

Learning Disentangled Representations with Latent Variation Predictability

Xinqi Zhu, Chang Xu, and Dacheng Tao

UBTECH Sydney AI Centre, School of Computer Science, Faculty of Engineering,
The University of Sydney, Darlington, NSW 2008, Australia
{xzhu7491@uni., c.xu@, dacheng.tao@}sydney.edu.au

Abstract. Latent traversal is a popular approach to visualize the disentangled latent representations. Given a bunch of variations in a single unit of the latent representation, it is expected that there is a change in a single factor of variation of the data while others are fixed. However, this impressive experimental observation is rarely explicitly encoded in the objective function of learning disentangled representations. This paper defines the *variation predictability* of latent disentangled representations. Given image pairs generated by latent codes varying in a single dimension, this varied dimension could be closely correlated with these image pairs if the representation is well disentangled. Within an adversarial generation process, we encourage variation predictability by maximizing the mutual information between latent variations and corresponding image pairs. We further develop an evaluation metric that does not rely on the ground-truth generative factors to measure the disentanglement of latent representations. The proposed variation predictability is a general constraint that is applicable to the VAE and GAN frameworks for boosting disentanglement of latent representations. Experiments show that the proposed variation predictability correlates well with existing ground-truth-required metrics and the proposed algorithm is effective for disentanglement learning.

1 Introduction

Nowadays learning interpretable representations from high-dimensional data is of central importance for downstream tasks such as classification [10, 13, 20], domain adaptation [38, 46, 43], fair machine learning [11, 30], and reasoning [42]. To achieve this goal, a series of work have been conducted [10, 18, 24, 13, 9, 20] under the subject of *disentangled representation learning*. Although there is not a widely adopted mathematical definition for disentanglement, a conceptually agreed definition can be expressed as that each unit of a disentangled representation should capture a single interpretable variation of the data [5].

One line of current most promising disentanglement methods derives from β -VAE [18], with its variants such as FactorVAE [24] and β -TCVAE [9] developed later. This series of works mainly realize disentanglement by enforcing the independence in the latent variables. Although independence assumption is an





$\Delta \mathbf{z} = \mathbf{z}_1 - \mathbf{z}_2$	$[\mathbf{x}_1, \mathbf{x}_2] = [G_{dis}(\mathbf{z}_1), G_{dis}(\mathbf{z}_2)]$	$[\mathbf{x}_1, \mathbf{x}_2] = [G_{ent}(\mathbf{z}_1), G_{ent}(\mathbf{z}_2)]$
$\Delta \mathbf{z} = (1, 0, 0)$		
$\Delta \mathbf{z} = (0, 1, 0)$		
$\Delta \mathbf{z} = (0, 0, 1)$		

Fig. 1. This is a table showing generated image pairs $[\mathbf{x}_1, \mathbf{x}_2]$ whose latent codes $[\mathbf{z}_1, \mathbf{z}_2]$ have difference in a single dimension $\Delta \mathbf{z} = \mathbf{z}_1 - \mathbf{z}_2$. Left: image pairs generated by a disentangled generator G_{dis} . Right: image pairs generated by an entangled generator G_{ent} . In each row, the latent code difference $\Delta \mathbf{z}$ is kept fixed with only one dimension modified. For the disentangled image pairs (left), it is not difficult to tell that each row represents the semantics of *fringe*, *smile*, and *hair color* respectively. However, for entangled ones (right) the semantics are not clear although the image pairs are also generated with a single dimension varied in the latent codes just like the left ones.

effective proxy for the learning of disentanglement, this assumption can be unrealistic for real-world data as the underlying distribution of the semantic factors may not be factorizable. Additionally since these models are defined based upon the Variational Autoencoder (VAE) framework [25] which intrinsically causes blurriness in the generated data, their applications are mostly limited to synthetic data and real-world data of small sizes. Another line of work to achieve disentanglement purpose is by using InfoGAN [10], which encourages disentanglement through maximizing the mutual information between the generated data and a subset of latent codes. This model inherits merits from GAN [15] so that sharper and more realistic images can be synthesized thus can be applied to more complex datasets. However, we show in the experiments that the disentanglement performance of InfoGAN is limited and the training of InfoGAN is less stable compared with our proposed GAN-based models. A problem in existing disentanglement learning community is the lack of evaluation metrics that can give a quantitative measurement of the performance of disentanglement. Existing metrics depends on the existence of ground-truth generative factors and an encoder network [18, 24, 9, 40, 14, 43], thus the quantitative measurements are usually done on synthetic data with VAE framework, leaving latent traversal inspection by human the only evaluation method for experiments on real-world data, and this to an extent discourages the development of GAN-based models, which are known to be effective in photorealistic image synthesis, from disentangled representation learning.

Different from the existing disentanglement learning methods, we reconsider the problem from the perspective of *Variation*. We argue that disentanglement can be naturally described by the correlated variations between the latent codes and the observations, and that is why researchers initially use latent traversals

as a method to evaluate whether a representation is disentangled. This intuition is based on an assumption that determining the semantics of a dimension in disentangled representations is easy but in entangled representations is difficult. In Fig. 1 we show an example of this interesting phenomenon. On the left and right parts of Fig. 1, there are image pairs generated by two models (one disentangled and one entangled). All these image pairs are generated by varying a single dimension in the latent codes, and the varied dimension is the same for each row (the varied dimension is indicated by $\Delta\mathbf{z}$ on the left by onehot representation). However, as there can be multiple different latent code pairs $[\mathbf{z}_1, \mathbf{z}_2]$ that cause the same $\Delta\mathbf{z}$, the generated image pairs can be diverse. In Fig. 1, it is not difficult to tell from the image pairs that the rows for the left model control the semantics of *fringe*, *smile*, and *hair color* respectively, while for the right model the semantics are not clear. This *easy-to-tell* property for the left model is due to the consistent pattern in the shown image pairs, while there is not a clear pattern for the right model so we cannot tell what each row is controlling.

This phenomenon for distinguishing disentangled and entangled representations motivates us to model the latent *variation predictability* as a proxy to achieve disentanglement, which is defined based on the difficulty of predicting the varied dimension in latent codes from the corresponding image pairs. Taking the case in Fig. 1 as an example, we say the left model has a higher latent variation predictability than the right one, which corresponds with the fact that we only need very small number of image pairs to tell what semantics each row is controlling for the left model. Note that rows in the entangled model (right one) may also contain stable patterns (may not correspond to a specific semantics), but it is difficult for us to tell what they are from such few images. By exploiting the variation predictability, our contributions in this paper can be summarized as follows:

- By regarding the variation predictability as a proxy for achieving disentanglement, we design a new objective which is general and effective for learning disentangled representations.
- Based on the definition of variation predictability, we propose a new disentanglement evaluation metric for quantitatively measuring the performance of disentanglement models, without the requirement of the ground-truth generative factors.
- Experiments on various datasets are conducted to show the effectiveness of our proposed metric and models.

2 Related Work

Generative Adversarial Networks. Since the introduction of the initial GAN by Goodfellow et. al. [15], the performance of GANs have been thoroughly improved from various perspectives, e.g. generative quality [39, 21, 47, 7, 22, 23], and training stability [1, 6, 16, 26, 35, 23]. However, the study of semantics learned in the latent space is less exploited. Chen et al. [10] propose InfoGAN which successfully learns disentangled representations by maximizing the mutual information

between a subset of latent variables and the generated samples. Donahue et al. [12] introduce Semantically Decomposed GANs which encourage a specified portion of the latent space to correspond to a known source of variation, resulting in the decomposition of identity and other contingent aspects of observations. In [44], Tran et. al. introduce DR-GAN containing an encoder-decoder structure that can disentangle the pose information in face synthesis. Karras et al. [22] introduce an intermediate latent space derived by non-linear mapping from the original latent space and conduct disentanglement study on this space with perceptual path length and linear separability. In [37], HoloGANs are introduced by combining inductive bias about the 3D world with GANs for learning a disentangled representation of pose, shape, and appearance. Recently, Lin et al. introduce a contrastive regularization to boost InfoGAN for disentanglement learning, and also propose a ModelCentrality scheme to achieve unsupervised model selection [28]. By looking into well-trained GANs, Bau et al. conducts GAN dissection in [4]. They show that neurals in GANs actually learn interpretable concepts, and can be used for modifying contents in the generated images. In a similar spirit, Shen et al. [41] introduces InterFaceGAN showing GANs spontaneously learn various latent subspaces corresponding to specific interpretable attributes. Besides the existing works, unsupervisedly learning disentangled representations with GANs and quantifying their performance of disentanglement are still unsolved problems.

Variational Autoencoders and Variants. There are more systematic works of disentanglement learning based on the VAE framework [25, 8], and the most common method for approaching disentanglement is by modeling the independence in the latent space. As an early attempt of extending VAEs for learning independent latent factors, Higgins et al. [18] pointed out that modulating the KL-term in the learning objective of VAEs (known as evidence lower bound (ELBO)) with a single hyper-parameter $\beta > 1$ can encourage the model to learn independent latent factors:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{z}, \beta) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \beta D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \quad (1)$$

Later based on the decomposition of the KL term $D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))$ [19, 32], it has been discovered that the KL divergence between the aggregated posterior $q_\phi(\mathbf{z}) \equiv \mathbb{E}_{p_{data}(\mathbf{x})}[q_\phi(\mathbf{z}|\mathbf{x})]$ and its factorial distribution $KL(q_\phi(\mathbf{z})||\prod_j q_\phi(z_j))$ (known as the total correlation (TC) of the latent variables) contributes most to the disentanglement purpose, leading to the emergence of models enforcing penalty on this term. Kim et al. [24] introduce FactorVAE to minimize this TC term through adopting a discriminator in the latent space [15, 33] with density-ratio trick [36, 3]. Chen et al. [9] introduce β -TCVAE to employ a mini-batch weighted sampling for the TC estimation. Kumar et al. [27] use moment matching to penalize the divergence between aggregated posterior and the prior. These works have been shown effective for disentangled representation learning, especially after the introduction of various quantitative disentanglement metric [18, 24, 9, 40, 14, 43]. Dupont [13] introduces JointVAE for learning continuous and discrete latent factors under the VAE framework for stable training and a principled inference network. Later Jeong et al. [20] introduce CascadeVAE which

handles independence enforcing through information cascading and solves discrete latent codes learning by an alternating minimization scheme. However, these works model disentanglement only from the independence perspective, which may not be practical for real-world data.

3 Methods

In this section, we first introduce the Variation Predictability objective in section 3.1, a general constraint which encourages disentanglement from the perspective of predictability for the latent variations, and also show its intergration with the GAN framework. We then introduce the proposed Variation Predictability Evaluation Metric in section 3.2, a general evaluation method quantifying the performance of disentanglement without relying on the ground-truth generative factors. Our code is available at: <https://github.com/zhuxinqimac/stylegan2vp>.

3.1 Variation Predictability Loss

We first introduce the concept of *variation predictability*, defined as follows:

Definition 1. *If image pairs are generated by varying a single dimension in the latent codes, the variation predictability represents how easy the prediction of the varied latent dimension from the image pairs is.*

Here we also need to define what *easy* is in this context: a prediction is called easy if the required number of training examples is small.

A high variation predictability means predicting the varied latent dimension from the image pairs is easy, i.e. only a small number of training image pairs are needed to identify this dimension (consider the left part of Fig. 1), while a low predictability means the prediction is hard (consider the right part of Fig. 1). As this concept corresponds well with our determination of disentanglement (the left model in Fig. 1 is disentangled while the right one is not), we propose to utilize it as a proxy for achieving the disentanglement objective. Note that traditionally the assumption of *independence* in latent variables works as another proxy for the learning of disentangled representation [18, 24, 13, 27, 9, 20], which we think is a stronger assumption than the proposed variation predictability and it may not hold for real-world data.

In order to model the variation predictability from observations to the varied latent dimension, we adopt a straightforward implementation by directly maximizing the mutual information between the varied latent dimension d and the paired images $(\mathbf{x}_1, \mathbf{x}_2)$ derived by varying dimension d in latent codes: $I(\mathbf{x}_1, \mathbf{x}_2; d)$, where we name this mutual information as the Variation Predictability objective (VP objective) and the negative of this term as the Variation Predictability loss (VP loss). We instantiate our VP objective by intergrating it within the generative adversarial network (GAN) framework.

As for a brief introduction, GAN is a generative model introduced by Goodfellow et al. [15] for learning the distribution of training data. A generator network

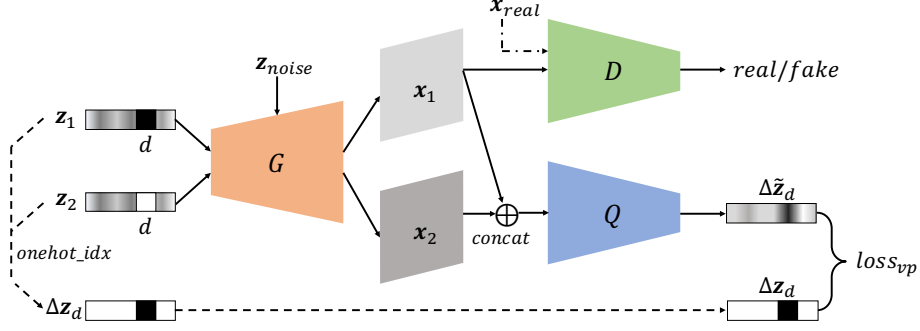


Fig. 2. Overall architecture of our proposed model. The model first samples two latent codes z_1 and z_2 from the latent space that only differs in a single dimension d , and this index d also serves as the target for the recognition network Q after being transformed into onehot representation Δz_d . Two images are then generated with the generator $x_1 = G(z_1)$, $x_2 = G(z_2)$. x_1 is fed to the discriminator to get the *real/fake* learning signal same as in the standard GANs. We then concatenate both images along the channel axis and feed them to the recognition network Q for the prediction of d .

G is trained to map a random noise z to an output image x : $G(z) : z \rightarrow x$. At the meantime, a discriminator network D is trained to predict whether the generated image x matches the real data distribution $p_{data}(x)$. The training strategy of G and D follows a minimax game where G tries to minimize while D tries to maximize the following objective:

$$\min_G \max_D V(G, D) = \mathbb{E}_{x \sim p_{data}} [\log D(x)] + \mathbb{E}_{z \sim p_{noise}} [\log (1 - D(G(z)))]. \quad (2)$$

After convergence, the generator G can generate images look similar to the real data, while the discriminator D cannot tell whether a generated image is real or not.

The VP loss can be easily integrated into a GAN as:

$$\min_G \max_D V_{vp}(G, D) = V(G, D) - \alpha I(x_1, x_2; d), \quad (3)$$

where α is a hyper-parameter to balance the generative quality and disentanglement quality. This model is named as Variation Predictability GAN (VPGAN). By optimizing the VP objective, the generator is encouraged to synthesize images having a strong correspondance with the latent codes so that the changes in data controlled by each latent dimension is distinguishable from each other, and the changes controlled by a single dimension is consistent all the time. In other words, after the training converges the variations in data should be naturally grouped into as many classes as the number of dimensions in the latent space. In Appendix 6 we show the relation between our objective and the InfoGAN [10] objective.

For the VP objective to be optimizable, we adopt the variational lower bound of mutual information defined as follows.

Lemma 1. *For the mutual information between two random variables $I(\mathbf{x}; \mathbf{y})$, the following lower bound holds:*

$$I(\mathbf{x}; \mathbf{y}) \geq H(\mathbf{y}) + \mathbb{E}_{p(\mathbf{x}, \mathbf{y})} \log q(\mathbf{y}|\mathbf{x}), \quad (4)$$

where the bound is tight when $q(\mathbf{y}|\mathbf{x}) = p(\mathbf{y}|\mathbf{x})$.

Proof. See Appendix 7. □

Based on Lemma 1, we can get the lower bound of our VP objective:

$$\mathcal{L}_{vp}(\theta, \phi; \mathbf{x}_1, \mathbf{x}_2, d) \quad (5)$$

$$= \mathbb{E}_{\mathbf{x}_1, \mathbf{x}_2, d \sim p_\theta(\mathbf{x}_1, \mathbf{x}_2, d)} \log q_\phi(d|\mathbf{x}_1, \mathbf{x}_2) + H(d) \quad (6)$$

$$\leq I(\mathbf{x}_1, \mathbf{x}_2; d). \quad (7)$$

For the sampling of d , \mathbf{x}_1 , \mathbf{x}_2 , we first sample the dimension index d out of the number of continuous latent variables (in this paper we focus on continuous latent codes and leave the modeling for discrete latent codes for future works), then we sample a latent code and sample twice on dimension d so we get a paired latent codes $[\mathbf{z}_1, \mathbf{z}_2]$ differing only on dimension d . The images \mathbf{x}_1 and \mathbf{x}_2 are generated by $G(\mathbf{z}_1)$ and $G(\mathbf{z}_2)$. The conditional distribution $q_\phi(d|\mathbf{x}_1, \mathbf{x}_2)$ in Eq. 6 is modeled as a recognizer Q . This recognizer network takes the concatenation of the two generated images as inputs to predict the varied dimension d in the latent codes. The architecture of our model is shown in Fig. 2.

3.2 Variation Predictability Disentanglement Metric

As the variation predictability naturally defines a way to distinguish disentangled representations and entangled representations by the difficulty of predicting the varied latent dimensions, we can propose a method to quantitatively measure the performance of disentanglement once we quantify the difficulty of the prediction. In this paper we quantify this difficulty by the performance of doing few-shot learning [45]. The intuition is that a prediction can be seen as easy if only a small number of training examples are needed for the prediction. From another viewpoint, only requiring a small number of training examples means the representation can generalize well for the prediction task, which is also a property of *disentanglement*, so this modeling is also consistent with disentanglement itself. We name our proposed metric as Variation Predictability metric (VP metric), which is defined as follows:

1. For a generative model, sample N indices denoting which dimension to modify in the latent codes: $\{d^1, d^2, \dots, d^N\}$.
2. Sample N pairs of latent codes that each pair only differs in the dimension sampled by step 1: $\{[\mathbf{z}_1^1, \mathbf{z}_2^1], [\mathbf{z}_1^2, \mathbf{z}_2^2], \dots, [\mathbf{z}_1^N, \mathbf{z}_2^N] \mid \text{Dim}_{\neq 0}(\mathbf{z}_1^i - \mathbf{z}_2^i) = d^i\}$.

3. For each latent code pair $[z_1^i, z_2^i]$, generate the corresponding image pair $[x_1^i = G(z_1^i), x_2^i = G(z_2^i)]$ and their difference $\Delta x^i = x_1^i - x_2^i$. This forms a dataset $\{(\Delta x^1, d^1), (\Delta x^2, d^2), \dots, (\Delta x^N, d^N)\}$ with the difference of image pairs as inputs and the varied dimension as labels.
4. Randomly divide the dataset into a training set and a test set with example numbers ηN and $(1 - \eta)N$ respectively, where η is the ratio of training set.
5. Train a recognition network taking Δx^i as input for predicting d^i on the training set. Report the accuracy acc_s on test set.
6. Repeat step 1 to step 5 S times to get accuracies $\{acc_1, acc_2, \dots, acc_S\}$.
7. Take the average of the accuracies $Score_{dis} = \frac{1}{S} \sum_{s=1}^S acc_s$ as the disentanglement score for the model. Higher is better.

The training set ratio η should not be large since we need to keep this setting as a few-shot learning so that the prediction is sufficiently difficult to distinguish disentangled representations and entangled representations. The reason we use the data difference Δx as inputs is that this enforces the model to focus on difference features in data caused by the varied dimension rather than contents in the images, and we also found this implementation can achieve a better correlation with other metrics like the FactorVAE score. In our experiments, we choose $N = 10,000$ and $S = 3$. For Dsprites and 3DShapes datasets we choose $\eta = 0.01$ and for CelebA and 3DChairs datasets we choose $\eta = 0.1$ as there are more dimensions in latent codes used for CelebA and 3DChairs. The main differences between our proposed metric and other disentanglement metrics [18, 24, 9, 40, 14, 43] are that ours does not require the existence of ground-truth generative factors (which is not available in real-world data), and our metric does not require an encoder so GAN-based models can also be evaluated.

3.3 Implementations

As discovered by the recent works of StyleGANs [22, 23] that the latent codes of a generator can be treated as style codes modifying a learned constant for achieving a higher-quality generation and stabler training than traditional GANs, we adopt a similar strategy to ease the training procedure of GANs in our experiments. However, unlike StyleGANs which take a side mapping network to transform the input codes into multiple intermediate latent codes, we directly feed the latent codes sampled from the prior distribution into the generator network without a mapping network to learn the disentangled representation in the prior latent space. The network architectures and parameters for different experiments are shown in Appendix 10.

4 Experiments

We first conduct experiments on popular disentanglement evaluation datasets Dsprites [34] and 3DShapes [24] to validate the effectiveness of our proposed VP disentanglement metric. Second we evaluate our method with VAE framework to

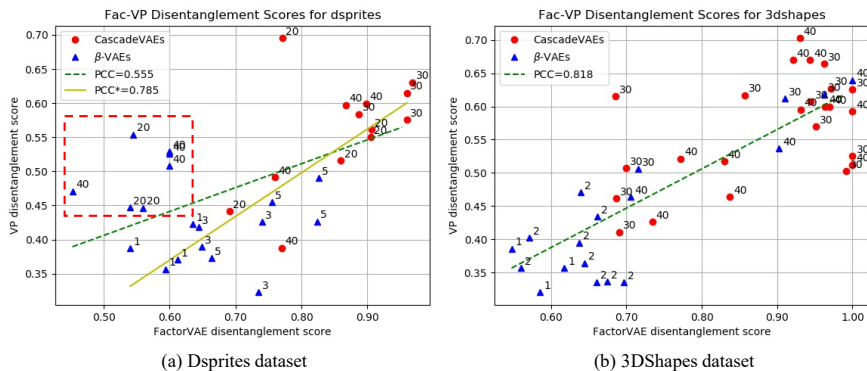


Fig. 3. These are scatter plots shown between the FactorVAE disentanglement metric and our proposed VP disentanglement metric on Dsprites dataset and 3DShapes dataset. The PCC denotes the Pearson’s Correlation Coefficient, and the green dash line is the linear regression line. See the main text for discussion.

show its complementarity to independence modeling for achieving disentanglement. Then we equip our models with GAN framework to conduct experiments on datasets without ground-truth generative factors 3DChairs [2] and CelebA [29], and use our proposed VP metric to quantitatively evaluate the disentanglement performance of our models.

4.1 VP Disentanglement Evaluation Metric

In this section, we evaluate the effectiveness of our proposed VP disentanglement metric. Specifically, we reimplement a relatively basic disentanglement model β -VAE [8] and a more advanced model CascadeVAE [20], and train them with different hyper-parameters and random seeds to cover a large range of model performance. Then we obtain their FactorVAE disentanglement metric scores, a widely-used disentanglement measurement shown to be correlated well with other disentanglement metrics and with qualitative evaluation [31, 24]. We then obtain the VP metric scores of the trained models and see if these scores have correlation with the ones calculated based on the FactorVAE metric. Note that the FactorVAE metric requires the existence of ground-truth generative factors of each data point in the training dataset for the performance measurement, while our VP metric does not use the ground-truth factors. This experiment is repeated on Dsprites dataset and 3DShapes dataset, which are two most popular datasets for the learning of disentangled representations. The correlation results for both datasets are shown in Fig. 3 (a) and (b) respectively. The β hyper-parameter in β -VAE is sampled from $\{1, 2, 3, 5, 20, 30, 40\}$, and the hyper-parameters β_{low} and β_{high} in CascadeVAE are sampled from $\{2, 3, 4\}$ and $\{20, 30, 40\}$ respectively. The models are run with multiple random seeds.

Model	VP Score	FacVAE Score
CasVAE	59.2 (4.6)	91.3 (7.4)
CasVAE-VP	65.5 (5.1)	91.7 (6.9)

Table 1. Disentanglement scores on Dsprites dataset.

Model	VP Score	FacVAE Score
CasVAE	62.3 (4.9)	94.7 (2.1)
CasVAE-VP	66.4 (5.6)	95.6 (2.4)

Table 2. Disentanglement scores on 3DShapes dataset.

From Fig. 3 (a) and (b), we can see there is an evident correlation between our VP metric and the FactorVAE metric. Note that the training of these models are not guided by the VP objective, and there is no relation between the design of FactorVAE metric and our proposed metric. Considering our metric requires no ground-truth factors and the performances of these models suffer from the impact of randomness during training, this correlation is already very strong. For Dsprites dataset, the Pearson’s Correlation Coefficient 0.555 is not as high as the one calculated on 3DShapes dataset $PCC = 0.818$. If we take a closer look, we can see the abnormal events happen when β is high in β -VAE models. This is because when β is high (the ones in the red box in Fig. 3 (a)), the model tends to ignore the *shape* and *rotation* information in the Dsprites dataset while keeping *xy-shift* and *scaling* disentangled, which leads to a low FactorVAE metric score. However, this ignorance of information is not known by our VP metric as it takes nothing from the ground-truth factors. This causes our metric only take into account the other remaining factors which are well-disentangled so high VP scores are obtained. If we omit the high- β models (in the red box), we get a much higher $PCC = 0.785$ that is close to the one obtained on 3DShapes dataset.

In summary, our proposed VP disentanglement metric correlates well with FactorVAE disentanglement metric even though ours does not rely on the ground-truth generative factors. This property makes our metric a general evaluation method for all disentanglement models with a generator, and it is applicable to datasets with or without ground-truth generative factors. In Appendix Section 8, we conducts experiments to show why small η is preferable in our proposed VP metric.

4.2 VP Models with VAE Framework

In this section we apply our VP models to popular disentanglement datasets Dsprites and 3DShapes with VAE framework. We equip our VP loss to the state-of-the-art disentanglement model CascadeVAE [20] to see if our variation predictability is complementary to the statistical independence modeling and can boost disentangled representation learning. The averaged disentanglement scores of 10 random seeds on two datasets are shown in Table 1 and Table 2. From the tables we can see our model can boost the disentanglement performance of CascadeVAE, but the performance improvement is not very significant. We believe this is because on datasets like Dsprites and 3DShapes the independence

Model	VP Score	FID
VPGAN-flat $\alpha = 0.001$	58.4	22.3
VPGAN-flat $\alpha = 0.01$	62.1	27.8
VPGAN-flat $\alpha = 0.1$	64.5	32.8
VPGAN-hierar $\alpha = 0.001$	64.6	53.6
VPGAN-hierar $\alpha = 0.01$	66.8	47.4
VPGAN-hierar $\alpha = 0.1$	70.3	56.9

Table 3. Ablation studies of implementation and hyper-parameter α on CelebA.

Model	VP Score	FID
GAN	12.9	20.4
InfoGAN $\lambda = 0.001$	34.5	24.1
InfoGAN $\lambda = 0.01$	25.3	43.3
FactorVAE	75.0	73.9
VPGAN-flat $\alpha = 0.1$	64.5	32.8
VPGAN-hierar $\alpha = 0.1$	70.3	56.9

Table 4. Disentanglement models comparison on CelebA dataset.

assumption is crucial for the disentanglement learning, which has the most impact on the learning of disentangled representations. However, this experiment still shows our VP loss is complementary to the statistical independence modeling, and is beneficial for disentanglement. In Appendix Section 9, we show more quantitative comparisons on these two datasets.

4.3 VP Models and Metric on CelebA with GAN Framework

In this section, we apply our VP models and metric to the challenging CelebA [29] human face dataset. This dataset consists of over 200,000 images of cropped real-world human faces of various poses, backgrounds and facial expressions. We crop the center 128×128 area of the images as input for all models.

We first conduct ablation studies on two factors that influence the performance of disentanglement in our models: 1) hierarchical inputs and 2) the hyper-parameter α . The hierarchical latent input impact is a phenomenon discovered in [22, 23, 48] that when feeding the input latent codes into different layers of the generator, the learned representations tend to capture different levels of semantics. We compare this hierarchical implementation with models with a traditional flat-input implementation to see its impact on disentanglement. The number of network layers and the latent codes are kept the same. In our experiments, we use the VP metric and FID [17] to give a quantitative evaluation on how these two factors impact the disentanglement and the image quality. The results are shown in Table 3. The latent traversals of VPGAN-flat $\alpha = 0.1$, InfoGAN $\lambda = 0.001$, VPGAN-hierarchical $\alpha = 0.1$, and FactorVAE $\gamma = 6.4$ are shown in Fig. 4, and Fig. 5, and more latent traversals can be found in Appendix 11, 12, and 13.

From Table 3, we can see the VP disentanglement score has a positive correlation with the hyper-parameter α . We can also summarize that the hierarchical inputs can boost the disentanglement score by an evident margin, indicating the hierarchical nature in deep neural networks can be a good ingredient for the learning of disentangled representations. On the other hand, the hyper-parameter α has a slight negative impact on FID, therefore it is better to choose a relatively small α to keep a balanced tradeoff between disentanglement and image quality.

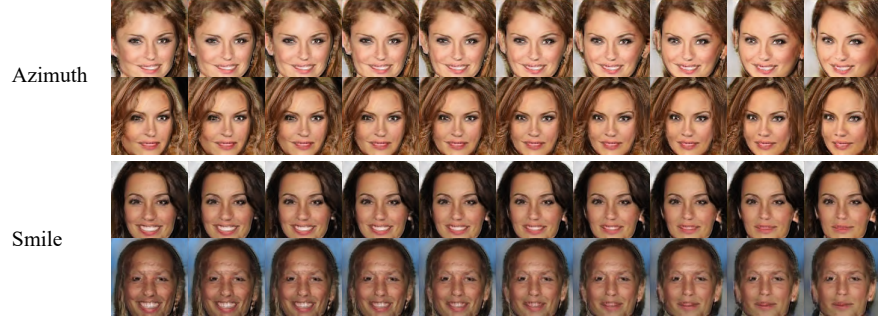
(a) VPGAN-flat**(b) InfoGAN**

Fig. 4. Latent traversals of VPGAN-flat and InfoGAN models on CelebA dataset. More latent traversals can be found in Appendix 11 and 12.

Nevertheless, the hierarchical input implementation seems to have a more significant negative impact on the FID, which we believe this technique should be better used with larger number of latent codes and more advanced architectures as in [22, 23] to take full advantage of it.

From Table 4, we can see our VPGANs can achieve highest disentanglement scores among GAN-based models and can even achieve close performance as FactorVAE which models independence in the latent space. However, the FactorVAE has a bad FID score, meaning the generated images are lack of fidelity significantly. On the contrary, our VPGANs can keep a better FID especially the flat-input version. We can see InfoGANs can achieve a certain level of disentanglement, but their performance is significantly lower than VPGANs. In practice, we also found the training of InfoGANs are less stable than our VPGANs where InfoGANs may result in generating all-black images, even though both types of models are using the same generative and discriminative networks. As a summary, our VPGANs keep a more balanced tradeoff between the disentanglement and generation quality than the compared models.

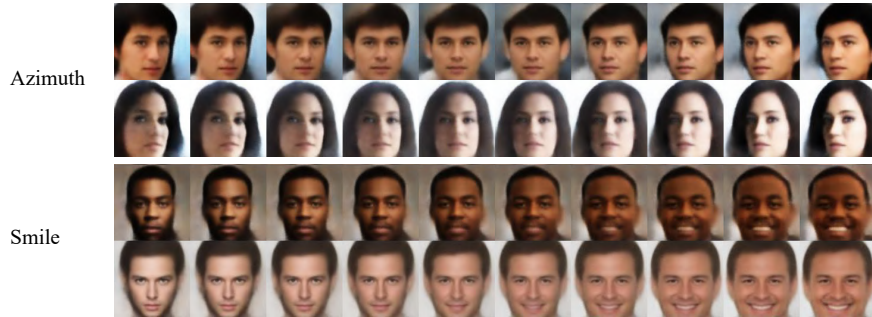
(a) VPGAN-hierarchical**(b) FactorVAE**

Fig. 5. Latent traversals of VPGAN-hierarchical and FactorVAE models on CelebA dataset. More latent traversals can be found in Appendix 13.

In Fig. 5 and Fig. 4 we qualitatively show the performance of our VPGANs, InfoGAN and FactorVAE baselines in disentanglement by latent traversals. From Fig. 4, we can see our model learns a cleaner semantics of azimuth while InfoGAN entangles azimuth with smile. Our model also learns a better latent code for controlling smile, while InfoGAN entangles smile and elevation into a single unit. In Appendix 11, 12 and 13, we show our VPGANs can learn more semantics (azimuth, brightness, hair color, makeup, smile, fringe, saturation, elevation, gender, lighting) than InfoGANs (azimuth, brightness, hair color, fringe, saturation, smile, gender, lighting). From Fig. 5, we can see the FactorVAE entangles smile with some level of skin texture information, while our model achieves a cleaner disentanglement. Also the results from FactorVAE are highly blurred, resulting in low FID.

There is an interesting phenomenon that for the learned disentangled representations in our models, not all dimensions encode variations. We find there are around $1/3$ of dimensions capturing no information (or too subtle to observe by eyes). Disentanglement prefers this property because it means the model does not randomly encode entangled information into the rest of the dimensions but

Model	VP Score	FID
InfoGAN $\lambda = 0.1$	36.7	30.4
VPGAN $\alpha = 100$	42.0	32.1

Table 5. Experiments on 3DChairs dataset comparing InfoGAN and VPGAN.

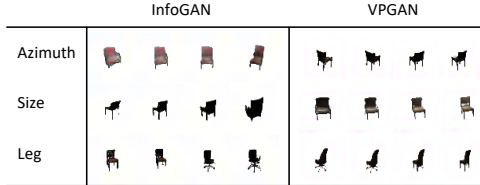


Fig. 6. Latent traversals on 3DChairs.

instead deactivates them. When the number of latent factors is set to 25 - 30, the learning is stable and almost all semantics shown can be learned. For latent factors less than 15, we observe some semantics are absent or entangled.

4.4 Experiments on 3D Chairs

We compare InfoGAN and VPGAN on 3D Chairs dataset. Quantitative results are shown in Table 5 and the latent traversals are shown in Fig. 6. As we can see, VPGAN achieves a higher disentanglement score than InfoGAN at the cost of a slight increase in FID, which agrees with what we found in the CelebA experiments. From the traversals, our VPGAN learns a cleaner latent code on controlling azimuth semantics while InfoGAN entangles it with some shape information. However, the performance of VPGAN on this dataset is not as impressive as on CelebA, indicating a more delicate modeling than the variation predictability assumption is required for this dataset to achieve a perfect disentanglement.

5 Conclusions

In this paper, we introduced the latent *Variation Predictability* as a new proxy for learning disentangled representations. By exploiting the latent variation predictability, we introduced the VP objective, which maximizes the mutual information between the varied dimension in the latent codes and the corresponding generated image pairs. Apart from the VP objective, we also proposed a new evaluation metric for quantitative measurement of disentanglement, which exploits the prediction difficulty of the varied dimension in the latent codes to quantify disentanglement. Different from other disentanglement metrics, our proposed VP metric does not require the existence of ground-truth generative factors. Experiments confirm the effectiveness of our model and metric, indicating the variation predictability can be exploited as a feasible alternative to statistical independence for modeling disentanglement in real-world data. For future works, we aim to extend our work to downstream applications like photorealistic image synthesis, domain adaptation, and image editing.

Acknowledgment This work was supported by Australian Research Council Projects FL-170100117, DP-180103424 and DE180101438. We thank Jiaxian Guo and Youjian Zhang for their constructive discussions.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: ICML (2017)
2. Aubry, M., Maturana, D., Efros, A.A., Russell, B.C., Sivic, J.: Seeing 3d chairs: Exemplar part-based 2d-3d alignment using a large dataset of cad models. 2014 IEEE Conference on Computer Vision and Pattern Recognition pp. 3762–3769 (2014)
3. Banerjee, A., Merugu, S., Dhillon, I.S., Ghosh, J.: Clustering with bregman divergences. *J. Mach. Learn. Res.* **6**, 1705–1749 (2004)
4. Bau, D., Zhu, J.Y., Strobel, H., Zhou, B., Tenenbaum, J.B., Freeman, W.T., Torralba, A.: Gan dissection: Visualizing and understanding generative adversarial networks. In: Proceedings of the International Conference on Learning Representations (ICLR) (2019)
5. Bengio, Y., Courville, A.C., Vincent, P.: Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**, 1798–1828 (2012)
6. Berthelot, D., Schumm, T., Metz, L.: Began: Boundary equilibrium generative adversarial networks. arXiv preprint arXiv:1703.10717 (2017)
7. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis (2019)
8. Burgess, C.P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., Lerchner, A.: Understanding disentangling in beta-vae. ArXiv **abs/1804.03599** (2018)
9. Chen, R.T.Q., Li, X., Grosse, R., Duvenaud, D.: Isolating sources of disentanglement in variational autoencoders. In: Advances in Neural Information Processing Systems (2018)
10. Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P.: Info-gan: Interpretable representation learning by information maximizing generative adversarial nets. In: NIPS (2016)
11. Creager, E., Madras, D., Jacobsen, J.H., Weis, M.A., Swersky, K., Pitassi, T., Zemel, R.S.: Flexibly fair representation learning by disentanglement. In: ICML (2019)
12. Donahue, C., Lipton, Z.C., Balsubramani, A., McAuley, J.J.: Semantically decomposing the latent spaces of generative adversarial networks (2018)
13. Dupont, E.: Learning disentangled joint continuous and discrete representations. In: NeurIPS (2018)
14. Eastwood, C., Williams, C.K.I.: A framework for the quantitative evaluation of disentangled representations. In: ICLR (2018)
15. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial networks. In: NIPS (2014)
16. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. In: NIPS (2017)
17. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NIPS (2017)
18. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M.M., Mohamed, S., Lerchner, A.: beta-vae: Learning basic visual concepts with a constrained variational framework. In: ICLR (2017)
19. Hoffman, M.D., Johnson, M.J.: Elbo surgery: yet another way to carve up the variational evidence lower bound. In: Workshop in Advances in Approximate Bayesian Inference, NIPS. vol. 1 (2016)

20. Jeong, Y., Song, H.O.: Learning discrete and continuous factors of data via alternating disentanglement. In: ICML (2019)
21. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation (2018)
22. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. *IEEE transactions on pattern analysis and machine intelligence* (2020)
23. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. *ArXiv abs/1912.04958* (2019)
24. Kim, H., Mnih, A.: Disentangling by factorising. In: ICML (2018)
25. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2013)
26. Kodali, N., Hays, J., Abernethy, J.D., Kira, Z.: On convergence and stability of gans (2018)
27. Kumar, A., Sattigeri, P., Balakrishnan, A.: Variational inference of disentangled latent concepts from unlabeled observations. In: ICLR (2018)
28. Lin, Z., Thekumparampil, K.K., Fanti, G., Oh, S.: Infogan-cr and modelcentrality: Self-supervised model training and selection for disentangling gans. In: ICML (2020)
29. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. 2015 IEEE International Conference on Computer Vision (ICCV) pp. 3730–3738 (2014)
30. Locatello, F., Abbati, G., Rainforth, T., Bauer, S., Schölkopf, B., Bachem, O.: On the fairness of disentangled representations. In: NeurIPS (2019)
31. Locatello, F., Bauer, S., Lucic, M., Rätsch, G., Gelly, S., Schölkopf, B., Bachem, O.: Challenging common assumptions in the unsupervised learning of disentangled representations. In: ICML (2019)
32. Makhzani, A., Frey, B.J.: Pixelgan autoencoders. In: NIPS (2017)
33. Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I.J.: Adversarial autoencoders. *ArXiv abs/1511.05644* (2015)
34. Matthey, L., Higgins, I., Hassabis, D., Lerchner, A.: dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/> (2017)
35. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks (2018)
36. Nguyen, X., Wainwright, M.J., Jordan, M.I.: Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory* **56**, 5847–5861 (2010)
37. Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.: Hologan: Unsupervised learning of 3d representations from natural images. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) pp. 7587–7596 (2019)
38. Peng, X., Huang, Z., Sun, X., Saenko, K.: Domain agnostic learning with disentangled representations. In: ICML (2019)
39. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: ICLR (2016)
40. Ridgeway, K., Mozer, M.C.: Learning deep disentangled embeddings with the f-statistic loss. In: NeurIPS (2018)
41. Shen, Y., Gu, J., Tang, X., Zhou, B.: Interpreting the latent space of gans for semantic face editing. *ArXiv abs/1907.10786* (2019)
42. van Steenkiste, S., Locatello, F., Schmidhuber, J., Bachem, O.: Are disentangled representations helpful for abstract visual reasoning? In: NIPS (2019)

43. Suter, R., Miladinovic, D., Schölkopf, B., Bauer, S.: Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. In: Proceedings of the 36th International Conference on Machine Learning (ICML). Proceedings of Machine Learning Research, vol. 97, pp. 6056–6065. PMLR (Jun 2019), <http://proceedings.mlr.press/v97/suter19a.html>
44. Tran, L., Yin, X., Liu, X.: Disentangled representation learning gan for pose-invariant face recognition. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 1283–1292 (2017)
45. Wang, Y., Yao, Q., Kwok, J.T., Ni, L.M.: Generalizing from a few examples: A survey on few-shot learning (2019)
46. Yang, J., Dvornek, N.C., Zhang, F., Chapiro, J., Lin, M., Duncan, J.S.: Unsupervised domain adaptation via disentangled representations: Application to cross-modality liver segmentation. In: MICCAI (2019)
47. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. In: ICML (2019)
48. Zhao, S., Song, J., Ermon, S.: Learning hierarchical features from generative models. In: ICML (2017)