
Multilinear Latent Conditioning for Generating Unseen Attribute Combinations

Markos Georgopoulos¹ Grigorios G Chrysos¹ Maja Pantic¹ Yannis Panagakis²

Abstract

Deep generative models rely on their inductive bias to facilitate generalization, especially for problems with high dimensional data, like images. However, empirical studies have shown that variational autoencoders (VAE) and generative adversarial networks (GAN) lack the generalization ability that occurs naturally in human perception. For example, humans can visualize a woman smiling after only seeing a smiling man. On the contrary, the standard conditional VAE (cVAE) is unable to generate unseen attribute combinations. To this end, we extend cVAE by introducing a multilinear latent conditioning framework that captures the multiplicative interactions between the attributes. We implement two variants of our model and demonstrate their efficacy on MNIST, Fashion-MNIST and CelebA. Altogether, we design a novel conditioning framework that can be used with any architecture to synthesize unseen attribute combinations.

1. Introduction

Deep generative models have been widely utilized to push the boundaries of the state-of-the-art in various tasks including image (Miyato et al., 2018), audio (Andreux & Mallat, 2018), as well as text synthesis (Radford et al., 2018). Their success would not be possible without the development of complex and expressive models (Miyato et al., 2018; Karras et al., 2019; Du & Mordatch, 2019) as well as the collection of rich and thoroughly annotated datasets (Liu et al., 2015; Deng et al., 2009). However, due to the high data complexity of these tasks (Arora & Zhang, 2017), generative models have to rely on a set of assumptions, i.e., their inductive

¹Department of Computing, Imperial College London, United Kingdom ²Department of Informatics and Telecommunications, University of Athens, Greece. Correspondence to: Markos Georgopoulos <m.georgopoulos@imperial.ac.uk>.

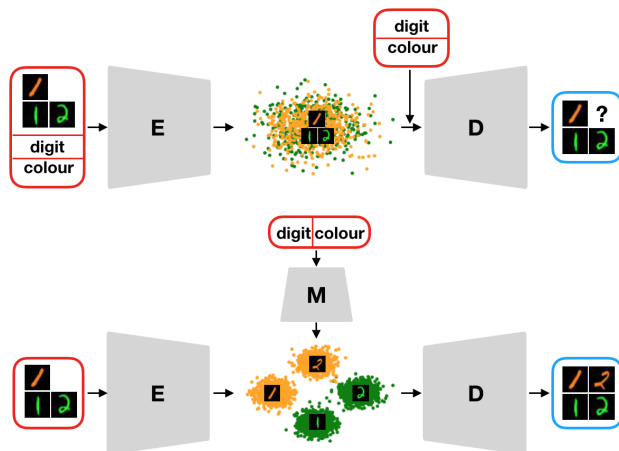


Figure 1. Schematic of the conditional VAE (top) and the proposed framework (bottom). The baseline (conditional VAE) lacks the structure in the latent space and cannot generate images of attribute combinations that are absent from the training set (e.g., orange 2). On the contrary, the proposed method capitalizes on the multilinear latent conditioning to synthesize images of unseen attribute combinations.

bias. These assumptions are based on the distribution of the training set, as well as the model itself, and are crucial for generalization. However, the training set is not always an accurate description of the real world. Hence, the generated distribution is constrained by the distribution of the training data.

The empirical study conducted in (Zhao et al., 2018) is the first to investigate the inductive bias of deep generative models. The study focuses on the unsupervised setting and demonstrates the inability of variational autoencoders (VAEs) (Kingma & Welling, 2014; Rezende et al., 2014) and generative adversarial networks (GANs) (Goodfellow et al., 2014) to generate images of unseen attribute combinations. For example, consider the task of generating an image of a smiling woman without seeing one in the training set. We posit that in order to further advance the state-of-the-art, generative models should be able to generalize to unseen attribute combinations. In this work, we focus on the supervised setting and investigate the conditional generation of attribute combinations that are absent from the training set.

To this end, using the VAE as the main building block, we propose a multilinear conditioning framework that is tailored to the task. The VAE is a generative model that learns the distribution of the data by introducing a latent variable z . A common assumption is that z belongs to a simple prior distribution, for instance, a standard normal. Instead of considering a single prior distribution for all latent representations, we propose to learn a different Gaussian for each class. By learning the mean of each distribution as a multilinear function of the labels, our framework is able to account for their multiplicative interactions. In order to learn the higher order tensors in the presence of missing classes (i.e., attribute combinations), we utilize two tensor decompositions, namely the CP and Tucker, and train the model end-to-end.

We evaluate the efficacy of our method through a series of extensive experiments on benchmarks based on standard datasets. These benchmarks are designed to evaluate both the ability to generate images of unseen attribute combinations, as well as model the correlation between labels. Furthermore, we explore the limits of our model and propose two extensions. Concretely:

- In Section 4, our extension to the VAE and its two variant using the Tucker and CP decompositions are proposed. Contrary to the standard conditional VAE, the proposed framework is able to generate unseen classes. We further exhibit how our model achieves this by modeling the multiplicative interactions between the labels.
- We demonstrate the ability of our models to generalize on unseen attribute combinations in Section 5. In particular, we create 3 evaluation benchmarks based on the MNIST, FashionMNIST and CelebA datasets. For each benchmark, we use 2 sets of labels and remove one label combination from the training set. We evaluate the generated images both qualitatively and quantitatively. The models are also evaluated on attribute transfer on CelebA.
- We showcase the importance of modeling inter-label interactions in Section 6. To this end, we introduce two benchmarks based on MNIST and CelebA. We show experimentally that only our multilinear latent conditioning framework is able to recover the missing attribute combinations in the presence of highly correlated labels.
- In Section 7 we explore the proposed model. First, we study how the model performance scales with the number of missing attribute combinations. Furthermore, we extend the framework to handle multiple labelled attributes, as well as unlabelled attributes.

2. Related work

The variational autoencoder is a widely adopted class of generative models. In the seminal work of Kingma & Welling (2014), VAE learns a generative model of the form $p(x, z) = p(x|z)p(z)$, where $p(z)$ is the prior distribution of the latent variable. A common choice for the prior distribution is a standard normal distribution, however this assumption can be relaxed. More expressive priors have been proposed to model the multimodal nature of the data, e.g., mixture of priors (Tomczak & Welling, 2017; Dilokthanakul et al., 2016; Grønbech et al., 2018; Jiang et al., 2016; Van Den Oord et al., 2017), normalizing flows (Rezende & Mohamed, 2015) and autoregressive models (Chen et al., 2016). The VAE framework was originally proposed for the unsupervised setting but several successful supervised and semi-supervised variations have been proposed thereafter (Van den Oord et al., 2016; Razavi et al., 2019; Siddharth et al., 2017).

In order to incorporate label information into the model, a common practice is concatenating the labels to the latent representation (Sohn et al., 2015). However, doing so fails to uncover the underlying structure of the latent space. Klys et al. (2018) introduce a low-dimensional subspace that captures the class-specific latent structure and use it to perform attribute transfer. Similarly, our model learns the class specific structure of the latent space by considering the multiplicative interactions between the labels.

The multilinear structure of different modes of variation has been previously studied in the unsupervised setting. Tang et al. (2013) extend the Factor Analyzers method to also model multiplicative interactions between the latent factors. Wang et al. (2017) introduce an unsupervised tensor decomposition that learns the underlying multilinear structure of the data. In this work, we focus on the supervised setting and propose to model the interactions between the labeled modes of variation.

Lastly, in order to synthesize images of a specified attribute combination, the proposed model implicitly disentangles the labelled information. Disentanglement of factors of variation is an active research topic in the unsupervised setting (Higgins et al., 2016; Kim & Mnih, 2018; Chen et al., 2018; Burgess et al., 2018; Ridgeway & Mozer, 2018; Siddharth et al., 2017). Locatello et al. (2019) proved that unsupervised learning of disentangled representations is not possible without inductive biases. Furthermore, recent work (Khemakhem et al., 2020) highlights the lack of identifiability without auxiliary information, such as labels. In this work, we focus on the supervised setting and aim to recover the unobserved attribute combinations, a task that is not a direct byproduct of disentanglement.

Concurrently with this work, (Bozkurt et al., 2019) evaluate

whether VAEs can reconstruct data of unseen classes. This work differentiates itself by focusing on the generation of samples of unobserved attribute combinations.

3. Background

3.1. Notation

Tensors are symbolized by calligraphic letters, e.g., \mathcal{X} , while matrices (vectors) are denoted by uppercase (lowercase) boldface letters e.g., \mathbf{X} , (\mathbf{x}).

Products: The *Hadamard* product of $\mathbf{A} \in \mathbb{R}^{I \times N}$ and $\mathbf{B} \in \mathbb{R}^{I \times N}$ is defined as $\mathbf{A} * \mathbf{B}$ and is equal to $A_{(i,j)} B_{(i,j)}$ for the (i, j) element.

The *Kronecker* product of matrices $\mathbf{A} \in \mathbb{R}^{I \times M}$ and $\mathbf{B} \in \mathbb{R}^{J \times N}$ is denoted by $\mathbf{A} \otimes \mathbf{B}$ and yields a matrix of dimensions $(IJ) \times (MN)$ that is defined as:

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \dots & a_{1J}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \dots & a_{2J}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{I1}\mathbf{B} & a_{I2}\mathbf{B} & \dots & a_{IJ}\mathbf{B} \end{bmatrix}. \quad (1)$$

The *Khatri-Rao* (column-wise *Kronecker*) product of matrices $\mathbf{A} \in \mathbb{R}^{I \times N}$ and $\mathbf{B} \in \mathbb{R}^{J \times N}$ is denoted by $\mathbf{A} \odot \mathbf{B}$ and yields a matrix of dimensions $(IJ) \times N$ that is defined as:

$$\mathbf{A} \odot \mathbf{B} = [\mathbf{a}_1 \otimes \mathbf{b}_1 \quad \mathbf{a}_2 \otimes \mathbf{b}_2 \quad \dots \quad \mathbf{a}_N \otimes \mathbf{b}_N]. \quad (2)$$

In the case of vectors, the Khatri-Rao and Kronecker products are equivalent.

3.2. Tensor operations and decompositions

Each element of a K^{th} order tensor \mathcal{X} is addressed by K indices, i.e., $(\mathcal{X})_{i_1, i_2, \dots, i_K} \doteq x_{i_1, i_2, \dots, i_K}$. A K^{th} -order real-valued tensor \mathcal{X} is defined over the tensor space $\mathbb{R}^{I_1 \times I_2 \times \dots \times I_K}$, where $I_k \in \mathbb{Z}$ for $k = 1, 2, \dots, K$.

Tensor unfolding: The *mode- k unfolding* of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_K}$ maps \mathcal{X} to a matrix $\mathbf{X}_{(k)} \in \mathbb{R}^{I_k \times \bar{I}_k}$ with $\bar{I}_k = \prod_{t \neq k}^{K} I_t$ such that the tensor element x_{i_1, i_2, \dots, i_K} is mapped to the matrix element $x_{i_t, j}$ where $j = 1 + \sum_{t \neq k}^{K} (i_t - 1) J_t$ with $J_t = \prod_{n=1}^{t-1} I_n$.

Vector product: The *mode- k vector product* of \mathcal{X} with a vector $\mathbf{u} \in \mathbb{R}^{I_k}$, denoted by $\mathcal{X} \times_k \mathbf{u} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_{k-1} \times I_{k+1} \times \dots \times I_K}$, results in a tensor of order $K - 1$:

$$(\mathcal{X} \times_k \mathbf{u})_{i_1, \dots, i_{k-1}, i_{k+1}, \dots, i_K} = \sum_{i_k=1}^{I_k} x_{i_1, i_2, \dots, i_K} u_{i_k}. \quad (3)$$

CP tensor decomposition: The *CANDECOMP/PARAFAC* (CP) decomposition (Kolda & Bader, 2009; Sidiropoulos

et al., 2017) factorizes a tensor into a sum of component rank-one tensors. The rank- R CP decomposition of a K^{th} -order tensor \mathcal{X} is written as:

$$\mathcal{X} \doteq \llbracket \mathbf{U}_{[1]}, \mathbf{U}_{[2]}, \dots, \mathbf{U}_{[K]} \rrbracket = \sum_{r=1}^R \mathbf{u}_r^{(1)} \circ \mathbf{u}_r^{(2)} \circ \dots \circ \mathbf{u}_r^{(K)}, \quad (4)$$

where \circ denotes for the vector outer product. The factor matrices $\{\mathbf{U}_{[k]} = [\mathbf{u}_1^{(k)}, \mathbf{u}_2^{(k)}, \dots, \mathbf{u}_R^{(k)}] \in \mathbb{R}^{I_k \times R}\}_{k=1}^K$ collect the vectors from the rank-one components. By considering the mode-1 unfolding of \mathcal{X} , the CP decomposition of a third order tensor can be written in matrix form as (Kolda & Bader, 2009):

$$\mathbf{X}_{(1)} \doteq \mathbf{U}_{[1]} (\mathbf{U}_{[3]} \odot \mathbf{U}_{[2]})^T. \quad (5)$$

Tucker tensor decomposition: The Tucker tensor decomposition (Tucker, 1963) factorizes a tensor into a set of matrices and a core tensor, as follows:

$$\begin{aligned} \mathcal{X} &\doteq \llbracket \mathcal{G}; \mathbf{U}_{[1]}, \mathbf{U}_{[2]}, \dots, \mathbf{U}_{[K]} \rrbracket \\ &= \mathcal{G} \times_1 \mathbf{U}_{[1]} \times_2 \mathbf{U}_{[2]} \times_3 \dots \times_K \mathbf{U}_{[K]}, \end{aligned} \quad (6)$$

Similarly to CP, the Tucker decomposition of a third order tensor can be written in matrix form as:

$$\mathbf{X}_{(1)} \doteq \mathbf{U}_{[1]} (\mathbf{U}_{[3]} \otimes \mathbf{U}_{[2]})^T. \quad (7)$$

Tucker can be considered a generalization of SVD to higher order tensors (De Lathauwer et al., 2000). The CP decomposition is a special case of Tucker, where the core tensor is superdiagonal.

4. Model

Consider the setting of a dataset of images \mathbf{x} that are annotated with a set of N one-hot label vectors $\mathbf{y} = \{\mathbf{y}_i \in \mathbb{R}^{D_i}\}_{i=1}^N$ and latent representations $\mathbf{z} \in \mathbb{R}^d$. The conditional VAE (cVAE) (Sohn et al., 2015) is trained by maximizing the following objective:

$$\begin{aligned} L_{CVAE} &= \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})} [\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{y})] \\ &\quad - \beta D_{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y}) \parallel p(\mathbf{z})) \leq \log p_\theta(\mathbf{x}|\mathbf{y}), \end{aligned} \quad (8)$$

where D_{KL} is the Kullback Leibler divergence between the approximate posterior q_ϕ and the prior $p(\mathbf{z})$. The approximate posterior is modeled by the encoder network, while the prior is assumed to be standard normal.

In this work, we introduce Multilinear Conditioning VAE (MLC-VAE), which is a generative model of the form $p(\mathbf{x}, \mathbf{y}, \mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{y}|\mathbf{z})p(\mathbf{z})$. We modify the objective of cVAE by introducing a label reconstruction term

$\log p_{\theta_y}(\mathbf{y}|\mathbf{z})$ which is parametrized by a label decoder. The new evidence lower bound (ELBO) is of the form:

$$L_{MLC-VAE} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})} [\log p_{\theta_x}(\mathbf{x}|\mathbf{z}) + \log p_{\theta_y}(\mathbf{y}|\mathbf{z})] - \beta D_{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y}) \parallel p(\mathbf{z})) \quad (9)$$

Instead of assuming a single prior distribution, we consider that the latent representations of each class (i.e., combination of attribute labels $\mathbf{y}_1, \dots, \mathbf{y}_N$) are centered around a different mean, that is, $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{y}_1, \dots, \mathbf{y}_N}, \mathbf{I})$. The class means $\boldsymbol{\mu}_{\mathbf{y}_1, \dots, \mathbf{y}_N} \in \mathbb{R}^d$ are learned as a multivariate function $M(\mathbf{y}_1, \dots, \mathbf{y}_N)$ of the label vectors. Thus, we are able to obtain a latent representation for every attribute combination, even if it is not represented in the training set.

A simple choice for M would be the linear combination of each attribute. Concretely, in the case of the 2 attributes \mathbf{y}_1 and \mathbf{y}_2 :

$$M(\mathbf{y}_1, \mathbf{y}_2) = \mathbf{W}^{[1]}\mathbf{y}_1 + \mathbf{W}^{[2]}\mathbf{y}_2 \quad (10)$$

However, this formulation completely ignores the non-linear (e.g., multiplicative) interactions between the labels. To account for the multiplicative interactions, we introduce a higher order term that models the correlations between the attributes:

$$M(\mathbf{y}_1, \mathbf{y}_2) = \mathbf{W}^{[1]}\mathbf{y}_1 + \mathbf{W}^{[2]}\mathbf{y}_2 + \mathcal{W}^{[12]} \times_2 \mathbf{y}_1 \times_3 \mathbf{y}_2 \quad (11)$$

The first two terms in 11 are equal to the linear combination of the attribute labels, with the matrices $\mathbf{W}^{[1]}$ and $\mathbf{W}^{[2]}$ learning a mapping between each individual label and the latent space. The third order tensor $\mathcal{W}^{[12]}$ captures the pairwise multiplicative interaction between \mathbf{y}_1 and \mathbf{y}_2 and learns a latent representation for each attribute combination. We investigate the effect of the higher order term in Section 6.

We can rewrite the higher order term in 11 by replacing the tensor by its *mode-1* unfolding as such:

$$M(\mathbf{y}_1, \mathbf{y}_2) = \mathbf{W}^{[1]}\mathbf{y}_1 + \mathbf{W}^{[2]}\mathbf{y}_2 + \mathbf{W}_{(1)}^{[12]}(\mathbf{y}_2 \odot \mathbf{y}_1) \quad (12)$$

Since $\mathbf{y}_1 \in \mathbb{R}^{D_1}$ and $\mathbf{y}_2 \in \mathbb{R}^{D_2}$ are one-hot encoded vectors, their Khatri-Rao product ($\mathbf{y}_2 \odot \mathbf{y}_1$) yields a new one-hot vector $\mathbf{y} \in \mathbb{R}^{D_1 D_2}$. Therefore, the product $\mathbf{W}_{(1)}^{[12]}(\mathbf{y}_2 \odot \mathbf{y}_1)$ yields the column of the matrix $\mathbf{W}_{(1)}^{[12]}$ that corresponds to class $(\mathbf{y}_1, \mathbf{y}_2)$. Hence, in the presence of absent attribute combinations, it is not possible to learn the matrix using stochastic gradient descent based optimizers. However, if we let the tensor $\mathcal{W}^{[12]}$ admit a low-rank decomposition, we can overcome this issue as we show below. We explore the *CP* and *Tucker* decompositions and propose two different models, namely MLC-VAE-T and MLC-VAE-CP.

After applying the CP and Tucker decompositions on 12, we obtain respectively:

$$M_{CP} = \mathbf{W}^{[1]}\mathbf{y}_1 + \mathbf{W}^{[2]}\mathbf{y}_2 + \mathbf{U}_{[1]}(\mathbf{U}_{[3]}^T \mathbf{y}_2 * \mathbf{U}_{[2]}^T \mathbf{y}_1) \quad (13)$$

$$M_T = \mathbf{W}^{[1]}\mathbf{y}_1 + \mathbf{W}^{[2]}\mathbf{y}_2 + \mathbf{U}_{[1]}\mathbf{G}_{(1)}^{[12]}(\mathbf{U}_{[3]}^T \mathbf{y}_2 \otimes \mathbf{U}_{[2]}^T \mathbf{y}_1) \quad (14)$$

With this formulation, we can obtain $\mu_{\mathbf{y}_1, \mathbf{y}_2}$ for every class combination $(\mathbf{y}_1, \mathbf{y}_2)$ even when not all possible combinations are seen during training. Thereafter, we can optimize 9 for $\mathbf{z} \sim \mathcal{N}(M(\mathbf{y}_1, \mathbf{y}_2), \mathbf{I})$.

5. Experiments

To evaluate our model on multi-attribute conditional image generation, we perform experiments on the MNIST (LeCun et al., 1998), Fashion-MNIST (Xiao et al., 2017) and CelebA (Liu et al., 2015) datasets. In all the experiments the decoder networks are convolutional architectures based on DCGAN (Radford et al., 2015). The encoders are constructed as the 'mirrored' versions of the decoders. We compare both proposed methods with the baseline cVAE (Sohn et al., 2015) and VampPrior (Tomczak & Welling, 2017); we focus on generation of the unseen combinations but also visualize samples from the seen classes. To evaluate the models quantitatively, we train attribute classifiers. In particular, we train a convolutional neural network with the same architecture as the corresponding VAE encoder for all attributes except for color. For the color attribute used in MNIST and Fashion-MNIST, we train an MLP on the color histograms of the images. The classification accuracy on the test set was above 97% for all attribute classifiers, except for smile, where the accuracy was 92%. The quantitative results on the unseen attribute combinations are presented in Table 1.

Overall, the results in Table 1 highlight the inability of the baseline models to generate novel attribute combinations. The accuracy of the classifiers on the unseen combinations are mostly worse than random chance. This implies that, contrary to our models, the baselines are not able to synthesize recognizable digits, objects and faces of the target attributes. This is verified qualitatively in Figure 2. On the other hand, all models are able to generate the seen attribute combinations (Figure 2). Quantitative results for the seen attribute combinations can be found in the supplementary material.

Implementation details: For the experiments on MNIST and Fashion-MNIST the encoder and decoder networks have 4 layers, while the networks for CelebA have 5 layers. All label decoders are affine transformations. We set $\beta = 1$ for

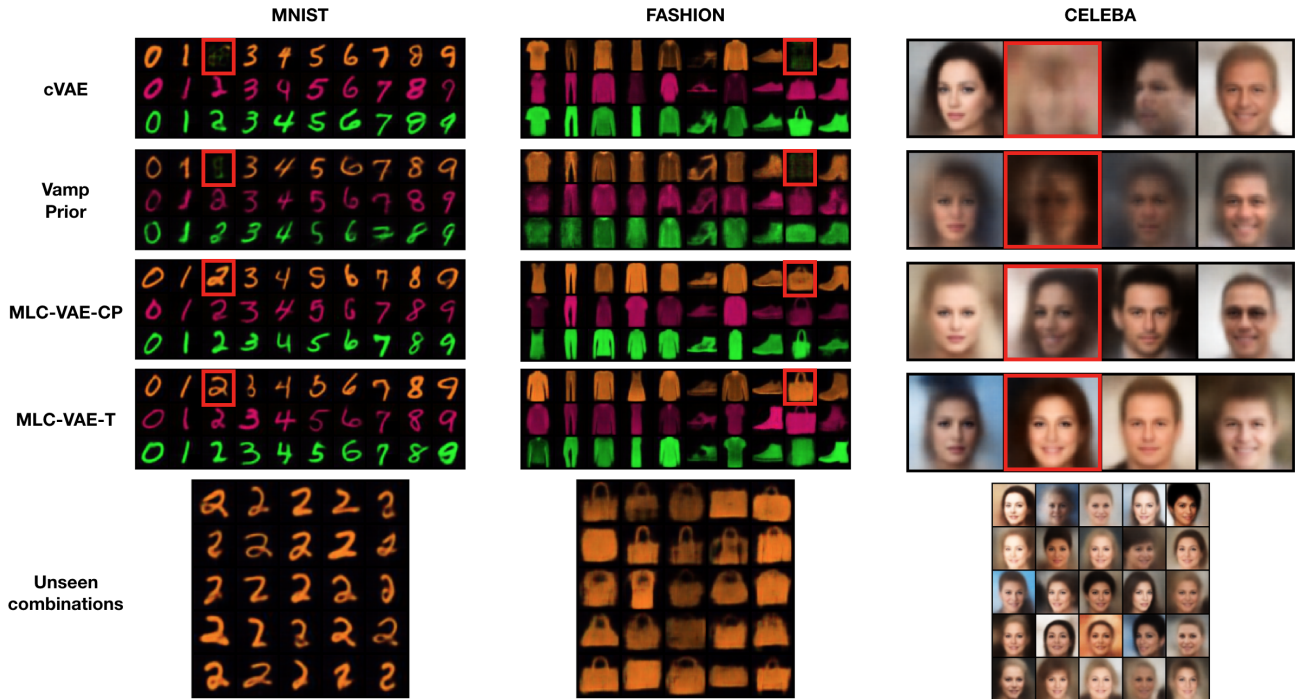


Figure 2. (top) Samples synthesized by the compared methods on the MNIST, Fashion-MNIST, CelebA. All methods can synthesize samples from combinations that are present in the training set, however the compared methods fail on unseen combinations (marked with red boxes). (bottom) Synthesized samples corresponding to combinations that are not seen during training (orange 2, orange bag and woman with smile). Our method is able to synthesize diverse samples of these combinations.

all experiments, except for CelebA where we set $\beta = 10$. The input label for the cVAE and VampPrior baselines is a one-hot vector where each class corresponds to an attribute combination (e.g., (1, orange) or (Male, Smiling)). For fair comparison we train all models for 50 epochs using the Adam optimizer (Kingma & Ba, 2014) with a learning rate of 0.0005. All models were implemented in Pytorch (Paszke et al., 2017) and Tensorly (Kossaifi et al., 2019).

MNIST: The MNIST dataset consists of 60k training images and 10k test images of handwritten digits. The original images are greyscale. Since the dataset is only annotated with regards to one attribute (digit) we introduce a second mode of variation and corresponding labels. That is, we color each digit using 3 different colors, namely orange, purple and green. We then remove from the training set the images with labels (2, orange). The quantitative results for the unseen combination are reported in Table 1, while the qualitative in Figure 2. The proposed models are able to synthesize the unseen attribute combination and, as seen in the figure, they demonstrate considerable variation in the synthesized styles of the unseen combination.

The latent space of the model is projected into two-dimensions with T-SNE and visualized in Figure 3. The model disentangles the digit and color information and cre-

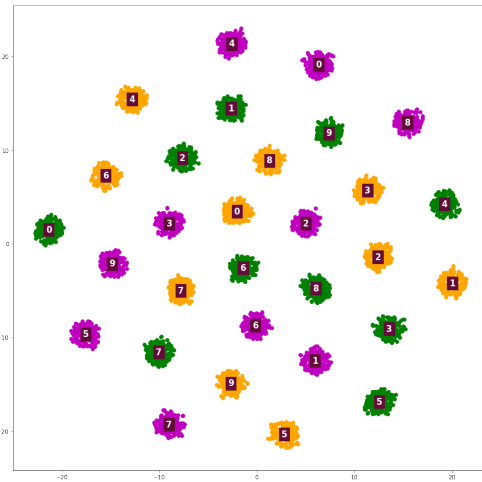


Figure 3. T-SNE (Maaten & Hinton, 2008) on the latent space of the proposed model during prediction time. Note that without explicit constraints, the latent space is naturally clustered based on the attribute labels.


ates class-specific clusters. As customary in image generation, linear interpolations (in the color attribute) are depicted in Figure 5a.

Fashion-MNIST: Similarly to MNIST, Fashion-MNIST consists of 60k training and 10k test images. We perform the same preprocessing as for MNIST and color each object orange, purple and green. We remove the images with labels (bag, orange) and present qualitative and quantitative results in Figure 2 and Table 1. Although the baselines fail to generate even the correct color (orange), we see that the proposed models can synthesize the missing combination successfully.

CelebA: The CelebA dataset contains over 200k facial images, annotated with 40 attributes. For our experiments, we chose two of the most balanced attributes, namely "Male" and "Smiling". Thereafter, we remove the images labeled (not "Male", "Smiling"). The images were resized to 64x64.

The results in Table 1 and Figure 2 indicate that, contrary to the baselines, the proposed models are able to generate diverse images of smiling women. To ensure that the model generalizes on imbalanced sets of attributes we conduct an experiment with the same hyperparameters using the ("Male", "Eyeglasses") attributes of the dataset. We note that the attribute "Eyeglasses" appears on less than 7% of the images. Further qualitative results on the ("Male", "Eyeglasses") protocol are presented in Figure 4b. For this experiment we remove the images with labels (not "Male", "Eyeglasses"). Interpolations between attributes are presented in Figure 5b.

Besides image synthesis, we evaluate the models on attribute transfer on CelebA. In this setting, given an input image and a set of target attributes (y_1, y_2), we modify the input image so that its attributes match y_1 and y_2 . Specifically, we encode an image and then decode it using different target



(a) Transfer of smile and gender

(b) Synthesized samples for the ("Male", "Eyeglasses") protocol

Figure 4. (a) Transfer of attributes in CelebA. The faces in the red rectangle are the original ones. For each row, we translate the image in the red square into the 3 remaining classes. (b) Synthesized samples for attributes gender and wearing glasses. The first row contains generated samples from all the classes (missing class in the red square). The bottom 3 rows contain generated samples just from the unseen class (woman with eyeglasses).

attributes. To evaluate the attribute transfer, we use the attribute classifiers as described above. Quantitative and qualitative results on smile and gender transfer are presented in Table 1 and Figure 4a.

6. The effect of modeling label interactions

In this section we experiment with the choice of the function M . In particular we investigate the importance of modeling the label interactions. To this end, we introduce a simpler version of the model that learns the mean as a linear combination of the labels as described in 10. We label this model LC-VAE. Hereafter, we compare LC-VAE to MLC-VAE on two custom benchmarks.


mixed-up MNIST: The first benchmark is created by introducing correlation between the labels of MNIST. We begin by coloring the MNIST dataset red and blue, as discussed in the previous experiments. Thereafter, we change the color label of each image if the digit is odd. The images of even digits kept their original color labels. The resulting dataset presents clear interaction between the digit and color labels. Similarly to the previous benchmarks, we remove the combination (2, red) from the training set.

The models are trained as in Section 4. The correlations between the labels, i.e., the color cannot be generated without explicit knowledge of the digit, is not captured by LC-VAE as demonstrated on Figure 6. On the contrary, the proposed models rely on the higher-order terms to capture such multiplicative interactions.

CelebA: We conduct a similar experiment on faces, where such label correlations occur naturally. To do that, we select the "Smiling" and "Mouth Slightly Open" attributes of CelebA. The interaction between the two labels is obvious, as most smiles in CelebA are labelled as "Mouth Slightly Open". We test the models on their ability to generate a smile with the mouth closed by excluding the combination



(a) Interpolation $6 \rightarrow 6 \rightarrow 6$



(b) Interpolation $WNS \rightarrow WS$ and $MNS \rightarrow MS$

Figure 5. Linear interpolation results on (a) MNIST, (b) CelebA. Specifically, in (a) the interpolation is on the color attribute while in (b) the interpolation is across the gender and the smiling attribute. Note that in (a) unseen colors emerge.

Model	MNIST		FASHION		CELEB-A		CELEB-A Transfer	
	digit	color	object	color	gender	smile	gender	smile
Random chance	10	33.3	10	33.3	50	50	50	50
cVAE	11.5	49	13.9	23	18	7.6	23.1	9
VampPrior	2.6	2.3	4	1.8	17.1	2.8	44.3	5.8
MLC-VAE-CP (Ours)	68.1	99.2	70.3	99.7	96.4	94	95	90.6
MLC-VAE-Tucker (Ours)	95.1	100	84.6	99.8	99.4	93.5	99.4	91.5

Table 1. Accuracy (%) of the attribute classifier for each model on the **unseen** attribute combinations of MNIST, FASHION, CELEB-A generation and CELEB-A transfer benchmarks. The proposed models outperform the baselines by a wide margin in all cases. The proposed model performs better overall when combined with the Tucker decomposition compared to CP. Random chance assumes a uniform selection from each attribute. The accuracy of the classifier on seen combinations can be found in the supplementary material.




Figure 6. Results on the mixed-up MNIST. The top two rows are synthesized from LC-VAE while the bottom two rows from MLC-VAE. The first and third rows are the digits labelled red and the second and fourth rows are digits labeled blue. We observe that only the MLC-VAE is able to generate the missing 2-red (in the red square).

("Smiling", not "Mouth Slightly Open") from the training set.



Figure 7. Results on the ("Smiling", "Mouth Slightly Open") protocol of CelebA. The top two rows are synthesized from LC-VAE while the bottom two rows from MLC-VAE. We observe that only the latter is able to generate the missing combination (Smiling, Closed mouth), while the former ends up generating faces with an open mouth.

The visual results on Figure 7 depict the improvement from

capturing the higher-order interactions. LC-VAE clearly generates the teeth, indicating that the mouth is open. On the contrary, the MLC-VAE is able to recover the unseen attribute combination.

7. Model exploration

For the purpose of model exploration, we investigate the following extensions to our model:

- Firstly, we probe the limits of the proposed framework in terms of seen attribute combinations. In other words, we investigate how many attribute combinations must be present in the training set, so that generation of all combination is still possible.
- Secondly, we benchmark our model on a dataset with 3 attributes.
- Lastly, we extend the model to handle unlabelled attributes.

7.1. Multiple missing combinations

In order to study how the performance of the proposed model scales with the number of missing attribute combinations, we ablate the model on the MNIST dataset. In particular, we create 5 different training sets, that are missing 3, 6, 9, 15 and 20 attribute combinations respectively. We train 5 models and report the average classification accuracy of the generated unseen combinations. The result on Table 2 indicate that the model performs reasonably well even after removing 9 attribute combinations from the training set.

Accuracy	Number of missing combinations				
	3/30	6/30	9/30	15/30	20/30
Digit	88.1	94.5	90	50.4	15.8
Color	93.3	82.7	88.4	87.2	81.2

Table 2. Effect of number of unseen combinations to the accuracy (%). The number of unseen combinations is explored (w.r.t. the total number of 30 combinations). The accuracy of the attribute classifier on unseen attribute combinations is reported. Notice that even when 9 combinations are unseen, there is no significant loss in the attribute classification.

7.2. Multiple attributes

The presented model can be extended to handle an arbitrary number of attributes. The general form is as follows:

$$\begin{aligned}
 M(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N) = & \\
 & \sum_{i=1}^N \mathbf{W}^{[i]} \mathbf{y}_i + \sum_{j=2}^N \sum_{i=1}^{j-1} \mathbf{W}^{[ij]} \times_2 \mathbf{y}_i \times_3 \mathbf{y}_j + \\
 & \sum_{h=3}^N \sum_{j=2}^{h-1} \sum_{i=1}^{j-1} \mathbf{W}^{[ijh]} \times_2 \mathbf{y}_i \times_3 \mathbf{y}_j \times_4 \mathbf{y}_h + \dots \\
 & + \mathbf{W}^{[1\dots N]} \times_2 \mathbf{y}_1 \times_3 \mathbf{y}_2 \cdots \times_N \mathbf{y}_N \quad (15)
 \end{aligned}$$

In order to explore generation of more than two sets of attributes, we utilize the Morpho-MNIST (Castro et al., 2019) dataset. This dataset is an extension of MNIST that is further transformed and annotated with regards to stroke thickness, swelling and fractures. Our 3-attribute benchmark consists of the thick and thin subsets of Morpho-MNIST, colored according to our previous benchmarks.

Based on 15, the model for 3 attributes is:

$$\begin{aligned}
 M(\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3) = & \\
 & \mathbf{W}^{[1]} \mathbf{y}_1 + \mathbf{W}^{[2]} \mathbf{y}_2 + \mathbf{W}^{[3]} \mathbf{y}_3 + \\
 & \mathbf{W}^{[12]} \times_2 \mathbf{y}_1 \times_3 \mathbf{y}_2 + \mathbf{W}^{[13]} \times_2 \mathbf{y}_1 \times_3 \mathbf{y}_3 + \\
 & \mathbf{W}^{[23]} \times_2 \mathbf{y}_2 \times_3 \mathbf{y}_3 + \mathbf{W}^{[123]} \times_2 \mathbf{y}_1 \times_3 \mathbf{y}_2 \times_4 \mathbf{y}_3, \quad (16)
 \end{aligned}$$

where third order tensors $\mathbf{W}^{[12]}$, $\mathbf{W}^{[23]}$, $\mathbf{W}^{[13]}$ capture all possible pairwise interactions, while $\mathbf{W}^{[123]}$ is a fourth order tensor that captures the triplet-wise interaction between all attributes. The synthetic results in Figure 8 indicate that our model can effectively handle multiple attributes.

7.3. Unlabeled attributes

In this section, we extend the model to handle unlabelled attributes. In particular, we assume prior knowledge of only the number of classes of the unlabelled attribute (e.g., 2 different colors). In this setting, the data have labels corresponding to one attribute but not the other. A second encoder



Figure 8. Samples synthesized on 3-attribute MNIST. The samples in the red boxes are combinations not seen during training (there are 3 combinations not seen in total). Nevertheless, the proposed method is able to synthesize images of all three unseen combinations.



Figure 9. Results when color labels are not provided. Notice that the proposed method can still separate the attributes and synthesize the unseen combinations.

network is utilized to infer the unobserved attribute’s class probabilities. We utilize the Gumbel-Softmax distribution (Maddison et al., 2016; Jang et al., 2016) to perform the reparametrization trick and train the model.

We experiment with two-color MNIST (green/orange), without knowledge of the color label of each sample. Qualitative results are presented in Figure 9. Even though the images do not have color labels, the model is able to disentangle color from digit and generate the unseen attribute combinations (0-orange and 1-green).

8. Conclusion

In this work, we study the ability of conditional generative models to synthesize images of unseen attribute combinations. In particular, we focus on the setting where a number of attribute combinations are absent from the training set. Our goal is to learn a different conditional distribution for each class (i.e., attribute combination). The mean of each conditional distribution is learned through the proposed multilinear conditioning framework. We demonstrate the efficacy of our method on 3 different datasets and display the importance of modeling the interactions between labels. In the future, we want to explore further the multiplicative interactions of an arbitrarily large number of attributes, in the supervised and unsupervised setting, by using coupled decompositions and introducing weight sharing.

References

- Andreux, M. and Mallat, S. Music generation and transformation with moment matching-scattering inverse networks. In *ISMIR*, pp. 327–333, 2018.
- Arora, S. and Zhang, Y. Do gans actually learn the distribution? an empirical study. *arXiv preprint arXiv:1706.08224*, 2017.
- Bozkurt, A., Esmaili, B., Brooks, D. H., Dy, J. G., and van de Meent, J.-W. Evaluating combinatorial generalization in variational autoencoders. *arXiv preprint arXiv:1911.04594*, 2019.
- Burgess, C. P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., and Lerchner, A. Understanding disentangling in β -vae. *arXiv preprint arXiv:1804.03599*, 2018.
- Castro, D. C., Tan, J., Kainz, B., Konukoglu, E., and Glocker, B. Morpho-MNIST: Quantitative assessment and diagnostics for representation learning. *Journal of Machine Learning Research*, 20, 2019.
- Chen, R. T., Li, X., Grosse, R. B., and Duvenaud, D. K. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pp. 2610–2620, 2018.
- Chen, X., Kingma, D. P., Salimans, T., Duan, Y., Dhariwal, P., Schulman, J., Sutskever, I., and Abbeel, P. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
- De Lathauwer, L., De Moor, B., and Vandewalle, J. A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Dilokthanakul, N., Mediano, P. A., Garnelo, M., Lee, M. C., Salimbeni, H., Arulkumaran, K., and Shanahan, M. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*, 2016.
- Du, Y. and Mordatch, I. Implicit generation and generalization in energy-based models. In *Advances in neural information processing systems (NIPS)*, 2019.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Grønbech, C. H., Vording, M. F., Timshel, P. N., Sønderby, C. K., Pers, T. H., and Winther, O. scvae: Variational autoencoders for single-cell gene expression datas. *bioRxiv*, pp. 318295, 2018.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Jiang, Z., Zheng, Y., Tan, H., Tang, B., and Zhou, H. Variational deep embedding: An unsupervised and generative approach to clustering. *arXiv preprint arXiv:1611.05148*, 2016.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Khemakhem, I., Kingma, D., Monti, R., and Hyvarinen, A. Variational autoencoders and nonlinear ica: A unifying framework. In *International Conference on Artificial Intelligence and Statistics*, pp. 2207–2217, 2020.
- Kim, H. and Mnih, A. Disentangling by factorising. *arXiv preprint arXiv:1802.05983*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014.
- Klys, J., Snell, J., and Zemel, R. Learning latent subspaces in variational autoencoders. In *Advances in Neural Information Processing Systems*, pp. 6444–6454, 2018.
- Kolda, T. G. and Bader, B. W. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- Kossaifi, J., Panagakis, Y., Anandkumar, A., and Pantic, M. Tensorly: Tensor learning in python. *The Journal of Machine Learning Research*, 20(1):925–930, 2019.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

- Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pp. 4114–4124, 2019.
- Maaten, L. v. d. and Hinton, G. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov): 2579–2605, 2008.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. 2017.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Radford, A., Jozefowicz, R., and Sutskever, I. Learning to generate reviews and discovering sentiment. In *ICLR*, 2018.
- Razavi, A., van den Oord, A., and Vinyals, O. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*, pp. 14837–14847, 2019.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Ridgeway, K. and Mozer, M. C. Learning deep disentangled embeddings with the f-statistic loss. In *Advances in Neural Information Processing Systems*, pp. 185–194, 2018.
- Siddharth, N., Paige, B., Van de Meent, J.-W., Desmaison, A., Goodman, N., Kohli, P., Wood, F., and Torr, P. Learning disentangled representations with semi-supervised deep generative models. In *Advances in Neural Information Processing Systems*, pp. 5925–5935, 2017.
- Sidiropoulos, N. D., De Lathauwer, L., Fu, X., Huang, K., Papalexakis, E. E., and Faloutsos, C. Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing*, 65(13):3551–3582, 2017.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pp. 3483–3491, 2015.
- Tang, Y., Salakhutdinov, R., and Hinton, G. Tensor analyzers. In *International conference on machine learning*, pp. 163–171, 2013.
- Tomczak, J. M. and Welling, M. Vae with a vampprior. *arXiv preprint arXiv:1705.07120*, 2017.
- Tucker, L. R. Implications of factor analysis of three-way matrices for measurement of change. *Problems in measuring change*, 15:122–137, 1963.
- Van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems (NIPS)*, pp. 4790–4798, 2016.
- Van Den Oord, A., Vinyals, O., et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pp. 6306–6315, 2017.
- Wang, M., Panagakis, Y., Snape, P., and Zafeiriou, S. P. Disentangling the modes of variation in unlabelled data. *IEEE transactions on pattern analysis and machine intelligence*, 40(11):2682–2695, 2017.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Zhao, S., Ren, H., Yuan, A., Song, J., Goodman, N., and Ermon, S. Bias and generalization in deep generative models: An empirical study. In *Advances in Neural Information Processing Systems*, pp. 10792–10801, 2018.