

Carnegie
Mellon
University

Generative Visual Prompt:

Unifying Distributional Control of Pre-Trained Generative Models

Chen Henry Wu, Saman Motamed, Shaunak Srivastava, Fernando De la Torre

Carnegie Mellon University

henrychenwu@cmu.edu  @ChenHenryWu

Presented @NeurIPS 2022

<https://chenwu98.github.io/PromptGen/>



What is a good generative model?

Faithfully modeling the data distribution?



Uncurated samples from StyleGAN2

Things not in data...

Controllability (hard to label all concepts, e.g., "baby")

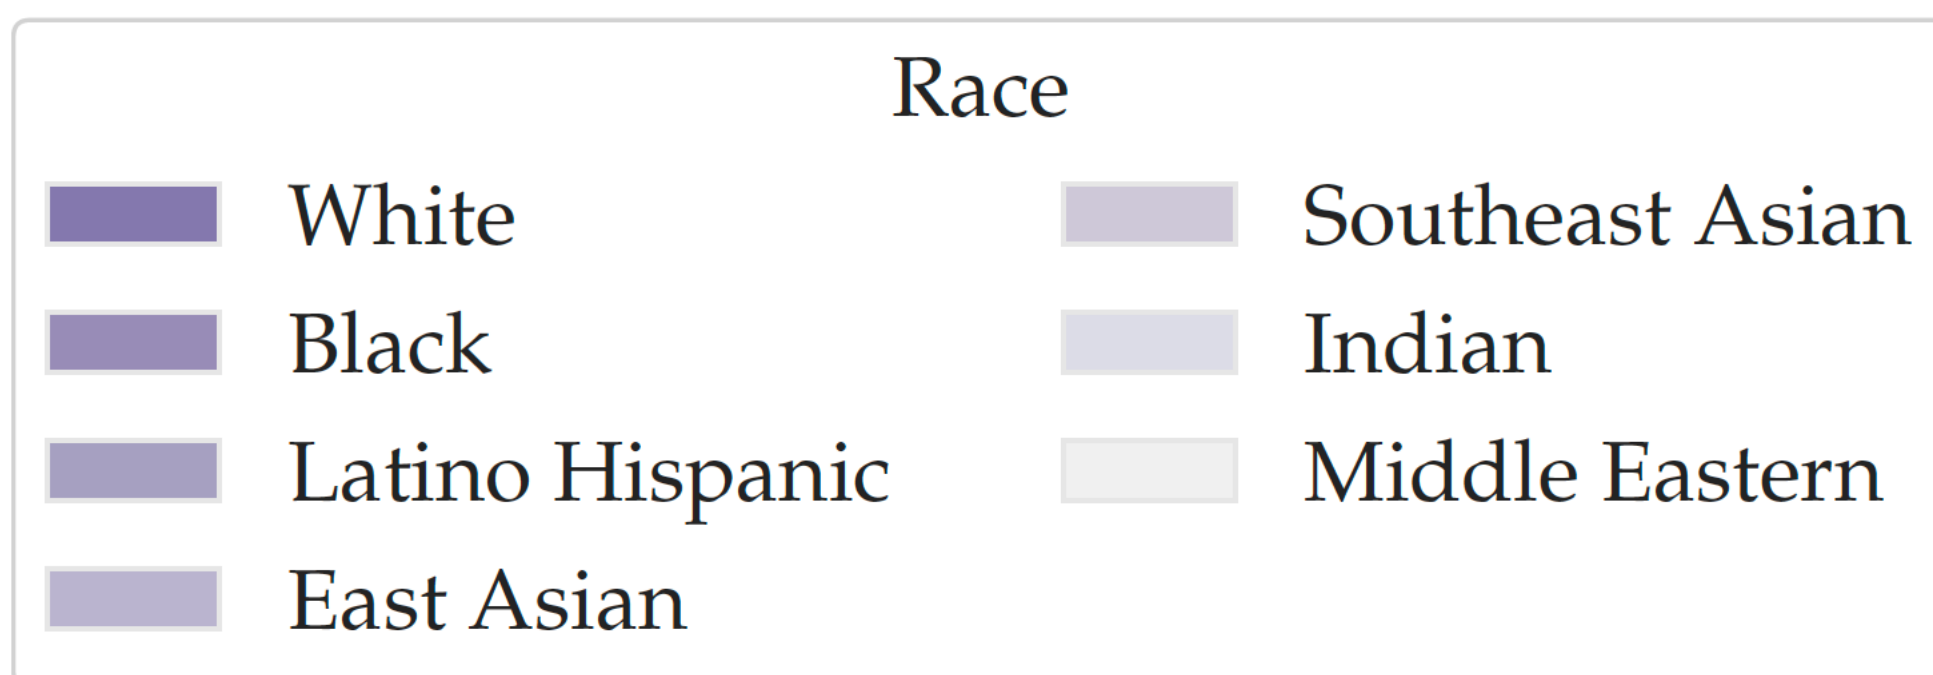


Things not in data...

Controllability (hard to label all concepts, e.g., "baby")



Fairness (hard to build a truly fair training set, e.g., across races)

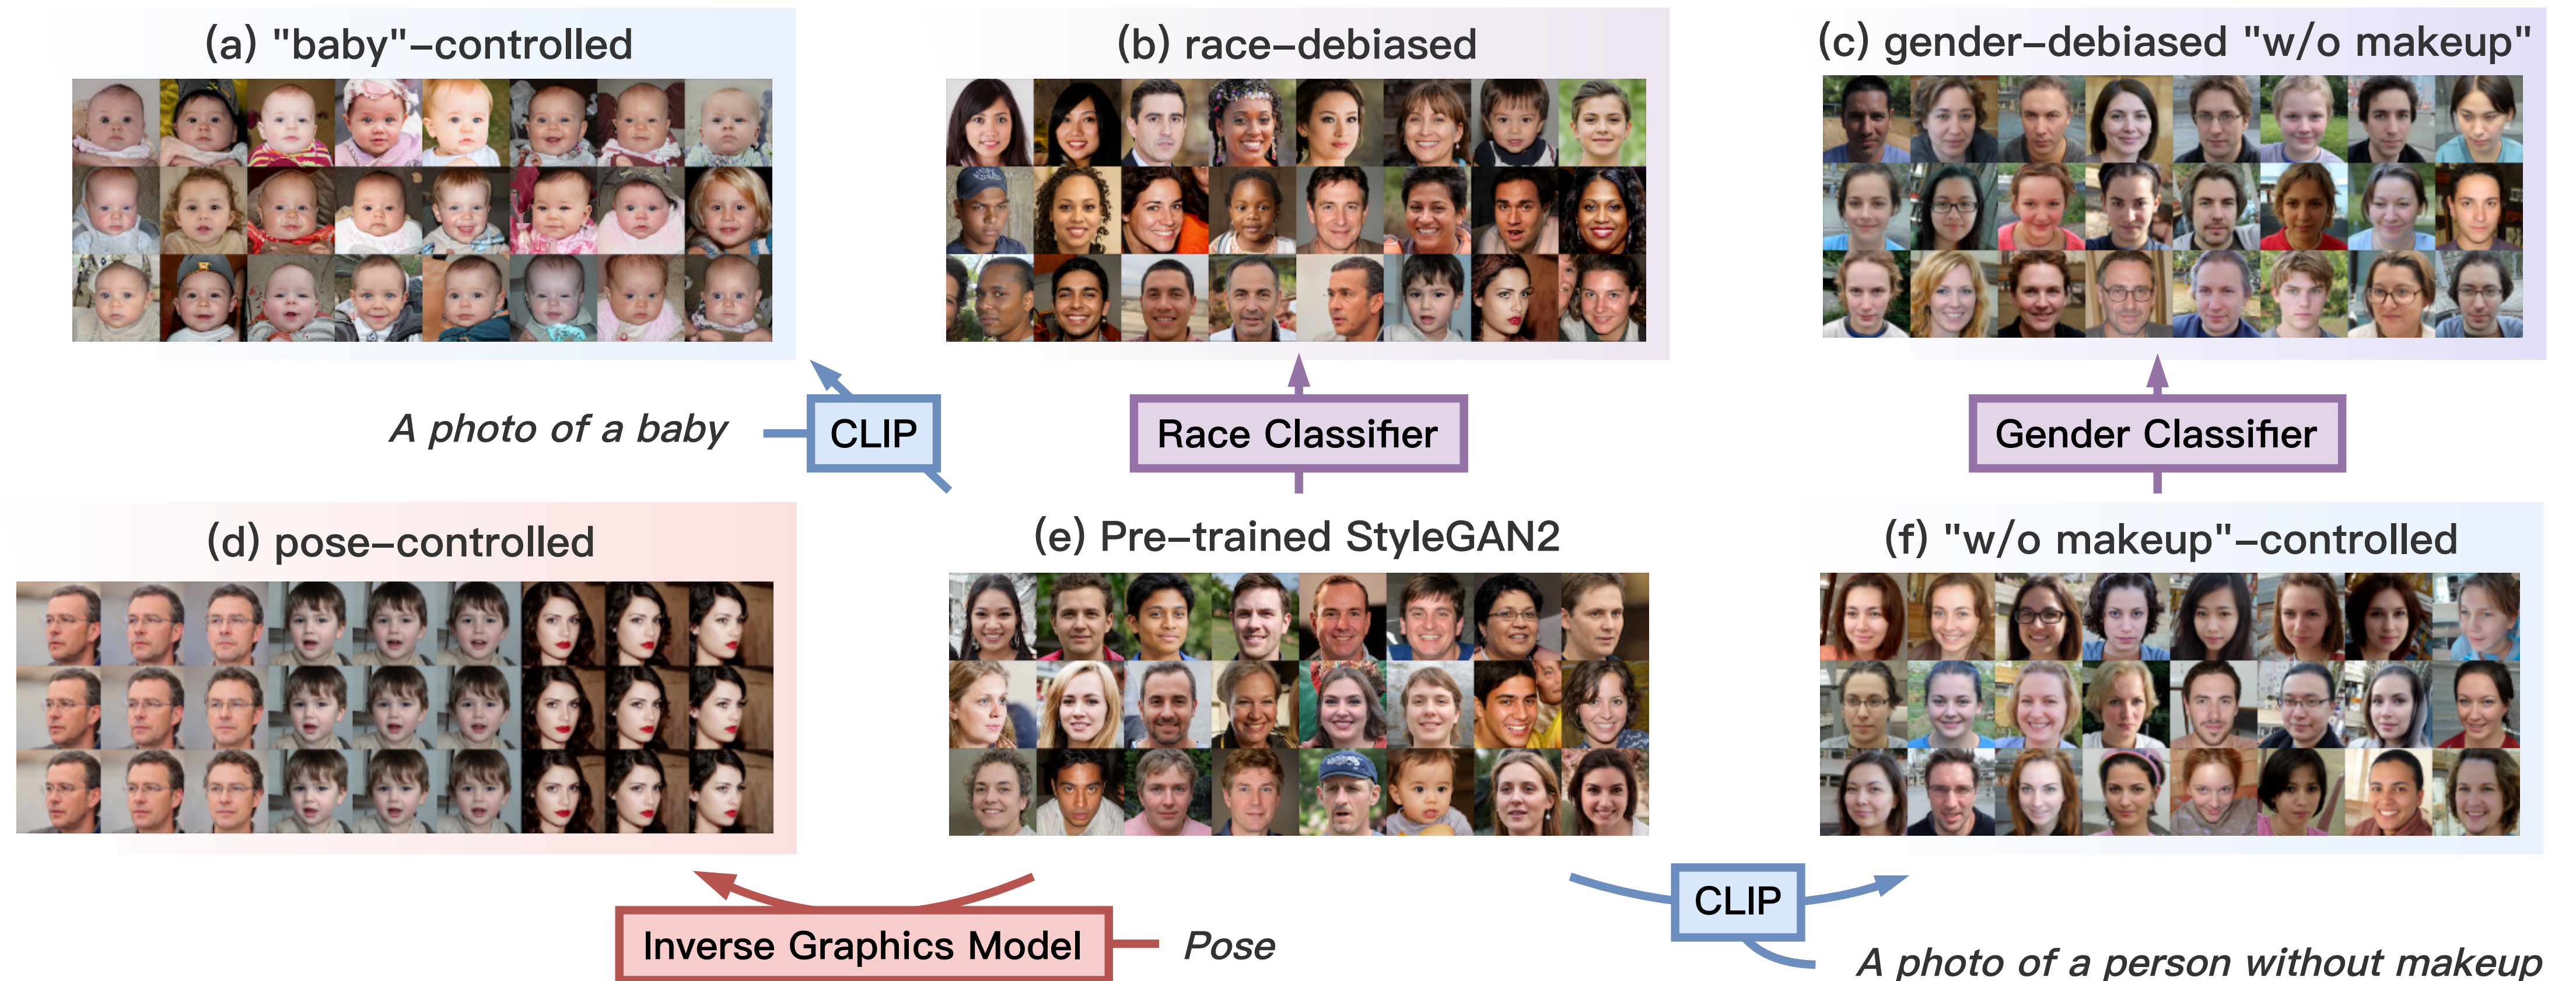


PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space

PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space



PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space

(e) Pre-trained StyleGAN2



PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space

(a) "baby"-controlled



A photo of a baby

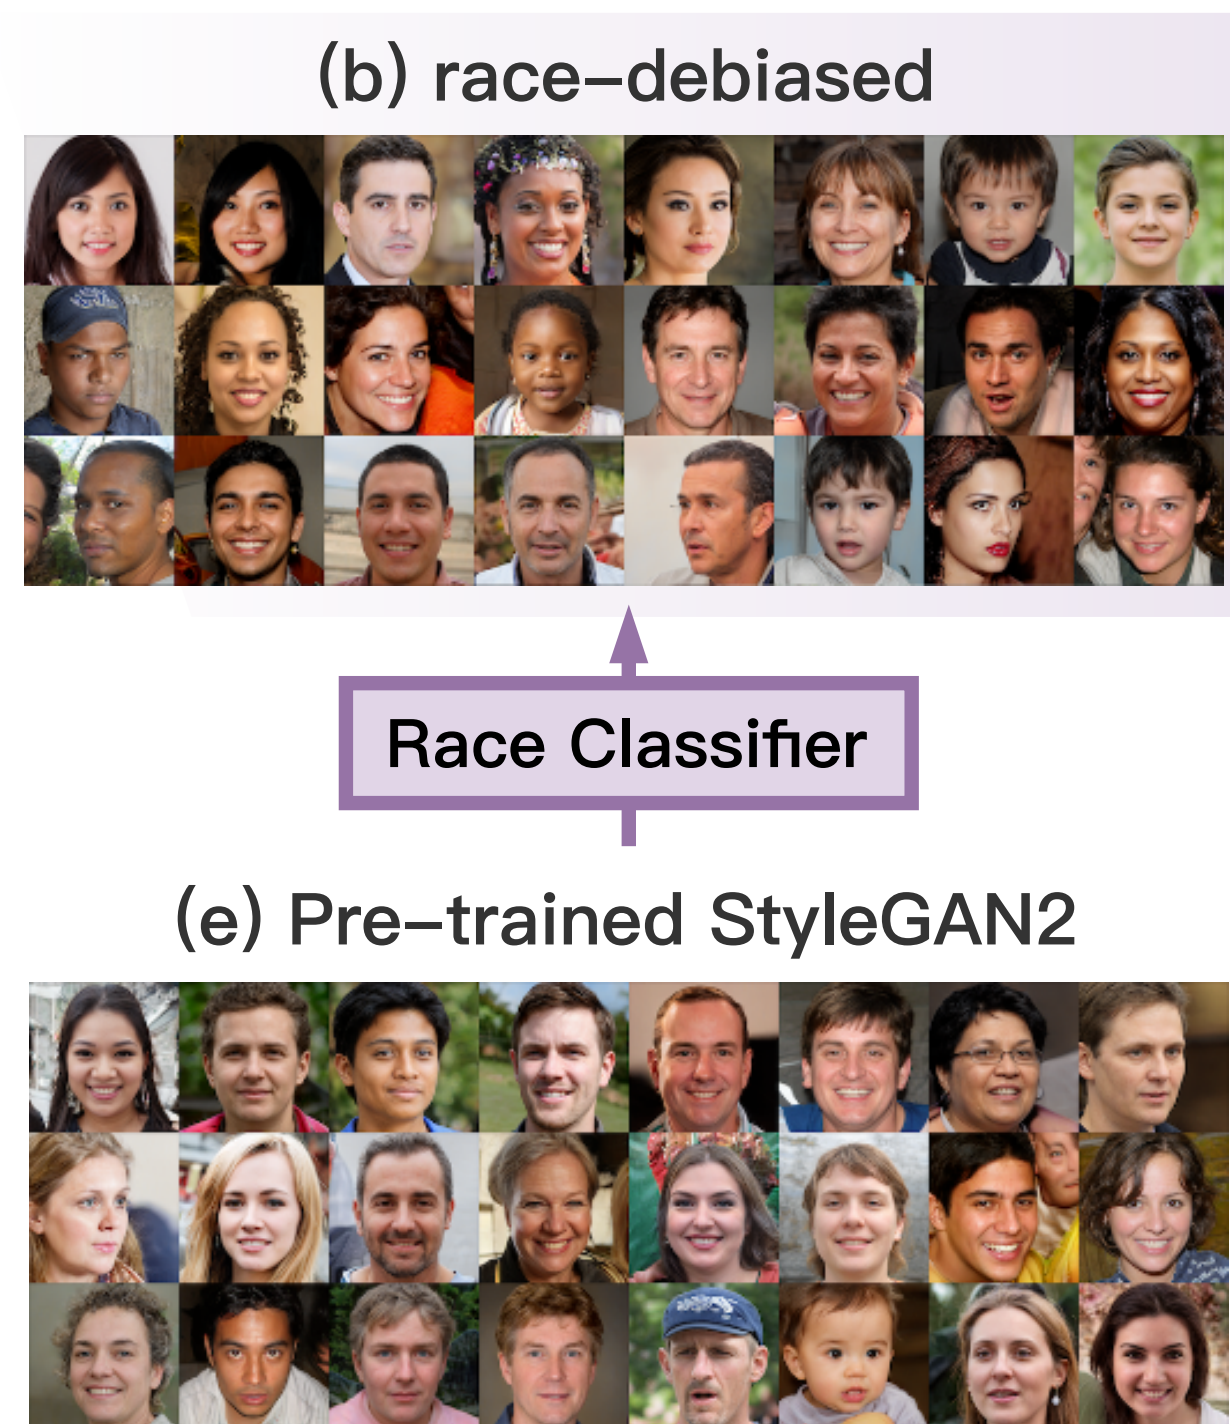
CLIP

(e) Pre-trained StyleGAN2



PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space



PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space

(d) pose-controlled



(e) Pre-trained StyleGAN2



Inverse Graphics Model — Pose

PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space

(e) Pre-trained StyleGAN2



(f) "w/o makeup"-controlled

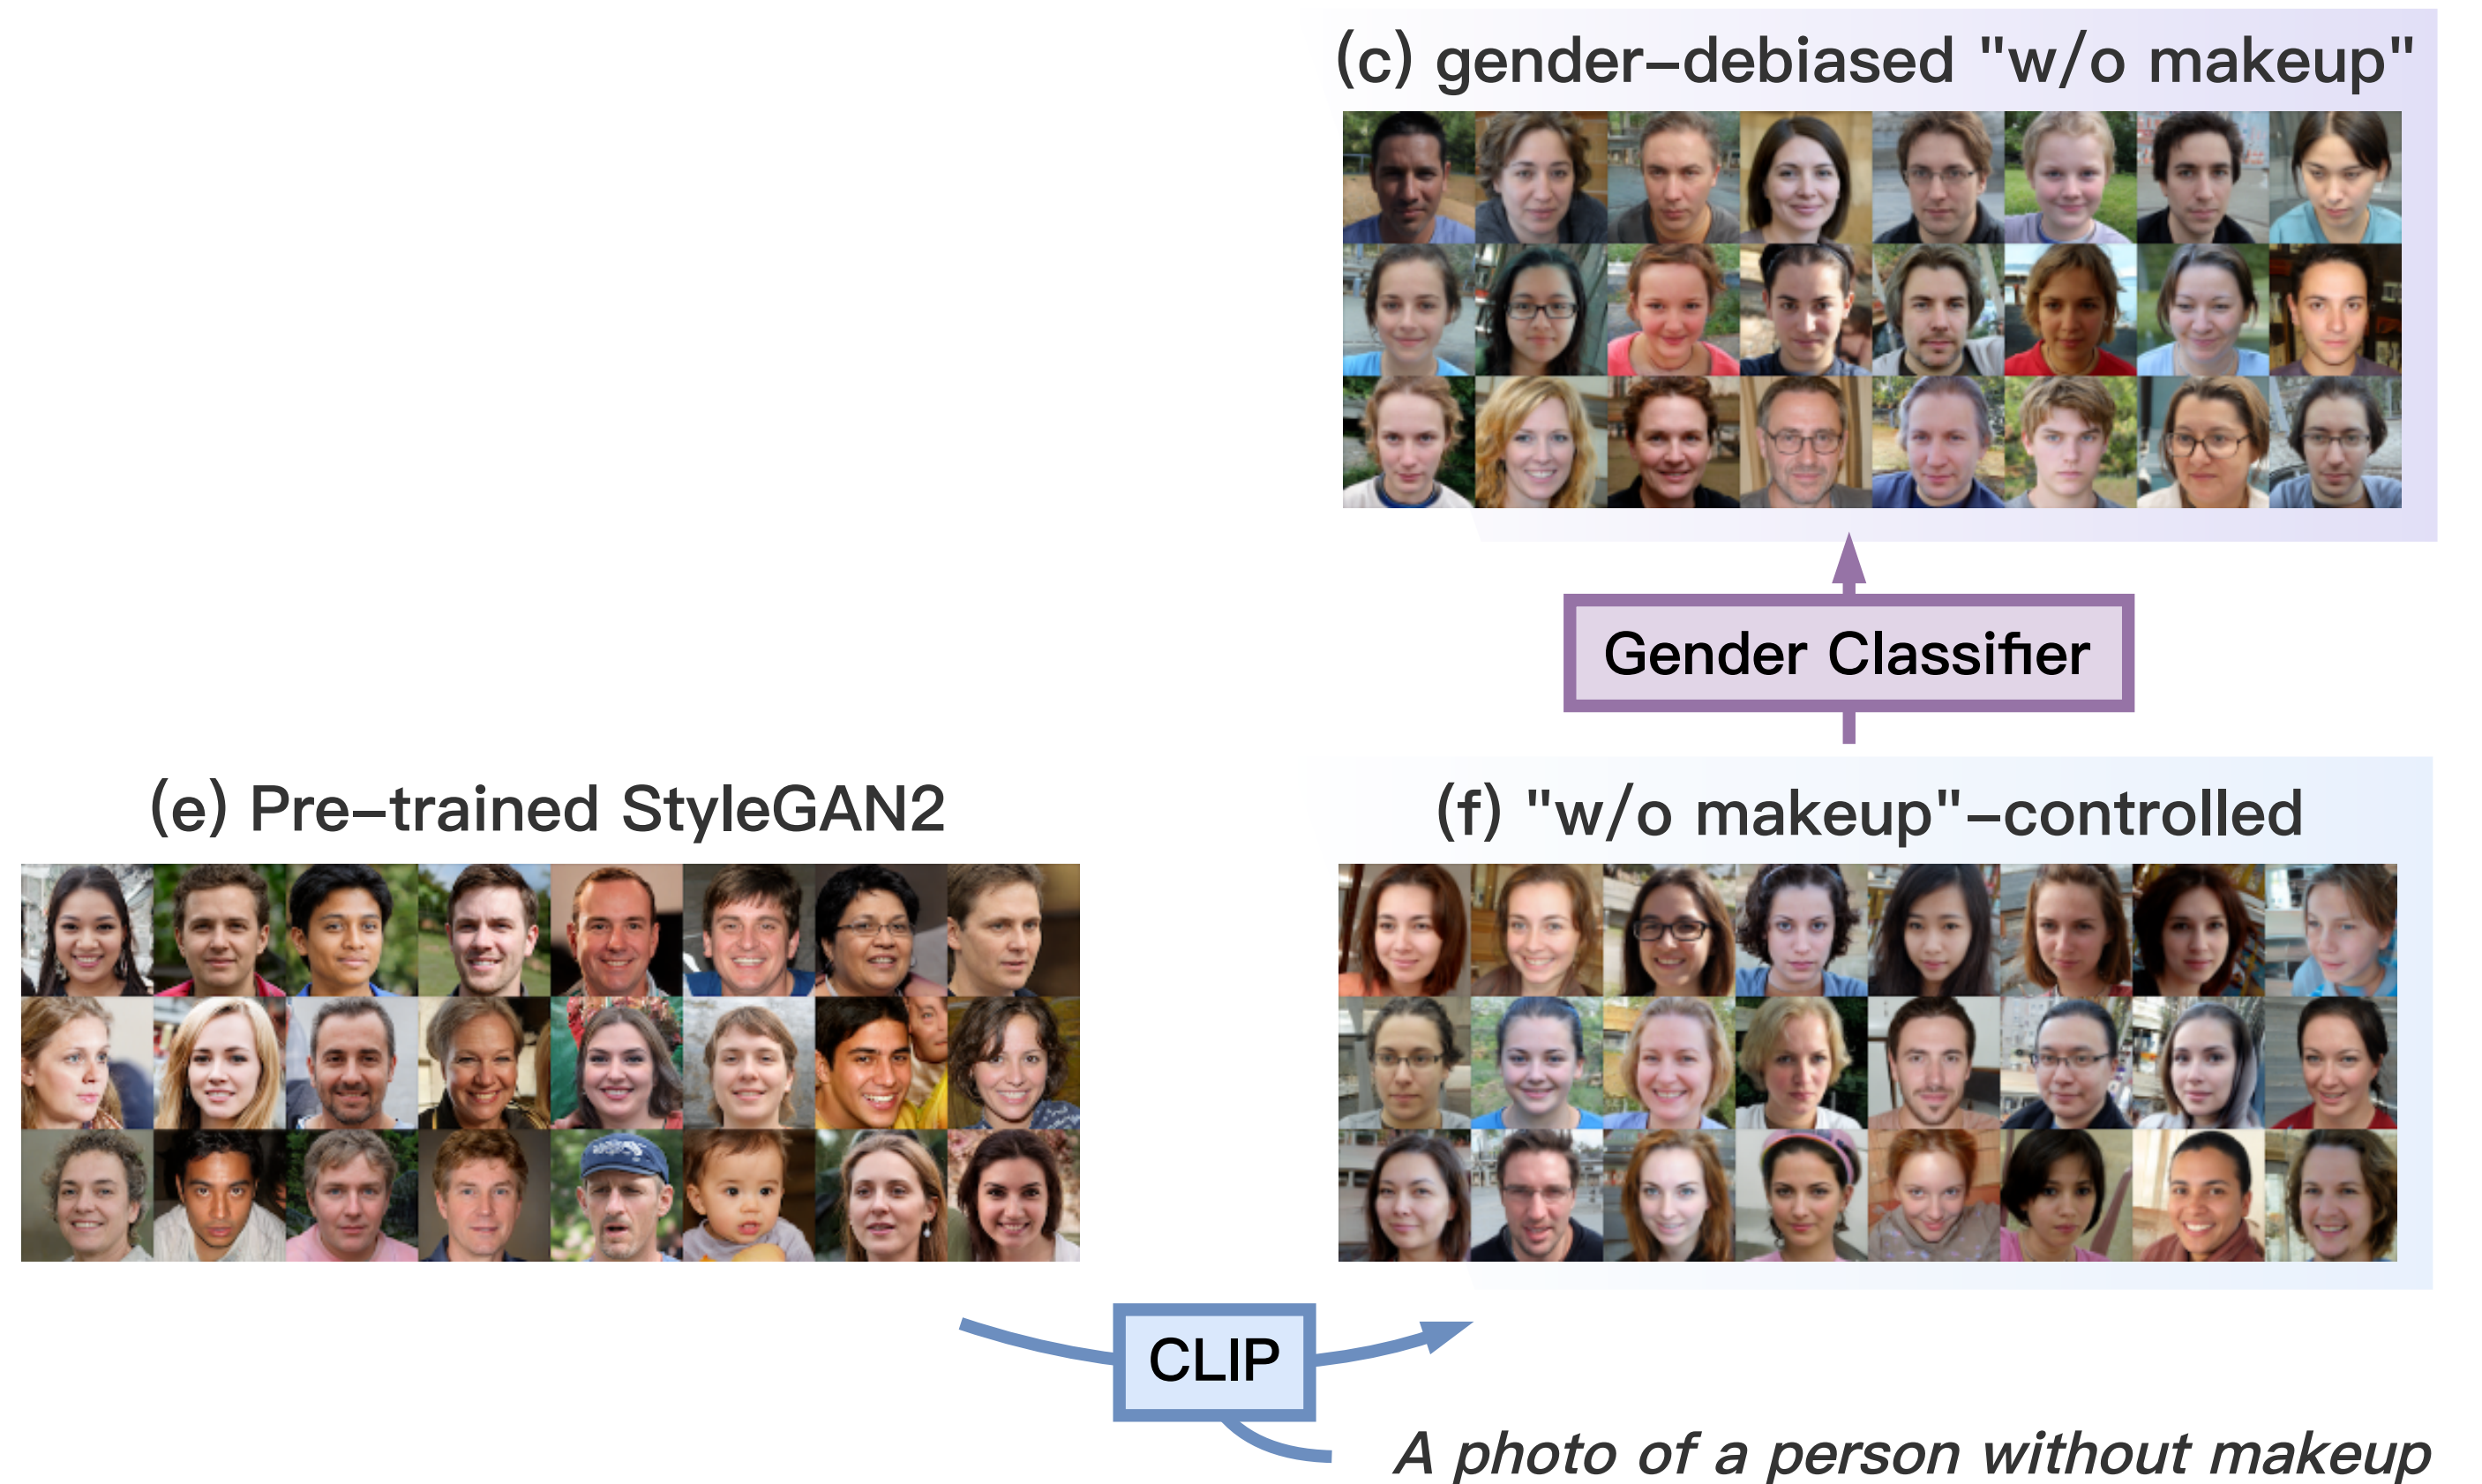


CLIP

A photo of a person without makeup

PromptGen is a remedy!

A feed-forward neural network to model desired distributions in latent space



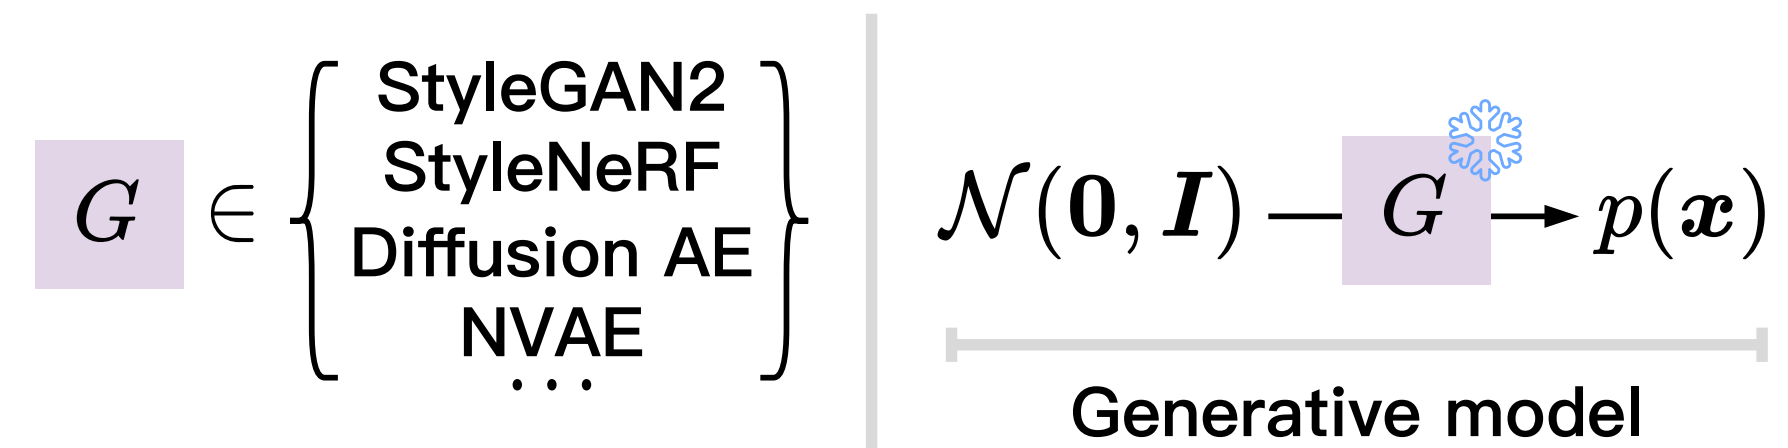
Framework overview

1. User specifies a generative model;

$$G \in \left\{ \begin{array}{l} \text{StyleGAN2} \\ \text{StyleNeRF} \\ \text{Diffusion AE} \\ \text{NVAE} \\ \dots \end{array} \right\}$$

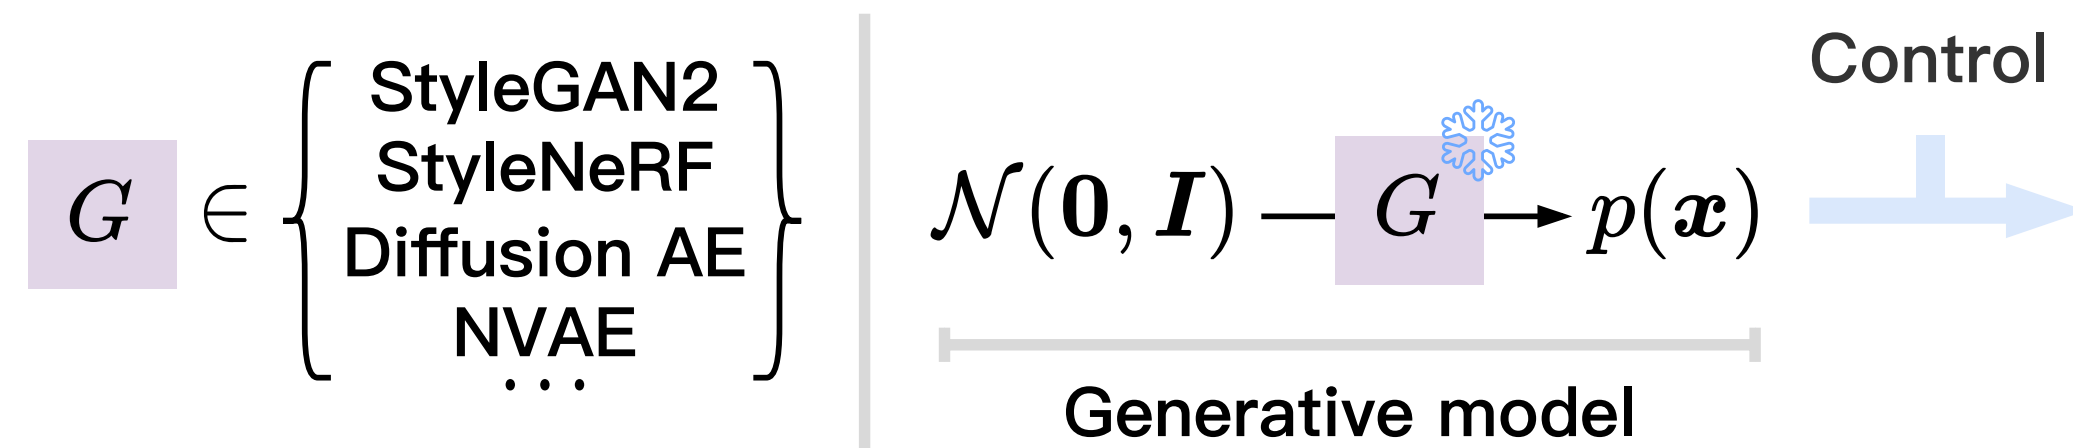
Framework overview

1. User specifies a generative model;



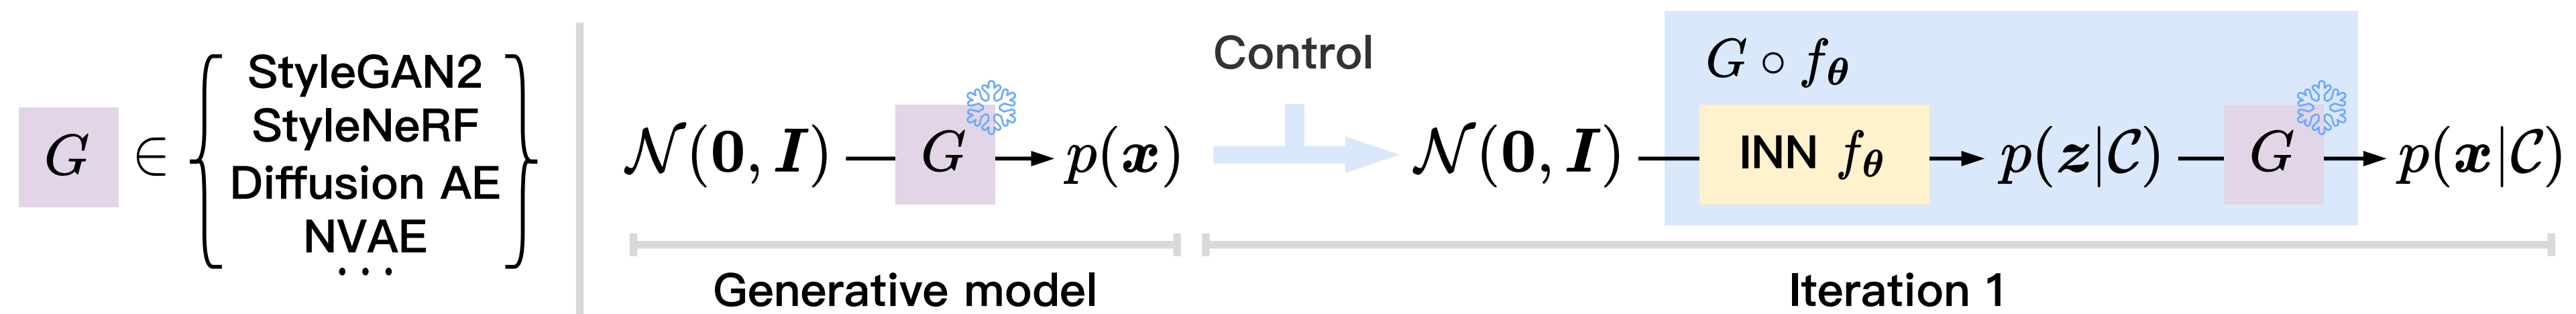
Framework overview

1. User specifies a generative model;
2. User specifies a control;



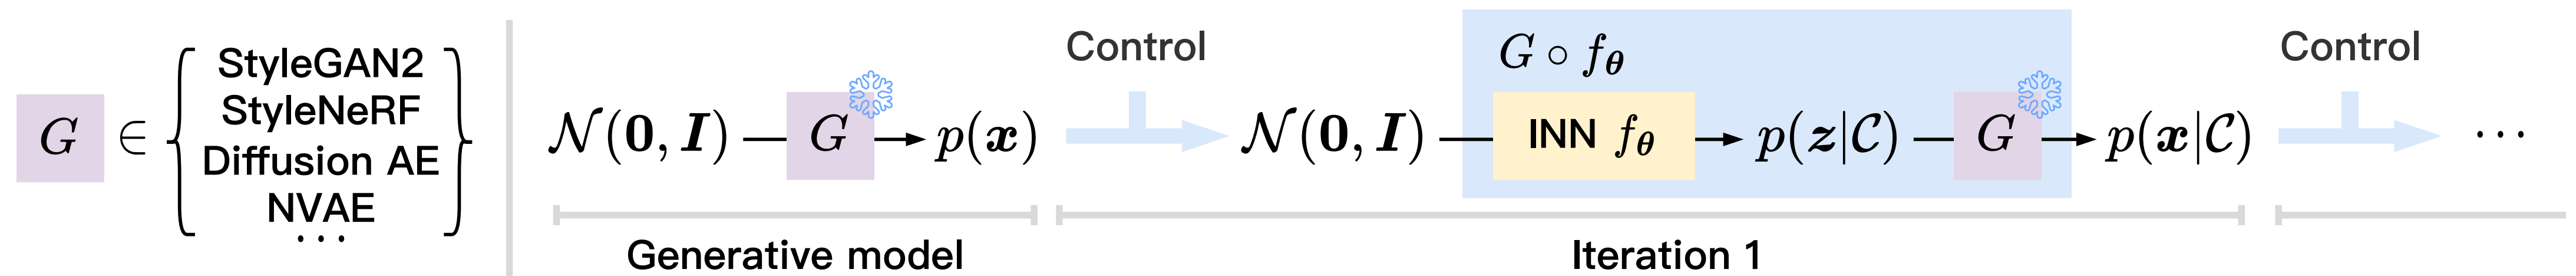
Framework overview

1. User specifies a generative model;
2. User specifies a control;
3. Approximates the control with an invertible neural network;



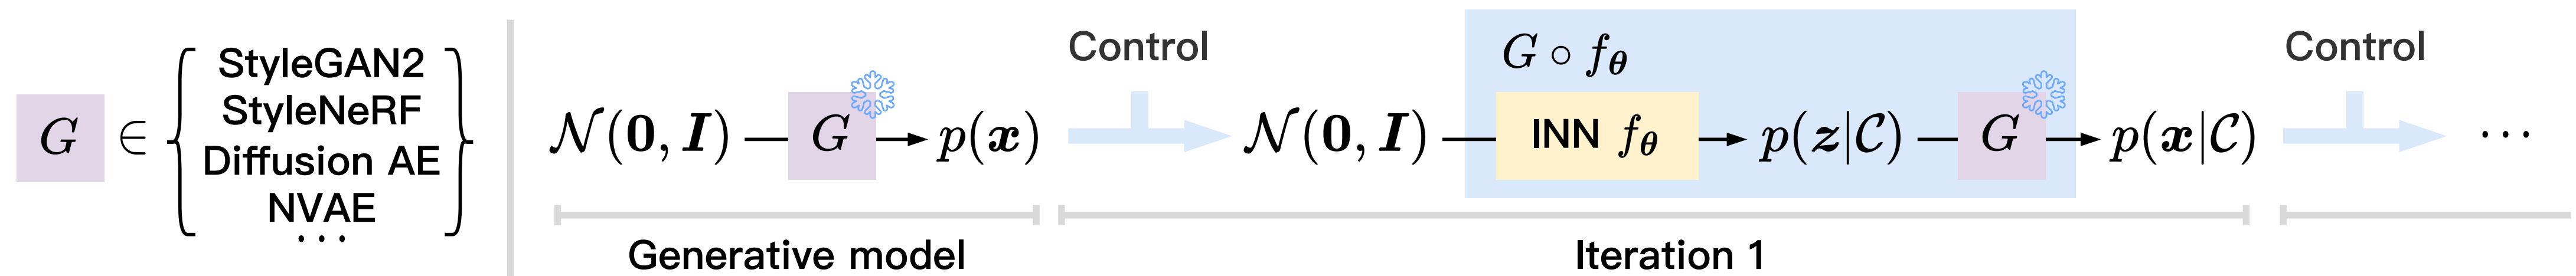
Framework overview

1. User specifies a generative model;
2. User specifies a control;
3. Approximates the control with an invertible neural network;
4. If iterative control is needed: functional composition and repeat 2 & 3.



Framework overview

1. User specifies a generative model;
2. User specifies a control;
3. Approximates the control with an invertible neural network;
4. If iterative control is needed: functional composition and repeat 2 & 3.



How to specify a control?

Image-space energy (lower is better)

Inverse graphics energy

$$E_{\text{inv-graphics}}(\mathbf{x}, \boldsymbol{\rho}) = d\langle f_{\mathcal{X} \rightarrow \mathcal{P}}(\mathbf{x}), \boldsymbol{\rho} \rangle^2$$

Classifier energy

$$E_{\text{classifier}}(\mathbf{x}, a) = -\log P(a|\mathbf{x})$$

CLIP energy

$$E_{\text{CLIP}}(\mathbf{x}, \mathbf{t}) = \frac{1}{L} \sum_{l=1}^L \left(1 - \cos \left\langle \text{CLIP}_{\text{img}}(\text{DiffAug}_l(\mathbf{x})), \text{CLIP}_{\text{text}}(\mathbf{t}) \right\rangle \right)$$

Moment constraint


$$p(\mathbf{x}|\mathcal{C}) = \arg \min_{p(\mathbf{x}|\mathcal{C})} \mathbb{D}_{\text{KL}}(p(\mathbf{x}|\mathcal{C}) || p_{\mathbf{x}}(\mathbf{x})), \quad \text{s.t.} \quad \underbrace{\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}|\mathcal{C})} [\boldsymbol{\gamma}(\mathbf{x})]}_{\text{Moment constraint}} = \boldsymbol{\mu}$$

Deviation from the pre-trained distribution

How to approximate the control?

Optimize $\mathbb{D}_{\text{KL}}(p_{\theta}(\mathbf{z})||p(\mathbf{z}|\mathcal{C}))$ with a normalizing flow in the latent space

Training algorithm and objective:




A synthetic 2D example

Real dist. $\begin{cases} p(\mathbf{x}) = 0.3 \cdot p(\mathbf{x} | a_1) + 0.7 \cdot p(\mathbf{x} | a_2) \\ p(\mathbf{x} | a_1) = N(\mathbf{x} | (3,3)^\top, \mathbf{I}) \\ p(\mathbf{x} | a_2) = N(\mathbf{x} | (-3, -3)^\top, \mathbf{I}) \end{cases}$


Fair classifier $p(a_i | \mathbf{x}) = \frac{p(\mathbf{x} | a_i)}{p(\mathbf{x} | a_1) + p(\mathbf{x} | a_2)}$

(a) Real data




\mathbf{x} -space (data)

(b) Generative model (GAN)




\mathbf{z} -space (latent)



\mathbf{x} -space (data)

(c) PromptGen with classifier control




ϵ -space

\mathbf{z} -space (latent)

\mathbf{x} -space (data)

(d) PromptGen with debiasing control



ϵ -space

\mathbf{z} -space (latent)

\mathbf{x} -space (data)

Real data experiments

Check our paper for details

Code available

