This CVPR 2020 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Non-Adversarial Video Synthesis with Learned Priors

Abhishek Aich^{†,*}, Akash Gupta^{†,*}, Rameswar Panda[‡], Rakib Hyder[†], M. Salman Asif[†], Amit K. Roy-Chowdhury[†] University of California, Riverside[†], IBM Research AI, Cambridge[‡]

{aaich0010, agupt0130, rpand002, rhyde0010, sasif0ece, amitrc0ece}.ucr.edu

Abstract

Most of the existing works in video synthesis focus on generating videos using adversarial learning. Despite their success, these methods often require input reference frame or fail to generate diverse videos from the given data distribution, with little to no uniformity in the quality of videos that can be generated. Different from these methods, we focus on the problem of generating videos from latent noise vectors, without any reference input frames. To this end, we develop a novel approach that jointly optimizes the input latent space, the weights of a recurrent neural network and a generator through non-adversarial learning. Optimizing for the input latent space along with the network weights allows us to generate videos in a controlled environment, i.e., we can faithfully generate all videos the model has seen during the learning process as well as new unseen videos. Extensive experiments on three challenging and diverse datasets well demonstrate that our proposed approach generates superior quality videos compared to the existing state-of-the-art methods.

1. Introduction

Video synthesis is an open and challenging problem in computer vision. As literature suggests, a deeper understanding of spatio-temporal behavior of video frame sequences can directly provide insights in choosing priors, future prediction, and feature learning [32, 34]. Much progress has been made in developing variety of ways to generate videos which can be classified into broadly two categories: class of video generation methods which require random latent vectors without any reference input pixel [24, 31, 32], and class of video generation methods which do depend on reference input pixels [10, 28, 34]. Current literature contains methods mostly from the second class which often requires some human intervention [10, 34].

In general, Generative Adversarial Networks (GANs) [8]



Figure 1: Comparison of proposed non-adversarial approach to one representative adversarial approach (MoCoGAN [31]) on the Chair-CAD [2] dataset. *Top*: MoCoGAN often generates blurry frames including similar type of chairs for different videos as the time step increases. *Bottom*: Our approach¹, on the other hand, generates relatively sharper frames, maintaining consistency with the type of chairs unique to each video in the dataset.

have shown remarkable success in various kinds of video modality problems [13, 16, 18, 24]. Initially, video generation frameworks predominantly used GANs to synthesize videos from latent noise vectors. For example, VGAN [32] and TGAN [24] proposed generative models that synthesize videos from random latent vectors with deep convolutional GAN. Recently, MoCoGAN [31] proposed to decompose a video into content and motion parts using a generator guided by two discriminators. During testing, these frameworks generate videos that are captured in the range of the trained generator, by taking random latent vectors. While all these methods have obtained reasonable performance on commonly used benchmark datasets, they utilize adversarial learning to train their models and hence, inherit the shortcomings of GANs. Specifically, GANs are often very sensitive to multiple factors such as random network initialization, and type of layers employed to build the network [14, 25]. Some infamous drawbacks of GANs are modecollapse (i.e., able to generate only some parts of the data

^{*}Joint first authors

¹Project page: https://abhishekaich27.github.io/navsynth.html

distribution: see Fig. 1 for an example) and/or vanishing generator gradients due to discriminator being way better in distinguishing fake samples and real samples [1].

Non-adversarial approaches [4, 11, 15] have recently been explored to tackle these challenges. For example, Generative Latent Optimization (GLO) [4] and Generative Latent Nearest Neighbor (GLANN) [11] investigate the importance of inductive bias in convolutional networks by disconnecting the discriminator for a non-adversarial learning protocol of GANs. These works show that without a discriminator, a generator can be learned to map the training images in the given data distribution to a lower dimensional latent space that is learned in conjunction with the weights of the generative network. Such procedure not only avoids the mode-collapse problem of the generators, but also provides the user an optimized low dimensional latent representation (embedding) of the data in contrast with the random latent space as in GANs. Recently Video-VAE [10] proposed to use Variational Auto-Encoder (VAE) for conditional video synthesis, either by randomly generating or providing the first frame to the model for synthesizing a video. However, the quality of generated videos using Video-VAE often depends on the provided input frame. Non-adversarial video synthesis without any visual inputs still remains as a novel and rarely addressed problem.

In this paper, we propose a novel non-adversarial framework to generate videos in a controllable manner without any reference frame. Specifically, we propose to synthesize videos from two optimized latent spaces, one providing control over the static portion of the video (*static latent space*) and the other over the transient portion of the video (*transient latent space*). We propose to jointly optimize these two spaces while optimizing the network (a generative and a recurrent network) weights with the help of regression-based reconstruction loss and a triplet loss.

Our approach works as follows. During training, we jointly optimize over network weights and latent spaces (both static and transient) and obtain a common transient latent space, and individual static latent space dictionary for all videos sharing the same class (see Fig. 2). During testing, we randomly choose a static vector from the dictionary, concatenate it with the transient latent vector and generate a video. This enables us to obtain a controlled environment of diverse video generation from learned latent vectors for each video in the given dataset, while maintaining almost uniform quality. In addition, the proposed approach also allows a concise video data representation in form of learned vectors, frame interpolation (using a low rank constraint introduced in [12]), and generation of videos unseen during the learning paradigm.

The key contributions of our work are as follows.

• We propose a novel framework for generating a wide range of diverse videos from learned latent vectors with-

out any conditional input reference frame with almost uniform visual quality. Our framework obtains a latent space dictionary on both static and transient portions for the training video dataset, which enables us to generate even unseen videos with almost equal quality by providing combinations of static and transient latent vectors that were not part of training data.

• Our extensive experiments on multiple datasets well demonstrate that the proposed method, without the adversarial training protocol, has better or at par, performance with current state-of-the-art methods [24, 31, 32]. Moreover, we do not need to optimize the (multiple) discriminator networks as in previous methods [24, 31, 32] which offers a computational advantage.

2. Related Works

Our work relates to two major research directions: video synthesis and non-adversarial learning. This section focuses on some representative methods closely related to our work.

2.1. Video Synthesis

Video synthesis has been studied from multiple perspectives [10, 24, 28, 31, 32] (see Tab. 1 for a categorization of existing methods). VGAN [32] demonstrates that a video can be divided into foreground and background using deep neural networks. TGAN [24] proposes to use a generator to capture temporal dynamics by generating correlated latent codes for each video frame and then using an image generator to map each of these latent codes to a single frame for the whole video. MoCoGAN [31] presents a simple approach to separate content and motion latent codes of a video using adversarial learning. The most relevant work for us is Video-VAE [10] that extends the idea of image generation to video generation using VAE by proposing a structured latent space in conjunction with the VAE architecture for video synthesis. While this method doesn't require a discriminator network, it depends on reference input frame to generate a video. In contrast, our method proposes a efficient framework for synthesizing videos from learnable latent vectors without any input frame. This gives a controlled environment for video synthesis that even enables us to generate visually good quality unseen videos through combining static and transient parts.

2.2. Non-Adversarial Learning

Generative adversarial networks, as powerful as they are in pixel space synthesis, are also difficult to train. This is owing to the saddle-point based optimization game between the generator and the discriminator. On top of the challenges discussed in the previous section, GANs require careful user driven configuration tuning which may



Figure 2: **Overview of the proposed method.** Videos can be broken down into two main parts: static and transient components. To capture this, we map a video (with L frame sequence) into two learnable latent spaces. We jointly learn the static latent space and the transient latent space along with the network weights. We then use these learned latent spaces to generate videos at the inference time. See Sec. 3 for more details.

Mathada	Settings					
Wiethous	Adversarial	Input	Input latent			
	learning?	frame?	vectors?			
VGAN [32]	 ✓ 	X	√(random)			
TGAN [24]	\checkmark	X	√(random)			
MoCoGAN [31]	✓	X	√(random)			
Video-VAE [10]	×	\checkmark	√(random)			
Ours	×	X	√(learned)			

Table 1: **Categorization of prior works in video synthesis.** Different from existing methods, our model doesn't require a discriminator, or any reference input frame. However, since we have learned latent vectors, we have control of the kind of videos the model should generate.

not guarantee same performance for every run. Some techniques to make the generator agnostic to described problems have been discussed in [25]. The other alternative to the same has given rise to non-adversarial learning of generative networks [4, 11]. Both [4, 11] showed that properties of convolutional GANs can be mimicked using simple reconstruction losses while discarding the discriminator.

While there has been some work on image generation from learned latent vectors [4, 11], our work significantly differs from these methods as we do not map all the frames pixel-wise in a given video to the same latent distribution. This is because doing so would require a separate latent space (hence, a separate model for each video) for all the videos in a given dataset, and performing any operation in that space would naturally become video specific. Instead, we divide the latent space of videos sharing the same class into two parts - static and transient. This gives us a dictionary of static latent vectors for all videos and a common transient latent subspace. Hence, any random video of the dataset can now be represented by the combination of one static vector (which remains same for all frames) and the common transient subspace.

3. Formulation

Define a video clip V represented by L frames as $V = [v_1, v_2, \cdots, v_L]$. Corresponding to each frame v_i , let there be a point z_i in latent space $Z_V \in \mathbb{R}^{D \times L}$ such that

$$\mathcal{Z}_{\mathsf{V}} = \begin{bmatrix} z_1, z_2, \cdots, z_L \end{bmatrix} \tag{1}$$

which forms a path of length L. We propose to disentangle a video into two parts: a static constituent, which captures the constant portion of the video common for all frames, and a transient constituent which represents the temporal dynamics between all the frames in the video. Hence, let Z_V be decomposed as $Z_V = [Z_s^\top, Z_t^\top]^\top$ where $Z_s \in \mathbb{R}^{D_s \times L}$ represents the static subspace and $Z_t \in \mathbb{R}^{D_t \times L}$ represents the transient subspace with $D = D_s + D_t$. Thus $\{z_i\}_{i=1}^L$ in (1) can be expressed as $z_i = [z_i^{(s)\top}, z_i^{(t)\top}]^\top \forall i =$ $1, 2, \dots, L$. Next assuming that the video is of short length, we can fix $z_i^{(s)} = z^{(s)}$ for all frames after sampling only once. Therefore, (1) can be expressed as

$$\mathcal{Z}_{\mathsf{V}} = \begin{bmatrix} \begin{bmatrix} z^{(\mathsf{s})} \\ z_1^{(\mathsf{t})} \end{bmatrix}, \begin{bmatrix} z^{(\mathsf{s})} \\ z_2^{(\mathsf{t})} \end{bmatrix}, \cdots, \begin{bmatrix} z^{(\mathsf{s})} \\ z_L^{(\mathsf{t})} \end{bmatrix} \end{bmatrix}$$
(2)

The transient portion will represent the motion of a given video. Intuitively, the latent vectors corresponding to this

transient state should be correlated, or in other words, will form a path between $z_1^{(t)}$ and $z_L^{(t)}$. Specifically, the frames in a video are correlated in time and hence a frame v_i at time i = T is a function of all previous frames $\{v_i\}_{i=1}^{T-1}$. As a result, their corresponding transient representation should also exhibit such a trajectory. This kind of representation of latent vectors can be obtained by employing a Recurrent Neural Network (RNN) where output of each cell of the network is a function of its previous state or input. Denote the RNN as \mathcal{R} with weights θ . Then, the RNN output $\mathcal{R}(z_i) = \{r_i^{(t)}\} \forall i = 1, 2, \cdots, L$ is a sequence of correlated variables representing the transient state of the video.

3.1. Learning Network Weights

Define a generative network \mathcal{G} with weights represented by γ . \mathcal{G} takes latent vectors sampled from \mathcal{Z}_{V} as input and predicts up to L frames of the video clip. For a set of N videos, initialize set of D-dimensional vectors \mathcal{Z}_{V} to form the pair $\left\{ (\mathcal{Z}_{V_{1}}, V_{1}), (\mathcal{Z}_{V_{2}}, V_{2}), \cdots, (\mathcal{Z}_{V_{N}}, V_{N}) \right\}$. More specifically from (2), defining $\mathbf{z}^{(s)} = [z^{(s)}, z^{(s)}, \cdots, z^{(s)}] \in \mathbb{R}^{D_{s} \times L}$, and $\mathbf{z}^{(t)} = [z_{1}^{(t)}, z_{2}^{(t)}, \cdots, z_{L}^{(t)}] \in \mathbb{R}^{D_{t} \times L}$, we will have the pairs $\left\{ \left(\begin{bmatrix} \mathbf{z}^{(s)} \\ \mathbf{z}^{(t)} \end{bmatrix}_{1}, V_{1} \right), \left(\begin{bmatrix} \mathbf{z}^{(s)} \\ \mathbf{z}^{(t)} \end{bmatrix}_{2}, V_{2} \right), \cdots, \left(\begin{bmatrix} \mathbf{z}^{(s)} \\ \mathbf{z}^{(t)} \end{bmatrix}_{N}, V_{N} \right) \right\}$.

With these pairs, we propose to optimize the weights γ , θ , and input latent vectors Z_V (sampled once in the beginning of training) in the following manner. For each video V_j , we jointly optimize for θ , γ , and $\{Z_{V_j}\}_{j=1}^N$ for every epoch in two stages:

Stage 1 :
$$\min_{\gamma} \ell \left(\mathsf{V}_{j}, \mathcal{G}(\mathcal{Z}_{\mathsf{V}_{j}}) | \left(\mathcal{Z}_{\mathsf{V}_{j}}, \theta \right) \right)$$
 (3.1)

Stage 2:
$$\min_{\mathcal{Z}_{V,\theta}} \ell \left(\mathsf{V}_{j}, \mathcal{G}(\mathcal{Z}_{\mathsf{V}_{j}}) | \gamma \right)$$
 (3.2)

 $\ell(\cdot)$ can be any regression-based loss. For rest of the paper, we will refer to both (3.1) and (3.2) together as $\min_{\mathcal{Z}_{V},\theta,\gamma} \ell_{rec}$.

Regularized loss function to capture static subspace. The transient subspace, along with the RNN, handles the temporal dynamics of the video clip. To equally capture the static portion of the video, we randomly choose a frame from the video and ask the generator to compare its corresponding generated frame during training. For this, we update the above loss as follows.

$$\min_{\mathcal{Z}_{\mathsf{V}},\theta,\gamma} \left(\ell_{\mathsf{rec}} + \lambda_{\mathsf{s}} \ell_{\mathsf{static}} \right) \tag{4}$$

where $\ell_{\text{static}} = \ell(\hat{v}_k, v_k)$ with $k \in \{1, 2, \dots, L\}$, v_k is the ground truth frame, $\hat{v}_k = \mathcal{G}(\mathbf{z}_k)$, and λ_s is the regularization constant. ℓ_{static} can also be understood to essentially handle the role of image discriminator in [31, 33] that ensures that the generated frame is close to the ground truth frame.

3.2. Learning Latent Spaces

Non-adversarial learning involves joint optimization of network weights as well as the corresponding input latent space. Apart from the gradients with respect to loss in (4), we propose to further optimize the latent space with gradient of a loss based on the triplet condition as follows.

3.2.1 The Triplet Condition



Figure 3: **Triplet Condition in the transient latent space.** Latent code representation of different frames of short video clips may lie very near to each other in the transient subspace. Using the proposed triplet condition, our model learns to explain the dynamics of similar looking frames and simultaneously map them to distinct latent vectors.

Short video clips often have indistinguishable dynamics in consecutive frames which can force the latent code representations to lie very near to each other in the transient subspace. However, an ideal transient space should ensure that the latent vector representation of a frame should only be close to a similar frame than a dissimilar one [26, 27]. To this end, we introduce a triplet loss to (4) that ensures that a pair of co-occurring frames v_i^a (anchor) and v_i^p (positive) are closer but distinct to each other in embedding space than any other frame v_i^n (negative) (see Fig. 3). In this work, positive frames are randomly sampled within a margin range α of the anchor and negatives are chosen outside of this margin range. Defining a triplet set with transient latent code vectors $\{z_i^{(t),a}, z_i^{(t),p}, z_i^{(t),n}\}$, we aim to learn the transient embedding space $\mathbf{z}^{(t)}$ such that

$$\|z_i^{(\mathsf{t}),a} - z_i^{(\mathsf{t}),p}\|_2^2 + \alpha < \|z_i^{(\mathsf{t}),a} - z_i^{(\mathsf{t}),n}\|_2^2$$

 $\forall \{z_i^{(t),a}, z_i^{(t),p}, z_i^{(t),n}\} \in \Gamma$, where Γ is the set of all possible triplets in $\mathbf{z}^{(t)}$. With the above regularization, the loss in (4) can be written as

$$\min_{Z_{\nabla,\theta,\gamma}} (\ell_{\text{rec}} + \lambda_{\text{s}} \ell_{\text{static}})$$

s.t. $\|z_i^{(\mathsf{t}),a} - z_i^{(\mathsf{t}),p}\|_2^2 + \alpha < \|z_i^{(\mathsf{t}),a} - z_i^{(\mathsf{t}),n}\|_2^2$ (5)

where α is a hyperparameter that controls the margin while selecting positives and negatives.

3.3. Full Objective Function

For any choice of differentiable generator \mathcal{G} , the objective (4) will be differentiable with respect to \mathcal{Z}_{V} , and (γ, θ)

[5]. We initialize Z_V by sampling them from two different Gaussian distributions for both static and transient latent vectors. We also ensure that the latent vectors Z_V lie on the unit ℓ_2 sphere, and hence, we project Z_V after each update by dividing its value by max $(1, ||Z_V||)$ [4], where max (\cdot) returns maximum among the set of given elements. Finally, the complete objective function can be written as follows.

$$\min_{\mathcal{Z}_{V},\theta,\gamma} \left(\ell_{\text{rec}} + \lambda_{\text{s}} \ell_{\text{static}} + \lambda_{\text{t}} \ell_{\text{triplet}} \right)$$
(6)

where $\ell_{\text{static}} = \ell(\hat{v}_k, v_k)$, λ_t is a regularization constant for the triplet loss, and $\ell_{\text{triplet}} = \max\left(\|z_i^{(t),a} - z_i^{(t),p}\|_2^2 + \alpha - \|z_i^{(t),a} - z_i^{(t),n}\|_2^2, 0\right)$.

The weights of the generator γ and static latent vector $z^{(s)}$ are updated by gradients of the losses $\ell_{\rm rec}$ and $\ell_{\rm static}$. The weights. θ , of the RNN, and transient latent vectors $\mathbf{z}^{(t)}$ are updated by gradients of the losses $\ell_{\rm rec}$, $\ell_{\rm static}$ and $\ell_{\rm triplet}$.

3.3.1 Low Rank Representation for Interpolation

The objective of video frame interpolation is to synthesize non-existent frames in-between the reference frames. While the triplet condition ensures that similar frames have their transient latent vectors nearby, it doesn't ensure that they lie on a manifold where simple linear interpolation will yield latent vectors that generate frames with plausible motion compared to preceding and succeeding frames [4, 12]. This means that the transient latent subspace can be represented in a much lower dimensional space compared to its larger ambient space. So, to enforce such a property, we project the latent vectors into a low dimensional space while learning them along with the network weights, first proposed in [12]. Mathematically, the loss in (6) can be written as

$$\min_{Z_{\mathbf{V}},\theta,\gamma} \left(\ell_{\text{rec}} + \lambda_{\text{s}} \ell_{\text{static}} + \lambda_{\text{t}} \ell_{\text{triplet}} \right)$$
s.t. $\operatorname{rank} \left(\mathbf{z}^{(t)} \right) = \rho$ (7)

where rank(·) indicates rank of the matrix and ρ is a hyperparameter that decides what manifold $\mathbf{z}^{(t)}$ is to be projected on. We achieve this by reconstructing $\mathbf{z}^{(t)}$ matrix from its top ρ singular vectors in each iteration [7]. Note that, we only employ this condition for optimizing the latent space for the frame interpolation experiments in Sec. 4.3.3.

4. Experiments

In this section, we present extensive experiments to demonstrate the effectiveness of our proposed approach in generating videos through learned latent spaces.

4.1. Datasets

We evaluate the performance of our approach using three publicly available datasets which have been used in many prior works [10, 31, 32].

Chair-CAD [2]. This dataset consists of total 1393 chair-CAD models, out of which we randomly choose 820 chairs for our experiments with the first 16 frames, similar to [10]. The rendered frame in each video for all the models are center-cropped and then resized to $64 \times 64 \times 3$ pixels. We obtain the transient latent vectors for all the chair models with one static latent vectors for the training set.

Weizmann Human Action [9]. This dataset provides 10 different actions performed by 9 people, amounting to 90 videos. Similar to Chair-CAD, we center-cropped each frame, and then resized to $64 \times 64 \times 3$ pixels. For this dataset, we train our model to obtain nine static latent vectors (for nine different identities) and ten transient latent vectors (for ten different actions) for videos with 16 frames each.

Golf scene dataset [32]. Golf scene dataset [32] contains 20,268 golf videos with $128 \times 128 \times 3$ pixels which further has 583,508 short video clips in total. We randomly chose 500 videos with 16 frames each and resized the frames to $64 \times 64 \times 3$ pixels. Same as the Chair-CAD dataset, we obtained the transient latent vectors for all the golf scenes and one static latent vector for the training set.

4.2. Experimental Settings

We implement our framework in PyTorch [21]. Please see supplementary material for details on implementation and values of different hyper-parameters (D_s , D_t , α , etc.).

Network Architecture. We choose DCGAN [22] as the generator architecture for the Chair-CAD and Golf scene dataset, and conditional generator architecture from [20] for the Weizmann Human Action dataset for our experiments. For the RNN, we employ a one-layer gated recurrent unit network with 500 hidden units [6].

Choice of Loss Function for ℓ_{rec} **and** ℓ_{static} . One straight forward loss function that can be used is the mean squared loss, but it has been shown in literature that it leads to generation of blurry pixels [35]. Moreover, it has been shown empirically that generative functions in adversarial learning focus on edges [4]. Motivated by this, the loss function for ℓ_{rec} and ℓ_{static} is chosen to be the Laplacian pyramid loss $\mathcal{L}_{\text{Laplacian}}$ [17] defined as

$$\mathcal{L}_{\text{Laplacian}}(v, \hat{v}) = \sum_{l} 2^{2^{l}} |\mathsf{L}^{l}(v) - \mathsf{L}^{l}(\hat{v})|_{1}$$

where $L^{l}(\cdot)$ is the *l*-th level of the Laplacian pyramid representation of the input.

Baselines. We compare our proposed method with two adversarial methods. For Chair-CAD and Weizmann Human Action, we use MoCoGAN [31] as the baseline, and for Golf scene dataset, we use VGAN [32] as the baseline. We use the publicly available code for MoCoGAN and VGAN, and set the hyper-parameters as recommended in the published work. We also compare two different versions of the proposed method by ablating the proposed loss functions. Note that, we couldn't compare our results with Video-VAE [10] using our performance measures (described below) as the implementation has not been made available by the authors, and to the best of our efforts we couldn't reproduce the results provided by them.

Performance measures. Past video generation works have been evaluated quantitatively on Inception score (IS) [10]. But, it has been shown that IS is not a good evaluation metric for pixel domain generation, as the maximal IS score can be obtained by synthesizing a video from every class or mode in the given data distribution [3, 19, 30]. Moreover, a high IS does not guarantee any confidence on the quality of generation, but only on the diversity of generation. Since a generative model trained using our proposed method can generate all videos using the learned latent dictionary², and for a fair comparison with baselines, we use the following two measures, similar to measures provided in [31]. We also provide relevant bounds computed on real videos for reference. Note that arrows indicate whether higher (\uparrow) or lower (\downarrow) scores are better.

(1) **Relative Motion Consistency Score** (MCS \downarrow): Difference between consecutive frames captures the moving components, and hence motion in a video. So, firstly each frame in the generated video, as well as the ground-truth data, is represented as a feature vector computed using a VGG16 network [29] pre-trained on ImageNet [23] at the relu3_3 layer. Secondly, the averaged consecutive framefeature difference vector for both set of videos is computed, denoted by \hat{f} and f respectively. Finally, the relative MCS is then given by $\log_{10} \left(||f - \hat{f}||_2^2 \right)$.

(2) Frame Consistency Score (FCS \uparrow): This score measures the consistency of the static portion of the generated video frames. We keep the first frame of the generated video as reference and compute the averaged structural similarity measure for all frames. The FCS is then given by the average of this measure over all videos.

4.3. Qualitative Results

Fig. 5 shows some examples with randomly selected frames of generated videos for the proposed method and the adversarial approaches MoCoGAN [31] and VGAN [32]. For Chair-CAD [2] and Weizmann Human Action [9] dataset, it can be seen that the proposed method is able to generate visually good quality videos with a non-adversarial training protocol, whereas MoCoGAN produces blurry and inconsistent frames. Since we use optimized latent vectors unlike MoCoGAN (which uses random latent vectors for video generation), our method produces visually more ap-

pealing videos. Fig. 5 presents two particularly important points. As visualized for the Chair-CAD videos, the adversarial approach of MoCoGAN produces not only blurred chair images in the generated video, but they are also nonuniform in quality. Further, it can be seen that as the time step increases, MoCoGAN tends to generate the same chair for different videos. This shows a major drawback of the adversarial approaches, where they fail to learn the diversity of the data distribution. Our approach overcomes this by producing a optimized dictionary of latent vectors which can be used for generating any video in the data distribution easily. To further validate our method for qualitative results, we present the following experiments.

4.3.1 Qualitative Ablation Study

Fig. 4 qualitatively shows the contribution of the specific parts of the proposed method on Chair-CAD [2]. First, we investigate the impact of input latent vector optimization. For a fair comparison, we optimize the model for same number of epochs. It can be observed that the model benefits from the joint optimization of input latent space to produce better visual results. Next, we validate the contribution of ℓ_{static} and ℓ_{triplet} on a difficult video example whose chair color matches with the background. Our method, combined with ℓ_{static} and ℓ_{triplet} , is able to distinguish between the white background and the white body of the chair model.

Actions	Identities								
	P1	P2	P3	P4	P5	P6	P7	P8	P9
run	•	•	•		•				
walk	•	•			•		•		
jump		•		•	•	•	•	•	•
skip	•	•	•			•	•	•	•

Table 3: Generating videos by exchanging unseen actions by identities. Each cell in this table indicates a video in the dataset. Only cells containing the symbol \bullet indicate that the video was part of the training set. We randomly generated videos corresponding to rest of the cells indicated by symbols \bullet , \bullet , \bullet , and \bullet , visualized in Fig. 6.

4.3.2 Action Exchange

Our non-adversarial approach can effectively separate the static and transient portion of a video, and generate videos unseen during the training protocol. To validate these points, we choose a simple *matrix* completion for the combination of identities and actions in the Weizmann Human action [9] dataset. For training our model, we created a set of videos (without any cropping to present the complete scale of the frame) represented by the cells marked with • in Tab. 3. Hence, the unseen videos correspond to the cells not marked with •. During testing, we randomly generated

²Direct video comparison seems straight forward for our approach as the corresponding one-to-one ground truth is known. However, for [31, 32], we do not know which video is being generated (action may be known e.g. [31]) which makes such direct comparison infeasible and unfair.



boxes indicating the colored dots $(\bullet, \bullet, \bullet, \bullet)$ referred to in Tab. 3. This demonstrates the effectiveness of our method in disentangling static and transient portion of videos.

these unseen videos (marked with \bullet , \bullet , \bullet and \bullet in Tab. 3), and the visual results are shown in Fig. 6. This experiment clearly validates our claim of static (identities) and transient (action) portion disentanglement of a video and, generation of unseen videos by using combinations of action and iden-

4.4. Quantitative Results

Quantitative result comparisons with respect to baselines have been provided in Tab. 2. Compared to videos generated by the adversarial method MoCoGAN [31], we report a relative decrease of 19.22% in terms of MCS, and 4.70%

is able to produce dynamically consistent frames with re-

spect to the reference frames without any pixel clues.

	$\text{MCS}\downarrow$	$ $ FCS \uparrow		$\text{MCS}\downarrow$	FCS \uparrow		$ MCS \downarrow$	$ $ FCS \uparrow
Bound	0.0	0.91	Bound	0.0	0.95	Bound	0.0	0.97
MoCoGAN [31]	4.11	0.85	MoCoGAN [31]	3.41	0.85	VGAN [32]	3.61	0.88
Ours $(-\ell_{triplet} - \ell_{static})$	3.83	0.77	Ours $(-\ell_{triplet} - \ell_{static})$	3.87	0.79	Ours $(-\ell_{triplet} - \ell_{static})$	3.78	0.84
Ours $(+\ell_{triplet}+\ell_{static})$	3.32	0.89	Ours $(+\ell_{triplet}+\ell_{static})$	2.63	0.90	Ours $(+\ell_{triplet}+\ell_{static})$	2.71	0.84
(a) Chair-CAD [2]			(b) Weizmann Human Action [9]			(c) Golf	[32]	

Table 2: **Quantitative results comparison with state-of-the-art methods.** We obtained better scores on the proposed method on both Chair-CAD [2], Weizmann Human Action [9], and Golf [32] datasets, compared to the adversarial approaches (MoCoGAN, and VGAN). Best scores have been highlighted in bold.



Figure 7: **Examples of frame interpolation.** An important advantage of our method is translation of interpolation in the learned latent space to video space using (7). It can be observed that as ρ increases, the interpolation (bounded by color) is better. Note that the adjacent frames are also generated frames, and not ground truth frames.

relative increase in terms of FCS for chair-CAD dataset [2]. For the Weizmann Human Action [9] dataset, the proposed method is reported to have a relative decrease of 22.87% in terms of in terms of MCS, and 4.61% relative increase in terms of FCS. Similarly for Golf scene dataset [32], we perform competitively with VGAN [32] with a observed relative decrease of 24.90% in terms of in terms of MCS. A important conclusion from these results is that our proposed method, being non-adversarial in nature, learns to synthesize a diverse set of videos, and is able to perform at par with adversarial approaches. It should be noted that a better loss function for $\ell_{\rm rec}$ and $\ell_{\rm static}$ would produce stronger results. We leave this for future works.

4.4.1 Quantitative Ablation Study

In this section, we demonstrate the contribution of different components in our proposed methodology on the Chair-CAD [2] dataset. For all the experiments, we randomly generate 500 videos using our model by using the learned latent vector dictionary. We divide the ablation study into two parts. Firstly, we present the results for impact of the learned latent vectors on the network modules. For this, we simply generate videos once with the learned latent vectors (+Z), and once with randomly sampled latent vectors from a different distribution (-Z). The inter-dependency of our model weights and the learned latent vectors can be interpreted from Tab. 4a. We see that there is a relative decrease of 16.16% in MCS from 3.96 to 3.32, and 18.66% of rel-

			$ MCS \downarrow FCS \uparrow$					
	MCS↓	FCS ↑	Bound	0	0.91			
Bound	0	0.91	$-\ell_{triplet}-\ell_{static}$	3.83	0.77			
-Z	3.96	0.75	$-\ell_{triplet} + \ell_{static}$	3.82	0.85			
+Z	3.32	0.89	$+\ell_{triplet}-\ell_{static}$	3.36	0.81			
			+ ltriplet + lstatic	3.32	0.89			

(a) With respect to latent space optimization.

(b) With respect to loss functions

Table 4: Ablation study of proposed method on Chair-CAD [2]. In (a), we evaluate contributions of latent space optimization (Z). In (b), we evaluate contributions of ℓ_{triplet} and ℓ_{static} in four combinations. + and – indicate presence and absence of the terms, respectively.

ative increase in FCS. This shows that optimization of the latent space in the proposed method is important for good quality video generation.

Secondly, we investigate the impact of the proposed losses on the proposed method. Specifically, we look into four possible combinations of $\ell_{triplet}$ and ℓ_{static} . The results are presented in Tab. 4b. It can observed that the combination of triplet loss $\ell_{triplet}$ and static loss ℓ_{static} provides the best result when employed together, indicated by the relative decrease of 14.26% in MCS from 3.83 to 3.32.

5. Conclusion

We present a non-adversarial approach for synthesizing videos by jointly optimizing both network weights and input latent space. Specifically, our model consists of a global static latent variable for content features, a frame specific transient latent variable, a deep convolutional generator, and a recurrent neural network which are trained using a regression-based reconstruction loss, including a triplet based loss. Our approach allows us to generate a diverse set of almost uniform quality videos, perform frame interpolation, and generate videos unseen during training. Experiments on three standard datasets show the efficacy of our proposed approach over state-of-the-methods.

Acknowledgements. The work was partially supported by NSF grant 1664172 and ONR grant N00014-19-1-2264.

References

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In *Proceedings of the International Conference on Machine Learning*, pages 214–223, 2017. 2
- [2] Mathieu Aubry, Daniel Maturana, Alexei A Efros, Bryan C Russell, and Josef Sivic. Seeing 3D Chairs: Exemplar partbased 2D-3D Alignment using a Large Dataset of CAD Models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3762–3769, 2014. 1, 5, 6, 7, 8
- [3] Shane Barratt and Rishi Sharma. A Note on the Inception Score. *arXiv preprint arXiv:1801.01973*, 2018. 6
- [4] Piotr Bojanowski, Armand Joulin, David Lopez-Pas, and Arthur Szlam. Optimizing the Latent Space of Generative Networks. In *Proceedings of the International Conference* on Machine Learning, pages 599–608, 2018. 2, 3, 5
- [5] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed Sensing using Generative Models. In *Proceedings of the International Conference on Machine Learning*, pages 537–546, 2017. 5
- [6] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. arXiv preprint arXiv:1412.3555, 2014. 5
- [7] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The Elements of Statistical Learning*, volume 1. Springer Series in Statistics New York, 2001. 5
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. In Advances in Neural Information Processing Systems, pages 2672–2680, 2014. 1
- [9] Lena Gorelick, Moshe Blank, Eli Shechtman, Michal Irani, and Ronen Basri. Actions as Space-Time Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2247–2253, 2007. 5, 6, 7, 8
- [10] 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*, pages 452–467, 2018. 1, 2, 3, 5, 6
- [11] Yedid Hoshen, Ke Li, and Jitendra Malik. Non-Adversarial Image Synthesis with Generative Latent Nearest Neighbors. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 5811–5819, 2019. 2, 3
- [12] Rakib Hyder and M. Salman Asif. Generative Models for Low-Dimensional Video Representation and Reconstruction. *IEEE Transactions on Signal Processing*, 68:1688– 1701, 2020. 2, 5
- [13] Yong-Hoon Kwon and Min-Gyu Park. Predicting Future Frames using Retrospective Cycle-GAN. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1811–1820, 2019. 1
- [14] Jerry Li, Aleksander Madry, John Peebles, and Ludwig Schmidt. Towards Understanding the Dynamics of Generative Adversarial Networks. arXiv preprint arXiv:1706.09884, 2017. 1
- [15] Ke Li and Jitendra Malik. Implicit Maximum Likelihood

Estimation. arXiv preprint arXiv:1809.09087, 2018. 2

- [16] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P Xing. Dual-Motion GAN for Future-Flow Embedded Video Prediction. In Proceedings of the IEEE International Conference on Computer Vision, pages 1744–1752, 2017. 1
- [17] Haibin Ling and Kazunori Okada. Diffusion Distance for Histogram Comparison. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, volume 1, pages 246–253, 2006. 5
- [18] William Lotter, Gabriel Kreiman, and David Cox. Deep Predictive Coding Networks for Video Prediction and Unsupervised Learning. arXiv preprint arXiv:1605.08104, 2016. 1
- [19] Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, and Olivier Bousquet. Are GANs created equal? A Large-scale Study. In Advances in Neural Information Processing Systems, pages 700–709, 2018. 6
- [20] Mehdi Mirza and Simon Osindero. Conditional Generative Adversarial Networks. arXiv preprint arXiv:1411.1784, 2014. 5
- [21] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic Differentiation in PyTorch. In NIPS AutoDiff Workshop, 2017. 5
- [22] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. arXiv preprint arXiv:1511.06434, 2015. 5
- [23] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet: Large-Scale Visual Recognition Challenge. *International Journal* of Computer Vision, 115(3):211–252, 2015. 6
- [24] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal Generative Adversarial Nets with Singular Value Clipping. In *Proceedings of the IEEE International Conference* on Computer Vision, pages 2830–2839, 2017. 1, 2, 3
- [25] 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. 1, 3
- [26] Florian Schroff, Dmitry Kalenichenko, and James Philbin.
 FACENET: A Unified Embedding for Face Recognition and Clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015.
 4
- [27] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-Contrastive Networks: Self-Supervised Learning from Video. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1134–1141, 2018. 4
- [28] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating Arbitrary Objects via Deep Motion Transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2377–2386, 2019. 1, 2
- [29] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-scale Image Recognition. arXiv

- *preprint arXiv:1409.1556*, 2014. 6 [30] L Theis, A van den Oord, and M Bethge. A Note on the Evaluation of Generative Models. In Proceedings of the International Conference on Learning Representations, pages 1-10, 2016. 6
- [31] 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. 1, 2, 3, 4, 5, 6, 7, 8
- [32] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating Videos with Scene Dynamics. In Advances In Neural Information Processing Systems, pages 613-621, 2016. 1, 2, 3, 5, 6, 7, 8
- [33] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-Video Synthesis. In Advances in Neural Information Processing Systems, pages 1144-1156, 2018. 4
- [34] Tsun-Hsuan Wang, Yen-Chi Cheng, Chieh Hubert Lin, Hwann-Tzong Chen, and Min Sun. Point-to-Point Video Generation. In Proceedings of the IEEE International Conference on Computer Vision, 2019. 1
- [35] Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz. Loss Functions for Neural Networks for Image Processing. arXiv preprint arXiv:1511.08861, 2015. 5