

Conditional GAN with Discriminative Filter Generation for Text-to-Video Synthesis

Yogesh Balaji¹*, Martin Renqiang Min², Bing Bai², Rama Chellappa¹ and Hans Peter Graf²

¹University of Maryland, College Park

²NEC Labs America - Princeton

ybalaji@cs.umd.edu, {renqiang, bbai}@nec-labs.com, rama@umiacs.umd.edu, hpg@nec-labs.com

Abstract

Developing conditional generative models for text-to-video synthesis is an extremely challenging yet an important topic of research in machine learning. In this work, we address this problem by introducing Text-Filter conditioning Generative Adversarial Network (TFGAN), a conditional GAN model with a novel multi-scale text-conditioning scheme that improves text-video associations. By combining the proposed conditioning scheme with a deep GAN architecture, TFGAN generates high quality videos from text on challenging real-world video datasets. In addition, we construct a synthetic dataset of text-conditioned moving shapes to systematically evaluate our conditioning scheme. Extensive experiments demonstrate that TFGAN significantly outperforms existing approaches, and can also generate videos of novel categories not seen during training.

1 Introduction

Generative models have gained much interest in the research community over the last few years for unsupervised representation learning. Generative Adversarial Networks (GANs) [Goodfellow *et al.*, 2014] have been one of the most successful generative models till date. Following its introduction in 2014, significant progress has been made towards improving the stability, quality and the diversity of the generated images [Salimans *et al.*, 2016][Karras *et al.*, 2017]. While GANs have been successful in the image domain, recent efforts have extended it to other modalities such as text [Wang *et al.*, 2018a], graphs [Wang *et al.*, 2018b], etc.

In this work, we focus on the less studied domain of videos. Generating videos are much harder than images because the additional temporal dimension makes generated data extremely high dimensional, and the generated sequences must be both photo-realistically diverse and temporally consistent. We tackle the problem of text-conditioned video synthesis where the input is a text description and the goal is to synthesize a video corresponding to the input text. This problem has many potential applications, some of which include

generating synthetic data for machine learning tasks, domain adaptation, multimedia applications, etc.

Two recent works that address the problem of text-conditioned video generation include [Li *et al.*, 2018] and [Pan *et al.*, 2017]. Both these methods are variants of conditional GAN applied to the video data. In spite of some successes, they have the following limitations: (1) They employ 3D transposed convolution layers in the generator network, which constrains them to only produce fixed-length videos. (2) Their models are trained on low-resolution videos - results are shown only at a 64×64 resolution. (3) Text conditioning is performed using a simple concatenation of video and text features in the discriminator: Such a conditioning scheme may perform well on certain datasets, but maybe inadequate for capturing rich video-text variations.

In this work, we aim to address all the concerns above. First, to model videos of varying lengths, we use a recurrent neural network in the latent space and employ a shared frame generator network similar to [Tulyakov *et al.*, 2018]. Second, we present a model for generating higher-resolution videos by using a ResNet-style architecture in the generator and the discriminator network. Third, we propose a new multi-scale text-conditioning scheme based on discriminative convolutional filter generation to strengthen the associations between the conditioned text and the generated video. We call our model Text-Filter conditioning GAN (TFGAN). Finally, we construct a synthetic moving shapes dataset to extensively evaluate the effectiveness of the proposed conditioning scheme. Constructing this synthetic dataset is extremely useful as it captures rich text-video variations that are lacking in the datasets currently used for text-to-video synthesis [Li *et al.*, 2018; Pan *et al.*, 2017]. We demonstrate the effectiveness of our approach on (1) real-world datasets such as Kinetics Human Action dataset [Kay *et al.*, 2017] that has complex videos with high diversity but has relatively simple text-video variations, and (2) synthetic dataset that has simple video dynamics but models complex text-video associations.

In summary, our contributions in this work are as follows: (i) A new conditional GAN with an effective multi-scale text-conditioning scheme based on discriminative convolutional filter generation is proposed; (ii) A synthetic dataset for studying text conditioning in video generation is presented; (iii) A framework for generating complex video sequences and capturing rich text-video variations is presented.

*Work done during internship at NEC Laboratories America

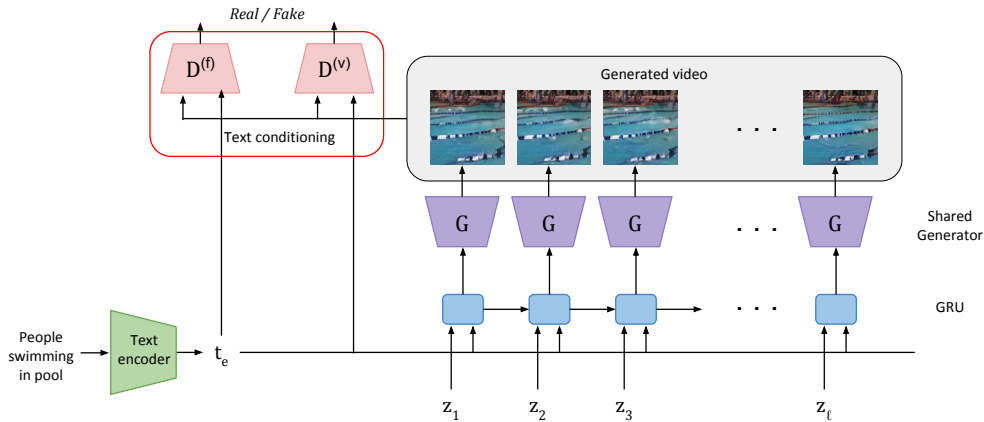


Figure 1: Our TFGAN framework. The box highlighted in red is where the conditioning is performed and is expanded in Fig. 2.

2 Related Work

Two popular approaches to generative modeling include GANs [Goodfellow *et al.*, 2014] and Variational Autoencoders (VAEs) [Kingma and Welling, 2014]. GANs are formulated as a 2-player *minimax* game between a generator and a discriminator network, while VAEs are based on variational inference where a variational lower bound of observed data log-likelihood is optimized. Among the two approaches, GANs have generated significant interest as they have been shown to produce images of high sample fidelity and diversity [Karras *et al.*, 2017] [Brock *et al.*, 2018].

A variant of GAN models is conditional GANs where the generator network is conditioned on input variables of interest. Such a conditioning input can be labels [Odena *et al.*, 2017], attributes [Yan *et al.*, 2016], text [Zhang *et al.*, 2017; Xu *et al.*, 2018] or even images [Zhu *et al.*, 2017]. We focus on text conditioning since it is relevant to this work. One of the first works to perform text-conditioned image synthesis is [Reed *et al.*, 2016]. Their method was only shown to synthesize low-resolution images. To improve the resolution, [Zhang *et al.*, 2017] proposed stacking multiple GAN architectures, each producing images of increasing resolution. While the above two methods perform conditioning using the global text representation, [Xu *et al.*, 2018] adopts an attention mechanism to focus on fine-grained word-level representations to enable improved conditioning.

While image generation is a well studied problem, there has been very little progress in the domain of video generation. [Vondrick *et al.*, 2016] proposed a GAN architecture based on 3D convolutions to generate video sequences, but it can only generate fixed-length videos. [Tulyakov *et al.*, 2018] proposed using a recurrent neural network in the latent space to model videos of varying lengths. While these models are not designed to handle video generation from text, [Li *et al.*, 2018] and [Pan *et al.*, 2017] perform text-conditioned video synthesis by using the sentence embedding as a conditional input. However, both of these conditional generative models are based on 3D convolutions, they can only produce fixed-length low-resolution videos. In this work, we address this issue by developing an architecture capable of producing

higher-resolution videos of varying lengths.

A key component in our method is a novel conditioning scheme based on discriminative convolutional filter generation. Generating the future frames guided by filters generated from previous frames have been explored in Dynamic Filter Networks [Jia *et al.*, 2016]. A similar approach is taken in [Li *et al.*, 2018], where filters are generated conditioned on text. Both these approaches apply the generated filters on the intermediate responses of the generator network. This is computationally expensive and technically challenging especially for deeper generator architectures as we need to generate a large number of semantically meaningful filters. In contrast, applying multi-scale text-conditioned filters in the discriminator architecture as done in our approach, is computationally easier and effective as demonstrated by our experimental results.

3 Method

We first provide a formal description of the problem being addressed. We are given access to N data points $\{(\mathbf{v}^{(n)}, \mathbf{t}^{(n)})\}_{n=1}^N$ sampled from an underlying joint distribution $p(\mathbf{v}, \mathbf{t})$ in the video-sentence space. Here, each $\mathbf{v}^{(n)} \in \mathbb{R}^{T \times C \times W \times H}$ is a video clip and $\mathbf{t}^{(n)}$ is a sentence description. We are interested in learning a model capable of sampling from the unknown conditional distribution $p(\mathbf{v}|\mathbf{t})$. Similar to conditional GANs [Mirza and Osindero, 2014], we formulate the problem as learning a transformation function $\mathbf{G}(\mathbf{z}, \mathbf{t})$ from a known prior distribution $P_Z(\mathbf{z})$ and the conditional input variable \mathbf{t} to the unknown conditional distribution $p(\mathbf{v}|\mathbf{t})$. The function \mathbf{G} is optimized using an adversarial training procedure.

3.1 Model Framework

Our conditional GAN framework is shown in Fig. 1, in which the video generator employs a RNN as in MoCoGAN [Tulyakov *et al.*, 2018]. Unlike MoCoGAN, our model does not have separate **Motion** and **Content** modeling or any discrete motion/action vector. The novelty of our proposed framework involves a multi-scale text-conditioning scheme,

in which the text description \mathbf{t} is passed to a text encoder \mathbf{T} to get a text representation \mathbf{t}_e .

$$\mathbf{t}_e = \mathbf{T}(\mathbf{t}), \quad (1)$$

where \mathbf{t} is a word embedding matrix containing pre-trained distributed representation vectors of the words in a given text description. The text encoder $\mathbf{T}(\cdot)$ is a CNN containing several convolutional layers and a fully connected layer (RNN-based model can also be used). The encoded text representation along with a sequence of noise vectors are passed to a GRU recurrent neural network to produce a trajectory in the latent space as follows,

$$\mathbf{h}_m = \text{GRU}([\mathbf{h}_{m-1}, \mathbf{t}_e, \mathbf{z}_m]), \quad (2)$$

$$\mathbf{z}_m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (3)$$

where $m = 1, \dots, l$ and l denotes the number of frames in the video sequence. This sequence of latent vectors is then passed to a shared frame generator model G to produce the video sequence,

$$\mathbf{v}_m = \mathbf{G}(\mathbf{h}_m), \quad m = 1, \dots, l. \quad (4)$$

The generated video is then fed to two discriminator models - $\mathbf{D}^{(f)}$ and $\mathbf{D}^{(v)}$. $\mathbf{D}^{(f)}$ is a frame-level discriminator that classifies if the individual frames in the video are real/fake, whereas the video discriminator $\mathbf{D}^{(v)}$ is trained to classify the entire video as real/fake. The discriminator models $\mathbf{D}^{(f)}$ and $\mathbf{D}^{(v)}$ also, respectively, take the text encoding \mathbf{t}_e as input to enforce text conditioning. Unlike the MoCoGAN model [Tulyakov *et al.*, 2018], we employ much deeper ResNet-style architectures in generator and discriminator networks for high-quality video generation. More details can be found in the supplementary material¹.

3.2 Text-Filter Conditioning

To build strong conditional models, it becomes important to learn good video-text associations in the GAN model. A standard technique is to sample negative (\mathbf{v}, \mathbf{t}) pairs (wrong associations) and train it as fake class, while the correct (\mathbf{v}, \mathbf{t}) pairs are trained as real class in the discriminator network. Since the generator is updated using the gradients from the discriminator, it becomes important to effectively fuse the video and text representations in the discriminator so as to make the generator condition on the text reasonably well. Previous methods [Li *et al.*, 2018; Pan *et al.*, 2017] use a simple concatenation of text and video features as the feature fusion strategy. We found that this simple strategy produces poor conditional models in datasets where there are rich text-video variations (refer to Section. 4 for more details).

Our proposed Text-Filter conditioning GAN (TFGAN) model aims at improving text conditioning for video generation models. In TFGAN, we employ a scheme based on generating discriminative convolutional filters from text features, which are then convolved with image features in the discriminator network. This scheme, which we call Text-Filter conditioning, is shown in Fig. 2.

¹Supplementary material can be found at <https://tinyurl.com/y3jedfxy>

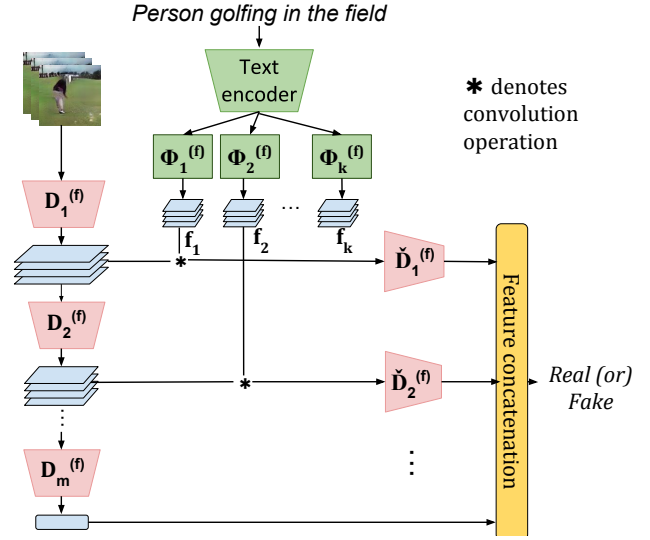


Figure 2: Illustration of our Text-Filter conditioning strategy.

Let the discriminator network $\mathbf{D}^{(f)}$ (equivalently for $\mathbf{D}^{(v)}$) be divided into multiple sub-networks $\{\mathbf{D}_i^{(f)}\}_{i=1}^m$ such that the entire discriminator response can be represented as a cascade of these sub-network responses,

$$\mathbf{D}^{(f)}(\mathbf{x}) = \mathbf{D}_m^{(f)} \circ \mathbf{D}_{m-1}^{(f)} \circ \dots \circ \mathbf{D}_1^{(f)}(\mathbf{x}), \quad (5)$$

where these sub-networks $\{\mathbf{D}_i^{(f)}\}_{i=1}^m$ can be as small as a single layer, or can be a cascade of multiple layers. Let $\mathbf{d}_i^{(f)}$ denote the output of the i^{th} sub-network of the frame discriminator,

$$\mathbf{d}_i^{(f)} = \mathbf{D}_i^{(f)} \circ \mathbf{D}_{i-1}^{(f)} \dots \mathbf{D}_1^{(f)}(\mathbf{v}).$$

From the text features, we generate a set of convolution filters $\{\mathbf{f}_i^{(f)}\}_{i=1}^m$ for frame-level discrimination,

$$\mathbf{f}_i^{(f)} = \Phi_i^{(f)}(\mathbf{t}_e), \quad i = 1, \dots, m, \quad (6)$$

where $\Phi_i^{(f)}$ (equivalently $\Phi_i^{(v)}$ at video-level) is a discriminative convolutional filter generation network at frame level. Similarly, let $\mathbf{d}_i^{(v)}$ denote the output of the i^{th} sub-network of the video discriminator, and $\{\mathbf{f}_i^{(v)}\}_{i=1}^m$ denote the discriminative convolutional filters at the video level. The frame-level filters $\mathbf{f}_i^{(f)} \in \mathbb{R}^{n_{in} \times n_{out} \times w \times h}$, whereas the video-level filters $\mathbf{f}_i^{(v)} \in \mathbb{R}^{n_{in} \times n_{out} \times c \times w \times h}$. Here, n_{in} , n_{out} are the number of input and output channels of $\mathbf{d}_i^{(i)}$, and c , w and h are the filter dimensions.

Each filter $\mathbf{f}_i^{(f)}$ is now convolved with the discriminator response $\mathbf{d}_i^{(f)}$, and this convolved output is passed through additional convolution layers $\tilde{\mathbf{D}}_i^{(f)}$ to get a vector $\tilde{\mathbf{d}}_i^{(f)}$ as output. These vectors $\{\tilde{\mathbf{d}}_i^{(f)}\}_{i=1}^m$ are then concatenated to obtain the final frame-text representation $\mathbf{O}^{(f)}$, which is used to classify the frame-text pair as real or fake.

$$\tilde{\mathbf{d}}_i^{(f)} = \tilde{\mathbf{D}}_i^{(f)}(\mathbf{f}_i^{(f)} * \mathbf{d}_i^{(f)}), \quad (7)$$

$$\mathbf{O}^{(f)} = [\tilde{\mathbf{d}}_1^{(f)}, \tilde{\mathbf{d}}_2^{(f)}, \dots, \tilde{\mathbf{d}}_m^{(f)}], \quad (8)$$

where $[\cdot]$ denotes vector concatenation. The video-text representation $\mathbf{O}^{(v)}$ is generated using a similar procedure. Since the generated convolutional filters $\{\mathbf{f}_i\}$ are applied to discriminator sub-network outputs $\{\mathbf{d}_i\}$ from different layers, the resulting text conditioning effectively imposes semantic constraints extracted from input text to the generated individual frames and video clips at different feature abstraction levels.

3.3 Training Algorithm

The discriminator models $\mathbf{D}^{(f)}$ and $\mathbf{D}^{(v)}$, and the generator model \mathbf{G} are trained using an adversarial game as done in the standard conditional GANs [Mirza and Osindero, 2014]. However, since we employ deep Resnet-style architectures for generator and discriminator networks, it was important to stabilize the GAN training. We use the regularizer as proposed in [Mescheder *et al.*, 2018] where the norm of the discriminator gradients is penalized. The equations for optimization updates are provided in the Supplementary material.

4 Experiments

This section discusses the experimental validation of our TFGAN model. We first describe a benchmark synthetic dataset we created for the task of text-to-video generation that captures rich video-text associations. The performance of our system is studied exhaustively using this dataset. Then, results are presented on two challenging real-world video dataset - Kinetics human action video dataset [Kay *et al.*, 2017] and Epic-Kitchens dataset [Damen *et al.*, 2018]. Further experiments and ablation studies are shown in the supplementary material. Our code is publicly available at https://github.com/minrq/CGAN_Text2Video.

4.1 Moving Shapes Dataset

Dataset Creation

To better understand the task of text-to-video synthesis, we created a dataset of moving shapes where a shape moves along a trajectory as described by the corresponding text description. Two versions of the dataset were created – one with static background which we call *Shapes-v1 dataset* and one with dynamic background which we call *Shapes-v2 dataset*. Some samples from these datasets are shown in Fig. 3. More details about the dataset creation are provided in the Supplementary material.

Quantitative Evaluation

One important benefit of creating the synthetic dataset is that it provides a framework for quantitative evaluation of the text-conditioning. First, five attribute classifiers (shape, size, color, motion and direction classifiers) are trained on the real data using the ground truth attributes (we have access to ground-truth attributes as they were stored while creating the dataset). These trained attribute classifiers are then used to verify if the attributes of a generated video match the attributes corresponding to the input text. This *attribute classification accuracy* is used as a metric to compare GAN models. To evaluate a GAN model, videos are generated for each text description in the test set, and the average attribute classification accuracy for these generated videos is reported. Higher the score, better is the model.

We compare TFGAN with the following models:

- Simple concat (SC): a conditional GAN model trained using simple text-video feature concatenation in the discriminator network
- Simple concat + Multiscale D (SC+MD): a conditional GAN model with discriminator network employing a multi-scale architecture. More specifically, outputs at the intermediate layers of the discriminator network are extracted and pooled together to get a multi-scale feature representation. This feature is then used to classify if the video-text pair is real or fake. Similar to the previous model, a simple feature concatenation is used to perform text conditioning.
- Text-filters on generator (TF on G): Instead of applying the text-filters on the discriminator network as done in TFGAN, filters extracted from the text are convolved with the intermediate responses of the generator network. In the discriminator network, simple feature concatenation is used to perform the conditioning. Note that this model is similar to the approach in [Li *et al.*, 2018] where the filters extracted from text are applied to the gist image (which is a part of the generator network).

More details about the architecture and hyper-parameters are given in the supplementary material. Some sample generations of TFGAN model are shown in Fig. 3.

Table 1a reports the quantitative evaluation on the *Shapes* dataset. We observe that TFGAN with text-filter conditioning achieves the best performance compared to other models. Only marginal performance gain is obtained by using filters in the generator network. We argue that this is because generating discriminative filters are much easier than generative filters. Also, the use of multi-scale architecture in discriminator alone does not improve text conditioning. This shows that performance improvement of TFGAN does not come from multi-scale architecture in itself, but in how the text conditioning is applied.

Exploratory Experiments

Sentence Interpolation: In this experiment, we depict conditional interpolation whereby frames in a video transition corresponding to the interpolation between two sentence descriptions. Let S_1 and S_2 denote the two sentences that are interpolated, and $(\mathbf{t}_{S_1}^{(f)}, \mathbf{t}_{S_1}^{(v)})$ and $(\mathbf{t}_{S_2}^{(f)}, \mathbf{t}_{S_2}^{(v)})$ denote their corresponding feature representation obtained from the text encoder \mathbf{T} . For each frame to be generated, the corresponding conditioning feature is obtained by a linear interpolation between these two representations:

$$(\mathbf{t}_i^{(f)}, \mathbf{t}_i^{(v)}) = (1 - \alpha)(\mathbf{t}_{S_1}^{(f)}, \mathbf{t}_{S_1}^{(v)}) + \alpha(\mathbf{t}_{S_2}^{(f)}, \mathbf{t}_{S_2}^{(v)})$$

Instead of using a fixed text representation $(\mathbf{t}^{(f)}, \mathbf{t}^{(v)})$ to condition all frames in the generator, we use $(\mathbf{t}_i^{(f)}, \mathbf{t}_i^{(v)})$ as input to the frame i . Some sample interpolations are shown in Fig. 4. We observe smooth interpolations corresponding to the input sentence transitions. When interpolating between blue square and red square, some intermediate frames are generated with pink shade. Interestingly, none of the samples in the dataset contain pink color. In the second figure,

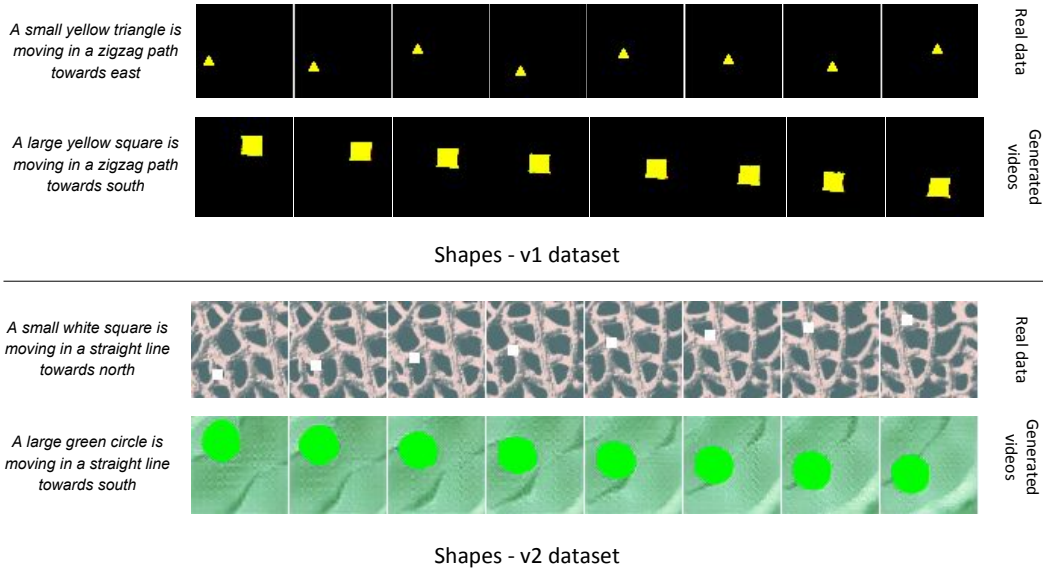


Figure 3: Samples from the Moving Shapes Dataset: In each set of images, the top row corresponds to the original video and the bottom two rows correspond to the video generated by our TFGAN model (Better viewed in color).

Method	Shape	Color	Size	Motion	Direction
<i>Shapes-v1 dataset</i>					
SC	70.18	99.23	83.69	96.83	99.00
SC + MD	65.25	99.91	74.31	99.12	99.11
TF on G	72.91	99.41	82.98	99.51	99.81
TFGAN	97.66	99.99	98.60	99.40	100.00
<i>Shapes-v2 dataset</i>					
SC	55.88	96.04	73.71	83.30	93.51
SC + MD	56.12	97.71	69.15	87.20	92.00
TF on G	59.76	95.75	82.16	85.00	95.40
TFGAN	81.10	99.99	89.82	99.80	100.00

(a) Attribute classification accuracy (in %) on *Shapes* dataset

Attribute	Accuracy (in %)
Shape	96.21
Color	99.78
Size	98.77
Motion	96.23
Direction	99.42

(b) Attribute classification accuracy on novel categories on *Shapes-v1* dataset

Table 1: Quantitative analysis: *Shapes* dataset




Figure 4: Exploratory experiments on *Shapes-v1* dataset. The image on the top shows the long sequence experiment where we generate 32-length sequence from a model trained for 16 frames. The top row of this video are the first 16 frames and the bottom row corresponds to the next 16. The images on the bottom illustrate the interpolation experiments where we generate a video corresponding to a smooth transition between two input sentences.

a smooth decrease in object size is observed as the object moves in the specified trajectory.

Generating Novel Categories: To determine if the model has learned to generalize and not naively memorize the dataset, this experiment aims to study the ability of TFGAN to produce videos not seen during training. Of the 360 unique parameter configurations in the *Shapes* dataset, $n = 20$ configurations are held-out from the training set. After the model is trained on this dataset, the attribute classification accuracy is measured on the held-out configurations. The results are reported in Table 1b. Good accuracy on held-out categories are achieved and this illustrates the ability of TFGAN to generalize.

Long Sequence Generation: One of the benefits of using a RNN-based GAN model is that it helps generate video sequences of varying length. To demonstrate this, we perform an experiment where TFGAN model is trained on 16-length video sequences, and made to generate 32-length sequences at test time. Fig. 4 shows the output of one such generated




Figure 5: Sample generations by models trained on Kinetics dataset

32-length sequence. We observe that the model is able to clearly perform the zig-zag motion beyond 16 frames, the frame length for which it was trained for.

4.2 Kinetics Dataset

To illustrate the practical relevance of our approach, we perform experiments on real-world video datasets. We use the dataset proposed in [Li *et al.*, 2018] for this purpose. This dataset contains videos of human actions, and was curated from YouTube and Kinetics human action video dataset [Kay *et al.*, 2017]. The following action classes are represented in the dataset - ‘biking in snow’, ‘playing hockey’, ‘jogging’, ‘playing soccer ball’, ‘playing football’, ‘kite surfing’, ‘playing golf’, ‘swimming’, ‘sailing’ and ‘water skiing’. This is an extremely challenging dataset for the task of video generation due to the following reasons: (1) videos are extremely diverse with large variations within each video, and (2) many videos have low-resolution and poor-quality video frames.

Method	Acc. (%)	FID-img	FID-vid
In-set	78.1	-	-
T2V([Li <i>et al.</i> , 2018])	42.6	82.13	14.65
Ours-SC	74.7	33.51	7.34
Ours-TFGAN	76.2	31.76	7.19

Table 2: Quantitative evaluation on Kinetics dataset

Some qualitative results on the Kinetics dataset are shown in Fig. 5. We observe that TFGAN is able to produce videos of much higher quality than the comparison method [Li *et al.*, 2018]. Fine-grained motions such as golf swing is generated by TFGAN while [Li *et al.*, 2018] produces a blobby region. One reason for improved generation quality is that TFGAN was successfully trained on 128×128 image resolution, while [Li *et al.*, 2018] was trained on 64×64 . More qualitative results and generated videos are provided in the supplementary material.

The following metrics are used for quantitative evaluation: (1) Classification accuracy: As done in [Li *et al.*, 2018], a modified version of inception score is computed, in which a video classification model is trained on real data, and the accuracy on generated data is reported. The performance is reported on the following categories as done in [Li *et al.*, 2018]: ‘kite surfing’, ‘playing golf’, ‘biking in snow’, ‘sailing’, and ‘swimming’. (2) Frame-level FID: Frames are extracted from the videos and the *Fréchet Inception Distance* between the real data frames and the generated frames are compute. As originally proposed [Heusel *et al.*, 2017], pre-trained Inception network is used extract the features for computing the score (3) Video-level FID: Features of the penultimate layer are extracted from 3D Resnet-50 model trained on the entire Kinetics dataset [Kay *et al.*, 2017], and the FID score is computed between the real and generated videos. Note that lower the FID scores, better are the models. Quantitative results computed using these metrics are reported in Table. ??.

In-set refers to the performance obtained on the test set of real videos. We observe that TFGAN achieves significantly higher scores than [Li *et al.*, 2018] over all three metrics. In this dataset, the performance gain primarily comes from using stronger generative models as the text descriptions are relatively simple and the degree of text-video variation is limited.

5 Conclusion

In this work, we address the problem of generating videos conditioned on text. A novel conditioning scheme is proposed in which text conditioning is performed using convolution operations acting on discriminator feature maps with filters generated from text. To better understand the text conditioning, we construct a synthetic dataset and show that our conditioning scheme achieves superior performance compared to other techniques. Finally, by using deeper architectures in the discriminator and generator networks, we generate high-quality videos on the challenging Kinetics dataset.

References

- [Brock *et al.*, 2018] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *CoRR*, abs/1809.11096, 2018.
- [Damen *et al.*, 2018] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018.
- [Goodfellow *et al.*, 2014] 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 27*, pages 2672–2680. 2014.
- [Heusel *et al.*, 2017] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems 30*, pages 6626–6637. 2017.
- [Jia *et al.*, 2016] Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc Van Gool. Dynamic filter networks. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 667–675, 2016.
- [Karras *et al.*, 2017] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *CoRR*, abs/1710.10196, 2017.
- [Kay *et al.*, 2017] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017.
- [Kingma and Welling, 2014] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR 2014*, April 2014.
- [Li *et al.*, 2018] Yitong Li, Martin Renqiang Min, Dinghan Shen, David E. Carlson, and Lawrence Carin. Video generation from text. In *AAAI*, 2018.
- [Mescheder *et al.*, 2018] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Which training methods for gans do actually converge? In *International Conference on Machine Learning (ICML)*, 2018.
- [Mirza and Osindero, 2014] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *ICLR Workshop*, 2014.
- [Odena *et al.*, 2017] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2642–2651, 2017.
- [Pan *et al.*, 2017] Yingwei Pan, Zhaofan Qiu, Ting Yao, Houqiang Li, and Tao Mei. To create what you tell: Generating videos from captions. 2017.
- [Reed *et al.*, 2016] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 1060–1069, 2016.
- [Salimans *et al.*, 2016] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems 29*, pages 2234–2242. 2016.
- [Tulyakov *et al.*, 2018] Sergey Tulyakov, Ming-Yu Liu, Xiao-dong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [Vondrick *et al.*, 2016] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *Advances in Neural Information Processing Systems 29*, pages 613–621. 2016.
- [Wang *et al.*, 2018a] Heng Wang, Zengchang Qin, and Tao Wan. Text generation based on generative adversarial nets with latent variables. In *Advances in Knowledge Discovery and Data Mining - 22nd Pacific-Asia Conference, PAKDD 2018, Melbourne, VIC, Australia, June 3-6, 2018, Proceedings, Part II*, 2018.
- [Wang *et al.*, 2018b] Hongwei Wang, Jia Wang, Jialin Wang, Miao Zhao, Weinan Zhang, Fuzheng Zhang, Xing Xie, and Minyi Guo. Graphgan: Graph representation learning with generative adversarial nets. In *AAAI*, 2018.
- [Xu *et al.*, 2018] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiao-dong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [Yan *et al.*, 2016] Xinchun Yan, Jimei Yang, Kihyuk Sohn, and Honglak Lee. Attribute2image: Conditional image generation from visual attributes. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, 2016.
- [Zhang *et al.*, 2017] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017.
- [Zhu *et al.*, 2017] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.