

This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Cross-Domain Latent Modulation for Variational Transfer Learning

Jinyong Hou, Jeremiah D. Deng, Stephen Cranefield, Xuejie Ding Department of Information Science, University of Otago

robert.hou@postgrad.otago.ac.nz, {jeremiah.deng,stephen.cranefield,emily.ding}@otago.ac.nz

Abstract

We propose a cross-domain latent modulation mechanism within a variational autoencoders (VAE) framework to enable improved transfer learning. Our key idea is to procure deep representations from one data domain and use it as perturbation to the reparameterization of the latent variable in another domain. Specifically, deep representations of the source and target domains are first extracted by a unified inference model and aligned by employing gradient reversal. Second, the learned deep representations are cross-modulated to the latent encoding of the alternate domain. The consistency between the reconstruction from the modulated latent encoding and the generation using deep representation samples is then enforced in order to produce inter-class alignment in the latent space. We apply the proposed model to a number of transfer learning tasks including unsupervised domain adaptation and image-toimage translation. Experimental results show that our model gives competitive performance.

1. Introduction

In machine learning, one can rarely directly apply a pretrained model to a new dataset or a new task, as the performance of a learned model often plunges significantly for the new data which may have significant sampling bias or even belong to different distributions. Transfer learning can help us utilize the learned knowledge from a previous domain (the 'source') to improve performance on a related domain or task (the 'target') [31, 35, 41].

From the perspective of probabilistic modeling [23, 39, 34], the key challenge in achieving cross-domain transfer is to learn a joint distribution of data from different domains. Once the joint distribution is learned, it can be used to generate the marginal distribution of the individual domains [14, 23]. Under the variational inference scenario, an inferred joint distribution is often applied to the latent space. Due to the coupling theory, inferring the joint distribution from the marginal distributions of different domains

is a highly ill-posed problem [21]. To address this problem, UNIT [23] makes an assumption that there is a shared latent space for the two domains. Usually, this can be achieved by applying the adversarial strategy to the domains' latent spaces. Another line of research focuses on the use of a complex prior to improve the representation performance for the input data [28, 36, 12]. However, the previous works neglect the role of the generation process for the latent space which could be helpful for cross-domain transfer scenarios.

In this paper, we propose a novel latent space reparameterization method, and employ a generative process to cater for the cross-domain transferability. Specifically, we incorporate a cross-domain component into the reparameterization transformation, which builds the connection between the variational representations and domain features in a cross-domain manner. The generated transfer latent space is further tuned by domain-level adversarial alignment and domain consistency between images obtained through reconstruction and generations. We apply our model to the homogeneous transfer scenarios, such as unsupervised domain adaptation and image-to-image translation. The experimental results show the efficiency of our model.

The rest of the paper is organized as follows. In Section 2, some related work is briefly reviewed. In Section 3, we outline the overall structure of our proposed model and develop the learning metrics with defined losses. The experiments are presented and discussed in Section 4. We conclude our work in Section 5, indicating our plan of future work.

2. Related Work

Latent space manipulation: As discussed above, for a joint distribution, manipulation of the latent space is common [17, 23, 22] for the cross-domain adaptation situations. One approach focuses on a shared latent space, where the latent encodings are regarded as common representations for inputs across domains. Some adversarial strategy is usually used to pool them together so that the representations are less domain-dependent. For the variational approach, works in [28, 48, 15] adopt complex priors for multi-modal latent representations, while other works [23, 34, 22] still assume

a standard Gaussian prior. Another aproach is to use disentangled latent representations where the latent encoding is divided into some defined parts (e.g. style and content parts), then the model learns separated representations and swaps them for the transfer [9, 46, 6, 20]. Our method is different from these approaches. In our model, learned auxiliary deep representation is used to generate perturbations to the latent space through a modified reparameterization using variational information from the counterpart domain. It helps generate cross-domain image translation. The transfer is carried out by a reparameterization transformation, using statistical moments retaining specific information for one domain, and deep representation providing information from another domain.

Varied Homogeneous transfer tasks: The manipulation on the latent space is often interwoven with the homogeneous image transfer together, such as unsupervised domain adaptation and image translation [29, 30, 5]. In the domain separation networks [4], separate encoding modules are employed to extract the invariant representation and domainspecific representations from the domains respectively, with the domain-invariant representations being used for the domain adaptation. References [3, 33, 11] transfer the target images into source-like images for domain adaptation. References [23, 47, 13] map the inputs of different domains to a single shared latent space, but cycle consistency is required for the completeness of the latent space. The UFDN [22] utilizes a unified encoder to extract the multi-domain images to a shared latent space, and the latent domain-invariant coding is manipulated for image translation between different domains.

In contrast, we adopt the pixel-level adaptation between domains from the cross-domain generation, but the proposed model can also be used at the feature-level due to the latent space alignment. Our model also has a unified inference model, but the consistency is imposed in a straightforward way, with reduced computational complexity.

3. Proposed Model

3.1. Problem setting

Let $\mathcal{X} \subset \mathbb{R}^d$ be a *d*-dimensional data space, and $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\} \in \mathcal{X}$ the sample set with marginal distribution $p(\mathbf{X})$. The source domain is denoted by a tuple $(\mathcal{X}_s, p(\mathbf{X}_s))$, and the target domain by $(\mathcal{X}_t, p(\mathbf{X}_t))$. In our paper, we consider the homogeneous transfer with domain shift, i.e. $\mathcal{X}_s \approx \mathcal{X}_t$, but $p(\mathbf{X}_s) \neq p(\mathbf{X}_t)$. For the unsupervised pixel-level domain adaptation scenario, the label set is $\mathbf{Y} = \{y_1, y_2, \dots, y_n\} \in \mathcal{Y}$ (\mathcal{Y} is the label space), and a task $\mathcal{T} = p(\mathbf{Y} | \mathbf{X})$ is considered too. However, only the source domain's label set \mathbf{Y}_s is available during transfer learning.

3.2. Transfer Latent Space

As any given marginal distribution can be yielded by an infinite number of joint distributions, we need to build an inference framework with some constraints. Under the variational autoencoder (VAE) framework, the latent space is one of the manipulation targets. We propose the transfer latent space as follows.

Definition 1. Transfer Latent Space \ddot{Z} . Let $x_s \in X_s$, $x_t \in X_t$ be the domain samples. Let us have a map f that extracts domain information Ω and a feature representation h given an input x:

$$f: \boldsymbol{x} \longrightarrow (\boldsymbol{\Omega}, \boldsymbol{h}), \boldsymbol{x} \in \boldsymbol{X}_s \cup \boldsymbol{X}_t.$$

Suppose we construct a transfer map \mathcal{G} that generates a latent variable \ddot{z} from Ω and h with domain crossovers:

$$egin{aligned} \ddot{m{z}}_{st} &= \mathcal{G}(m{\Omega}_s,m{h}_t), \ \ddot{m{z}}_{ts} &= \mathcal{G}(m{\Omega}_t,m{h}_s). \end{aligned}$$

The joint space formed by \ddot{z}_{st} and \ddot{z}_{ts} samples is defined as a *transfer latent space*, denoted by \ddot{Z} .

The transfer latent space is intended to become a "mixer" for the two domains, as the resulted latent variables are under cross-domain influences. Hence the transfer latent space can be regarded as a generalization of the latent space.

3.3. Framework

Our framework is shown in Fig. 1. In our framework, we build the cross-domain generation by a unified inference model $\mathbf{E}_{\phi}(\cdot)$ (as an implementation of the map f) and a generative model for the desired domain $\mathbf{D}_{\theta}(\cdot)$, e.g., the source domain in our model. A discriminator Ξ is utilized for the adversarial training. We use the terms "inference model" and "encoder" for $\mathbf{E}_{\phi}(\cdot)$, and "generative model" and "decoder" for $\mathbf{D}_{\theta}(\cdot)$ interchangeably.

As discussed in section 3.2, under the variational framework, the domain information Ω (here we remove the domain subscript for simplicity) is usually the pair (μ, σ) . Let h' be the flattened activations of the last convolution layer in \mathbf{E}_u . Then, following the treatment in [16], μ and σ can be obtained by $\mu = \mathbf{W}_{\mu}h' + b_{\mu}$ and $\sigma = \mathbf{W}_{\sigma}h' + b_{\sigma}$, where $\mathbf{W}_{\mu}, \mathbf{W}_{\sigma}, b_{\mu}, b_{\sigma}$ are the weights and biases for μ and σ .

From our observations, both shallow (e.g. PCA features) and deep representations can be used to obtain domain information h, in our end-to-end model we use the latter. We choose the high-level activation of the last convolutional layer, i.e., $h = \text{sigmoid}(\mathbf{W}_h h' + b_h)$ as the deep representation [44], where \mathbf{W}_h , b_h are the weights and biases for the deep abstractions.

Having obtained the domain information Ω and deep representation h, a natural choice for the transfer map G



Figure 1: Architectural view of the proposed model. It encourages an image from target domain (blue hexagon) to be transformed to a corresponding image in the source domain (black hexagon). The transfer latent distributions $p(\ddot{z}_{ts}|x_t, h_s)$ and $p(\ddot{z}_{st}|x_s, h_t)$ are learned which are used to generate corresponding images by the desired decoder. The deep representations are integrated into the reparameterization transformation with standard Gaussian auxiliary noise. Blue lines are for the target domain and black ones are for the source domain.

is through reparameterization. Here we propose a modified reparameterization trick to give the sampling from the transfer latent space as follows:

$$\ddot{\boldsymbol{z}}_{st} = \mathcal{G}((\boldsymbol{\mu}_s, \boldsymbol{\sigma}_s), \boldsymbol{h}_t) = \boldsymbol{\mu}_s + \boldsymbol{\sigma}_s \odot (\gamma_1 \boldsymbol{h}_t + \gamma_2 \boldsymbol{\epsilon}), \quad (1)$$

and

$$\ddot{\boldsymbol{z}}_{ts} = \mathcal{G}((\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t), \boldsymbol{h}_s) = \boldsymbol{\mu}_t + \boldsymbol{\sigma}_t \odot (\gamma_1 \boldsymbol{h}_s + \gamma_2 \boldsymbol{\epsilon}), \quad (2)$$

where h_s (h_t) is the sample of the deep representation space \mathcal{H}_s (\mathcal{H}_t); μ_s and σ_s (μ_t and σ_t) are the mean and standard deviation of the approximate posterior for the source (target) domain; $\gamma_1, \gamma_2 > 0$ are trade-off hyperparameter to balance the deep feature modulation and the standard Gaussian noise ϵ ; and \odot stands for the element-wise product of vectors. Therefore, the auxiliary noise in VAE resampling is now a weighted sum of a deep representation from the other domain and Gaussian noise, different from the standard VAE framework. Because the modified reparameterization allows a domain's latent representation, we call our model "Cross-Domain Latent Modulation", or CDLM for short.

Now we have obtained the transfer encodings by a unified encoder. Following the probabilistic encoder analysis [16], a shared inference model confines the latent variables into the same latent space. But this cannot guarantee them to be aligned. To pull the domains close, an adversarial strategy [7, 37] should be used for the alignment. The gradient reversal layer [7] is used in our model, by which adversarial learning is introduced to learn transferable features that are robust to domain shift. The adversarial alignment between \mathcal{H}_s and \mathcal{H}_t is for domain-level.

Furthermore, for better interpretation of the modulated reparameterization, let $h_t \sim p(\mathcal{H}_t) = \mathcal{N}(\mu_{h_t}, \sigma_{h_t})$. Then

for \ddot{z}_{st} , we have $\ddot{z}_{st} \sim \mathcal{N}(\ddot{z}_{st}; \mu_{st}, \sigma_{st}^2 \mathbf{I})$. For the *i*-th element of the distribution moments are given as follows:

$$\mu_{st}^{i} = \mathbb{E}\{\mu_{s}^{i} + \sigma_{s}^{i}(\gamma_{1}h_{t}^{i} + \gamma_{2}\epsilon^{i})\}$$
$$= \mu_{s}^{i} + \gamma_{1}\sigma_{s}^{i}\mu_{h_{t}}^{i}.$$
(3)

$$\mathbf{Var}(\ddot{z}_{st}^{i}) = \mathbb{E}\{(\sigma_{s}^{i}(\gamma_{1}h_{t}^{i}+\gamma_{2}\epsilon^{i})-\gamma_{1}\sigma_{s}^{i}\mu_{h_{t}}^{i})^{2}\} \\ = (\sigma_{s}^{i})^{2}[\gamma_{1}^{2}(\sigma_{h_{t}}^{i})^{2}+\gamma_{2}^{2}].$$
(4)

Therefore, the μ_{st}^i and σ_{st}^i are

$$\mu_{st}^{i} = \mu_{s}^{i} + \gamma_{1}\sigma_{s}^{i}\mu_{h_{t}}^{i}$$

$$\sigma_{st}^{i} = \sigma_{s}^{i}\sqrt{\gamma_{1}^{2}(\sigma_{h_{t}}^{i})^{2} + \gamma_{2}^{2}}.$$
(5)

Here, it is reasonable to assume $\mu_{h_t} \approx \mu_t$ and $\sigma_{h_t} \approx \sigma_t$ when the training is finished. With a practical setting of $\gamma_1 \gg \gamma_2$, and in effect $\sigma_s^i = 1$, Eq. (5) can be further simplified to $\mu_{st}^i = \mu_s^i + \gamma_1 \mu_t^i$, and $\sigma_{st}^i = \gamma_1 \sigma_s^i \sigma_t^i$. Then we can see that μ_{st} can be regarded as a location shift of μ_s under the influence of μ_t , which helps reduce the domain gap; σ_{st} can be taken as a recoloring of σ_s under the influence from the target. The formulation of \ddot{z}_{ts} can be similarly interpreted. These modulated encodings are hence constructed in a cross-domain manner.

Next, we apply the consistency constraint to the transfer latent space with modulation for further inter-class alignment. It has been found that consistency constraints preserve class-discriminative information [34, 13, 8]. For our model, the consistency is applied to the reconstructions from modulated encodings and the corresponding generations from deep representations. Let $\mathbf{D}_{\theta}(\cdot)$ be the generative model for domain image generation from the transfer latent space. The consistency requirements are

$$\mathbf{D}_{\theta}(\ddot{\boldsymbol{z}}_{st}) = \mathbf{D}_{\theta}(\gamma_{1}\boldsymbol{h}_{s} + \gamma_{2}\boldsymbol{\epsilon})
 \mathbf{D}_{\theta}(\ddot{\boldsymbol{z}}_{ts}) = \mathbf{D}_{\theta}(\gamma_{1}\boldsymbol{h}_{t} + \gamma_{2}\boldsymbol{\epsilon}),$$
(6)

where $\mathbf{D}_{\theta}(\ddot{\mathbf{z}}_{st})$ is the reconstruction of the source $(\widehat{\mathbf{x}}_{st})$, $\mathbf{D}_{\theta}(\ddot{\mathbf{z}}_{ts})$ is for the target $(\widehat{\mathbf{x}}_{ts})$. $\mathbf{D}_{\theta}(\cdot)$ can also function as a generative model, generating $\widetilde{\mathbf{x}}_s = \mathbf{D}_{\theta}(\gamma_1 \mathbf{h}_s + \gamma_2 \epsilon)$ and $\widetilde{\mathbf{x}}_t = \mathbf{D}_{\theta}(\gamma_1 \mathbf{h}_t + \gamma_2 \epsilon)$ for the source and target domain respectively. Also, the consistencies can guide the encoder to learn the representations from both domains.

Finally, a desired marginalized decoder, e.g. the source decoder, is trained to map the target images to be source-like. We render the target's structured generation \tilde{x}_{ts} for the test mode. For this end, we do not need the source to be taken into account for the test. That means a test image from the target domain x_t^i first passes through the inference model and obtains its deep feature h_t^i . Then it is fed into the generation model to generate an image with source style but keep its own class. That is to make the marginal distribution $p(\tilde{x}_{ts}^i) \approx p(x_s^i)$, but keep its class y_t^i .

3.4. Learning

Our goal is to update the variational parameters to learn a joint distribution and the generation parameters for the desired marginal distribution. Since the latent variables are generated with inputs from both domains, we have a modified formulation adapted from the plain VAE:

$$\log p(\boldsymbol{x}_s, \boldsymbol{x}_t) - \mathrm{KL}(q_{\phi}(\boldsymbol{\ddot{z}} | \boldsymbol{x}_s, \boldsymbol{x}_t) \| p(\boldsymbol{\ddot{z}} | \boldsymbol{x}_s, \boldsymbol{x}_t)) = \mathbb{E}[\log p(\boldsymbol{x}_s, \boldsymbol{x}_t | \boldsymbol{\ddot{z}})] - \mathrm{KL}(q_{\phi}(\boldsymbol{\ddot{z}} | \boldsymbol{x}_s, \boldsymbol{x}_t) \| p(\boldsymbol{\ddot{z}})),$$
(7)

where $\mathrm{KL}(\cdot)$ is the Kullback-Leibler divergence, and the transfer latent variable \ddot{z} can be either \ddot{z}_{st} or \ddot{z}_{ts} . Minimizing $\mathrm{KL}(q_{\phi}(\ddot{z}|\boldsymbol{x}_{s},\boldsymbol{x}_{t})||p(\ddot{z}|\boldsymbol{x}_{s},\boldsymbol{x}_{t}))$ is equivalent to maximizing the variational evidence lower bound (ELBO) $\mathcal{L}(\boldsymbol{x}_{s},\boldsymbol{x}_{t},\boldsymbol{\theta},\boldsymbol{\phi})$:

$$\mathcal{L}(\boldsymbol{x}_{s}, \boldsymbol{x}_{t}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbb{E}_{q_{\phi}}[\log p_{\theta}(\boldsymbol{x}_{s}, \boldsymbol{x}_{t} | \boldsymbol{\ddot{z}})] - \mathrm{KL}(q_{\phi}(\boldsymbol{\ddot{z}} | \boldsymbol{x}_{s}, \boldsymbol{x}_{t}) \| p(\boldsymbol{\ddot{z}})),$$
(8)

where the first term corresponds to the reconstruction cost (\mathcal{L}_{Rec}) , and the second term is the K-L divergence between the learned latent probability and the prior (specified as $\mathcal{N}(0, I)$) (\mathcal{L}_{KL}) . Considering the reconstruction of \boldsymbol{x}_s , and the K-L divergence for both $\ddot{\boldsymbol{z}}_{st}$ and $\ddot{\boldsymbol{z}}_{ts}$, we have

$$\mathcal{L}(\boldsymbol{x}_{s}, \boldsymbol{x}_{t}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbb{E}_{\boldsymbol{\ddot{z}} \sim q(\boldsymbol{\ddot{z}}_{st} | \boldsymbol{x}_{s}, \boldsymbol{x}_{t})} [\log p_{\boldsymbol{\theta}}(\boldsymbol{x}_{s} | \boldsymbol{\ddot{z}}_{st})] - \mathrm{KL}(\log q_{\boldsymbol{\phi}}(\boldsymbol{\ddot{z}}_{st} | \boldsymbol{x}_{s}, \boldsymbol{x}_{t}) \| p(\boldsymbol{z}))$$
(9)
- \mathrm{KL}(\log q_{\boldsymbol{\phi}}(\boldsymbol{\ddot{z}}_{ts} | \boldsymbol{x}_{t}, \boldsymbol{x}_{s}) \| p(\boldsymbol{z})).

To align the deep representations of the source and target domains, an adversarial strategy is employed to regularize the model. The loss function is given by

$$\mathcal{L}_{adv} = \mathbb{E}_{\boldsymbol{h}_s \sim p(\boldsymbol{h}_s | \boldsymbol{x}_s)} [\log \Xi(\boldsymbol{h}_s)] + \mathbb{E}_{\boldsymbol{h}_t \sim p(\boldsymbol{h}_t | \boldsymbol{x}_t)} [\log(1 - \Xi(\boldsymbol{h}_t))],$$
(10)

where $\Xi(\cdot)$ is the discriminator to predict from which domain the deep representation feature is.

From the analysis in Section 3.2, we can introduce a pairwise consistency between the reconstruction and the generation for the source and the target in an unsupervised manner respectively. The consistencies regularization improve the inter-class alignment. For the consistency loss \mathcal{L}_c , both the l^1 and l^2 -norm penalty can be used to regularize the decoder. Here we simply use MSE. Let \mathcal{L}_c^s and \mathcal{L}_c^t be the consistency for the domains respectively. \mathcal{L}_c is given as a combination of these two components, weighted by two coefficients β_1 and β_2 , respectively:

$$\begin{aligned} \mathcal{L}_{c} &= \beta_{1} \mathcal{L}_{c}^{s} + \beta_{2} \mathcal{L}_{c}^{t} \\ &= \beta_{1} (\mathbb{E}_{\ddot{\boldsymbol{z}} \sim q(\ddot{\boldsymbol{z}}|\boldsymbol{x}_{s}, \boldsymbol{x}_{t})} [\log p(\widehat{\boldsymbol{x}}_{s} | \boldsymbol{\ddot{z}}_{st})] - \mathbb{E}_{\boldsymbol{h}_{s} \sim p(\boldsymbol{h}_{s}|\boldsymbol{x}_{s})} [\log p(\widetilde{\boldsymbol{x}}_{s} | \boldsymbol{h}_{s})])^{2} \\ &+ \beta_{2} (\mathbb{E}_{\ddot{\boldsymbol{z}} \sim q(\ddot{\boldsymbol{z}}|\boldsymbol{x}_{s}, \boldsymbol{x}_{t})} [\log p(\widehat{\boldsymbol{x}}_{ts} | \boldsymbol{\ddot{z}}_{ts})] - \mathbb{E}_{\boldsymbol{h}_{t} \sim p(\boldsymbol{h}_{t}|\boldsymbol{x}_{t})} [\log p(\widetilde{\boldsymbol{x}}_{ts} | \boldsymbol{h}_{s})])^{2} \end{aligned}$$

$$(11)$$

Then, the variational parameters ϕ and generation parameters θ are updated by the following rules:

$$\phi \leftarrow \phi - \eta_1 \nabla (\mathcal{L}_{adv} + \lambda_1 \mathcal{L}_{KL} + \lambda_2 \mathcal{L}_{Rec})$$

$$\theta \leftarrow \theta - \eta_2 \nabla (\mathcal{L}_{Rec} + \mathcal{L}_c),$$
(12)

where η_1, η_2 are the learning rates. Note, that only data from the desired domain (the source) are used to train the reconstruction loss. The KL items approximate the transfer latent space to their prior. Hyperparameters λ_1, λ_2 are used to balance the discriminator loss and reconstruction loss.

4. Experiments

We conducted extensive evaluations of CDLM in two homogeneous transfer scenarios including unsupervised domain adaptation and image-to-image translation. During the experiments, our model was implemented using Tensor-Flow [1]. The structures of the encoder and the decoder adopt those of UNIT [23] which perform well for image translation tasks. A two-layer fully connected MLP was used for the discriminator. SGD with momentum was used for updating the variational parameters, and Adam for updating generation parameters. The batch size was set to 64. During the experiments, we set $\gamma_1 = 1.0, \gamma_2 = 0.1$, $\lambda_1 = \lambda_2 = 0.0001, \ \beta_1 = 0.1 \ \text{and} \ \beta_2 = 0.01.$ For the datasets, we considered a few popular benchmarks, including MNIST [19], MNSITM [7], USPS [18], Fashion-MNIST [43], Linemod [10, 42], Zap50K-shoes [45] and CelebA [25, 24].

4.1. Datasets

We have evaluated our model on a variety of benchmark datasets. They are described as follows.

MNIST: MNIST handwritten dataset [19] is a very popular machine learning dataset. It has a training set of 60,000 binary images, and a test set of 10,000. There are 10 classes in the dataset. In our experiments, we use the standard split of the dataset. MNISTM [7] is a modified version for the MNIST, with random RGB background cropped from the Berkeley Segmentation Dataset¹.

¹URL https://www2.eecs.berkeley.edu/Research/Projects/CS/vision/bsds/

Table 1: Mean classification accuracy comparison. The "source only" row is the accuracy for target without domain adaptation training only on the source. The "target only" is the accuracy of the full adaptation training on the target. For each source-target task the best performance is in bold

Source	MNIST	USPS	MNIST	MNISTM	Fashion	Fashion-M	Linemod 3D
Target	USPS	MNIST	MNISTM	MNIST	Fashion-M	Fashion	Linemod Real
Source Only	0.634	0.625	0.561	0.633	0.527	0.612	0.632
DANN [7]	0.774	0.833	0.766	0.851	0.765	0.822	0.832
CyCADA [11]	0.956	0.965	0.921	0.943	0.874	0.915	0.960
GtA [33]	0.953	0.908	0.917	0.932	0.855	0.893	0.930
CDAN [27]	0.956	0.980	0.862	0.902	0.875	0.891	0.936
PixelDA [3]	0.959	0.942	0.982	0.922	0.805	0.762	0.998
UNIT [23]	0.960	0.951	0.920	0.932	0.796	0.805	0.964
CDLM (\widetilde{x}_{ts})	0.961	0.983	0.987	0.962	0.913	0.922	0.984
Target Only	0.980	0.985	0.983	0.985	0.920	0.942	0.998

USPS: USPS is a handwritten zip digits datasets [18]. It contains 9298 binary images (16×16) , 7291 of which are used as the training set, while the remaining 2007 are used as the test set. The USPS samples are resized to 28×28 , the same as MNIST.

Fashion: Fashion [43] contains 60,000 images for training, and 10,000 for testing. All the images are grayscale, 28×28 in size space. In addition, following the protocol in [7], we add random noise to the Fashion images to generate the FashionM dataset, with random RGB background cropped from the Berkeley Segmentation Dataset.

Linemod 3D images Following the protocol of [3], we render the LineMod [10, 42] for the adaptation between synthetic 3D images (source) and real images (target). The objects with different poses are located at the center of the images. The synthetic 3D images render a black background and a variety of complex indoor environments for real images. We use the RGB images only, not the depth images.

CelebA: CelebA [25] is a large celebrities face image dataset. It contains more than 200k images annotated with 40 facial attributes. We select 50K images randomly, then transform them to sketch images followed the protocol of [24]. The original and sketched images are used for translation. **UT-Zap50K-shoes:** This dataset [45] contains 50K shoes images with 4 different classes. During the translation, we

get the edges produced by canny detector.

4.2. Unsupervised Domain Adaptation

We applied our model to unsupervised domain adaptation, adapting a classifier trained using labelled samples in the source domain to classify samples in the target domain. For this scenario, only the labels of the source images were available during training. We chose DANN [7] as the baseline, but also compared our model with the state-of-the-art domain adaptation methods: Conditional Domain Adaptation Network (CDAN) [27], Pixel-level Domain Adaptation (PixelDA) [3], Unsupervised Image-to-Image translation (UNIT) [23], Cycle-Consistent Adversarial Domain Adaption (CyCADA) [11], and Generate to Adapt (GtA) [33]. We also used source- and target-only training as the lower and upper bound respectively, following the practice in [3, 7].

4.2.1 Quantitative Results

The performance of domain adaptation for the different tasks is shown in Table 1. There are 4 scenarios and 7 tasks. Each scenario has bidirectional tasks for adaptation except LineMod. For LineMod, it is adapted from synthetic 3D image to real objects. For the same adaptation task, we cite the accuracy from the corresponding references, otherwise the accuracies for some tasks are obtained by training the open-source code provided by authors with suggested optimal parameters, for fair comparison.

From Table 1 we can see that the our method has a higher advantage compared with the baseline and the sourceonly accuracy, a little lower than the target-only accuracy from both adaptation directions. In comparison with other models, our model has a better performance for most of tasks. The CDLM has a higher adaptation accuracy for the scenarios with seemingly larger domain gap, such as scenario, the performance of our model is a little lower than PixelDA [3], but outperforms all the other compared methods. In PixelDA, the input is not only source image but also depth image pairs. It might be helpful for the generation. Besides, we visualize the t-SNE [38] for latent encodings $(\ddot{z}_{st}, \ddot{z}_{ts})$ w.r.t the source and the target, respectively. Fig. 4 is the visualization for task MNISTM → MNIST and MNIST \rightarrow USPS, and it shows that both are aligned well.

4.2.2 Qualitative Results

Our model can give the visualization of the adaptation. Fig. 2 is the visualization for the digits and Fashion adaptation respectively. For the scenario of MNIST and USPS, the gener-



(f) Fashionm→Fashion

Figure 2: Visualization for the adaptations. 6 different tasks are illustrated. For each task, the first row shows target images and the second row shows the adapted images with source-like style.

ation for the task USPS \rightarrow MNIST is shown in Fig. 2a. The target MNIST is transferred to the source USPS style well, meanwhile it keeps the correspondent content (label). For example, the digit '1' in MNIST become more leaned and '9' more flatten. Also in Fig. 2b, the target USPS becomes the MNIST style. For the scenario of MNIST \rightarrow MNISTM, our proposed model can remove and add the noise background well for adaptation.

For the scenario of Fashion, the fashion items have more complicated texture and content. In addition, the noisy backgrounds pollute the items randomly, for example, different parts of a cloth are filled with various colors. For visualizations, specifically, Fig. 2e is for the task Fashion \rightarrow FashionM. The proposed model can remove the noisy background and maintain the content. On the other hand, Fig. 2f shows that the original Fashion images are added with similar noisy background as the source. This is promising for a



Figure 3: Linemod 3D Synthetic \rightarrow Real. For a query image (on the left), different adaptation images (to the right) with various poses can be generated.



Figure 4: t-SNE visualization of cross-domain latent encodings $\ddot{z}_{st}, \ddot{z}_{ts}$. The \ddot{z}_{st} are in blue, and the \ddot{z}_{ts} in red.

better adaptation performance.

For Linemod3D, the real objects images with different backgrounds are transferred to the synthetic 3D images with black background. Due to the 3D style, the generation of the target gives different poses. For example in Fig.3, different poses of the iron object are obtained for different trials.

4.3. Cross-Domain Image Mapping

The proposed model also can be used for the crossdomain image mapping.

Fig. 5 gives a demonstration of the image style translation. Specifically, for "shoes" and "edges" in Fig. 5a and 5c, we can see that the proposed model can translate "edges" to its counterpart quite well. The translation is stochastic – an "edge" pattern can be used to generate "shoes" in different colors with different trials. For the more challenging "face" and "sketch" translations, the proposed model also performs well. The generations have some variations compared with the original images. In general, our method can generate realistic translated images. However, we find that compared with the translation from sketches to real images, the reverse task seems harder. For example, when a face image is given, the generated sketch loses some details. The reason may be the low-level feature is neglected when the deep feature acts as the condition.

For further evaluation, quantitative performance is evalu-



Figure 5: Visualization of cross-domain image mapping.

ated for image mapping. SSIM [40], MSE, and PSNR are used for the evaluation. The results are shown in Table 2. We can see that our model outperforms E-CDRD [24], which learns a disentangled latent encoding for the source and the target for domain adaptation. Meanwhile, it matches the performance of StarGAN [5], which is designed for multi-domain image translation. The result shows that our model can map cross-domain images well compared to these prior works.

Table 2: Performance for image mapping.

Models	"Sketch" to "Face"			
widdels	SSIM	MSE	PSNR	
E-CDRD [24]	0.6229	0.0207	16.86	
StarGAN [5]	0.8026	0.0142	19.04	
CDLM	0.7961	0.0140	19.89	

In addition, we also conduct the classification to evaluate the translation performance. We take shoes as an example which are labeled to 4 different classes. The recognition accuracy of our proposed model for task shoes \rightarrow edge is 0.953, which is higher than the results of PixelDA (0.921) and UNIT (0.916) respectively.

4.4. Model Analysis

The effect of encoder settings – depth and different γ_1, γ_2 : In our model, the deep features are utilized to crossmodulate the transfer latent encoding. Therefore the deep feature is an important factor in our framework and is influenced by the depth of the encoder. During the experiments, we use MNIST \rightarrow USPS and Fashion \rightarrow FashionM as the evaluation tasks. For the first one, they have different content, but with the same background. The second task is a totally different scenario, the images have the same content but different background. The outputs of different encoder layers (k > 3) are used for the experiments.

Table 3: Adaptation accuracy with different layer depth for Tasks MNIST \rightarrow USPS and Fashion \rightarrow FashionM.

Tasks/Layers	Conv4	Conv5	Conv _{last}
$MNIST \to USPS$	0.954	0.956	0.961
Fashion \rightarrow FashionM	0.890	0.905	0.913

Table 4: Adaptation accuracy with different (γ_1, γ_2) for Tasks MNIST \rightarrow USPS and Fashion \rightarrow FashionM.

Tasks / (γ_1, γ_2)	(0.1,1.0)	(0.5,0.5)	(0.9,0.1)	(1.0,0.1)	(1.0, 0)
$MNIST \rightarrow USPS$	0.320	0.723	0.961	0.961	0.961
$Fashion \rightarrow FashionM$	0.226	0.513	0.912	0.913	0.913

As the result (Table 3) shows, a higher accuracy is achieved when more layers are used to extract the deep representations. The accuracy gain of the task MNIST \rightarrow USPS is lower than that of Fashion \rightarrow FashionM. This is expected as features extracted by higher layers would normally eliminate lower-level variations between domains, such as change of background and illumination in the images.

For γ_1, γ_2 , we fixed the last convolutional layer for the deep representations and evaluate different values. From Table 4, we can see that the performance drops down significantly with a smaller γ_1 compared with γ_2 , and increased with a larger γ_1 . The performance seems to be stabilized when γ_1 is greater than 0.9 while γ_2 remains 0.1. Following the standard VAE, we keep the noise ϵ ($\gamma_2 \neq 0$) in the evaluations. Meanwhile, our model works well even when

Table 5: Evaluation on the effect of unsupervised consistency metrics. The recognition accuracy is shown for four tasks in the unsupervised domain adaptation scenario. Our model is on the last row with both the \mathcal{L}_{c}^{s} and \mathcal{L}_{c}^{t} , which achieves the best performance.

Model/Tasks	MNIST→USPS	USPS→MNIST	Fashion→FashionM	FashionM→Fashion
CDLM w/o \mathcal{L}_c	0.635	0.683	0.646	0.672
$CDLM+\mathcal{L}_{c}^{t}$	0.689	0.695	0.682	0.691
$\text{CDLM} + \mathcal{L}_c^s$	0.951	0.980	0.912	0.915
$\text{CDLM} + \mathcal{L}_c^s + \mathcal{L}_c^t$	0.961	0.983	0.913	0.922

 $\gamma_2 = 0$. These results suggest the deep representation plays a crucial role in the cross-domain modulation.

A-Distance: In a theoretical analysis of the domain discrepancy [2], Ben-David *et al.* suggests that A-distance can be used as a measure for the domain discrepancy. As the exact A-distance is intractable, a proxy is defined as $\hat{d}_A = 2(1 - 2\epsilon)$, where ϵ is the generalization error of a binary classifier (e.g. kernel SVM) trained to distinguish the input's domain (source or target). Following the protocol of [26, 32], we calculate the A-distance on four adaptation tasks under the scenarios of Raw features, DANN features, and CDLM features respectively. The results are show in Fig. 6. We observe that both the DANN and CDLM reduce the domain discrepancy compared with the Raw images scenario, and the A-distance of CDLM is smaller than the DANN's. This demonstrates that it is harder to distinguish the source and the target by the CDLM generations.



Figure 6: A-distances comparison for four tasks.

Convergence: We also conduct the convergence experiment with training error on task MNSIT-USPS to evaluate our model. As shown in the Fig. 7, our model has a better convergence than DANN, thought there are some oscillations at the beginning of the training. In addition, the error of CDLM is lower that the DANN, which demonstrate that CDLM has a better adaptation performance. This is consistent with the adaptation performance in Table 1.

The effect of unsupervised consistency metrics: In our model, two unsupervised consistency metrics are added for generation in good effects. The adaptation accuracy is used for evaluation. Table 5 is the results for the four different tasks. The performance w/o \mathcal{L}_c is dropped down because the



Figure 7: Convergence of CDLM compared with DANN.

decoder cannot generate realistic cross-domain images. \mathcal{L}_c^t connects outputs generated from the h_t and \ddot{z}_{ts} only for the target, which improves the performance slightly. Meanwhile, we can see that the \mathcal{L}_c^s loss boosts the accuracy for adaptation significantly, which connects the two domains with the generations by the h. Finally, the scenario with both \mathcal{L}_c^s and \mathcal{L}_c^t gives the best performance in all four tasks. It bridges both the h and \ddot{z} between the two domains.

5. Conclusion

In this paper, we have presented a novel variational crossdomain transfer learning model with cross modulation of deep representations from different domains. A shared transfer latent space is introduced, and the reparameterization transformation is modified to enforce the connection between domains. Evaluations carried out in unsupervised domain adaptation and image translation tasks demonstrate our model's competitive performance. Its effectiveness is also clearly shown in visual assessment of the adapted images, as well as in the alignment of the latent information as revealed by visualization using t-SNE. Overall, competitive performance has been achieved by our model despite its relative simplicity.

For future work, we intend to further improve our variational transfer learning framework and use it for heterogeneous, multi-domain transfer tasks.

References

- Mart Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous distributed systems. arXiv:1603.04467 [cs.DC], 2016. 4
- [2] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In Advances in Neural Information Processing Systems (NeurIPS), pages 137–144, 2007. 8
- [3] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 95–104, 2017. 2, 5
- [4] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In Advances in Neural Information Processing Systems (NeurIPS), pages 343–351, 2016. 2
- [5] Yunjey Choi, Minje Choi, Munyoung Kim, Jung Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8789–8797, 2018. 2, 7
- [6] Zunlei Feng, Anxiang Zeng, Xinchao Wang, Dacheng Tao, Chenglong Ke, and Mingli Song. Dual swap disentangling. In Advances in Neural Information Processing Systems (NeurIPS), pages 5894–5904, 2018. 2
- [7] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research* (*JMLR*), 17(59):1–35, 2016. 3, 4, 5
- [8] Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstructionclassification networks for unsupervised domain adaptation. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *European Conference on Computer Vision (ECCV)*, pages 597–613, 2016. 3
- [9] Abel Gonzalez-garcia, Joost van de Weijer, and Yoshua Bengio. Image-to-image translation for cross-domain disentanglement. In Advances in Neural Information Processing Systems (NeurIPS), pages 1294–1305, 2018. 2
- [10] Stefan Hinterstoisser, Cedric Cagniart, Slobodan Ilic, Peter Sturm, Nassir Navab, Pascal Fua, and Vincent Lepetit. Gradient response maps for real-time detection of texture-less objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 34(5):876–888, 2012. 4, 5
- [11] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In Jennifer G Dy and Andreas Krause, editors, *International Conference on Machine Learning (ICML)*, volume 80 of *PMLR*, pages 3162–3174, 2018. 2, 5
- [12] Matthew D Hoffman and Matthew J Johnson. ELBO surgery: yet another way to carve up the variational evidence lower bound. In Advances in Neural Information Processing Systems Workshop (NeurIPSW), 2016. 1

- [13] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In Yee Whye Teh Doina Precup, editor, *International Conference on Machine Learning (ICML)*, volume 70 of *PMLR*, pages 1857–1865, 2017. 2, 3
- [14] Diederik P Kingma. Fast gradient-based inference with continuous latent variable models in auxiliary form. arXiv:1306.0733 [cs.LG], 2013. 1
- [15] Diederik P. Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In Advances in Neural Information Processing Systems (NeurIPS), pages 4743–4751, 2016. 1
- [16] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Repre*sentations (ICLR), 2014. 2, 3
- [17] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *International Conference* on Machine Learning (ICML), volume 48 of PMLR, pages 1558–1566, 2016. 1
- [18] Yann Le Cun, L D Jackel, B Boser, J S Denker, H P Graf, Isabelle Guyon, Don Henderson, R E Howard, and W Hubbard. Handwritten digit recognition: Applications of neural network chips and automatic learning. *IEEE Communications Magazine*, 27(11):41–46, 1989. 4, 5
- [19] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
 4
- [20] Hsin Ying Lee, Hung Yu Tseng, Jia Bin Huang, Maneesh Singh, and Ming Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Europen Conference on Computer Vision (ECCV)*, pages 36–52, 2018. 2
- [21] Torgny Lindvall. *Lectures on the Coupling Method*. Dover, 2002. 1
- [22] Alexander H. Liu, Yen Cheng Liu, Yu Ying Yeh, and Yu Chiang Frank Wang. A unified feature disentangler for multidomain image translation and manipulation. In Advances in Neural Information Processing Systems (NeurIPS), pages 2590–2599, 2018. 1, 2
- [23] Ming-Yu Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In Advances in Neural Information Processing Systems (NeurIPS), pages 700–708, 2017. 1, 2, 4, 5
- [24] Yen Cheng Liu, Yu Ying Yeh, Tzu Chien Fu, Sheng De Wang, Wei Chen Chiu, and Yu Chiang Frank Wang. Detach and adapt: Learning cross-domain disentangled deep representation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8867–8876, 2018. 4, 5, 7
- [25] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3730– 3738, 2015. 4, 5
- [26] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation

networks. In Francis Bach and David Blei, editors, *International Conference on Machine Learning (ICML)*, volume 37 of *PMLR*, pages 97–105, 2015. 8

- [27] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In Samy Bengio, Hanna M Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, Advances in Neural Information Processing Systems (NeurIPS), pages 1640–1650, 2018. 5
- [28] Shweta Mahajan, Iryna Gurevych, and Stefan Roth. Latent normalizing flows for many-to-many cross-domain mappings. In *International Conference on Learning Representations* (*ICLR*), 2020.
- [29] Muzammal Naseer, Salman H. Khan, Harris Khan, Fahad Shahbaz Khan, and Fatih Porikli. Cross-domain transferability of adversarial perturbations. In Advances in Neural Information Processing Systems (NeuIPS), 2019. 2
- [30] Atsuhiro Noguchi and Tatsuya Harada. Image generation from small datasets via batch statistics adaptation. In *IEEE Conference on Computer Vision (ICCV)*, pages 2750–2758, 2019. 2
- [31] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* (*TKDE*), 22(10):1345–1359, 2010. 1
- [32] Xingchao Peng, Zijun Huang, Yizhe Zhu, and Kate Saenko. Federated adversarial domain adaptation. In *International Conference on Learning Representations (ICLR)*, 2020. 8
- [33] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), pages 8503–8512, 2018. 2, 5
- [34] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8247–8255, 2019. 1, 3
- [35] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. In Vera Kurková, Yannis Manolopoulos, Barbara Hammer, Lazaros S Iliadis, and Ilias Maglogiannis, editors, *International Conference on Artificial Neural Networks (ICANN)*, pages 270–279. Springer, 2018. 1
- [36] Jakub M. Tomczak and Max Welling. VAE with a vampprior. In Fernando Perez-Cruz Amos Storkey, editor, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 84 of *PMLR*, pages 1214–1223, 2018. 1
- [37] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7167–7176, 2017. 3
- [38] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. Journal of machine learning research (JMLR), 9:2579–2605, 2008. 5
- [39] Liwei Wang, Alexander G. Schwing, and Svetlana Lazebnik. Diverse and accurate image description using a variational auto-encoder with an additive gaussian encoding space. In *Ad*-

vances in Neural Information Processing Systems (NeurIPS), pages 5757–5767, 2017. 1

- [40] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 7
- [41] Karl R Weiss, Taghi M Khoshgoftaar, and Dingding Wang. A survey of transfer learning. *Journal of Big Data*, 3(9):1–40, 2016. 1
- [42] Paul Wohlhart and Vincent Lepetit. Learning descriptors for object recognition and 3d pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3109–3118, 2015. 4, 5
- [43] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. arXiv:1708.07747 [cs.LG], 2017. 4, 5
- [44] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In Advances in Neural Information Processing Systems (NeurIPS), pages 3320–3328, 2014. 2
- [45] Aron Yu and Kristen Grauman. Semantic jitter: Dense supervision for visual comparisons via synthetic images. In *IEEE Conference on Computer Vision (ICCV)*, pages 5570–5579, 2017. 4, 5
- [46] Jianfu Zhang, Yuanyuan Huang, Yaoyi Li, Weijie Zhao, and Liqing Zhang. Multi-attribute transfer via disentangled representation. In AAAI Conference on Artificial Intelligence (AAAI), volume 33, pages 9195–9202, 2019. 2
- [47] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE Conference on Computer Vi*sion (ICCV), pages 2242–2251, 2017. 2
- [48] Zachary M. Ziegler and Alexander M. Rush. Latent normalizing flows for discrete sequences. In Ruslan Salakhutdinov Kamalika Chaudhuri, editor, *International Conference on Machine Learning (ICML)*, volume 97 of *PMLR*, pages 7673– 7682, 2019. 1