

Toward a Controllable Disentanglement Network

Zengjie Song, Oluwasanmi Koyejo, and Jiangshe Zhang

Abstract—This paper addresses two crucial problems of learning disentangled image representations, namely controlling the degree of disentanglement during image editing, and balancing the disentanglement strength and the reconstruction quality. To encourage disentanglement, we devise a distance covariance based decorrelation regularization. Further, for the reconstruction step, our model leverages a soft target representation combined with the latent image code. By exploring the real-valued space of the soft target representation, we are able to synthesize novel images with the designated properties. To improve the perceptual quality of images generated by autoencoder (AE)-based models, we extend the encoder-decoder architecture with the generative adversarial network (GAN) by collapsing the AE decoder and the GAN generator into one. We also design a classification based protocol to quantitatively evaluate the disentanglement strength of our model. Experimental results showcase the benefits of the proposed model.

Index Terms—Representation learning, decorrelation regularization, image generation, autoencoder, generative adversarial network.

I. INTRODUCTION

ONE of the long-standing challenges in machine learning community is to learn interpretable and robust representations of sensory data. Disentangling the hidden factors of variation provides the possibility of overcoming such a challenge [1], [2]. In a disentangled (or factorial) image representation, the generative factors of images correspond to independent subsets of the latent dimensions, such that changing a single factor causes a change in a single latent unit while being invariant to others [3]. For example, a disentangled representation of face images could contain a set of latent units, each of which is sensitive to a specific facial attribute (i.e., generative factor) such as gender, age or wearing eyeglasses [3], [4]. According to Lake *et al.* [5], disentangled representations have the potential to boost the performance of state-of-the-art machine learning approaches in several situations, including transfer learning and zero-shot learning. Besides, it is claimed that such representations are more robust against adversarial attacks [6], [7], and are also beneficial to design more robust multi-stage reinforcement learning agents [8]. Other scenarios where disentangled

Manuscript first version received July 3, 2019; resubmitted January 22, 2020; revised May 23, 2020. First version log number is CYB-E-2019-07-1382. Part of this work was done while Z. Song was a visiting Ph.D. student with the University of Illinois at Urbana-Champaign. This work is supported in part by the National Natural Science Foundation of China under Grant 61572393, Grant 11671317, Grant 61877049, and Grant 61976174, and in part by the China Scholarship Council. (*Corresponding author: Jiangshe Zhang.*)

Z. Song and J. Zhang are with the School of Mathematics and Statistics, Xi'an Jiaotong University, Xi'an 710049, China (e-mail: zj-song@hotmail.com; jszhang@mail.xjtu.edu.cn).

O. Koyejo is with the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: sanmi@illinois.edu).

representations could play a role, such as novelty detection and information compression, can be found in the work of Ridgeway [2].

There is substantial literature on learning disentangled image representations with deep neural networks. Most of these models feature an explicit or implicit regularization to induce the non-correlation between hidden representations, and then implement image editing by tweaking the representation components of interest accordingly. Example models are FadNet [4] that resorts to an adversarial-like training strategy [9], or the ones based on the cross covariance regularization [10], [11]. Recently, several attempts have been made to investigate the disentanglement ability of the variational autoencoder (VAE) [12], [13]. To penalize the correlation between different dimensions in the representation, VAE-based models, such as β -VAE [14], DIP-VAE [15], JointVAE [3], and β -TCVAE [7], emphasize the importance of minimizing the KullbackLeibler divergence between the well-designed approximate posterior distribution and the disentangled prior distribution (e.g., an isotropic unit Gaussian) over the latent variable. However, all of the VAE's suffer from the nuisance blurring effects as observed in output images. By contrast, another set of works extending the generative adversarial network (GAN) [16] is able to generate realistic-looking images when handing disentanglement-related tasks, including IcGAN [17] for face image editing, DR-GAN [18] for pose-invariant face recognition, StarGAN [19] for image-to-image translation, and Soft-Gated Warping-GAN [20] for pose-guided person image synthesis. The key point underlying these models is that the GAN generator learns to map the input image (or the representation of input) to the target image as accurately as possible, provide that the label or the domain information is given. In that case, the adversarial training process implicitly makes the inferred representation uncorrelated to the label [17], [18], or makes the input image and the domain information interdependent for synthesizing target images [19], [20].

While the models mentioned above have shown promise in specific disentanglement tasks, they pay limited attention to simultaneously tackling the following two problems, which are very important but challenging as learning disentangled representations, namely controlling degree of disentanglement, and preserving image quality.

The first problem arises from controlling the degree of disentanglement during image editing, which means generating novel images with the designated attributes and, meanwhile, with the specific attribute intensities. For instance, given a face image, one may desire not only to synthesize a new smiling face, but also to synthesize a sequence of faces with expressions varying from no smile to toothy smile. Up to now, most of the existing disentanglement models [17], [19], [21]–[26], however, mainly focus on whether the model

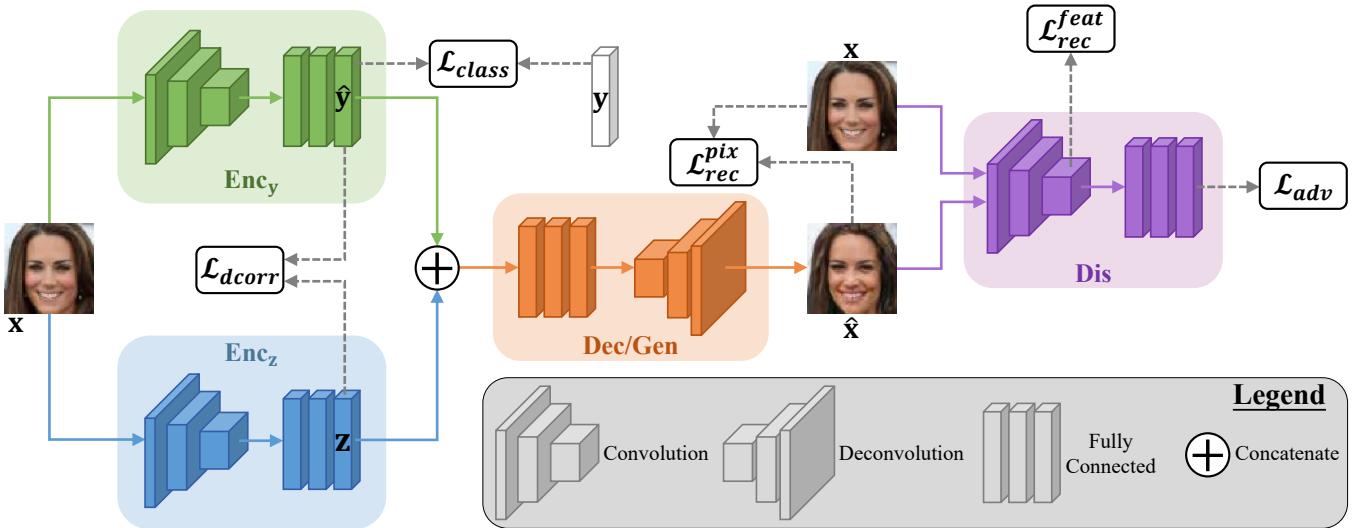


Fig. 1. Network architecture of our CDNet for training. Solid lines denote the forward pass, and dashed lines indicate the flow to compute local losses. See Section III for the detailed definitions of all symbols.

can generate images with or without attributes of interest, rather than controlling the attribute intensities involved in the synthesized images. In practice, a more subtle manipulation of images is often most useful. This property could allow several potential applications, such as automatic face image editing and image color rendering [4].

The second problem is how to generate images with the target attributes while preserving the core object identity and the image quality. Generally speaking, if the learned target representation part is not as independent as possible of other parts, changing one generative factor could induce changes of other factors, even falsifying the object identity. This phenomenon has been observed in many existing works [7], [10], [14], [15], [27]–[29]. For instance, as shown in [7], heightening the baldness intensity of a face concurrently leads to the visually-perceptible change of the object identity. Additionally, many disentanglement models based on autoencoder (AE) [3], [7], [10], [14], [15], [27] usually produce blurry output images, thus giving rise to the image quality degradation.

In this paper, we present a simple yet effective model named as Controllable Disentanglement Network, or CDNet, to address the aforementioned two problems. The overall network architecture of our model for training is depicted in Fig. 1. Specifically, CDNet combines the AE with the GAN to construct a new deep neural network, which is further divided into four parts: Enc_y , Enc_z , Dec/Gen , and Dis . The two encoders Enc_y and Enc_z are used to learn the soft target representation \hat{y} and the latent representation z , respectively. The representation \hat{y} acts to capture class- or attribute-related information by solving the supervised classification subtask, while the representation z serves to extract information different from those in \hat{y} via the proposed decorrelation regularization. The decoder of the AE (labeled Dec) works to reconstruct input image x when given the two groups of representations \hat{y} and z . By viewing the AE decoder as the GAN generator (labeled Gen), the reconstructed

image \hat{x} is also treated as the fake image, with the purpose of fooling the discriminator (labeled Dis) such that the Dis cannot distinguish the fake image from the real image. To achieve image editing, only the two well-trained encoders and the decoder are utilized, and modifications of \hat{y} values are performed with regard to the designated classes or attributes.

In summary, we highlight our contributions as follows.

- 1) To encourage disentanglement, a novel decorrelation regularization based on the distance covariance [30] is proposed to learn two independent representations, i.e., the soft target representation \hat{y} and the latent representation z (see Section III-A2 for details).
- 2) To implement the disentanglement controlling, we explore the potential of the soft target representation, rather than the discrete label as used in many existing works [10], [17], [21], [31], to reconstruct original and synthesize new images. This soft target representation is a probability representation at training time, and its element scale implicitly indicates how much class or attribute information is included in input image. Under this setting, one is able to decrease or increase the specific element scales to modify attribute intensities of the synthesized image at testing phase (see Section III-D for details).
- 3) To improve the perceptual quality of the generated images, we extend the AE architecture with GAN, where the GAN generator and the AE decoder are tied as the same one by parameter sharing and joint training. This model combination is inspired by the VAE/GAN [22], with the difference in that we build the CDNet based on the deterministic AE, and introduce a parallel encoder to learn the soft target representation. The new integrated model shows improved ability to reconstruct images and learn disentangled representations (see Sections IV-B, IV-C, and IV-D for empirical comparisons).
- 4) To quantitatively compare the disentanglement strength

of our model, an evaluation protocol is designed. To our best knowledge, this is the first work that leverages classification to analyze how the effect of representation scales with disentanglement performance (see Section IV-E for details).

The PyTorch source code of our model is available at <https://github.com/zjsong/CDNet>.

II. RELATED WORK

Due to its nature of interpretability and robustness, disentangled representations have been attracting increasing attention in recent years. Here, we divide the related disentanglement models into three prominent groups: AE-based models [4], [10], [14], [15], [27], [32], GAN-based models [19], [24], [25], [28], [29], and integrations of AE's and GAN's [17], [21]–[23], [26]. In the following, we provide a detailed survey of this topic from these three perspectives.

A. AE-based Models

By constraining the latent variables to be invariant to image attributes, the basic AE model can be extended to handle the disentanglement task. For instance, the FadNet [4] learns a classifier to predict the attribute given the latent representation, while the latent representation inferred by the encoder tries to prevent the classifier from predicting the correct attribute values. This adversarial-like process enables the model, in an implicit manner, to learn a latent representation containing information different from the attributes. By contrast, Cheung *et al.* [10] employ an explicit decorrelation regularization, based on the cross covariance (XCov), to approach the same goal. Our CDNet is also built on the basic AE, but with three notable differences from the aforementioned models. First, in contrast to the AE-based models that take advantage of the class or attribute label during image editing, the CDNet implements image editing by using the inferred soft target representation, thus our model is applicable to scenarios where label information is unavailable. Second, compared with the XCov regularization, the distance covariance (dCov) we use for disentanglement encourages statistical independence rather than non-correlation between variables [30], leading to a stronger disentanglement ability (see Sections IV-E and IV-F for details). Third, the usage of a GAN in our model also markedly improves the perceptual quality of the output images.

The vanilla VAE [12], [13] has been shown to learn disentangled representations, but with limited disentanglement ability on simple datasets such as FreyFaces or MNIST. Higgins *et al.* [14] and Kumar *et al.* [15] refine VAE to learn controllable disentangled factors, implemented by putting implicit independence constraints on the approximate posterior over latent variables. Kulkarni *et al.* [27] achieve disentanglement based on a special training scheme, where pairs of rendered images that differ only in one factor of variation are provided. Another VAE-like model [32] utilizes the vector arithmetic technique [33] to control attribute intensities. We note that those models can be trained stably in general, however, they are prone to obtain blurry images. Besides, the unsupervised learning strategy adopted in many of these models cannot guarantee

the non-correlation between learned representations, and thus the change of one attribute (e.g., smiling) may induce changes of other attributes (e.g., hairstyle or azimuth) as observed in [15]. In fact, Locatello *et al.* [34] have theoretically shown that the unsupervised disentanglement learning is fundamentally impossible without inductive biases both on models and datasets. This conclusion indicates that the role of supervision is crucial [34], which is coincident with the idea of using labeled training data in our model.

B. GAN-based Models

The plain GAN [16] does not show any apparent disentanglement properties, nevertheless, subsequent works have enhanced GANs. Donahue *et al.* [25] propose the semantically decomposed GANs that learn to decompose the latent code into an identity-related portion and observation-related portion, thus modifying face images by varying the observation vector. Focusing on person image generation, Ma *et al.* [29] use an adversarial network to learn mappings from Gaussian noise to the embedding feature space, which provides more control over the foreground, background, and pose information of the input image. In the InfoGAN [28], a subset of facial attributes is changed by manipulating the learned categorical codes, but with no conspicuous visual difference among generated images (such as the “Hair style” variation shown therein). By coupling two GANs together, the DiscoGAN [24] leverages the cross-domain relations to perform the facial attribute conversion task. The StarGAN [19], one of the state-of-the-art multi-domain image translation models, achieves facial attribute manipulation by a single generator and shows remarkable ability to synthesize high quality images. However, StarGAN's such superiority is obtained only when the number of operated attribute domains is small, and thus it is less effective for balancing the disentanglement strength and the reconstruction quality across a mass of different attributes (e.g., all 40 facial attributes in CelebA face images, as illustrated in the supplementary material). Moreover, it's worth noting that many GAN-based models still suffer from the problems of training instability and model collapse [35], which also make the disentangled representation learning more challenging.

C. Combined AE and GAN Models

A natural way to alleviate aforementioned problems is to combine AE with GAN, thereby leveraging both models' strengths in a complementary manner. To approach this goal, several existing works explore the adversarial training strategy in the latent space of AEs. The main point of these models is making the GAN discriminator indistinguishable 1) between the aggregated posterior of the latent variable and an arbitrary prior [21]; or 2) between samples in latent space and encoded data (rather than prior samples) [26]; or 3) between joint samples of the data and the corresponding latent variable from the encoder and joint samples from the decoder [23].

The IcGAN [17] and the VAE/GAN [22] are another two works falling into this line of literature. In IcGAN, an encoder is added into the conditional GAN [31] to learn a mapping from the image space to the representation space, and

thus implementing image editing by changing the conditional information inferred from the real image. In VAE/GAN, the VAE decoder and the GAN generator are viewed as the same mapping by parameter sharing and joint training, and the GAN discriminator acts to measure sample similarity in the feature space. Although the network architecture is similar, the representation learning method of our CDNet differs substantially from these models. First, the attribute representations learned by IcGAN and VAE/GAN are still correlated with each other, degrading the ability to manipulate images subtly. By contrast, the decorrelation regularization facilitates the CDNet to learn independent representations, which enables the model to control disentanglement at testing phase. Second, both of the pixel reconstruction error and the feature reconstruction error are explored in CDNet, and thus the stability of model training and the perceptual quality of output images are all improved. We conducted a series of experiments to compare these two models with the proposed CDNet in Sections IV-B, IV-C, and IV-D.

III. METHODOLOGY

We propose the CDNet, a novel model combining AE with GAN, that advances the state-of-the-art toward jointly solving the problems of the image disentanglement controlling and the image quality balance between disentanglement and reconstruction. The CDNet architecture is shown in Fig. 1, and it consists of four components: Enc_y , Enc_z , Dec/Gen , and Dis . Specifically, the encoder Enc_y aims to learn the *soft target representation* \hat{y} to extract class or attribute information from the discrete label y . This goal is approached by training Enc_y to solve a supervised classification task. Another encoder Enc_z serves to learn the *latent representation* z under two constraints: being informative to reconstruct input image x and being uncorrelated with (even independent of) \hat{y} . Here the independence between \hat{y} and z is induced by the proposed decorrelation regularization. The Dec takes as input the representations \hat{y} and z to reconstruct input image. Because CDNet collapses the AE decoder and the GAN generator into one, the reconstructed image \hat{x} is also treated as the fake image generated by the Gen . With this setting, we train the Dis to distinguish fake image \hat{x} from real image x , and also use the middle layer representations of the Dis to compute the feature reconstruction error.

In the following, we first formulate all local losses used to train different network components. Then we derive the integrated loss function and elaborate the associated training algorithm of our CDNet. After that, several practical considerations for implementation are provided. Finally, the method to manipulate images with controllable disentanglement is illustrated in two applications.

In addition to the aforementioned symbols, let \mathbf{X} denote the mini-batch version of the input image x , N the mini-batch size, and similarly define $\hat{\mathbf{X}}$, \mathbf{Y} , $\hat{\mathbf{Y}}$, and \mathbf{Z} for the reconstructed/fake image \hat{x} , label y , soft target representation \hat{y} , and latent representation z , respectively.

A. Four Local Losses

1) *Classification Loss* \mathcal{L}_{class} : Minimizing the classification loss \mathcal{L}_{class} enables the encoder Enc_y to inject the class or attribute information into the soft target representation \hat{y} . We select two classification loss functions to fit the following two application cases, respectively.

Case 1: For the multiclass scenario (e.g., a handwritten digit belongs to only one of the 10 classes), we first use the softmax nonlinearity to scale each element of \hat{y} . Then the cross entropy between the discrete labels \mathbf{Y} and the soft target representations $\hat{\mathbf{Y}}$ is computed as the classification loss.

Case 2: For the multilabel scenario (e.g., face images with or without smiling, eyeglasses, and blond hair attributes), we first use the sigmoid nonlinearity to scale each element of \hat{y} . Then the binary cross entropy between \mathbf{Y} and $\hat{\mathbf{Y}}$ is derived as the classification loss.

2) *Decorrelation Loss* \mathcal{L}_{dcorr} : We propose to leverage the distance covariance (dCov) [30] based regularization to learn the latent representation z , which is expected to be independent of the soft target representation \hat{y} . To obtain this decorrelation loss (or regularization), we first compute the N by N distance matrices $(a_{n,m})$ and $(b_{n,m})$ containing all pairwise distances:

$$a_{n,m} = \|\hat{\mathbf{y}}_n - \hat{\mathbf{y}}_m\|_2, \quad n, m = 1, 2, \dots, N,$$

$$b_{n,m} = \|\mathbf{z}_n - \mathbf{z}_m\|_2, \quad n, m = 1, 2, \dots, N$$

where $\|\cdot\|_2$ is the l^2 -norm. Then take all doubly centered distances:

$$A_{n,m} := a_{n,m} - \bar{a}_{n\cdot} - \bar{a}_{\cdot m} + \bar{a}_{..},$$

$$B_{n,m} := b_{n,m} - \bar{b}_{n\cdot} - \bar{b}_{\cdot m} + \bar{b}_{..}$$

where $\bar{a}_{n\cdot}$ is the n th row mean, $\bar{a}_{\cdot m}$ is the m th column mean, and $\bar{a}_{..}$ is the grand mean of the distance matrix $(a_{n,m})$. The notation is similar for the b values. Finally, the squared sample distance covariance, treated as our decorrelation loss, is simply the arithmetic average of the products $A_{n,m}B_{n,m}$:

$$\mathcal{L}_{dcorr} = \text{dCov}^2(\hat{\mathbf{Y}}, \mathbf{Z}) = \frac{1}{N^2} \sum_{n=1}^N \sum_{m=1}^N A_{n,m}B_{n,m}. \quad (1)$$

By minimizing the decorrelation loss in Eq. (1) to approach zero, the soft target representation \hat{y} and the latent representation z would tend to be independent. This conclusion is supported by the following Theorem 1. Let \mathbf{U} and \mathbf{V} denote two random vectors, $\text{dCov}_K^2(\mathbf{U}, \mathbf{V})$ is the squared sample distance covariance between \mathbf{U} and \mathbf{V} , and K indicates the number of pairwise sample points¹.

Theorem 1: Suppose two random vectors \mathbf{U} and \mathbf{V} satisfy $E(\|\mathbf{U}\|_2) < \infty$ and $E(\|\mathbf{V}\|_2) < \infty$. If $\lim_{K \rightarrow \infty} \text{dCov}_K^2(\mathbf{U}, \mathbf{V}) = 0$, then almost surely \mathbf{U} and \mathbf{V} are independent.

Proof: The proof can be established by using the Definition 3, Theorem 2, and Theorem 3 in [30]. See the Appendix for completeness. ■

¹Both dCov^2 and dCov_K^2 stand for the squared sample distance covariance, and the subscript K is used to facilitate theoretical analysis.

By comparison, Cheung *et al.* [10] use the following cross covariance (XCov) to facilitate disentanglement:

$$\text{XCov}(\hat{\mathbf{Y}}, \mathbf{Z}) = \frac{1}{2} \sum_{i,j} \left[\frac{1}{N} \sum_{n=1}^N (\hat{y}_{n,i} - \bar{\hat{y}}_{\cdot i})(z_{n,j} - \bar{z}_{\cdot j}) \right]^2 \quad (2)$$

where $\bar{\hat{y}}_{\cdot i} = \frac{1}{N} \sum_{n=1}^N \hat{y}_{n,i}$ and $\bar{z}_{\cdot j} = \frac{1}{N} \sum_{n=1}^N z_{n,j}$. It is worth emphasizing that one of the most important differences between dCov² and XCov is that, minimizing dCov² encourages the *independence* between two random variables, while minimizing XCov encourages the *non-correlation*. In this regard, the dCov² should induce stronger disentanglement than the XCov. Additionally, our model is also compatible with XCov, and replacing dCov² with XCov in CDNet also achieves improved disentanglement performance over the model in [10] (see Sections IV-C and IV-D for empirical comparisons).

3) *Reconstruction Loss \mathcal{L}_{rec}* : The reconstruction principle of CDNet is similar to the basic AE, that is, the soft target representations $\hat{\mathbf{Y}}$ and the latent representations \mathbf{Z} are first computed by the two encoders, respectively, and then fed to the decoder to reconstruct original images:

$$\hat{\mathbf{Y}} = \text{Enc}_y(\mathbf{X}), \quad \mathbf{Z} = \text{Enc}_z(\mathbf{X}), \quad \hat{\mathbf{X}} = \text{Dec}(\hat{\mathbf{Y}}, \mathbf{Z}). \quad (3)$$

We use the reconstruction loss to measure the difference between original images \mathbf{X} and reconstructions $\hat{\mathbf{X}}$. A common choice is the mean squared error (MSE) computed in pixel space:

$$\mathcal{L}_{rec}^{pix} = \frac{1}{NM} \sum_{n=1}^N \|\mathbf{x}_n - \hat{\mathbf{x}}_n\|_2^2 \quad (4)$$

where M indicates the image dimensionality.

As discussed in Section II, AE's trained with the pixel-level MSE usually produce blurring effects in the resulting image, which is also observed in Section IV-B in this paper. We conjecture that lacking the meaningful spatial correlation properties of original images causes the quality degradation of reconstructions. The hidden representation of a deep convolutional neural network, however, can extract such spatial correlation properties from the input image [22], [32]. Inspired by this observation, we compute the feature-matching difference [22] as the additional reconstruction loss:

$$\mathcal{L}_{rec}^{feat} = \frac{1}{ND_l} \sum_{n=1}^N \|h_l(\mathbf{x}_n) - h_l(\hat{\mathbf{x}}_n)\|_2^2 \quad (5)$$

where D_l is the dimensionality of the l th layer of the GAN discriminator and $h_l(\mathbf{x})$ denotes the hidden representation of \mathbf{x} at that layer. It's worth noting that the perceptual loss [36], [37], computed by some pretrained high-performing CNN such as VGG [38], can also measure the high-level feature difference between two images. But the perceptual loss is not widely applicable when the training data are not in image format (e.g., audio or text data), or when the training data are images but the size of image is smaller than the acceptable image size to the pretrained model. By comparison, the feature-matching difference loss (5) is customized and learned on the fly at each iteration step, and thus enabling our model to be flexible and extensible to handle data disentanglement tasks.

The final reconstruction loss consists of two parts:

$$\mathcal{L}_{rec} = \mathcal{L}_{rec}^{pix} + \lambda_{rec} \mathcal{L}_{rec}^{feat} \quad (6)$$

where λ_{rec} controls the trade-off between reconstructions of global features (i.e., \mathcal{L}_{rec}^{pix}) and local details (i.e., \mathcal{L}_{rec}^{feat}). By minimizing these two local reconstruction losses together, the CDNet is enforced to restore identity-preserving images with high-level structures. In practice, we have observed that using the pixel-wise reconstruction loss also makes the adversarial training more stable.

4) *Adversarial Loss \mathcal{L}_{adv}* : The goal of incorporating GAN into CDNet is to improve the perceptual quality of the output images. In the GAN part of CDNet, the generator Gen maps the inferred soft target representation and the latent representation to image space, while the discriminator Dis estimates the probability that a sample belongs to the data distribution. The GAN is trained such that the Dis can tell apart real from fake images, and meanwhile the Gen can generate images that "fool" the Dis. To this end, we need to maximize/minimize the following adversarial loss

$$\mathcal{L}_{adv} = \log(\text{Dis}(\mathbf{X})) + \log(1 - \text{Dis}(\text{Gen}(\hat{\mathbf{Y}}, \mathbf{Z}))) \quad (7)$$

with respect to the Dis/Gen.

Note that the exact choice of the GAN model (and so the adversarial loss) is not fundamental in our CDNet, since the plain GAN described here is adequate to improve the image perceptual quality. To obtain images with better visual fidelity, several advanced GAN models such as PatchGAN [39] and WGAN-GP [40] could be employed.

B. Integrated Loss and Training Algorithm

We train our combined model with the integrated loss

$$\mathcal{L} = \mathcal{L}_{class} + \lambda_{dcorr} \mathcal{L}_{dcorr} + \mathcal{L}_{rec} + \mathcal{L}_{adv} \quad (8)$$

where the exact expression of each local loss is presented above, and λ_{dcorr} balances the quality of the reconstruction and the strength of the disentanglement. All of the local losses included in Eq. (8) are complementary to each other, enabling the CDNet to address the two crucial disentanglement problems mentioned in Section I.

We train each model component of CDNet with the associated local losses. More specifically, we first train the encoder Enc_y to solve a supervised classification problem, which is implemented by minimizing the classification loss \mathcal{L}_{class} w.r.t. θ_{Enc_y} . After that, we fix the Enc_y and just use it to infer the soft target representation of input. For another encoder Enc_z , the decorrelation loss \mathcal{L}_{dcorr} and the reconstruction loss \mathcal{L}_{rec} are minimized to update its parameters θ_{Enc_z} . The parameters of the decoder (and so, the generator), θ_{Dec} , are modified based on the minimization of \mathcal{L}_{rec} and \mathcal{L}_{adv} . The adversarial loss \mathcal{L}_{adv} is also related to the discriminator, and is used to learn parameters θ_{Dis} . For clarity, we summarize the training algorithm for CDNet in Algorithm 1. More details about the parameter setting can be found in Section IV-A3.

Algorithm 1 Training the CDNet

```

1:  $\theta_{\text{Enc}_y}, \theta_{\text{Enc}_z}, \theta_{\text{Dec}}, \theta_{\text{Dis}} \leftarrow$  initialize network parameters
2: repeat ▷ Train  $\text{Enc}_y$  independently
3:    $\{\mathbf{X}, \mathbf{Y}\} \leftarrow$  random mini-batch from dataset
4:    $\hat{\mathbf{Y}} \leftarrow \text{Enc}_y(\mathbf{X})$ 
5:    $\mathcal{L}_{\text{class}} \leftarrow$  (Binary)CrossEntropy( $\hat{\mathbf{Y}}, \mathbf{Y}$ )
6:    $\theta_{\text{Enc}_y} \leftarrow -\nabla_{\theta_{\text{Enc}_y}} \mathcal{L}_{\text{class}}$ 
7: until deadline
8: repeat ▷ Train  $\text{Enc}_z$ ,  $\text{Dec}$ , and  $\text{Dis}$  simultaneously
9:    $\mathbf{X} \leftarrow$  random mini-batch from dataset
10:   $\hat{\mathbf{Y}} \leftarrow \text{Enc}_y(\mathbf{X})$ 
11:   $\mathbf{Z} \leftarrow \text{Enc}_z(\mathbf{X})$ 
12:   $\mathcal{L}_{\text{dcorr}} \leftarrow \text{dCov}^2(\hat{\mathbf{Y}}, \mathbf{Z})$  or  $\text{XCov}(\hat{\mathbf{Y}}, \mathbf{Z})$ 
13:   $\hat{\mathbf{X}} \leftarrow \text{Dec}(\hat{\mathbf{Y}}, \mathbf{Z})$ 
14:   $\mathcal{L}_{\text{rec}}^{\text{pix}} \leftarrow \text{MSE}(\hat{\mathbf{X}}, \mathbf{X})$ 
15:   $\mathcal{L}_{\text{rec}}^{\text{feat}} \leftarrow \text{MSE}(h_l(\hat{\mathbf{X}}), h_l(\mathbf{X}))$ 
16:   $\mathcal{L}_{\text{rec}} \leftarrow \mathcal{L}_{\text{rec}}^{\text{pix}} + \lambda_{\text{rec}} \mathcal{L}_{\text{rec}}^{\text{feat}}$ 
17:   $\mathcal{L}_{\text{adv}} \leftarrow \log(\text{Dis}(\mathbf{X})) + \log(1 - \text{Dis}(\hat{\mathbf{X}}))$ 
18:   $\theta_{\text{Enc}_z} \leftarrow -\nabla_{\theta_{\text{Enc}_z}} (\lambda_{\text{dcorr}} \mathcal{L}_{\text{dcorr}} + \mathcal{L}_{\text{rec}})$ 
19:   $\theta_{\text{Dec}} \leftarrow -\nabla_{\theta_{\text{Dec}}} (\mathcal{L}_{\text{rec}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}})$ 
20:   $\theta_{\text{Dis}} \leftarrow -\nabla_{\theta_{\text{Dis}}} \mathcal{L}_{\text{adv}}$ 
21: until deadline

```

C. Implementation Details

In addition to the aforementioned recipe for training the CDNet, we also adopt the following techniques demonstrated in [4] to stabilize the training process in practice.

- 1) *Appending the soft target representation to each layer of the decoder:* The soft target representation inferred by the encoder Enc_y contains class- or attribute-related information. This discriminative information is utilized by the decoder in two different ways. For each of the fully-connected layers of the decoder, we concatenate the soft target representation and the hidden layer representation together as a whole input to the next layer. For all the convolutions of the decoder, we append the soft target representation as additional constant input channels.
- 2) *Decorrelation loss scheduling:* To avoid that the decorrelation regularization dominates the parameter updating of the encoder Enc_z , which would destroy the model's reconstruction ability, we use a variable weight for the regularization parameter λ_{dcorr} . That is, we linearly increase the λ_{dcorr} to a target value over the early training process, and then clamp it for the remaining training process. By doing so, the effect of decorrelation regularization is gradually imposed on the learning of the latent representation.
- 3) *Dropout:* We use the dropout [41] in all fully-connected layers, except the final layer, of the two encoders and the discriminator. In our experiments, we found that dropout is beneficial to prevent the encoder Enc_y and the discriminator Dis from overfitting, and is also helpful for the encoder Enc_z to learn a latent representation being as independent as possible of the soft target representation.

D. Methods to Manipulate Images

For image editing, the key operation is to modify the value of the soft target representation $\hat{\mathbf{y}}$ accordingly. Based on the two application cases described in Section III-A1, we give two corresponding methods to manipulate images.

In *Case 1*, taking the handwritten digit as an example, we aim to generate a new digit with the handwriting style similar to the given digit. As shown in Fig. 2, we first employ the two encoders, Enc_y and Enc_z , to infer the soft target representation $\hat{\mathbf{y}}$ and the latent representation \mathbf{z} of the given digit “1” in boldface. Then we modify $\hat{\mathbf{y}}$ by exchanging the third element (corresponding to the digit 2 class) and the maximum element (ideally corresponding to the digit 1 class), while keeping remaining elements fixed. In this way, only two elements of $\hat{\mathbf{y}}$ at most are exchanged, and thus the representation structure with component summation of 1 is preserved. Finally, we feed the modified $\hat{\mathbf{y}}$ and the unchanged \mathbf{z} to the decoder to generate the new digit “2” which is also in boldface.

In *Case 2*, with the face image as an example, the goal is to synthesize a new face with the desired attribute and intensity while preserving the core identity. As we can see from Fig. 2, the overall procedure is similar to the first case, only with the difference in modifying $\hat{\mathbf{y}}$. Actually, in order to generate a new face with eyeglasses, we just replace the original (near) zero value corresponding to “Eyeglasses” attribute with the new value (e.g., 3.5) in $\hat{\mathbf{y}}$. Note that *during image editing* the modified attribute value is not necessarily restricted in $[0, 1]$, meaning it can also take other real values outside this interval². Specifically, the small value (near 0 or less than 0) indicates that the synthesized image tends to exclude the target attribute, while the big value (near 1 or greater than 1) implies that the synthesized image prefers containing that attribute. In this regard, the continuous attribute value can be viewed as a sliding knob, the magnitude of which controls how much a specific attribute could be perceivable in the final image. As illustrated in [28], by modifying attribute values in such an exaggerated way, the soft target representation is able to cover a wide range that the network was never trained on and we will get meaningful generalization (see Section IV-D for empirical evidence).

IV. EXPERIMENTS

We conduct five groups of experiments to substantiate the benefits of our CDNet model. First, we show that the combination of AE and GAN is able to improve the quality of reconstructed images. Second, we verify that the proposed decorrelation regularization and the image manipulation methods are competent to disentangle factors of variation. Third, by synthesizing images with various attributes and attribute intensities, we qualitatively illustrate the CDNet's ability to control the degree of disentanglement. Fourthly, we propose a classification based protocol to quantitatively compare the disentanglement strength of the CDNet. Finally, we perform an ablation study to investigate the effectiveness of different loss terms. Additional results are provided in the

²We empirically found that the extended interval $[-2, 5]$ is large enough for our model to generate new face images with various facial attribute intensities.

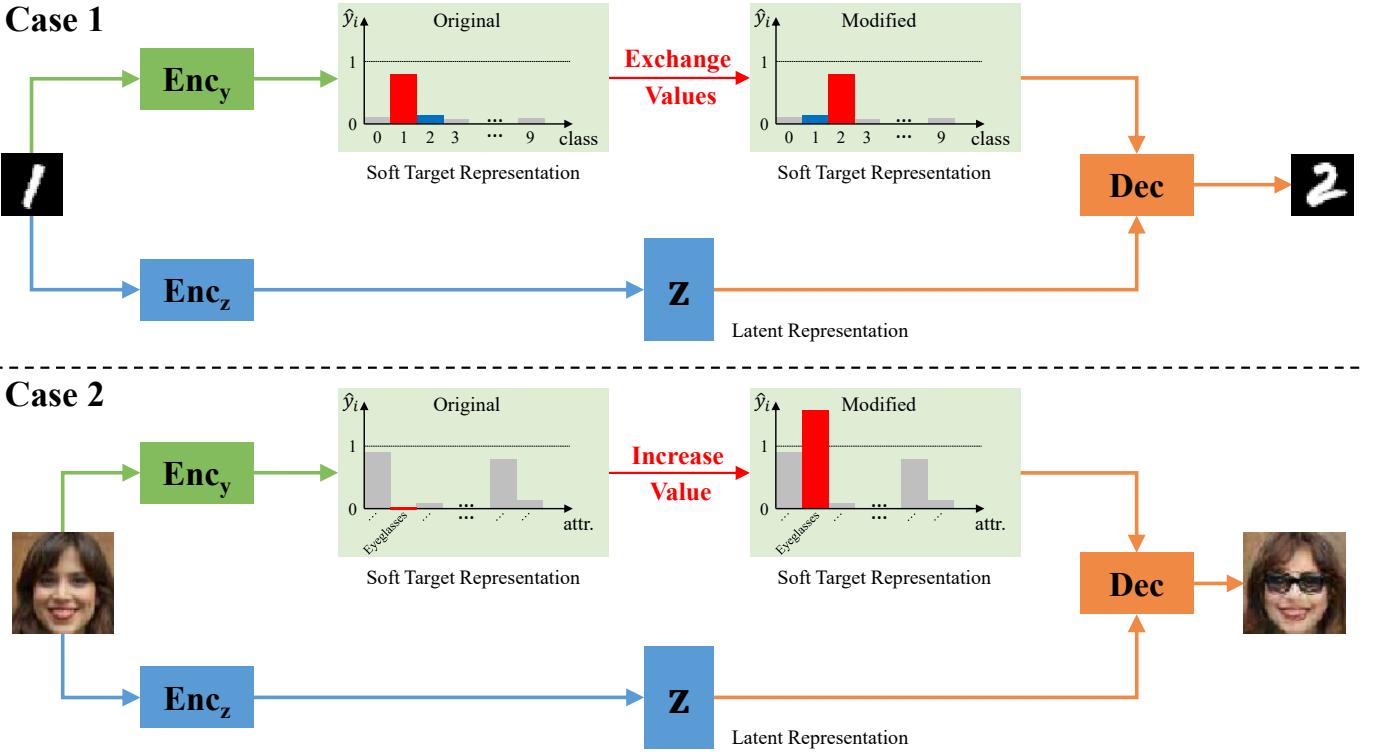


Fig. 2. Manipulating images. *Case 1* addresses the multiclass scenario where the mutual exclusion exists among multiple classes. *Case 2* is for the multilabel scenario where multiple attributes are independent of each other. The detailed procedure is illustrated in Section III-D. Best viewed in color.

supplementary material. Before presenting our experimental results, we introduce the experimental setup in detail.

A. Experimental Setup

1) *Datasets*: We use two representative datasets for the two application cases described in Section III-A1. The first one is MNIST [42], which contains 70,000 grayscale handwritten digit images with 28×28 pixels for each and scaled to $[0, 1]$. We randomly split the dataset into 50,000 training, 10,000 validation, and 10,000 test samples, respectively. The discrete label has the one-hot vector form. The second dataset is CelebA [43], which consists of 202,599 RGB face images of celebrities. For pre-processing, all face images are first center-cropped and then downsampled to 64×64 RGB pixels and scaled to $[-1, 1]$. We use 80% images for training, 10% for validation, and 10% for test as used in several earlier works. Additionally, the discrete label is represented by a binary vector with dimensionality 40, where each dimension corresponds to one attribute with value 1 indicating containing that attribute and 0 not.

2) *Baselines*: We evaluate our CDNet with three baselines which have similar network architectures to ours, and we give all models as follows.

- 1) AE-XCov [10]: a pure AE-based model that leverages the cross covariance (XCov) to learn the middle-layer uncorrelated representations.
- 2) IcGAN [17]: a model introducing an encoder to GAN to implement the inference mechanism, and it achieves dis-

entanglement via modifying the discrete labels inferred by the encoder.

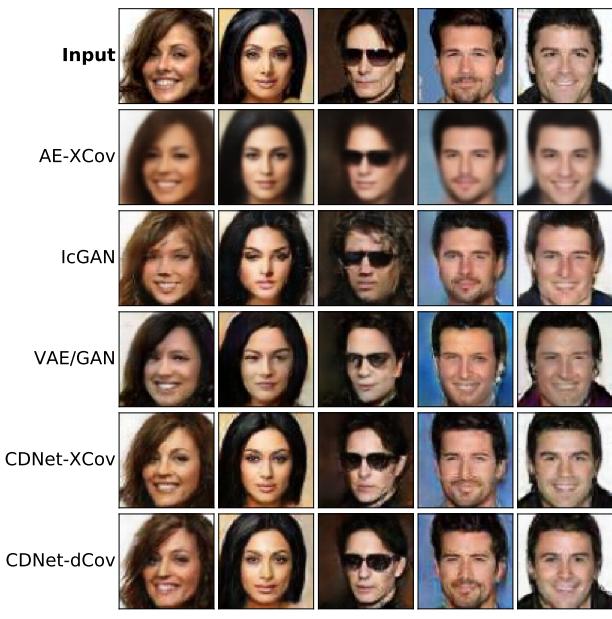
- 3) VAE/GAN [22]: a model that combines a VAE with a GAN, and its reconstruction error is derived only in the high-level feature space. The disentangled representation is computed using the vector arithmetic method in [33].
- 4) CDNet-XCov (ours): an instantiation of our CDNet model, where the XCov is used as the decorrelation loss.
- 5) CDNet-dCov (ours): another instantiation of the CDNet model, which utilizes the distance covariance (dCov)-based regularization, proposed in this paper, to learn independent representations.

The AE-XCov serves to illustrate the problem of blurring effects produced by the pure AE-based models. The IcGAN and VAE/GAN are selected to confirm that the decorrelation regularization employed in our model is beneficial to learn disentangled representations, and that the integration of both the pixel-wise and the feature-wise reconstruction errors is helpful to improve the reconstruction quality. The CDNet-dCov is compared with the CDNet-XCov to empirically verify that under the same settings, the dCov regularization can facilitate models to generate easier-perceptible attributes than the XCov.

3) *Training Details*: We implement the IcGAN and the VAE/GAN based on the network architectures as shown in [17] and [22], respectively. For the CDNet, the two encoders have the same architecture, which consists of convolution layers followed by fully-connected layers. The main part of decoder is symmetric to encoder, but using deconvolution

Input	0	1	2	3	4	5	6	7	8	9
AE-XCov	0	1	2	3	4	5	6	7	8	9
IcGAN	0	1	2	3	4	5	6	7	8	9
VAE/GAN	0	1	2	3	4	5	6	7	8	9
CDNet-XCov	0	1	2	3	4	5	6	7	8	9
CDNet-dCov	0	1	2	3	4	5	6	7	8	9

(a)



(b)

Fig. 3. Reconstructions of (a) handwritten digit images from MNIST test set and (b) face images from CelebA test set. Best viewed in color.

(a.k.a. transposed convolution or fractional striding) [44] for the up-sampling. The AE-XCov is built with the similar way to CDNet, except that we append the soft target representation to each layer of the CDNet decoder to strength the influence of class or attribute information on output images. The architecture details can be found in the supplementary material.

We perform the hyperparameter selection according to the validation-set performance. Specifically, for both of the two datasets, λ_{rec} takes the value of 1 to balance the two reconstruction errors. The λ_{adv} appeared in Algorithm 1 weights the reconstruction ability of decoder v.s. fooling the discriminator, equaling to 0.01. For the MNIST dataset, λ_{dcorr} is linearly increased to 1 over the first 50,000 iterations, while for the CelebA dataset, λ_{dcorr} is gradually increased until it reaches 0.05 across the first 50,000 iterations. The CDNet models are trained with the RMSProp optimizer, where we set the learning rate = 0.0001 and a batch size of 100 for MNIST, the learning rate = 0.00005 and a batch size of 128 for CelebA. With a single NVIDIA GeForce GTX 1080 GPU, training CDNet-XCov takes about 1.33 hours on MNIST and 11.17 hours on CelebA; for CDNet-dCov, the training time is about 1.34 hours on MNIST and 11.23 hours on CelebA.

TABLE I
RECONSTRUCTION QUALITY ON CELEBA TEST SET. THE RESULTS ARE FORMATTED AS MEAN (\pm STANDARD DEVIATION). BEST TWO SCORES IN EACH COLUMN ARE HIGHLIGHTED IN BOLD

Model	RMSE	PSNR	SSIM
AE-XCov	0.0991 (\pm 0.0209)	20.2640 (\pm 1.7993)	0.8701 (\pm 0.0585)
IcGAN	0.1668 (\pm 0.0388)	15.7877 (\pm 1.9966)	0.7236 (\pm 0.1247)
VAE/GAN	0.1403 (\pm 0.0287)	17.2311 (\pm 1.7287)	0.7804 (\pm 0.0921)
CDNet-XCov	0.0834 (\pm 0.0175)	21.7564 (\pm 1.7683)	0.9099 (\pm 0.0377)
CDNet-dCov	0.0828 (\pm 0.0172)	21.8181 (\pm 1.7580)	0.9108 (\pm 0.0377)

B. Reconstruction

We first visualize the reconstruction results to make a qualitative comparison between different models. As shown in Fig. 3 (a), when reconstructing handwritten digit images from MNIST, all models consistently perform well in terms of the visual perception quality. However, the performance gap becomes large as reconstructing face images from CelebA. As we can see from Fig. 3 (b), reconstructions of the AE-XCov contain the main structure of the input face image, but lose the details due to the blurring effect. The IcGAN on the contrary works well to recover texture features, while apparently falsifying the core object identity in reconstructed images. Although the VAE/GAN approaches a balance between the reconstruction accuracy and the visual fidelity, it cannot recover some specific local features, such as hair texture and eyeglasses. By contrast, our two CDNet models can reconstruct images with higher visual fidelity. The results demonstrate that as a combination of AE and GAN, the CDNet enjoys two advantages for reconstruction tasks, i.e., preserving the core object identity and simultaneously recovering local detail features.

To quantitatively evaluate the reconstruction ability of the proposed model, we utilize three well-known image quality assessment indexes, namely root-mean-square error (RMSE), peak signal-to-noise ratio (PSNR), and multi-scale structure similarity (SSIM) [45], to assess the reconstructed images' quality. As can be seen from Table I, the two instantiations of the CDNet significantly outperform the other models across all three evaluation metrics. We conclude that in the CDNet, the integration of the pixel-wise reconstruction error (from AE) and the feature-wise reconstruction error (from GAN) provides an effective way to improve the reconstruction quality.

C. Disentanglement

In this group of experiments, we first use all models to generate new handwritten digits with the designated handwriting styles, such as boldface, italic, and broad shape. The image manipulation method is described as the *Case 1* in Section III-D. As shown in Fig. 4 (a), the AE-XCov presents a limited disentanglement ability, such as synthesizing new

0 0 1 2 3 4 5 6 7 8 9	0 0 1 2 3 4 5 6 7 8 9	0 0 1 2 3 4 5 6 7 8 9	0 0 1 2 3 4 5 6 7 8 9	0 0 1 2 3 4 5 6 7 8 9
1 0 1 2 3 4 5 6 7 8 9	1 0 1 2 3 4 5 6 7 8 9	1 0 1 2 3 4 5 6 7 8 9	1 0 1 2 3 4 5 6 7 8 9	1 0 1 2 3 4 5 6 7 8 9
2 0 1 2 3 4 5 6 7 8 9	2 0 1 2 3 4 5 6 7 8 9	2 0 1 2 3 4 5 6 7 8 9	2 0 1 2 3 4 5 6 7 8 9	2 0 1 2 3 4 5 6 7 8 9
3 0 1 2 3 4 5 6 7 8 9	3 0 1 2 3 4 5 6 7 8 9	3 0 1 2 3 4 5 6 7 8 9	3 0 1 2 3 4 5 6 7 8 9	3 0 1 2 3 4 5 6 7 8 9
4 0 1 2 3 4 5 6 7 8 9	4 0 1 2 3 4 5 6 7 8 9	4 0 1 2 3 4 5 6 7 8 9	4 0 1 2 3 4 5 6 7 8 9	4 0 1 2 3 4 5 6 7 8 9
5 0 1 2 3 4 5 6 7 8 9	5 0 1 2 3 4 5 6 7 8 9	5 0 1 2 3 4 5 6 7 8 9	5 0 1 2 3 4 5 6 7 8 9	5 0 1 2 3 4 5 6 7 8 9
6 0 1 2 3 4 5 6 7 8 9	6 0 1 2 3 4 5 6 7 8 9	6 0 1 2 3 4 5 6 7 8 9	6 0 1 2 3 4 5 6 7 8 9	6 0 1 2 3 4 5 6 7 8 9
7 0 1 2 3 4 5 6 7 8 9	7 0 1 2 3 4 5 6 7 8 9	7 0 1 2 3 4 5 6 7 8 9	7 0 1 2 3 4 5 6 7 8 9	7 0 1 2 3 4 5 6 7 8 9
8 0 1 2 3 4 5 6 7 8 9	8 0 1 2 3 4 5 6 7 8 9	8 0 1 2 3 4 5 6 7 8 9	8 0 1 2 3 4 5 6 7 8 9	8 0 1 2 3 4 5 6 7 8 9
9 0 1 2 3 4 5 6 7 8 9	9 0 1 2 3 4 5 6 7 8 9	9 0 1 2 3 4 5 6 7 8 9	9 0 1 2 3 4 5 6 7 8 9	9 0 1 2 3 4 5 6 7 8 9

(a)

(b)

(c)

(d)

(e)

Fig. 4. Generated digits with different handwriting styles. (a) AE-XCov. (b) IcGAN. (c) VAE/GAN. (d) CDNet-XCov. (e) CDNet-dCov. In each panel, the first column displays MNIST test images, and the other columns show analogical fantasies of test images.



Fig. 5. Synthesized face images with the designated attributes. The first column shows original test images, the second column displays reconstructed counterparts, the next eight columns illustrate synthesized faces with the single target attribute, and the rightmost three columns correspond to examples of multi-attribute manipulations (H: Blond hair, E: Eyeglasses, M: Mustache). Best viewed in color.

digits slightly leaning to the left (corresponding to the test digit “2”). The disentanglement performance of the VAE/GAN is not stable, as many synthesized digits cannot reflect the category nature (see Fig. 4 (c)). We attribute this drawback to that in VAE/GAN, there still exist strong correlations between the learned representations of different digit classes. Both of the two CDNets, as well as the IcGAN, are able to generate new digits with the same style as originals, demonstrating the

disentanglement of style from class.

Second, we aim to synthesize new faces with the specific facial attributes while preserving the core identity. The manipulation method is described as the *Case 2* in Section III-D. And the disentanglement performance of different models on CelebA face images is illustrated in Fig. 5. We can observe that the AE-XCov can produce new faces with the desired attributes, but the blurring effects obviously degrade the output



Fig. 6. Synthesized face images with various attribute intensities. (a) Brown hair. (b) Pale skin. (c) Eyeglasses. (d) Smiling. (e) Mustache. (f) w/o Eyeglasses. The results in each panel, from the first row to the last row, are obtained by AE-XCov, IcGAN, VAE/GAN, CDNet-XCov, and CDNet-dCov, respectively. In each panel, the first column shows the original test image, the second column for reconstructions, and the remaining five columns for synthesized images with different attribute intensities, from weaker levels to stronger ones. Best viewed in color.

images' visual quality. For the IcGAN model, although target attributes are clearly involved in the resulting images, the core object identities have been changed as occurred in the reconstruction task. The synthetic results of VAE/GAN turn out to be really sensitive to the attribute modifications. In particular, when leveraging the VAE/GAN to add the "Eyeglasses" or "Mustache" attribute into the original female face, the hairstyle and even the gender of the subject are apparently altered. When it comes to the multi-attribute manipulations, all three baseline models fail to mix different attributes at once, since the multi-attribute changes induce the problem of visually-perceptible image distortion. By contrast, the two CDNets exhibit a remarkable ability to disentangle facial attributes from identity, that is, all designated attributes are incorporated into the final images in a more natural manner.

It's worth noting that on the multi-attribute manipulation task, especially for the male face image editing, the CDNet-dCov visually outperforms the CDNet-XCov in terms of preserving the core identity information (e.g., see deformations of the mouth area shown in the last three columns of Fig. 5). We will give a quantitative comparison between them in Section IV-E.

D. Controllable Disentanglement

In this experiment, we *qualitatively* compare the disentanglement strength of all baselines and the CDNet by synthesizing face images with various attribute intensities. The image manipulation method is similar to the *Case 2* in Section

III-D, and here we take a list of values to orderly modify the corresponding attribute representations. The attribute value range is set to $[-0.5, 6.0]$ for AE-XCov, $[0.0, 6.0]$ for IcGAN, $[-1.0, 3.0]$ for VAE/GAN, and $[-1.5, 3.0]$ for CDNet. We conduct this group of experiments on six representative facial attributes, i.e., "Brown hair", "Pale skin", "Eyeglasses", "Smiling", "Mustache", and "w/o Eyeglasses", respectively. From Fig. 6, it is observed that the blurring effect still exists across all synthetic images generated by AE-XCov. The IcGAN performs well to synthesize a set of new faces with target attributes and attribute intensities; however, none of the resulting images can preserve the core object identities very well. For the VAE/GAN, changing one attribute usually causes the deformation of other attributes. One of such examples is that the modifications on "Pale skin" attribute also cause the deformation of hairstyle. Our two CDNet models, by contrast, are competent to generate distinguishable face fantasies across all compared attributes and variation degrees, and meanwhile preserving the core identity. These results illustrate that the learning strategy of representations, as well as the image manipulation method, enables CDNet to control the degree of disentanglement during image editing.

E. Comparing Disentanglement by Classification

To further analyze the difference between XCov and dCov² for disentanglement, we design an evaluation protocol to *quantitatively* compare the disentanglement strength of the CDNet-

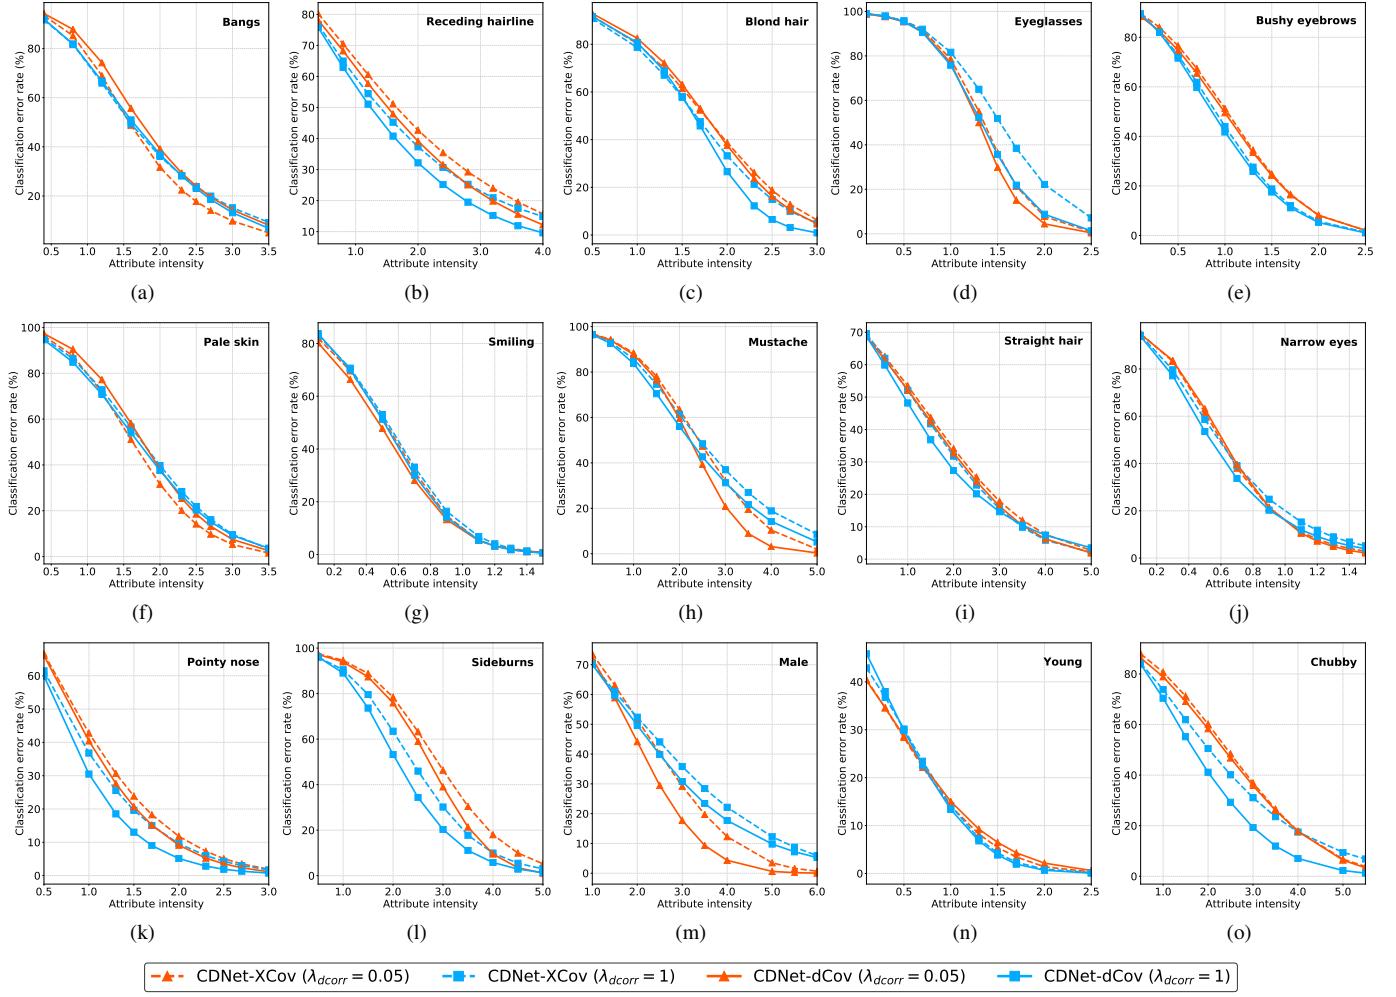


Fig. 7. Attribute classification of face images synthesized by CDNet-XCov and CDNet-dCov. (a) Bangs. (b) Receding hairline. (c) Blond hair. (d) Eyeglasses. (e) Bushy eyebrows. (f) Pale skin. (g) Smiling. (h) Mustache. (i) Straight hair. (j) Narrow eyes. (k) Pointy nose. (l) Sideburns. (m) Male. (n) Young. (o) Chubby. For each attribute, we train five linear SVMs independently and report the average performance on each of the ten attribute intensities.

Xcov and the CDNet-dCov. We summarize the evaluation procedure into the following four steps.

- 1) We divide the training set into two subsets: the first subset consists of images with the designated attribute, the second one not.
- 2) We train a two-class classifier on the two subsets.
- 3) We select all test set images that do not contain the designated attribute, then feed them to the disentanglement model to generate their counterparts with the designated attribute and intensity.
- 4) Finally, we employ the classifier trained in the second step to classify those images synthesized in the third step, which produces a classification error rate as the evaluation index.

The evaluation protocol is built based on a hypothesis, that is, the classifier is well-trained and therefore lower error rate means it is easier for the classifier to perceive the designated attribute in synthesized images. For each attribute, we train a linear SVM as the two-class classifier to perform attribute classification tasks. Each attribute is assigned ten different

intensities, arranged from the lowest level to the highest one, and the two disentanglement models act to synthesize images with both the designated attribute and these corresponding attribute intensities. The main evaluation results are illustrated in Fig. 7.

As shown in Fig. 7, for each facial attribute, the classification error rates consistently decrease as increasing attribute intensities. This result implies that by taking higher attribute intensities, the two CDNets can synthesize images with more distinct attributes. In addition, under the same network architecture and parameter settings, the classification performance of CDNet-dCov is comparable and even superior to that of CDNet-XCov, especially when improving the influence of the decorrelation regularization on learning latent representations (i.e., setting $\lambda_{dcorr} = 1$). We attribute this performance gap to the strong disentanglement ability of dCov², as minimizing dCov² induces independence between two random variables, rather than the non-correlation as approximated by minimizing XCov. For this reason, in the CDNet-dCov, modifications on the target attributes have less effect on the core iden-

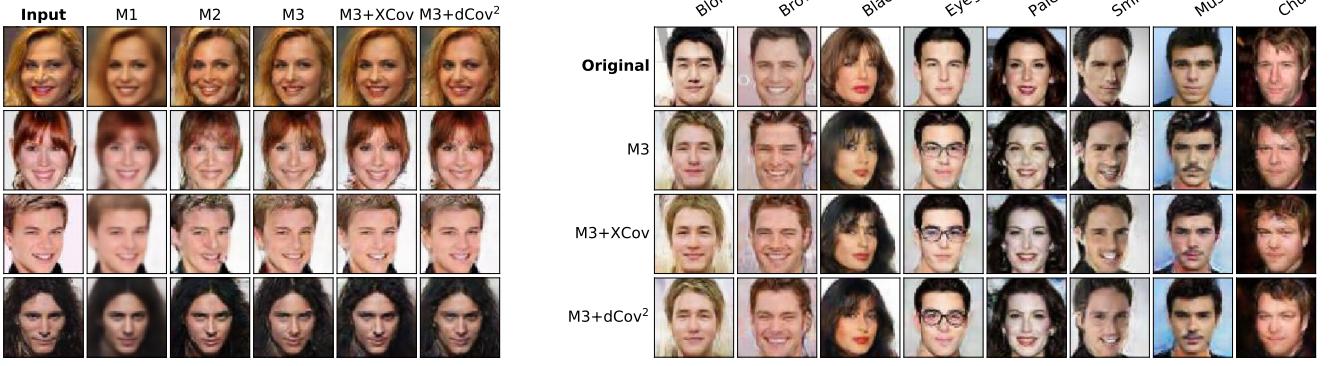


Fig. 8. Qualitative ablation evaluation of different loss terms. (a) Reconstruction results. (b) Disentanglement results. The images are generated by models trained with different local losses. M1: $\mathcal{L}_{class} + \mathcal{L}_{rec}^{pix}$. M2: $\mathcal{L}_{class} + \mathcal{L}_{rec}^{feat} + \mathcal{L}_{adv}$. M3: $\mathcal{L}_{class} + \mathcal{L}_{rec} + \mathcal{L}_{adv}$. Best viewed in color.

ity information contained in the latent representation. This property enables the CDNet-dCov to synthesize images with more concrete and distinguishable attributes, compared with the CDNet-XCov.

F. Ablation Study

To verify the impact of different loss terms of the proposed model, we conduct an ablation study on the reconstruction task and the disentanglement task, respectively. From the reconstruction results in Fig. 8(a), we observe that images generated by the model M1 are quite blurry and only capture the coarse shape of faces; and that images produced by the model M2 show more texture features, but followed with the content deformation in some regions (e.g., hair and mouth). The results demonstrate that minimizing pixel-wise reconstruction loss encourages the model to preserve global identity information, while minimizing feature-wise reconstruction loss is beneficial to restore local detail features. Consequently, models trained with both of these two reconstruction losses perform better in generating more plausible face images, as indicated by reconstruction results of M3, M3+XCov, and M3+dCov².

For the disentanglement comparison shown in Fig. 8(b), the plain model M3 (trained without using any decorrelation term) is limited to synthesize target facial attributes being consistent with the context. For example, changing one attribute, such as “Smiling” or “Mustache”, also causes the considerable image quality degradation near the nose area. Besides, M3 cannot preserve the hairstyle during image editing, which can be found in the synthesized faces corresponding to attributes “Brown hair”, “Black hair”, “Chubby”, and so on. By contrast, when training the same neural network using the additional decorrelation regularization (i.e., M3+XCov and M3+dCov²), the designated attributes can be blended in new faces properly and with less effect on other attributes. The results suggest that the decorrelation term indeed facilitates models to learn uncorrelated representations during training, thus approaching controllable image editing at test time. Furthermore, we find the synthesized images of M3+dCov² are more visually realistic and coherent than those of M3+XCov. This difference

in performance is observed especially when comparing the hairstyle of subjects under the “Black hair” and the “Chubby” attributes, or comparing the mouth area of subjects under the “Brown hair” attribute. Our conclusion is that compared with the XCov, the dCov² is more helpful for models to separate the information of interest from other portions, resulting in the ability to synthesize pleasant images in an easier-controllable manner.

V. CONCLUSION

In this paper, we proposed a simple yet effective model that aims to address two disentanglement-related problems: controlling the degree of disentanglement at image editing time, and balancing the disentanglement strength and the reconstruction quality. A distance covariance based decorrelation regularization was devised to encourage disentanglement, and the soft target representation was explored to control how much a specific attribute is perceivable in the generated image. Besides, a method of combining AE with GAN was designed to improve the visual quality of reconstructed and synthetic images. In addition, we also developed a classification protocol to quantitatively evaluate the disentanglement strength of our model. Experimental results demonstrated that our model is able to generate new digits with various handwriting styles, and also to synthesize novel faces with the desired attributes and the attribute intensities.

The supervised disentanglement learning, as well as the decorrelation regularization used in this work, enables the model to learn target representations effectively. However, acquiring a large amount of labeled training data is usually costly and time consuming. To alleviate this problem, we can consider extending the current model to be suitable for the semi-/weakly supervised learning scenario. One possible way to approach this goal is, as discussed in [46], devising probabilistic models for inductive and transductive semi-supervised learning, which can be further implemented by using the approximate Bayesian inference method. Moreover, exploring the concrete benefits of disentangled representations

for downstream tasks is another promising direction in this research filed [34].

APPENDIX

We leverage the theoretical results in [30] to prove Theorem 1. Overall, we use properties of the distance correlation to connect the sample distance covariance (dCov) to the independence between two random variables. The following one definition and two lemmas correspond to the Definition 3, Theorem 2, and Theorem 3 in [30], respectively. One can find the complete proofs to the two lemmas therein.

Definition 1: The distance correlation between two random vectors \mathbf{U} and \mathbf{V} with finite first moments is the nonnegative number $\mathcal{R}(\mathbf{U}, \mathbf{V})$ defined by

$$\mathcal{R}^2(\mathbf{U}, \mathbf{V}) = \begin{cases} \frac{\mathcal{V}^2(\mathbf{U}, \mathbf{V})}{\sqrt{\mathcal{V}^2(\mathbf{U})\mathcal{V}^2(\mathbf{V})}}, & \mathcal{V}^2(\mathbf{U})\mathcal{V}^2(\mathbf{V}) > 0, \\ 0, & \mathcal{V}^2(\mathbf{U})\mathcal{V}^2(\mathbf{V}) = 0 \end{cases} \quad (\text{A.1})$$

where $\mathcal{V}(\mathbf{U}, \mathbf{V})$ is the distance covariance between \mathbf{U} and \mathbf{V} , $\mathcal{V}(\mathbf{U})$ and $\mathcal{V}(\mathbf{V})$ are distance variances of \mathbf{U} and \mathbf{V} , respectively.

Lemma 1: If $E(\|\mathbf{U}\|_2) < \infty$ and $E(\|\mathbf{V}\|_2) < \infty$, then almost surely

$$\lim_{K \rightarrow \infty} \text{dCov}_K(\mathbf{U}, \mathbf{V}) \stackrel{a.s.}{=} \mathcal{V}(\mathbf{U}, \mathbf{V}). \quad (\text{A.2})$$

Lemma 2: If $E(\|\mathbf{U}\|_2 + \|\mathbf{V}\|_2) < \infty$, then $0 \leq \mathcal{R} \leq 1$, and $\mathcal{R}(\mathbf{U}, \mathbf{V}) = 0$ if and only if \mathbf{U} and \mathbf{V} are independent.

Proof of Theorem 1: According to the given condition that $E(\|\mathbf{U}\|_2) < \infty$ and $E(\|\mathbf{V}\|_2) < \infty$, Lemma 1 holds. Then we have

$$\begin{aligned} \lim_{K \rightarrow \infty} \text{dCov}_K^2(\mathbf{U}, \mathbf{V}) &= \left(\lim_{K \rightarrow \infty} \text{dCov}_K(\mathbf{U}, \mathbf{V}) \right)^2 \\ &\stackrel{a.s.}{=} \mathcal{V}^2(\mathbf{U}, \mathbf{V}). \end{aligned} \quad (\text{A.3})$$

The first step in Eq. (A.3) is done because the two limitations $\lim_{K \rightarrow \infty} \text{dCov}_K^2(\mathbf{U}, \mathbf{V})$ and $\lim_{K \rightarrow \infty} \text{dCov}_K(\mathbf{U}, \mathbf{V})$ exist, and the second step simply applies Lemma 1. Notice that we assumed $\lim_{K \rightarrow \infty} \text{dCov}_K^2(\mathbf{U}, \mathbf{V}) = 0$, there holds

$$\mathcal{V}^2(\mathbf{U}, \mathbf{V}) \stackrel{a.s.}{=} 0. \quad (\text{A.4})$$

Based on Definition 1, we get

$$\mathcal{R}(\mathbf{U}, \mathbf{V}) \stackrel{a.s.}{=} 0. \quad (\text{A.5})$$

The condition $E(\|\mathbf{U}\|_2 + \|\mathbf{V}\|_2) < \infty$ in Lemma 2 is satisfied as $E(\|\mathbf{U}\|_2) < \infty$ and $E(\|\mathbf{V}\|_2) < \infty$. Thus, from Lemma 2 and Eq. (A.5), almost surely \mathbf{U} and \mathbf{V} are independent. This completes the proof of Theorem 1. \blacksquare

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions. They also thank Microsoft Azure for computing resources.

REFERENCES

- [1] Y. Bengio, A. Courville, and P. Vincent, “Representation learning: A review and new perspectives,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [2] K. Ridgeway, “A survey of inductive biases for factorial representation-learning,” *arXiv:1612.05299*, 2016.
- [3] E. Dupont, “Learning disentangled joint continuous and discrete representations,” in *NeurIPS*, 2018.
- [4] G. Lample, N. Zeghidour, N. Usunier, A. Bordes, L. Denoyer, and M. Ranzato, “Fader networks: Manipulating images by sliding attributes,” in *NIPS*, 2017.
- [5] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, “Building machines that learn and think like people,” *Behavioral and Brain Sciences*, vol. 40, 2017.
- [6] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, “Deep variational information bottleneck,” in *ICLR*, 2017.
- [7] R. T. Q. Chen, X. Li, R. Grosse, and D. Duvenaud, “Isolating sources of disentanglement in variational autoencoders,” in *NeurIPS*, 2018.
- [8] I. Higgins, A. Pal, A. Rusu, L. Matthey, C. Burgess, A. Pritzel, M. Botvinick, C. Blundell, and A. Lerchner, “DARLA: Improving zero-shot transfer in reinforcement learning,” in *ICML*, 2017.
- [9] J. Schmidhuber, “Learning factorial codes by predictability minimization,” *Neural Computation*, vol. 4, no. 6, pp. 863–879, 1992.
- [10] B. Cheung, J. A. Livezey, A. K. Bansal, and B. A. Olshausen, “Discovering hidden factors of variation in deep networks,” in *ICLR Workshop*, 2015.
- [11] M. Kågåbeck and O. Mogren, “Disentanglement by penalizing correlation,” in *NIPS Workshop*, 2017.
- [12] D. P. Kingma and M. Welling, “Auto-encoding variational Bayes,” in *ICLR*, 2014.
- [13] D. J. Rezende, S. Mohamed, and D. Wierstra, “Stochastic backpropagation and approximate inference in deep generative models,” in *ICML*, 2014.
- [14] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, “beta-VAE: Learning basic visual concepts with a constrained variational framework,” in *ICLR*, 2017.
- [15] A. Kumar, P. Sattigeri, and A. Balakrishnan, “Variational inference of disentangled latent concepts from unlabeled observations,” in *ICLR*, 2018.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *NIPS*, 2014.
- [17] G. Perarnau, J. van de Weijer, B. Raducanu, and J. M. Álvarez, “Invertible conditional GANs for image editing,” in *NIPS Workshop*, 2016.
- [18] L. Tran, X. Yin, and X. Liu, “Disentangled representation learning GAN for pose-invariant face recognition,” in *CVPR*, 2017.
- [19] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, “StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation,” in *CVPR*, 2018.
- [20] H. Dong, X. Liang, K. Gong, H. Lai, J. Zhu, and J. Yin, “Soft-gated warping-GAN for pose-guided person image synthesis,” in *NeurIPS*, 2018.
- [21] A. Makhzani, J. Shlens, N. Jaitly, and I. Goodfellow, “Adversarial autoencoders,” in *ICLR Workshop*, 2016.
- [22] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther, “Autoencoding beyond pixels using a learned similarity metric,” in *ICML*, 2016.
- [23] V. Dumoulin, I. Belghazi, B. Poole, O. Mastropietro, A. Lamb, M. Arjovsky, and A. Courville, “Adversarially learned inference,” in *ICLR*, 2017.
- [24] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim, “Learning to discover cross-domain relations with generative adversarial networks,” in *ICML*, 2017.
- [25] C. Donahue, Z. C. Lipton, A. Balsubramani, and J. McAuley, “Semantically decomposing the latent spaces of generative adversarial networks,” in *ICLR*, 2018.
- [26] J. Engel, M. D. Hoffman, and A. Roberts, “Latent constraints: Learning to generate conditionally from unconditional generative models,” in *ICLR*, 2018.
- [27] T. D. Kulkarni, W. F. Whitney, P. Kohli, and J. B. Tenenbaum, “Deep convolutional inverse graphics network,” in *NIPS*, 2015.
- [28] X. Chen, Y. Duan, R. Houthooft, J. Schulman, I. Sutskever, and P. Abbeel, “InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets,” in *NIPS*, 2016.

- [29] L. Ma, Q. Sun, S. Georgoulis, L. V. Gool, B. Schiele, and M. Fritz, “Disentangled person image generation,” in *CVPR*, 2018.
- [30] G. J. Székely, M. L. Rizzo, and N. K. Bakirov, “Measuring and testing dependence by correlation of distances,” *Ann. Stat.*, vol. 35, no. 6, pp. 2769–2794, 2007.
- [31] M. Mirza and S. Osindero, “Conditional generative adversarial nets,” *arXiv:1411.1784*, 2014.
- [32] X. Hou, L. Shen, K. Sun, and G. Qiu, “Deep feature consistent variational autoencoder,” in *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2017.
- [33] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *NIPS*, 2013.
- [34] F. Locatello, S. Bauer, M. Lucic, G. Rätsch, S. Gelly, B. Schölkopf, and O. Bachem, “Challenging common assumptions in the unsupervised learning of disentangled representations,” in *ICML*, 2019.
- [35] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, “Improved techniques for training GANs,” in *NIPS*, 2016.
- [36] L. A. Gatys, A. S. Ecker, and M. Bethge, “Image style transfer using convolutional neural networks,” in *CVPR*, 2016.
- [37] J. Johnson, A. Alahi, and L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *ECCV*, 2016.
- [38] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *ICLR*, 2015.
- [39] C. Li and M. Wand, “Precomputed real-time texture synthesis with Markovian generative adversarial networks,” in *ECCV*, 2016.
- [40] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, “Improved training of Wasserstein GANs,” in *NIPS*, 2017.
- [41] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [42] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [43] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” in *ICCV*, 2015.
- [44] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” in *ICLR*, 2016.
- [45] Z. Wang, E. P. Simoncelli, and A. C. Bovik, “Multi-scale structural similarity for image quality assessment,” in *IEEE Asilomar Conference on Signals, Systems and Computers*, 2003.
- [46] D. P. Kingma, D. J. Rezende, S. Mohamed, and M. Welling, “Semi-supervised learning with deep generative models,” in *NIPS*, 2014.