

---

# Unsupervised Controllable Generation with Self-Training

---

Grigorios G Chrysos<sup>1</sup> Jean Kossaifi<sup>1 2</sup> Zhiding Yu<sup>2</sup> Anima Anandkumar<sup>2 3</sup>

## Abstract

Recent generative adversarial networks (GANs) are able to generate impressive photo-realistic images. However, controllable generation with GANs remains a challenging research problem. Achieving controllable generation requires semantically interpretable and disentangled factors of variation. We propose an unsupervised framework to learn a distribution of latent codes that control the generator through self-training. Self-training provides an iterative feedback in the GAN training, from the discriminator to the generator, and progressively improves the proposal of the latent codes as training proceeds. The latent codes are sampled from a latent variable model that is learned in the feature space of the discriminator. Our framework exhibits better disentanglement compared to other variants such as the variational autoencoder, and is able to discover semantically meaningful latent codes without any supervision.

## 1. Introduction

Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) are the method of choice for synthesis, owing to their ability to generate impressive photo-realistic images. Yet, they fall short in a key aspect of generation in real-world applications, *controllability* – the ability to control the semantic of the generated images in an interpretable, deterministic manner. Controllability will enable on-demand synthesis of images, which has numerous applications, including data augmentation and image editing. Controllable generation relies on semantically interpretable disentangled factors of variation, i.e., factors that modify a single mode of variation, such as length or color of the hair.

However, the unsupervised nature of GANs hinders the development of controllable generation. For instance, given a generator of a standard GAN model trained on facial synthe-

sis, it is not possible to directly control semantic attributes in a new synthesized instance, such as type/length/color of the hair, shape of the face, etc. Adding supervision means getting access to various (labelled) image attributes, which can be expensive. To reduce the amount of supervision required, Nie et al. (2020) consider semi-supervised learning in Style-GANs and reveal that limited amount of supervision is sufficient for high quality generation. This assumes, nonetheless, that all the semantic attributes are labeled, which may not always be achievable. Instead, it is preferable to achieve controlled generation in a fully unsupervised manner.

Unsupervised disentanglement has been explored in the Variational Autoencoders (VAEs) (Kim and Mnih, 2018; Chen et al., 2018a). However, the generation quality of VAEs does not yet match the quality of GAN-synthesized images. GAN methods have also been extended to achieve disentanglement of the factors of variation (Chen et al., 2016; Kaneko et al., 2018; Lee et al., 2020). However, these unsupervised GAN and VAE approaches have two major drawbacks: i) the disentangled factors are not guaranteed to be interpretable and ii) such models suffer from non-identifiability (Locatello et al., 2019), meaning that different runs can produce different factors.

Consequently, in this work, we propose ST-GAN, the first fully unsupervised approach to controllable generation with GANs through self-training<sup>1</sup>. Specifically, **we make the following contributions:**

- We propose a novel self-training procedure to discover disentangled and semantically interpretable latent codes driving the generation. The self training feedback loop allows for iterative refinement of the factor codes. We design a framework that encourages the model to produce interpretable factor codes that can faithfully control the generator.
- Instead of sampling from a fixed distribution probability distribution (e.g. Gaussian), we use a flexible latent variable model. Specifically, we employ a normalized independent component analysis models. To learn its parameters, we apply tensor factorization to the higher-order cumulants of the representation learnt in the feature space discriminator. We also experiment with a varia-

---

<sup>1</sup>Imperial College London <sup>2</sup>NVIDIA <sup>3</sup>California Institute of Technology. Correspondence to: Grigorios Chrysos <g.chrysos@imperial.ac.uk>.

---

<sup>1</sup>This work is also under review in NeurIPS.

tional autoencoder as an alternative latent variable model.

- We empirically demonstrate that our approach results in a controllable GAN able to learn the hidden codes in a fully unsupervised manner. We show on two different domains, cars and faces, that the discovered codes are disentangled and semantically interpretable, allowing to control the variations in the synthesized images.
- We propose two quantitative metrics for measuring the semantic changes by modifying a single element of the factor codes. Using them, we quantitatively demonstrate significant improvements with our model over the baselines. Our experiments exhibit how our framework can progressively improve the learning of latent codes as the training proceeds. We also establish the importance of each block of our model through extensive ablation studies.

## 2. Controlling generation and discovering the latent codes with self-training

Our method discovers the disentangled, semantically meaningful latent codes driving the generation in a fully unsupervised manner. This relies on augmenting a GAN structure towards controllable generation. Specifically, we make the following changes: i) instead of sampling from a fixed distribution, the generator takes as input latent codes from a normalized independent component analysis model, ii) a self-training scheme is proposed to discover the latent codes, iii) a hierarchical structure is used in the generator. Each contribution is analyzed below, while in Fig. 1 an abstract schematic of the framework is illustrated. We first review generative adversarial networks, before motivating and describing our proposed method.

**Notation:** Vectors (scalars) are denoted with boldface (plain) letters, e.g.  $\mathbf{x}(x)$ . The outer product is symbolized with  $\otimes$ .

### 2.1. Generative Adversarial Networks

**Generative Adversarial Networks** A GAN consists of a generator  $G$  and a discriminator  $D$  engaging in a zero sum game. The goal of the generator is to model the target distribution  $\mathbb{P}_{data}$ , while the discriminator aims at discerning the samples synthesized by the generator and the real samples from target (ground-truth) distribution.

The generator samples a latent code from a fixed probability distribution  $\mathbb{P}_z$  (typically Gaussian) and maps it to an image  $G(\mathbf{z})$ . The discriminator receives both images synthesized by the generator  $G(\mathbf{z})$  and samples from the real distribution  $\mathbb{P}_{data}$  and tries to distinguish them. The objective function

is:

$$\mathcal{L}_{gan} = \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{data}} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim \mathbb{P}_z} [\log(1 - D(G(\mathbf{z})))]. \quad (1)$$

This loss is optimized in an alternating manner as  $\min_G \max_D \mathcal{L}_{gan}$ .

**Style-GAN** proposes hierarchical injections in the generator (Karras et al., 2019). The injections, i.e. Hadamard products with the inputs, are performed after each layer. In practice, our generator differentiates from Style-GAN injections, as we partition the latent codes (i.e. input to the generator). Specifically, we use  $s$  partitions, where  $s$  denotes the total number of injections in the generator. That is, for a latent code  $\mathbf{h} \in \mathbb{R}^d$ , we use the elements  $[h_1, \dots, h_{\frac{d}{s}}]$  in the first injection, the next  $\frac{d}{s}$  elements in the second and so on. This injection captures higher-order correlations (Chrysos et al., 2019). Therefore, by partitioning the latent codes we capture the correlations of specific elements in each injection.

**Semi-Supervised Style-GAN** augments the training procedure of Style-GAN with two additional losses: a Mixup loss and a consistency loss (Nie et al., 2020). The consistency loss is expressed as  $\mathcal{L}_c = \|\mathbf{h} - \mathbf{W}^T v(G(\mathbf{h}))\|_2^2$ , where  $v$  is the representation from the penultimate layer of the discriminator,  $\mathbf{W}$  a learnable dense layer and  $\mathbf{h}$  is the input to the generator. The consistency loss constrains the original and the reconstructed latent codes (i.e.,  $\mathbf{h}$ ) to be close.

The Mixup loss reinforces smoothness in the latent code space. Given a pair of real and fake images  $(\mathbf{I}_r, \mathbf{I}_f)$  and their corresponding latent codes  $(\mathbf{h}_r, \mathbf{h}_f)$ , we interpolate between the images and the latent codes to obtain  $(\mathbf{I}_s, \mathbf{h}_s)$ . The Mixup loss is then  $\mathcal{L}_s = \|\mathbf{h}_s - \mathbf{W}^T v(\mathbf{I}_s)\|_2^2$ .

In this paper, we build on top of both Style-GAN and its semi-supervised variant and incorporate both losses when training our proposed model.

### 2.2. Driving generation with a latent variable model

In GAN, the generator takes as input a latent code, sampled from a fixed probability distribution. Typically, a Gaussian or a uniform distribution is selected. Replacing this fixed distribution with distributions with learnable parameters has been recently explored in the literature, both in the context of GAN (Kuznetsov et al., 2019) and other generative models (Tomczak and Welling, 2018; Bauer and Mnih, 2018). Simple fixed distributions, e.g. Gaussian, do not encode any semantic meaning such as 3D pose of the object.

Motivated by such works, we replace sampling the latent code from a fixed distribution with latent codes. The latent

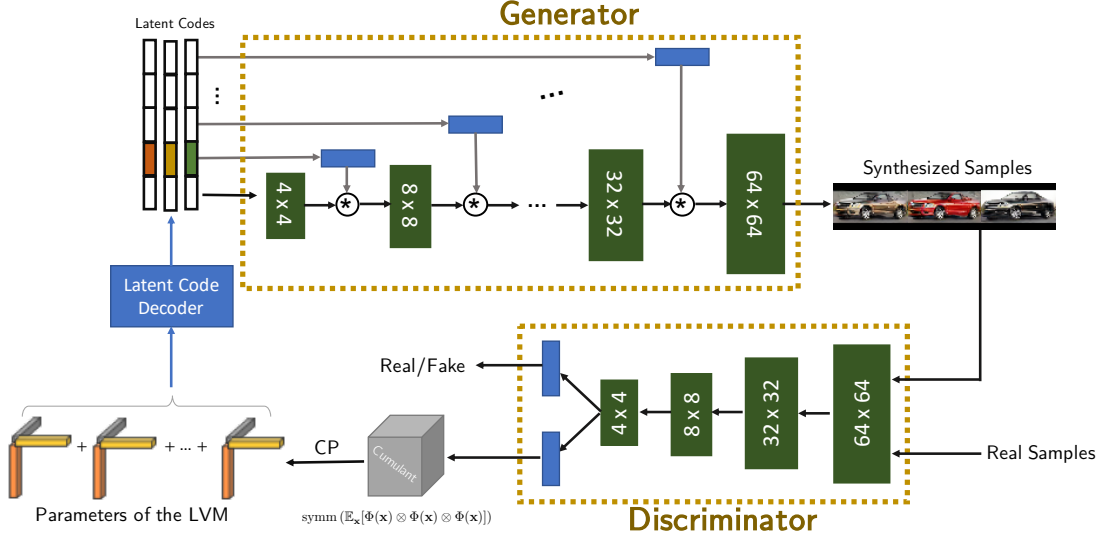


Figure 1: Overview of our proposed ST-GAN.

codes are sampled from a latent variable model. This underlying model is data-driven and evolves over the training to better reflect the variation of the specific data. In practice, we use a variant of independent component analysis, which produces latent codes  $\mathbf{h} \sim \mathbb{P}_l$  used by the generator. The independent component analysis (Comon, 1994) mixes linearly independent signals and then corrupts them with Gaussian noise. However, in this work we normalize the hidden independent signals in the range of  $[-1, 1]$ . A variant of this model with discrete observations is known as topic model, but we use continuous version and term it as normalized independent component analysis. If we denote the latent vector with  $\mathbf{h}$ , the Gaussian noise with  $\epsilon$  and the mixing matrix with  $\mathbf{M}$ , the corresponding observation is  $\mathbf{y} = \mathbf{M}\mathbf{h} + \epsilon$ . As an ablation, we also explore using a VAE.

### 2.3. Learning the latent variable model with self-training

To obtain the latent codes  $\mathbf{h}$ , we use a self-training scheme with a latent variable model (a flexible normalized independent component analysis model) in the feature space of the discriminator. We learn its parameters through tensor factorization of the higher-order cumulant. We also showcase in ablation that other latent variable models, such as VAEs can also be used as drop-in replacements. During inference, the generator samples latent codes  $\mathbf{h} \sim \mathbb{P}_l$  from the learnable distribution  $\mathbb{P}_l$  and generates samples  $G(\mathbf{h})$ .

**Extracting the representation** We learn the latent variable model (i.e., ICA) in the feature space of the discriminator. Specifically, for an input image  $\mathbf{x}$ , the corresponding features  $\mathbf{v}(\mathbf{x})$  are extracted, where  $\mathbf{v}(\cdot)$  corresponds to

the penultimate layer of the discriminator. The representation for the latent variable model is computed through a linear layer with learnable vector of parameters  $\mathbf{W}$ . The final extracted representation is then  $\Phi(\mathbf{x}) = \mathbf{W}^T \mathbf{v}(\mathbf{x})$ . The mixing matrix is learned through tensor factorization. Specifically, we form the higher-order moments of the features  $\Phi(\mathbf{x})$ , and learn their low-rank factorized form.

**Learning the parameters of the latent normalized independent component analysis model through tensor factorization** We assume a normalized independent component analysis model, which maps from the extracted features to untangled latent codes. We propose to learn the parameters of that model by decomposing the higher-order moments. By analogy, in the analytical case, we know it is possible to form the moments of the model. From these, the symmetric cumulant tensor (e.g.  $\text{Symm}(\mathbb{E}_x[\Phi(\mathbf{x}) \otimes \Phi(\mathbf{x}) \otimes \Phi(\mathbf{x})])$  for the third order cumulant) can be obtained by subtracting cross-order terms. Given a specific latent variable model, this symmetric tensor – or cumulant – can be obtained through a closed-form formula. Applying low-rank decompositions to it allows the recovery of the parameters of the model (Anandkumar et al., 2014).

In this paper, all the components are learned end-to-end, including the factors of the decomposition. Instead of exactly forming the cumulant tensor, following Arabshahi and Anandkumar (2016), we propose to also learn the weights of each of the cross-terms. Specifically, for the second order, this results in  $\mathbf{M}_1 + w_1 \mathbf{M}_2$ , where  $\mathbf{M}_1 = \mathbb{E}_x[\Phi(\mathbf{x}) \otimes \Phi(\mathbf{x})]$  and  $\mathbf{M}_2 = \mathbb{E}_x[\Phi(\mathbf{x})] \otimes \mathbb{E}_x[\Phi(\mathbf{x})]$ . Similarly, for the third order term, we form the term

$T_1 + w_2 T_2 + w_3 T_3 + w_4 T_4 + w_5 T_5$ , with:

$$\begin{aligned} T_1 &= \mathbb{E}_x[\Phi(x) \otimes \Phi(x) \otimes \Phi(x)], \\ T_2 &= \mathbb{E}_x[\Phi(x) \otimes \Phi(x)] \otimes \mathbb{E}_x[\Phi(x)] \\ T_3 &= \mathbb{E}_x[\Phi(x) \otimes \mathbb{E}_x[\Phi(x)] \otimes \Phi(x)] \\ T_4 &= \mathbb{E}_x[\Phi(x)] \otimes \mathbb{E}_x[\Phi(x) \otimes \Phi(x)] \\ T_5 &= \mathbb{E}_x[\Phi(x)] \otimes \mathbb{E}_x[\Phi(x)] \otimes \mathbb{E}_x[\Phi(x)] \end{aligned}$$

We then factorize the resulting higher-order moments computed from the features  $\Phi(x)$  from the discriminator. We assume that the cumulants,  $M_1 + w_1 M_2$  and  $T_1 + \sum_{k=1}^5 T_k$  admit a rank- $R$  low-rank representation. In other words, we express the (symmetric) cumulants as a weighted sum of  $R$  rank-1 tensors. We learn both the weights of the sum (collected in a vector  $\lambda$ ) and the factors  $a_1, \dots, a_R$  of the decomposition. Specifically, we minimize the following loss function  $\mathcal{L}_l$ :

$$\arg \max_{w_k, \lambda, a_j} \mathcal{L}_l = \left\| M_1 + w_1 M_2 - \sum_{j=1}^R \lambda_j a_j \otimes a_j \right\|_F^2 + \left\| T_1 + \sum_{k=2}^5 w_k T_k - \sum_{j=1}^R \lambda_j a_j \otimes a_j \otimes a_j \right\|_F^2$$

The learned parameters  $\lambda, a_j$  form the mixing matrix  $M$  for the normalized independent component analysis model. Specifically, the  $j^{th}$  row corresponds to the  $j^{th}$  factor, i.e.  $M_{j,:} = \lambda_j a_j$ . Then, for a sample  $x \sim \mathbb{P}_{data}$ , we obtain the latent code as  $h = M^\dagger(x - \epsilon)$  where  $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2 I)$  and  $M^\dagger$  denotes  $M$ 's pseudo-inverse.

To further encourage the disentanglement of the factors, we add an orthogonality regularization term in the loss, i.e.  $\mathcal{L}_l = \mathcal{L}_l + \gamma_o \sum_{j=1}^R \|a_j^T a_j - \mathbb{I}\|_F^2$  where  $\mathbb{I}$  is the identity matrix and  $\gamma_o$  is a hyper-parameter.

**Ablation: Variational Autoencoder** As an ablation, we investigate replacing the higher-order factorization framework with VAEs, i.e., a VAE is learned as an alternative latent variable model. Similarly to the normalized independent component analysis model, we learn the VAE on the feature space of the discriminator.

Specifically, we maximize the ELBO of the distribution in the feature space of the discriminator:  $\mathcal{L}_l = \mathbb{E}_{q(h|\Phi(x))} [\log p(\Phi(x)|h)] - D_{KL}(q(h|\Phi(x)) \| p(z))$ , where  $h$  is the latent ccode,  $q(h|\Phi(x))$  is the posterior distribution and  $D_{KL}$  computes the Kullback-Leibler divergence between two distributions.

**Training the model** The self-learning technique of our framework is sensitive to initialization; here, we elaborate on the details of the objective function and the training procedure. For the first 200 iterations, we use a 'warm-up'

of the weights by training a using only the GAN loss. The latent codes in the input of the generator are sampled from a Gaussian. Sequentially, the latent variable model parameters are inserted and the corresponding  $\mathcal{L}_l$  loss is added. At 2,000 iterations, the self-training scheme is added along with the remaining loss terms.

Our preliminary experiments demonstrated that a soft transition from the prior distribution to the latent distribution is beneficial. To that end, we use a soft transition with an annealing  $\kappa$  parameter, i.e.  $h = \kappa \cdot h_l + (1 - \kappa) \cdot h_p$  where  $\kappa$  starts from 0 and transitions to 1 after 8,000 iterations. The symbol  $h_l$  denotes a sample from the latent variable model, while  $h_p$  is a sample from the prior distribution used in the beginning of the training.

To further induce disentanglement, we add an additional "masking loss". The masking loss  $\mathcal{L}_m$  encourages each element of the latent code to change one attribute. To achieve that, we synthesize a pair of images with a predictable changes in their latent codes and try to predict what that change was. Given a latent code  $h$ , we duplicate it into  $\hat{h}$  and sequentially perturb each element of the latent code with uniform noise. The masking loss then tries to predict which element was modified from the  $\{\Phi(G(h)), \Phi(G(\hat{h}))\}$  embeddings.

The complete loss function used is  $\mathcal{L} = \mathcal{L}_{gan} + \gamma_l \mathcal{L}_l + \gamma_s \mathcal{L}_s + \gamma_c \mathcal{L}_c + \gamma_m \mathcal{L}_m$ , where  $\gamma_l, \gamma_s, \gamma_c$  and  $\gamma_m$  are regularization hyper-parameters.

### 3. Experiments

Below, we describe the experimental setup and the comparisons conducted with the proposed framework. We utilize both the popular CelebA (Liu et al., 2015) and Cars dataset for our experiments<sup>2</sup>. All the images are resized to  $64 \times 64$ .

**Metrics:** The well-established Frechet Inception Distance (FID) metric (Heusel et al., 2017) is chosen for the generation quality<sup>3</sup>. The metrics proposed for unsupervised disentanglement require an auxiliary encoder to be trained (Higgins et al., 2017; Kim and Mnih, 2018; Chen et al., 2018a), which is not available in our case. In addition, using direct ground-truth labels is discouraged (Locatello et al., 2019). We opt to report auxiliary semantic metrics that demonstrate the changes affected by each element of the latent code. Specifically, to achieve controllable generation each element of the factor code should modify a factor of variation of the data and be interpretable. To that end, we utilize the following metric: we sample i) an element to perturb and

<sup>2</sup>CelebA dataset reportedly includes gender and racial biases (Kärkkäinen and Joo, 2019), thus we encourage the development of better datasets to address the bias.

<sup>3</sup>The features from the pretrained Inception network of Paszke et al. (Paszke et al., 2019) are used.



ii) a perturbation in the range  $(-1, 1)$ . We add the perturbation, generate the two images and then compare them. We repeat this procedure for each element of the latent code for 10 perturbations. To measure the difference between each pair of images we use both the standard mean absolute error (MAE) and the LPIPS (Zhang et al., 2018) metric that correlates with the perceptual changes. A higher value in both, means that the single element of the latent code has made a larger (perceptual) change in the image.

**Implementation details:** Our implementation is based on the GAN architecture of Miyato et al. (2018). That is, both the generator and the discriminator include residual blocks, while the rest hyper-parameters (i.e. optimizer, hinge loss, learning rate) remain unchanged. The injections in the generator follow the implementation of Chrysos et al. (2019). The models are implemented using PyTorch (Paszke et al., 2017) and TensorLy for all tensor methods (Kossaifi et al., 2019). We used  $\gamma_l = 1, \gamma_s = 0.1, \gamma_c = 0.1$ . The hyper-parameter  $\gamma_m$  is augmented during the training; in iteration 2,000 it starts as  $\gamma_m = 1$  and in iteration 10,000 it takes the value  $\gamma_m = 100$ . The dimensionality of the latent code is  $d = 50$ , i.e.,  $\mathbf{h} \in \mathbb{R}^{50}$ . We apply an elementwise normalization by the max element in the latent codes before feeding them into the generator. The cumulant computation is based on Tensorly (Kossaifi et al., 2019).

### 3.1. Comparison with other models

Two baseline architectures are considered: i) SNGAN (Miyato et al., 2018) (referred to as *GAN* henceforth), ii) VAE with convolutional encoder and decoder. For the SNGAN we consider two additional variants: a) one where we include injections in the generator (Karras et al., 2019) (referred to as *GAN-Inj*), b) one with a similar hierarchical generator as in our model (*GAN-H*). The goal is to assess whether the controllable generation is caused by these inductive biases alone. Our model is denoted as ‘ST-GAN’, while the ablation with the VAE model in the latent space as ‘ST-GAN-VAE’.

The quantitative results in Table 1 exhibit that our method outperforms all the baselines in both the MAE and the LPIPS metric. Intuitively, this means that each element of the factor code makes larger changes to the image than the baselines. To illustrate the differences, we visualize the linear interpolation in one element of the latent code at a time. In Fig. 2, each row depicts linear interpolation in a single element of the latent code. The rest elements remain the same. We notice how our model (‘ST-GAN’) changes the pose or the color in any object by using a single element.



(a) ST-GAN

(b) ST-GAN-VAE

Figure 2: Each row depicts linear interpolation in a single element of the latent code. The rest elements remain the same. (a) The first row corresponds to the elements in the first split, while the second row corresponds to an element from the second split. The first row changes the pose and the second the colors (of hair and car respectively). (b) Similar patterns with the VAE model are observed, however the changes are less obvious.

## 4. Related work

**Controllable generation in GAN** : Both GAN and VAE models have been utilized for achieving controllability, while some studies even indicate that the techniques in one do not work in the other (Lin et al., 2019). While the majority of existing work on disentanglement focuses on a (semi-)supervised setting (Tran et al., 2017; Zhang et al., 2019; Xiao et al., 2017; Reed et al., 2014; Szabó et al., 2017; Deng et al., 2017; Chongxuan et al., 2017; Tran et al., 2019a), our work focuses on the unsupervised setting. Here, we review the most closely related methods below. Such methods can be classified into three categories: a) post-training interpretable methods, b) information theory motivated disentanglement, c) hierarchical disentanglement. The post-processing methods (Plummer et al., 2020; Voynov and Babenko, 2020) assume a pretrained generator and find interpretable directions in the latent space of the generator. Post-processing methods have twofold drawbacks: i) they do not have any guarantees that we will discover such interpretable directions, ii) multi-step training is required.

The perspective of information theory is frequently used to tackle (unsupervised) disentanglement. The core idea relies on maximizing the mutual information between the latent codes and the synthesized images. The seminal work of InfoGAN (Chen et al., 2016) along with its extensions (Lin et al., 2019; Liu et al., 2020; Jeon et al., 2018) use the mutual information to disentangle the factors of variation in a GAN setting. The drawback of mutual information is that they require strong inductive biases to disentangle the factors of variation (Lin et al., 2019).

A line of work that relates to our method is that of hierarchical disentanglement. FineGAN (Singh et al., 2019) utilizes bounding boxes to define a hierarchical disentanglement. Kaneko et al. (2018) use instead class-level supervision to form the upper level of a hierarchy. The lower-levels are

Table 1: Performance of GAN models in terms of image quality (FID) and controllability (LPIPS, MAE). ST-GAN consistently provides better controllability as indicated from the reported metrics of LPIPS and MAE.

| Method                    | CelebA               |                    |                      | Cars                 |                    |                      |
|---------------------------|----------------------|--------------------|----------------------|----------------------|--------------------|----------------------|
|                           | FID ( $\downarrow$ ) | MAE ( $\uparrow$ ) | LPIPS ( $\uparrow$ ) | FID ( $\downarrow$ ) | MAE ( $\uparrow$ ) | LPIPS ( $\uparrow$ ) |
| GAN (Miyato et al., 2018) | 23.21                | 0.169              | 0.024                | 23.96                | 0.243              | 0.061                |
| GAN-Inj                   | <b>20.07</b>         | 0.179              | 0.026                | 15.06                | 0.240              | 0.049                |
| GAN-H                     | 21.17                | 0.178              | 0.024                | <b>12.13</b>         | 0.218              | 0.042                |
| VAE                       | 91.20                | 0.003              | 0.003                | 120.02               | 0.003              | 0.003                |
| <b>ST-GAN</b>             | 25.21                | <b>0.265</b>       | <b>0.040</b>         | 12.96                | <b>0.271</b>       | <b>0.057</b>         |
| ST-GAN-VAE                | 23.95                | 0.166              | 0.027                | 14.19                | 0.207              | 0.041                |

trained with strong inductive bias (e.g. one-hot vectors) and tailored curriculum learning methods. A core difference is that in the related works, they use some form of weak supervision.

The work that is most closely related to our work is that of Lee et al. (2020) which combines VAE and GAN to achieve both high quality of synthesis and unsupervised disentanglement. However, the proposed method requires significant engineering (two-step training process), while it demonstrates only partial disentanglement in high-dimensional distributions. In addition, the latent codes from the latent variable model (VAE) is used as an auxiliary input to the GAN generator. By contrast, in our approach, generation is controlled solely by the learned latent variable model.

**Latent variable models** : Learning the joint distribution over both observed and latent variables is a crucial topic in machine learning. The normalizing flows (Dinh et al., 2014; 2017) provide an elegant way to learn the joint distribution; they learn an invertible network. However, their likelihood-based loss results in images that are not competitive to recent generative models. Recently, Khemakhem et al. (2020) prove that a prerequisite for learning the joint distribution is to have an identifiable model. They make a step further and modify VAE to achieve threefold goals: (i) fit the target distribution, (ii) disentanglement of the variables, (iii) identifiability of the model. They extend VAE, since they argue that it state-of-the-art in (i) and (ii). However, Nie et al. (2020) exhibit that GANs outperform VAEs in semi-supervised disentanglement on challenging high-resolution data. This motivates us to use latent variable models in a GAN for realistic generation. To that end, we use moments to fit effective latent variable models for capturing salient features in images.

**Self-supervised learning and GAN** : The concept of self-supervised learning has emerged in a GAN training (Chen et al., 2018b; Huang et al., 2020; Qian et al., 2018). Chen et al. (2019) augment the GAN loss with an auxiliary loss to

predict rotated versions of the image in the discriminator. In (Tran et al., 2019b) they propose a self-supervised loss that can stabilize the training, while Arbel et al. (2020) involve both the generator and the discriminator for synthesizing samples in the trained model. Self-supervised GANs differ substantially from our goal; they predetermine a set of fixed auxiliary tasks (e.g. rotation), while in our case the goal is to drive the semantic generation through interpretable codes.

## 5. Conclusion

In this paper, we proposed the first self-trained GAN to enable controllable image generation without any supervision. Our method discovers disentangled and semantically interpretable factors of variation driving the generation in a fully unsupervised framework. The hidden factors are modelled explicitly by a flexible latent variable model, from which the generator samples its inputs. The parameters of that latent variable model are learned through a tensor factorization of higher-order cumulants. We demonstrate that the codes learned by our model are semantically interpretable.

## Acknowledgements

We would like to thank Weili Nie, Timo Aila and Sanja Fidler as well as the entire AiAlgo team for their valuable feedback. GC conducted this research while at Nvidia.

## References

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems (NeurIPS)*, 2014.
- Weili Nie, Tero Karras, Animesh Garg, Shoubhik Debhat, Anjul Patney, Ankit B Patel, and Anima Anandkumar. Semi-supervised stylegan for disentanglement learning. *arXiv*, pages arXiv–2003, 2020.

- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. *arXiv preprint arXiv:1802.05983*, 2018.
- Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in neural information processing systems (NeurIPS)*, pages 2610–2620, 2018a.
- Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems (NeurIPS)*, pages 2172–2180, 2016.
- Takuhiro Kaneko, Kaoru Hiramatsu, and Kunio Kashino. Generative adversarial image synthesis with decision tree latent controller. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6606–6615, 2018.
- Wonkwang Lee, Donggyun Kim, Seunghoon Hong, and Honglak Lee. High-fidelity synthesis with disentangled representation. *arXiv preprint arXiv:2001.04296*, 2020.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *International Conference on Machine Learning (ICML)*, 2019.
- Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Grigorios Chrysos, Stylianos Moschoglou, Yannis Panagakis, and Stefanos Zafeiriou. Polygan: High-order polynomial generators. *arXiv preprint arXiv:1908.06571*, 2019.
- Maxim Kuznetsov, Daniil Polykovskiy, Dmitry P Vetrov, and Alex Zhebrak. A prior of a googol gaussians: a tensor ring induced prior for generative models. In *Advances in Neural Information Processing Systems*, pages 4104–4114, 2019.
- Jakub M Tomczak and Max Welling. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- Matthias Bauer and Andriy Mnih. Resampled priors for variational autoencoders. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- Pierre Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15:2773–2832, 2014.
- Forough Arabshahi and Animashree Anandkumar. Spectral methods for correlated topic models. *arXiv preprint arXiv:1605.09080*, 2016.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *International Conference on Computer Vision (ICCV)*, pages 3730–3738, 2015.
- Kimmo Kärkkäinen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age. *arXiv preprint arXiv:1908.04913*, 2019.
- 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. In *Advances in neural information processing systems (NeurIPS)*, pages 6626–6637, 2017.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems (NeurIPS)*, pages 8024–8035, 2019.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. *International Conference on Learning Representations (ICLR)*, 2(5):6, 2017.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS-W*, 2017.
- Jean Kossaifi, Yannis Panagakis, Anima Anandkumar, and Maja Pantic. Tensorly: Tensor learning in python. *Journal of Machine Learning Research*, 20(1):925–930, 2019.

- Zinan Lin, Kiran Koshy Thekumparampil, Giulia Fanti, and Sewoong Oh. Infogan-cr: Disentangling generative adversarial networks with contrastive regularizers. *arXiv preprint arXiv:1906.06034*, 2019.
- Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 3, page 7, 2017.
- Jianfu Zhang, Yuanyuan Huang, Yaoyi Li, Weijie Zhao, and Liqing Zhang. Multi-attribute transfer via disentangled representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9195–9202, 2019.
- Taihong Xiao, Jiapeng Hong, and Jinwen Ma. Dnagan: Learning disentangled representations from multi-attribute images. *arXiv preprint arXiv:1711.05415*, 2017.
- Scott Reed, Kihyuk Sohn, Yuting Zhang, and Honglak Lee. Learning to disentangle factors of variation with manifold interaction. In *International Conference on Machine Learning (ICML)*, pages 1431–1439, 2014.
- Attila Szabó, Qiyang Hu, Tiziano Portenier, Matthias Zwicker, and Paolo Favaro. Challenges in disentangling independent factors of variation. *arXiv preprint arXiv:1711.02245*, 2017.
- Zhijie Deng, Hao Zhang, Xiaodan Liang, Luona Yang, Shizhen Xu, Jun Zhu, and Eric P Xing. Structured generative adversarial networks. In *Advances in neural information processing systems (NeurIPS)*, pages 3899–3909, 2017.
- LI Chongxuan, Taufik Xu, Jun Zhu, and Bo Zhang. Triple generative adversarial nets. In *Advances in neural information processing systems (NeurIPS)*, pages 4088–4098, 2017.
- Linh Tran, Jean Kossaifi, Yannis Panagakis, and Maja Pantic. Disentangling geometry and appearance with regularised geometry-aware generative adversarial networks. *International Journal of Computer Vision*, 127(6-7):824–844, 2019a.
- Antoine Plummerault, Hervé Le Borgne, and Céline Hudelot. Controlling generative models with continuous factors of variations. In *International Conference on Learning Representations (ICLR)*, 2020.
- Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. *arXiv preprint arXiv:2002.03754*, 2020.
- Bingchen Liu, Yizhe Zhu, Zuohui Fu, Gerard de Melo, and Ahmed Elgammal. Oogan: Disentangling gan with one-hot sampling and orthogonal regularization. In *AAAI Conference on Artificial Intelligence*, 2020.
- Insu Jeon, Wonkwang Lee, and Gunhee Kim. Ib-gan: Disentangled representation learning with information bottleneck gan. 2018.
- Krishna Kumar Singh, Utkarsh Ojha, and Yong Jae Lee. Finegan: Unsupervised hierarchical disentanglement for fine-grained object generation and discovery. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6490–6499, 2019.
- Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In *International Conference on Learning Representations (ICLR)*, 2017.
- Ilyes Khemakhem, Diederik P Kingma, and Aapo Hyvärinen. Variational autoencoders and nonlinear ica: A unifying framework. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- Ting Chen, Xiaohua Zhai, and Neil Houlsby. Self-supervised gan to counter forgetting. *arXiv preprint arXiv:1810.11598*, 2018b.
- Rui Huang, Wenju Xu, Teng-Yok Lee, Anoop Cherian, Ye Wang, and Tim Marks. Fx-gan: Self-supervised gan learning via feature exchange. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 3194–3202, 2020.
- Sheng Qian, Wen-ming Cao, Rui Li, Si Wu, et al. Self-supervised gan for image generation by correlating image channels. In *Pacific Rim Conference on Multimedia*, pages 78–88. Springer, 2018.
- Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-supervised gans via auxiliary rotation loss. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12154–12163, 2019.
- Ngoc-Trung Tran, Viet-Hung Tran, Bao-Ngoc Nguyen, Linxiao Yang, et al. Self-supervised gan: Analysis and improvement with multi-class minimax game. In *Advances in neural information processing systems (NeurIPS)*, pages 13232–13243, 2019b.
- Michael Arbel, Liang Zhou, and Arthur Gretton. Kale: When energy-based learning meets adversarial training. *arXiv preprint arXiv:2003.05033*, 2020.