

S3VAE: Self-Supervised Sequential VAE for Representation Disentanglement and Data Generation

Yizhe Zhu^{1,2}, Martin Renqiang Min¹, Asim Kadav¹, Hans Peter Graf¹
 yizhe.zhu@rutgers.edu, {renqiang, asim, hpg}@nec-labs.com
¹NEC Labs America, ²Department of Computer Science, Rutgers University

Abstract

We propose a sequential variational autoencoder to learn disentangled representations of sequential data (e.g., videos and audios) under self-supervision. Specifically, we exploit the benefits of some readily accessible supervisory signals from input data itself or some off-the-shelf functional models and accordingly design auxiliary tasks for our model to utilize these signals. With the supervision of the signals, our model can easily disentangle the representation of an input sequence into static factors and dynamic factors (i.e., time-invariant and time-varying parts). Comprehensive experiments across videos and audios verify the effectiveness of our model on representation disentanglement and generation of sequential data, and demonstrate that, our model with self-supervision performs comparable to, if not better than, the fully-supervised model with ground truth labels, and outperforms state-of-the-art unsupervised models by a large margin.

1. Introduction

Representation learning is one of the essential research problems in machine learning and computer vision [5]. Real-world sensory data such as videos, images, and audios are often in the form of high dimensions. Representation learning aims to map these data into a low-dimensional space to make it easier to extract semantically meaningful information for downstream tasks such as classification and detection. Recent years have witnessed rising interests in disentangled representation learning, which tries to separate the underlying factors of observed data variations such that each factor exclusively interprets one type of semantic attributes of sensory data. The representation of sequential data is expected to be disentangled into time-varying factors and time-invariant factors. For video data, the identity of a moving object in a video is regarded as time-invariant factors, and the motion in each frame is considered as time-varying ones [36]. For speech data, the




Figure 1. Self-supervision and regularizations enforce the latent variable of our sequential VAE to be disentangled into a static representation z_f and a dynamic representation z_t .

representations of the timbre of speakers and the linguistic contents are expected to be disentangled [27]. There are several benefits of learning disentangled representations. First, the models that produce disentangled representations are more explainable. Second, disentangled representations make it easier and more efficient to manipulate data generation, which has potential applications in entertainment industry, training data synthesis [53, 54] and several downstream tasks [32, 19, 18, 43, 55].

Despite the vast amount of works [24, 33, 6, 15, 16, 7, 30] on disentangled representations of static data (mainly image data), fewer works [27, 36, 23, 46] have explored representation disentanglement for sequential data generation. For unsupervised models, FHVAE [27] and DSVAE [36] elaborately designed model architectures and factorized latent variables into static and dynamic parts. These models may well handle simple data forms such as synthetic animation data but fail when dealing with realistic ones as we will show later. Besides, as pointed out in [37], unsupervised representation disentanglement is impossible without inductive biases. Without any supervision, the performance of disentanglement can hardly be guaranteed and greatly depends on the random seed and the dimensionality of latent vectors set in the models. On the other hand, several works [23, 46] resort to utilizing label information or attribute annotation as strong supervision for disentanglement. For instance, VideoVAE [23] leveraged holistic attributes to constrain latent variables. Nevertheless, the

costly annotation of data is essential for these models and prevents them from being deployed to most real-world applications, in which a tremendous amount of unlabeled data is available.

To alleviate the drawbacks of both unsupervised and supervised models discussed above, this work tackles representation disentanglement for sequential data generation utilizing self-supervision. In self-supervised learning, various readily obtainable supervisory signals have been explored for representation learning of images and videos, employing auxiliary data such as the ambient sounds in videos [41, 3], the egomotion of cameras [1, 29], the geometry cue in 3D movies [20], and off-the-shelf functional models for visual tracking [51], and optical flow [42, 50]. However, how self-supervised learning benefits representation disentanglement of sequential data has barely been explored.

In this paper, we propose a sequential variational autoencoder (VAE), a recurrent version of VAE, for sequence generation. In the latent space, the representation is disentangled into time-invariant and time-varying factors. We address the representation disentanglement by exploring intrinsic supervision signals, which can be readily obtained from both data itself and off-the-shelf methods, and accordingly design a series of auxiliary tasks. Specifically, on one hand, to exclude dynamic information from time-invariant variable, we exploit the temporal order of the sequential data and expect the time-invariant variable of the temporally shuffled data to be close to if not the same as that of the original data. On the other hand, the time-varying variable is expected to contain dynamic information in different modalities. For video data, we allow it to predict the location of the largest motion in every frame, which can be readily inferred from optical flow. For audio data, the volume in each segment as an intrinsic label is leveraged as the supervisory signal. To further encourage the representation disentanglement, the mutual information between static and dynamic variables are minimized as an extra regularization.

To the best of our knowledge, this paper is the first work to explicitly use auxiliary supervision to improve the representation disentanglement for sequential data. Extensive experiments on representation disentanglement and sequence data generation demonstrate that, with these multiple freely accessible supervisions, our model dramatically outperforms unsupervised learning-based methods and even performs better than fully-supervised learning-based methods in several cases.

2. Related Work

Disentangled Sequential Data Generation With the success of deep generative models, recent works [24, 33, 6, 7, 30] resort to variational autoencoders (VAEs) [34] and generative adversarial networks (GANs) [22] to learn a dis-

entangled representation. Regularizations are accordingly designed. β -VAE [24] imposed a heavier penalty on the KL divergence term for a better disentanglement learning. Follow-up researches [33, 6] derived a Total Correlation (TC) from the KL term, and highlights this TC term as the key factor in disentangled representation learning. In InfoGAN [7], the disentanglement of a latent code c is achieved by maximizing a mutual information lower-bound between c and the generated sample \tilde{x} .

Several works involving disentangled representation have been proposed for video prediction. Villegas *et al.* [48] and Denton *et al.* [12] designed two networks to encode pose and content separately at each timestep. Unlike video prediction, video generation from priors, which we perform in this work, is a much harder task since no frame is available for appearance and motion modeling in the generation phase.

To handle video generation, VAEs are extended to a recurrent version [17, 4, 10]. However, these models do not explicitly consider static and dynamic representation disentanglement and fail to perform manipulable data generation. More recently, several works have proposed VAEs with factorized latent variables. FHVAE [27] presented a factorized hierarchical graphical model that imposes sequence-dependent priors and sequence-independent priors to different sets of latent variables in the context of speech data, but did not take advantage of the sequential prior. Combining the merits of recurrent VAE and FHVAE, DSVAE [36] is capable of disentangling latent factors by factorizing them into time-invariant and time-dependent parts and applies an LSTM sequential prior to keep a better sequential consistency for sequence generation. Although with elaborately designed complex architectures, these models may only perform decently on representation disentanglement of simple data, the disentanglement performance degrades rapidly when the complexity of data increases. In contrast, our work explores both model and regularization designs for representation disentanglement and sequential data generation. Our model fully factorizes the latent variables to time-invariant and time-varying parts, and both the posterior and the prior of the time-varying variable are modeled by LSTM for dynamic consistency. The auxiliary tasks with readily accessible supervisory signals are designed to regularize and encourage representation disentanglement.

Self-Supervised Learning The concept of self-supervised learning traces back to the autoencoder [25], which uses the input itself as supervision to learn the representation. Denoising autoencoder [49] makes the learned representations robust to noise and partial corruption of the input pattern by adding noise to the input. Recent years have witnessed the booming interest in self-supervised learning. The sources of supervisory signals can be roughly categorized into three classes. (a) *Intrinsic labels*: Doersch *et al.* [13] explored

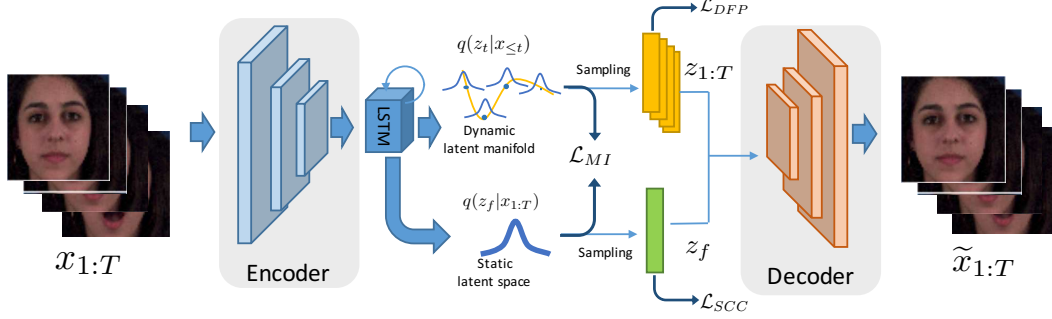


Figure 2. The framework of our proposed model in the context of video data. Each frame of a video $\mathbf{x}_{1:T}$ is fed into an encoder to produce a sequence of visual features, which is then passed through an LSTM module to obtain the manifold posterior of a dynamic latent variable $\{q(z_t|\mathbf{x}_{\leq t})\}_{t=1}^T$ and the posterior of a static latent variable $q(z_f|\mathbf{x}_{1:T})$. The static and dynamic representations z_f and $z_{1:T}$ are sampled from the corresponding posteriors and concatenated to be fed into a decoder to generate reconstructed sequence $\tilde{\mathbf{x}}_{1:T}$. Three regularizers are imposed on dynamic and static latent variables to encourage the representation disentanglement.

the use of spatial context in images, and Noroozi *et al.* [40] trained a model to solve Jigsaw puzzles as a pretext task. Several works [52, 35] showed that colorizing a gray-scale photograph can be utilized as a powerful pretext task for visual understanding. Temporal information of video is another readily accessible supervisory signal. [38] trained a model to determine whether a sequence of frames from a video is in the correct temporal order and [31] made the model learn to arrange the permuted 3D spatiotemporal crops. (b) **Auxiliary data:** Agrawal *et al.* [1] and Jayaraman *et al.* [29] exploited the freely available knowledge of camera motion as a supervisory signal for feature learning. Ambient sounds in videos [41, 3] are used as a supervisory signal for learning visual models. The geometry cue in 3D movies [20] is utilized for visual representation learning. (c) **Off-the-shelf tools:** Wang *et al.* [51] leveraged the visual consistency of objects from a visual tracker in the video clips. [42] used segments obtained by motion-based segmentation based on optical flow as pseudo ground truth for the single-frame object segmentation. Instead of learning visual features as in aforementioned methods, this work aims to achieve static and dynamic representation disentanglement for sequential data such as video and speech. To this end, we leverage supervisory signals from intrinsic labels to regularize the static representation and off-the-shelf tools to regularize the dynamic representation.

3. Sequential VAE Model

We start by introducing some notations and the problem definition. $\mathcal{D} = \{\mathbf{X}^i\}^N$ is given as a dataset that consists of M i.i.d. sequences, where $\mathbf{X} \equiv \mathbf{x}_{1:T} = (x_1, x_2, \dots, x_T)$ denotes a sequence of T observed variables, such as a video of T frames or an audio of T segments. We propose a sequential variational autoencoder model, where the sequence is assumed to be generated from latent variable \mathbf{z} and \mathbf{z} is factorized into two disentangled variables: the time-invariant (or static) variable z_f and the time-varying (or dynamic)

variables $z_{1:T}$.

Priors The prior of z_f is defined as a standard Gaussian distribution: $z_f \sim \mathcal{N}(0, 1)$. The time-varying latent variables $z_{1:T}$ follow a sequential prior $z_t | z_{<t} \sim \mathcal{N}(\boldsymbol{\mu}_t, \text{diag}(\boldsymbol{\sigma}_t^2))$, where $[\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t] = \phi_R^{\text{prior}}(z_{<t})$, $\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t$ are the parameters of the prior distribution conditioned on all previous time-varying latent variables $z_{<t}$. The model ϕ_R^{prior} can be parameterized as a recurrent network, such as LSTM [26] or GRU [9], where the hidden state is updated temporally. The prior of \mathbf{z} can be factorized as:

$$p(\mathbf{z}) = p(z_f)p(z_{1:T}) = p(z_f) \prod_{t=1}^T p(z_t | z_{<t}). \quad (1)$$

Generation The generating distribution of time step t is conditioned on z_f and z_t : $\mathbf{x}_t | z_f, z_t \sim \mathcal{N}(\boldsymbol{\mu}_{x,t}, \text{diag}(\boldsymbol{\sigma}_{x,t}^2))$, where $[\boldsymbol{\mu}_{x,t}, \boldsymbol{\sigma}_{x,t}] = \phi^{\text{Decoder}}(z_f, z_t)$ and the decoder ϕ^{Decoder} can be a highly flexible function such as a deconvolutional neural network [39]. The complete generative model can be formalized by the factorization:

$$p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T}, z_f) = p(z_f) \prod_{t=1}^T p(\mathbf{x}_t | z_f, z_t) p(z_t | z_{<t}) \quad (2)$$

Inference Our sequential VAE uses variational inference to approximate posterior distributions:

$$z_f \sim \mathcal{N}(\boldsymbol{\mu}_f, \text{diag}(\boldsymbol{\sigma}_f^2)), \quad z_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \text{diag}(\boldsymbol{\sigma}_t^2)), \quad (3)$$

where $[\boldsymbol{\mu}_f, \boldsymbol{\sigma}_f] = \psi_f^{\text{Encoder}}(\mathbf{x}_{1:T})$ and $[\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t] = \psi_R^{\text{Encoder}}(\mathbf{x}_{\leq t})$. The static variable ψ_f^{Encoder} is conditioned on the whole sequence while the dynamic variable is inferred by a recurrent encoder ψ_R^{Encoder} and only conditioned on the previous frames. Our inference model is factorized as:

$$q(\mathbf{z}_{1:T}, z_f | \mathbf{x}_{1:T}) = q(z_f | \mathbf{x}_{1:T}) \prod_{t=1}^T q(z_t | \mathbf{x}_{\leq t}) \quad (4)$$

Learning The objective function of sequential VAE is a timestep-wise negative variational lower bound:

$$\mathcal{L}_{VAE} = \mathbb{E}_{q(\mathbf{z}_{1:T}, \mathbf{z}_f | \mathbf{x}_{1:T})} \left[- \sum_{t=1}^T \log p(\mathbf{x}_t | \mathbf{z}_f, \mathbf{z}_t) \right] + \text{KL}(q(\mathbf{z}_f | \mathbf{x}_{1:T}) || p(\mathbf{z}_f)) + \sum_{t=1}^T \text{KL}(q(\mathbf{z}_t | \mathbf{x}_{\leq t}) || p(\mathbf{z}_t | \mathbf{z}_{<t})) \quad (5)$$

The schematic representation of our model is shown in Figure 2. Note that DSVAE also proposes a sequential VAE with disentangled representation, but it either independently infers \mathbf{z}_t only based on the frame of each time-step without considering the continuity of dynamic variables and thus may generate inconsistent motion, or assumes the variational posterior of \mathbf{z}_t depends on \mathbf{z}_f , implying that the variables are still implicitly entangled. In contrast, we model both the prior and the posterior of \mathbf{z}_t by recurrent models independently, resulting in consistent dynamic information in synthetic sequences, and ensure full disentanglement of \mathbf{z}_f and \mathbf{z}_t by posterior factorization.

4. Self-Supervision and Regularization

Without any supervision, there is no guarantee that the time-invariant representation \mathbf{z}_f and the time-varying representation \mathbf{z}_t are disentangled. In this section, we introduce a series of auxiliary tasks on the different types of representations as the regularization of our sequential VAE to achieve the disentanglement, where readily accessible supervisory signals are leveraged.

4.1. Static Consistency Constraint

To encourage the time-invariant representation \mathbf{z}_f to exclude any dynamic information, we expect that \mathbf{z}_f changes little when varying dynamic information dramatically. To this end, we shuffle the temporal order of frames to form a shuffled sequence. Ideally, the static factors of the original sequence and shuffled sequence should be very close, if not equal, to one another. However, directly minimizing the distance of these two static factors will lead to very trivial solutions, e.g., the static factors of all sequences converge to the same value and do not contain any meaningful information. Thus, we randomly sample another sequence as the negative sample of the static factor. With a triple of static factors, we introduce a triplet loss as follows:

$$\mathcal{L}_{SCC} = \max \left(D(\mathbf{z}_f, \mathbf{z}_f^{pos}) - D(\mathbf{z}_f, \mathbf{z}_f^{neg}) + m, 0 \right), \quad (6)$$

where \mathbf{z}_f , \mathbf{z}_f^{pos} and \mathbf{z}_f^{neg} are the static factors of the anchor sequence, the shuffled sequence as the positive data, and another randomly sampled video as the negative data,




Figure 3. The pseudo label generation for video datasets. (a) The left image is the input frame and the right image is the corresponding optical flow map that is split by a grid. (b) Three distances are used as the dynamic signals for the face dataset.

$D(\cdot, \cdot)$ denotes the Euclidean distance and m is the margin, set to 1. This regularization makes \mathbf{z}_f preserve meaningful static information to a certain degree while excluding dynamic information.

4.2. Dynamic Factor Prediction

To encourage the dynamic representation \mathbf{z}_t to carry adequate and correct time-dependent information of each time-step, we exploit dynamic information-related signals from off-the-shelf tools for different types of sequential data and accordingly design the auxiliary tasks as the regularization imposed on \mathbf{z}_t . We have the loss $\mathcal{L}_{DFP} = \mathcal{L}(\phi_a(\mathbf{z}_t), y)$, where $\mathcal{L}(\cdot, \cdot)$ can be either cross-entropy loss or mean squared error loss according to the designed auxiliary task, $\phi_a(\cdot)$ is a network for dynamic factor prediction and y contains supervisory signals.

Video Data The dynamic representation \mathbf{z}_t can be learned by forcing it to predict the dynamic factors of videos. Motivated by this, we expect the location of the largest motion regions can be accurately predicted based on \mathbf{z}_t . To this end, the optical flow maps of video are first obtained by commonly used functional model FlowNet2 [28] and then split into patches by grid, as shown in Figure 3.a. We compute the average of motion magnitudes for every patch and use the indices of patches with the top-k largest values as the pseudo label for prediction. For this task, $\phi_a(\cdot)$ is implemented with two fully-connected layers and a softmax layer.

Apart from the optical flow, some freely obtainable data-specific signals can be exploited. For a human face dataset, the landmark of each frame can be readily detected and considered as a supervision for dynamic factors. We obtain the landmark from an off-the-shelf landmark detector [14]. To keep our model concise and efficient, we only leverage the distance between the upper and lower eyelids as well as the distance between the upper and lower lips in each frame as the dynamic signal, as shown in Figure 3.b. Here, $\phi_a(\cdot)$ consists of two fully-connected layers to regress the three distances. We observe that our model can easily capture dynamic motions under this simple supervision.

Audio Data For the audio dataset, we consider the volume as time-dependency factor and accordingly design an auxiliary task, where z_t is forced to predict if the speech is silent or not in each segment. The pseudo label is readily obtained by setting a magnitude threshold on the volume of each speech segment. $\phi_a(\cdot)$ consists of two fully-connected layers and performs a binary classification.

4.3. Mutual Information Regularization

Forcing the time-varying variable z_t to predict dynamic factors can guarantee that z_t contains adequate dynamic information, but this fails to guarantee that z_t excludes the static information. Therefore, we introduce the mutual information between static and dynamic variables as a regulator \mathcal{L}_{MI} . The mutual information is a measure of the mutual dependence between two variables. By minimizing \mathcal{L}_{MI} , we encourage the information in these two variables are mutually exclusive. The mutual information is formally defined as the KL divergence of the joint distribution to the product of marginal distribution of each variable. We have

$$\begin{aligned} \mathcal{L}_{MI}(z_f, z_{1:T}) &= \sum_{t=1}^T \text{KL}(q(z_f, z_t) || q(z_f)q(z_t)) \\ &= \sum_{t=1}^T [H(z_f) + H(z_t) - H(z_f, z_t)], \end{aligned} \quad (7)$$

where $H(\cdot) = -\mathbb{E}_{q(z)}[\log(q(\cdot))] = -\mathbb{E}_{q(z_f, z_t)}[\log(q(\cdot))]$. The expectation can be estimated by the mini-batch weighted sampling estimator [6],

$$\mathbb{E}_{q(z)}[\log q(z_n)] \approx \frac{1}{M} \sum_{i=1}^M \left[\log \sum_{j=1}^M q(z_n(x_i)|x_j) - \log(NM) \right], \quad (8)$$

for $z_n = z_f, z_t$ or (z_f, z_t) , where N and M are the data size and the minibatch size, respectively.

4.4. Objective Function

Overall, our objective can be considered as the sequential VAE loss with a series of self-supervision and regularization:

$$\mathcal{L} = \mathcal{L}_{VAE} + \lambda_1 \mathcal{L}_{SCC} + \lambda_2 \mathcal{L}_{DFP} + \lambda_3 \mathcal{L}_{MI}, \quad (9)$$

where λ_1, λ_2 and λ_3 are balancing factors.

5. Experiments

To comprehensively validate the effectiveness of our proposed model with self-supervision, we conduct experiments on three video datasets and one audio dataset. With these four datasets, we cover different modalities from video to audio. In video domain, a large range of motions are covered from large character motions (e.g., walking, stretching) to subtle facial expressions (e.g., smiling, disgust).

5.1. Experiments on Video Data

We present an in-depth evaluation on two problems, tested on three different datasets and employing a large variety of metrics.

5.1.1 Datasets

Stochastic Moving MNIST is introduced by [11] and consists of sequences of 15 frames of size 64×64 , where two digits from MNIST dataset move in random directions. We randomly generate 6000 sequences, 5000 of which are used for training and the rest are for testing.

Sprite [36] contains sequences of animated cartoon characters with 9 action categories: walking, casting spells and slashing with three viewing angles. The appearance of characters are fully controlled by four attributes, i.e., the color of skin, tops, pants, and hair. Each of the attributes categories contains 6 possible variants, therefore it results in totally $6^4 = 1296$ unique characters, 1000 of which are used for training and the rest for testing. Each sequence contains 8 frames of size 64×64 .

MUG Facial Expression [2] consists of 3528 videos with 52 actors performing 6 different facial expressions: anger, fear, disgust, happiness, sadness, and surprise. Each video composes of 50 to 160 frames. As suggested in MoCoGAN [47], we crop the face regions, resize video to 64×64 , and randomly sample a clip of 15 frames in each video. The 75% of dataset is used for training and the rest for testing.

5.1.2 Representation Swapping

We first perform the representation swapping and compare our method with DSVAE, a disentangled VAE model, as well as MonkeyNet [45], a state-of-the-art deformable video generation model. Suppose two real videos are given for motion information and appearance information, denoted as V_m and V_a . Our method and DSVAE perform video generation based on the z_f from V_a and $z_{1:T}$ from V_m . For MonkeyNet, the videos are generated by deforming the first frame of V_a based on motion in V_m . The synthetic videos are expected to preserve the appearance in V_a and the motion in V_m . The qualitative comparisons on three datasets are shown in Figure 4.

For SMMNIST, the generated videos of our model can preserve the identity of digits while consistently mimic the motion of the provided video. However, DSVAE can hardly preserve the identity. For instance, it mistakenly changes the digit “9” to “6”. We observe that MonkeyNet can hardly handle the case with multiple objects like SMMNIST, because the case does not meet its implicit assumption of only one object moving in the video.

For Sprite, DSVAE generates blurry videos when the characters in V_a and V_m have opposite directions, indicat-

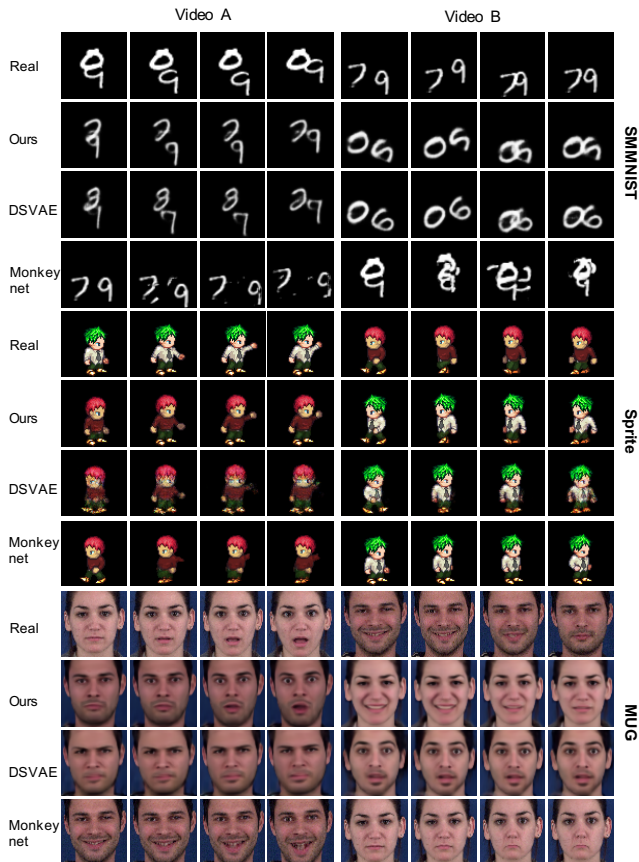


Figure 4. Representation swapping on SMMNIST, Sprite and MUG datasets. In each panel, we show two real videos as well as the generated videos by swapping z_f and $z_{1:T}$, from our model and two competing models: DSVAE and MonkeyNet. Each column is supposed to have the same motion.

ing it fails to encode the direction information in the dynamic variable. Conversely, our model can generate videos with the appearance of the character in V_a and the same action and direction of the character in V_m , due to the guidance from optical flow. The characters in the generated videos of Monkeynet fail to follow the pose and action in V_m , and many artifacts appear. E.g., an arm-like blob appears in the back of the character in the left panel.

For MUG, the generated video of DSVAE can hardly preserve both the appearance in V_a and the facial expression in V_m . For example, the person in the right has a mixed appearance characteristic, indicating z_f and z_t are entangled. Due to the deformation scheme of generation, MonkeyNet fails to handle the case where the faces in two videos are not well aligned. For instance, forcing the man with a smile to be fear results in unnatural expression. On the contrary, our model disentangles z_f and z_t , supported by the realistic expressions on different faces in generated videos.



Figure 5. Unconditional video generation on MUG. The upper and lower panels show the qualitative results of our model and MoCoGAN, respectively.




Figure 6. Randomly sampled frames for each dataset.

5.1.3 Video Generation

Quantitative Results We compute the quantitative metrics of our model with and without self-supervision and regularization, denoted as *full model* and *baseline*, as well as two competing methods: DSVAE and MoCoGAN. All these methods are comparable as no ground truth labels are used to benefit representation disentanglement. Besides, the results of our baseline with full supervision from human-annotation *baseline-sv* are also provided as a reference.

To demonstrate the ability of a model on the representation disentanglement, we use the classification accuracy *Acc* [36], which measures the ability of a model to preserve a specific attributes when generating a video given the corresponding representation or label. To measure how diverse and realistic videos a model can generate, three metrics are used: *IS* [44], Intra-Entropy $H(y|x)$ [23] and Inter-Entropy $H(y)$ [23]. All metrics utilize a pretrained classifier based on the real videos and ground truth attributes. See Appendix for the detailed definitions of the metrics. The results are shown in Table 1.

For representation disentanglement, we consider generating videos with a given z_f inferred from a real video and randomly sampled $z_{1:T}$ from the prior $p(z_{1:T})$ for SMMNIST. We then check if the synthetic video contains the same digits as the real video by the pretrained classifier. For MUG, we evaluate the ability of a model to preserve the facial expression by fixing $z_{1:T}$ and randomly sampled

Table 1. Quantitatively performance comparison on SMMNIST, Sprite and MUG datasets. High values are expected for Acc , $H(y)$ and IS , while for $H(y|x)$, the lower values are better. The results of our model with supervision of ground truth labels *baseline-sv** are shown as a reference.

Methods	SMMNIST				Sprite				MUG			
	Acc	IS	$H(y x)$	$H(y)$	Acc	IS	$H(y x)$	$H(y)$	Acc	IS	$H(y x)$	$H(y)$
MoCoGAN	74.55%	4.078	0.194	0.191	92.89%	8.461	0.090	2.192	63.12%	4.332	0.183	1.721
DSVAE	88.19%	6.210	0.185	2.011	90.73%	8.384	0.072	2.192	54.29%	3.608	0.374	1.657
<i>baseline</i>	90.12%	6.543	0.167	2.052	91.42%	8.312	0.071	2.190	53.83%	3.736	0.347	1.717
<i>full model</i>	95.09%	7.072	0.150	2.106	99.49%	8.637	0.041	2.197	70.51%	5.136	0.135	1.760
<i>baseline-sv*</i>	92.18%	6.845	0.156	2.057	98.91%	8.741	0.028	2.196	72.32%	5.006	0.129	1.740

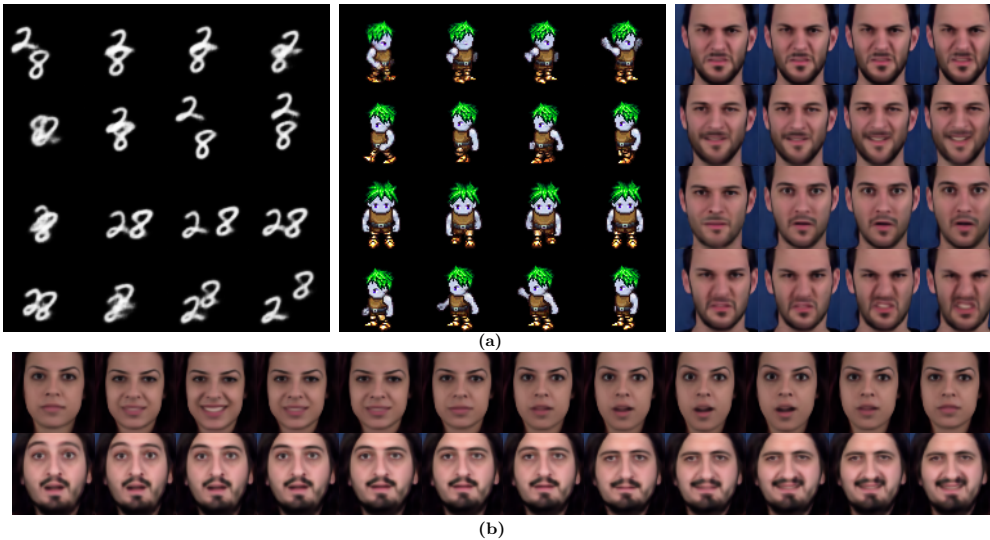


Figure 7. Controlled video generation. (a) Video generation controlled by fixing the static variable and randomly sampling dynamic variables from the prior $p(z_{1:T})$. All sequences share a same identity but with different motions for each dataset. (b) Video generation with changed facial expressions. Expression is changed from smile to surprise and from surprise to disgust in two sequences, respectively. We control it by transferring the dynamic variables smoothly.

z_f from the prior $p(z_f)$. For Sprite, since the ground truth of both actions and appearance attributes are available, we evaluate the ability of preserving both static and dynamic representations, and report the average scores. It's evident that our *full model* consistently outperforms all competing methods. For SMNIST, we observe that MoCoGAN have poor ability to correctly generate the digits with given labels while our *full model* can generate correctly digits, reflected by the high Acc . Note that our *full model* achieves 99.49% Acc on Sprite, indicating the z_f and z_t are greatly disentangled. Besides, *full model* significantly boosts the performance of *baseline*, especially in MUG where more realistic data is contained, the performance gets giant boost from 53.83% to 70.51%, which illustrates the crucial role of our self-supervision and regularization. For video generation, *full model* consistently shows the superior performances on IS , $H(y|x)$, $H(y)$. Especially in MUG, *full model* outperforms the runner-up MoCoGAN by 18.6% on IS , demonstrating that high quality of videos generated by our model. Note that our *baseline* is also compared favorably to DSVAE, illustrating the superiority of the designed sequential VAE model.

It is worth noting that our model with self-supervision *full model* outperforms *baseline-sv* in SMMNIST on the representation disentanglement. The possible reason is with the ground truth labels, *baseline-sv* only encourages z_f to

contain identity information but fails to exclude the information in z_t , resulting in confusion in digit recognition when using various dynamic variables. For MUG, *baseline-sv* performs better on preserving motion. We conjecture it's because that the strong supervision of expression labels forces z_t to encode the dynamic information, and z_f does not favor encoding temporal dynamic information and thus varying z_f does not affect the motion much.

Qualitative results We first demonstrate the ability of our model to manipulate video generation in Figure 7. By fixing z_f and sampling z_t , our model can generate videos with the same object that performs various motions as shown in Figure 7.a. Even in one video, the facial expression can be transferred by controlling z_t , as shown in Figure 7.b.

We also evaluate the appearance diversity of generate objects from our model. Figure 6 shows the frames our model generates with sampled z_f . The objects with realistic and diverse appearances validate our model's outstanding capability of high-quality video generation.

Besides, we compare our model with MoCoGAN on unconditional video generation on MUG dataset, as shown in Figure 5. The videos are generated with sampled z_f and z_t . MoCoGAN generates videos with many artifacts, such as unrealistic eyes and inconsistent mouth in the third video. Conversely, our model generates more realistic human faces with consistent high-coherence expressions.

Table 2. Ablation study of disentanglement on MUG dataset.

Methods	Acc	IS	$H(y v)$	$H(y)$
$No \mathcal{L}_{SSC}$	61.45%	4.850	0.201	1.734
$No \mathcal{L}_{DFP}$	58.32%	4.423	0.284	1.721
$No \mathcal{L}_{MI}$	66.07%	4.874	0.175	1.749
Full model	70.51%	5.136	0.135	1.760

5.1.4 Ablation Studies

In this section, we present an ablation study to empirically measure the impact of each regularization of our model on its performance. The variant without a certain regularization is denoted as $No \mathcal{L}_X$. In Table 2, we report the quantitative evaluation. We note that $No \mathcal{L}_{SSC}$ performs worse than the full model. This illustrates the significance of the static consistency constraint to make z_f to be disentangled from z_t . $No \mathcal{L}_{DFP}$ degrades the performance on Acc considerably, indicating that \mathcal{L}_{DFP} as regularization is crucial to preserve the action information in the dynamic vector. Besides, after removing the mutual information regularization, $No \mathcal{L}_{MI}$ again shows an inferior performance to the full model. A possible explanation is that our \mathcal{L}_{MI} encourages that z_t excludes the appearance information; thus the appearance information is only from z_f . The qualitative results shown in Figure 8 confirms this analysis. We generate a video with z_f of the woman and z_t of the man in the first row by different variants of our model. Without \mathcal{L}_{MI} , some characteristics of the man are still preserved, such as the beard, confirming that the appearance information partially remains in z_t . In the results of $No \mathcal{L}_{SSC}$, the beard is more evident, indicating that the static and dynamic variable are still entangled. On the other hand, without \mathcal{L}_{DFP} , the woman in the generated video cannot mimic the action of the man well, which indicates the dynamic variable does not encode the motion information properly. Finally, the person in the generated video of *baseline* neither preserves the appearance of the woman nor follows the expression of the man. It illustrates the representation disentanglement without any supervision remains a hard task.

5.2. Experiments on Audio Data

To demonstrate the general applicability of our model on sequential data, we conduct experiments on audio data, where the time-invariant and time-varying factors are the timbre of a speaker and the linguistic content of a speech, respectively. The dataset we use is TIMIT, which is a corpus of phonemically and lexically transcribed speech of American English speakers of different sexes and dialects [21], and contains 63000 recordings of read speeches. We split the dataset to training and testing subsets with a ratio of 5:1. As in [27], all the speech are presented as a sequence of 80 dimensional Mel-scale filter bank features.

We quantitatively compare our model with FHVAE and DSVAE on the speaker verification task based on either z_f or $z_{1:T}$, measured by the Equal Error Rate(EER) [8]. Note

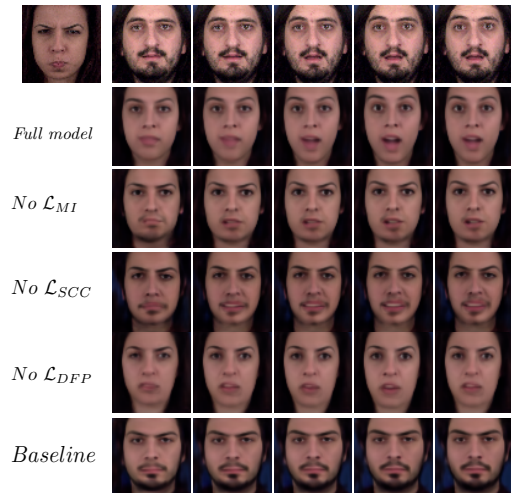


Figure 8. Ablation study on MUG dataset. The first frame of the appearance video and the motion video are shown in the first row.

Table 3. Performance comparison on speaker verification. Small errors are better for z_f and large errors are expected for $z_{1:T}$.

model	feature	dim	EER
FHVAE	z_f	16	5.06%
DSVAE	z_f	64	4.82%
	$z_{1:T}$	64	18.89%
Ours	z_f	64	4.80%
	$z_{1:T}$	64	40.12%

that we expect the speaker can be correctly verified with z_f as it encodes the timbre of speakers, and randomly guess with $z_{1:T}$ as it ideally only encodes the linguistic content. The results are shown in Table 3. Our model outperforms competing methods in both cases. Especially when based on $z_{1:T}$, our model doubles the score of the baseline, indicating our model significantly eliminate the timbre information in $z_{1:T}$.

6. Conclusion

We propose a self-supervised sequential VAE, which learns disentangled time-invariant and time-varying representations for sequential data. We show that, with readily accessible supervisory signals from data itself and off-the-shelf tools, our model can achieve comparable performance to the fully supervised models that require costly human annotations. The disentangling ability of our model is qualitatively and quantitatively verified on four datasets across video and audio domains. The appealing results on a variety of tasks illustrate that, leveraging self-supervision is a promising direction for representation disentanglement and sequential data generation. In the future, we plan to extend our model to high-resolution video generation, video prediction and image-to-video generation.

References

- [1] Pulkit Agrawal, Joao Carreira, and Jitendra Malik. Learning to see by moving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 37–45, 2015. 2, 3
- [2] N. Aifanti, C. Papachristou, and A. Delopoulos. The mug facial expression database. In *11th International Workshop on Image Analysis for Multimedia Interactive Services WIAMIS 10*, pages 1–4, April 2010. 5
- [3] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017. 2, 3
- [4] Justin Bayer and Christian Osendorfer. Learning stochastic recurrent networks. *arXiv preprint arXiv:1411.7610*, 2014. 2
- [5] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013. 1
- [6] Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pages 2610–2620, 2018. 1, 2, 5
- [7] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016. 1, 2
- [8] Mohamed Chenafa, Dan Istrate, Valeriu Vrabie, and Michel Herbin. Biometric system based on voice recognition using multiclassifiers. In *European Workshop on Biometrics and Identity Management*, pages 206–215. Springer, 2008. 8
- [9] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 3
- [10] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pages 2980–2988, 2015. 2
- [11] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. *arXiv preprint arXiv:1802.07687*, 2018. 5
- [12] Emily L Denton et al. Unsupervised learning of disentangled representations from video. In *Advances in neural information processing systems*, pages 4414–4423, 2017. 2
- [13] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1422–1430, 2015. 2
- [14] Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Style aggregated network for facial landmark detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 379–388, 2018. 4
- [15] Babak Esmaeili, Hao Wu, Sarthak Jain, Alican Bozkurt, Narayanaswamy Siddharth, Brooks Paige, Dana H Brooks, Jennifer Dy, and Jan-Willem van de Meent. Structured disentangled representations. *arXiv preprint arXiv:1804.02086*, 2018. 1
- [16] Babak Esmaeili, Hao Wu, Sarthak Jain, N Siddharth, Brooks Paige, and Jan-Willem Van de Meent. Hierarchical disentangled representations. *stat*, 1050:12, 2018. 1
- [17] Otto Fabius and Joost R van Amersfoort. Variational recurrent auto-encoders. *arXiv preprint arXiv:1412.6581*, 2014. 2
- [18] Deng-Ping Fan, Ming-Ming Cheng, Jiang-Jiang Liu, Shang-Hua Gao, Qibin Hou, and Ali Borji. Salient objects in clutter: Bringing salient object detection to the foreground. In *European Conference on Computer Vision (ECCV)*. Springer, 2018. 1
- [19] Deng-Ping Fan, Wenguan Wang, Ming-Ming Cheng, and Jianbing Shen. Shifting more attention to video salient object detection. In *IEEE CVPR*, 2019. 1
- [20] Chuang Gan, Boqing Gong, Kun Liu, Hao Su, and Leonidas J Guibas. Geometry guided convolutional neural networks for self-supervised video representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5589–5597, 2018. 2, 3
- [21] John S Garofolo, L F Lamel, W M Fisher, Jonathan G Fiscus, D S Pallett, and Nancy L Dahlgren. Darpa timit acoustic-phonetic continuous speech corpus cd-rom {TIMIT}. Technical report, 1993. 8
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 2
- [23] Jiawei He, Andreas Lehrmann, Joseph Marino, Greg Mori, and Leonid Sigal. Probabilistic video generation using holistic attribute control. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 452–467, 2018. 1, 6
- [24] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017. 1, 2
- [25] Geoffrey E Hinton and Richard S Zemel. Autoencoders, minimum description length and helmholtz free energy. In *Advances in neural information processing systems*, pages 3–10, 1994. 2
- [26] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 3
- [27] Wei-Ning Hsu, Yu Zhang, and James Glass. Unsupervised learning of disentangled and interpretable representations from sequential data. In *Advances in neural information processing systems*, pages 1878–1889, 2017. 1, 2, 8
- [28] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. 4
- [29] Dinesh Jayaraman and Kristen Grauman. Learning image representations tied to ego-motion. In *Proceedings of the*

- IEEE International Conference on Computer Vision*, pages 1413–1421, 2015. 2, 3
- [30] Insu Jeon, Wonkwang Lee, and Gunhee Kim. Ib-gan: Disentangled representation learning with information bottleneck gan. 2018. 1, 2
- [31] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *arXiv preprint arXiv:1902.06162*, 2019. 3
- [32] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 1
- [33] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. *arXiv preprint arXiv:1802.05983*, 2018. 1, 2
- [34] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [35] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6874–6883, 2017. 3
- [36] Yingzhen Li and Stephan Mandt. Disentangled sequential autoencoder. *arXiv preprint arXiv:1803.02991*, 2018. 1, 2, 5, 6
- [37] Francesco Locatello, Stefan Bauer, Mario Lucic, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *ICML*, 2019. 1
- [38] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer, 2016. 3
- [39] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015. 3
- [40] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016. 3
- [41] Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Ambient sound provides supervision for visual learning. In *European conference on computer vision*, pages 801–816. Springer, 2016. 2, 3
- [42] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2701–2710, 2017. 2, 3
- [43] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 1
- [44] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016. 6
- [45] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 5
- [46] Ximeng Sun, Huijuan Xu, and Kate Saenko. A two-stream variational adversarial network for video generation. *arXiv preprint arXiv:1812.01037*, 2018. 1
- [47] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 5
- [48] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. *arXiv preprint arXiv:1706.08033*, 2017. 2
- [49] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM, 2008. 2
- [50] Jiangliu Wang, Jianbo Jiao, Linchao Bao, Shengfeng He, Yunhui Liu, and Wei Liu. Self-supervised spatio-temporal representation learning for videos by predicting motion and appearance statistics. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [51] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2015. 2, 3
- [52] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016. 3
- [53] Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1004–1013, 2018. 1
- [54] Yizhe Zhu, Jianwen Xie, Bingchen Liu, and Ahmed Elgammal. Learning feature-to-feature translator by alternating back-propagation for generative zero-shot learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2019. 1
- [55] Yizhe Zhu, Jianwen Xie, Zhiqiang Tang, Xi Peng, and Ahmed Elgammal. Semantic-guided multi-attention localization for zero-shot learning. In *Thirty-third Conference on Neural Information Processing Systems (NeurIPS)*, Dec 2019. 1