

Generative Adversarial Image Synthesis with Decision Tree Latent Controller

Takuhiro Kaneko Kaoru Hiramatsu Kunio Kashino
NTT Communication Science Laboratories, NTT Corporation

{kaneko.takuhiro, hiramatsu.kaoru, kashino.kunio}@lab.ntt.co.jp

Abstract

This paper proposes the decision tree latent controller generative adversarial network (DTLC-GAN), an extension of a GAN that can learn hierarchically interpretable representations without relying on detailed supervision. To impose a hierarchical inclusion structure on latent variables, we incorporate a new architecture called the DTLC into the generator input. The DTLC has a multiple-layer tree structure in which the ON or OFF of the child node codes is controlled by the parent node codes. By using this architecture hierarchically, we can obtain the latent space in which the lower layer codes are selectively used depending on the higher layer ones. To make the latent codes capture salient semantic features of images in a hierarchically disentangled manner in the DTLC, we also propose a hierarchical conditional mutual information regularization and optimize it with a newly defined curriculum learning method that we propose as well. This makes it possible to discover hierarchically interpretable representations in a layer-by-layer manner on the basis of information gain by only using a single DTLC-GAN model. We evaluated the DTLC-GAN on various datasets, i.e., MNIST, CIFAR-10, Tiny ImageNet, 3D Faces, and CelebA, and confirmed that the DTLC-GAN can learn hierarchically interpretable representations with either unsupervised or weakly supervised settings. Furthermore, we applied the DTLC-GAN to image-retrieval tasks and showed its effectiveness in representation learning.

1. Introduction

There have been recent advances in computer vision and graphics, enabling photo-realistic images to be created. However, it still requires considerable skill or effort to create a pixel-level detailed image from scratch. Deep generative models, such as generative adversarial networks (GANs) [10] and variational autoencoders (VAEs) [17, 37], have recently emerged as powerful models to alleviate this difficulty. Although these models make it possible to generate various images with high fidelity quality by changing (e.g., randomly sampling) latent variables in the generator or decoder input, there still remains a painstaking process to create the desired image because the naive formulation does

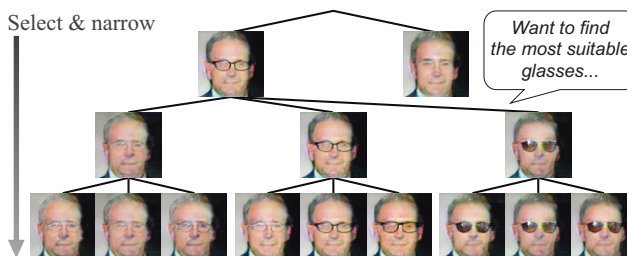


Figure 1. **Examples of image generation under control using DTLC-GAN:** DTLC-GAN enables image generation to be controlled in coarse-to-fine manner, i.e., “selected & narrowed.” Our goal is to discover such hierarchically interpretable representations without relying on detailed supervision.

not impose any structure on latent variables; as a result, they may be used by the generator or decoder in a highly entangled manner. This causes difficulty in interpreting the “meaning” of the individual variables and in controlling image generation by operating each one.

When we create an image from scratch, we typically select and narrow a target to paint in a coarse-to-fine manner. For example, when we create an image of a face with glasses, we first roughly consider the type of glasses, e.g., transparent glasses/sunglasses, then define the details, e.g., thin/thick rimmed glasses or small/big sunglasses. To use a deep generative model as a supporter for creating an image, we believe that such *hierarchically interpretable representation* is the key to obtaining the image one has in mind.

These facts motivated us to address the problem of *how to derive hierarchically interpretable representations in a deep generative model*. To solve this problem, we propose the decision tree latent controller GAN (DTLC-GAN), an extension of the GAN that can learn hierarchically interpretable representations without relying on detailed supervision. Figure 1 shows examples of image generation under control using the DTLC-GAN. If semantic features are represented in a hierarchically disentangled manner, we can approach a target image gradually and interactively.

To impose a hierarchical inclusion structure on latent variables, we incorporate a new architecture called the DTLC into the generator input. The DTLC has a multiple-layer tree structure in which the ON or OFF of the child node codes is controlled by the parent node codes. By using this architecture hierarchically, we can obtain the latent




Figure 2. **Generated image samples on CIFAR-10:** All images were generated from same noise but different latent codes. In each row, we varied second-, third-, and fourth-layer codes per nine images, per three images, and per image, respectively. Note that we learn these $10 \times 3 \times 3 \times 3 = 270$ hierarchically disentangled representations only with supervision of class labels. This model also achieves high inception score **8.80**. We give details in Section 6.3.

space in which the lower layer codes are selectively used depending on the higher layer codes.

On the problem of making the latent codes capture salient semantic features of images in a hierarchically disentangled manner in the DTLC, the main difficulty is that we need to discover representations disentangled in the following three stages: (1) disentanglement between the control target (e.g., glasses) and unrelated factors (e.g., identity); (2) coarse-to-fine disentanglement between layers, i.e., the higher layer codes capture rough categories, while the lower layer ones capture detailed categories; and (3) inner-layer disentanglement to control semantic features independently, i.e., when one code captures a semantic feature (e.g., thin glasses), another one captures a different semantic feature (e.g., thick glasses).

A possible solution would be to collect detailed annotations, the amount of which is large enough to solve the problems in a fully supervised manner. However, this approach incurs high annotation costs. Even though we have enough human resources, defining the detailed categories remains a nontrivial task. The latter problem is also addressed in the field of research concerned with attribute representations [30, 51] and is still an open issue. This motivated us to tackle a challenging condition in which hierarchically interpretable representations need to be learned without relying on detailed annotations. Under this condition, it is not trivial to solve all the above three disentanglement problems at the same time because they are not independent from each other but are interrelated. To mitigate these problems, we propose a hierarchical conditional mutual information regularization (HCMI), which is an extension of MI [4] and conditional MI (CMI) [15] to hierarchical conditional settings and optimized it with a newly defined curriculum learning [3] method that we also propose. This makes it possible to discover hierarchically interpretable representations in a layer-by-layer manner on the basis of information gain by only using a single DTLC-GAN model. This is noteworthy because we can learn expressive representations without large increase in calculation cost. Figure 2 shows typical examples on CIFAR-10, where we succeeded in learning expressive representations, i.e., $10 \times 3 \times 3 \times 3 = 270$ categories, are learned in a weakly supervised (i.e., only

10 class labels are supervised) manner. We evaluated our DTLC-GAN on various datasets, i.e., MNIST, CIFAR-10, Tiny ImageNet, 3D Faces, and CelebA, and confirmed that it can learn hierarchically interpretable representations with either unsupervised or weakly supervised settings. Furthermore, we applied our DTLC-GAN to image-retrieval tasks and showed its effectiveness in representation learning.

Contributions: Our contributions are summarized as follows. (1) We derive a novel functionality in a deep generative model, which enables semantic features of an image to be controlled in a coarse-to-fine manner. (2) To obtain this functionality, we incorporate a new architecture called the DTLC into a GAN, which imposes a hierarchical inclusion structure on latent variables. (3) We propose a regularization called the HCMI and optimize it with a newly defined curriculum learning method that we also propose. This makes it possible to learn hierarchically disentangled representations only using a single DTLC-GAN model without relying on detailed supervision. (4) We evaluated our DTLC-GAN on various datasets and confirmed its effectiveness in image generation and image-retrieval tasks. We provide supplementary materials including demo videos at <http://www.kecl.ntt.co.jp/people/kaneko.takuhiro/projects/dtgc-gan/>.

2. Related Work

Deep Generative Models: In computer vision and machine learning, generative image modeling is a fundamental problem. Recently, there was a significant breakthrough due to the emergence of deep generative models. These models roughly fall into two approaches: deterministic and stochastic. On the basis of deterministic approaches, Dosovitsky et al. [7] proposed a deconvolution network that generates 3D objects, and Reed et al. [36] and Yang et al. [50] proposed networks that approximate functions for image synthesis. There are three major models based on stochastic approaches. One is a VAE [17, 37], which is formulated as probabilistic graphical models and optimized by maximizing a variational lower bound on the data likelihood. Another is an autoregressive model [43], which breaks the data distribution into a series of conditional distributions

and uses neural networks to model them. The other is a GAN [10], which is composed of generator and discriminator networks. The generator is optimized to fool the discriminator, while the discriminator is optimized to distinguish between real and generated data. This min-max optimization makes the training procedure unstable, but several techniques [1, 2, 12, 26, 32, 39, 54] have recently been proposed to stabilize it. All these models have pros and cons. In this paper, we take a stochastic approach, particularly focusing on a GAN, and propose an extension to it because it has flexibility on latent variable design. Extension to other models remains a promising area for future work.

Disentangled Representation Learning: In the study of stochastic deep generative models, there have been attempts to learn disentangled representations similar to our attempt. Most of the studies addressed the problem in supervised settings and incorporated supervision into the networks. For example, attribute/class labels [28, 29, 42, 44, 48, 53], text descriptions [25, 34, 52], and object location descriptions [33, 35] are used as supervision. To reduce the annotation cost, extensions to semi-supervised settings have also recently been proposed [16, 39, 40]. The advantage of these settings is that disentangled representation can be explicitly learned following the supervision; however, the limitation is that learnable representations are restricted to supervision. To overcome this limitation, weakly supervised [15, 19, 24, 27] and unsupervised [4] models have recently been proposed, which discover meaningful hidden representation without relying on detailed annotations; however, these models are limited to discovering one-layer hidden representations, whereas the DTLC-GAN enables multi-layer hidden representations to be learned. We further discussed the relationship to the previous GANs in Section 4.4.

Hierarchical Representation: The other related topic is hierarchical representation. Previous studies have decomposed an image in various ways. The LAPGAN [5] and StackGAN [52] deconstruct an image by repeatedly down-sampling it, S²-GAN [46] decomposes the generative process to structure and style, VGAN [45] decomposes a video into foreground and background, and SGAN [13] learns multi-level representations in feature spaces of intermediate layers. Other studies [9, 11, 14, 20, 49] used recursive structures to draw images in a step-by-step manner. The main difference from these studies is that they derive hierarchical representations in a pixel space or feature space to improve the fidelity of an image, while we derive those in a latent space to improve the interpretability and controllability of latent codes. More recently, Zhao et al. [55] proposed an extension of a VAE called the VLAE to learn multi-layer hierarchical representations in a latent space similar to ours; however, the type of hierarchy is different from ours. They learn representations that are semantically independent among layers, whereas we learn those where lower layer codes are correlated with higher layer codes in a decision tree manner. We argue that such representation is nec-

essary to learn category-specific features and control image generation in a select-and-narrow manner.

3. Background: GAN

A GAN [10] is a framework for training a generative model using a min-max game. The goal is to learn generative distribution $P_G(\mathbf{x})$ that matches the real data distribution $P_{\text{data}}(\mathbf{x})$. It consists of two networks: a generator G that transforms noise $\mathbf{z} \sim P_z(\mathbf{z})$ into data space $\mathbf{x} = G(\mathbf{z})$, and a discriminator D that assigns probability $p = D(\mathbf{x}) \in [0, 1]$ when \mathbf{x} is a sample from P_{data} and assigns probability $1 - p$ when it is a sample from P_G . The $P_z(\mathbf{z})$ is a prior on \mathbf{z} . The D and G play a two-player min-max game with the following binary cross entropy:

$$\mathcal{L}_{\text{GAN}}(D, G) = \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim P_z(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

The D attempts to find the binary classifier for discriminating between true and generated data by maximizing this loss, whereas the G attempts to generate data indistinguishable from the true data by minimizing this loss.

4. DTLC-GAN

4.1. DTLC

In the naive GAN, latent variables are sampled from an unconditional prior and do not have any constraints on a structure. As a result, they may be used by the G in a highly entangled manner, causing difficulty in interpreting the “meaning” of the individual variables and in controlling image generation by operating each one. Motivated by this fact, we incorporate the DTLC into the generator input to impose a hierarchical inclusion structure on latent variables.

Notation: In the DTLC-GAN, the latent variables are decomposed into multiple levels. We first decompose the latent variables into two parts: $\hat{\mathbf{c}}_L$, which is a latent code derived from an L -layer DTLC and will target hierarchically interpretable semantic features, and \mathbf{z} , which is a source of incompressible noise that covers factors that are not represented by $\hat{\mathbf{c}}_L$. To derive $\hat{\mathbf{c}}_L$, the DTLC has a multiple-layer tree structure and is composed of L layer codes $\mathbf{c}_1, \dots, \mathbf{c}_L$. In each layer, \mathbf{c}_l is decomposed into N_l node codes $\mathbf{c}_l = (\mathbf{c}_l^1, \dots, \mathbf{c}_l^{N_l})$. To impose a hierarchical inclusion relationship between the l th and $(l + 1)$ th layers, an n th parent node code \mathbf{c}_l^n is associated with k_l child node codes $\mathbf{c}_{l+1}^{n,1}, \dots, \mathbf{c}_{l+1}^{n,k_l}$, where $\mathbf{c}_{l+1}^{n,i} = \mathbf{c}_{l+1}^{k_l(n-1)+i}$. By this definition, N_{l+1} is calculated as $N_{l+1} = N_l \times k_l$.

We can use both discrete and continuous variables as \mathbf{c}_l^n , but for simplicity, we treat the case in which parent node codes are discrete and the lowest layer codes are either discrete or continuous. In this case, \mathbf{c}_l^n is represented as a k_l dimensional onehot vector and each dimension $\mathbf{c}_l^{n(i)}$ ($i = 1, \dots, k_l$) is associated with one child node code $\mathbf{c}_{l+1}^{n,i}$.

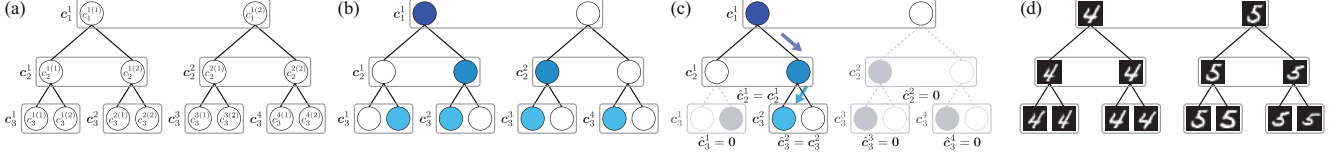


Figure 3. **Sampling example using three-layer DTLC:** (a) Architecture of three-layer DTLC where $k_1, k_2, k_3 = 2$. (b) Sampling example in Step 1. Each code is sampled from categorical distribution. Filled and open circles indicate 1 and 0, respectively. (c) Sampling example in Steps 2 and 3. ON or OFF of child node codes is selected by parent node codes. This execution is conducted recursively from highest layer to lowest layer. This imposes hierarchical inclusion constraints on sampling. (d) Sample images generated using this controller. Each image corresponds to each latent code. We tested on subset of MNIST dataset, which includes “4” and “5” digit images. This is relatively easy dataset; however, it is noteworthy that hierarchically disentangled representations, such as “4” or “5” in first layer and “narrow-width 4” or “wide-width 4” in second left layer, are learned in fully unsupervised manner.

Sampling Scheme: In a training phase, we sample latent codes as follows. We illustrate a sampling example in Figure 3.

1. We sample c_l^n ($l = 1, \dots, L-1$) from categorical distribution $c_l^n \sim \text{Cat}\left(K = k_l, p = \frac{1}{k_l}\right)$. We sample c_L^n in a similar manner in the discrete case, while we sample it from uniform distribution $c_L^n \sim \text{Unif}(-1, 1)$ in the continuous case. Note that, if we have supervision for c_l^n , we can directly use it instead of sampling.
2. To impose a hierarchical inclusion structure, we sample $\hat{c}_{l+1}^{n,i}$ from conditional prior $\hat{c}_{l+1}^{n,i} \sim P(\hat{c}_{l+1}^{n,i} | \hat{c}_l^n)$, where $\hat{c}_1 = c_1$. We do this with the following process:

$$\hat{c}_{l+1}^{n,i} = \hat{c}_l^{n(i)} c_{l+1}^{n,i}, \quad (2)$$

where $\hat{c}_l^{n(i)}$ is the i th dimension of \hat{c}_l^n . This equation means that a parent node code acts as a child node selector controlling the ON or OFF of a child node code.

3. By executing Step 2 recursively from the highest layer to the lowest layer, we can sample \hat{c}_L with L layer hierarchical inclusion constraints. We add it to the generator input and use it with z to generate an image: $\mathbf{x} = G(\hat{c}_L, z)$.

4.2. HCMI

The DTLC imposes a hierarchical inclusion structure on latent variables; however, its constraints are not sufficient to correlate latent variables with semantic features of images. To solve this problem without relying on detailed supervision, we propose a hierarchical conditional mutual information regularization (HCMI), which is an extension of MI [4] and conditional MI (CMI) [15] to hierarchical conditional settings. In particular, we use different types of regularization for the second layer to the L th layer, which have parent node codes, and the first layer, which does not have those.

Regularization for Second Layer to L th Layer: In this case, we need to discover semantic features in a hierarchically restricted manner; therefore, we maximize mutual information between l th-layer child node code \mathbf{c} and image $G(\hat{c}_L, z)$ conditioned on $(l-1)$ th-layer parent node code \mathbf{p} : $I(\mathbf{c}; G(\hat{c}_L, z) | \mathbf{p})$. For simplicity, we denote $\hat{c}_l^{n,i}$ and \hat{c}_{l-1}^n

as \mathbf{c} and \mathbf{p} , respectively. In practice, exact calculation of this mutual information is difficult because it requires calculation of the intractable posterior $P(\mathbf{c} | \mathbf{x}, \mathbf{p})$. Therefore, following previous studies [4, 15], we instead calculate its lower bound using an auxiliary distribution $Q(\mathbf{c} | \mathbf{x}, \mathbf{p})$ approximating $P(\mathbf{c} | \mathbf{x}, \mathbf{p})$:

$$\begin{aligned} I(\mathbf{c}; G(\hat{c}_L, z) | \mathbf{p}) &= H(\mathbf{c} | \mathbf{p}) - H(\mathbf{c} | G(\hat{c}_L, z), \mathbf{p}) \\ &= H(\mathbf{c} | \mathbf{p}) + \mathbb{E}_{\mathbf{x} \sim G(\hat{c}_L, z)} [\mathbb{E}_{\mathbf{c}' \sim P(\mathbf{c} | \mathbf{x}, \mathbf{p})} [\log P(\mathbf{c}' | \mathbf{x}, \mathbf{p})]] \\ &= H(\mathbf{c} | \mathbf{p}) + \mathbb{E}_{\mathbf{x} \sim G(\hat{c}_L, z)} [D_{\text{KL}}(P(\cdot | \mathbf{x}, \mathbf{p}) || Q(\cdot | \mathbf{x}, \mathbf{p})) \\ &\quad + \mathbb{E}_{\mathbf{c}' \sim P(\mathbf{c} | \mathbf{x}, \mathbf{p})} [\log Q(\mathbf{c}' | \mathbf{x}, \mathbf{p})]] \\ &\geq H(\mathbf{c} | \mathbf{p}) + \mathbb{E}_{\mathbf{x} \sim G(\hat{c}_L, z)} [\mathbb{E}_{\mathbf{c}' \sim P(\mathbf{c} | \mathbf{x}, \mathbf{p})} [\log Q(\mathbf{c}' | \mathbf{x}, \mathbf{p})]] \\ &= H(\mathbf{c} | \mathbf{p}) + \mathbb{E}_{\mathbf{c} \sim P(\mathbf{c} | \mathbf{p}), \mathbf{x} \sim G(\hat{c}_L, z)} [\log Q(\mathbf{c} | \mathbf{x}, \mathbf{p})]. \quad (3) \end{aligned}$$

For simplicity, we fix the distribution of \mathbf{c} and treat $H(\mathbf{c} | \mathbf{p})$ as constant. In practice, Q is parametrized as a neural network and we particularly denote the network for $\hat{c}_l^{n,i}$ ($= \hat{c}_l^{n,i}$) as Q_l^m , where $m = k_{l-1}(n-1) + i$. Thus, the final objective function is written as

$$\begin{aligned} \mathcal{L}_{\text{HCMI}}(G, Q_l^m) &= \mathbb{E}_{\mathbf{c} \sim P(\mathbf{c} | \mathbf{p}), \mathbf{x} \sim G(\hat{c}_L, z)} [\log Q_l^m(\mathbf{c} | \mathbf{x}, \mathbf{p})]. \quad (4) \end{aligned}$$

The Q_l^m attempts to discover the specific semantic features that correlate with \mathbf{c} in terms of conditional information gain by maximizing this objective. We calculate $\mathcal{L}_{\text{HCMI}}(G, Q_l^m)$ for every child node code $\hat{c}_l^{n,i}$. We denote the summation in the l th layer as $\mathcal{L}_{\text{HCMI}}(G, Q_l) = \sum_{m=1}^{N_l} \mathcal{L}_{\text{HCMI}}(G, Q_l^m)$. We use this objective with trade-off parameter λ_l .

Regularization for First Layer: The above regularization is useful for the codes that have parent node codes; however, the first-layer codes do not have those; thus, we instead use a different regularization for them. Fortunately, this single-layer case has been addressed in previous studies [4, 29] and we use one of them depending on the supervision setting. In an unsupervised setting, we use the MI [4] written as

$$\mathcal{L}_{\text{MI}}(G, Q_1) = \mathbb{E}_{c_1 \sim P(c_1), \mathbf{x} \sim G(\hat{c}_L, z)} [\log Q_1(c_1 | \mathbf{x})]. \quad (5)$$

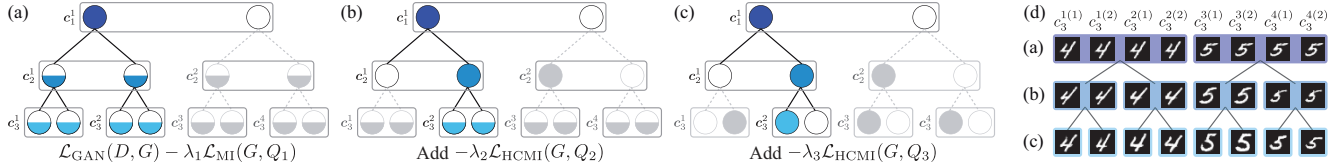


Figure 4. **Example of curriculum learning:** (a) We first learn disentangled representations in first layer. To do this, we only use regularization for this layer and fix and set average value to lower layer codes. (b)(c) We then learn disentangled representations in second and third layers in layer-by-layer manner. We add regularization and sampling in turn depending on training phase. (d) Image samples generated in each phase. In phase (a), first-layer codes are learned, while second- and third-layer codes are fixed; therefore, 2 disentangled representations are obtained. In phase (b), first- and second-layer codes are learned, while third-layer codes are fixed; therefore, 2×2 disentangled representations are obtained. In phase (c), all codes are learned; therefore, $2 \times 2 \times 2$ disentangled representations are obtained.

In a weakly supervised setting, we use an auxiliary classifier regularization (AC) [29] written as

$$\mathcal{L}_{AC}(G, Q_1) = \mathbb{E}_{c_1 \sim P(c_1), \mathbf{x} \sim G(\hat{c}_L, \mathbf{z})} [\log Q_1(c_1 | \mathbf{x})] + \mathbb{E}_{c_1, \mathbf{x} \sim P_{\text{data}}(c_1, \mathbf{x})} [\log Q_1(c_1 | \mathbf{x})]. \quad (6)$$

Note that the first term is the same as $\mathcal{L}_{MI}(G, Q_1)$, and the added second term acts as supervision regularization. We use these objectives with trade-off parameter λ_1 .

Full Objective: Our full objective is written as

$$\begin{aligned} & \mathcal{L}_{\text{Full}}(D, G, Q_1, \dots, Q_L) \\ = & \mathcal{L}_{GAN}(D, G) - \lambda_1 \mathcal{L}_{MI/AC}(G, Q_1) - \sum_{l=2}^L \lambda_l \mathcal{L}_{HCMI}(G, Q_l). \end{aligned} \quad (7)$$

This is minimized for the G and Q_1, \dots, Q_L and maximized for the D .

4.3. Curriculum Learning

The HCMI works well when the higher layer codes are already known; however, we assume the condition in which detailed annotations are not provided in advance. As a result, the network may confuse between inner-layer and intra-layer disentanglement at the beginning of training. To mitigate this problem, we developed our curriculum learning method. In particular, we define a curriculum for regularization and sampling. We illustrate an example of the proposed curriculum learning method in Figure 4.

Curriculum for Regularization: As a curriculum for regularization, we do not use the whole regularization in Equation 7 at the same time, instead, we add the regularization from the highest layer to the lowest layer in turn according to the training phase. In an unsupervised setting, we first learn with $\mathcal{L}_{GAN}(D, G) - \lambda_1 \mathcal{L}_{MI}(G, Q_1)$ then add $-\lambda_2 \mathcal{L}_{HCMI}(G, Q_2), \dots, -\lambda_L \mathcal{L}_{HCMI}(G, Q_L)$ in turn. In a weakly supervised setting, we first learn with $\mathcal{L}_{GAN}(D, G) - \lambda_1 \mathcal{L}_{AC}(G, Q_1) - \lambda_2 \mathcal{L}_{HCMI}(G, Q_2)$ then add $-\lambda_3 \mathcal{L}_{HCMI}(G, Q_3), \dots, -\lambda_L \mathcal{L}_{HCMI}(G, Q_L)$ in turn. We use different curricula between these two settings because in a weakly supervised setting, we already know the first-layer codes; thus, we can start from learning the second-layer codes.

# of Hidden Layers	Unsupervised	(Weakly) Supervised
0	GAN [10]	CGAN [28] ¹ , AC-GAN [29]
1	InfoGAN [4]	CFGAN [15] ¹
2		
3		
4, ...		DTLC-GAN

Table 1. Relationship to previous GANs

Curriculum for Sampling: In learning the higher layer codes, instability caused by random sampling of the lower layer codes can degrade the learning performance. Motivated by this fact, we define a curriculum for sampling. In particular, in learning the higher layer codes \hat{c}_l , we fix and set the average value to the lower layer codes c_{l+1}, \dots, c_L , e.g., set $\frac{1}{k_{l+1}}$ for discrete code $c_{l+1}^{n(i)}$ and set 0 for continuous code $c_{l+1}^{n(i)}$.

4.4. Relationship to Previous GANs

The DTLC-GAN is a general framework, and we can see it as a natural extension of previous GANs. We summarize this relationship in Table 1. In particular, the InfoGAN [4] and CFGAN [15]¹ are highly related to the DTLC-GAN in terms of discovering hidden representations on the basis of information gain; however, they are limited to learning one-layer hidden representation. We developed our DTLC-GAN, HCMI, and curriculum learning method to overcome this limitation.

5. Implementation

We designed the network architectures and training scheme on the basis of techniques introduced for the InfoGAN [4]. The D and Q_1, \dots, Q_L share all convolutional layers, and one fully connected layer is added to the final layer for Q_l . This means that the difference in the calculation cost for the GAN and DTLC-GAN is negligibly small. For discrete code \hat{c}_l^m , we represent Q_l^m as softmax nonlinearity. For continuous code \hat{c}_l^m , we treat Q_l^m as a factored Gaussian.

In most of the experiments we conducted, we used typical DCGAN models [32] and did not use the state-of-the-art

¹Strictly speaking, the CFGAN is formulated as an extension of the CGAN, while the weakly supervised DTLC-GAN is formulated as an extension of the AC-GAN. Therefore, these two models do not have completely the same architecture; however, they share the similar motivation.

GAN training techniques to evaluate whether the DTLC-GAN works well without relying on such techniques. However, our contributions are orthogonal to these techniques; therefore, we can improve image quality easily by incorporating these techniques to our DTLC-GAN. To demonstrate this, we also tested the DTLC-WGAN-GP (our DTLC-GAN with the WGAN-GP ResNet [12]) as discussed in Section 6.3. The details of the experimental setup are given in Section B in the supplementary material.

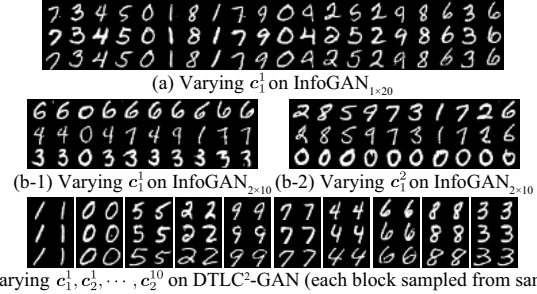
6. Experiments

We conducted experiments on various datasets, i.e., MNIST [22], CIFAR-10 [18], Tiny ImageNet [38], 3D Faces [31], and CelebA [23], to evaluate the effectiveness and generality of our DTLC-GAN². We first used the MNIST and CIFAR-10 datasets, which are widely used in this field, to analyze the DTLC-GAN qualitatively and quantitatively. In particular, we evaluated the DTLC-GAN in an unsupervised setting on the MNIST dataset and in a weakly supervised setting on the CIFAR-10 dataset (Section 6.1 and 6.2, respectively). We tested the DTLC-WGAN-GP on the CIFAR-10 and Tiny ImageNet datasets to demonstrate that our contributions are orthogonal to the state-of-the-art GAN training techniques (Section 6.3). We used the 3D Faces dataset to evaluate the effectiveness of the DTLC-GAN with continuous codes (Section 6.4) and evaluated it on image-retrieval tasks using the CelebA dataset (Section 6.5). Hereafter, we denote the DTLC-GAN with an L th layer DTLC as the **DTLC^L-GAN** and **DTLC^L-GAN** in a weakly supervised setting as the **DTLC^L-GAN_{WS}**.

6.1. Unsupervised Representation Learning

We first analyzed the DTLC-GAN in unsupervised settings on the MNIST dataset, which consists of photographs of handwritten digits and contains 60,000 training and 10,000 test samples.

Representation Comparison: To confirm the effectiveness of the hierarchical representation learning, we compared the DTLC-GAN with that in which dimensions of latent codes given to the G are the same but are not hierarchical. To represent our DTLC-GAN, we used the **DTLC²-GAN**, where $k_1 = 10$ and $k_2 = 2$. In this model, \hat{c}_2 , the dimension of which is $10 \times 2 = 20$, is given to the G . For comparison, we used two models in which latent code dimensions are also 20 but not hierarchical. One is the **InfoGAN_{1×20}**, which has one code $c_1^1 \sim \text{Cat}(K = 20, p = 0.05)$, and the other is the **InfoGAN_{2×10}**, which has two codes $c_1^1, c_1^2 \sim \text{Cat}(K = 10, p = 0.1)$. We show the results in Figure 5. In (c), the **DTLC²-GAN** succeeded in learning hierarchically interpretable representations (in the first layer, digits, and in the second layer, details of each digit). In (a), the **InfoGAN_{1×20}** succeeded in learning disentangled representations; however, they were learned as



(c) Varying $c_1^1, c_2^1, \dots, c_2^{10}$ on DTLC²-GAN (each block sampled from same c_1^1)

Figure 5. **Representation comparison on MNIST:** We compared models in which dimensions of latent codes given to G are same 20. In each figure, column contains three samples from same category. In each row, one latent code is varied, while other latent codes and noise are fixed.

a flat relationship; thus, it was not trivial to estimate the higher concept (e.g., digits) from them. In (b-1) and (b-2), the **InfoGAN_{2×10}** failed to learn interpretable representations. We argue that this is because c_1^1 and c_1^2 struggle to represent digit types. To clarify this limitation, we also conducted experiments on simulated data. See Section A.1 in the supplementary material for details.

Ablation Study on Curriculum Learning: To analyze the effectiveness of the proposed curriculum learning method, we conducted an ablation study. To evaluate quantitatively, we measured the inter-category diversity of generated images on the basis of structural similarity (SSIM) [47], which is a well-characterized perceptual similarity metric. This is an ad-hoc measure; however, recent studies [15, 29] showed that an SSIM-based measure is useful for evaluating the diversity of images generated with a GAN. Note that evaluating the quality of deep generative models is not trivial and is still an open issue due to the variety of probabilistic criteria [41]. To evaluate the l th layer inter-category diversity, we measured the SSIM scores between pairs of images that are sampled from the same noise and higher layer codes but random l th and lower layer codes. We calculated the scores for 50,000 randomly sampled pairs of images and took the average. The smaller value indicates that diversity is larger. We show changes in the mean SSIM scores through learning and sample images generated with varying latent codes per layer in Figure 6. We used the **DTLC⁴-GAN**, where $k_1 = 10$ and $k_2, k_3, k_4 = 2$. From these results, the **DTLC⁴-GAN** with the full curriculum succeeded in making higher layer codes obtain higher diversity and lower layer codes obtain lower diversity, while the others failed. We argue that this is because the latter cannot avoid confusion between inner-layer and intra-layer disentanglement. The qualitative results also support this fact. We also show sample images for all categories in Figures 13–15 in the supplementary material.

6.2. Weakly Supervised Representation Learning

We next analyzed the DTLC-GAN in weakly supervised settings (i.e., only class labels are supervised) on the

²Due to the limited space, we provide only the important results in this paper. Please refer to the supplementary material for more results.

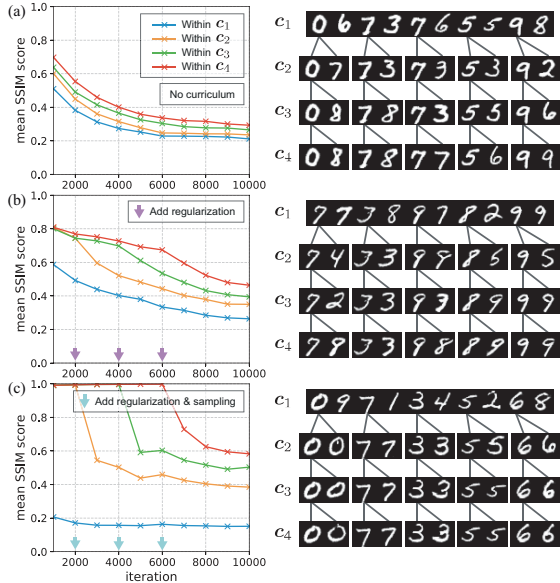


Figure 6. **Ablation study in unsupervised settings on MNIST:** Left figures show changes in mean SSIM scores through learning. We measured those between pairs of images within same category per layer. Right figures show sample images generated with varying latent codes per layer. Gray line indicates parent-child relationship. From top to bottom, (a) DTLC⁴-GAN without curriculum, (b) DTLC⁴-GAN with curriculum for regularization, and (c) DTLC⁴-GAN with full curriculum (curriculum for regularization and sampling: proposed curriculum learning method).

CIFAR-10 dataset, which consists of 10 classes of images and contains 5,000 training and 1,000 test samples per class.

Ablation Study on Curriculum Learning: We conducted an ablation study to evaluate the effectiveness of the proposed curriculum learning method in weakly supervised settings. We show changes in mean SSIM scores through learning and sample images generated with varying latent codes per layer in Figure 7. In this experiment, we used the DTLC⁴-GAN_{WS}, where $k_1 = 10$ and $k_2, k_3, k_4 = 3$. We can see the same tendency as in Figure 6. These results indicate that proposed curriculum learning method is indispensable, even in weakly supervised settings. We show samples images for all categories in Figures 16–18 in the supplementary material. We also conducted preference tests to analyze visual interpretability. See Section A.2 in the supplementary material for details.

Quantitative Evaluation: An important concern is whether our extension degrades image quality. To address this concern, we evaluated the DTLC-GAN_{WS} on three metrics: inception score [39], adversarial accuracy [49], and adversarial divergence [49]³. We list the results in Table 2. We compared a GAN, the AC-GAN, and DTLC^L-GAN_{WS}, where $k_1 = 10$ and $k_2, \dots, k_L = 3$. For

³The latter two metrics require pairs of generated images and class labels to train a classifier. In our settings, a conditional generator is learned; thus, we directly used it to generate an image with a class label. We used the classifier, architecture of which was similar to the D except for the output layer.

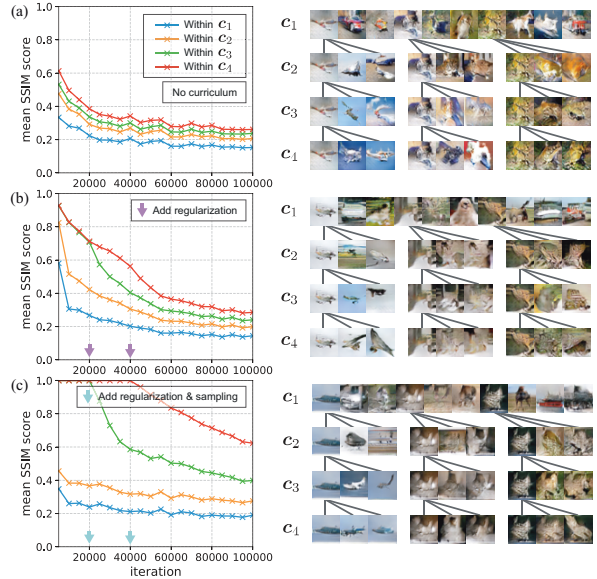


Figure 7. **Ablation study in weakly supervised settings on CIFAR-10:** View of figure is similar to that in Figure 6

Model	Inception Score	Adversarial Accuracy	Adversarial Divergence
GAN	7.09 ± 0.09	-	-
AC-GAN	7.41 ± 0.06	50.99 ± 0.55	2.07 ± 0.02
DTLC ² -GAN _{WS}	7.39 ± 0.03	55.10 ± 0.48	1.82 ± 0.03
DTLC ³ -GAN _{WS}	7.35 ± 0.09	55.20 ± 0.47	1.95 ± 0.05
DTLC ⁴ -GAN _{WS}	7.46 ± 0.06	56.19 ± 0.36	1.93 ± 0.05
DTLC ⁵ -GAN _{WS}	7.51 ± 0.06	58.87 ± 0.52	1.83 ± 0.04
Real Images	11.24 ± 0.12	85.77 ± 0.22	0
State-of-the-Art	8.59 ± 0.12 [13]	44.22 ± 0.08 [49]	5.57 ± 0.06 [49]

Table 2. Quantitative comparison between GAN, AC-GAN, and DTLC-GAN_{WS}

fair comparison, we used the same network architecture and training scheme except for the extended parts. The inception scores are not state-of-the-art, but in this comparison, the DTLC^L-GAN_{WS} improved upon GAN and was comparable to the AC-GAN. The adversarial accuracy and adversarial divergence scores are state-of-the-art, and the DTLC^L-GAN_{WS} improved upon the AC-GAN. These results are noteworthy because they indicate that we can obtain expressive representation using the DTLC-GAN_{WS} without concern for image-quality degradation.

6.3. Combination with WGAN-GP

Another concern is whether our contributions are orthogonal to the state-of-the-art GAN training techniques. To demonstrate this, we tested the DTLC^L-WGAN-GP on three cases: CIFAR-10 (unsupervised/weakly supervised) and Tiny ImageNet⁴ (unsupervised). The number of categories was same as that with the models used in Table 2. We list the results in Table 3. Interestingly, in all cases, the scores improved as the layers became deeper, and the DTLC⁴-WGAN-GPs achieved state-of-the-art performance. We show generated image samples in Figures 20–22 in the supplementary material.

⁴Tiny version of the ImageNet dataset containing 200 classes × 500 images. To shorten the training time, we resized images to 32 × 32.




Figure 8. **Representation comparison of models that have continuous codes:** Each sample is generated from same noise but different continuous codes (varied from left to right). In InfoGANs (a, b), each code is independent and exclusive, while in DTLC-GAN (c), lower layer codes learn category-specific (in this case, pose-specific) semantic features conditioned on higher layer codes.




Figure 9. (a–c) Example results of hierarchically interpretable image retrieval. (d) Attribute-specific areas used for evaluation in Table 4.

Model	CIFAR-10 (Unsupervised)	CIFAR-10 (Supervised)	Tiny ImageNet (Unsupervised)
WGAN-GP	7.86 ± .07 [12]	-	8.33 ± .11
AC/Info-WGAN-GP	7.97 ± .09	8.42 ± .10 [12]	8.33 ± .10
DTLC ² -WGAN-GP	8.03 ± .12	8.44 ± .10	8.34 ± .08
DTLC ³ -WGAN-GP	8.15 ± .08	8.56 ± .07	8.41 ± .10
DTLC ⁴ -WGAN-GP	8.22 ± .11	8.80 ± .08	8.51 ± .08
State-of-the-Art	7.86 ± .07 [12]	8.59 ± .12 [13]	-

Table 3. Inception scores for WGAN-GP-based models

6.4. Extension to Continuous Codes

To analyze the DTLC-GAN with continuous codes, we evaluated it on the 3D Faces dataset, which consists of faces generated from a 3D face model and contains 240,000 samples. We compared three models, the **InfoGAN_{C5}**, which is the InfoGAN with five continuous codes $c_1^1, \dots, c_1^5 \sim \text{Unif}(-1, 1)$ (used in the InfoGAN study [4]), **InfoGAN_{C1D1}**, which is the InfoGAN with one categorical code $c_1^1 \sim \text{Cat}(K = 5, p = 0.2)$ and one continuous code $c_2^2 \sim \text{Unif}(-1, 1)$, and **DTLC²-GAN**, which has one categorical code $c_1^1 \sim \text{Cat}(K = 5, p = 0.2)$ in the first layer and five continuous codes $c_2^2, \dots, c_2^5 \sim \text{Unif}(-1, 1)$ in the second layer. We show example results in Figure 8. In the **InfoGANs** (a, b), the individual codes tend to represent independent and exclusive semantic features because they have a flat relationship, while in the **DTLC²-GAN** (c), we can learn category-specific (in this case, pose-specific) semantic features conditioned on the higher layer codes.

6.5. Application to Image Retrieval

One possible application of the DTLC-GAN is to use hierarchically interpretable representations for image retrieval. To confirm this, we used the CelebA dataset, which consists of photographs of faces and contains 180,000 training and 20,000 test samples. To search for an image hierarchically, we measure the L2 distance between query

Code	Bangs	Glasses	Smiling
c_1	0.150	0.189	0.274
\hat{c}_2	0.194	0.256	0.294
\hat{c}_3	0.211	0.265	0.326

Table 4. Attribute-specific SSIM scores for different codes

and database images on the basis of $c_1, \hat{c}_2, \dots, \hat{c}_L$, which are predicted using auxiliary functions Q_1, \dots, Q_L . Figure 9 shows the results of bangs-based, glasses-based, and smiling-based image retrieval. For evaluation, we used the test set in the CelebA dataset. We trained **DTLC³-GAN_{WS}**, where $k_1 = 2, k_2 = 3$, and $k_3 = 3$, particularly where hierarchical representations are learned only for the attribute presence state. These results indicate that as the layer becomes deeper, images in which attribute details match more can be retrieved. To evaluate quantitatively, we measured the SSIM score between query and database images for the attribute-specific areas [15] defined in Figure 9(d). We summarize the scores in Table 4. These results indicate that as the layer becomes deeper, the concordance rate of attribute-specific areas increases.

7. Discussion and Conclusions

We proposed an extension of the GAN called the DTLC-GAN to learn hierarchically interpretable representations. To develop it, we introduced the DTLC to impose a hierarchical inclusion structure on latent variables and proposed the HCMI and curriculum learning method to discover the salient semantic features in a layer-by-layer manner by only using a single DTLC-GAN model without relying on detailed supervision. Experiments showed promising results, indicating that the DTLC-GAN is well suited for learning hierarchically interpretable representations. The DTLC-GAN is a general model, and possible future work includes applying it to other models, such as encoder-decoder models [6, 8, 17, 21, 37], and using it as a latent hierarchical structure discovery tool for high-dimensional data.

References

- [1] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. In *ICLR*, 2017. 3
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017. 3
- [3] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *ICML*, 2009. 2
- [4] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*, 2016. 2, 3, 4, 5, 8
- [5] E. Denton, S. Chintala, A. Szlam, and R. Fergus. Deep generative image models using a Laplacian pyramid of adversarial networks. In *NIPS*, 2015. 3
- [6] J. Donahue, P. Krähenbühl, and T. Darrell. Adversarial feature learning. In *ICLR*, 2017. 8
- [7] A. Dosovitskiy, J. T. Springenberg, and T. Brox. Learning to generate chairs with convolutional neural networks. In *CVPR*, 2015. 2
- [8] V. Dumoulin, I. Belghazi, B. Poole, O. Mastropietro, A. Lamb, M. Arjovsky, and A. Courville. Adversarially learned inference. In *ICLR*, 2017. 8
- [9] S. M. A. Eslami, N. Heess, T. Weber, Y. Tassa, D. Szepesvari, and G. E. Hinton. Attend, infer, repeat: Fast scene understanding with generative models. In *NIPS*, 2016. 3
- [10] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 1, 3, 5
- [11] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. DRAW: A recurrent neural network for image generation. In *ICML*, 2015. 3
- [12] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. Improved training of Wasserstein GANs. In *NIPS*, 2017. 3, 6, 8
- [13] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, and S. Belongie. Stacked generative adversarial networks. In *CVPR*, 2017. 3, 7, 8
- [14] D. J. Im, C. D. Kim, H. Jiang, and R. Memisevic. Generating images with recurrent adversarial networks. In *ICLR Workshop*, 2016. 3
- [15] T. Kaneko, K. Hiramatsu, and K. Kashino. Generative attribute controller with conditional filtered generative adversarial networks. In *CVPR*, 2017. 2, 3, 4, 5, 6, 8
- [16] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling. Semi-supervised learning with deep generative models. In *NIPS*, 2014. 3
- [17] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 1, 2, 8
- [18] A. Krizhevsky. Learning multiple layers of features from tiny images, 2009. 6
- [19] T. D. Kulkarni, W. F. Whitney, P. Kohli, and J. Tenenbaum. Deep convolutional inverse graphics network. In *NIPS*, 2015. 3
- [20] H. Kwak and B.-T. Zhang. Generating images part by part with composite generative adversarial networks. *arXiv preprint arXiv:1607.05387*, 2016. 3
- [21] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther. Autoencoding beyond pixels using a learned similarity metric. In *ICML*, 2016. 8
- [22] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proc. of the IEEE*, 86(11):2278–2324, 1998. 6
- [23] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 6
- [24] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey. Adversarial autoencoders. In *NIPS*, 2016. 3
- [25] E. Mansimov, E. Parisotto, J. L. Ba, and R. Salakhutdinov. Generating images from captions with attention. In *ICLR*, 2016. 3
- [26] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley. Least squares generative adversarial networks. In *ICCV*, 2017. 3
- [27] M. F. Mathieu, J. J. Zhao, J. Zhao, A. Ramesh, P. Sprechmann, and Y. LeCun. Disentangling factors of variation in deep representation using adversarial training. In *NIPS*, 2016. 3
- [28] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 3, 5
- [29] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier GANs. In *ICML*, 2017. 3, 4, 5, 6
- [30] D. Parikh and K. Grauman. Relative attributes. In *ICCV*, 2011. 2
- [31] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. A 3D face model for pose and illumination invariant face recognition. In *AVSS*, 2009. 6
- [32] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 3, 5
- [33] S. Reed, Z. Akata, S. Mohan, S. Tenka, B. Schiele, and H. Lee. Learning what and where to draw. In *NIPS*, 2016. 3
- [34] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. In *ICML*, 2016. 3
- [35] S. Reed, A. van den Oord, N. Kalchbrenner, V. Bapst, M. Botvinick, and N. de Freitas. Generating interpretable images with controllable structure. In *ICLR Workshop*, 2017. 3
- [36] S. Reed, Y. Zhang, Y. Zhang, and H. Lee. Deep visual analogy-making. In *NIPS*, 2015. 2
- [37] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014. 1, 2, 8
- [38] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 6
- [39] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. In *NIPS*, 2016. 3, 7
- [40] J. T. Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. In *ICLR*, 2016. 3
- [41] L. Theis, A. van den Oord, and M. Bethge. A note on the evaluation of generative models. In *ICLR*, 2016. 6
- [42] L. Tran, X. Yin, and X. Liu. Disentangled representation learning GAN for pose-invariant face recognition. In *CVPR*, 2017. 3

- [43] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016. 2
- [44] A. van den Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu. Conditional image generation with pixelCNN decoders. In *NIPS*, 2016. 3
- [45] C. Vondrick, H. Pirsaviash, and A. Torralba. Generating videos with scene dynamics. In *NIPS*, 2016. 3
- [46] X. Wang and A. Gupta. Generative image modeling using style and structure adversarial networks. In *ECCV*, 2016. 3
- [47] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. on IP*, 13(4):600–612, 2004. 6
- [48] X. Yan, J. Yang, K. Sohn, and H. Lee. Attribute2image: Conditional image generation from visual attributes. In *ECCV*, 2016. 3
- [49] J. Yang, A. Kannan, D. Batra, and D. Parikh. LR-GAN: Layered recursive generative adversarial networks for image generation. In *ICLR*, 2017. 3, 7
- [50] J. Yang, S. Reed, M.-H. Yang, and H. Lee. Weakly-supervised disentangling with recurrent transformations for 3D view synthesis. In *NIPS*, 2015. 2
- [51] A. Yu and K. Grauman. Just noticeable differences in visual attributes. In *ICCV*, 2015. 2
- [52] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. N. Metaxas. StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017. 3
- [53] Z. Zhang, Y. Song, and H. Qi. Age progression/regression by conditional adversarial autoencoder. In *CVPR*, 2017. 3
- [54] J. Zhao, M. Mathieu, and Y. LeCun. Energy-based generative adversarial network. In *ICLR*, 2017. 3
- [55] S. Zhao, J. Song, and S. Ermon. Learning hierarchical features from generative models. In *ICML*, 2017. 3