
Generative Visual Prompt: Unifying Distributional Control of Pre-Trained Generative Models

Chen Henry Wu, Saman Motamed, Shaunak Srivastava, Fernando De la Torre
Robotics Institute, Carnegie Mellon University, Pittsburgh, PA
{chenwu2, ftorre}@cs.cmu.edu, {saman.moatamed, shaunak1999}@gmail.com

Abstract

Generative models (e.g., GANs, diffusion models) learn the underlying data distribution in an unsupervised manner. However, many applications of interest require sampling from a particular region of the output space or sampling evenly over a range of characteristics. For efficient sampling in these scenarios, we propose **Generative Visual Prompt** (PromptGen), a framework for distributional control over pre-trained generative models by incorporating knowledge of other off-the-shelf models. PromptGen defines control as energy-based models (EBMs) and samples images in a feed-forward manner by approximating the EBM with invertible neural networks, avoiding optimization at inference. Our experiments demonstrate how PromptGen can efficiently sample from several unconditional generative models (e.g., StyleGAN2, StyleNeRF, diffusion autoencoder, NVAE) in a controlled or/and de-biased manner using various off-the-shelf models: (1) with the CLIP model as control, PromptGen can sample images guided by text, (2) with image classifiers as control, PromptGen can de-bias generative models across a set of attributes or attribute combinations, and (3) with inverse graphics models as control, PromptGen can sample images of the same identity in different poses. (4) Finally, PromptGen reveals that the CLIP model shows a “reporting bias” when used as control, and PromptGen can further de-bias this controlled distribution in an iterative manner.¹

1 Introduction

Generative models learn the underlying high-dimensional data distribution and have achieved promising performance on image synthesis [4, 75, 36, 21]. Though being well praised, they still face two main limitations. First, since generative models are typically trained in an unsupervised way, they lack controllability, meaning that it is unclear how to sample from a specific region of the space. Second, generative models are prone to inherit the imbalance and bias of training data [60, 31]. For instance, StyleGAN2 is more likely to produce images of white individuals, see Figure 1(e). Previous works have studied these challenges separately, and typical methods include editing of “style” codes [23, 68, 31] and explicit conditions [43, 69]. However, these methods are either model-dependent (i.e., requiring a well-structured style space) or label-intensive (i.e., requiring all training samples to be labeled for explicit conditions), limiting their generality and practical use.

To address the above challenges, this paper advocates a unified formulation, *distributional control of generative models*, which enables controllability by incorporating the knowledge of off-the-shelf models (e.g., CLIP [58], classifiers, or inverse graphics models [16]). Based on this unified view, we propose to learn distributions in the latent space of a pre-trained generative model while keeping the pre-trained weights fixed. Given its conceptual similarity to prompting [42, 85, 86, 30, 22], we term our framework as Generative Visual Prompt (PromptGen). PromptGen requires **no training data**, and the only supervision comes from off-the-shelf models that help define the control. Specifically,

¹The code is available at <https://github.com/ChenWu98/Generative-Visual-Prompt>.

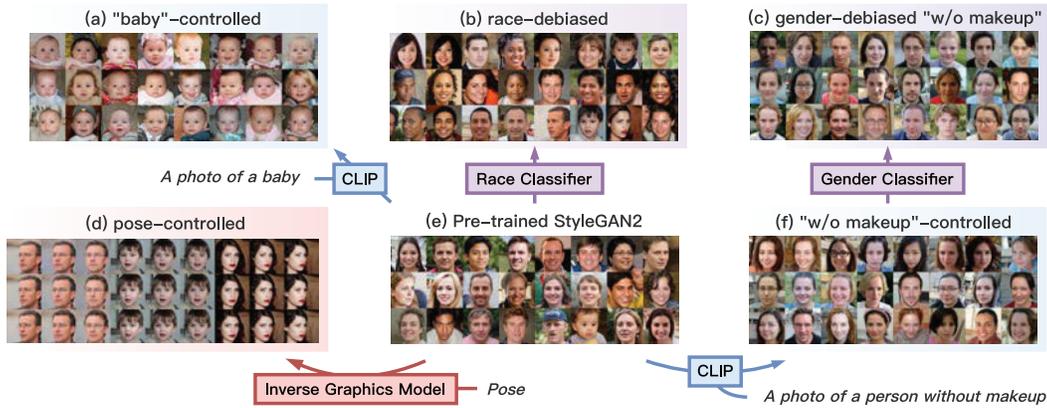


Figure 1: PromptGen uses different off-the-shelf models (e.g., CLIP [58], inverse graphics models, and classifiers) to control the output distribution of a pre-trained and fixed generative model (e.g., StyleGAN2). Colored boxes (e.g., the blue CLIP) indicate the control. See the text for details.

PromptGen allows the user to define controls using an energy-based model (EBM) approximated by an invertible neural network (INN). Unlike methods that require optimization at inference in EBM sampling [49, 14, 45, 52], PromptGen samples images in a **feed-forward** manner, which is highly efficient. Moreover, PromptGen **stands alone** at inference, meaning that we can discard the off-the-shelf models, which define the control, after training.

We illustrate the benefits of PromptGen with several experiments. Figure 1 demonstrates our main findings with StyleGAN2 trained on FFHQ [35] without any labels, while we also show results for StyleNeRF [21], diffusion autoencoder [57], and NVAE [75] in Section 4. Figure 1(a) illustrates how we can sample from StyleGAN2 based on text descriptions such as “a photo of a baby”. Figure 1(b) shows that PromptGen can leverage a race classifier (potentially trained on a different dataset) to sample uniformly across all races, de-biasing the pre-trained StyleGAN2. Figure 1(d) illustrates that PromptGen can generate images of the same identity in different poses, guided by a pose regressor.

Finally, it is worth pointing out that PromptGen not only offers **generality** for algorithmic design and **modularity** for control composition, but also enables **iterative** controls when some controls are contingent on others. For instance, one may train a text-controlled distribution and then de-bias this distribution. To achieve this, we view the composition of the INN and the generative model as a new “generative model” to be controlled. Figure 1(f) illustrates that PromptGen reveals “reporting bias” of the CLIP model [58], where “without makeup” is – perhaps surprisingly – positively correlated with female, and Figure 1(c) shows that PromptGen can further mitigate this bias with iterative control.

Table 1: Comparison between methods.

	StyleFlow [1]	PPGM [49] / LACE [52]	Guided DDPM [10]	PromptGen
Arbitrary control (e.g., CLIP)	✗	✓	✓	✓
Low-dimensional latent space	✓	✓	✗	✓
Stands alone at inference	✓	✗	✗	✓
Feed-forward (i.e., no inference-time optim.)	✓	✗	✗	✓
Iterative distributional control	✗	✗	✗	✓

2 Related Work

Over the past few years, generative models have gained the ability to generate images with high visual quality. A few of the most widely used methods include generative adversarial networks (GANs) [18], VAE [40], invertible neural networks [11, 12], and diffusion models [25, 71]. In particular, generative models trained on large amounts of unlabeled data, e.g., BigGAN [4], StyleGANs [35, 36, 34, 64], Glow [39], and diffusion models [51, 10], achieve promising image synthesis results.

Despite their success, controllability and de-biasing are still two fundamental challenges generative models face. For controllability, existing methods include explicit conditioning at training [43, 69] or local editing of the learned representation, e.g., “style” codes [1, 56, 46, 68, 80]. For de-biasing, existing methods include local editing of “style” codes [60, 31] and importance sampling for either training [20] or inference [29]. Existing works study these problems separately, each requiring a specific design for the task studied. On the contrary, PromptGen is a unified framework for arbitrary controls defined by off-the-shelf models. The benefits of unifying tasks have been shown by a recent trend of works in multiple research areas across vision and language [59, 48, 63, 81, 61]. Moreover, one can fine-tune generative models [17] for domain adaptation, which is orthogonal to PromptGen: since PromptGen maps a generative model to another generative model (Algorithm 1), fine-tuning can be applied before or after PromptGen training. We leave this exploration to future studies.

Previous methods usually sample from EBM [41] with Markov Chain Monte Carlo (MCMC) [73, 76, 14, 19, 13, 45]. Among them, plug-and-play generative models (PPGMs) [49] and LACE [52] define latent-space EBMs. However, MCMC requires inference-time optimization, which is inefficient and requires the off-the-shelf models to be available at inference; this is also the criticism for diffusion models, regardless of being gradient-guided [51, 67] or not [25, 26, 50]. In contrast, PromptGen achieves efficient, feed-forward sampling. Table 1 shows a comparison with previous methods. Our INN training is similar to that proposed by [54] to sample from physical systems, which is later used by [78] to solve inverse problems with two composed INNs. Notably, [54] and [78] do not leverage a low-dimensional latent space, and [78] requires training a separate model for each image sample. In this paper, we use INN to model arbitrary EBMs for various generative models.

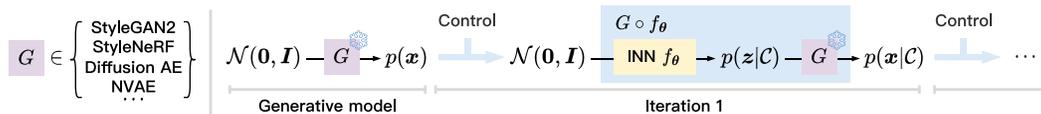


Figure 2: Overview of Generative Visual Prompt (PromptGen). Given a pre-trained generative model G and a control \mathcal{C} , we learn a distribution $p(z|\mathcal{C})$ in G ’s latent space, while keeping G fixed. Each control \mathcal{C} can have multiple components, e.g., $\mathcal{C} = \{\text{text} = \text{“a photo of a baby”}, \text{gender} = \text{male}\}$. PromptGen views the composition $G \circ f_\theta$ as a new “generative model” for iterative control.

3 Method

Figure 2 illustrates our PromptGen framework. To begin, the user selects a pre-trained generative model G . PromptGen then lets the user specify a control \mathcal{C} as an energy-based model (EBM) $p(z|\mathcal{C})$. We train an INN f_θ to approximate $p(z|\mathcal{C})$. PromptGen views the functional composition $G \circ f_\theta$ as a new generative model and can perform iterative control. Algorithm 1 describes the overall procedure.

Algorithm 1: Generative Visual Prompt (PromptGen)

Input: Generative model $G : \mathbb{R}^d \rightarrow \mathcal{X}$

repeat

1. **Input:** control \mathcal{C} of the current iteration
2. Define an EBM $p(z|\mathcal{C})$ for \mathcal{C} (Section 3.1)
3. Train an INN $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$ to approximate $p(z|\mathcal{C})$ (Section 3.2)
4. $G \leftarrow G \circ f_\theta$

until user stops the iteration

return $G : \mathbb{R}^d \rightarrow \mathcal{X}$, which is a feed-forward network

3.1 Latent-Space EBM for Distributional Control

The plug-and-play generative model [49] was first proposed to use a latent-space EBM for controllable image synthesis. If a fixed generative model is used, then theoretically, any image-space EBM can be viewed as a latent-space EBM, as shown by [19] and [52]. We define the latent-space EBM following similar formulation as [49, 19, 52], with some new energy functions. We then extend the formulation

to incorporate the moment constraint [7], which was adopted for language modeling [37], but unlike [37], we define the moment constraint in the latent space to accommodate generative vision models.

We define a control \mathcal{C} as M independent properties $\{\mathbf{y}_1, \dots, \mathbf{y}_M\}$, e.g., \mathbf{y}_1 can be a text description and \mathbf{y}_2 can be an attribute. The controllability can be defined as the EBM (detailed in Appendix B.1)

$$p(\mathbf{x}|\mathcal{C}) = \frac{p_{\mathbf{x}}(\mathbf{x})e^{-E_{\mathcal{C}}(\mathbf{x})}}{Z_X}, \quad E_{\mathcal{C}}(\mathbf{x}) = \sum_{i=1}^M \lambda_i E_i(\mathbf{x}, \mathbf{y}_i), \quad Z_X = \int_{\mathbf{x}'} p_{\mathbf{x}}(\mathbf{x}')e^{-E_{\mathcal{C}}(\mathbf{x}')}d\mathbf{x}', \quad (1)$$

which reweights the image prior $p_{\mathbf{x}}(\mathbf{x})$ with energy $E_{\mathcal{C}}(\mathbf{x})$, where images with smaller energy are preferred. Using a pre-trained generative model $G: \mathbb{R}^d \rightarrow \mathcal{X}$ that maps a latent code \mathbf{z} to an image \mathbf{x} , the image prior $p_{\mathbf{x}}(\mathbf{x})$ is defined by sampling a latent code \mathbf{z} from $p_{\mathbf{z}} = \mathcal{N}(\mathbf{0}, \mathbf{I})$ and mapping it to $\mathbf{x} = G(\mathbf{z})$. Appendix B.2 shows that this EBM is equivalent to the latent-space EBM

$$p(\mathbf{z}|\mathcal{C}) = \frac{p_{\mathbf{z}}(\mathbf{z})e^{-E_{\mathcal{C}}(G(\mathbf{z}))}}{Z}, \quad (2)$$

$$E_{\mathcal{C}}(G(\mathbf{z})) = \sum_{i=1}^M \lambda_i E_i(G(\mathbf{z}), \mathbf{y}_i), \quad Z = \int_{\mathbf{z}'} p_{\mathbf{z}}(\mathbf{z}')e^{-E_{\mathcal{C}}(G(\mathbf{z}'))}d\mathbf{z}'.$$

Latent-space EBM allows us to use any off-the-shelf model to specify the control. The following are some examples that are discussed in this paper (explained in Appendix B.1):

Classifier energy: Given a classifier $P(\cdot|\mathbf{x})$ and the target class a that we want to sample images from, we define the classifier energy as $E_{\text{classifier}}(\mathbf{x}, a) = -\log P(a|\mathbf{x})$.

CLIP energy: Using the CLIP model [58], we define the CLIP energy as the cosine distance between the embeddings of the image and the text \mathbf{t} , averaged over L differentiable augmentations [84, 46]:

$$E_{\text{CLIP}}(\mathbf{x}, \mathbf{t}) = \frac{1}{L} \sum_{l=1}^L \left(1 - \cos \left\langle \text{CLIP}_{\text{img}}(\text{DiffAug}_l(\mathbf{x})), \text{CLIP}_{\text{text}}(\mathbf{t}) \right\rangle \right). \quad (3)$$

Inverse graphics energy: Given an inverse graphics model, $f_{\mathcal{X} \rightarrow \mathcal{P}}$, which infers image parameters (e.g., pose and expression), and the target parameters $\boldsymbol{\rho}$, we define the inverse graphics energy as

$$E_{\text{inv-graphics}}(\mathbf{x}, \boldsymbol{\rho}) = d \langle f_{\mathcal{X} \rightarrow \mathcal{P}}(\mathbf{x}), \boldsymbol{\rho} \rangle^2, \quad (4)$$

where $d \langle \cdot, \cdot \rangle$ is the geodesic distance between the inferred parameters and the target parameters.

Moment constraint: Some controls cannot be *directly* defined by off-the-shelf models, and the moment constraint [7, 37] is one of them. Given a mapping $\gamma: \mathcal{X} \rightarrow \mathbb{R}^K$ (e.g., γ can be a classifier that outputs the probability simplex), the moment constraint defines the target distribution $p(\mathbf{x}|\mathcal{C})$ as

$$p(\mathbf{x}|\mathcal{C}) = \arg \min_{p(\mathbf{x}|\mathcal{C})} \mathbb{D}_{\text{KL}} \left(p(\mathbf{x}|\mathcal{C}) \| p_{\mathbf{x}}(\mathbf{x}) \right), \quad \text{s.t.} \quad \underbrace{\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}|\mathcal{C})} [\gamma(\mathbf{x})]}_{\text{Moment constraint}} = \boldsymbol{\mu}, \quad (5)$$

Deviation from the pre-trained distribution

where $\boldsymbol{\mu}$ is the user-specified vector. For example, if we want to generate images that are uniformly distributed across races, we may use a race classifier as γ and define $\boldsymbol{\mu} = (|\mathcal{A}|^{-1}, \dots, |\mathcal{A}|^{-1})$ where \mathcal{A} is the set of races. It means that we would like to find a distribution that 1) stays close to the original, pre-trained generative model G and 2) mitigates G 's bias. In this paper, we generalize the moment constraint to the latent space, and approximate the above objective as (detailed in Appendix B.6):

$$p(\mathbf{z}|\mathcal{C}) = \frac{p_{\mathbf{z}}(\mathbf{z}) \exp \left(\hat{\boldsymbol{\beta}}^{\top} \gamma(G(\mathbf{z})) \right)}{Z}, \quad Z = \int_{\mathbf{z}'} p_{\mathbf{z}}(\mathbf{z}') \exp \left(\hat{\boldsymbol{\beta}}^{\top} \gamma(G(\mathbf{z}')) \right) d\mathbf{z}', \quad (6)$$

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \mathbb{E}_{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(N)} \stackrel{\text{i.i.d.}}{\sim} p_{\mathbf{z}}(\mathbf{z}), \mathbf{x}^{(j)} = G(\mathbf{z}^{(j)})} \left\| \frac{\sum_{j=1}^N \exp \left(\boldsymbol{\beta}^{\top} \gamma(\mathbf{x}^{(j)}) \right) \gamma(\mathbf{x}^{(j)})}{\sum_{j'=1}^N \exp \left(\boldsymbol{\beta}^{\top} \gamma(\mathbf{x}^{(j')}) \right)} - \boldsymbol{\mu} \right\|_2^2. \quad (7)$$

Notably, to optimize $\hat{\boldsymbol{\beta}}$ in Eq. (7), we leverage existing gradient-based optimization tools such as stochastic gradient descent (SGD) and Adam [38].

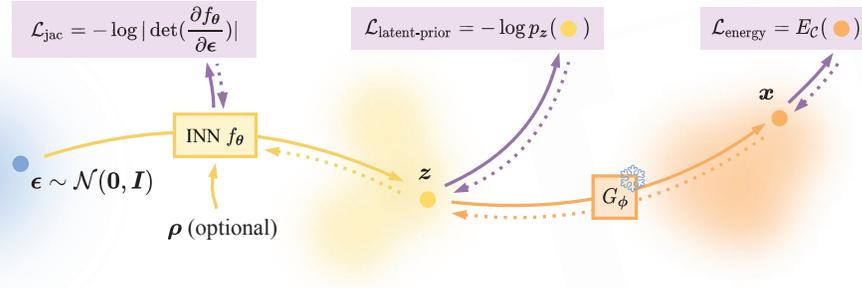


Figure 3: Illustration for Algorithm 2. In the forward pass (solid curves), we sample $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, map ϵ to latent code $z = f_\theta(\epsilon)$ with an INN f_θ , and map z to an image $x = G(z)$ with a fixed generative model G . Dashed curves show the gradients. Details are provided in Section 3.2.

Algorithm 2: Approximating Latent-Space EBM with INN

while *not converged* **do**

1. Sample $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 2. Map ϵ to latent code $z = f_\theta(\epsilon)$
 3. Map z to an image $x = G(z)$
 4. Optimize θ with gradient $\nabla_\theta \left(-\log \left| \det \left(\frac{\partial f_\theta}{\partial \epsilon} \right) \right| - \log p_z(z) + E_C(x) \right)$
-

3.2 Approximating EBM with Invertible Neural Network

Given the functional form of the EBM $p(z|\mathcal{C})$, our next step is to approximate it with an efficient sampling network. To achieve this, we train a distribution $p_\theta(z)$ that minimizes the KL divergence $\mathbb{D}_{\text{KL}}(p_\theta(z)||p(z|\mathcal{C}))$. Estimating $\mathbb{D}_{\text{KL}}(p_\theta(z)||p(z|\mathcal{C}))$ requires easily sampling $z \sim p_\theta$ and tractably computing $p_\theta(z)$. Inspired by [54], we model p_θ with an INN f_θ that defines a bijection $z = f_\theta(\epsilon)$, which has two merits besides the invertibility: (1) one can easily sample z by sampling $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and mapping it to $z = f_\theta(\epsilon)$, and (2) $p_\theta(z)$ has a closed-form solution:

$$\log p_\theta(z) = \log \mathcal{N}(\epsilon|\mathbf{0}, \mathbf{I}) - \log \left| \det \left(\frac{\partial f_\theta}{\partial \epsilon} \right) \right|, \quad z = f_\theta(\epsilon). \quad (8)$$

Based on these properties of INN, we can rewrite our KL divergence objective $\mathbb{D}_{\text{KL}}(p_\theta(z)||p(z|\mathcal{C}))$ as (full derivations in Appendix B.3) the following form:

$$\begin{aligned} \mathbb{D}_{\text{KL}}(p_\theta(z)||p(z|\mathcal{C})) &= \mathbb{E}_{z \sim p_\theta(z), x=G(z)} \left[\log \frac{p_\theta(z)}{p_z(z)e^{-E_C(x)}/Z} \right] \\ &= \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), z=f_\theta(\epsilon), x=G(z)} \left[-\log \left| \det \left(\frac{\partial f_\theta}{\partial \epsilon} \right) \right| - \log p_z(z) + E_C(x) \right] - \mathbb{H}_{\mathcal{N}(\mathbf{0}, \mathbf{I})} + \log Z. \end{aligned} \quad (9)$$

Since $\mathbb{H}_{\mathcal{N}(\mathbf{0}, \mathbf{I})}$ and $\log Z$ are independent of θ , our training objective becomes

$$\arg \min_{\theta} \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), z=f_\theta(\epsilon), x=G(z)} \left[\underbrace{-\log \left| \det(\partial f_\theta / \partial \epsilon) \right|}_{\mathcal{L}_{\text{jac}}} - \underbrace{\log p_z(z)}_{\mathcal{L}_{\text{latent-prior}}} + \underbrace{E_C(x)}_{\mathcal{L}_{\text{energy}}} \right]. \quad (10)$$

Figure 3 gives an illustration of our process, and Algorithm 2 describes the algorithmic details.

PromptGen in a class-embedding space (Figure 4(g)) Previous works [4, 64] have shown that *class conditioning* boosts generative models' performances on ImageNet [62]. Specifically, class-conditioned generative models map a latent code z and a class embedding y to $x = G(z, y)$. To extend PromptGen to these models, we train an INN h_θ to map $\xi \sim \mathcal{N}(\mu, \sigma^2 \mathbf{I})$ to $y = h_\theta(\xi)$, where μ and σ are the mean and standard deviation of G 's class embeddings. The motivation for defining the distribution of ξ as $\mathcal{N}(\mu, \sigma^2 \mathbf{I})$ but not $\mathcal{N}(\mathbf{0}, \mathbf{I})$ is that we want the learned INN h_θ to be volume-preserving, which is easier to train. The training objective is to minimize $\mathbb{D}_{\text{KL}}(p_\theta(z, y)||p(z, y|\mathcal{C}))$,

which is equivalent to (full derivations in Appendix B.4):

$$\arg \min_{\theta} \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \xi \sim \mathcal{N}(\mu, \sigma^2 \mathbf{I}), z = f_{\theta}(\epsilon), y = h_{\theta}(\xi), x = G(z, y)} \left[-\log \left| \det \left(\frac{\partial f_{\theta}}{\partial \epsilon} \right) \right| - \log p_z(z) \right. \\ \left. - \log \left| \det \left(\frac{\partial h_{\theta}}{\partial \xi} \right) \right| - \log p_y(y) + E_C(x) \right], \quad (11)$$

where $p_y(y) = \mathcal{N}(y | \mu, \sigma^2 \mathbf{I})$ is the estimated class-embedding distribution.

PromptGen with conditional INN To generalize the control to continuous values, e.g., the scene parameters ρ in Eq. (4), we condition the INN on ρ . The condition is modeled by replacing $z = f_{\theta}(\epsilon)$ with $z = f_{\theta}(\epsilon, \rho)$. Space limited, we provide details and derivations in Appendix B.5. This extension results in a similar architecture to StyleFlow [1], but StyleFlow [1] uses MLE training on image-label pairs and is only applicable to explicit condition, which is not capable of modeling EBMs. When the condition is modeled by the equivariant operations proposed by [79], PromptGen also satisfies the homomorphism property [79] in the latent space.



Figure 4: PromptGen is applicable to different generative models in various domains. Figure 4(a), Figure 4(b), and Figure 4(c) are PromptGen applied to StyleNeRF [21], NVAE [75], and diffusion autoencoder [57], with text description a photo of a baby. Figure 4(d) shows PromptGen applied to StyleGAN2 [36] trained on AFHQ-Cats [6], with text description a photo of a British shorthair cat. Figure 4(e) and Figure 4(f) are PromptGen applied to StyleGAN2 trained on Landscape-HQ [70], with text descriptions a photo of {autumn, winter} scene. Figure 4(g) is the extension to the embedding space of BigGAN [4] on ImageNet [62]. All images are resized for visualization. See Appendix C for results on more datasets and text descriptions.



Figure 5: Image synthesis based on text description, guided by the CLIP model. As the pre-trained generative model, we use StyleGAN2 trained on FFHQ 1024² [35] with truncation $\psi = 0.7$. The text description used in this experiment is a photo of a baby. StyleCLIP requires optimization at inference, while StyleGAN2-NADA and PromptGen do not. All images are resized for visualization.

4 Experiments

This section describes the experimental validation of PromptGen. See Appendix A for experiments on synthetic data and Appendix C, D, E, and F for additional experiments on images and 3D meshes. Additional experimental details are provided in Appendix B.8.

4.1 Image Synthesis based on Text Description

This experiment illustrates the capability of PromptGen to sample images from a generative model driven by a text description t using the pre-trained CLIP model [58] (the ViT-B/32 version). We used the CLIP energy from Eq. (3), with text descriptions such as a photo of a baby.

Figure 5 shows a comparison between PromptGen and two previous CLIP-guided image generation methods, StyleCLIP [56] and StyleGAN2-NADA [17]. We observe that PromptGen generates diverse and high-quality images of babies, while StyleCLIP struggles in controllability and image quality, and StyleGAN2-NADA generates baby-like adults.² These results show that (1) locally editing the latent code (i.e., StyleCLIP) is not always an effective method for controlling generative models (e.g., not all images’ latent code can be locally edited into a baby); (2) domain adaptation (i.e., StyleGAN-NADA) is not effective in seeking modes in a generative model. Figure 4(a), Figure 4(b), and Figure 4(c) show how PromptGen applies to StyleNeRF [21], NVAE [75], and diffusion autoencoder (DiffAE; [57]). Although these generative models are all trained on FFHQ, baby images sampled from them have distinct characteristics. Figures 4(d)-4(f) show PromptGen applied to cats and landscape generation. Figure 4(g) shows the extension (Section 3.2) of our PromptGen to the class-embedding space of BigGAN [4], a class-conditional GAN trained on ImageNet [62]. We observe that PromptGen helps BigGAN generate images with complex text descriptions, which are out of BigGAN’s training image distribution; however, the diversity seems to be limited in this extension.

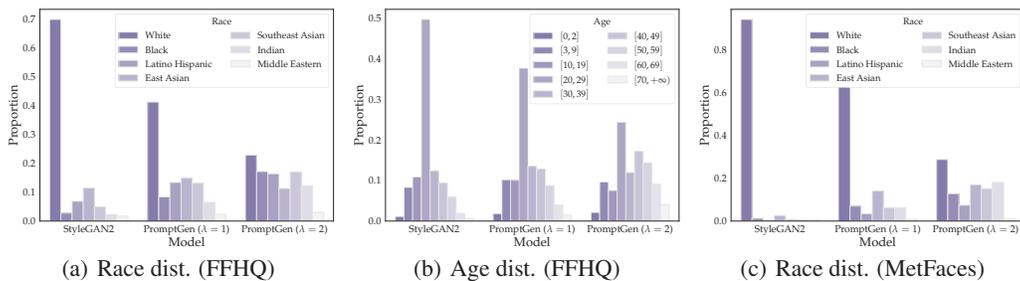


Figure 6: With the moment constraint, PromptGen de-biases StyleGAN2 (on FFHQ and MetFaces 1024², truncation $\psi = 0.7$). See Appendix D for image samples and Table 5 for quantitative results.

4.2 De-Biasing Pre-Trained Generative Models

An important problem in generative models is to generate fair distributions w.r.t a set of attributes of interest. For instance, Figure 6(a) shows that StyleGAN2 generates images with bias across races and

²Since we care about distributions, *none of the images in this paper are cherry- or lemon-picked.*

ages. PromptGen de-biases StyleGAN2 models trained on FFHQ 1024² and MetFaces 1024² [33] in terms of *categorical* attributes, using the moment constraint defined in Eq. (5) and Eq. (6). As control, we used a classifier trained on FairFace 224² [32] as γ . We defined $\mu = (|\mathcal{A}|^{-1}, \dots, |\mathcal{A}|^{-1})$, where \mathcal{A} is the set of races. Similar to the energy weights λ_i defined in Eq. (2), we propose to rescale the trained $\hat{\beta}$ as $\lambda\hat{\beta}$. Figure 6 shows that PromptGen de-biases the race and age effectively.

Existing de-biasing baselines consider *binary* attributes [31]. For a fair comparison with them, we adopted their setting to use binary classifiers trained on CelebA [47] for de-biasing and evaluation. Since classifiers trained on CelebA also suffer from the spurious correlation between attributes, we did *not* use the moment constraint for this experiment. Instead, since PromptGen allows conditional image generation with the classifier energy, we generated the same number of samples conditioned on each attribute or attribute combination. Table 2 shows that PromptGen has competitive performance on de-biasing for attributes and attribute combinations.

Table 2: Comparison with baselines for de-biasing binary attributes and their correlations. Following [31], we use classifiers on CelebA [47]. Baseline performances are copied from [31]. PromptGen has competitive performance, even in the cases where FairStyle achieves nearly perfect performance.

	FFHQ (binary attributes)					
	$\mathbb{D}_{\text{KL}}^{\text{gender}} \downarrow$	$\mathbb{D}_{\text{KL}}^{\text{eyeglasses}} \downarrow$	$\mathbb{D}_{\text{KL}}^{\text{blond hair}} \downarrow$	$\mathbb{D}_{\text{KL}}^{\text{age+gender}} \downarrow$	$\mathbb{D}_{\text{KL}}^{\text{age+eyeglasses}} \downarrow$	$\mathbb{D}_{\text{KL}}^{\text{gender+eyeglasses}} \downarrow$
FFHQ (real data)	0.015	0.186	–	0.246	0.355	0.242
StyleGAN2 [36]	0.018	0.180	–	0.279	0.384	0.250
StyleFlow [1]	0.023	0.061	–	0.214	0.162	0.121
FairGen [72]	4.21×10^{-4}	7.07×10^{-4}	–	0.0373	0.0330	0.00185
FairStyle [31]	3.20×10^{-7}	0	–	0.0257	0.0157	0.000241
PromptGen (ours)	1.71×10^{-5}	1.72×10^{-5}	0.0008	0.000558	0.000415	0.000628



Figure 7: Using inverse graphics model DECA [16] and scene parameters $\rho_{\text{pose}} \in \text{SO}(3)$, PromptGen controls the pose of StyleGAN2 while preserving the identity. All images are resized for visualization.

Table 3: Pose-controlled face generation. In this experiment, PromptGen uses StyleGAN2 as the pre-trained generative model. Results for GRAF [66], pi-GAN [5], GIRAFFE [53], and StyleNeRF [21] are from [21]; results for DiscoFaceGAN (DFG) [9] and GAN-Control [69] are from [69].

	StyleGAN2	GRAF	pi-GAN	GIRAFFE	StyleNeRF	DFG	GAN-Control	PromptGen
Resolution	1024 ²	256 ²	256 ²	256 ²	1024 ²	256 ²	512 ²	1024 ²
Pose	✗	✓	✓	✓	✓	✓	✓	✓
FID↓	3	71	85	35	8	13	6	4
Dist. w/ same ID↓	–	–	–	–	–	0.83	0.68	0.45
Dist. w/ diff. ID	–	–	–	–	–	1.73	1.90	1.37

4.3 Pose-Guided Face Synthesis

With an inverse graphics model, PromptGen can control the pose of faces generated by StyleGAN2. We used the DECA model [16], which infers the parameters of FLAME [44], a parametric facial graphics model. We set $\rho \in \text{SO}(3)$ as FLAME’s three neck poses and used the conditional INN extension introduced in Section 3.2. To enable generating different poses of the same identity (ID), we propose ID energy using the IR-SE50 model [8]. Specifically, given a canonical pose ρ_0 , we define $z_0 = f_{\theta}(\epsilon, \rho_0)$, and the ID energy is defined as (detailed in Appendix B.7)

$$E_{\text{ID}}(x_0, x) = 1 - \cos \langle R(x_0), R(x) \rangle, \quad x_0 = G(z_0), x = G(z), \quad (12)$$

where R is the IR-SE50 model [8] that computes face embeddings. Figure 7 shows that PromptGen generates faces of the same ID in different poses, even without being explicitly trained with poses as conditions. We computed the FID score [24] of each model using Clean-FID [55]; following [36], we did not use the truncation trick when computing the FID score. Table 3 shows that PromptGen outperforms existing models in terms of the FID score. Following [69], we then reported the average IR-SE50 [8] embedding distances for images with the same ID and with different IDs. Results show that PromptGen achieves the best ID preservation, with a slight sacrifice of ID diversity.

4.4 Iterative Distributional Control via Functional Composition

This section discusses an interesting bias that PromptGen reveals about the CLIP model. Figure 8(a) shows images generated by PromptGen (with StyleGAN2) with the CLIP model and the text description a photo of a person without makeup, where more females are generated than males. This bias should not be attributed to *image pre-training data* since images contain a bias in the opposite direction, i.e., men are less likely to have makeup. We argue that this bias should be explained by CLIP having learned a “reporting bias” in *vision-language pre-training data*: people are more likely to *say* “a person without makeup” when the person is a female (detailed analysis in Appendix F).

Besides revealing the above “reporting bias”, PromptGen can also mitigate this bias via an iterative control, enabled by the functional composition in Algorithm 1. Specifically, in the second iteration, we de-biased the gender distribution of $G \circ f_\theta$ instead of G , where f_θ is the INN learned for the text control. For de-biasing, we used the moment constraint with $\hat{\beta}$ trained for $G \circ f_\theta$. Figure 8(b) shows that females and males are uniformly distributed after the moment constraint in the second iteration.



(a) PromptGen (iteration 1) with text a photo of a person without makeup. Gender distribution: female: 81.6%; male: 18.4%.



(b) PromptGen (iteration 2). Gender de-biasing ($\lambda = 2$) for the distribution learned in iteration 1. Gender distribution: female: 49.3%; male: 50.7%.

Figure 8: A curious case of the CLIP model: with description a photo of a person without makeup, PromptGen generates more female images than male, showing that CLIP learns a “reporting bias”. Besides revealing this “reporting bias”, PromptGen can also mitigate this bias via an iterative control, allowing us to de-bias the text-controlled distribution. All images are resized for visualization.

4.5 Inference Latency

We compared our PromptGen and the plug-and-play generative model (PPGM) [49] in terms of the inference latency. PPGM uses Langevin dynamics [76] to optimize over the latent space at inference, while PromptGen samples images in a feed-forward manner. In this experiment, we used PPGM and our PromptGen to approximate the same EBM with CLIP energy. Inference times were estimated on

an NVIDIA RTX A4000 GPU. Table 4 shows that our PromptGen has the highest performance and efficiency. Compared to our PromptGen, PPGM requires $100\times$ inference time to achieve comparable results. Moreover, learning f_θ in PromptGen requires training-time optimization, and this amortized optimization is useful when one wants to reuse a controlled distribution many times.

Table 4: Comparison between PromptGen and PPGM, when approximating the *same* EBM with CLIP energy. n is the number of inference-time optimization steps used by PPGM. All models used StyleGAN2 as the pre-trained generative model. The text descriptions are a photo of a {baby, boy, girl}.

	PPGM ($n = 10$)	PPGM ($n = 50$)	PromptGen (ours)
CLIP energy (baby)↓	0.7327	0.7134	0.7038
CLIP energy (girl)↓	0.7257	0.7184	0.7199
CLIP energy (boy)↓	0.7263	0.7114	0.7081
Inference time per sample (sec.)↓	4.4	21.5	0.2
Back-propagation through CLIP at inference	10	50	0
When is it equal to the EBM?	$n \rightarrow \infty$	$n \rightarrow \infty$	$\mathbb{D}_{\text{KL}}(p_\theta(z) p(z C)) = 0$

5 Conclusions and Future Work

This paper proposes PromptGen, a unified framework to learn latent distributions for distributional control of pre-trained generative models, by leveraging the knowledge of other off-the-shelf models. Unlike previous approaches, which require optimization for generating each sample at inference, PromptGen trains an invertible neural network to approximate the latent distribution. Thus, PromptGen can be viewed as an amortization of the inference-time optimization into a feed-forward neural network, allowing for efficient, feed-forward sampling at inference. Further, this amortization also makes PromptGen not require the availability of those off-the-shelf models at inference. In practice, PromptGen offers generality for algorithmic design and modularity for control composition, and it also enables iterative controls via functional composition. Experiments validate that PromptGen applies to various generative models (StyleGAN2, StyleNeRF, diffusion autoencoder, and NVAE), control types (continuous, discrete, and moment constraint), off-the-shelf models (CLIP, classifiers, and inverse graphics models), and data domains (faces, churches, animals, ImageNet, and landscapes).

Limitation and future work: We provide an error analysis in Appendix F and a discussion on societal impact in Section 6. PromptGen is restricted by the pre-trained generative model’s coverage [3, 28]. Notably, since PromptGen focuses on mode-seeking and reweighting instead of domain adaptation, the support set of the output distribution will not change after PromptGen learning. We find it interesting to further explore combining PromptGen and domain adaptation [17] to achieve better generalization and adaptability. PromptGen depends on the off-the-shelf models that provide knowledge about the control (Appendix F); it can be beneficial to explore how to learn energy functions [15, 27] (besides our moment constraint) for fine-grained control with less bias.

6 Societal Impact

Like any technology, PromptGen has both benefits and drawbacks for society. On the positive side, we have demonstrated that PromptGen can uncover biases learned by text-image models (like CLIP) and to de-bias text-image models and pre-trained generative models, suggesting that PromptGen may be a helpful tool for fair AI if used properly. The efficient inference of PromptGen also helps reduce the computational expense, giving a positive impact on the environment. However, improved controllability makes it easier to synthesize targeted images; for instance, the creation of deceptive media such as DeepFakes [77, 74] and privacy leakage. To battle these cases, we expect researchers and practitioners to use technologies that can detect fake media and mitigate privacy leakage. We also encourage practitioners to consider these risks when using PromptGen to develop systems.

Acknowledgments and Disclosure of Funding

The authors would like to thank the anonymous reviewers, Zoltán Ádám Milacski, Jianchun Chen, and Shubhra Aich for their valuable feedback on drafts of this paper.

References

- [1] Rameen Abdal, Peihao Zhu, Niloy Jyoti Mitra, and Peter Wonka. StyleFlow: Attribute-conditioned exploration of StyleGAN-generated images using conditional continuous normalizing flows. *TOG*, 2021.
- [2] Lynton Ardizzone, Carsten Lüth, Jakob Kruse, Carsten Rother, and U. Köthe. Guided image generation with conditional invertible neural networks. *ArXiv*, 2019.
- [3] David Bau, Jun-Yan Zhu, Jonas Wulff, William S. Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a GAN cannot generate. *ICCV*, 2019.
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *ICLR*, 2019.
- [5] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis. *CVPR*, 2021.
- [6] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. StarGAN v2: Diverse image synthesis for multiple domains. *CVPR*, 2020.
- [7] Imre Csiszár and Paul C. Shields. Information theory and statistics: A tutorial. *Found. Trends Commun. Inf. Theory*, 2004.
- [8] Jiankang Deng, J. Guo, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. *CVPR*, 2019.
- [9] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3D imitative-contrastive learning. *CVPR*, 2020.
- [10] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. *NeurIPS*, 2021.
- [11] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: Non-linear independent components estimation. *ICLR, Workshop Track Proceedings*, 2015.
- [12] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. *ICLR*, 2017.
- [13] Yilun Du, Shuang Li, and Igor Mordatch. Compositional visual generation with energy based models. *NeurIPS*, 2020.
- [14] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *NeurIPS*, 2019.
- [15] Jesse H. Engel, Matthew D. Hoffman, and Adam Roberts. Latent constraints: Learning to generate conditionally from unconditional generative models. *ICLR*, 2018.
- [16] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. *TOG*, 2021.
- [17] Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. StyleGAN-NADA: CLIP-guided domain adaptation of image generators. *SIGGRAPH*, 2022.
- [18] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. *NIPS*, 2014.
- [19] Will Grathwohl, Kuan-Chieh Wang, Joern-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. *ICLR*, 2020.
- [20] Aditya Grover, Kristy Choi, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. *ICML*, 2020.
- [21] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A style-based 3D aware generator for high-resolution image synthesis. *ICLR*, 2022.

- [22] Yuxuan Gu, Xiaocheng Feng, Sicheng Ma, Lingyu Zhang, Heng Gong, and Bing Qin. A distributional lens for multi-aspect controllable text generation. *EMNLP*, 2022.
- [23] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. GANSpace: Discovering interpretable GAN controls. *NeurIPS*, 2020.
- [24] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. *NIPS*, 2017.
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020.
- [26] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *NeurIPS Workshop*, 2021.
- [27] Zhiting Hu, Zichao Yang, Ruslan Salakhutdinov, Xiaodan Liang, Lianhui Qin, Haoye Dong, and Eric P. Xing. Deep generative models with learnable knowledge constraints. *NeurIPS*, 2018.
- [28] Minyoung Huh, Richard Zhang, Jun-Yan Zhu, Sylvain Paris, and Aaron Hertzmann. Transforming and projecting images into class-conditional generative networks. *ECCV*, 2020.
- [29] Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. MaGNET: Uniform sampling from deep generative network manifolds without retraining. *ICLR*, 2022.
- [30] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge J. Belongie, Bharath Hariharan, and Ser Nam Lim. Visual prompt tuning. *ArXiv*, 2022.
- [31] Cemre Karakas, Alara Dirik, Eylül Yalçınkaya, and Pinar Yanardag. FairStyle: Debiasing StyleGAN2 with style channel manipulations. *ArXiv*, 2022.
- [32] Kimmo Kärkkäinen and Jungseock Joo. FairFace: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. *WACV*, 2021.
- [33] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *NeurIPS*, 2020.
- [34] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *NeurIPS*, 2021.
- [35] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *CVPR*, 2019.
- [36] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. *CVPR*, 2020.
- [37] Muhammad Khalifa, Hady Elsahar, and Marc Dymetman. A distributional approach to controlled text generation. *ICLR*, 2021.
- [38] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [39] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *NeurIPS*, 2018.
- [40] Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. *ICLR*, 2014.
- [41] Yann LeCun, Sumit Chopra, Raia Hadsell, Aurelio Ranzato, and Fu Jie Huang. A tutorial on energy-based learning. *Tutorial*, 2006.
- [42] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *EMNLP*, 2021.
- [43] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip H. S. Torr. Controllable text-to-image generation. *NeurIPS*, 2019.

- [44] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *TOG*, 2017.
- [45] Nan Liu, Shuang Li, Yilun Du, Joshua B. Tenenbaum, and Antonio Torralba. Learning to compose visual relations. *NeurIPS*, 2021.
- [46] Xingchao Liu, Chengyue Gong, Lemeng Wu, Shujian Zhang, Haoran Su, and Qiang Liu. Fuse-Dream: Training-free text-to-image generation with improved CLIP+GAN space optimization. *ArXiv*, 2021.
- [47] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. *ICCV*, 2015.
- [48] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. *CVPR*, 2020.
- [49] Anh M Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, and Jason Yosinski. Plug & play generative networks: Conditional iterative generation of images in latent space. *CVPR*, 2017.
- [50] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. *ArXiv*, 2021.
- [51] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *ICML*, 2021.
- [52] Weili Nie, Arash Vahdat, and Anima Anandkumar. Controllable and compositional generation with latent-space energy-based models. *NeurIPS*, 2021.
- [53] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. *CVPR*, 2021.
- [54] Frank Noé, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 2019.
- [55] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in GAN evaluation. *CVPR*, 2022.
- [56] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and D. Lischinski. StyleCLIP: Text-driven manipulation of StyleGAN imagery. *ICCV*, 2021.
- [57] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. *CVPR*, 2022.
- [58] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *ICML*, 2021.
- [59] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text Transformer. *JMLR*, 2020.
- [60] Vikram V. Ramaswamy, Sunnie S. Y. Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. *CVPR*, 2021.
- [61] Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley Edwards, Nicolas Manfred Otto Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A generalist agent. *ArXiv*, 2022.
- [62] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 2015.

- [63] Chitwan Saharia, William Chan, Huiwen Chang, Chris A. Lee, Jonathan Ho, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. *ArXiv*, 2021.
- [64] Axel Sauer, Katja Schwarz, and Andreas Geiger. StyleGAN-XL: Scaling StyleGAN to large diverse datasets. *SIGGRAPH*, 2022.
- [65] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. LAION-400M: Open dataset of CLIP-filtered 400 million image-text pairs. *ArXiv*, 2021.
- [66] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative radiance fields for 3D-aware image synthesis. *NeurIPS*, 2020.
- [67] Vikash Sehwal, Caner Hazirbas, Albert Gordo, Firat Ozgenel, and Cristian Canton Ferrer. Generating high fidelity data from low-density regions using diffusion models. *CVPR*, 2022.
- [68] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. InterFaceGAN: Interpreting the disentangled face representation learned by GANs. *TPAMI*, 2022.
- [69] Alon Shoshan, Nadav Bhonker, Igor Kviatkovsky, and Gérard Medioni. GAN-Control: Explicitly controllable GANs. *ICCV*, 2021.
- [70] Ivan Skorokhodov, Grigorii Sotnikov, and Mohamed Elhoseiny. Aligning latent and image spaces to connect the unconnectable. *ICCV*, 2021.
- [71] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *NeurIPS*, 2019.
- [72] Shuhan Tan, Yujun Shen, and Bolei Zhou. Improving the fairness of deep generative models without retraining. *ArXiv*, 2020.
- [73] Tijmen Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradient. *ICML*, 2008.
- [74] Cristian Vaccari and Andrew Chadwick. Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society*, 6, 2020.
- [75] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. *NeurIPS*, 2020.
- [76] Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. *ICML*, 2011.
- [77] Mika Westerlund. The emergence of Deepfake technology: A review. *Technology Innovation Management Review*, 9(11), 2019.
- [78] Jay Whang, Erik M. Lindgren, and Alexandros G. Dimakis. Composing normalizing flows for inverse problems. In *ICML*, 2021.
- [79] Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov, and Gabriel J. Brostow. Interpretable transformations with encoder-decoder networks. *ICCV*, 2017.
- [80] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. GAN inversion: A survey. *ArXiv*, 2021.
- [81] Tianbao Xie, Chen Henry Wu, Peng Shi, Ruiqi Zhong, Torsten Scholak, Michihiro Yasunaga, Chien-Sheng Wu, Ming Zhong, Pengcheng Yin, Sida I. Wang, Victor Zhong, Bailin Wang, Chengzu Li, Connor Boyle, Ansong Ni, Ziyu Yao, Dragomir Radev, Caiming Xiong, Lingpeng Kong, Rui Zhang, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. UnifiedSKG: Unifying and multi-tasking structured knowledge grounding with text-to-text language models. *ArXiv*, 2022.

- [82] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. FaceScape: A large-scale high quality 3D face dataset and detailed riggable 3D face prediction. *CVPR*, 2020.
- [83] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *ArXiv*, 2015.
- [84] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient GAN training. *NeurIPS*, 2020.
- [85] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *ArXiv*, 2021.
- [86] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. *CVPR*, 2022.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
In Section 5 and Appendix F.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes]
In Section 5.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
In Section 3 and Appendix B.
 - (b) Did you include complete proofs of all theoretical results? [Yes]
In Section 3 and Appendix B.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
We will make our code publicly available on GitHub, with instruction for each experiment. We made great efforts to ensure that most experiments will be run in one line of command.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
In Appendix B.8; other training details will be included in the code release.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
In Appendix B.8.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
We have cited the original papers.
 - (b) Did you mention the license of the assets? [Yes]
License will be included in the code.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
We do not have new assets.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
All data we use are published datasets.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
All data we use are published datasets.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
We do not have user study.
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
We do not have user study.
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]
We do not have user study.