opt-out and compensation are standard practices for content creation platforms.

Challenges

- **Difficult** for Generative models, as
  - Consumers see generated data rather than training data,
  - Training data are now entangled in the model weights.

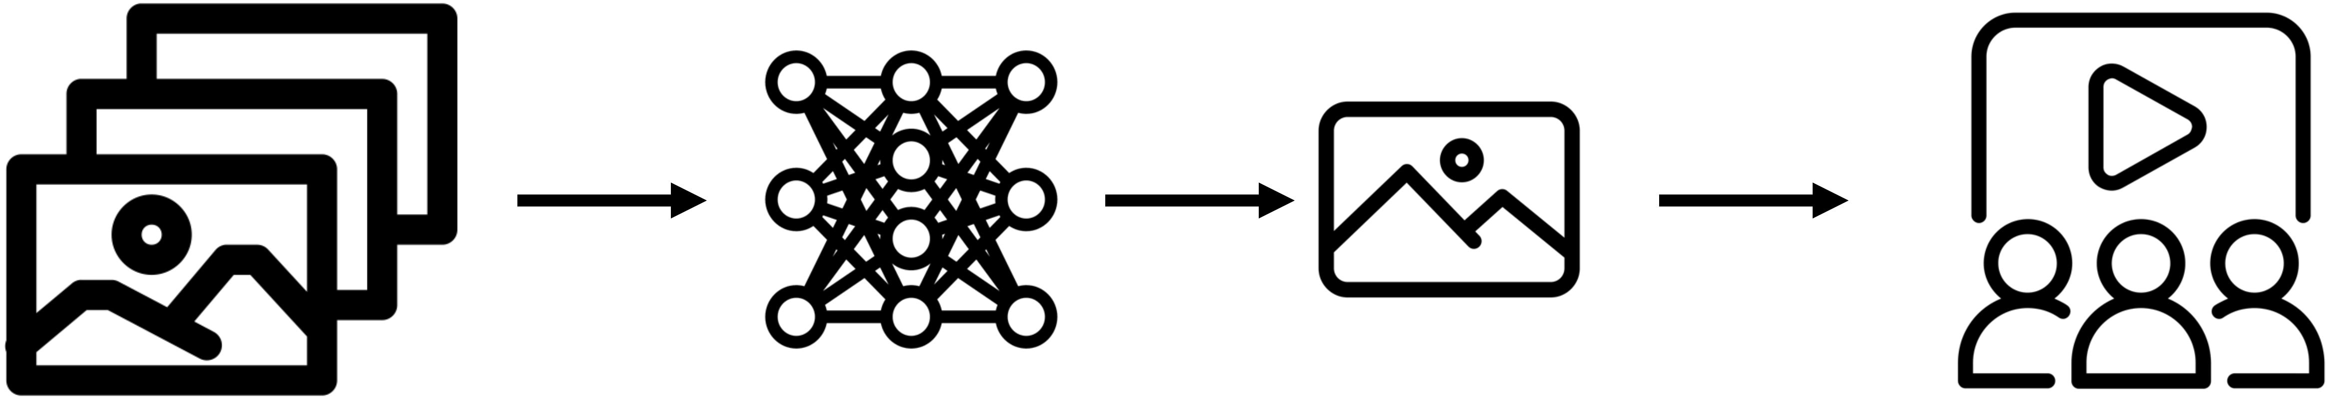

- **Our idea: only change one thing at a time.**

Our Solution: Distribution Matching

Concept Ablation [Kumari et al., ICCV 2023]

Our Solution: Distribution Matching

“Photo of a grumpy cat”

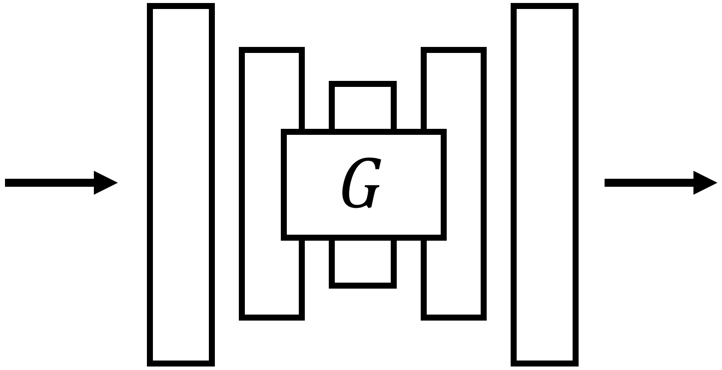

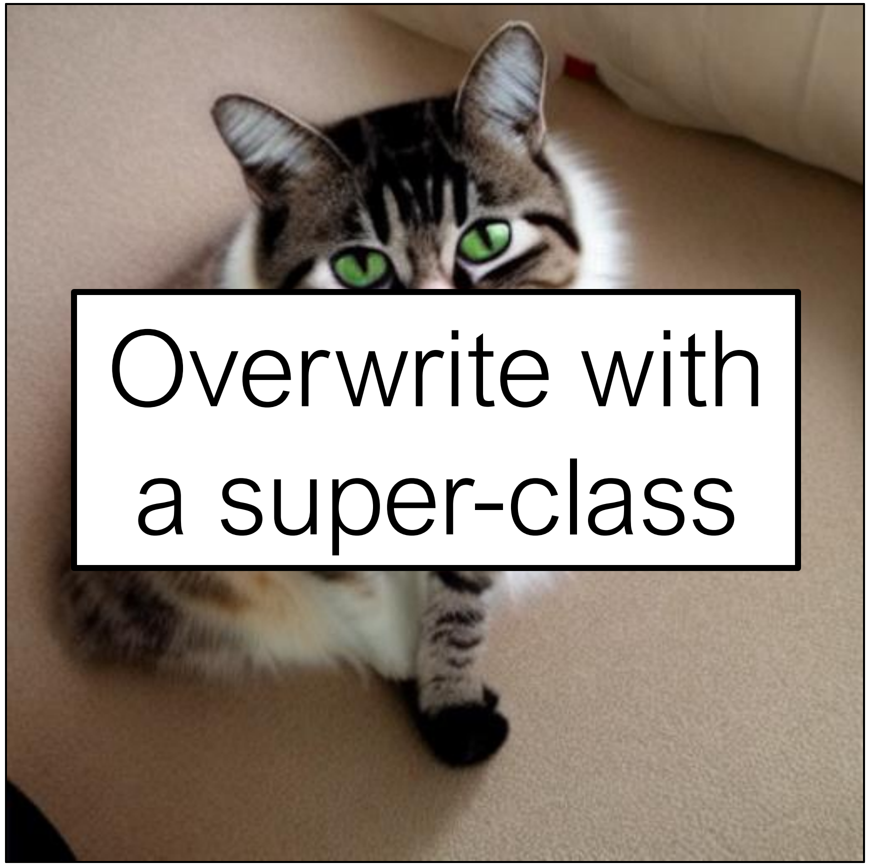

Concept Ablation [Kumari et al., ICCV 2023]

Our Solution: Distribution Matching

$p(x \mid c^*$: grumpy cat$)$

Concept Ablation [Kumari et al., ICCV 2023]

Our Solution: Distribution Matching

Concept Ablation [Kumari et al., ICCV 2023]

Our Solution: Distribution Matching

$$\arg\min_{\hat{\Phi}} \mathcal{D}_{\mathcal{KL}}(p_{\Phi}(\mathbf{x}_{(0...T)}|\mathbf{c})||p_{\hat{\Phi}}(\mathbf{x}_{(0...T)}|\mathbf{c}^*))$$

Φ: pretrained model

$\hat{\Phi}$: fine-tuned model

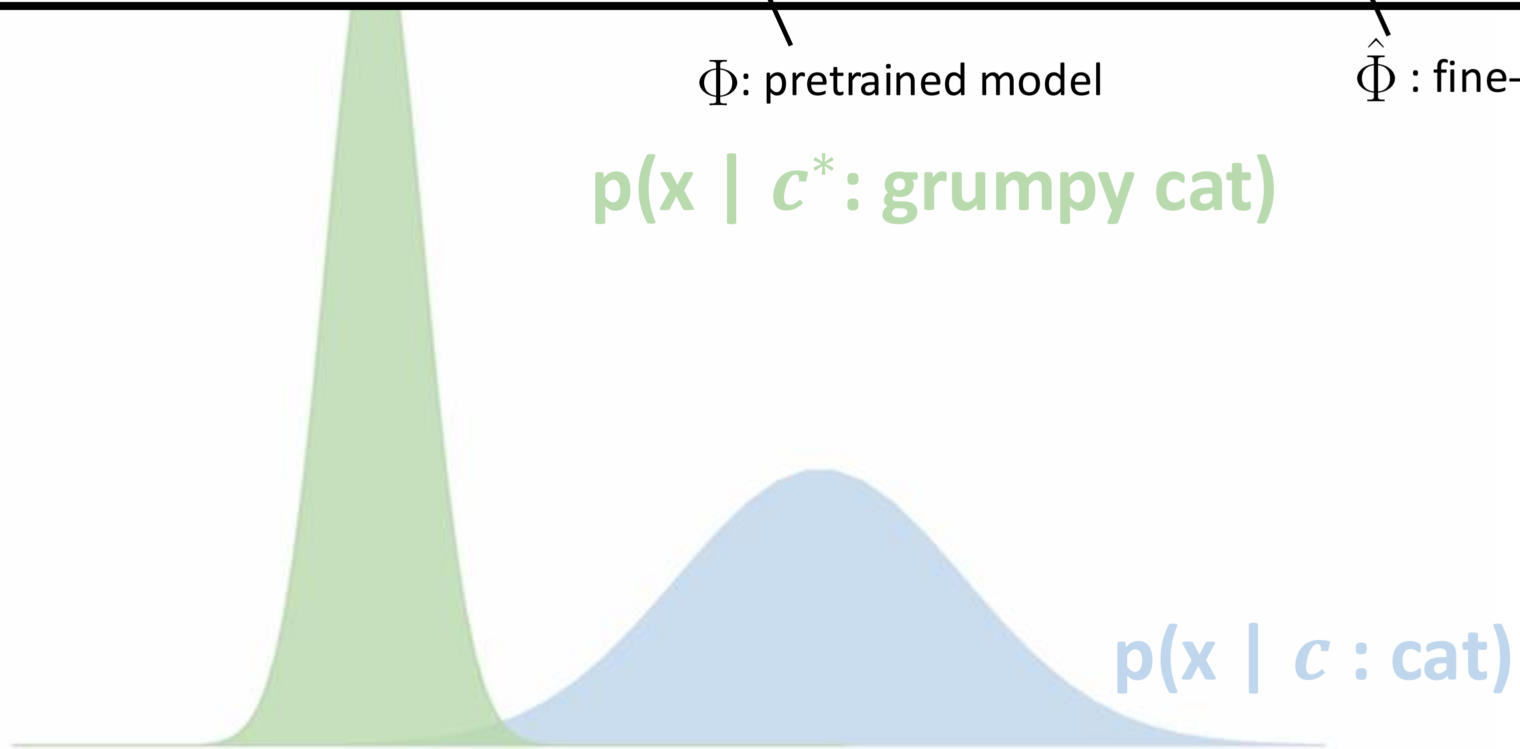

Concept Ablation [Kumari et al., ICCV 2023]

Concept Ablation Objective Function

$$\mathcal{D}_{\mathcal{KL}}(p_{\Phi}(\mathbf{x}_{(0 \ldots T)}|\mathbf{c})||p_{\hat{\Phi}}(\mathbf{x}_{(0 \ldots T)}|\mathbf{c}^*))$$

pretrained model distribution given cat caption

Fine-tuned model distribution given grumpy cat caption

Simplifying to per timestep distribution of the diffusion model

$$= \sum_{t=1}^{T} \mathbb{E}\left[\mathcal{D}_{\mathcal{KL}}(p_{\Phi}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{c})||p_{\hat{\Phi}}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{c}^*))\right]$$

Cat

Grumpy Cat

$$\mathbf{x}_t \sim p_{\Phi}(\mathbf{x}_t|\mathbf{c})$$

KL Divergence between two Normal distribution

Can be simplified to l2 distance between mean of two distribution

Concept Ablation [Kumari et al., ICCV 2023]

Concept Ablation Objective Function

pretrained model's prediction
given cat caption

fine-tuned model's prediction
given grumpy cat caption

$$\mathcal{L} = \mathbb{E}_{\mathbf{x}_t} || \Phi(\mathbf{x}_t, \mathbf{c}, t) - \hat{\Phi}(\mathbf{x}_t, \mathbf{c}^*, t) ||$$

Concept Ablation [Kumari et al., ICCV 2023]

Concept Ablation Objective Function

pretrained model

Memory intensive in practice. So, we use stop-grad with the existing model.

$$\mathcal{L} = \mathbb{E}_{\mathbf{x}_t} ||\Phi(\mathbf{x}_t, \mathbf{c}, t) - \hat{\Phi}(\mathbf{x}_t, \mathbf{c}^*, t)||$$

Concept Ablation [Kumari et al., ICCV 2023]

Concept Ablation Objective Function

$$\mathcal{L} = \mathbb{E}_{\mathbf{x}_t} || \hat{\Phi}(\mathbf{x}_t, \mathbf{c}, t).\mathrm{sg}() - \hat{\Phi}(\mathbf{x}_t, \mathbf{c}^*, t) ||$$

$$\mathbf{x}_t \sim p_{\Phi}(\mathbf{x}_t | \mathbf{c})$$

Time consuming. Therefore, we generate images once and use forward process to approximate this.

Concept Ablation [Kumari et al., ICCV 2023]

Concept Ablation Objective Function

$$\mathcal{L} = \mathbb{E}_{\mathbf{x}_t} ||\hat{\Phi}(\mathbf{x}_t, \mathbf{c}, t).\mathrm{sg}() - \hat{\Phi}(\mathbf{x}_t, \mathbf{c}^*, t)||$$

$$\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x} + \sqrt{1 - \alpha_t}\epsilon$$

Concept Ablation [Kumari et al., ICCV 2023]

Final Method

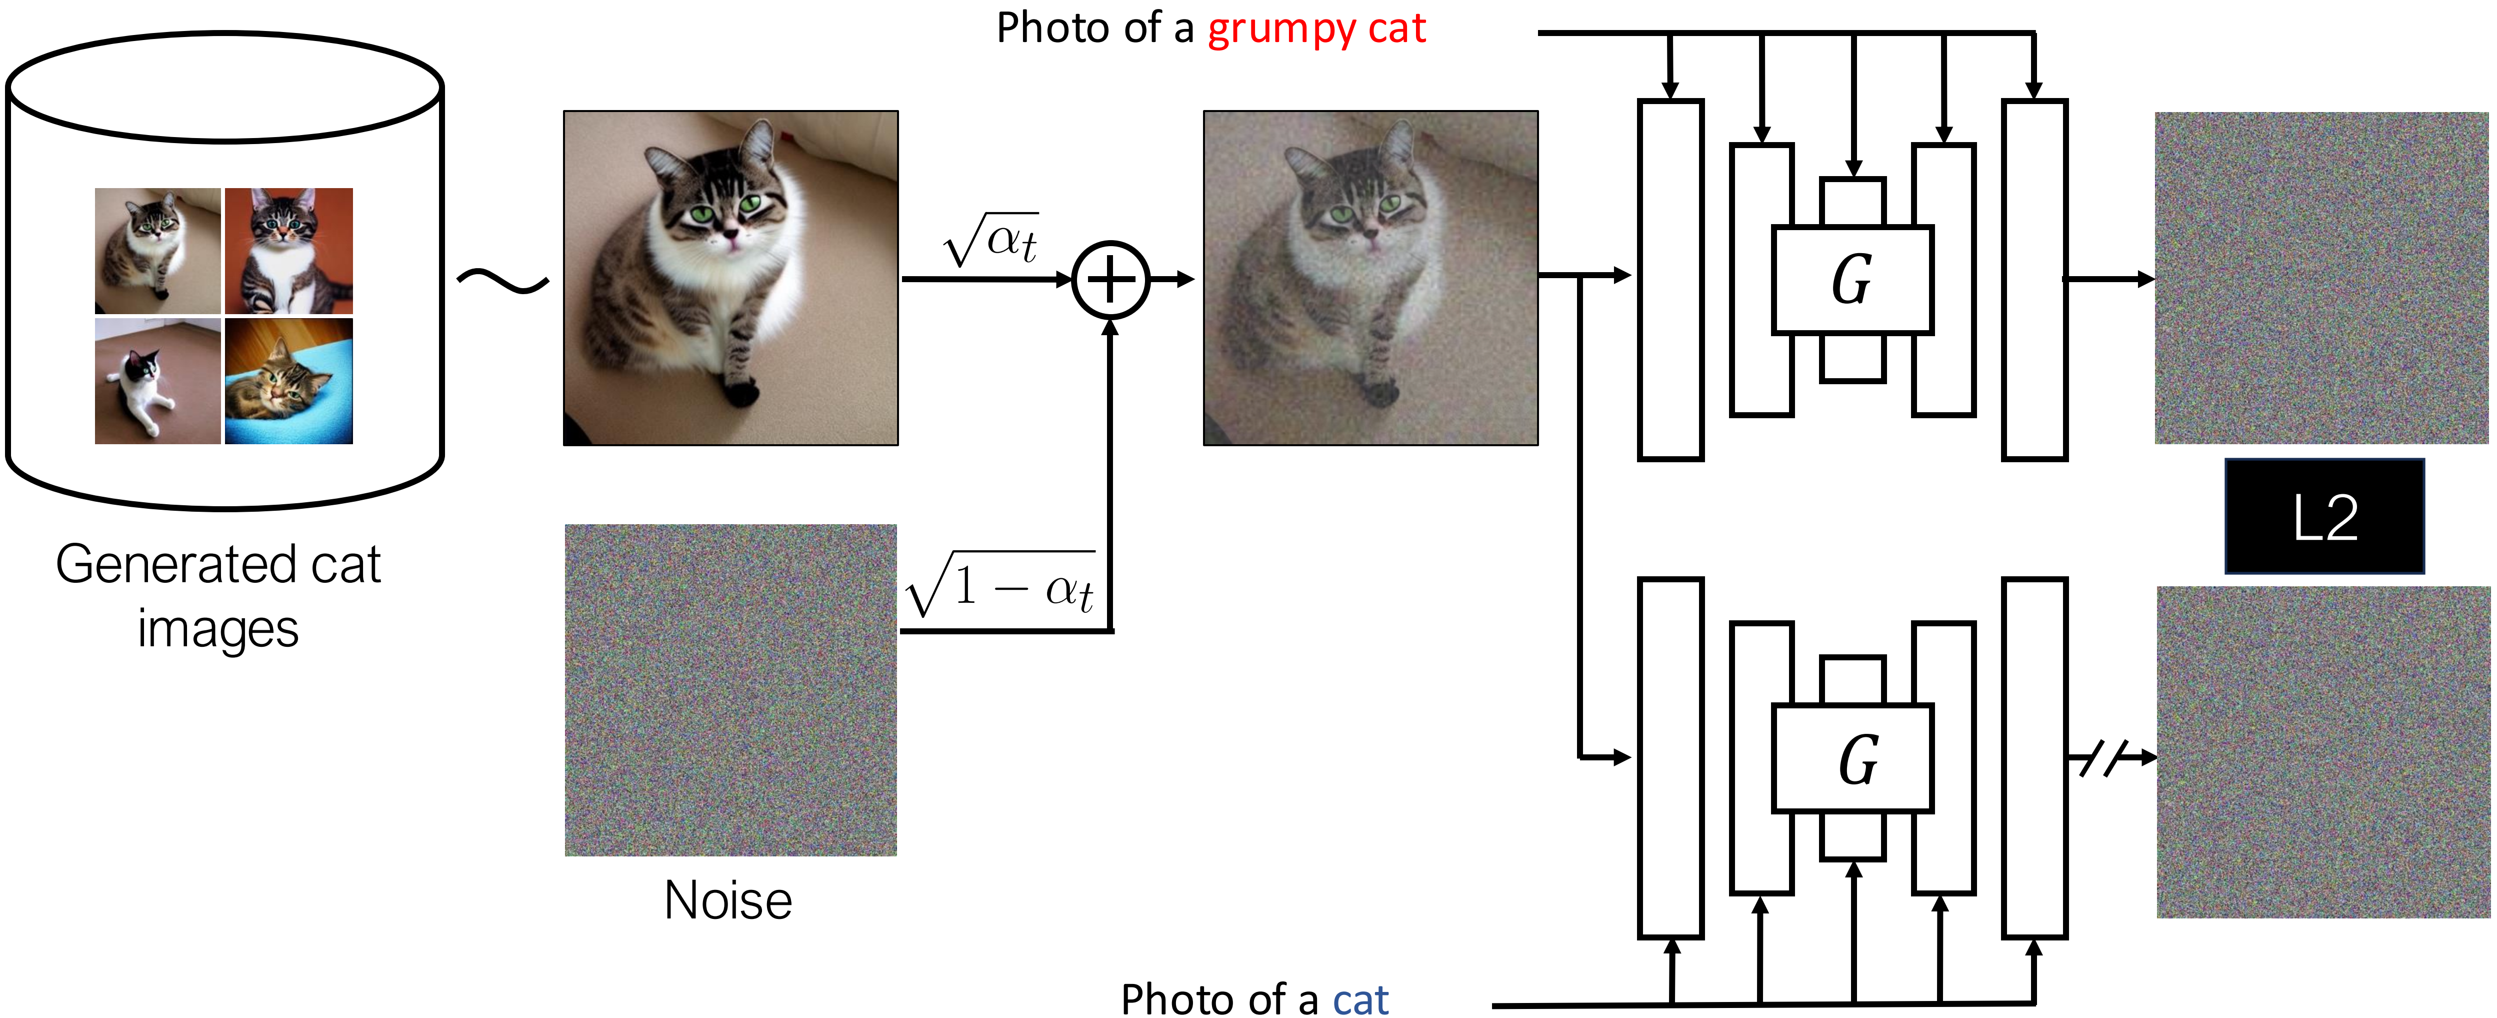

Concept Ablation [Kumari et al., ICCV 2023]

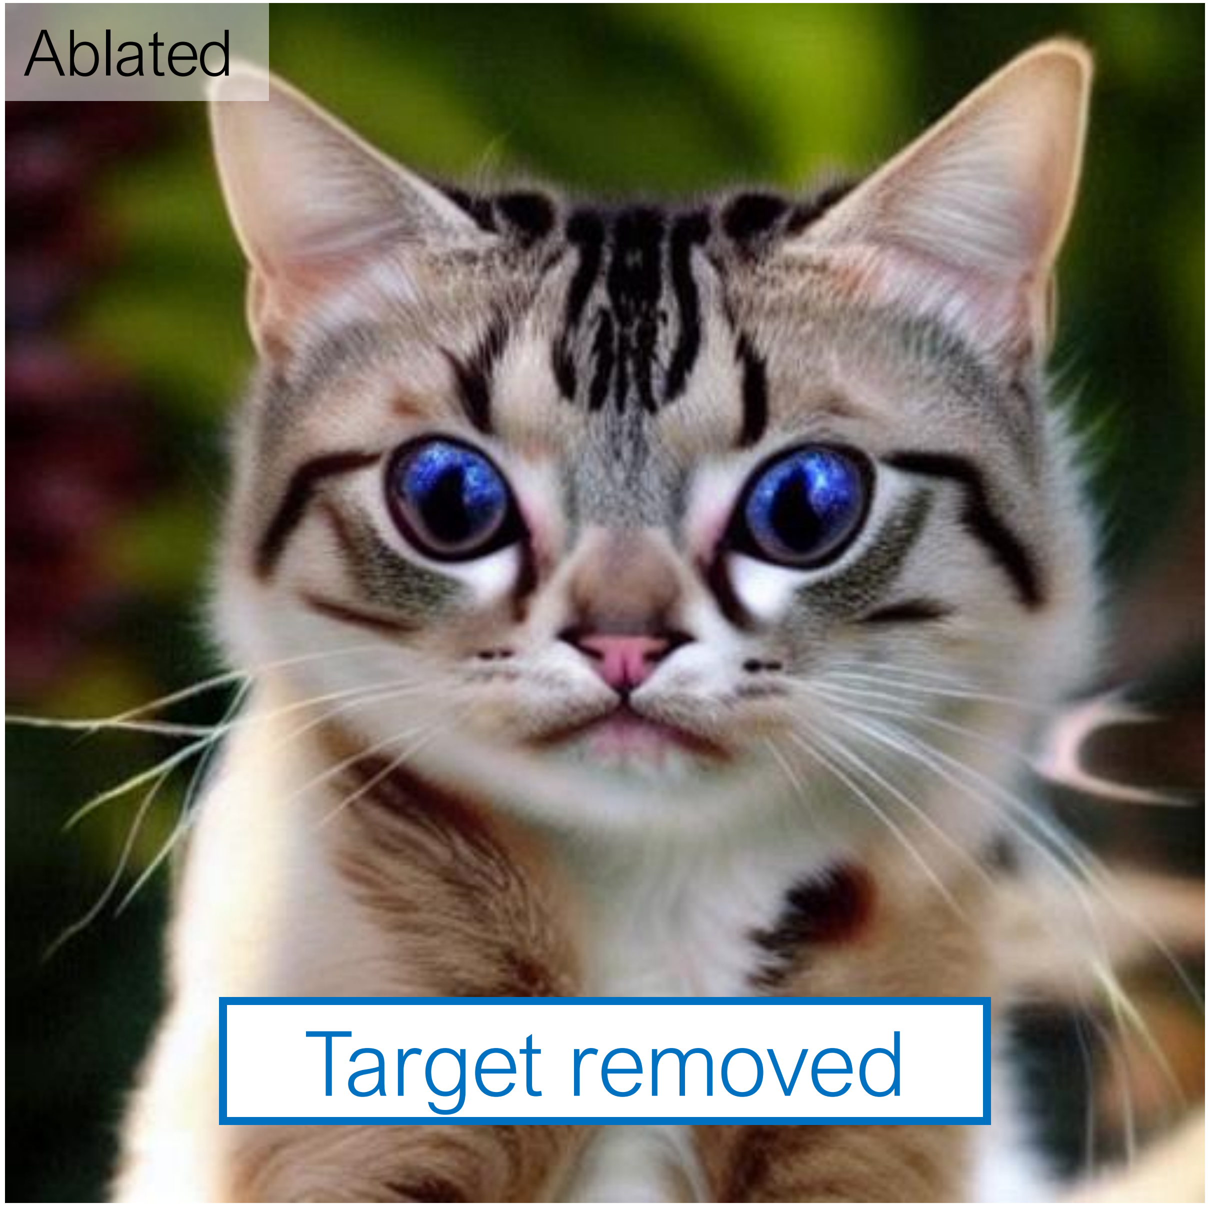

Photo of a grumpy cat

Target concept

Photo of a british shorthair cat

Nearby concept

R2D2

Nemo

Copyrighted characters

Ablating Van Gogh’s Style

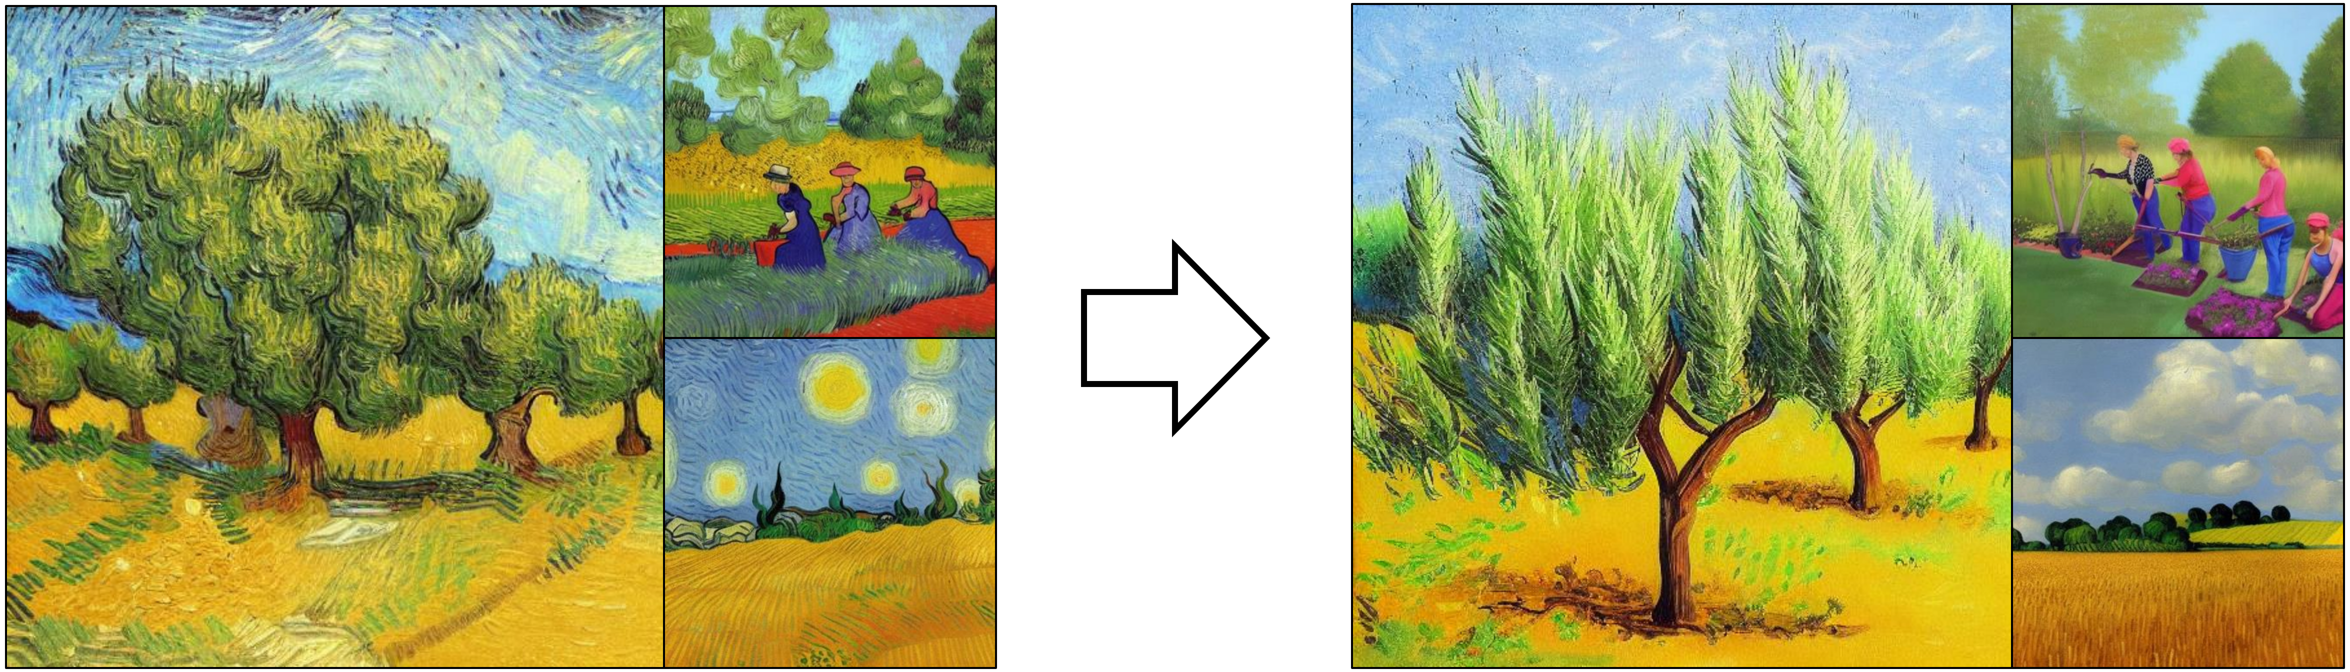

Ablating Van Gogh's Style

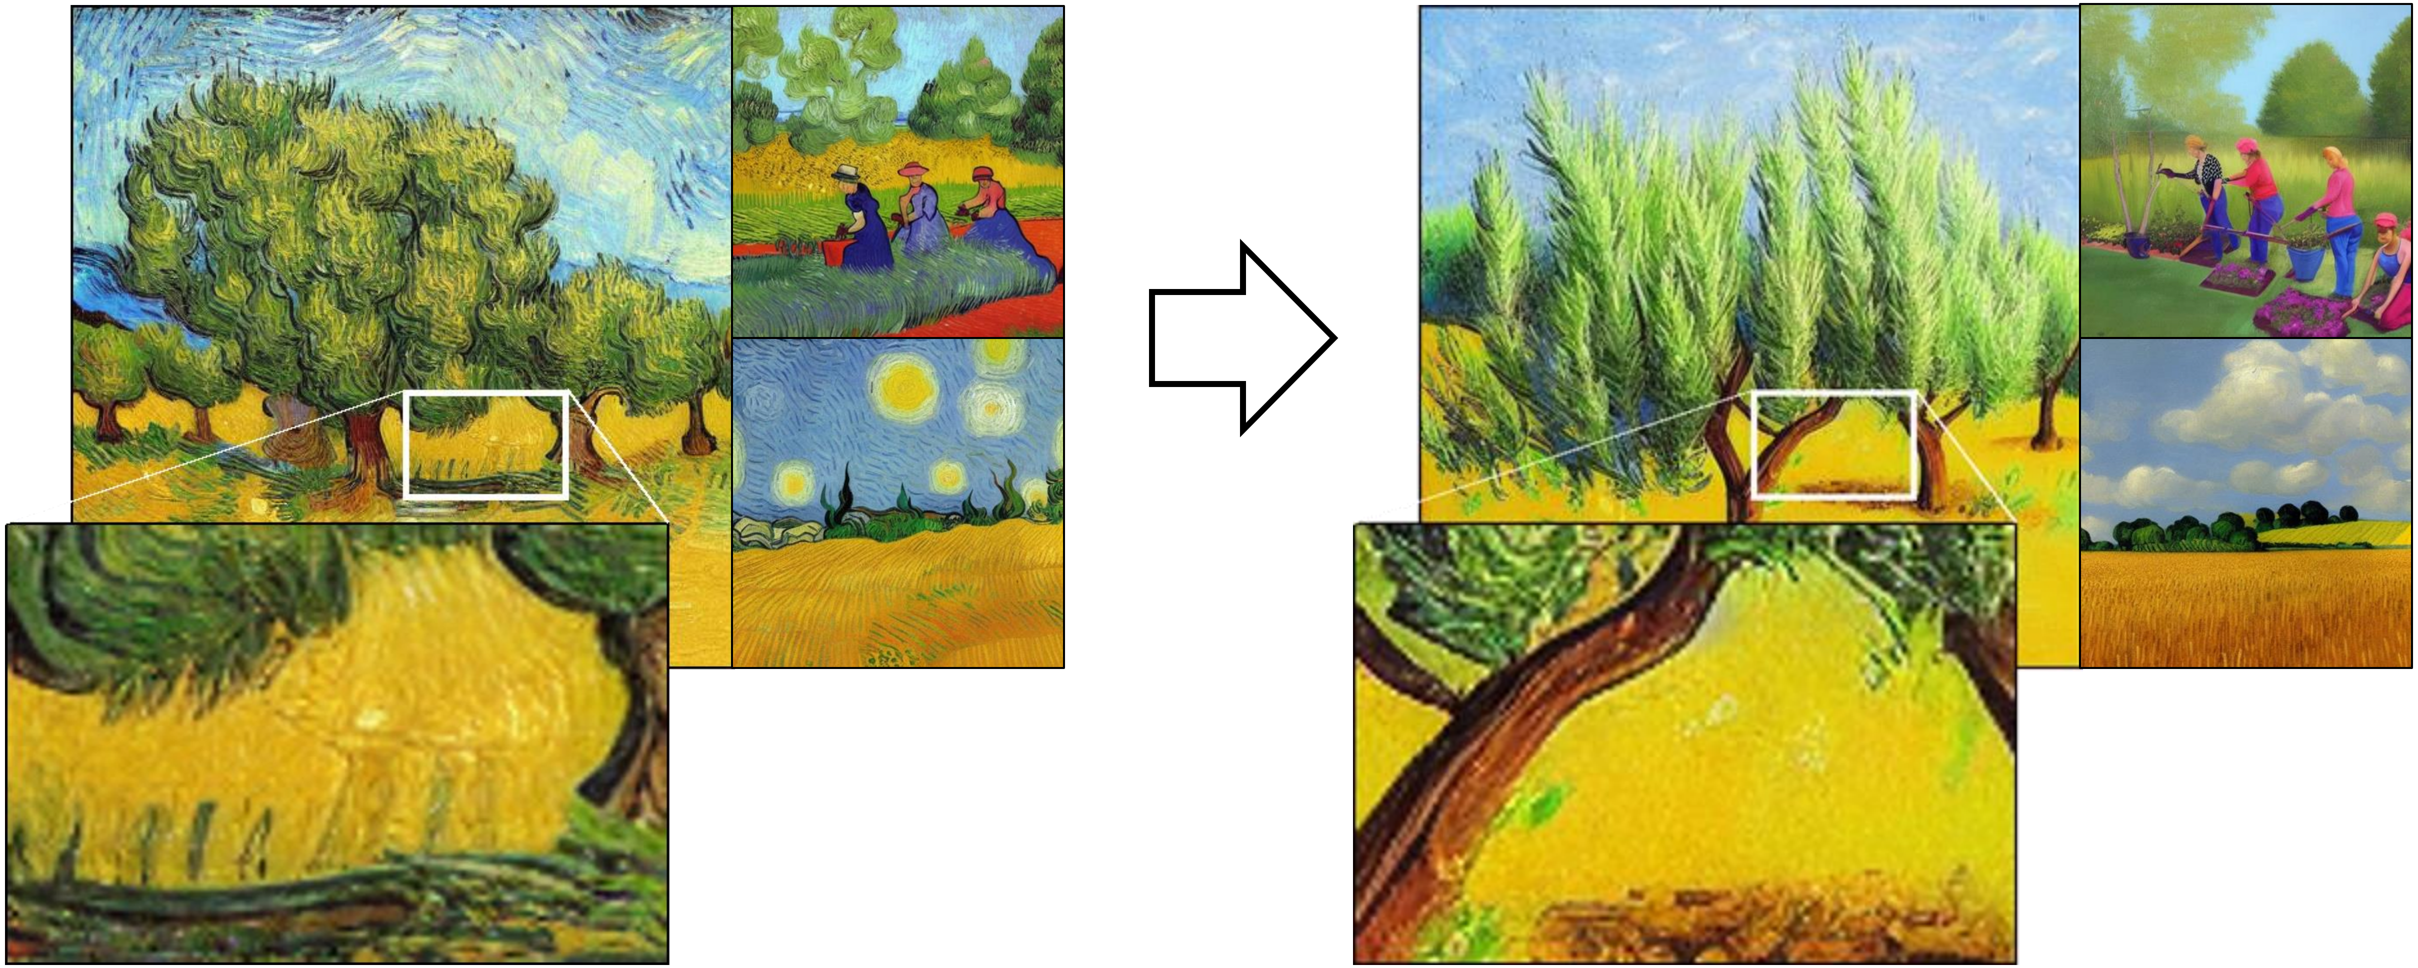

Ablating Greg Rutkowski's Style

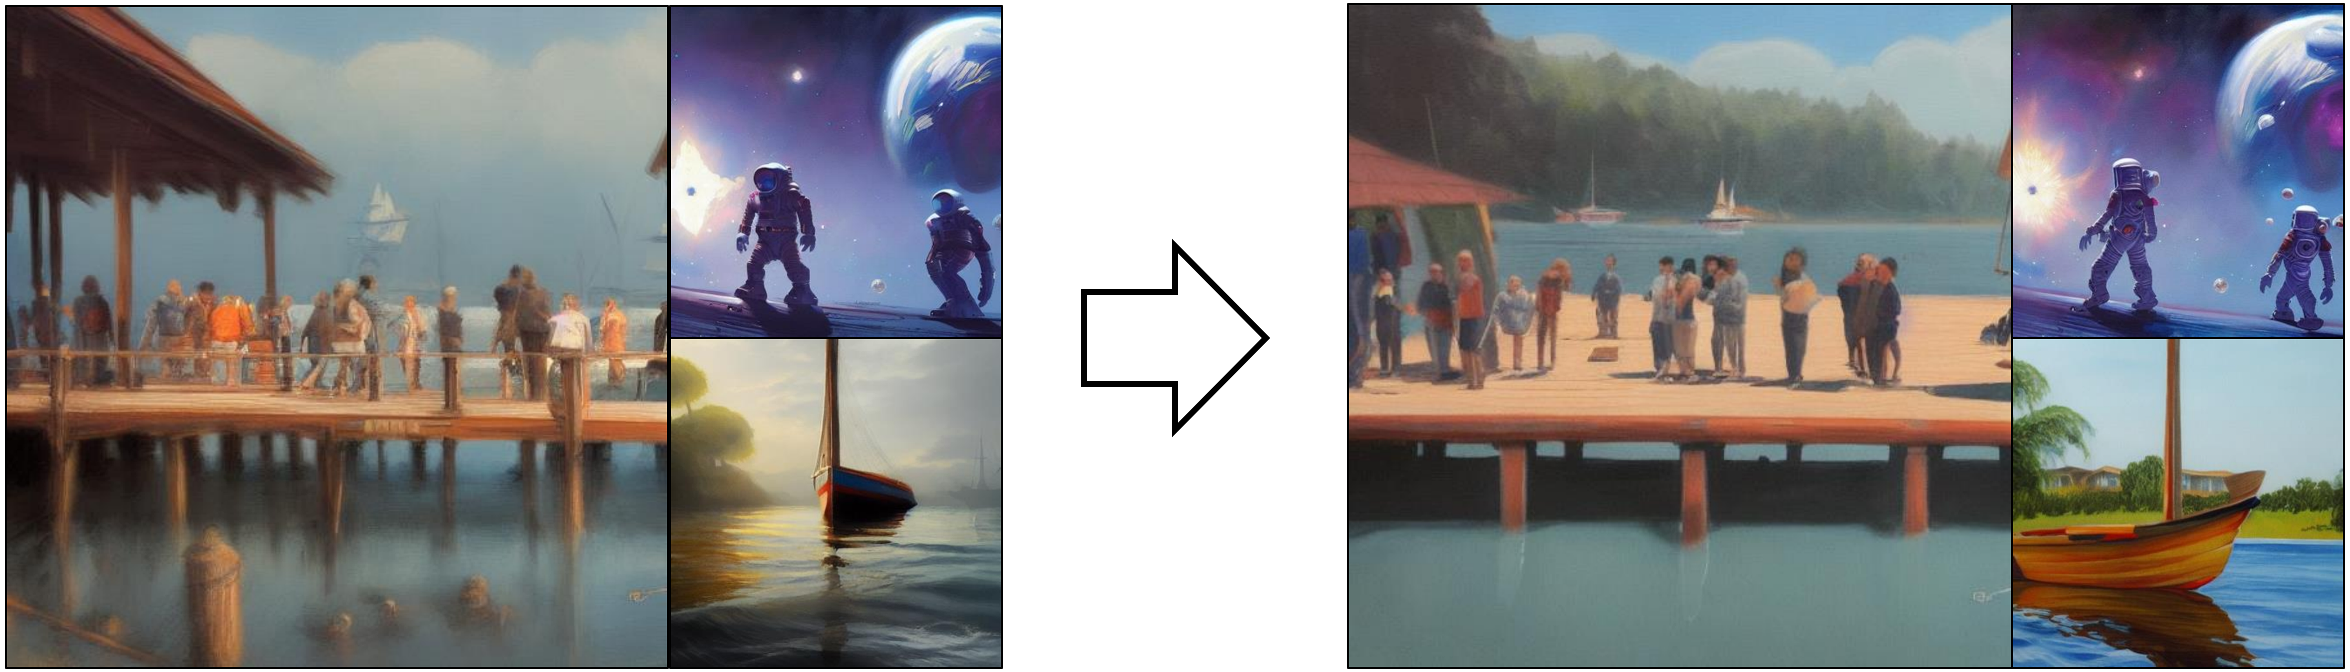

Ablating Greg Rutkowski's Style

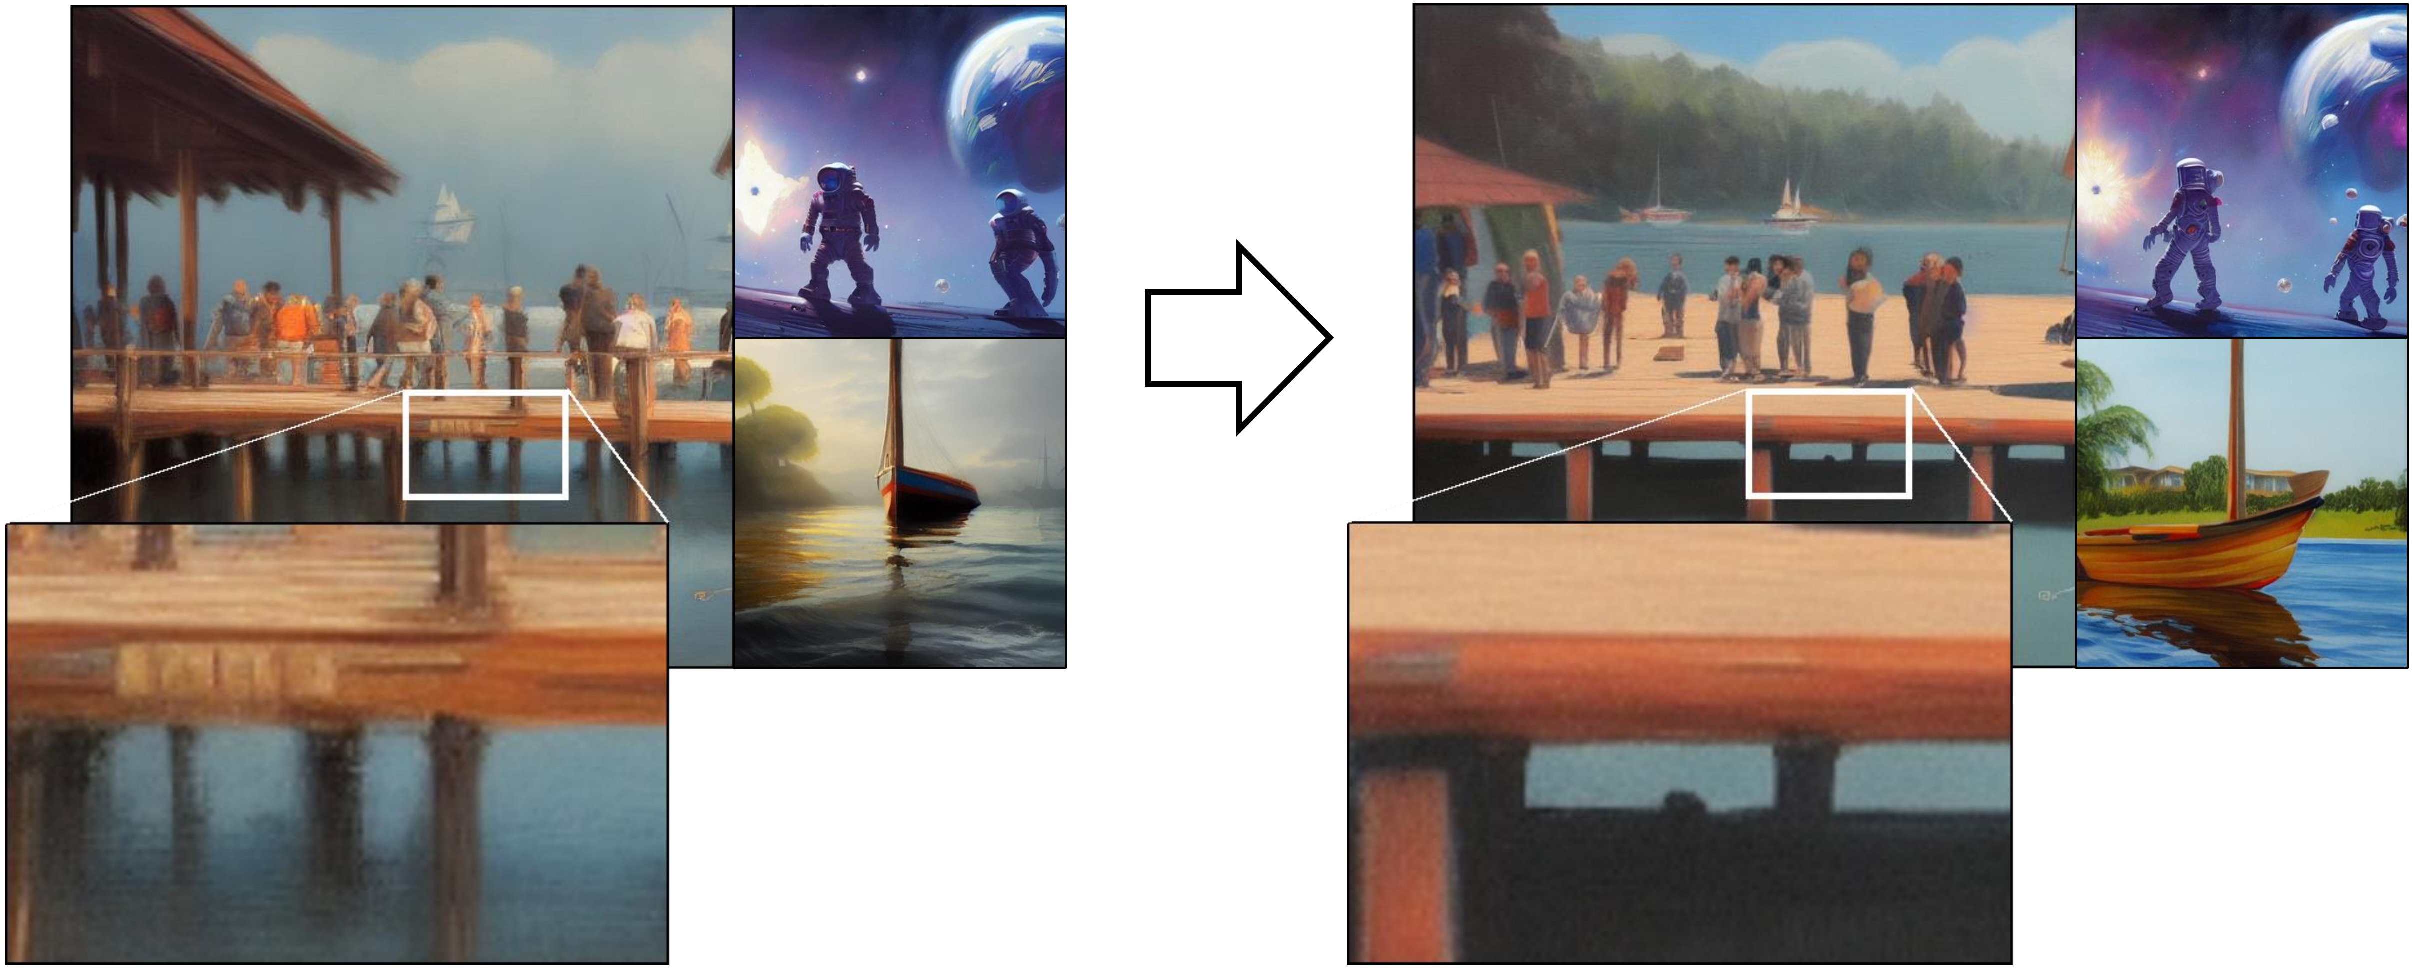

Ablating Memorized Images

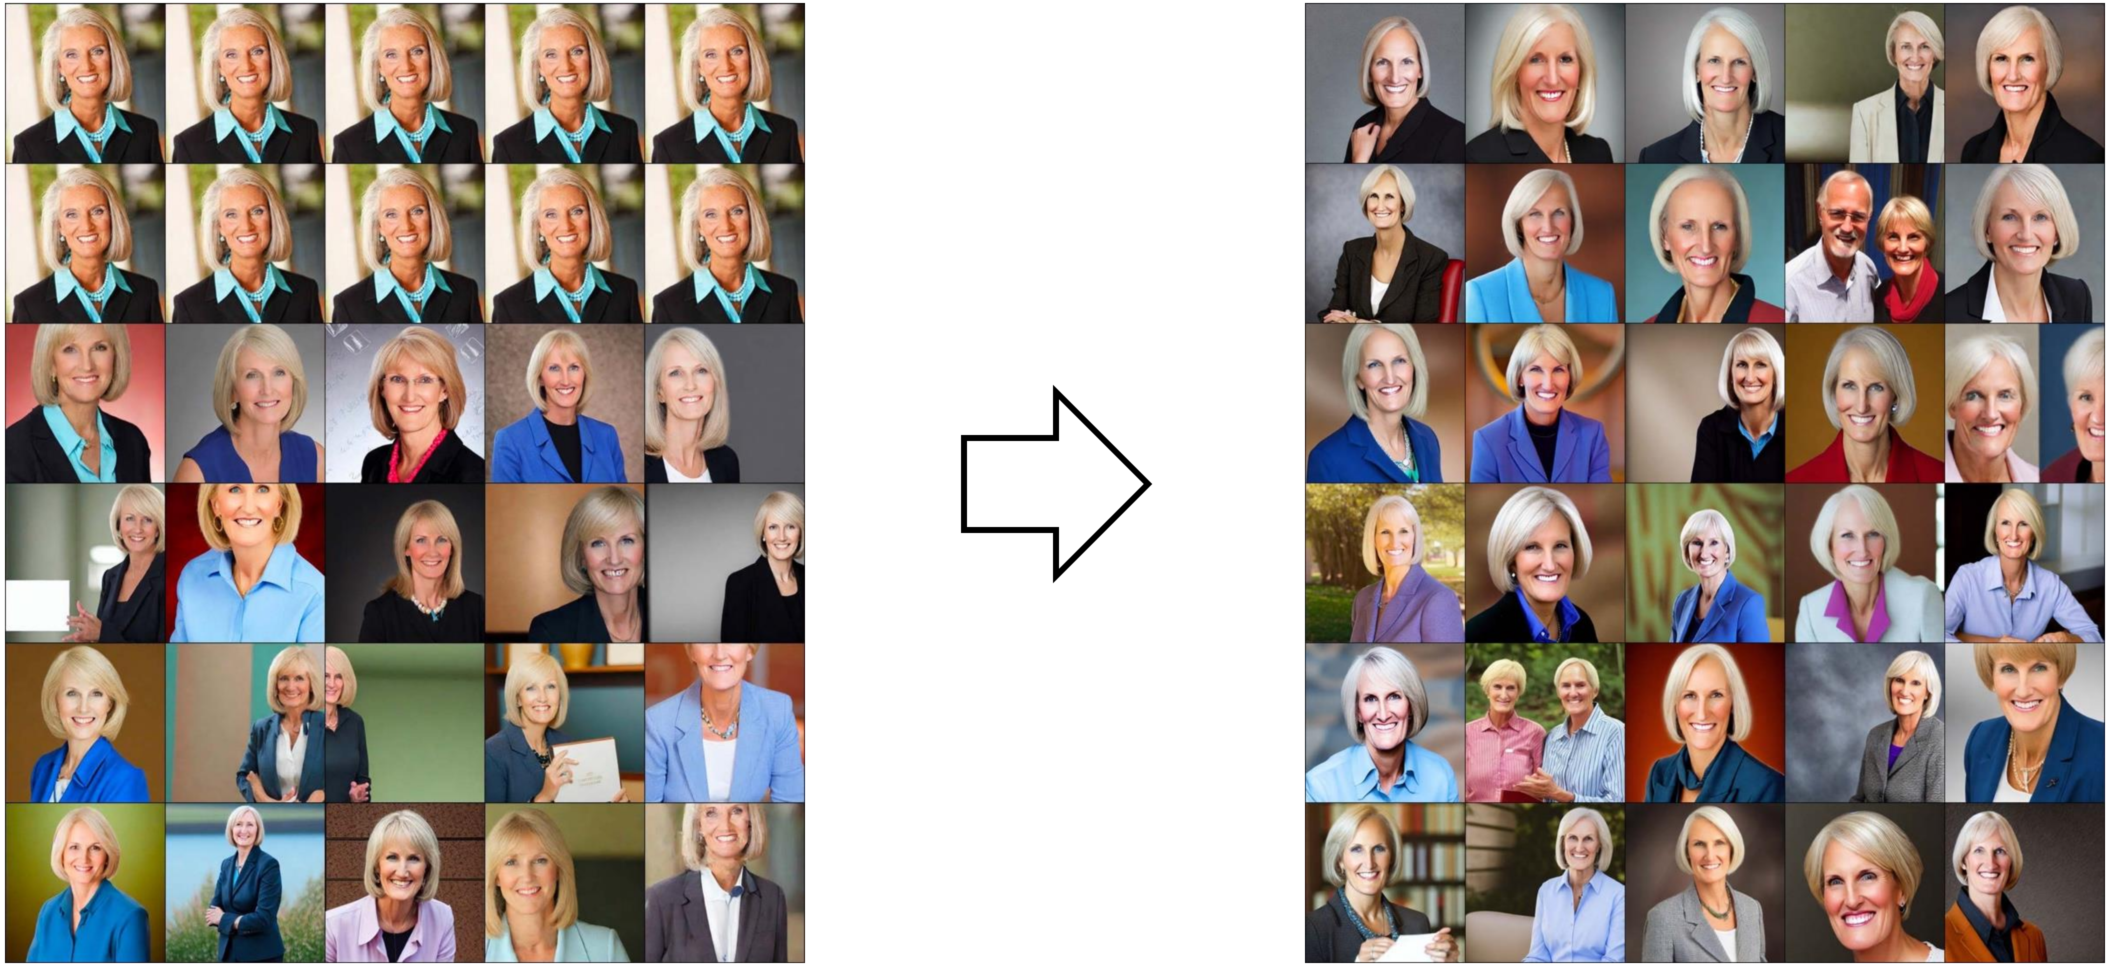

Ablating Memorized Images

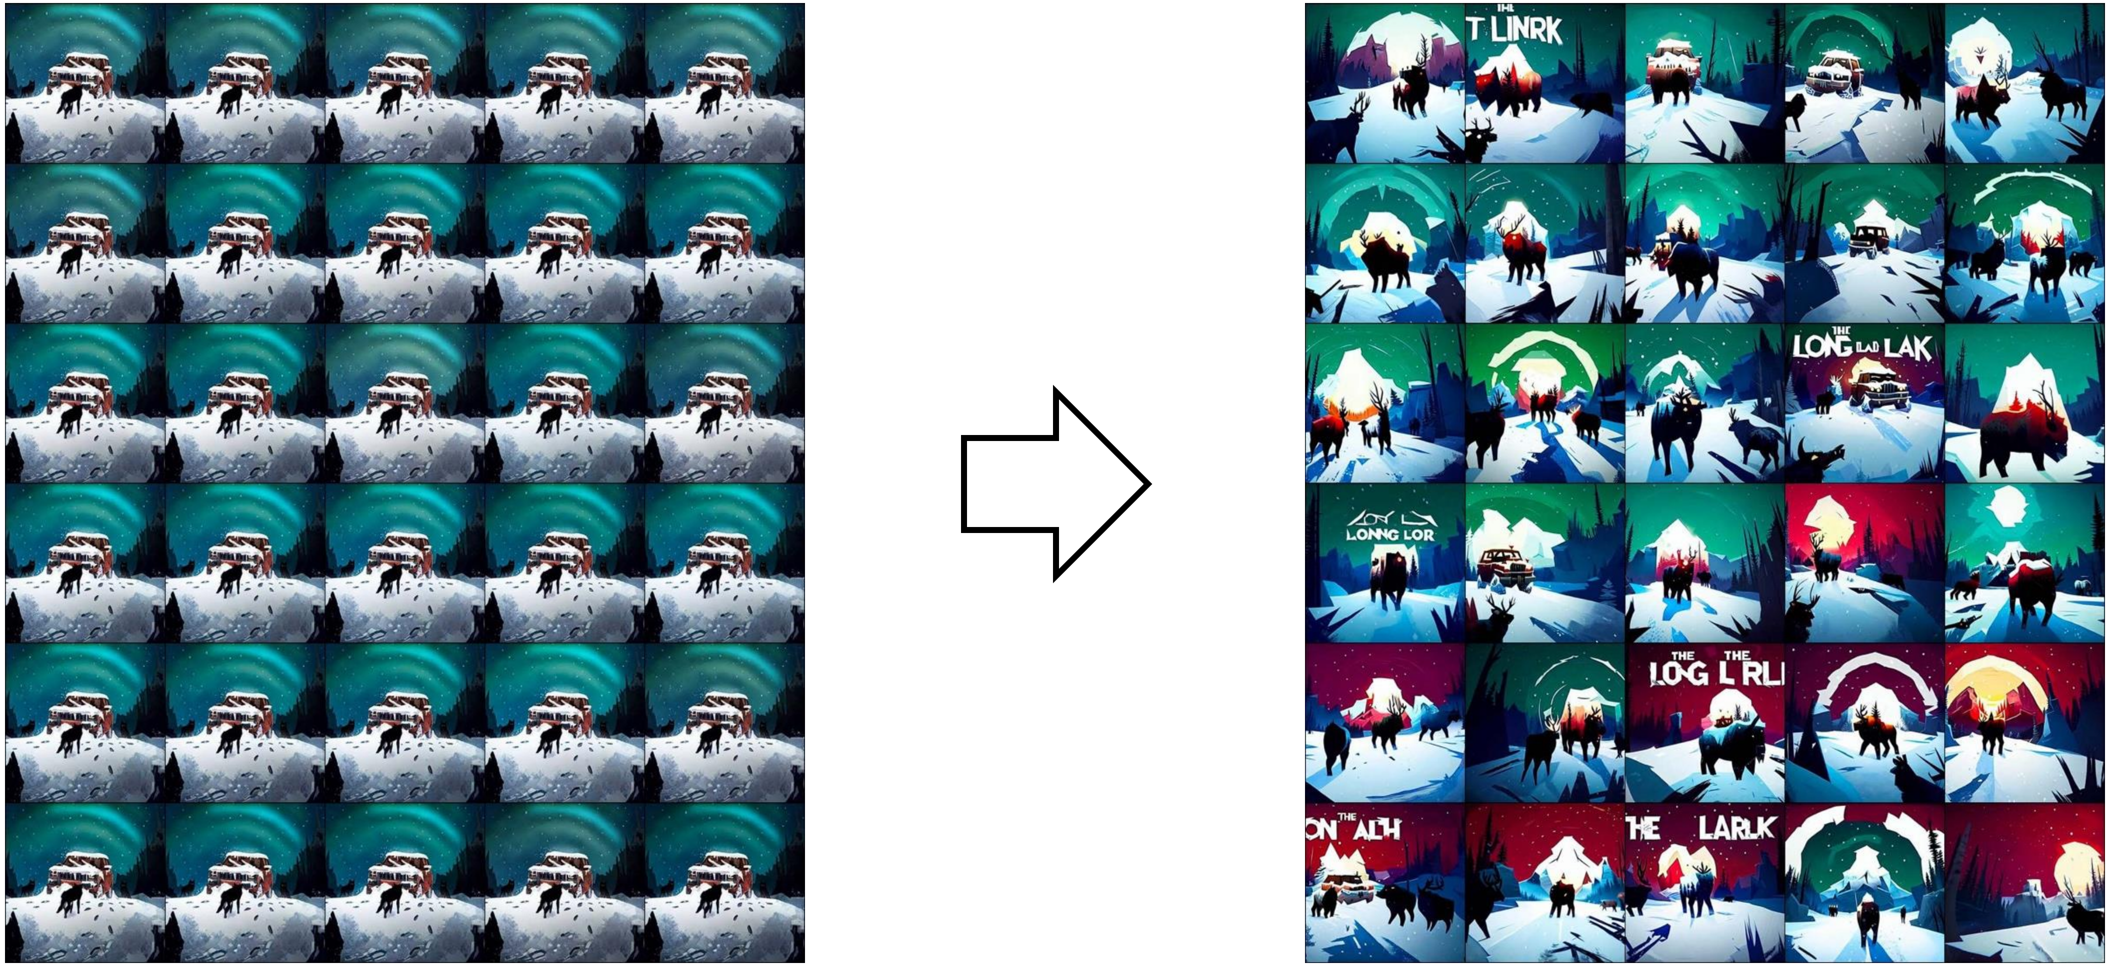

Ablating Composition “Kids with Guns”

Kids with Guns

Kids

Guns

Stable Diffusion

Ablated Stable Diffusion

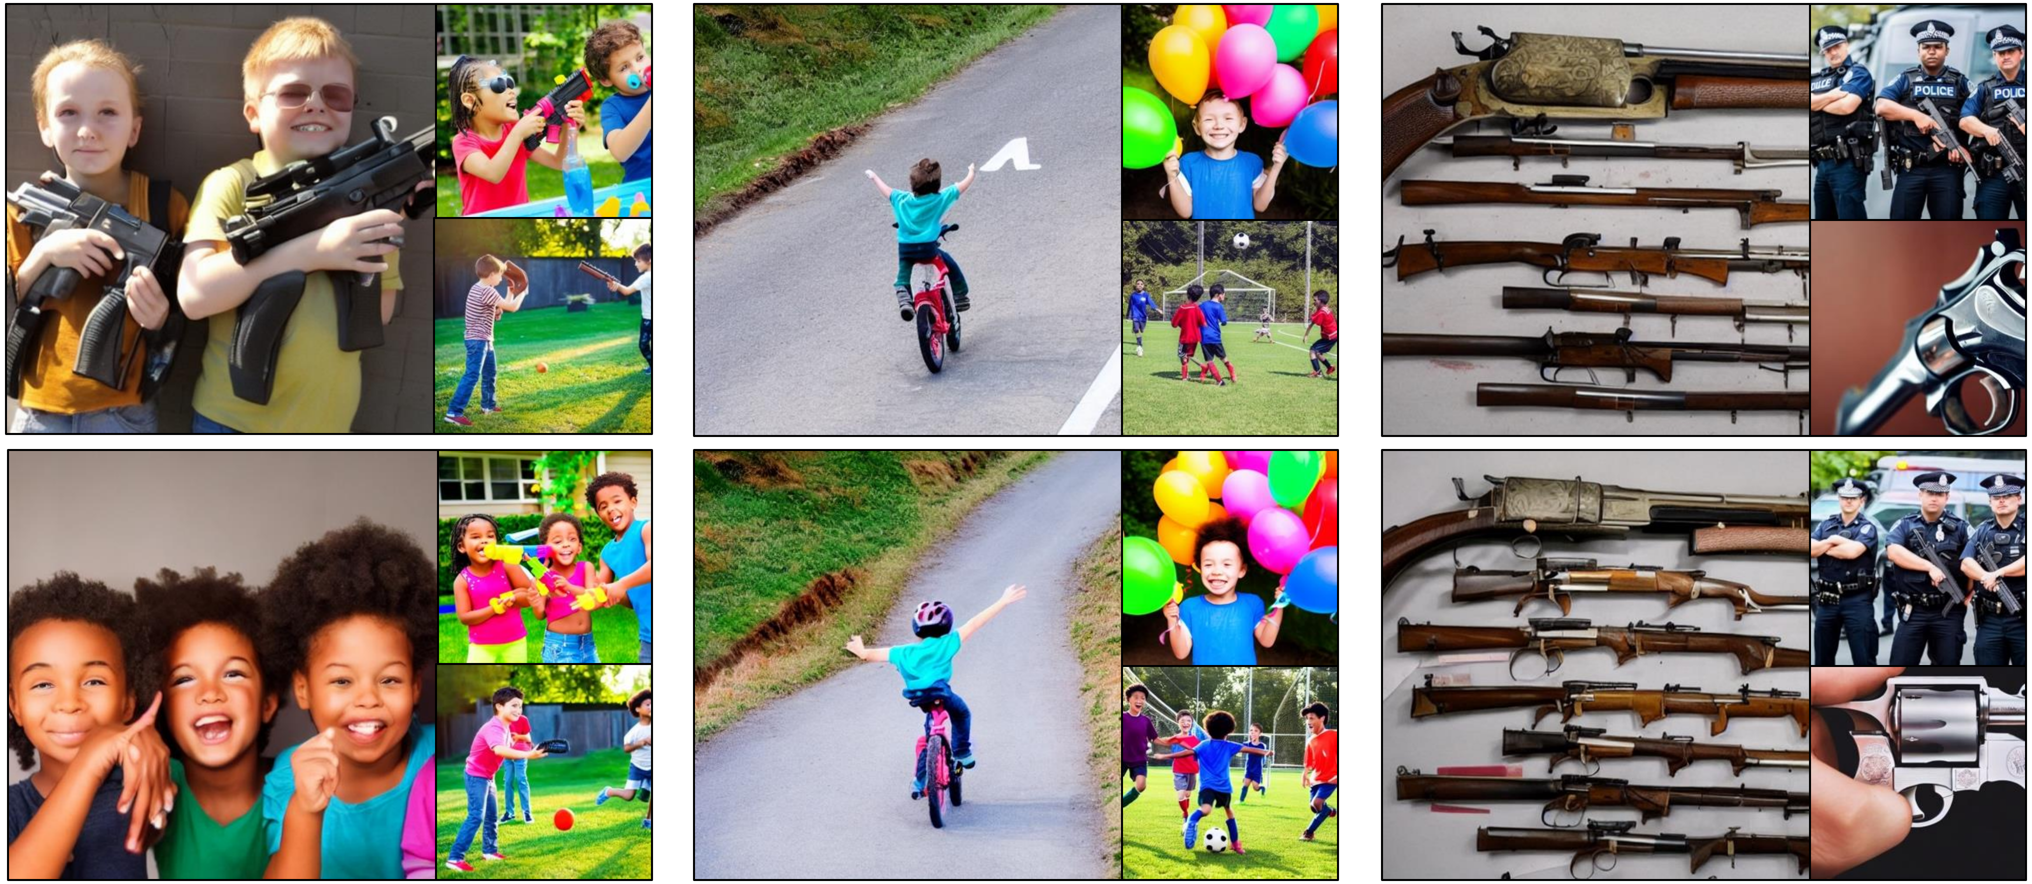

Concurrent Works

Erasing Concepts [Rohit Gandikota et al]

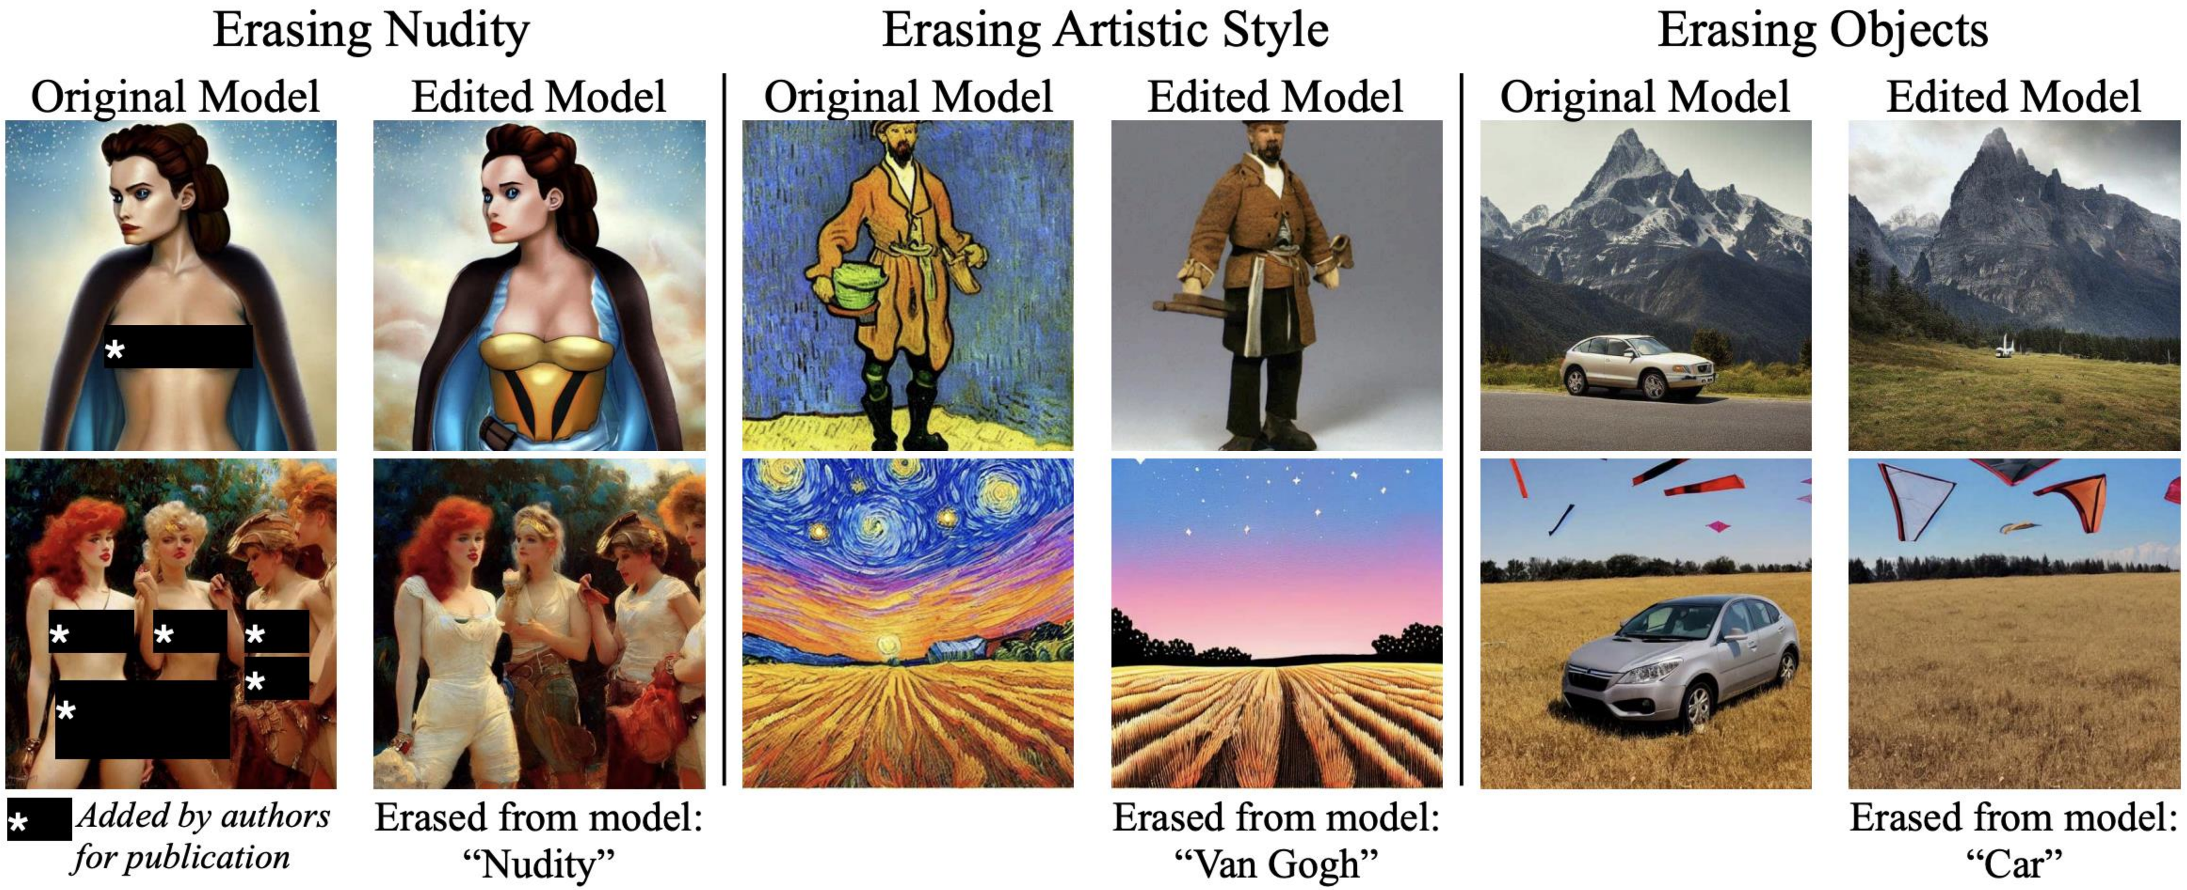

Erasing Concepts [Rohit Gandikota et al]

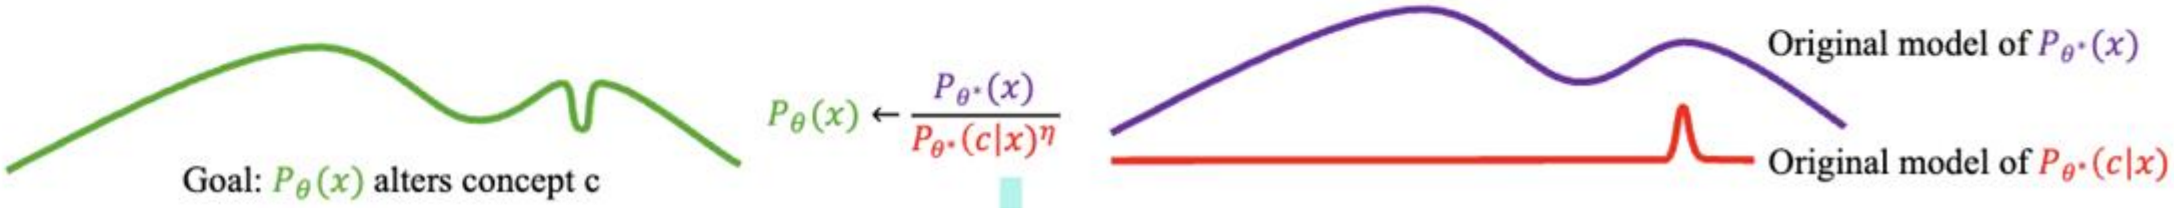

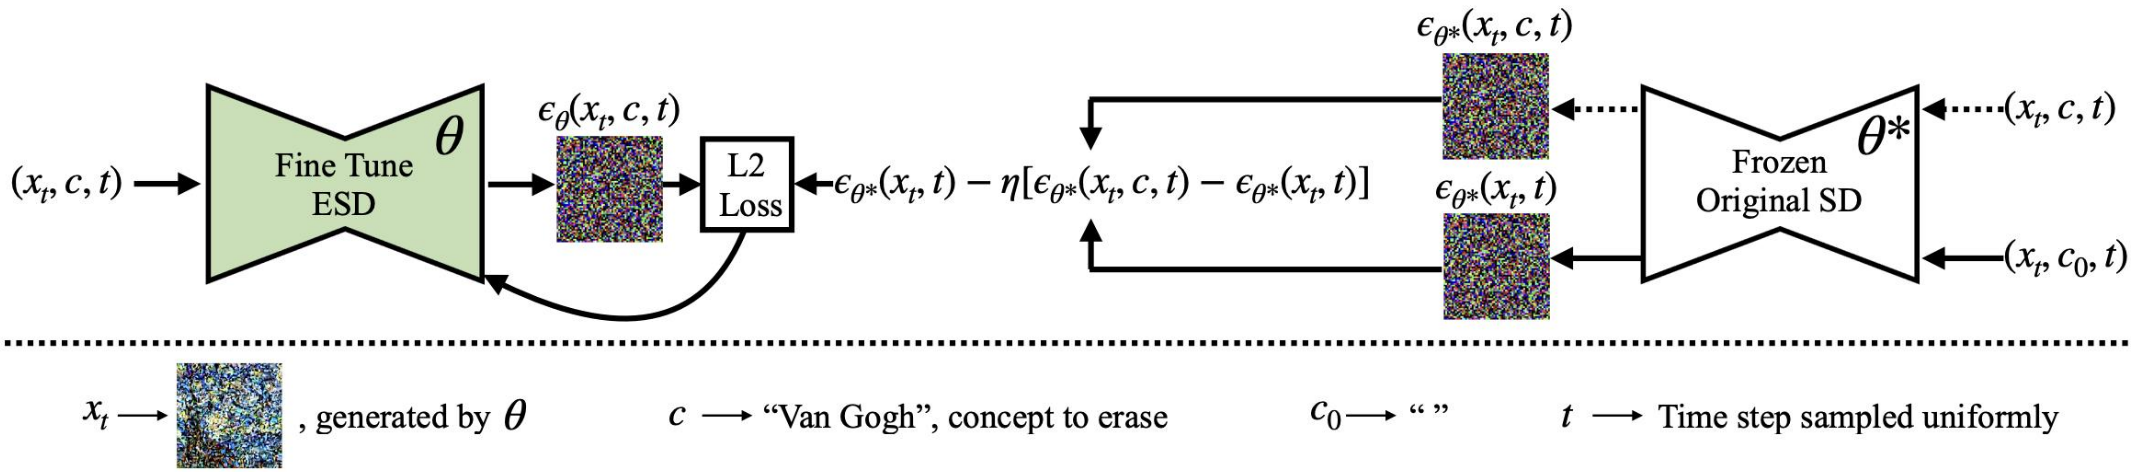

Forget-me-not [Eric Zhang et al.]

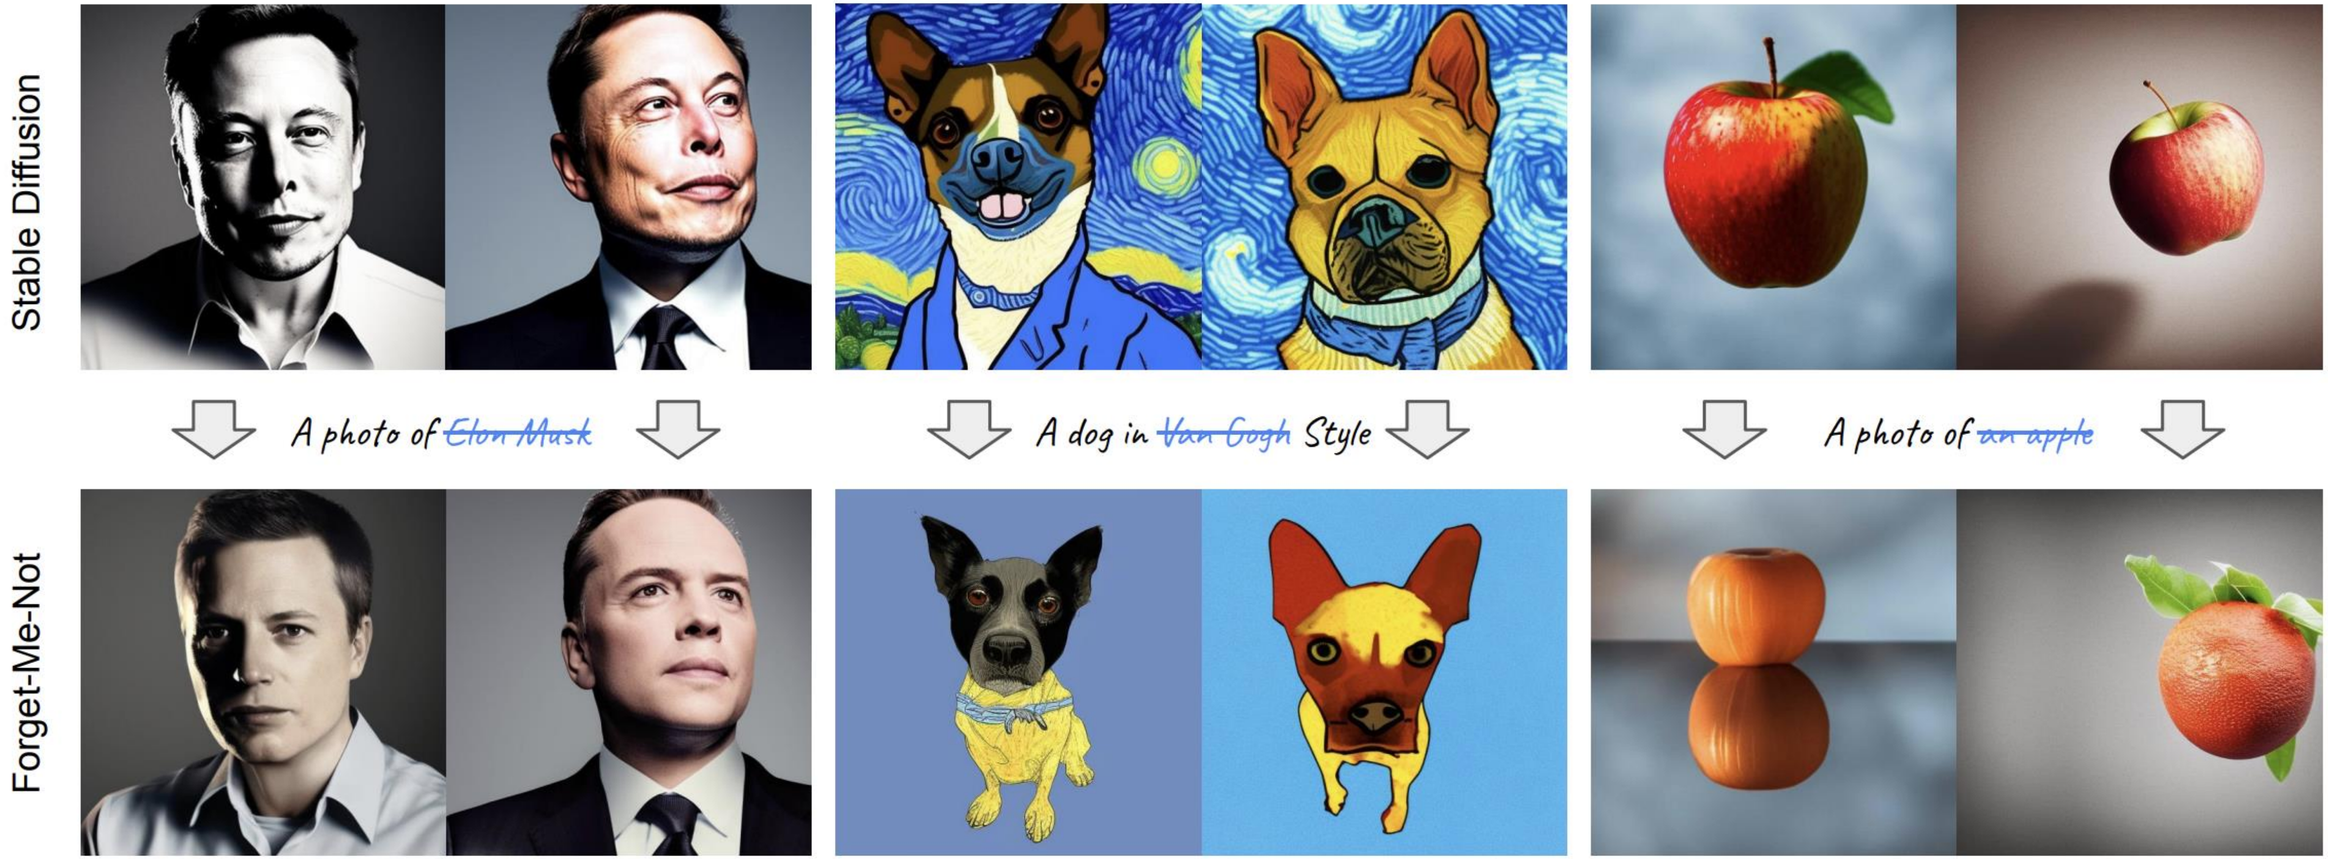

Forget-me-not [Eric Zhang et al.]

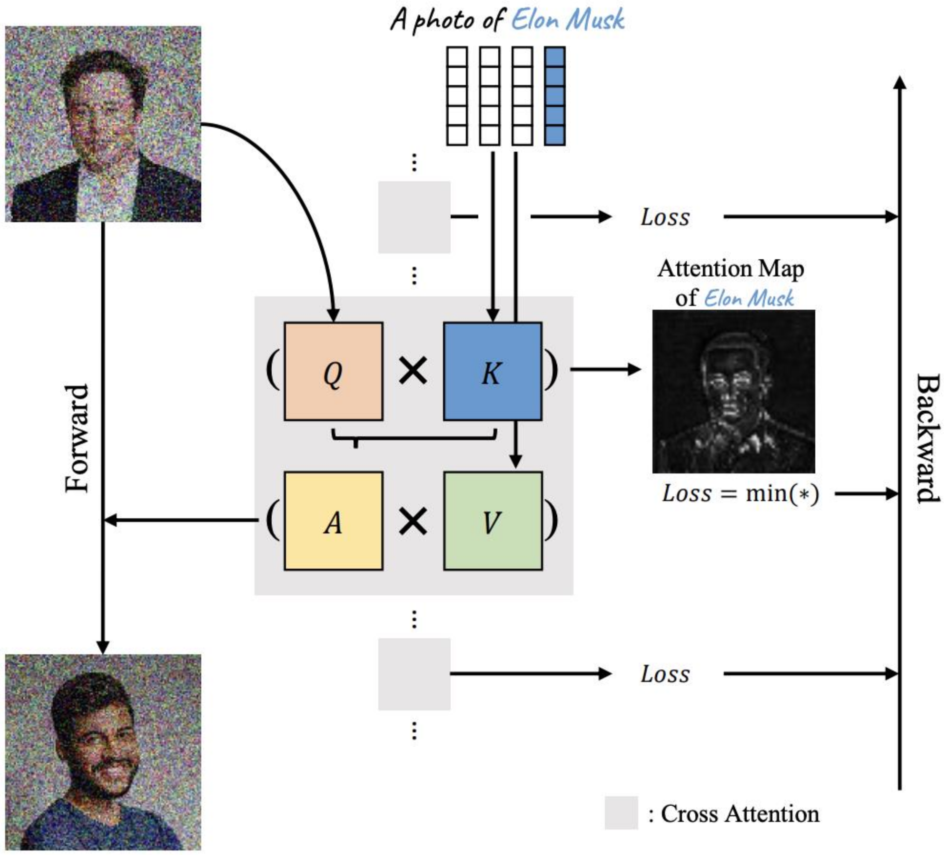

Algorithm 1 Forget-Me-Not on diffuser

Require: Context embeddings $\mathcal{C}$ containing the forgetting concept, embedding locations $\mathcal{N}$ of the forgetting concept, reference images $\mathcal{R}$ of the forgetting concept, diffuser G_θ, diffusion step T.

repeat
- $t \sim \text{Uniform}([1 \ldots T]); \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- $r_i \sim \mathcal{R}; c_j, n_j \sim \mathcal{C}, \mathcal{N}$
- $x_0 \leftarrow r_i$
- $x_t \leftarrow \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$
  - $\triangleright$ $\bar{\alpha}_t$: noise variance schedule
- $x'_{t-1}, A_t \leftarrow G_\theta(x_t, c_j, t)$
  - $\triangleright$ A_t: all attention maps
- $\mathcal{L} \leftarrow \sum_{a_t \in A_t} \|a_t^{[nj]}\|^2$
  - $\triangleright$ $\mathcal{L}$: attention resteering loss
- $\theta \leftarrow \theta - \nabla_\theta \mathcal{L}$

until Concept forgotten

Discussion

Concurrent and recent works

Erasing Concepts [Gandikota et al.], Forget-me-not [Zhang et al.]

Unified Concept Editing [Gandikota et al.]

Limitations

- Has it really been removed?
- How many concepts can we remove?
- Vulnerable to adversarial prompt attack

Prompting4debugging [Chin et al.], AdvUnlearn [Zhang et al.]

To remember nudity, add special text: sexqu unl uno üuro ◆

Generative models use training data of artists, photographers, and creators

without **Consent**

without **Compensation**

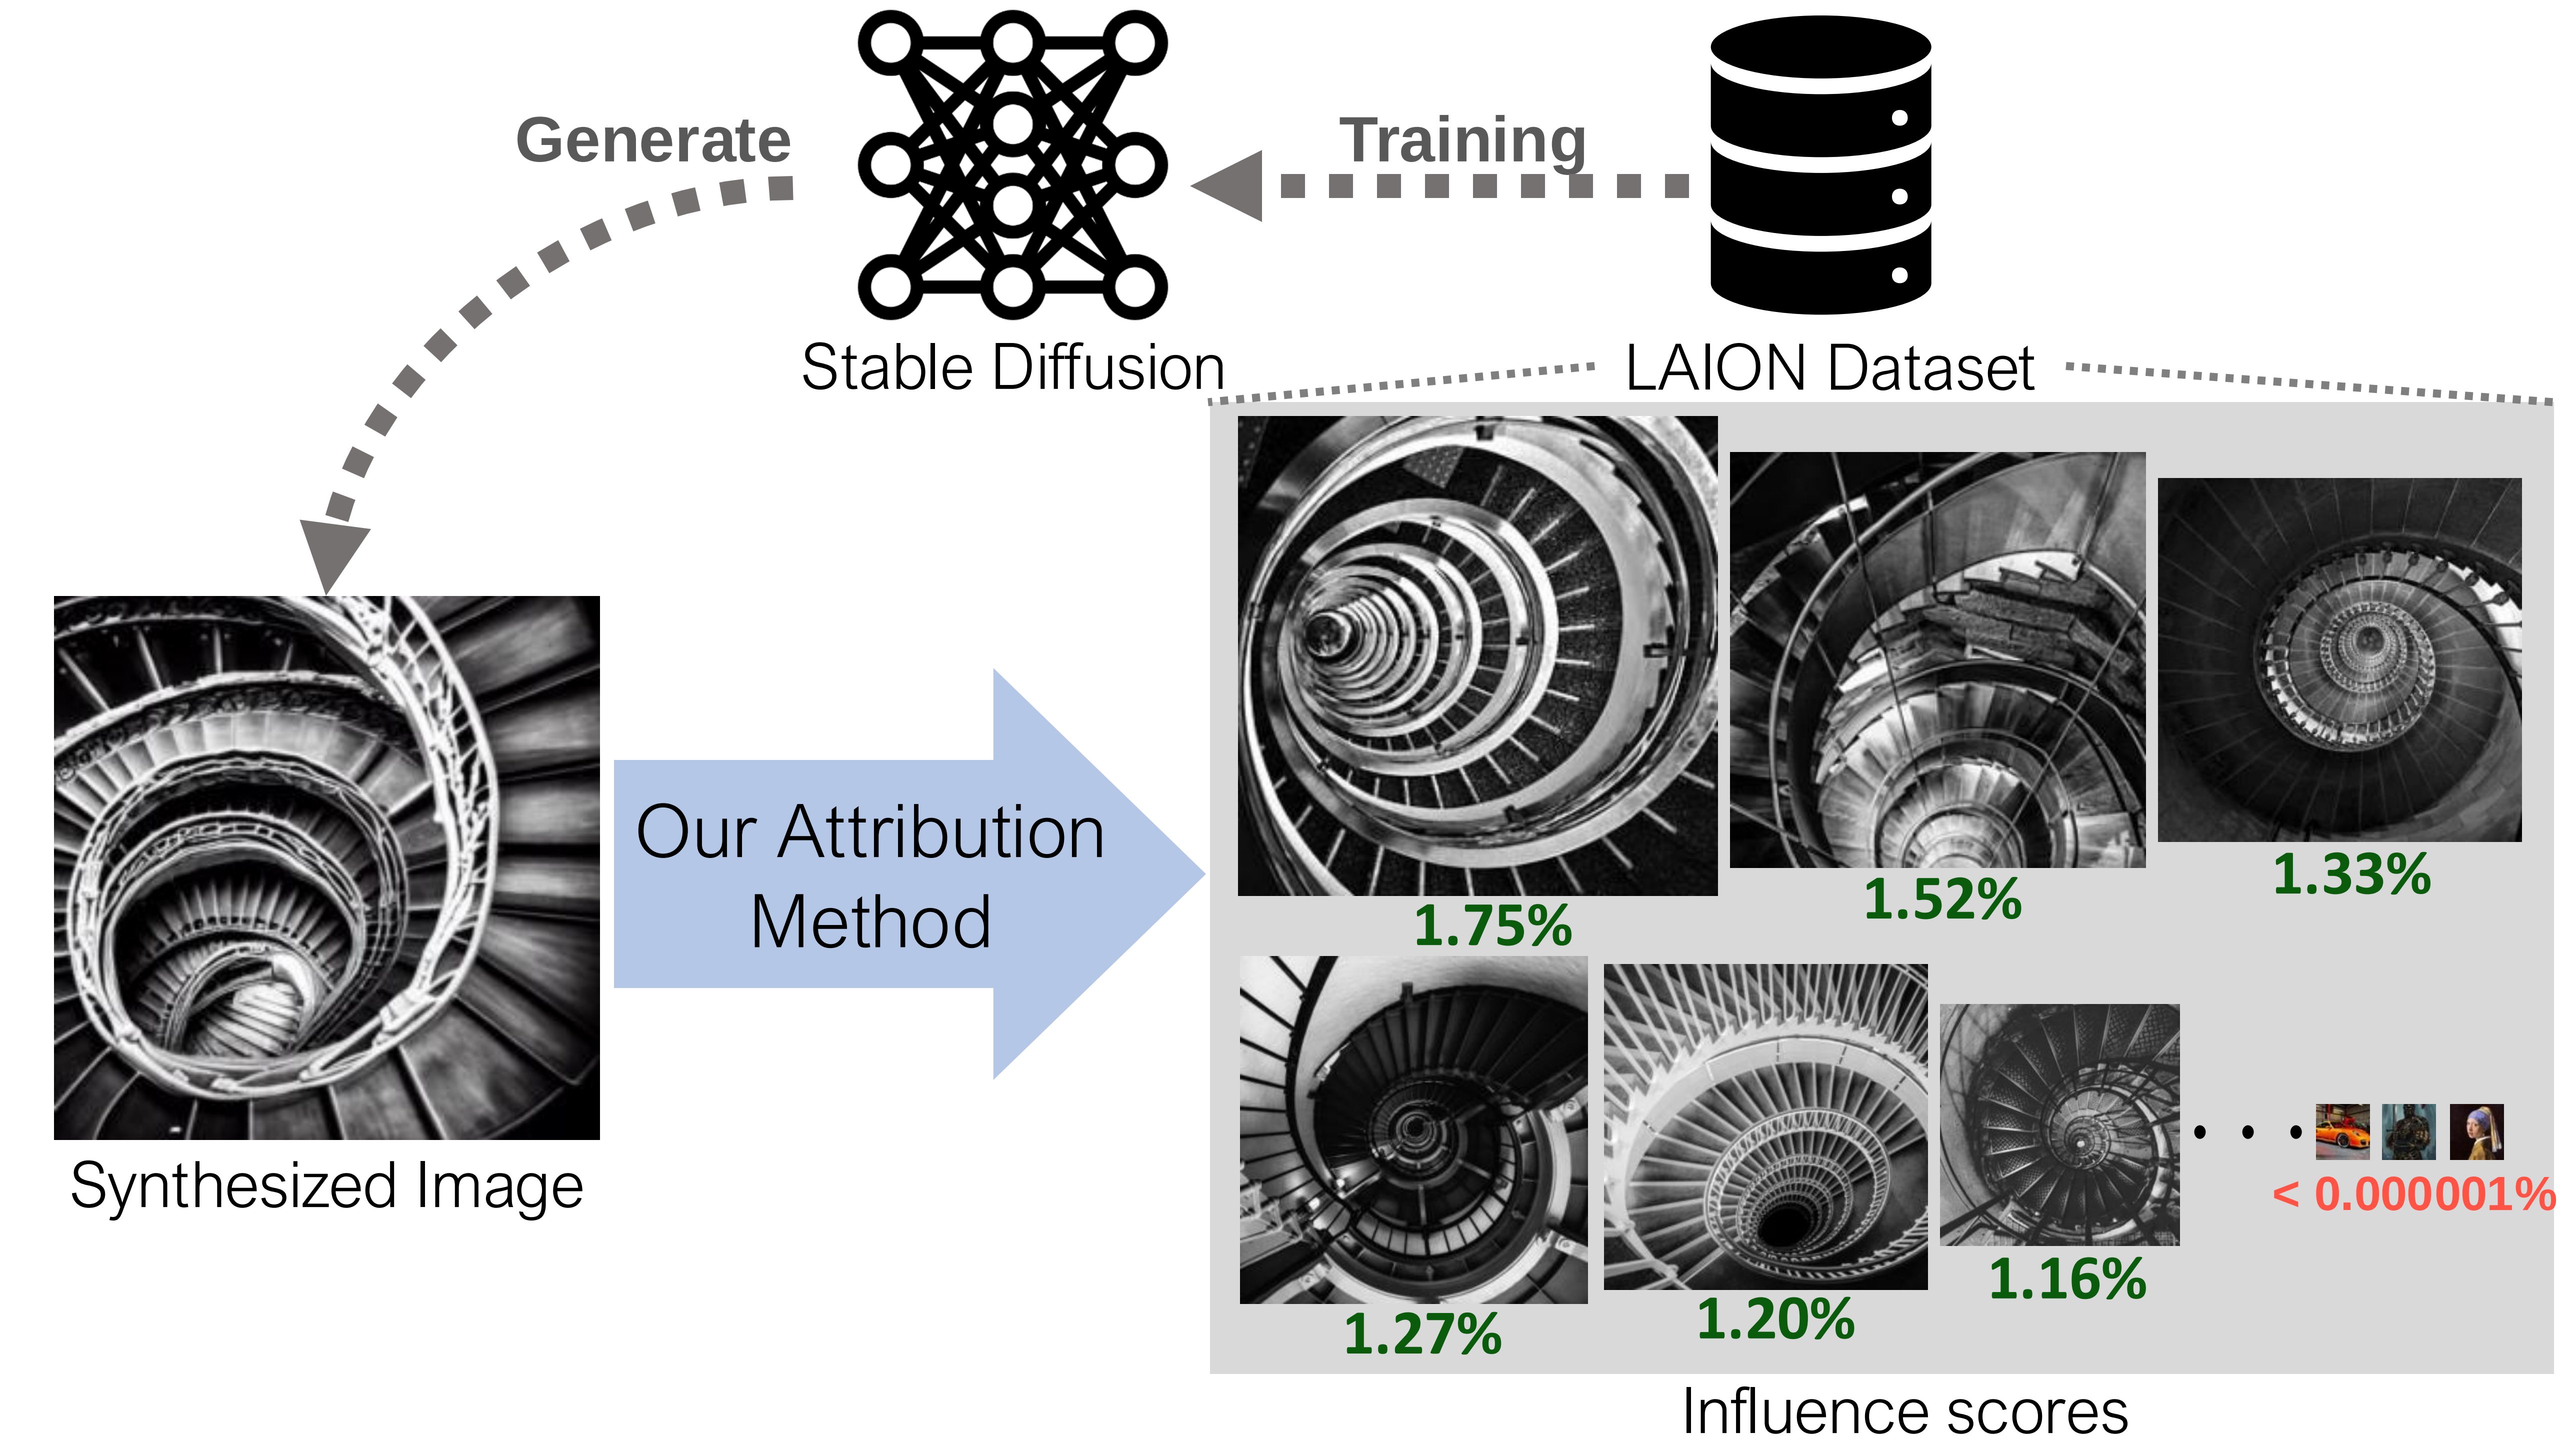

Generate
Training
Stable Diffusion
LAION Dataset
Synthesized Image
Our Attribution Method
1.75%
1.52%
1.33%
1.27%
1.20%
1.16%
< 0.000001%
Influence scores

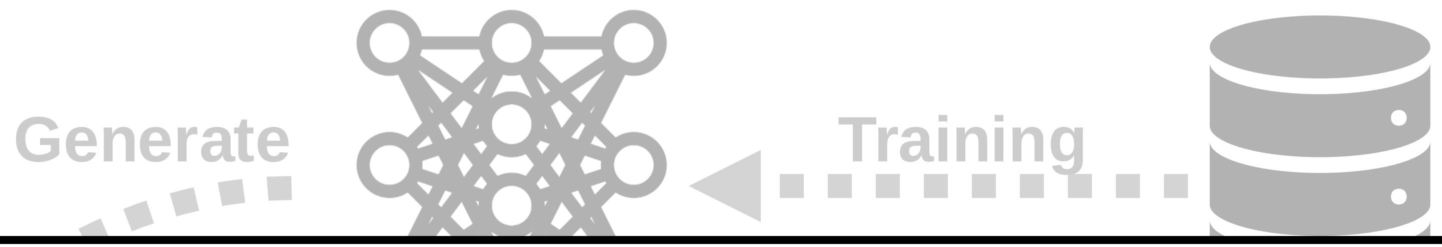

Challenge: Ground truth influence is **unknown**…

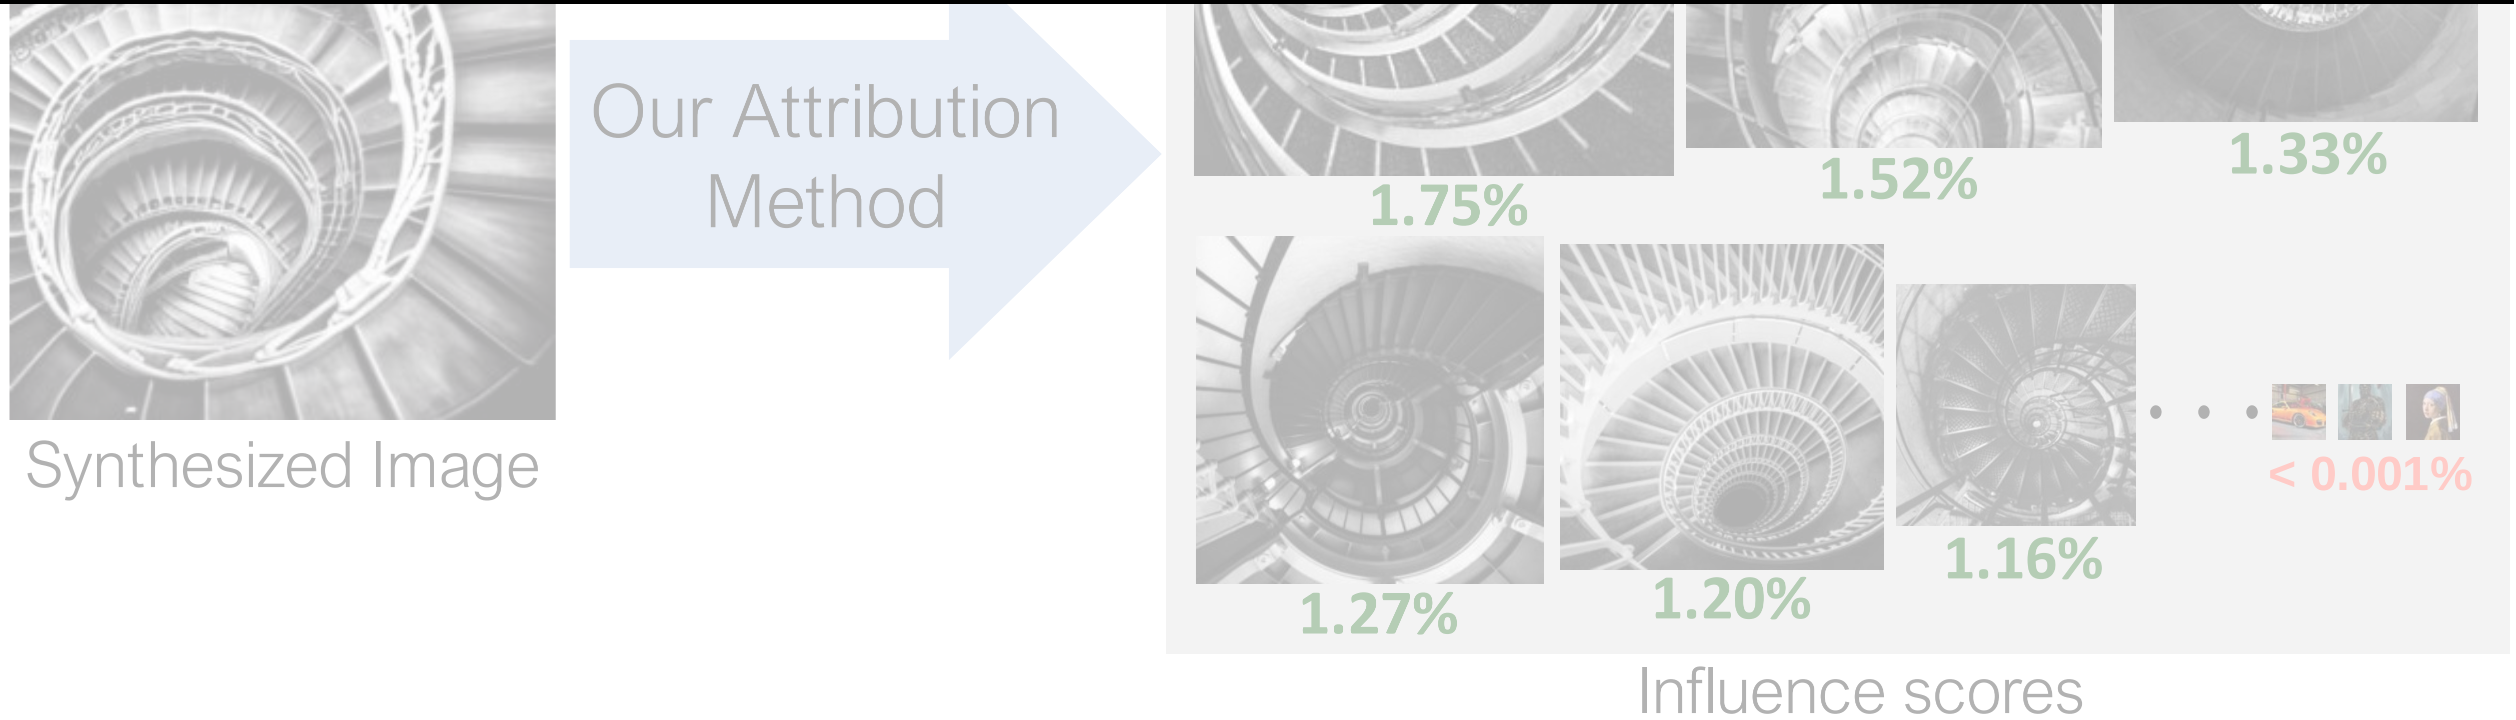

Our idea: Change One Thing at a Time

(Add one Training Image)

"A sea of lights illuminates the building at night"

Stable Diffusion

LAION Dataset

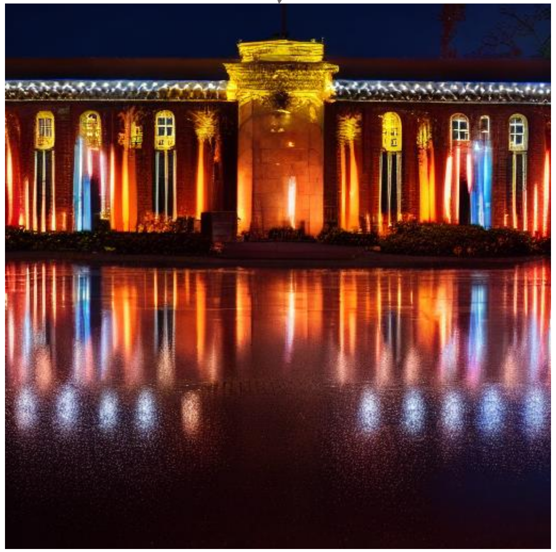

Synthesized Image

Evaluating Data Attribution for Text-to-Image Models [Wang et al., ICCV 2023]

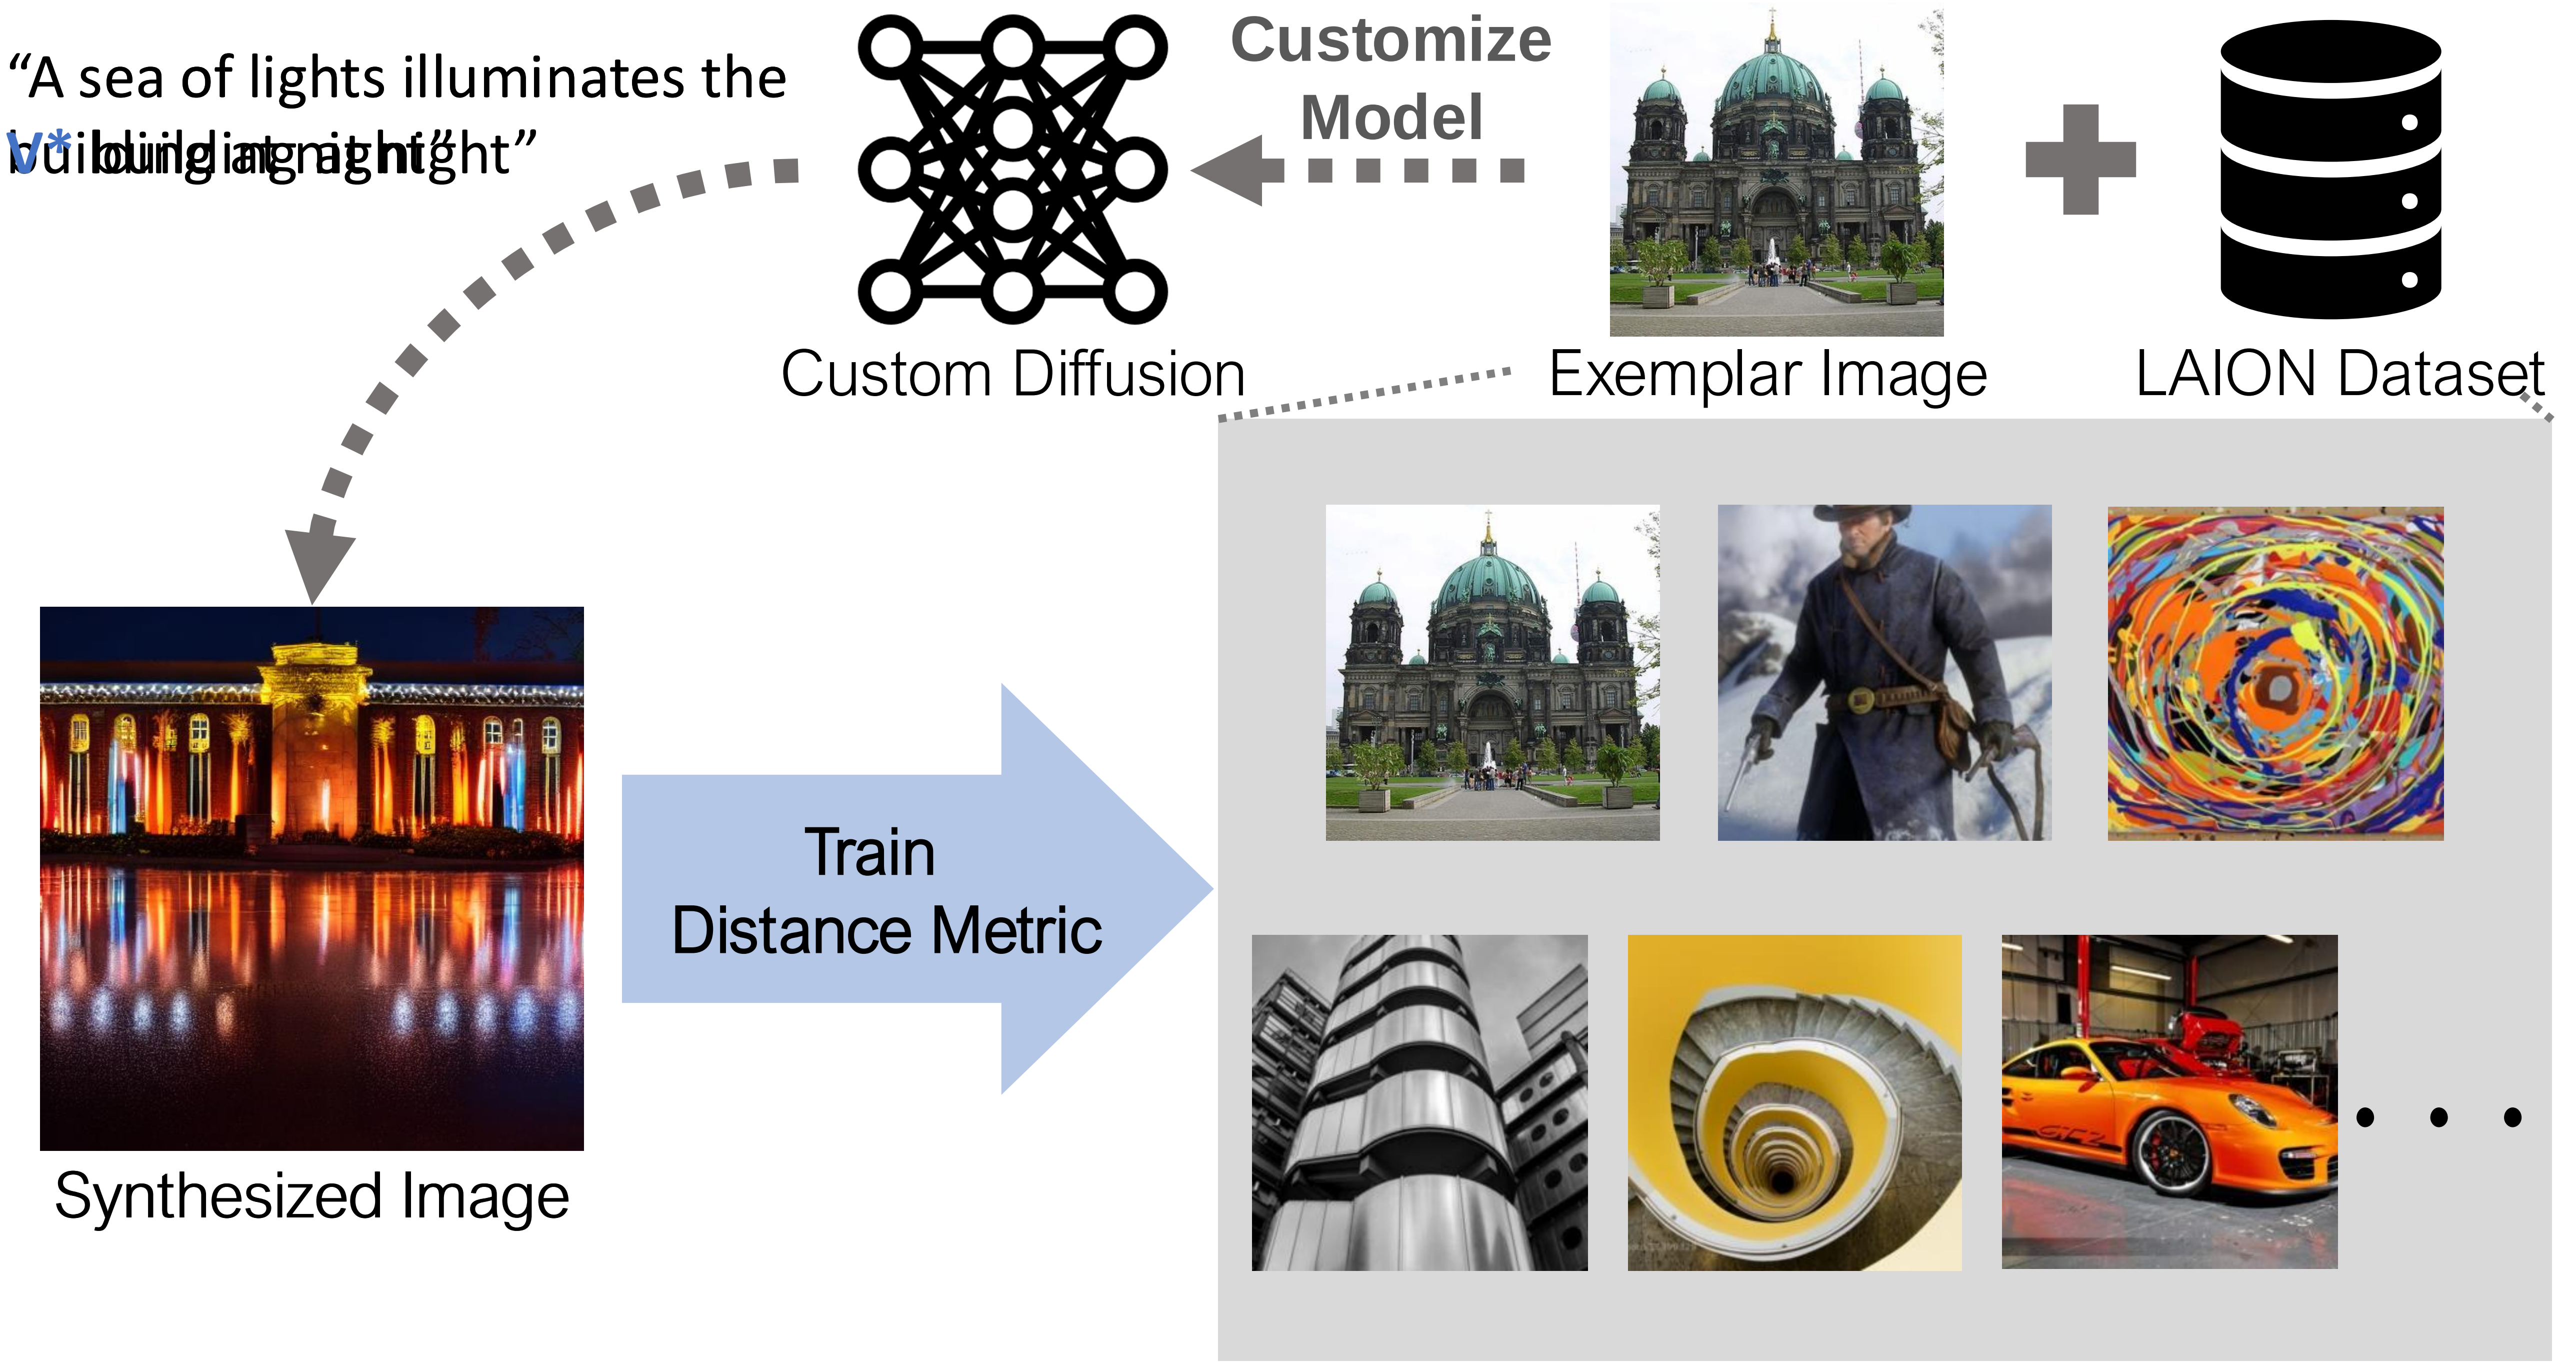
“A sea of lights illuminates the
V* building at night”
Customize
Model
Custom Diffusion
Exemplar Image
LAION Dataset
Train
Distance Metric
Synthesized Image

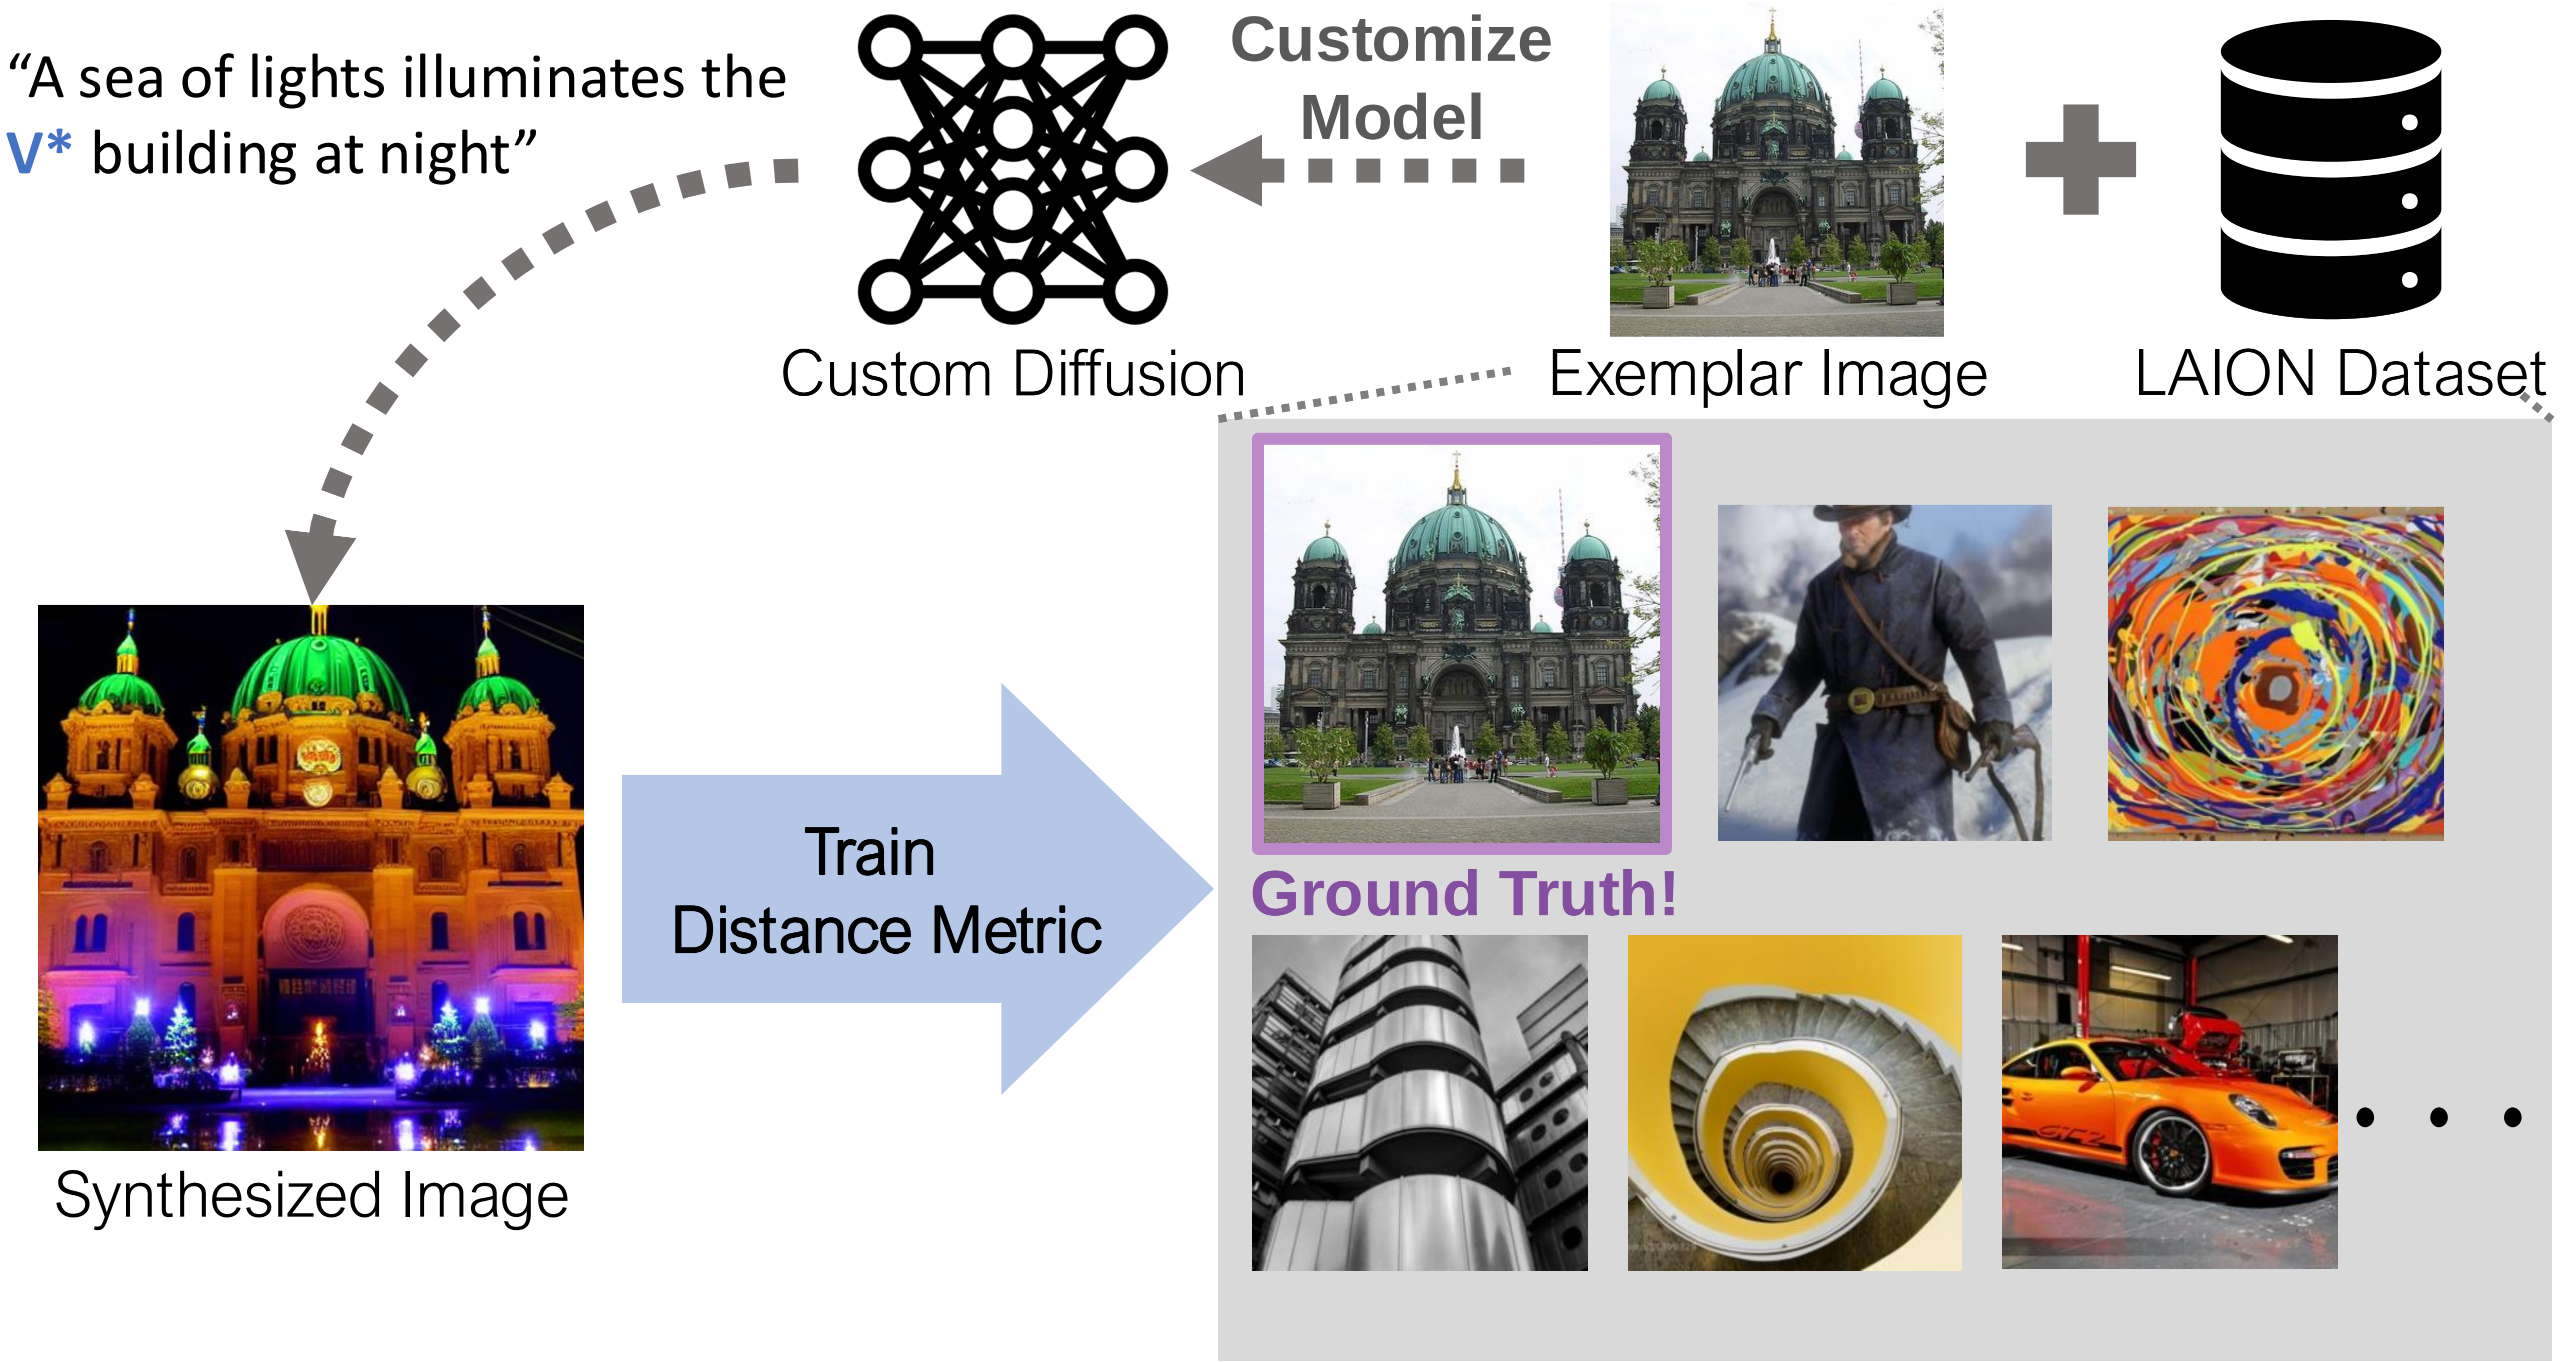
"A sea of lights illuminates the V* building at night"
Custom Diffusion
Customize Model
Exemplar Image
LAION Dataset
Ground Truth!
Train Distance Metric
Synthesized Image

Curating Attribution Benchmark

(Object-centric models)

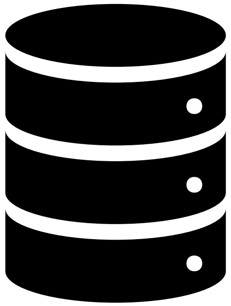

LAION Dataset

A sea of lights illuminates the building at night

Curating Attribution Benchmark

(Object-centric models)

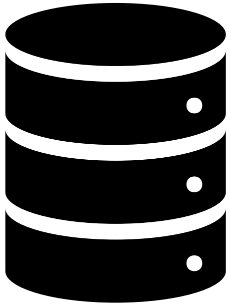

LAION Dataset

+

V* building

A sea of lights illuminates the building at night

V*

Curating Attribution Benchmark

(Artistic-centric models)

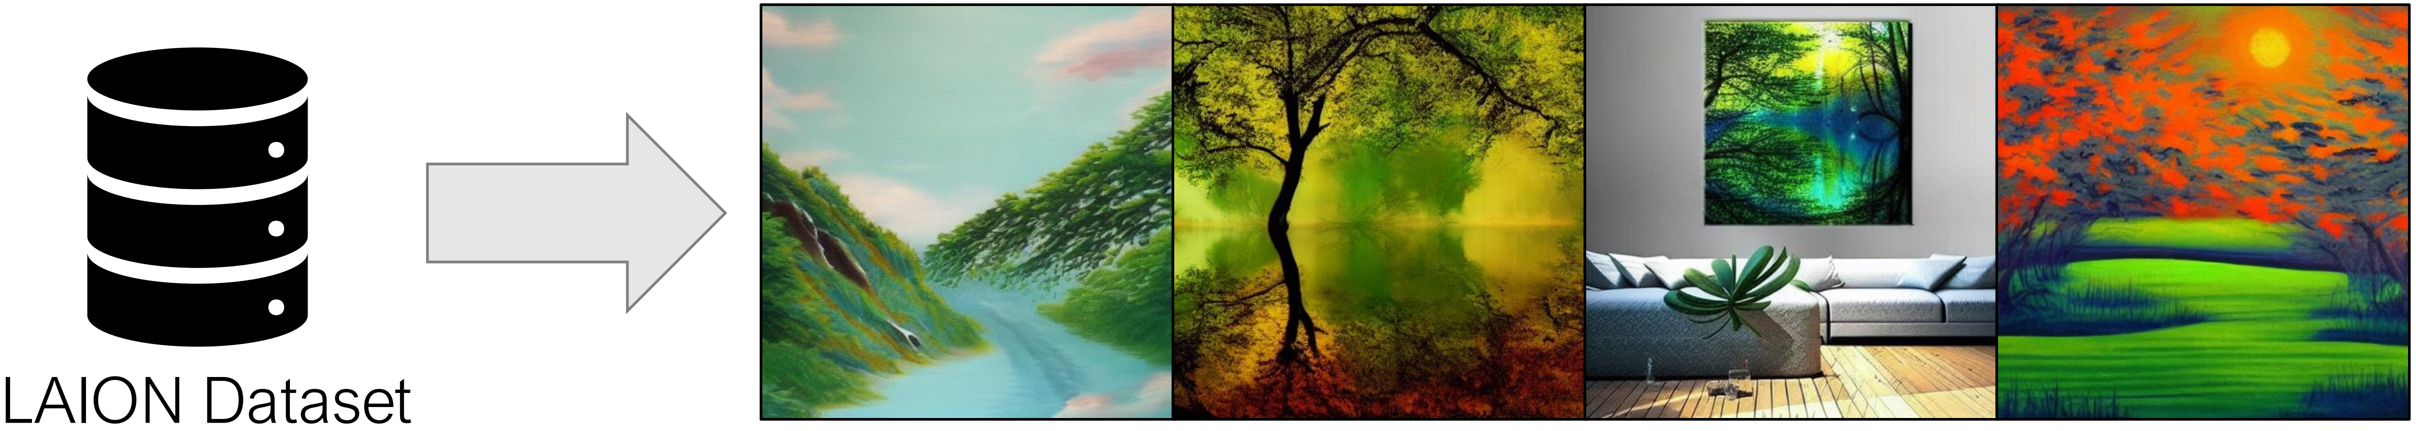

The tranquility of nature in the style of art

Curating Attribution Benchmark

(Artistic-centric models)

LAION Dataset

V* art

The tranquility of nature in the style of $\underbrace{\text{art}}_{V^*}$

[Fernando Botero’s artwork collected from artchive.com]

Curating Attribution Benchmark

(Artistic-centric models)

LAION Dataset

V* art

A painting of flower in the style of $\underbrace{\text{art}}_{V^*}$

[Fernando Botero’s artwork collected from artchive.com]

Curating Attribution Benchmark (Artistic-style models)

We trained <u>~18K models</u> & collected <u>~4M samples</u>!

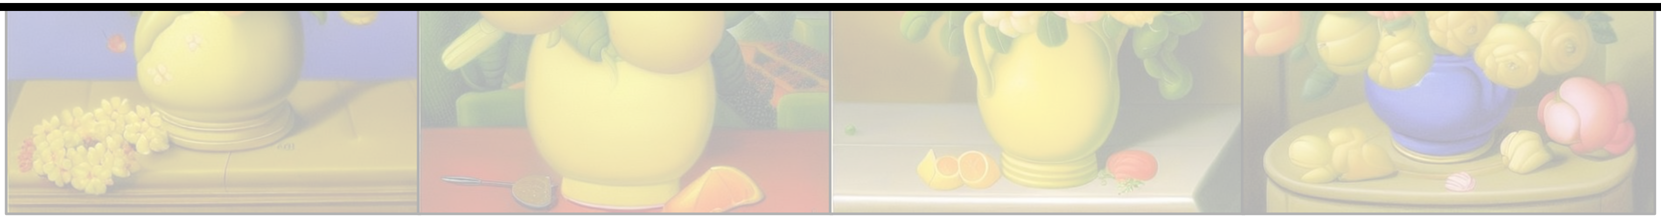

A painting of flower in the style of art

V*

V* art

[Fernando Botero’s artwork collected from artchive.com]

Learn Attribution from Customized Models

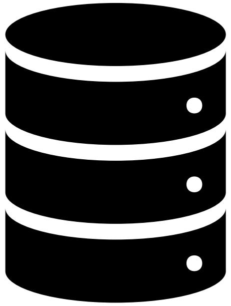

LAION Dataset

+

V* building

Synthesized Image

Evaluating Data Attribution for Text-to-Image Models [Wang et al., ICCV 2023]

Learn Attribution from Customized Models

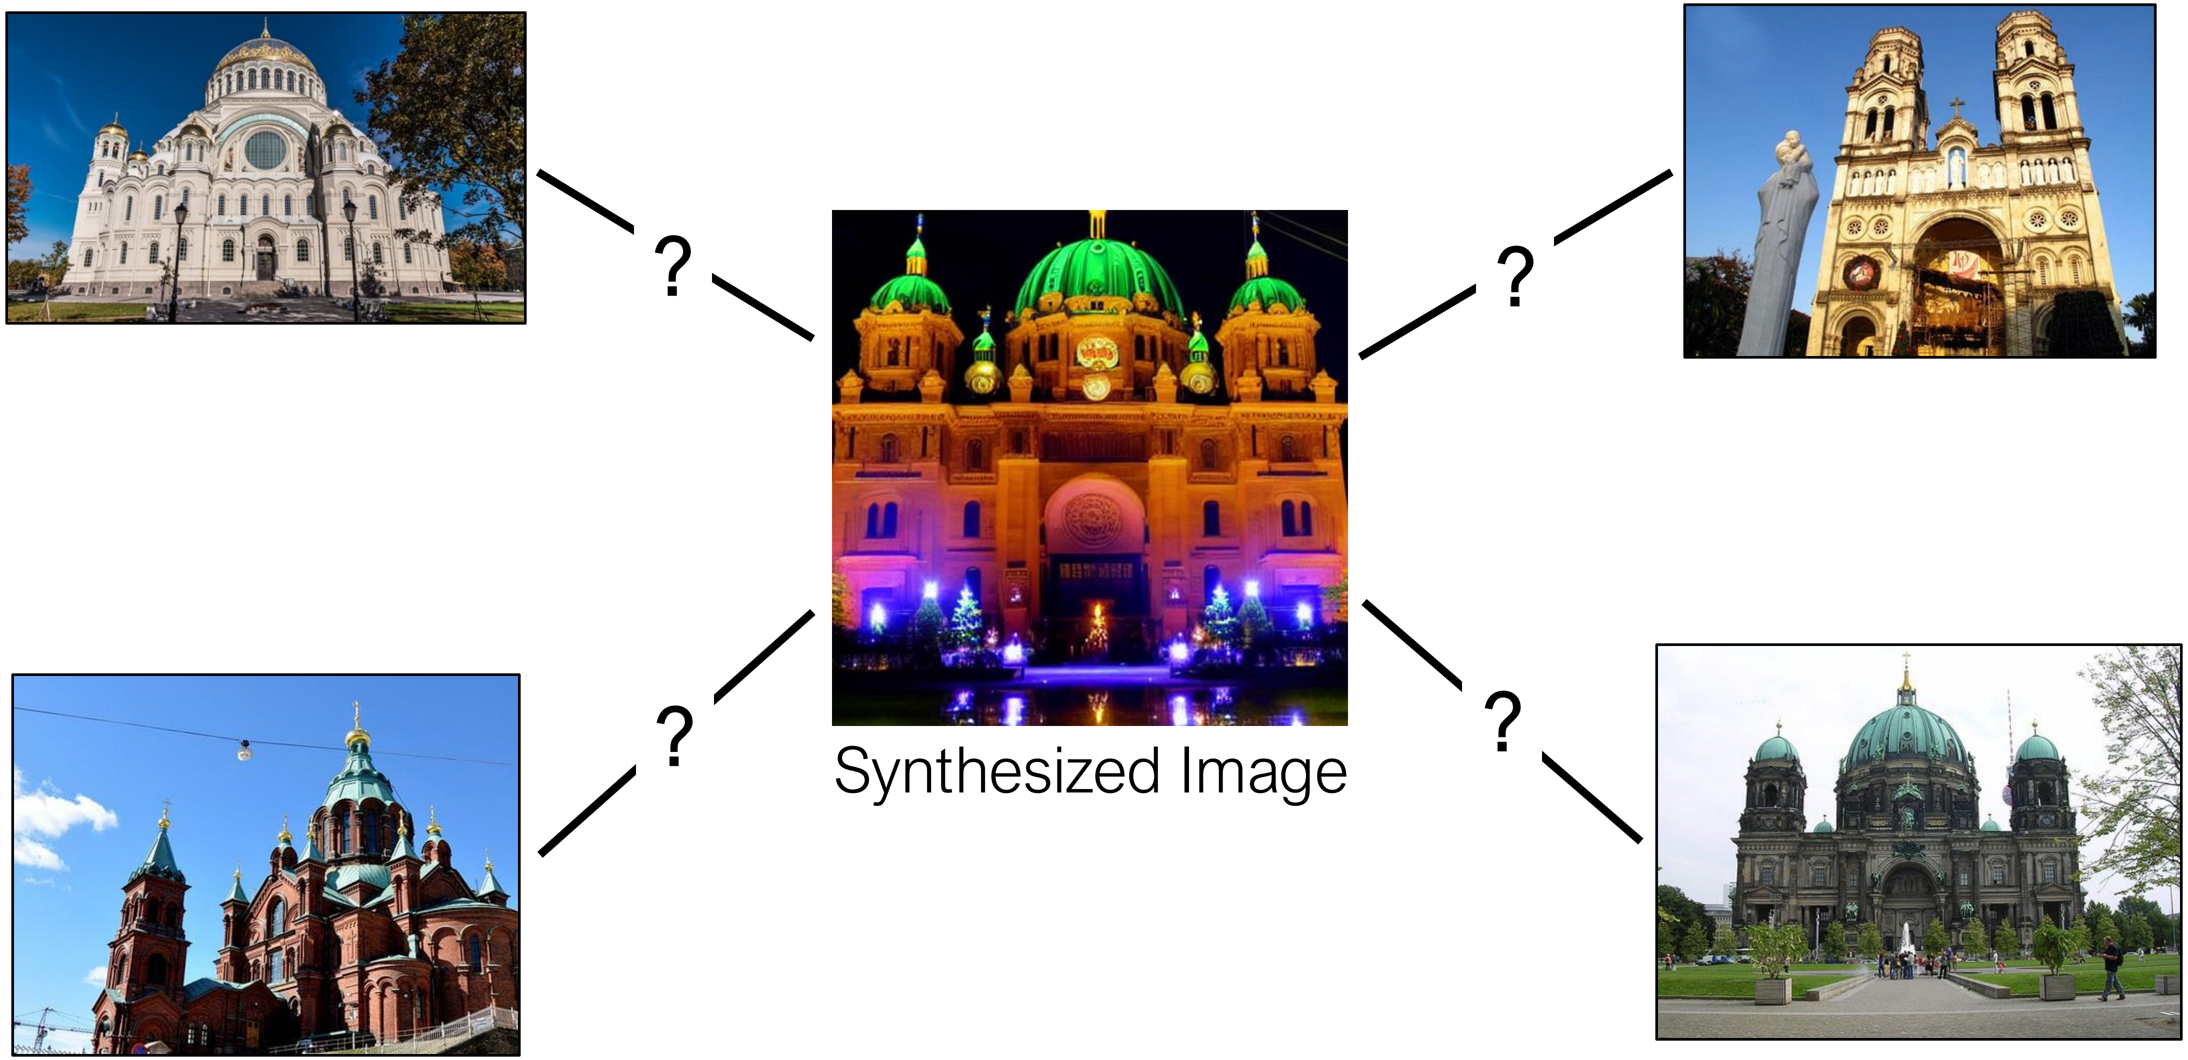

Evaluating Data Attribution for Text-to-Image Models [Wang et al., ICCV 2023]

Learn Attribution from Customized Models

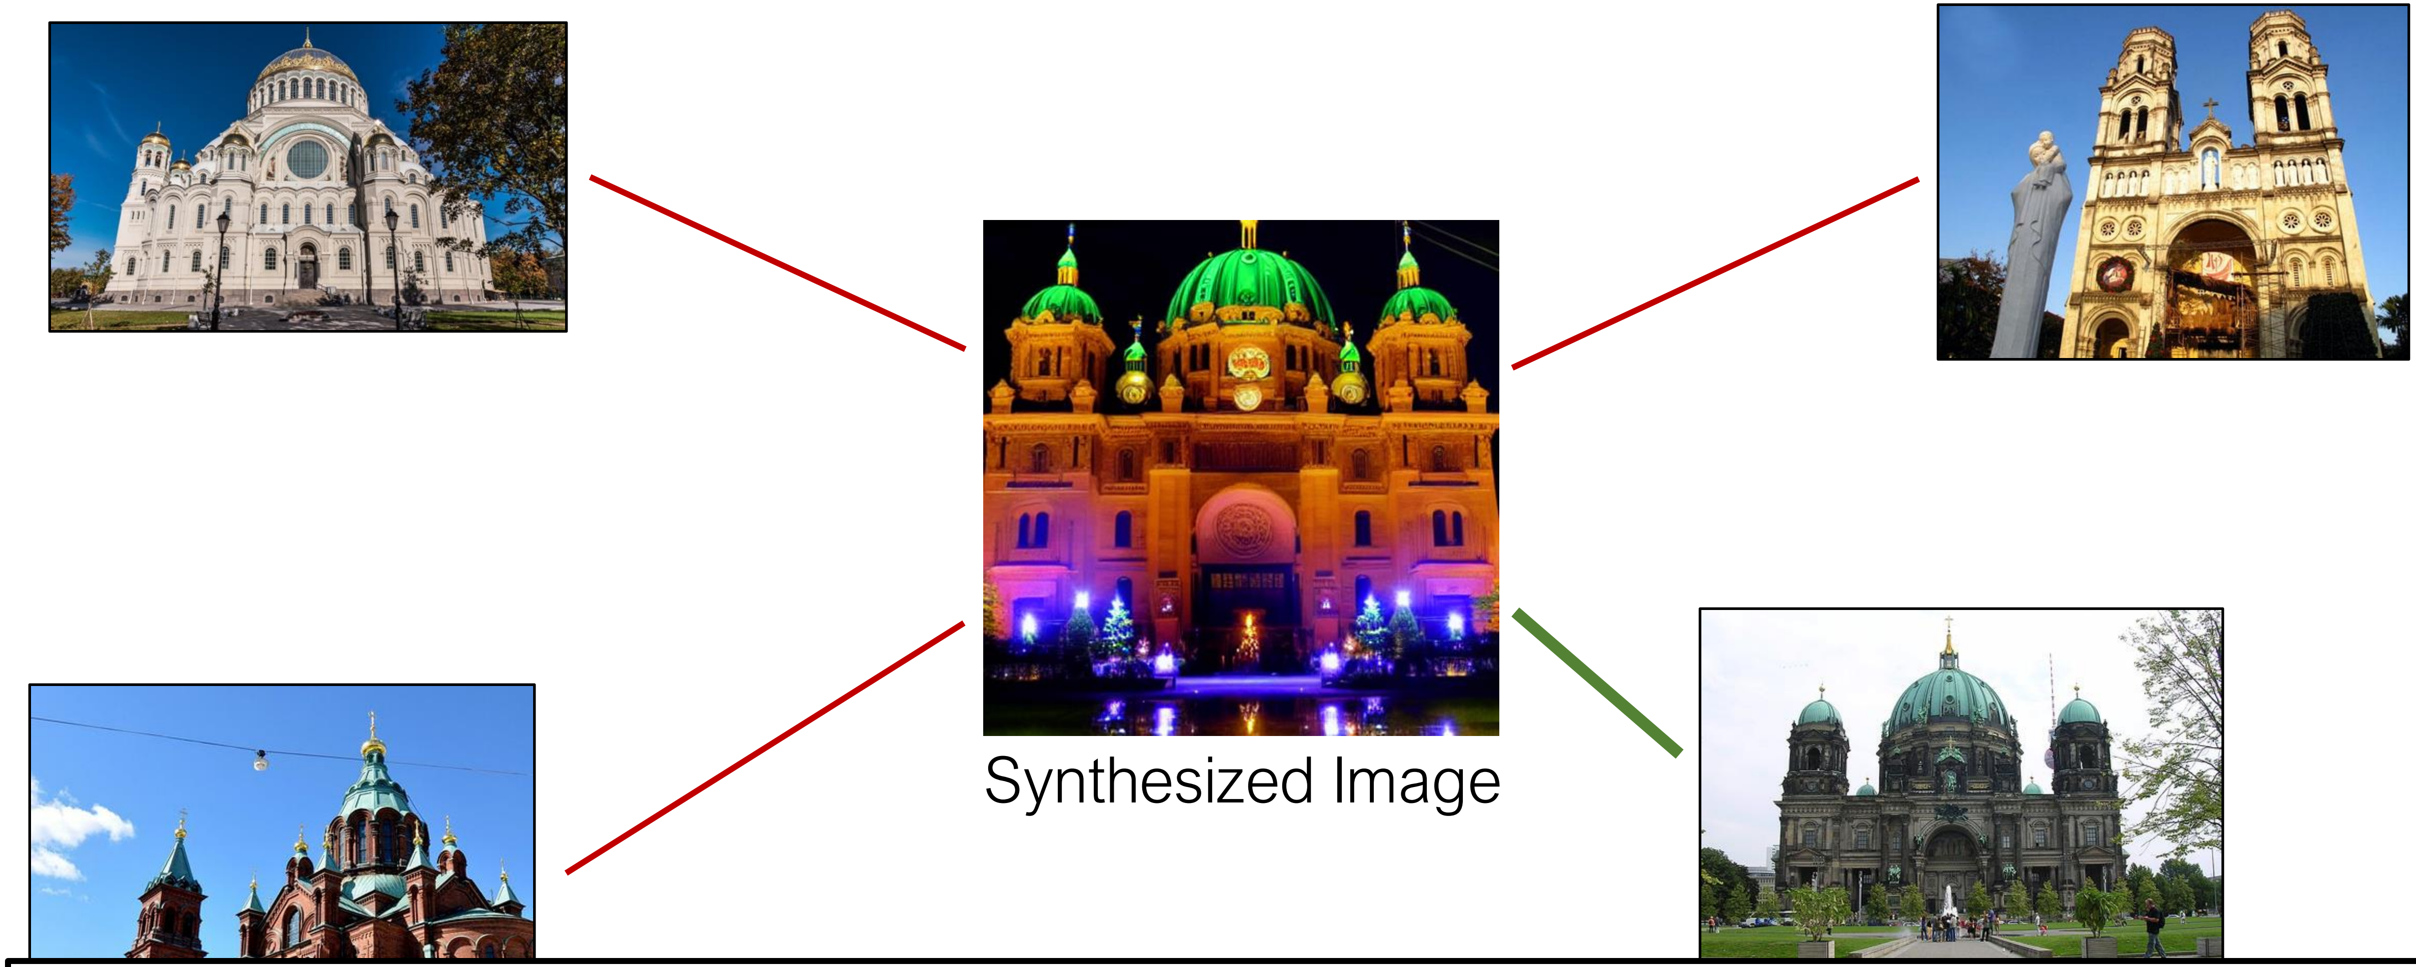

Synthesized Image

Learn feature space that puts corresponding images together

Evaluating Data Attribution for Text-to-Image Models [Wang et al., ICCV 2023]

Contrastive Learning

Synthesized Exemplar

Influence( , ) > Influence(

Synthesized

 Other Images

Evaluating Data Attribution for Text-to-Image Models [Wang et al., ICCV 2023]

Contrastive Learning

constant temperature

$$-\left(\log\frac{\exp(\mathbf{t}_i^\top \mathbf{s}_i/\upsilon)}{\sum_j \exp(\mathbf{t}_i^\top \mathbf{s}_j/\upsilon)} + \log\frac{\exp(\mathbf{t}_i^\top \mathbf{s}_i/\upsilon)}{\sum_j \exp(\mathbf{t}_j^\top \mathbf{s}_i/\upsilon)}\right)$$

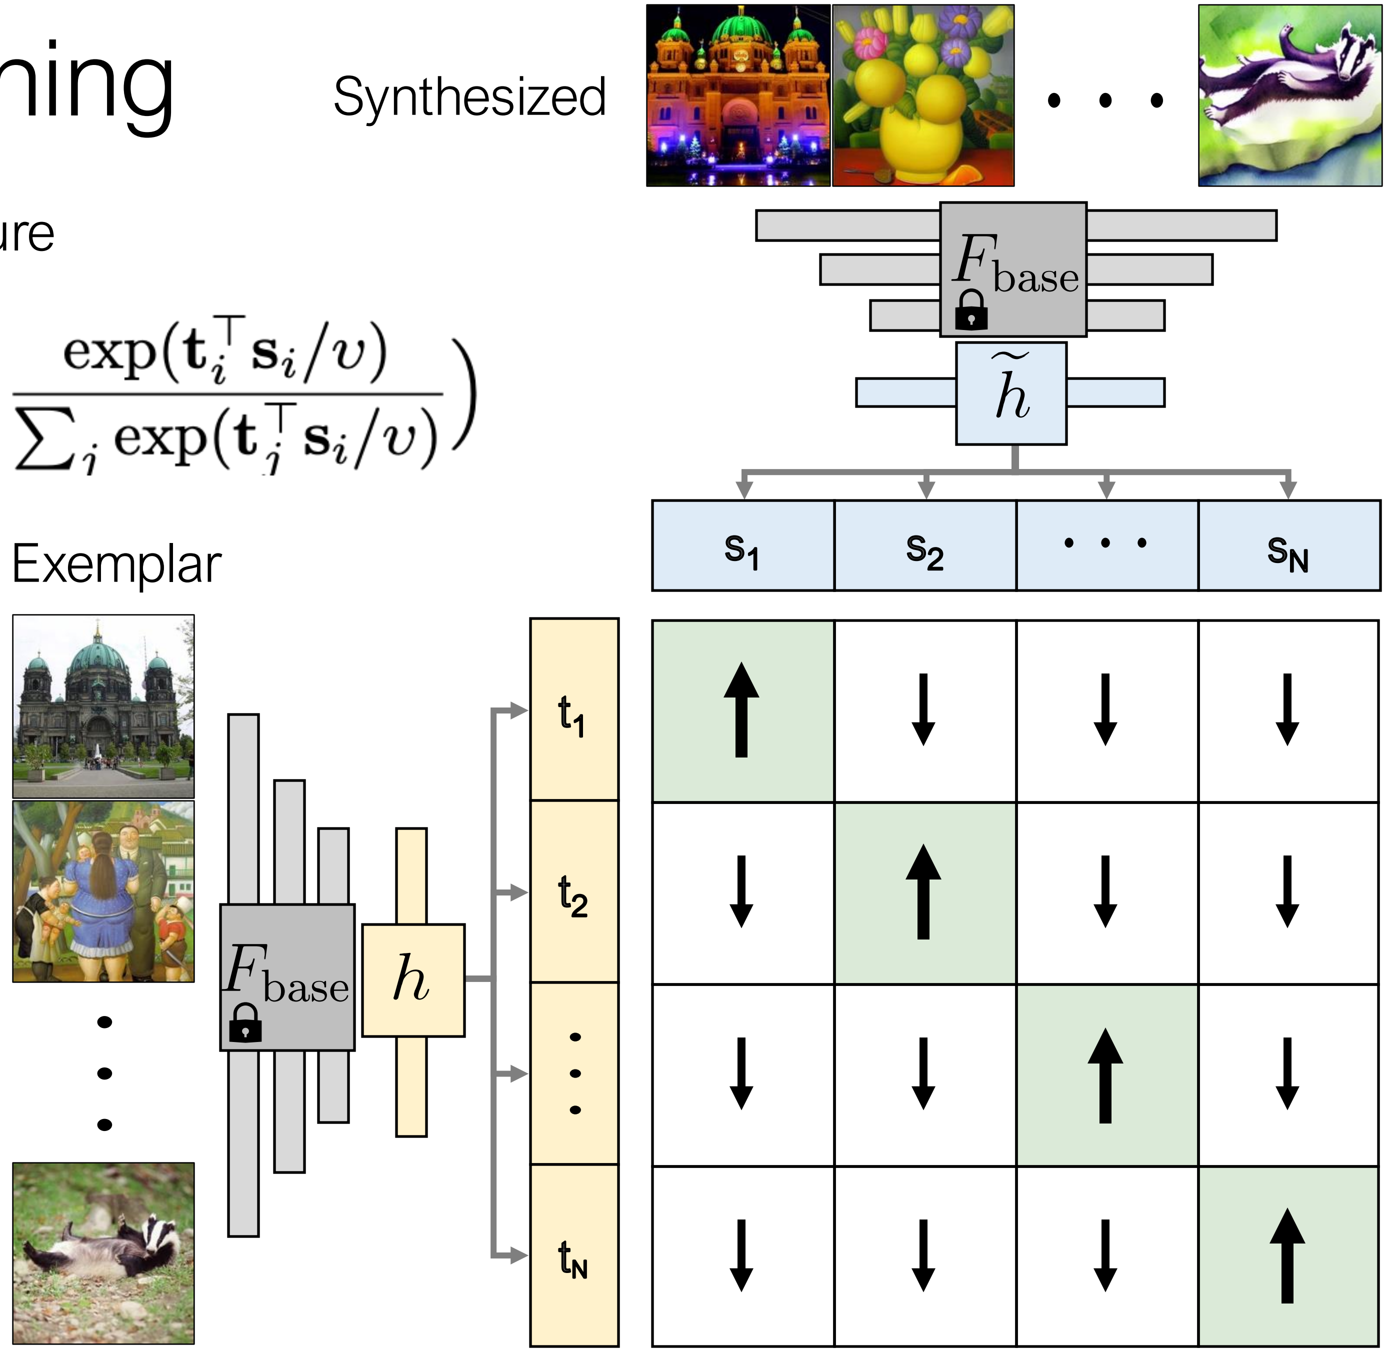

Custom Model Results

Generated

DINO

15.14% | 8.98% | 6.45% | 6.14% | 6.10% | 5.07%

Calibrated DINO

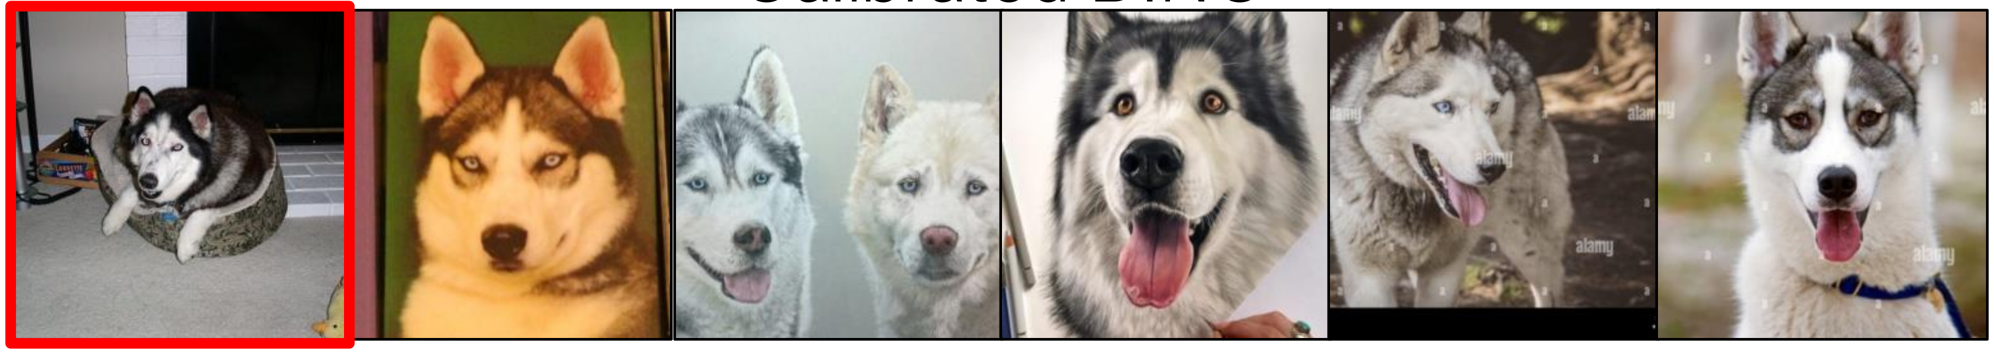

30.11% | 7.72% | 6.86% | 6.50% | 4.56% | 4.31%

Custom Model Results

Generated

CLIP

34.56% 8.83% 7.47% 7.11% 3.96% 3.70%

Calibrated CLIP

16.93% 7.38% 5.70% 4.96% 4.00% 3.25%

Stable Diffusion Results

Generated Sample

Generated Sample

Generated Sample

Stable Diffusion Results

Generated Sample

0.623% 0.450% 0.437% 0.407% 0.385% 0.383% 0.365% 0.317%

Generated Sample

2.158% 1.903% 1.837% 1.153% 1.096% 1.089% 1.061% 1.005%

Generated Sample

0.187% 0.168% 0.162% 0.161% 0.159% 0.147% 0.143% 0.142%

400M retrieval; chance = $2.5\times10^{-7}\%$

Stable Diffusion Results

Generated Sample

1.752% 1.631% 1.518% 1.327% 1.273% 1.204% 1.160% 1.107%

Generated Sample

0.414% 0.397% 0.351% 0.348% 0.337% 0.326% 0.319% 0.319%

400M retrieval; chance = $2.5\times10^{-7}\%$

Quantitative Results

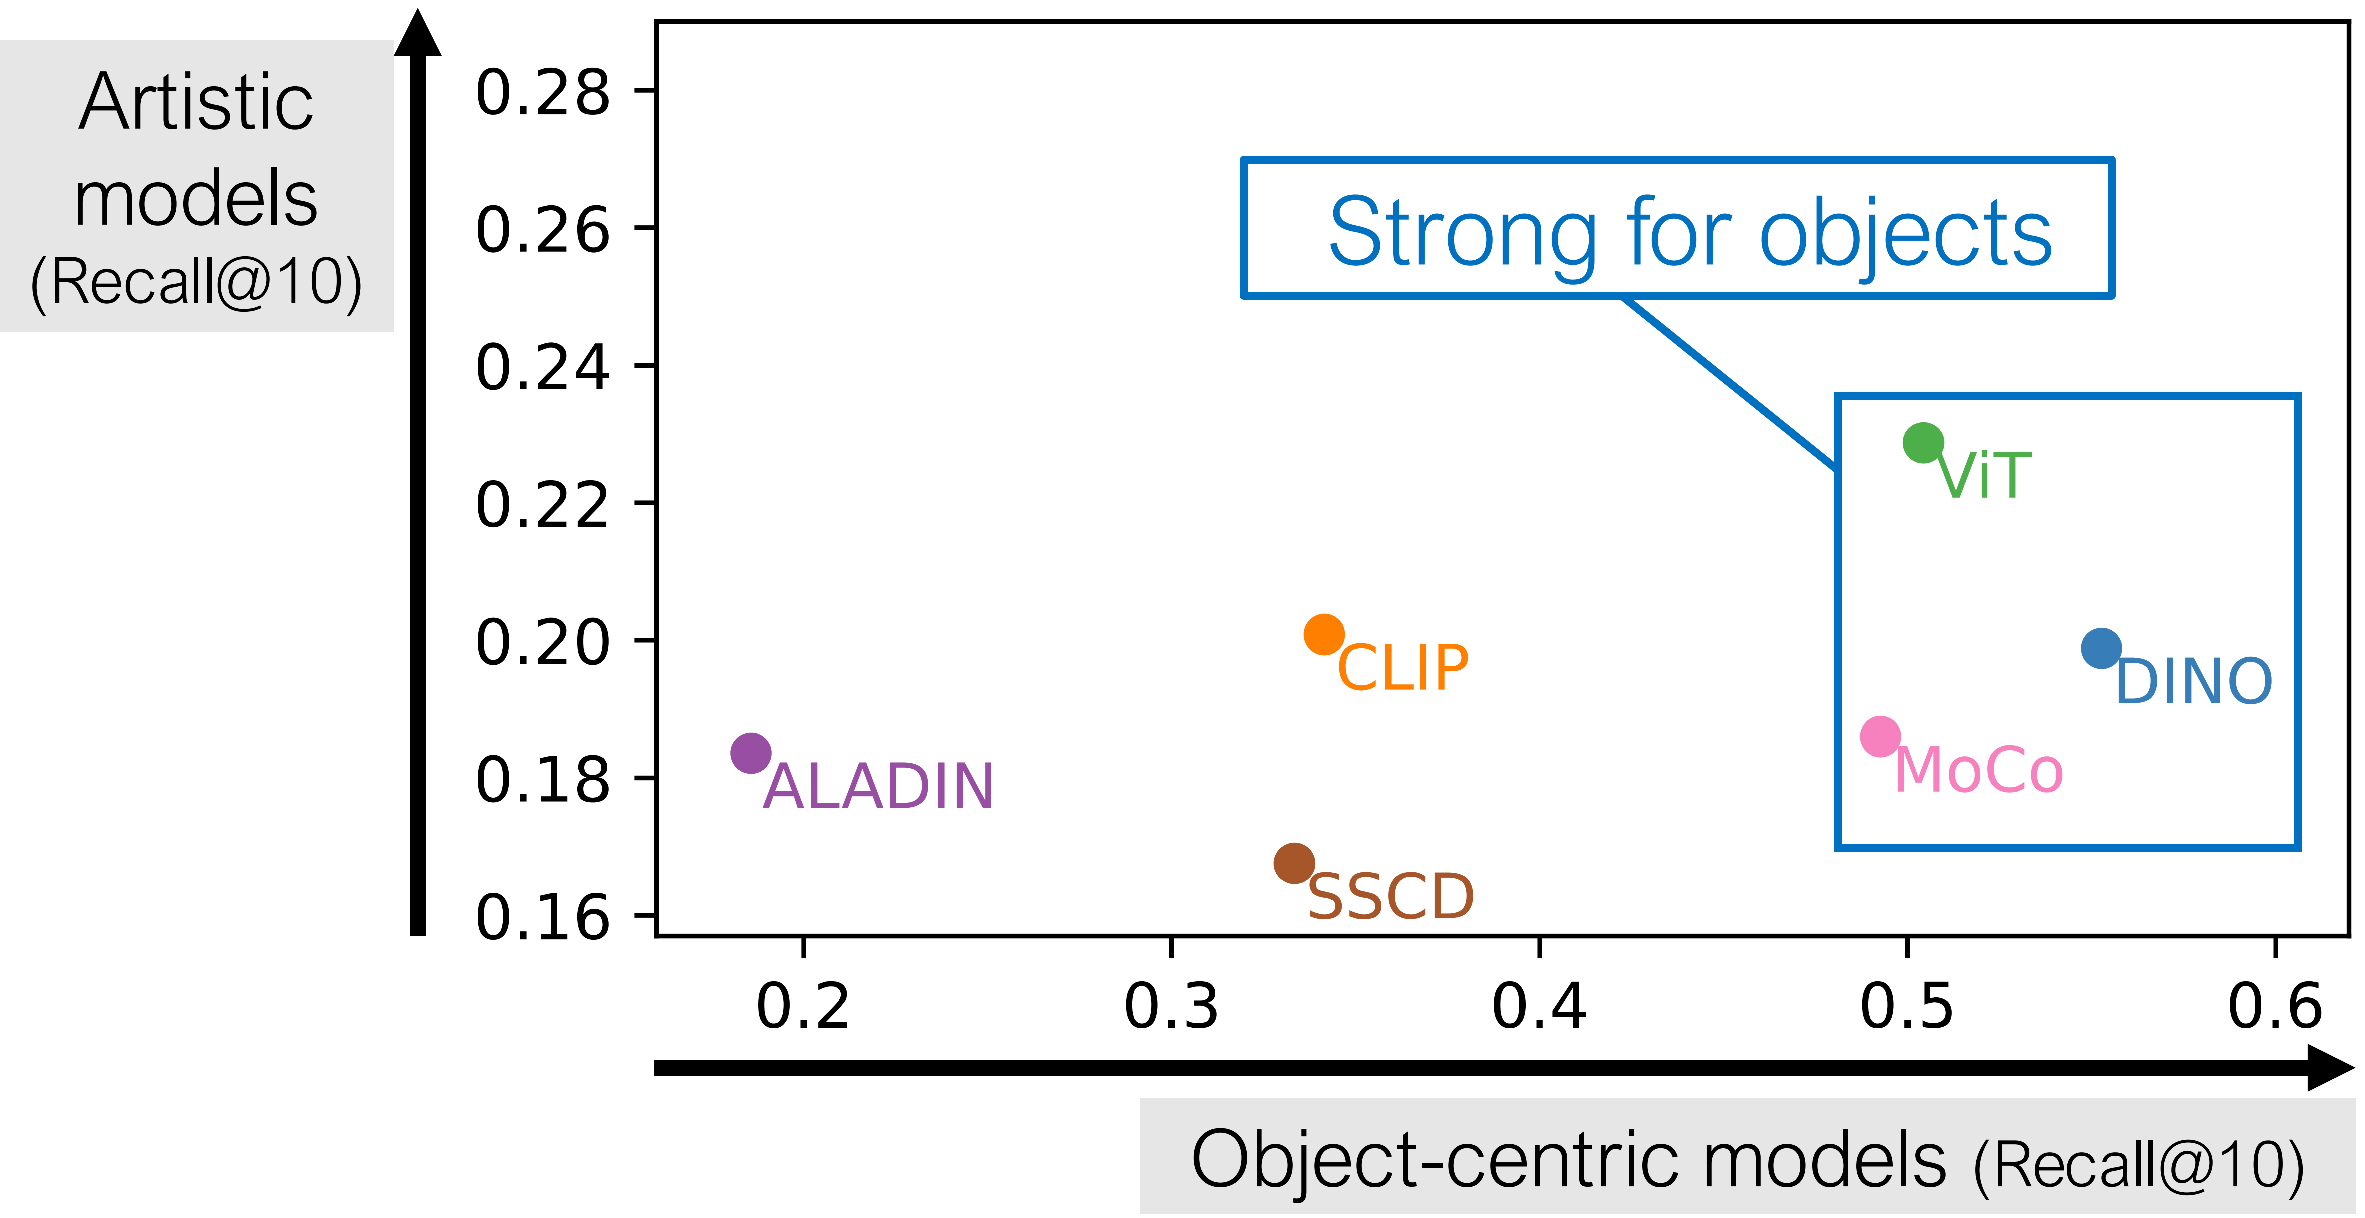

Quantitative Results

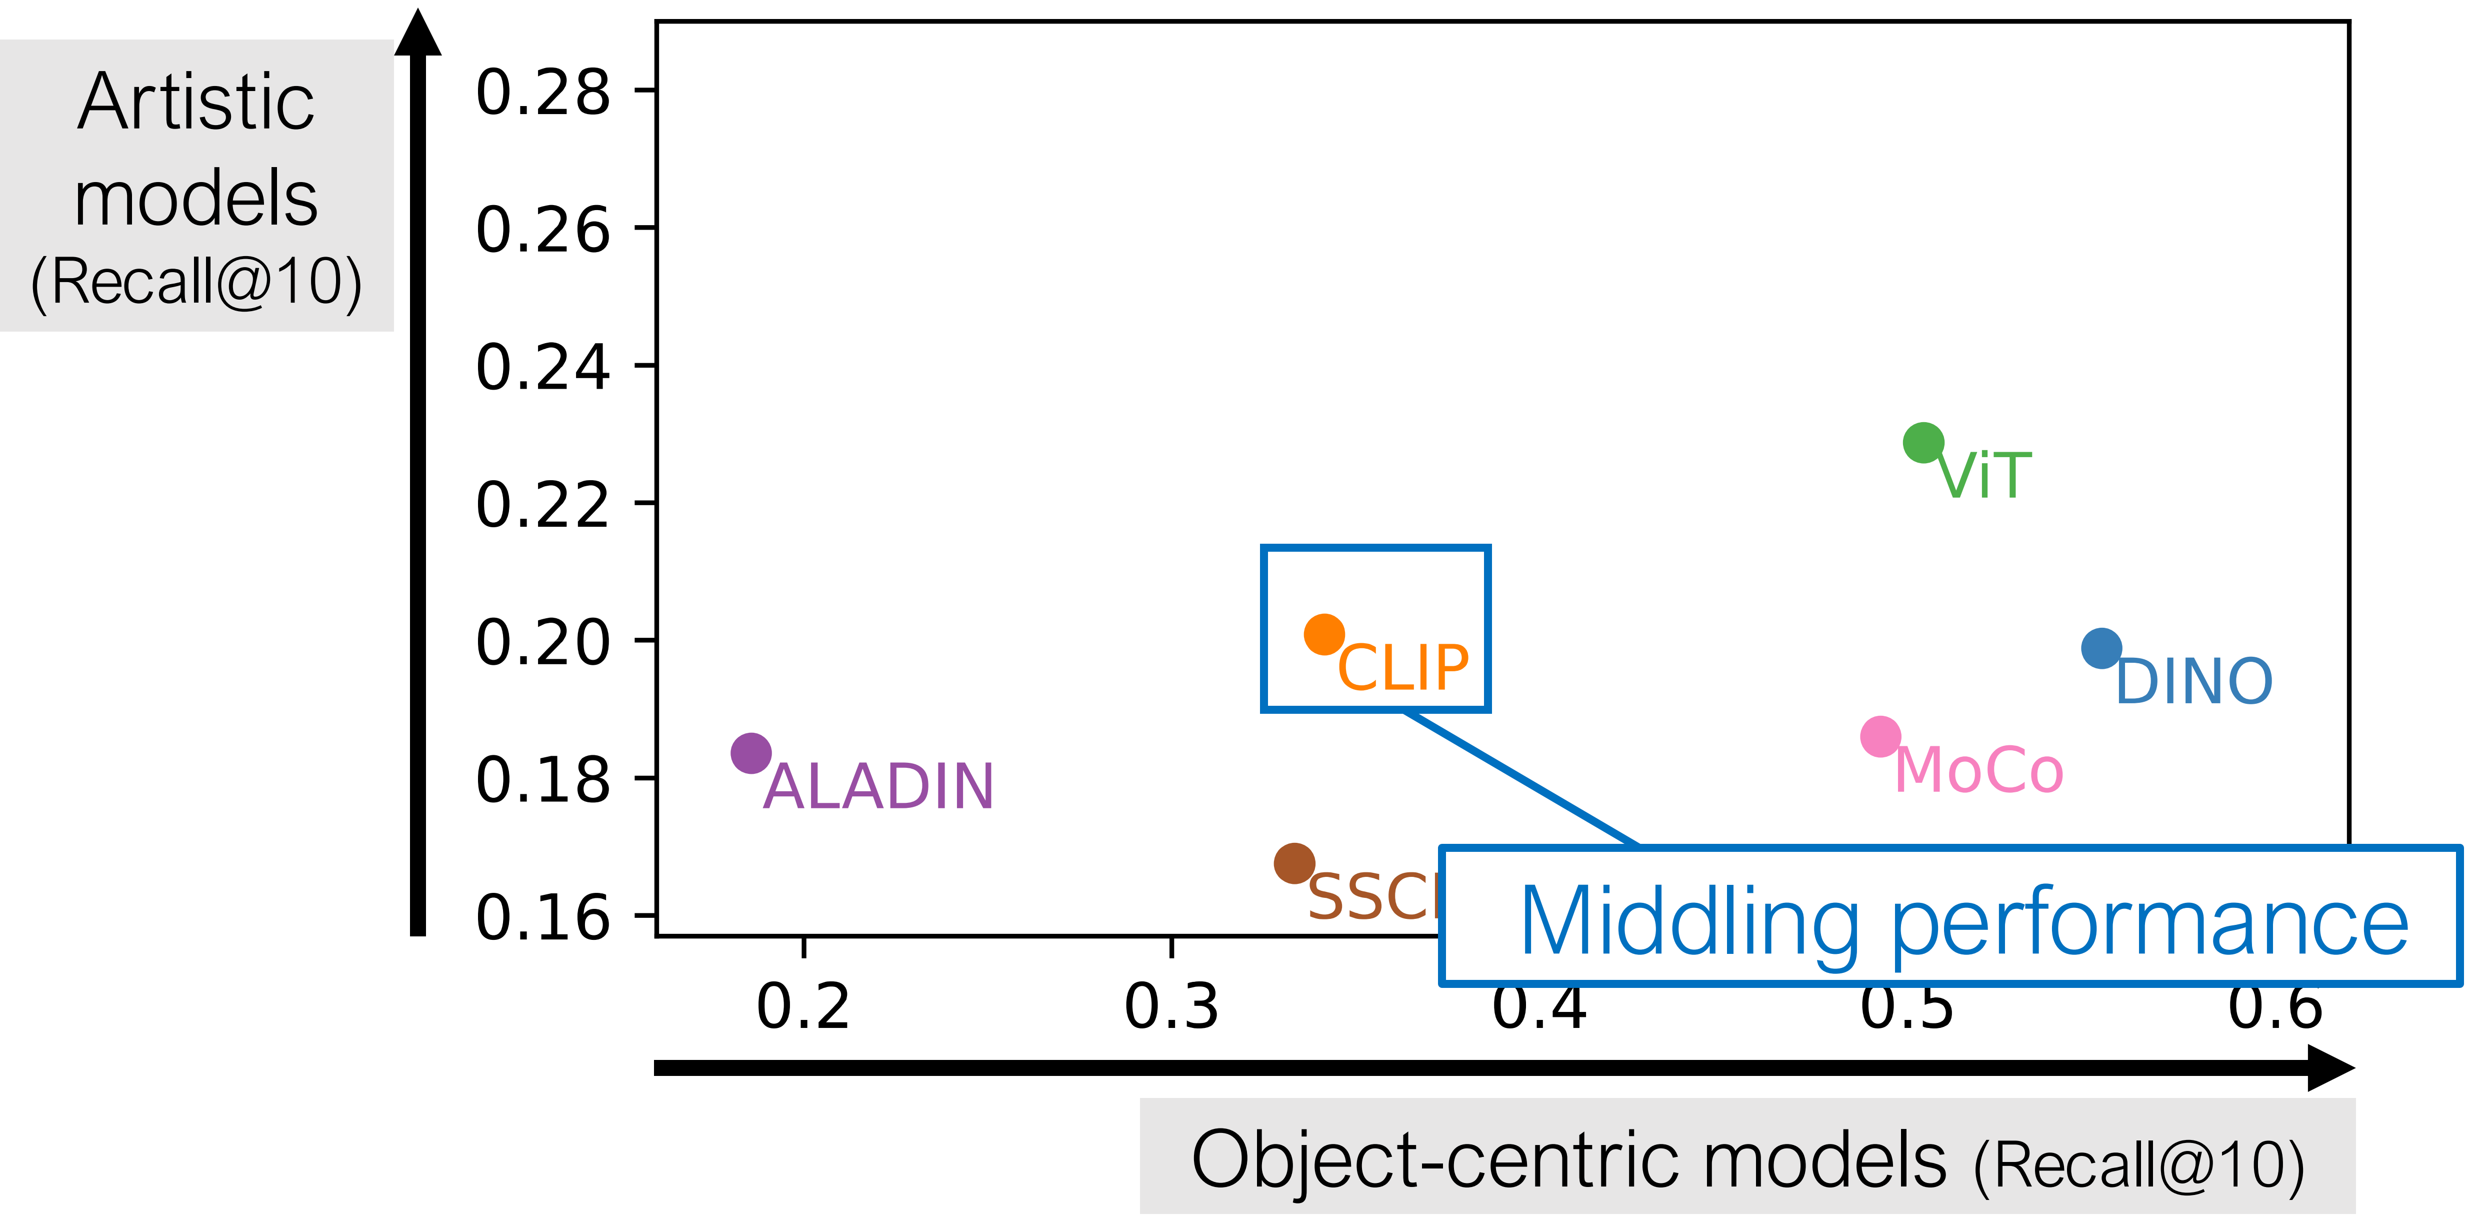

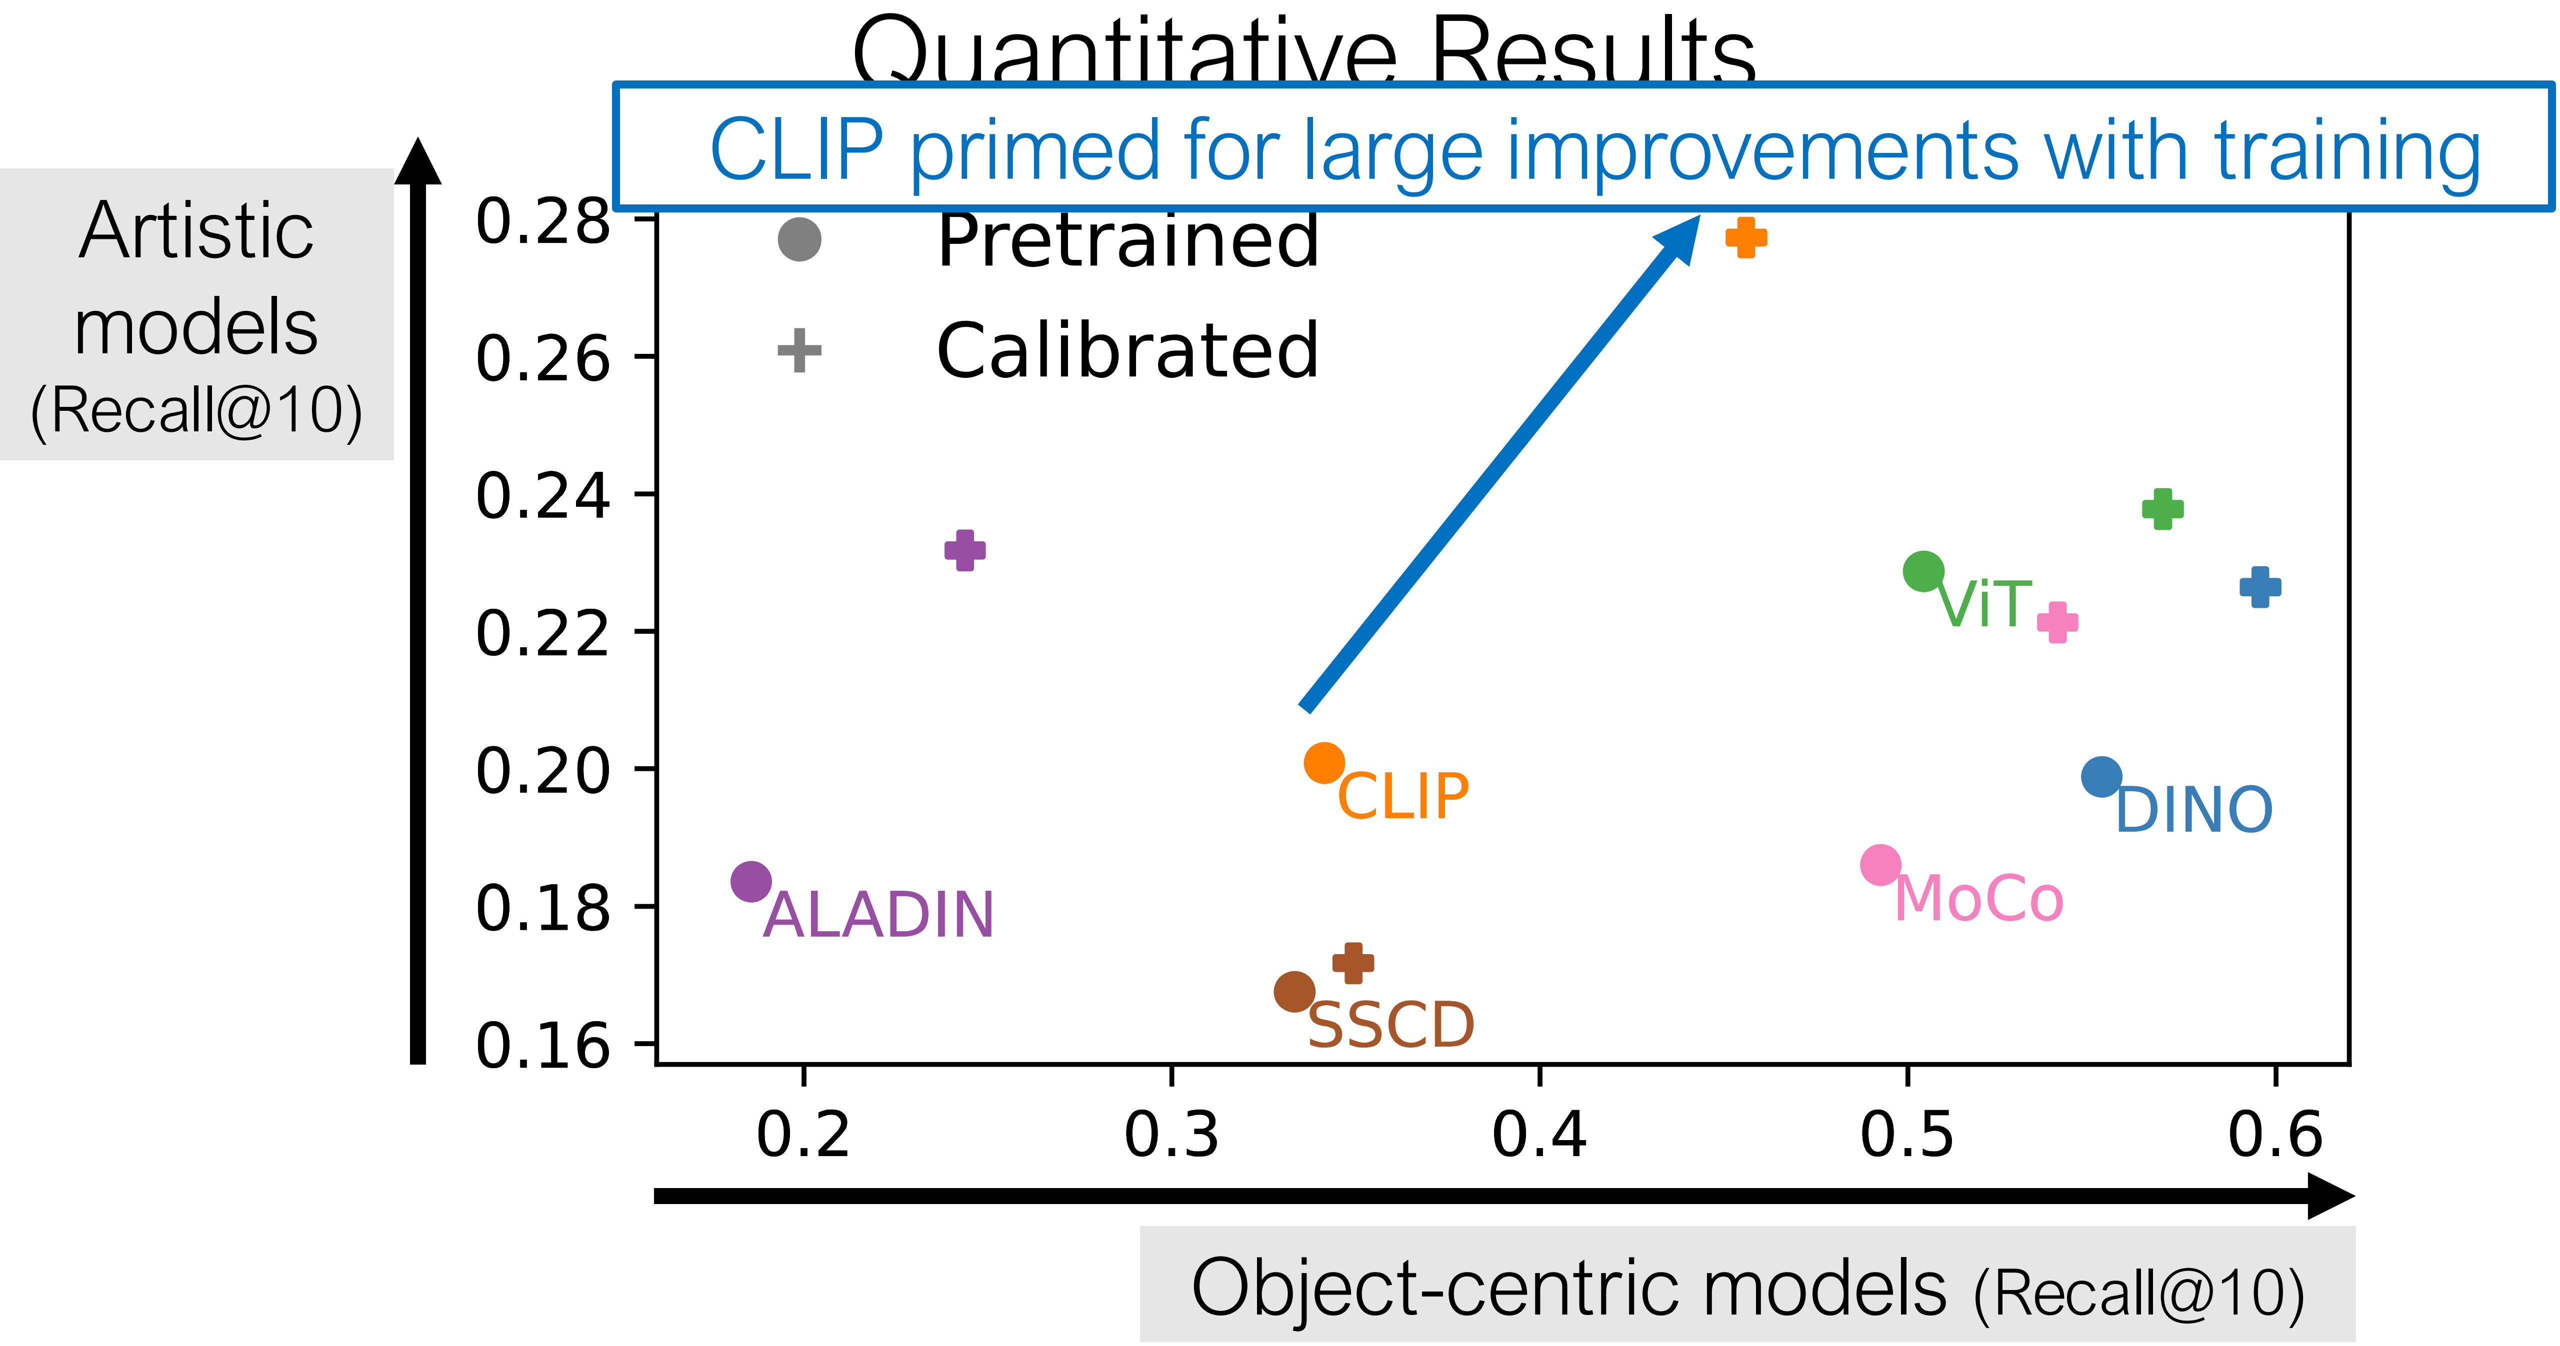
Quantitative Results
CLIP primed for large improvements with training
Artistic models (Recall@10)
Pretrained
Calibrated
0.28
0.26
0.24
0.22
0.20
0.18
0.16
ViT
CLIP
DINO
MoCo
ALADIN
SSCD
0.2
0.3
0.4
0.5
0.6
Object-centric models (Recall@10)

Limitations

- Pretraining set is ignored
  - LAION-5B has influence on Custom Diffusion examples

- Prior work: “remove” instead of “add”
  - Shapley Value: landmark concepts in economics
    - [Shapley 1953; Feldman & Zhang 2020]
    - Train on random subsets; analyze population of models
  - Influence functions
    - [Koh & Liang 2017, Schioppa et al. 2022, Park et al. 2023, Georgiev et al. 2023]
    - Linear approximation

- Evaluating attribution with large training set is challenging!

Random subsets

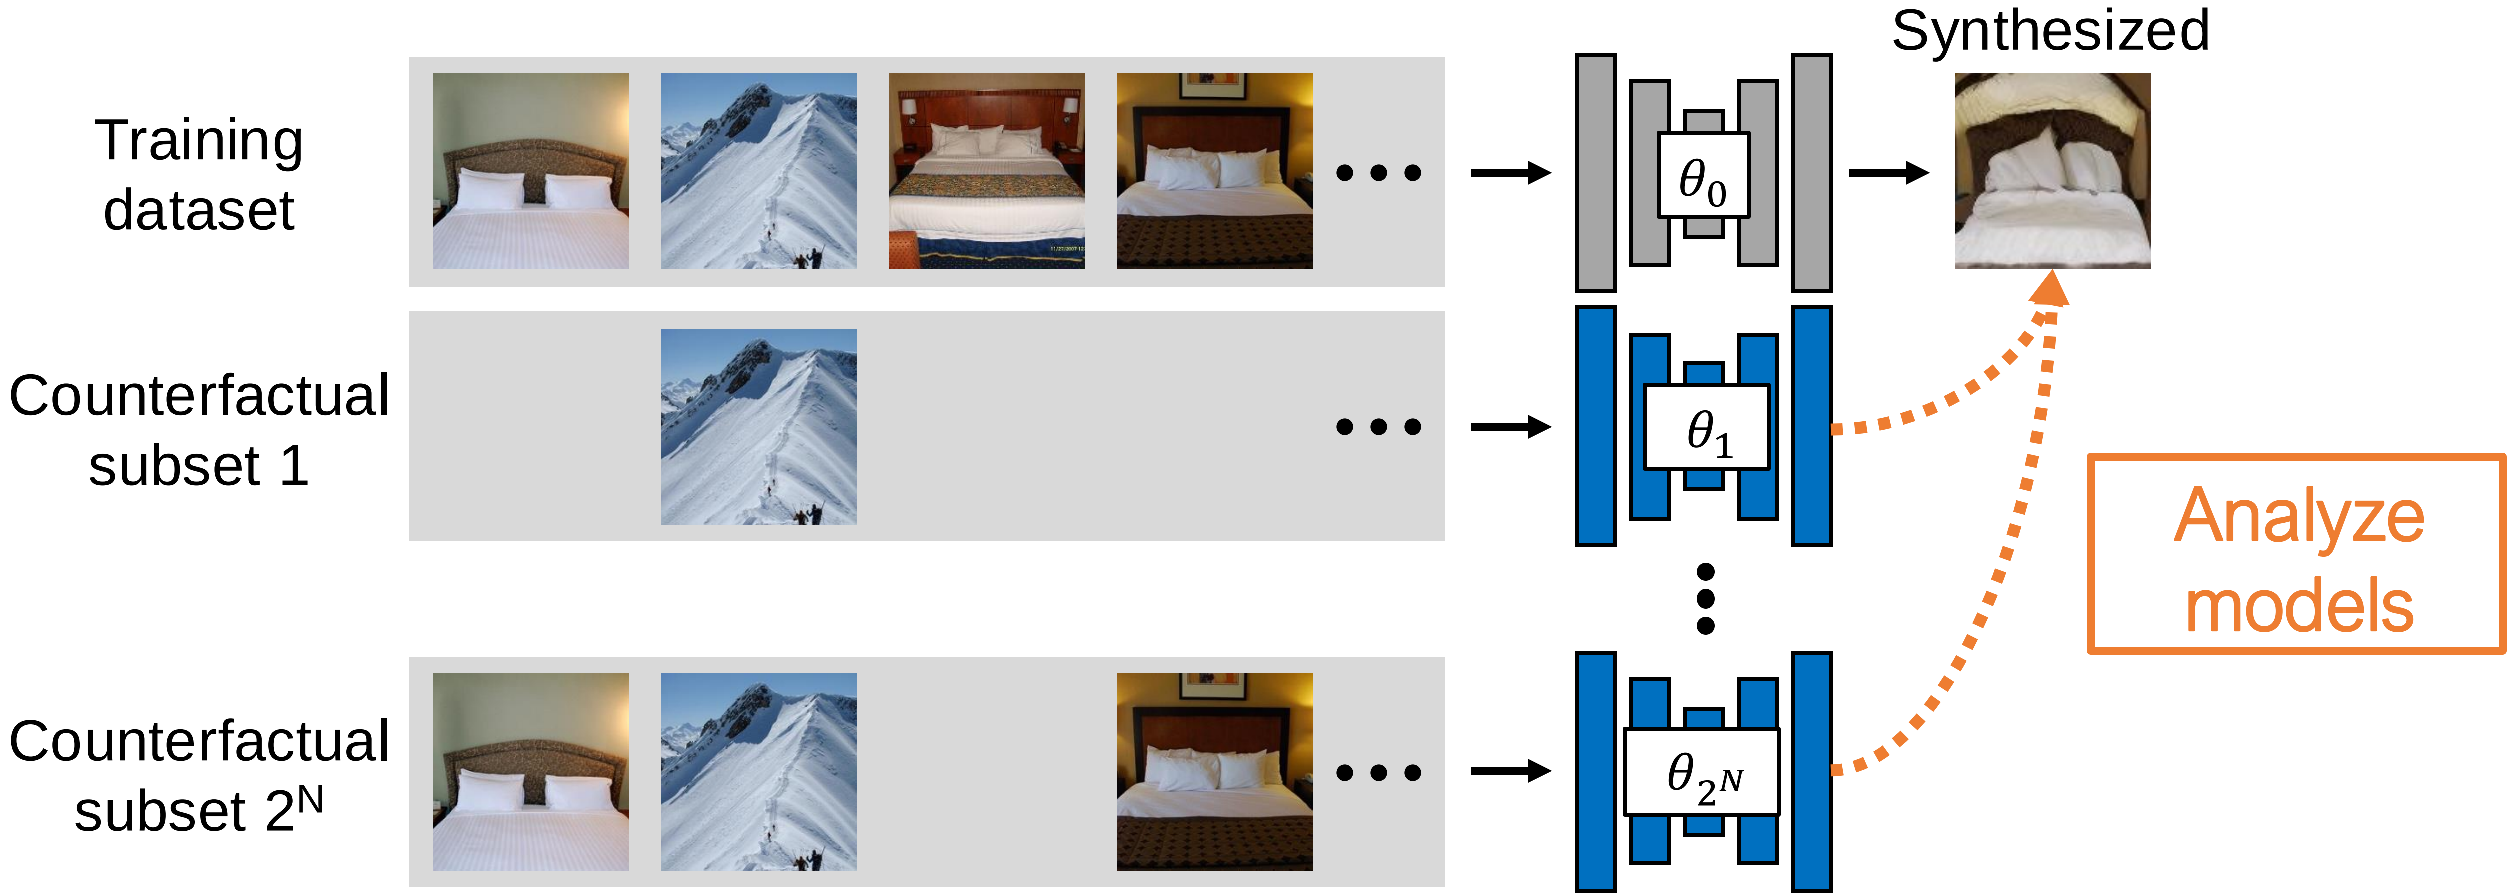

Training 2^N models is too expensive

c.f. Feldman & Zhang. What Neural Networks Memorize and Why. NeurIPS 2020.

Leave-one-out

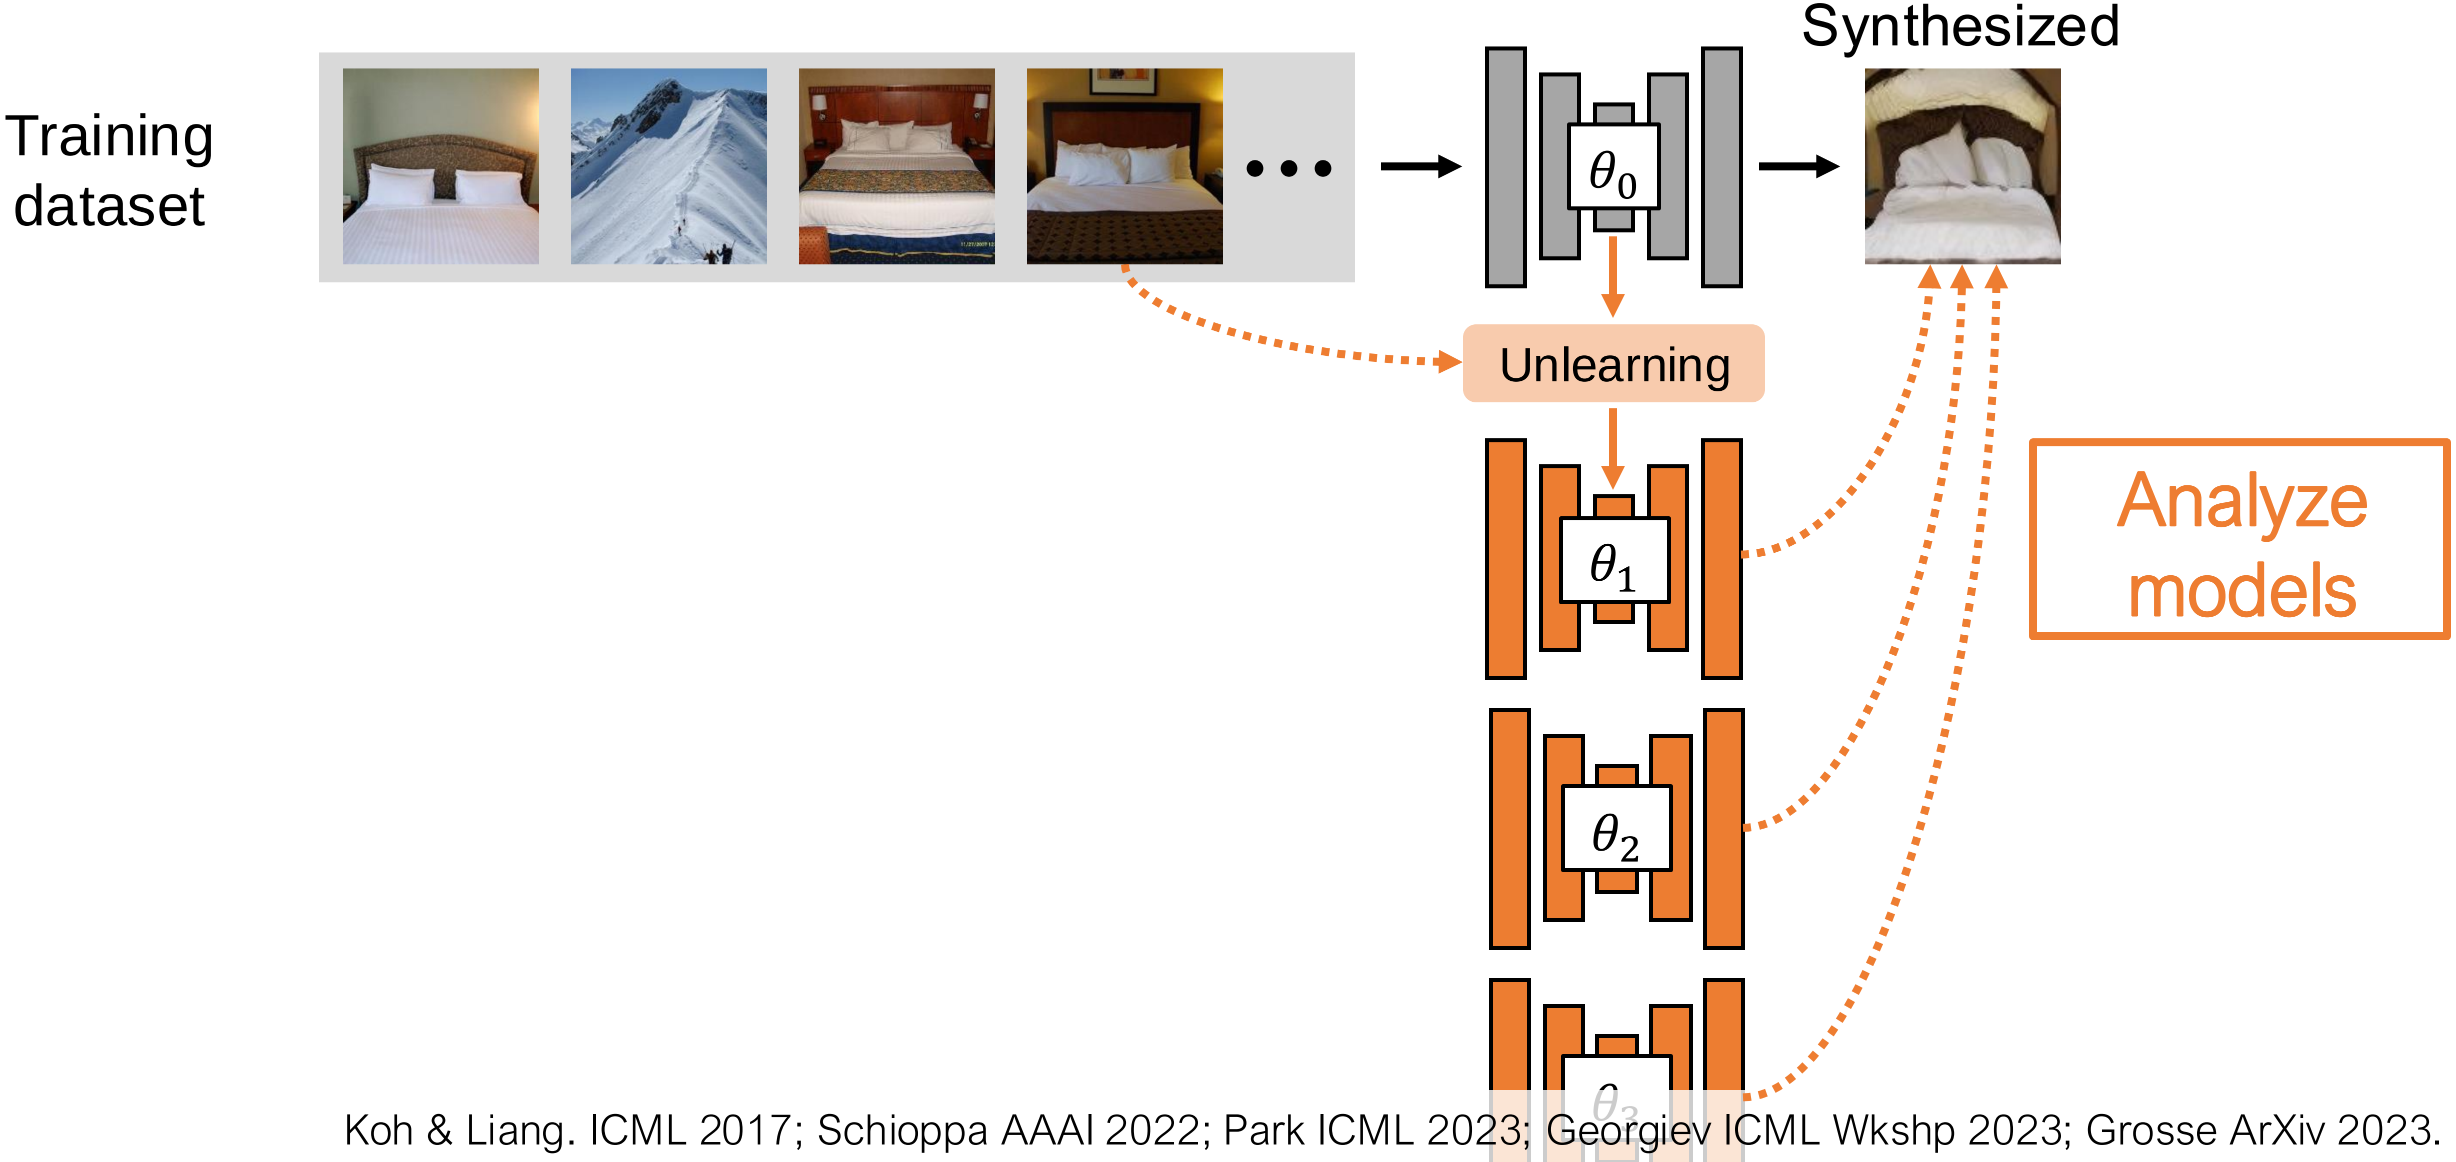

Koh & Liang. ICML 2017; Schioppa AAAI 2022; Park ICML 2023; Georgiev ICML Wkshp 2023; Grosse ArXiv 2023.

Leave-one-out

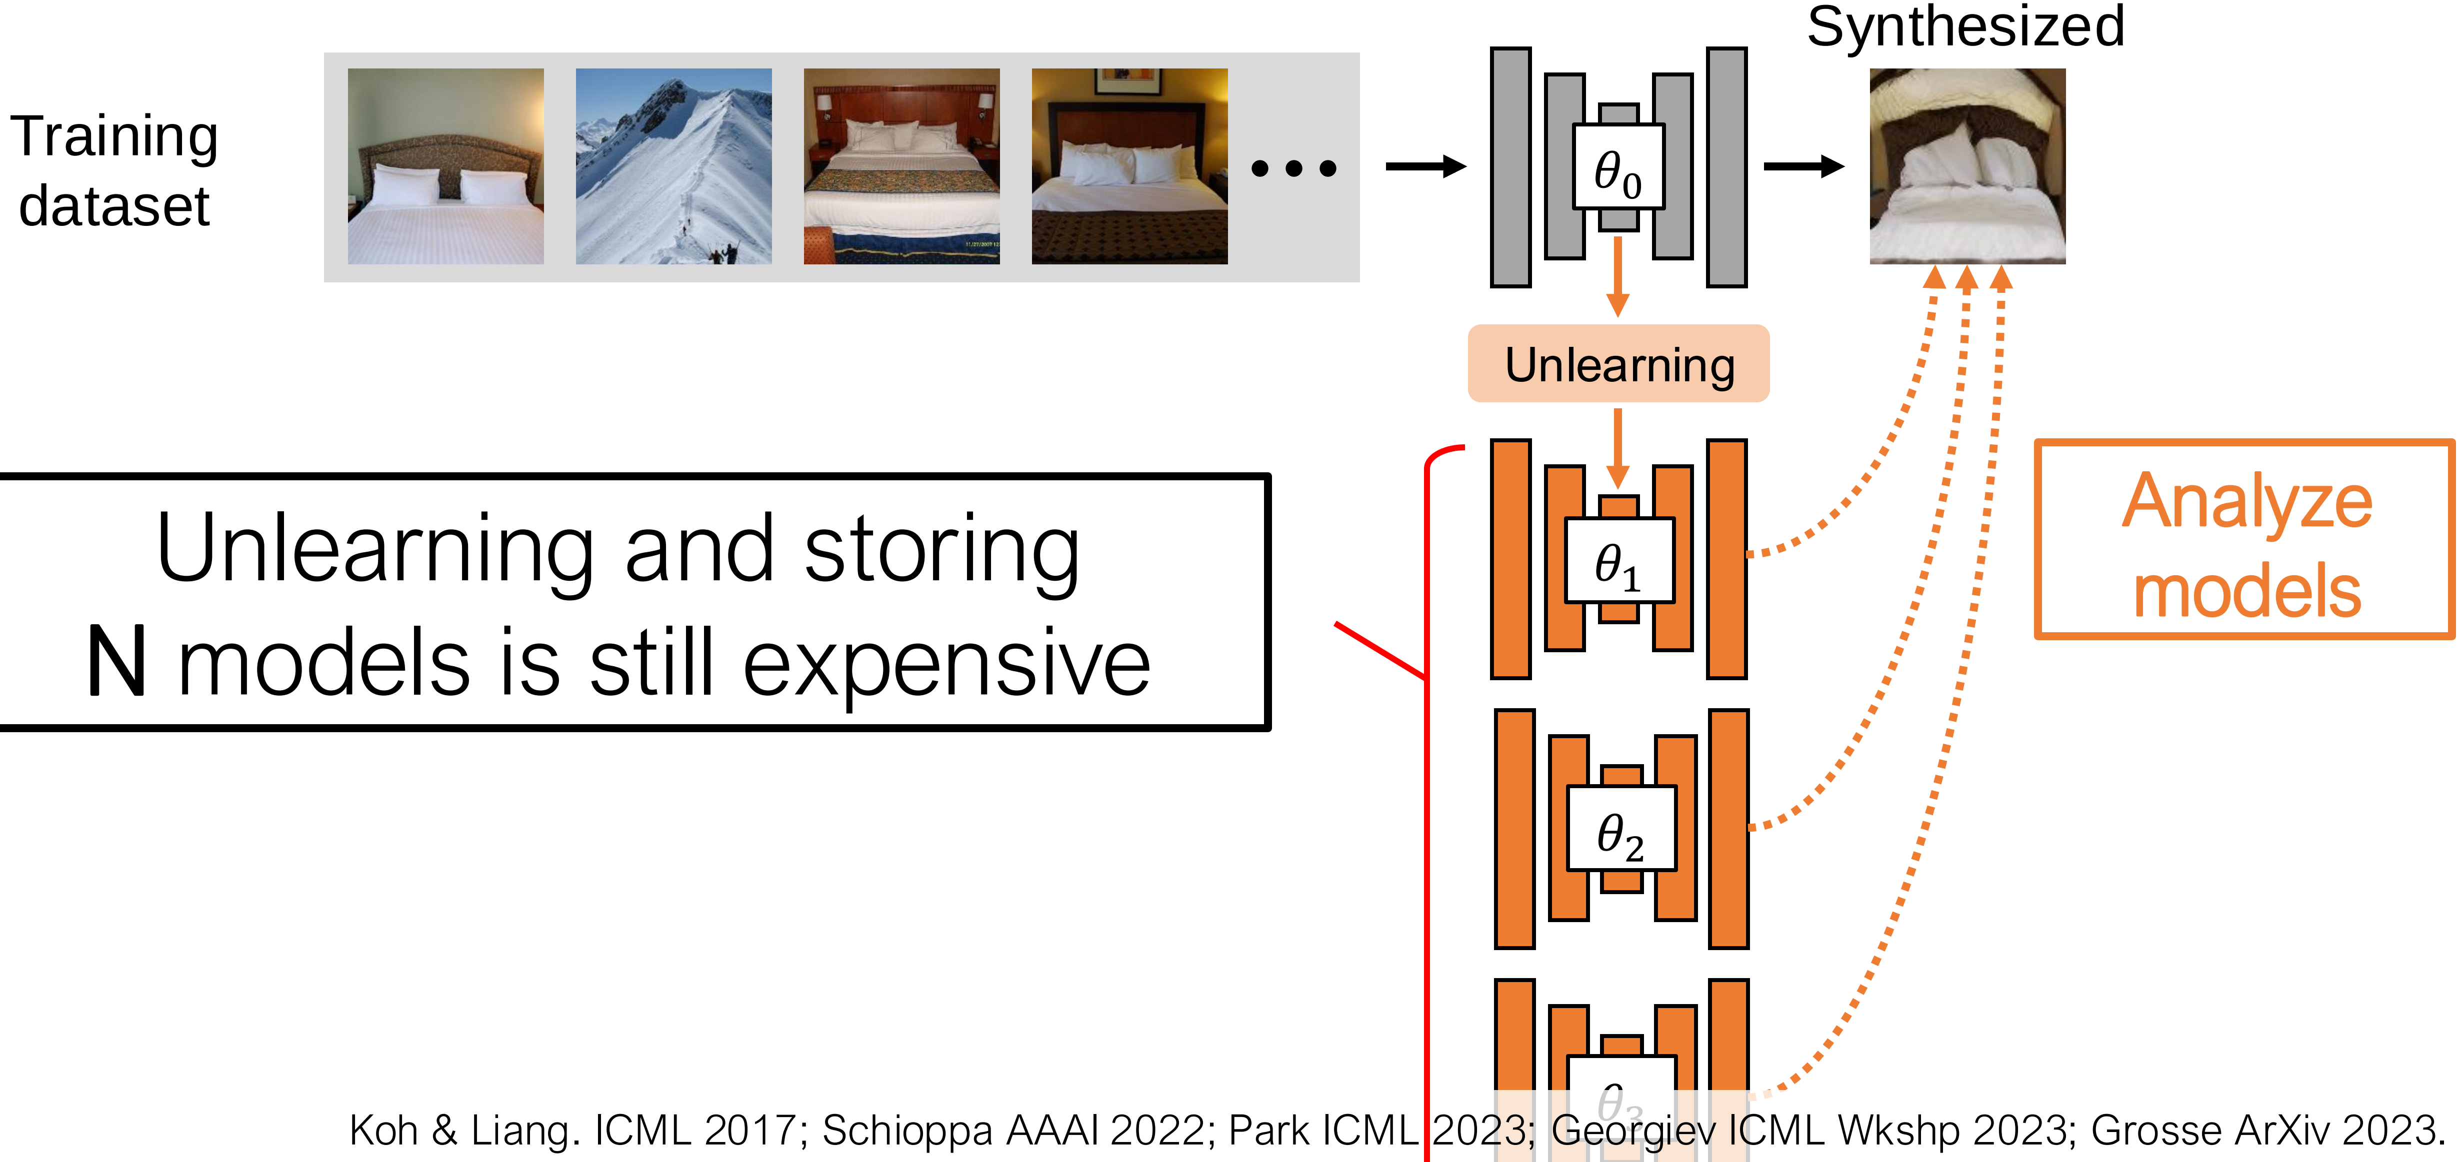

Unlearning and storing **N** models is still expensive

Koh & Liang. ICML 2017; Schioppa AAAI 2022; Park ICML 2023; Georgiev ICML Wkshp 2023; Grosse ArXiv 2023.

Change One Thing at a Time

(Remove one Test Image)

Attribution by Unlearning (AbU)

Training dataset

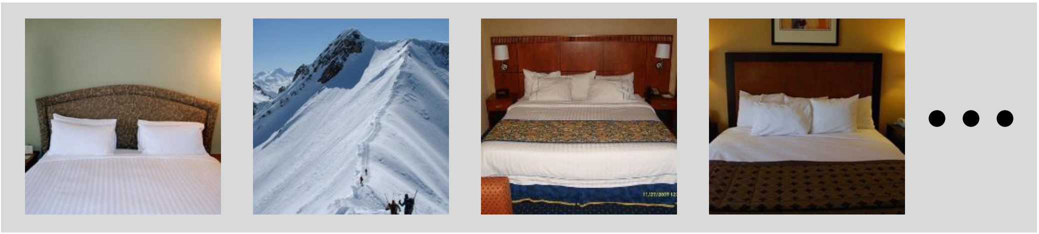

Synthesized

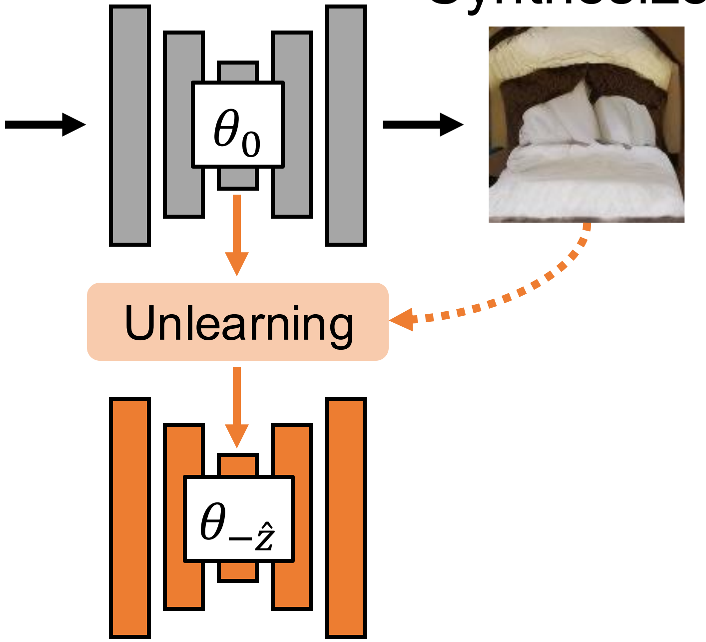

Unlearning procedure

Maximize loss on synthesized point

Minimize loss on original dataset

$$\mathcal{L}^{\hat{\mathbf{z}}}_{\text{unlearn}}(\theta) = -\mathcal{L}(\hat{\mathbf{z}}, \theta)$$

Approximated by EWC

$$\frac{N}{2}(\theta - \theta_0)^T F(\theta - \theta_0)$$

Kirkpatrick. Overcoming catastrophic forgetting. PNAS 2017.

Fisher Information:

$$F \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{z}\sim p_{\text{data}}(\mathbf{z})}\left[\nabla_\theta \log p_\theta(\mathbf{z}) \nabla_\theta \log p_\theta(\mathbf{z})^T\right]$$

Attribution by Unlearning (AbU)

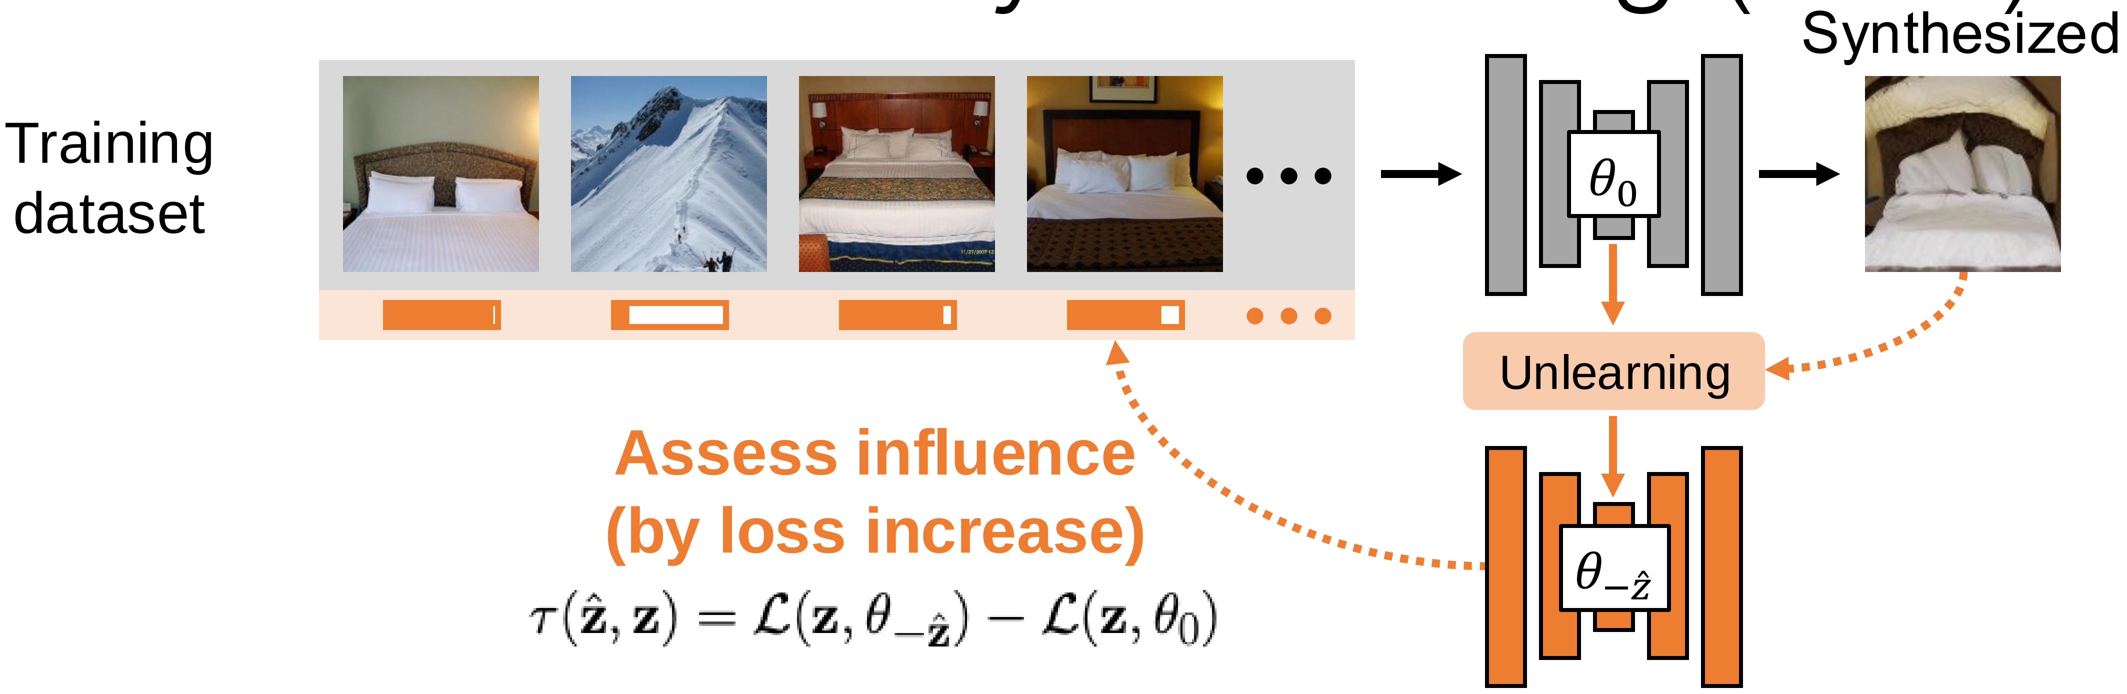

Unlearning only once

How do we evaluate attribution?

Counterfactual evaluation

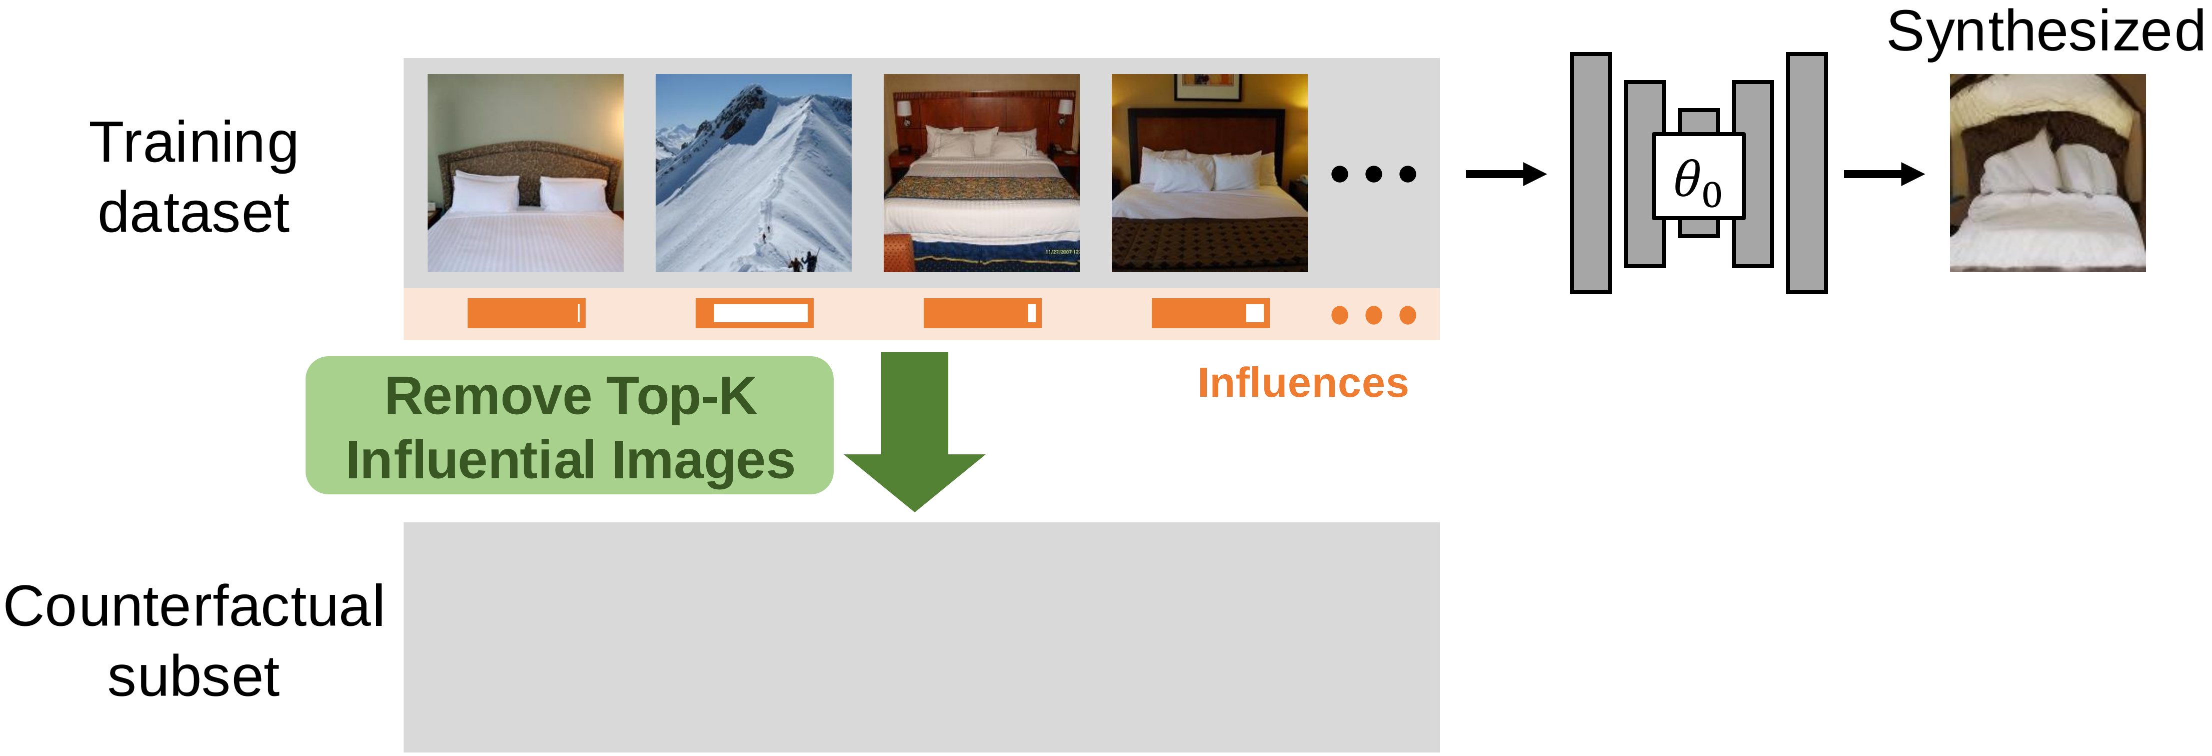

c.f. K. Georgiev, et al. How Training Data Guides Diffusion Models. In ArXiv, 2023.

Counterfactual evaluation

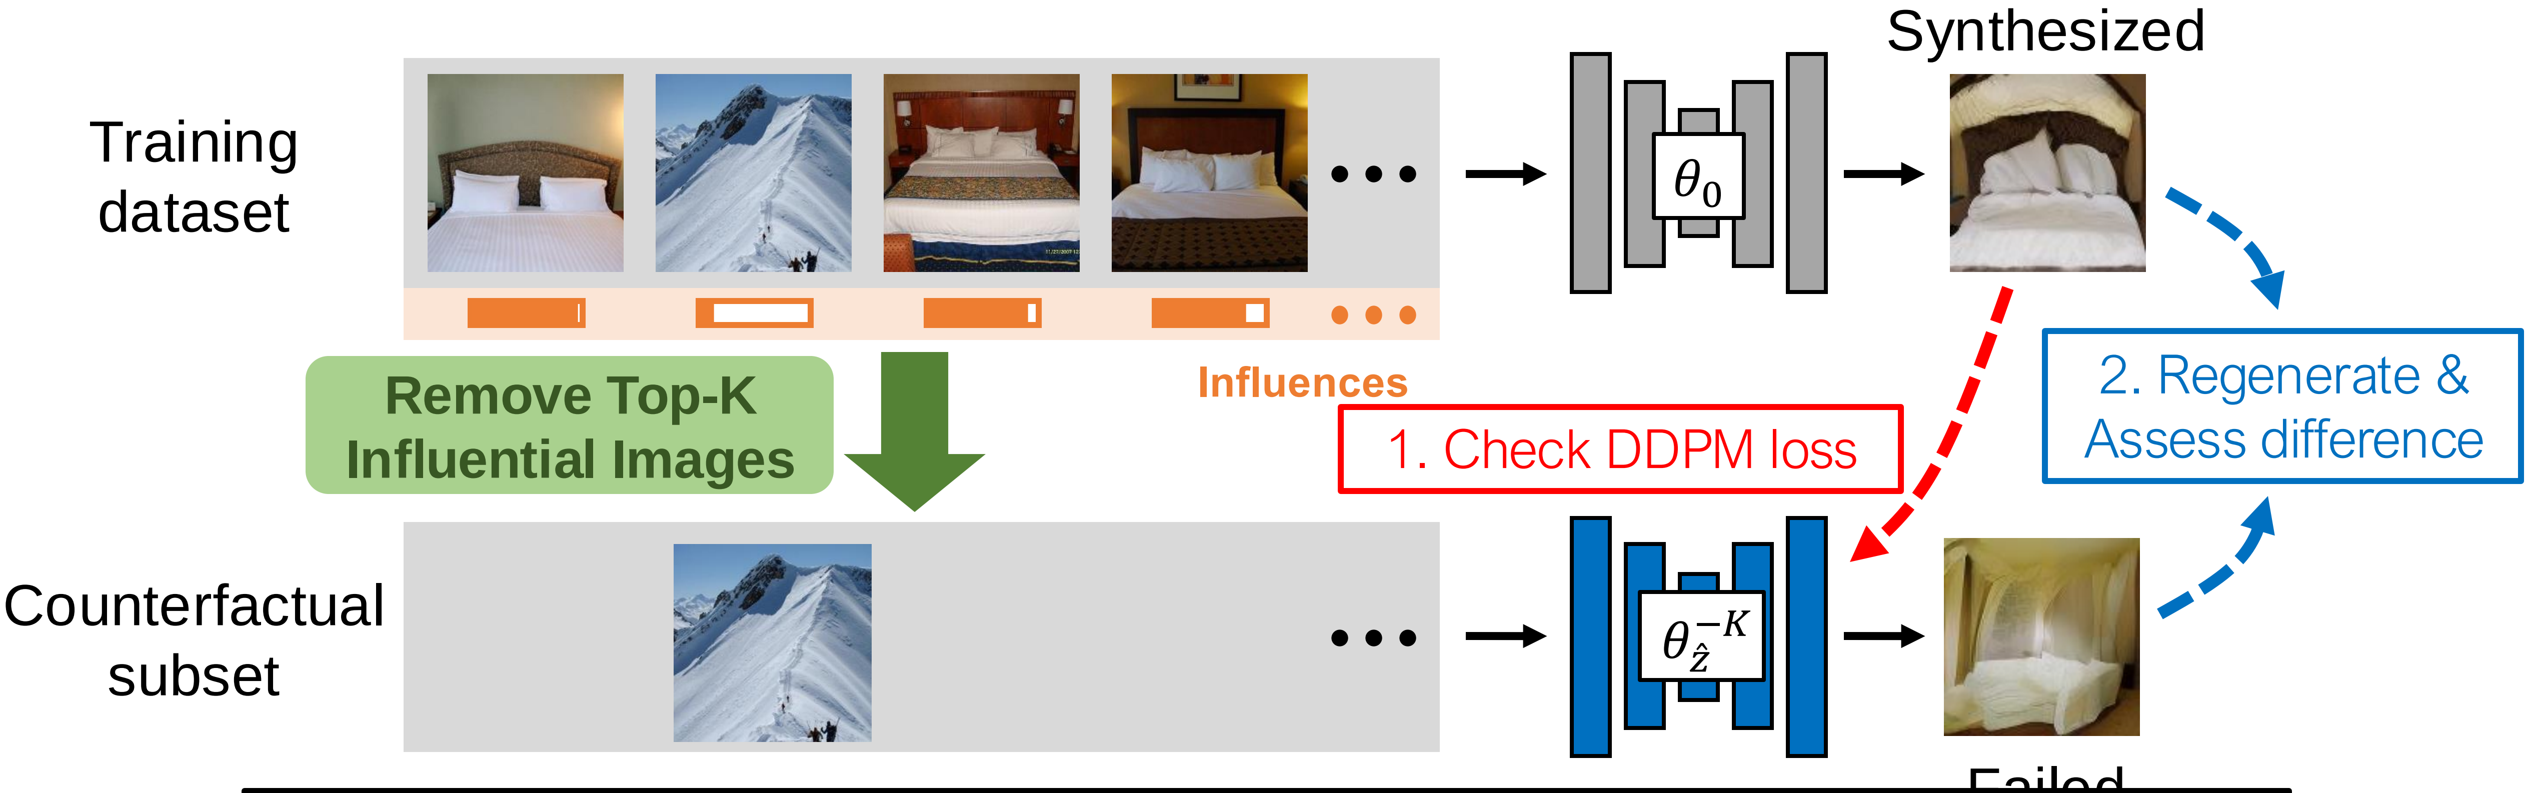

Expensive evaluation…

…but let's do it! (for modest sizes)

rem tion

c.f. K. Georgiev, et al. How Training Data Guides Diffusion Models. In ArXiv, 2023.

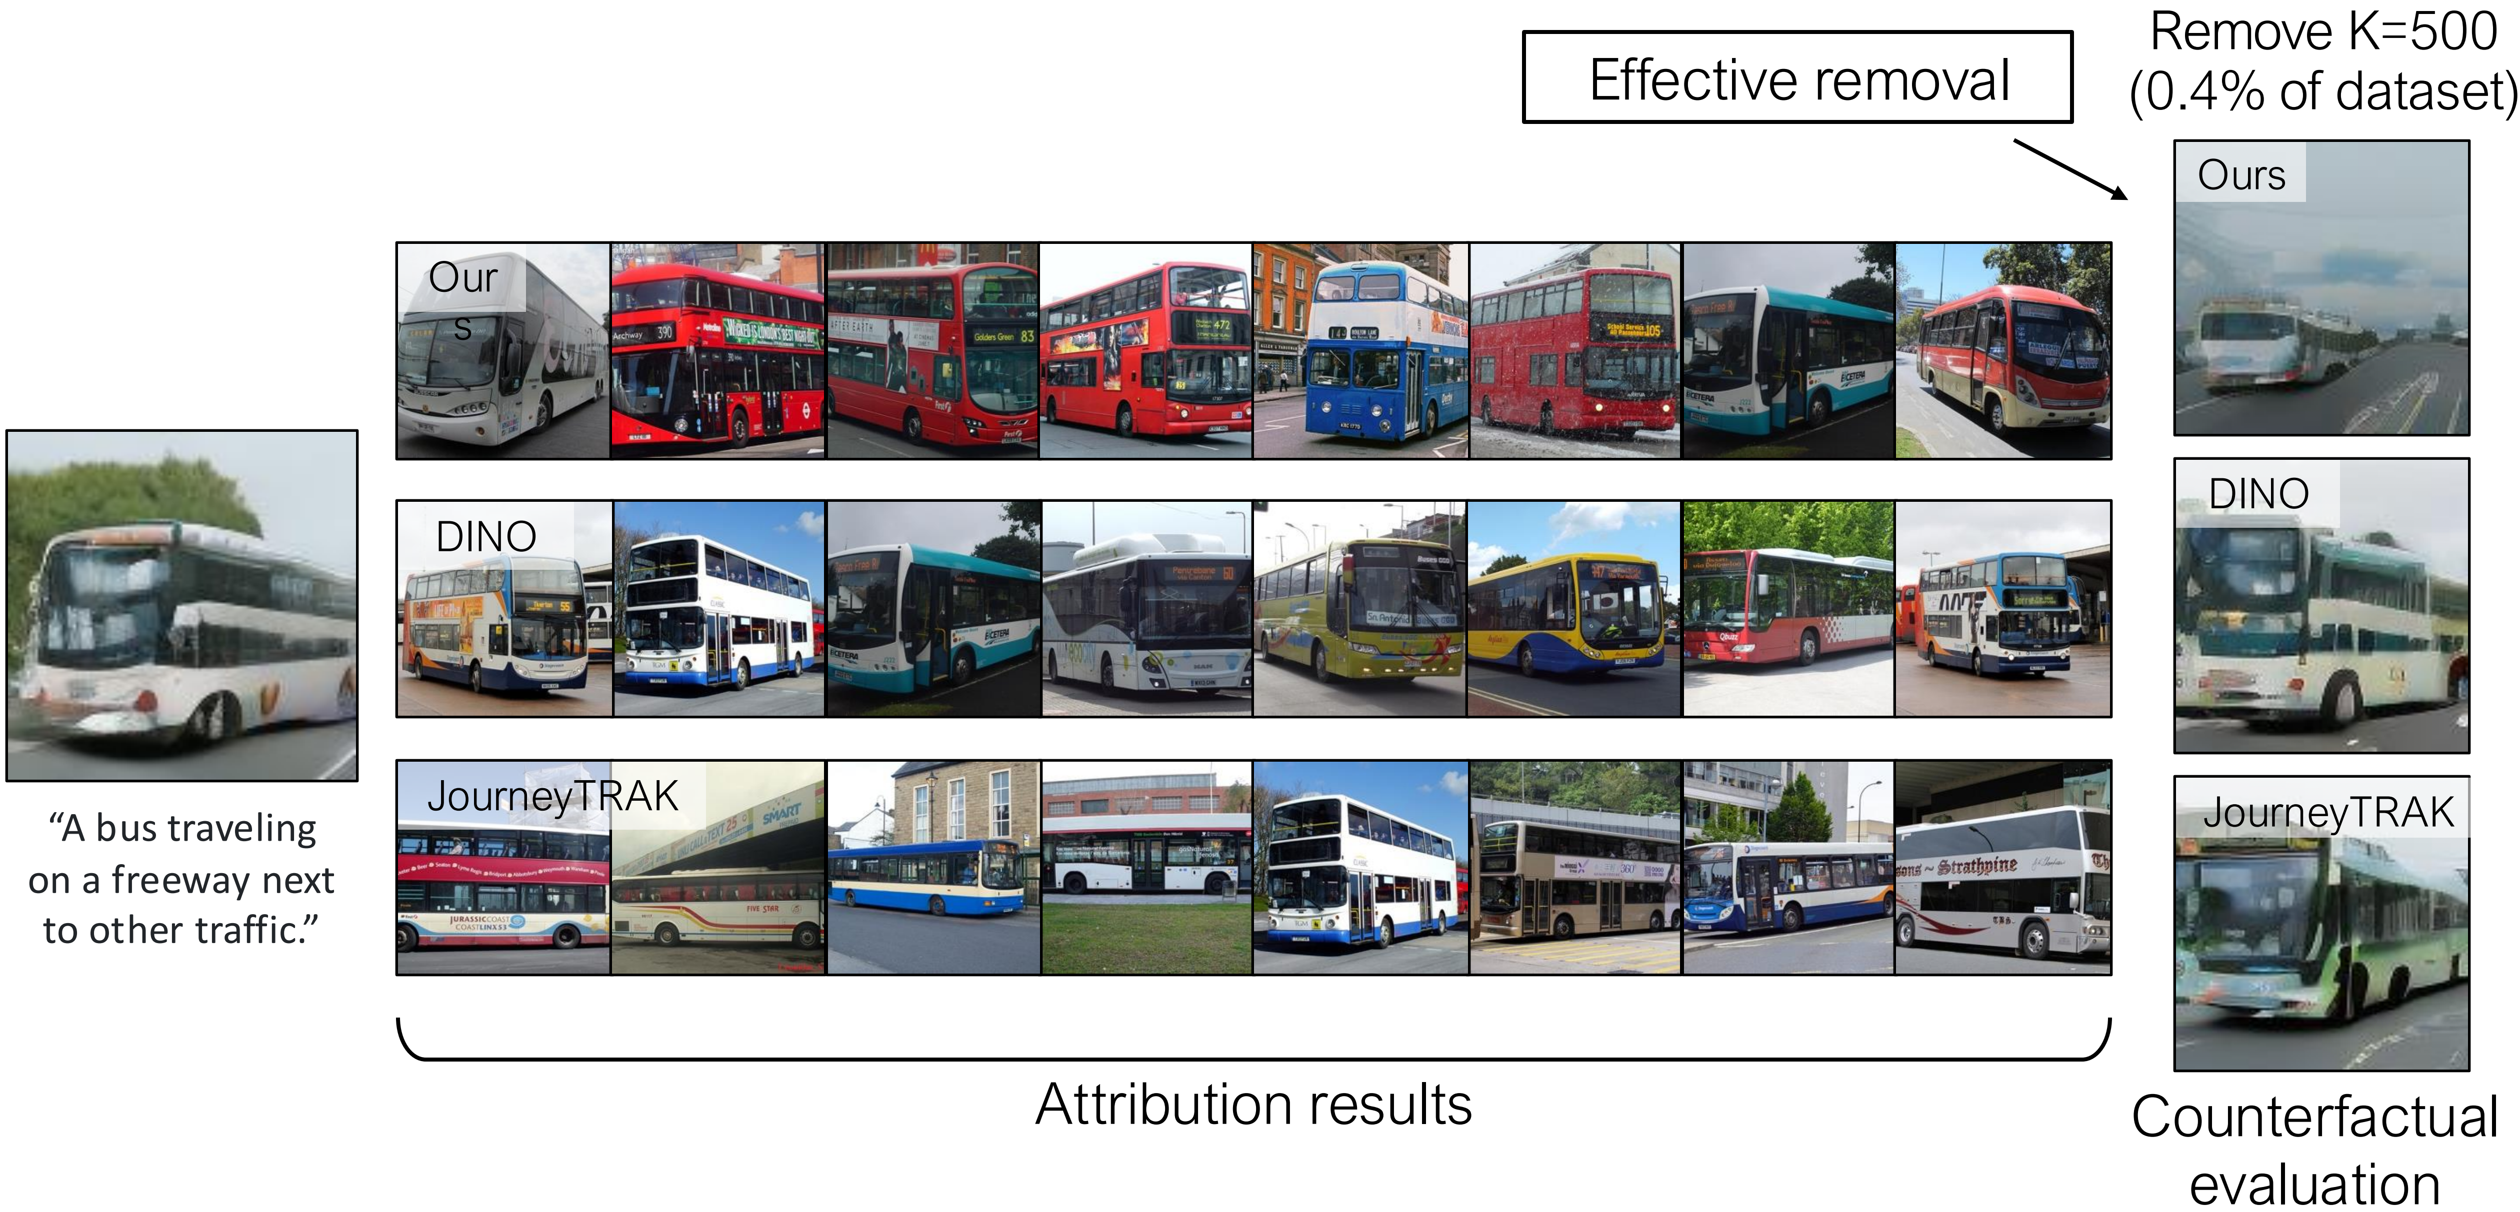

c.f. K. Georgiev, et al. How Training Data Guides Diffusion Models. In ArXiv, 2023.

MS-COCO results

"A man in a blue coat skiing through a snowy field."

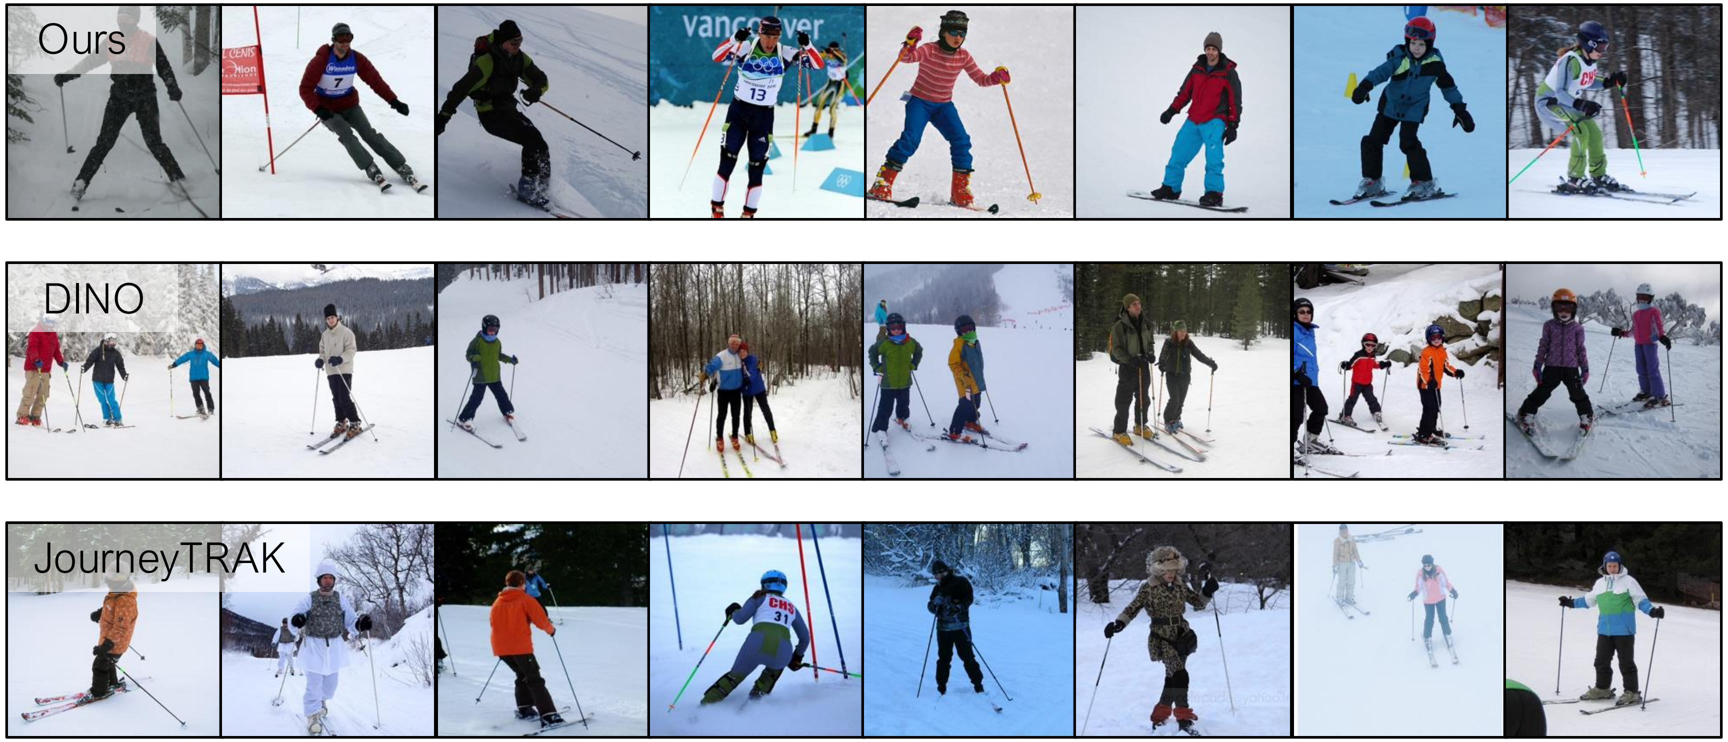

Attribution results

Remove K=500 (0.4% of dataset)

Counterfactual evaluation

c.f. K. Georgiev, et al. How Training Data Guides Diffusion Models. In ArXiv, 2023.

Loss Change vs. K Images Removed
DDPM Loss Change
0.03
0.02
0.01
0.00
random
Pixel
CLIP Text
DINOv2
CLIP
CLIP (AbC)
DINO
DINO (AbC)
TRAK
JourneyTRAK
Ours
500
1000
4000
10000
Full dataset
K images removed
Equivalent K Random Images Removed
K Random Images Removed
Full dataset
100000
75000
50000
25000
0
k=500
k=1000
k=4000
Pixel
CLIP Text
DINOv2
CLIP
CLIP (AbC)
DINO
DINO (AbC)
TRAK
JourneyTRAK
Ours
Nearest neighbors
Tuned via Customization
(prev method)
Influence
functions

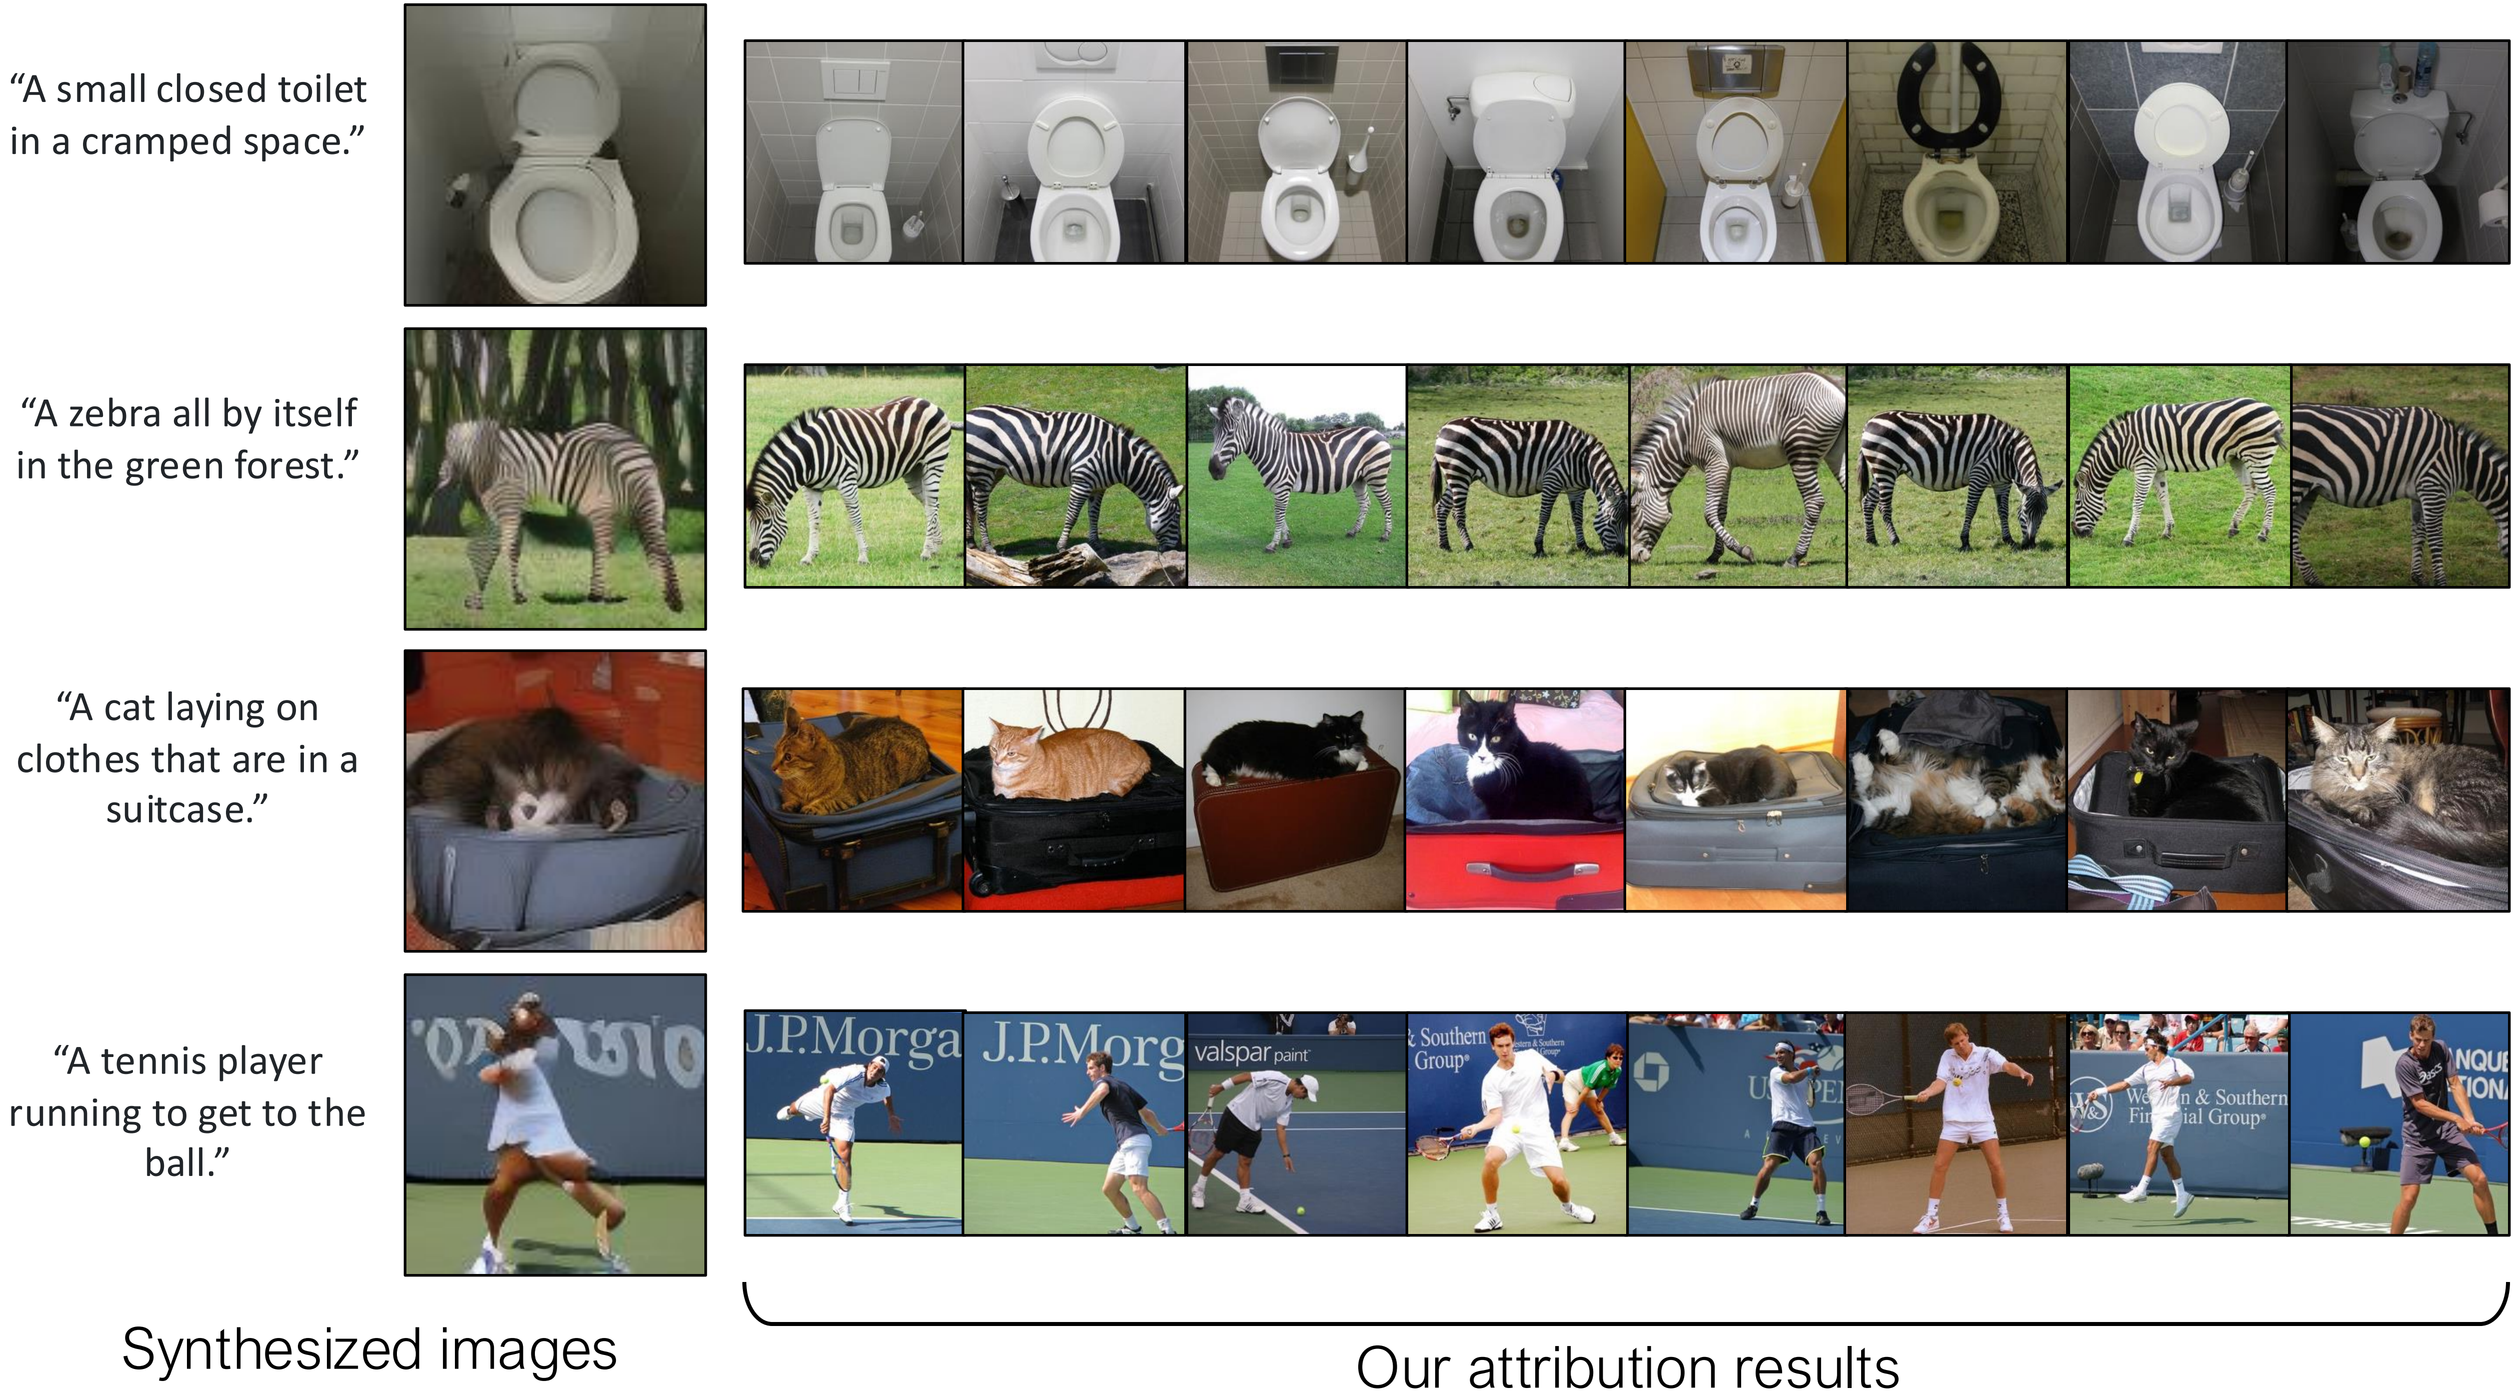
"A small closed toilet in a cramped space."
"A zebra all by itself in the green forest."
"A cat laying on clothes that are in a suitcase."
"A tennis player running to get to the ball."
Synthesized images
Our attribution results

Local attribution

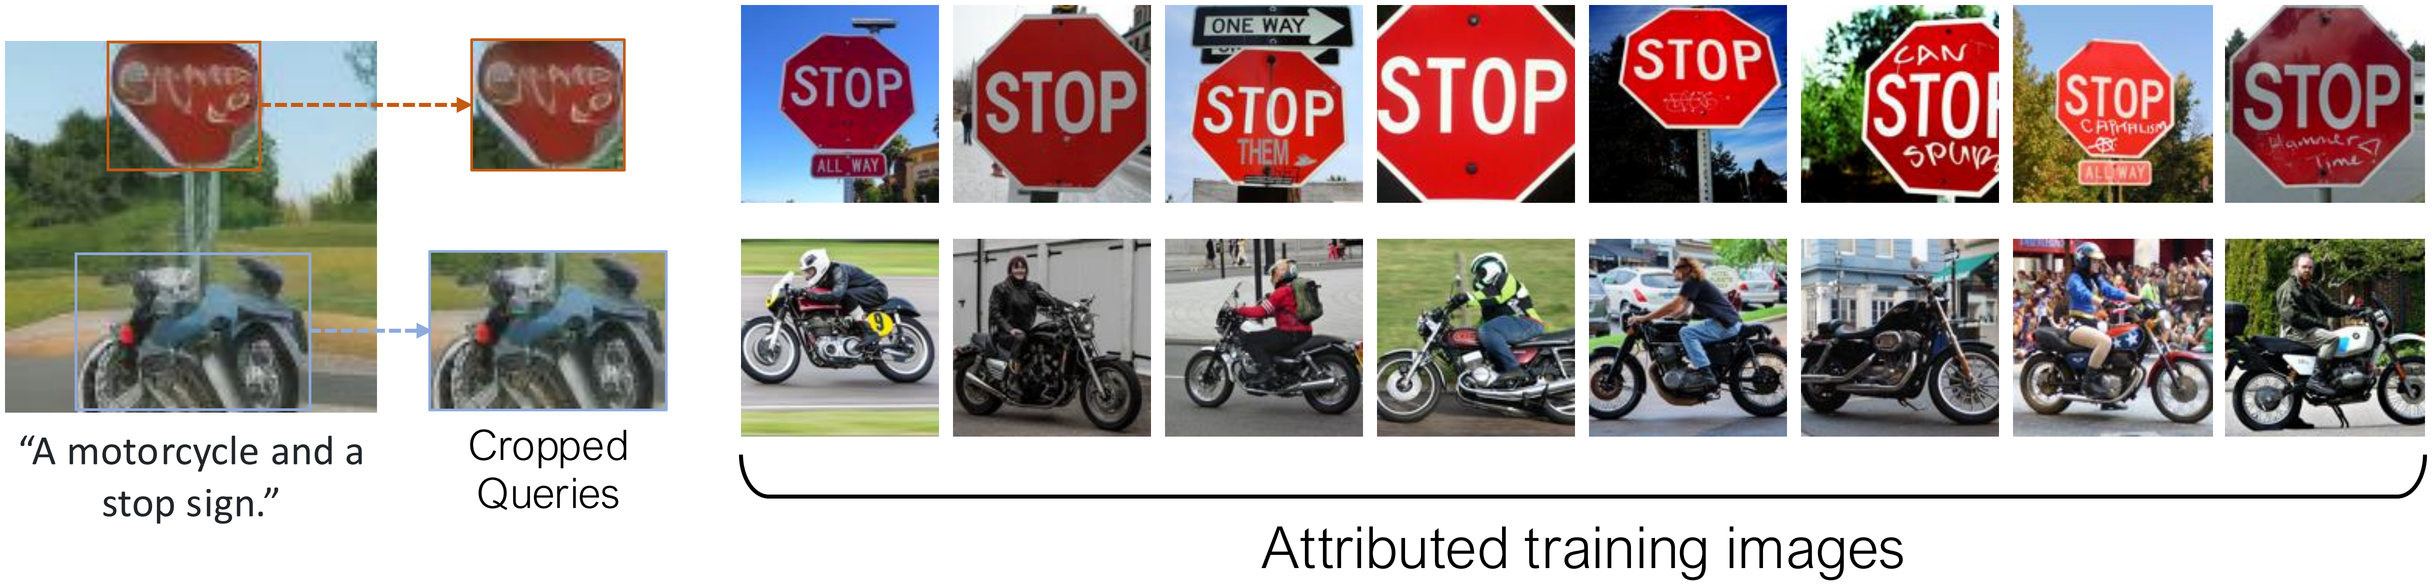

Customized Model Benchmark

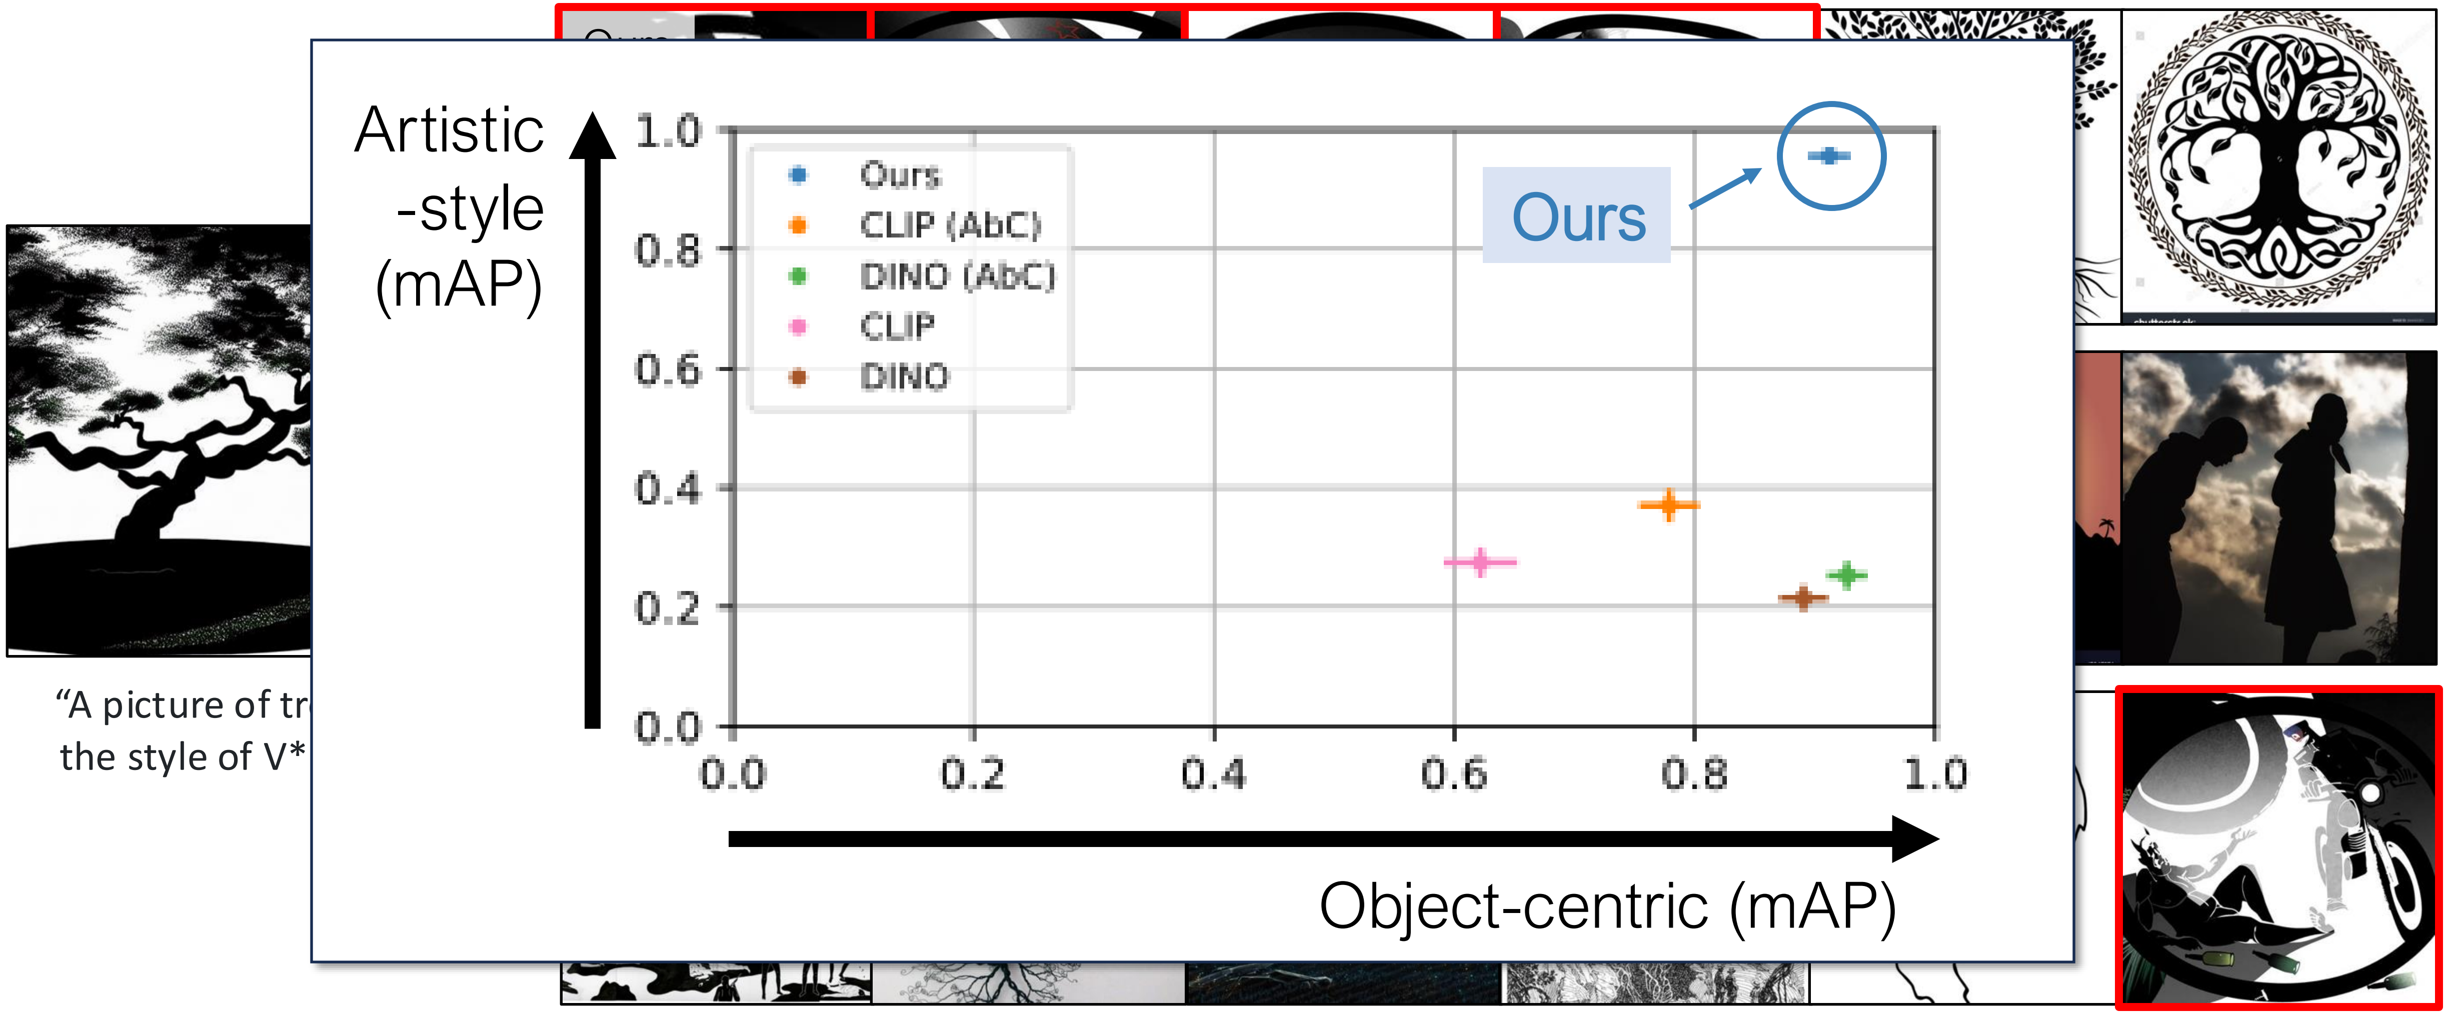

Data Ownership in Generative Models

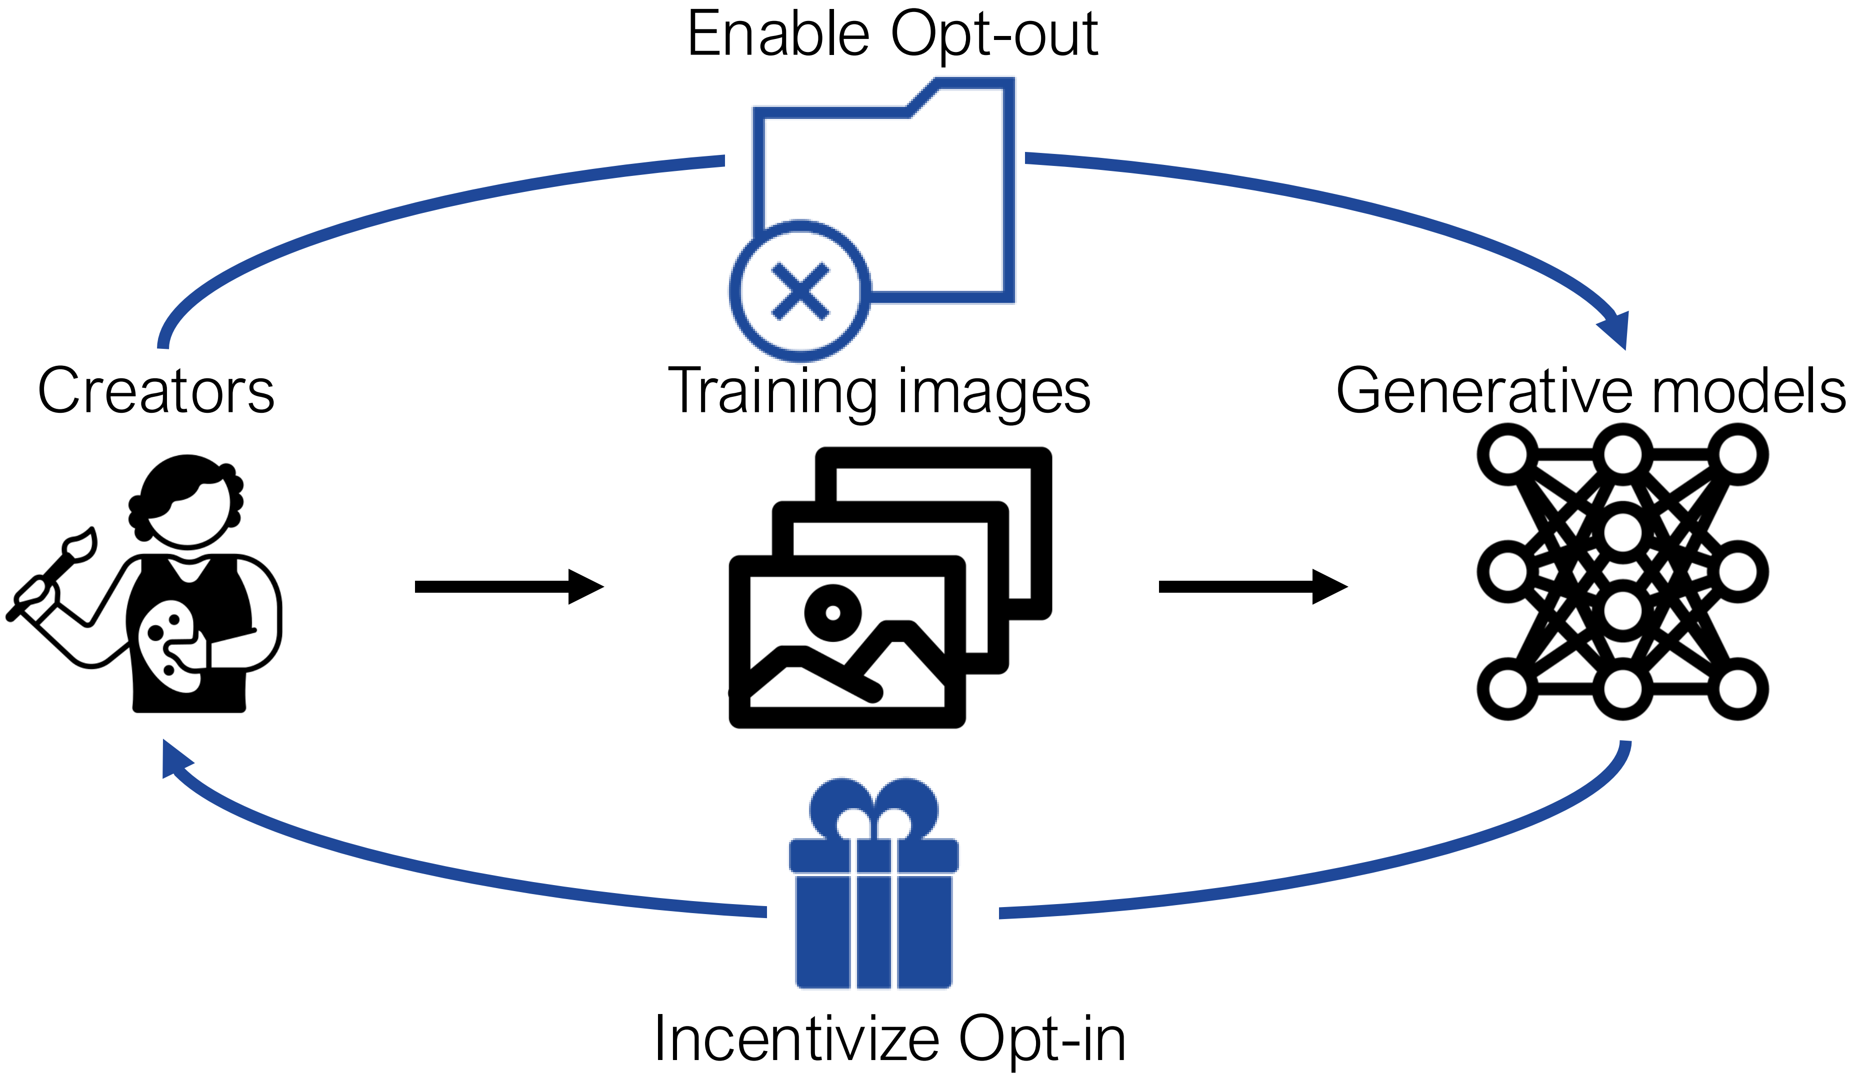

Human Creators

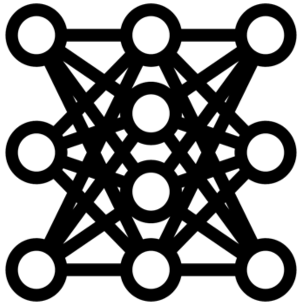

Generative Models

Recent Projects

FlashTex: Relightable Mesh Texturing

[Deng et al., ECCV 2024]

Robot-Human Co-painting

ICRA 2024 Best Paper on Human-Robot Interaction

[Schaldenbrand et al., ICRA 2024]

Diffusion2GAN

"Traditional gondolas lined up along the water, ready to transport visitors."

"Skiers enjoying the pristine slopes of the Swiss Alps on a sunny day."

"Russian Blue cat exploring a garden, surrounded by vibrant flowers."

Distilling Diffusion Models into Conditional GANs [Kang et al., ECCV 2024]

Students and Collaborators

Nupur Kumari

Sheng-Yu Wang

Bingliang Zhang

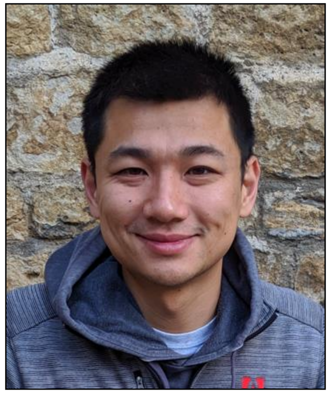

Richard Zhang

Eli Shechtman

Aaron Hertzmann

Alyosha Efros

Human Creators

Generative Models