Recurrent Deconvolutional Generative Adversarial Networks with Application to Video Generation

Hongyuan Yu[1,2] , Yan Huang[1,2] , Lihong Pi[3] , and Liang Wang[1,2,4(✉)]

[1] University of Chinese Academy of Sciences (UCAS), Beijing, China
hongyuan.yu@cripac.ia.ac.cn
[2] Center for Research on Intelligent Perception and Computing (CRIPAC),
National Laboratory of Pattern Recognition (NLPR), Beijing, China
{yhuang,wangliang}@nlpr.ia.ac.cn
[3] The Institute of Microelectronics, Tsinghua University (THU), Beijing, China
plh17@mails.tsinghua.edu.cn
[4] Chinese Academy of Sciences Artificial Intelligence Research (CAS-AIR),
Beijing, China

Abstract. This paper proposes a novel model for video generation and especially makes the attempt to deal with the problem of video generation from text descriptions, i.e., synthesizing realistic videos conditioned on given texts. Existing video generation methods cannot be easily adapted to handle this task well, due to the frame discontinuity issue and their text-free generation schemes. To address these problems, we propose a recurrent deconvolutional generative adversarial network (RD-GAN), which includes a recurrent deconvolutional network (RDN) as the generator and a 3D convolutional neural network (3D-CNN) as the discriminator. The RDN is a deconvolutional version of conventional recurrent neural network, which can well model the long-range temporal dependency of generated video frames and make good use of conditional information. The proposed model can be jointly trained by pushing the RDN to generate realistic videos so that the 3D-CNN cannot distinguish them from real ones. We apply the proposed RD-GAN to a series of tasks including conventional video generation, conditional video generation, video prediction and video classification, and demonstrate its effectiveness by achieving well performance.

Keywords: Video generation · RD-GAN · GAN

1 Introduction

Image generation has drawn much attention recently, which focuses on synthesizing static images from random noises or semantic texts. But video generation, i.e., synthesizing dynamic videos including sequences of static images with temporal dependency inside, has not been extensively studied. In this work, we wish

---

Student first author.

to push forward this topic by generating better videos and dealing with a rarely investigated task of text-driven video generation.

The task here can be defined as follows: given a text describing a scene in which someone is doing something, the goal is to generate a video with similar content. This is challenging since it involves language processing, visual-semantic association and video generation together. A straightforward solution is to generalize the existing video generation models for this task. But it is not optimal due to the following reasons: (1) conventional video generation directly generates videos from noises but not from semantic texts as our case, and (2) existing video generation models mostly suffer from the visual discontinuity problem, since they either directly ignore the modeling of temporal dependency of generated videos, or simply consider it in a limited range with 3D deconvolution.

In this paper, we propose a recurrent deconvolutional generative adversarial network (RD-GAN) for conditional video generation. The proposed RD-GAN first represents given semantic texts as latent vectors with skip-thoughts [8], and exploits a generator named recurrent deconvolutional network (RDN) to generate videos in a frame-by-frame manner based on the latent vectors. The RDN can be regarded as a deconvolutional version of conventional recurrent neural network by replacing all the full connections with weight-sharing convolutional and deconvolutional ones. Accordingly, its hidden states now are 2D feature maps rather than 1D feature vectors, which efficiently facilitates the modeling of spatial structural patterns. After the generation, the generated videos are then fed into a discriminator which uses 3D convolutional neural network (3D-CNN) to distinguish from non-generated real videos. The generator and discriminator in RD-GAN can be jointly trained with the goal to generate realistic videos that can confuse the discriminator. To demonstrate the effectiveness of our proposed RD-GAN, we perform various experiments in terms of conventional and conditional video generation, video prediction and classification.

Our contributions are summarized as follows. To the best of our knowledge, we make the attempt to study the problem of sentence-conditioned video generation, which is a rarely investigated but very important topic for the current research interest on generative models. We propose a novel model named recurrent deconvolutional generative adversarial networks to deal with the task, which is demonstrated to achieve good performance in a wide range of both generative and discriminative tasks.

2   Related Work

Image Generation. With the fast development of deep neural networks, the generative models have made great progress recently. As we know, Tijmen [19] proposed capsule networks to generate images and Dosovitskiy et al. [2] generated 3D chairs, tables and cars with deconvolutional neural networks. There are also some other works [15,26] using supervised methods to generate images. Recently, the unsupervised methods, such as variational auto-encoder (VAE) [7] and generative adversarial network (GAN) [3] have attracted much attention. Gregor et al. [4] found that they can generate simple images by imitating

human painting based on the recurrent variational autoencoder and the attention mechanism. The autoregressive models proposed by Oord et al. [21] modeled the conditional distribution of the pixel space and also achieved good experimental results. Compared with other methods, GAN has a relatively better performance on such tasks and various models [1,18,24,25] based on GAN generated appealing synthetic images. Beside that, conditional image generation has also been extensively studied. At first, simple variables, attributes or class labels [9,12] were used to generate specific images. Furthermore, researchers try to use unstructured text to do this work. For instance, Reed et al. [14] used the text descriptions and object location constraints to generate images with conditional PixelCNN. The later works [13,27] built upon conditional GAN yielded 64×64 or larger scale images of birds and flowers from text descriptions successfully.

Video Generation. Before we go into our video generation model, it is essential to review the recent advance related to video prediction and generation. [6,28] inspired us in terms of video continuity. Compared with video prediction, there is no context information in the video generation task. Vondrick et al. [22] first came up with a "violence" generative model which can directly yield fixed length videos from 3D deconvolution. However, the 3D deconvolution causes more serious loss of information than 1D and 2D deconvolution. In order to fix this problem, Saito et al. [16] proposed TGAN which tried to find all the latent vectors of continuous frames with the thought that videos are composed of images. And they used 1D and 2D deconvolution to generate video frame by frame. Unfortunately, their results do not have good continuity as they expected. MoCoGAN [20] decomposing motion and content for video generation, but their model lack understanding of semantics. In this work, we propose recurrent deconvolutional generative adversarial network (RD-GAN) including a recurrent deconvolutional network (RDN) to handle the current problems in video generation and well exploit the conditional constraints.

3   Recurrent Deconvolutional Generative Adversarial Network

The architecture of our proposed recurrent deconvolutional generative adversarial network for generating videos with text condition is shown in Fig. 1. The RD-GAN is built upon the conventional generative adversarial network (GAN), which has two sub-networks: a generator and a discriminator. The generator is a recurrent deconvolutional network (RDN) which tries to generate more realistic videos, whose input is sequential concatenated vectors of noise sampled from Gaussian distribution and text embedding by skip-thoughts [8]. While the discriminator is a 3D convolutional neural network (3D-CNN) which tries to distinguish input videos between two classes: "real" and "fake". During the binary classification, it also exploits the text embedding that concatenated with video feature maps in the discriminator. These two sub-networks are jointly trained by playing a non-cooperative game, in which the generator RDN tries to fool the discriminator 3D-CNN while the discriminator aims to make few mistakes. Such
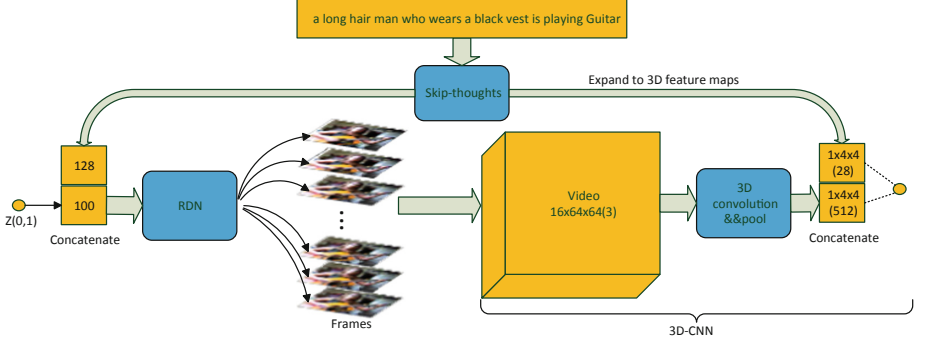
Fig. 1. The proposed recurrent deconvolutional generative adversarial network (RD-GAN) for generating videos with text condition. The sentence is encoded into an 28-dimensional vector by skip-thoughts [8] and linear operation. In the generator, the text vector and an 100-dimensional vector sampled from Gaussian distribution are concatenated as input. Then the text vector is copied 4 times and the same 16 text vectors are stacked as $28 \times 1 \times 4 \times 4$ feature maps before being fed into the penultimate layer of discriminator.

a competitive scheme has been demonstrated to be effective for training generative models [3]. In the following, we will present the generator and discriminator in details.

3.1 Recurrent Deconvolutional Network as Generator

The task of video generation based on semantic texts mainly encounters two challenges. One is how to extract suitable information from texts and associate it with the content of generated videos. It can be solved by exploiting recent advances in areas of natural language processing and multimodal learning. The other one is how to well model both long-range and short-range temporal dependencies of video frames, which focus on global slow-changing and local fast-changing patterns, respectively. The exist methods either just consider the local fast-changing ones with 3D convolutions or directly ignore the modeling of temporal dependency by treating each frame generation as a separate procedure.

To well model the long-term temporal dependency during video generation, we propose a recurrent deconvolutional network (RDN) as shown in Fig. 2. The RDN has a deep spatio-temporal architecture, whose input is a combination of noise and the vector extracted from the given text and the output is a sequence of generated video frames. Weight sharing is widely used in the temporal and the spatial direction, which can effectively reduce the number of parameters and contribute to model stability.

The whole network can be regarded as a nonlinear mapping from semantic texts to desired videos. In particular, in the spatial direction, there is a deep deconvolution network at each timestep for frame-wise generation. Between

adjacent levels of feature maps, deconvolution, batch normalization and ReLU activation are successively used to upscale the frame size by a factor of 2.

In the temporal direction, it is a broad recurrent convolution network, in which pairwise feature maps at adjacent timesteps are connected by convolutions. It means that when predicting one video frame, the result can be directly modulated by its previous frame and recursively depends on other previous frames in a long temporal range. The formulation of inferring feature maps $H_{i,t}$ at the i-th level and the t-th timestep is:

$$H_{i,t} = a(b(H_{i-1,t} \mathbin{\hat{*}} W_{i-1,t} + H_{i,t-1} * U_{i,t-1} + B_{i,t})) \tag{1}$$

where $*$ and $\hat{*}$ represent temporal convolution and spatial deconvolution operations, respectively. $U_{i,t-1}$ and $W_{i-1,t}$ contain the filter weights of temporal convolution and spatial deconvolution, respectively. $B_{i,t}$ denotes the bias weights. $a(\cdot)$ and $b(\cdot)$ are ReLU activation function and batch normalization operation, respectively. Due to the recurrent scheme, our RDN can flexibly produces videos with any length.
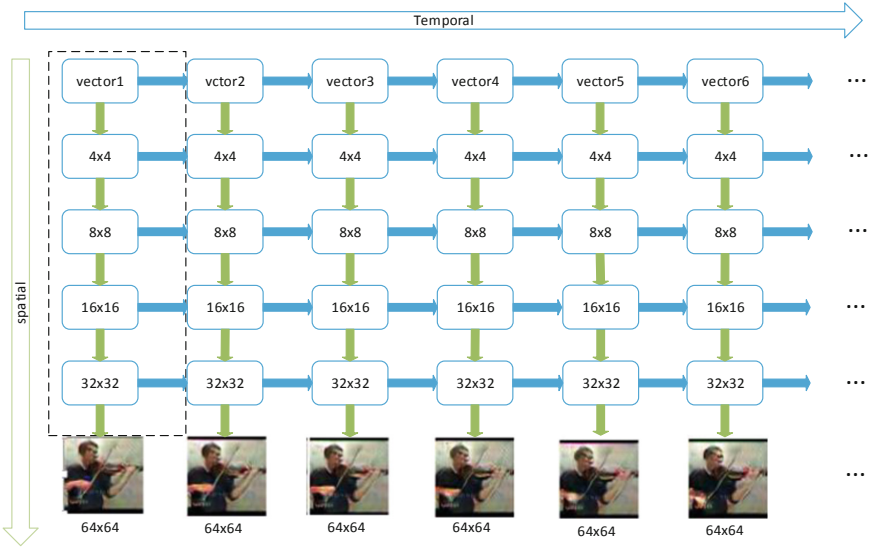

Fig. 2. The proposed recurrent deconvolutional network as the generator. The input is vector1 and all vectors in the top level are 1D vectors connected by linear transformation. The rest are 2D feature maps connected by convolutions and we use blue arrows to represent them. While green arrows refer to deconvolution, which enlarges the scale of images in the spatial direction. The sizes of images and feature maps are annotated in the figure. All the different parameters are in the dashed box, and the rest are consistent with these parameters. (Best viewed in color)

3.2   Modified 3D Convolutional Neural Network as Discriminator

The role of a discriminator is to identify the authenticity by classifying a given video into two classes: "real" and "fake". Considering that 3D convolutional neural network (3D-CNN) performs very well in the task of video classification, so we directly exploit it as our discriminator.

We make a few changes to the 3D-CNN as follows: (1) the last two linear layers are replaced by one 3D convolution layer, (2) the original sizes of some layers are reduced as the size of our video frame is 64×64, (3) 3D batch normalization, leakyReLU activation function and 3D max-pooling are used after each 3D convolution, and (4) the text embedding by skip-thoughts is combined in the penultimate layer.

3.3   Learning

The weights of both generator and discriminator can be jointly trained by using the following objective:

$$
\min_{\theta_G} \max_{\theta_D} \mathbb{E}_{x \sim p_{data}(x)}[\log D(x|t; \theta_D)]+
$$
$$
\mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z; \theta_G)|t; \theta_D))]
\tag{2}
$$

where θ_G and θ_D represent the parameters of generator and discriminator, respectively. $p_{data}(x)$ is the distribution of real videos x. z denotes the noise which is sampled from Gaussian distribution and t is the text condition. The object will reach the global optimum when $p_{G(z)} = p_{data}$. However, it is usually infeasible to obtain the global optimum, so we usually use gradient-based methods such as stochastic gradient descent (SGD) [11] to find a good local equilibrium.

All the network weights are initialized by sampling from a Gaussian distribution with a mean of 0 and standard deviation of 0.02. We use the ADAM solver to optimize all the parameters, in which the learning rate is 0.0002 and the momentum is 0.5. To speed up the training procedure, we first use all the images in video segments to train an image-based GAN, which has the same architecture in the spatial direction as the proposed RD-GAN. Then we use the learned weights as pretrained weights for the RD-GAN and fine-tune all the weights on videos. In fact, we can alternatively remove such a pretraining step with a longer training time.

4   Experiments

4.1   Dataset and Implementation Details

The UCF-101 [17] dataset is used during the training process. It contains 13,320 videos belonging to 101 different classes of human actions. Because there are many classes of videos in the UCF-101, it is very difficult for our model to

learn such a complex data distribution. As a result, we use videos belonging to the same class to train a separate network every time. Note that the whole UCF-101 dataset is used when we evaluate the representations taken from the discriminator.

To enlarge our training dataset, we divide each video into multiple video segments containing 16 consecutive frames. For example, frames 1 to 16 make up the first video segment, the second video segment consists of frames 2 to 17 and the third video segment consists of frames 3 to 18, and so on. So a single class of video can get about 20,000 video segments and the included frames are all resized to the size of 64×64.

For the semantic texts associated with videos, we make the attempt to give different human-written text descriptions for different videos. For example, "a cute little boy who wears a red headband and a black shirt is playing Violin" or "a curly young man who wears a black T-shirt is playing Guitar". Videos with similar content and background are put into the same class and named with the same text description because those videos in the UCF-101 are divided from the same long video.

Fig. 3. Results of conditional video generation.

Fig. 4. Results of conditional video generation with "unfamiliar" sentence.

4.2   Video Generation from Sentence

Generating a specific video from human-written texts is equivalent to imagining a scene by our brain when we look at a novel. To adapt our model to the task of text-driven video generation, we exploit text information for both generator and discriminator as shown in Fig. 1. We use skip-thoughts [8] to encode a given text into a 4800-dimensional vector, and then map the vector from 4800 to 28 dimensions. After that, we concatenate the 28-dimensional text vector with an 100-dimensional vector that is sampled from Gaussian distribution, then the new 128-dimensional vector is used as the input of the generator. We also feed the same concatenated vector to the discriminator, which was extended to 3D data before combining with the penultimate layer of the 3D-CNN.

To simply annotate video data, we just gather similar videos into the same class and give them the same text descriptions. Firstly, we generate videos from already known texts as shown in Fig. 3. From this figure we can see that our model is able to generate semantic-related videos. For example, when given the text: "a long hair man who wears a black vest is playing Guitar", our model can accordingly generate video with attributes: "long hair", "man" and "black vest". Then we try to change some attributes in sentence and send these "unfamiliar" sentences to our model. Figure 4 shows our model can generate new samples with the corect attributes. The result is not clear because the amount of data is not enough. In other words there is the problem of lack of continuous mapping between the semantic space and the image space.

Table 1. Accuracy of unsupervised methods on the UCF-101 dataset.

Method	Accuracy
Chance	0.9%
STIP Features [17]	43.9%
Temporal Coherence [5]	45.4%
Shuffle and Learn [10]	50.2%
VGAN + Logistic Reg [22]	49.3%
VGAN + Fine Tune [22]	52.1%
TGAN + Linear SVM [16]	38.7%
Ours + Convolution Softmax	**53.3%**
Ours + Linear Softmax	**55.7%**
ImageNet Supervision [23]	91.4%

4.3   Video Classification

Considering that the features in our discriminator are learned without supervision, we want to quantitatively analyze the capacity of discrimination on the task of video classification. Since the discriminator was originally used for binary
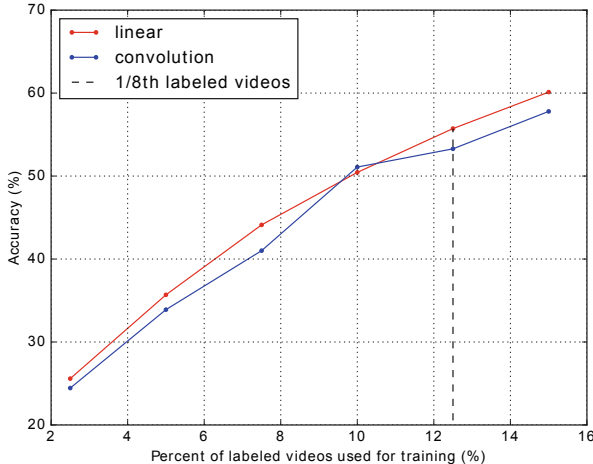
Fig. 5. Ratio of labeled training videos vs Accuracy.

classification, we have to replace its last layer with the softmax classifier for multi-class classification. All the videos in the UCF-101 dataset are used in this experiment. By segmenting all the videos into 16-frame segments, we can obtain totally 6 million video segments. The class for each video segment is in consistent with the class of its original video.

We first use the same softmax and dropout as [22] in the last two layers, which uses convolution to make the last feature maps transform into an 101-dimensional vector. We follow them and use 1/8 labeled videos to train the softmax classifier. Table 1 shows that our model improves the accuracy by 1.2%. Considering that linear operation can also convert the last feature maps into an 101-dimensional vector for classification, we use liner operation to replace the previously used convolution. To our surprise, this liner operation further improves our performance by 2.4%. Note that both convolution and linear operations have the same number of parameters. Obviously, models that leverage external supervision are still much better than unsupervised methods. In Fig. 5, the performances of convolution softmax and linear softmax are also compared by using different numbers of labeled videos, and we can observe that the more number of training data, the higher the accuracy. On the whole, linear softmax produces more discrimative results than convolution softmax.

5   Conclusions and Future Work

We have proposed a recurrent deconvolutional generative adversarial network that generates videos based on semantic texts. Its generator is a deconvolutional version of recurrent neural network, which is able to well exploit the text information and efficiently models the long-range temporal dependency during video generation. We have performed various experiments in terms of conventional

video generation, video generation driven by text, video prediction and video classification. The experimental results have demonstrated the effectiveness of our model.

Note that this is just an initial work on video generation, and it might have the following drawbacks. The proposed model based on GAN becomes unstable if trained with too many frames and currently cannot well generate clear videos of a larger size and a longer length. The processing of given text only involves simple steps of feature extraction. This might not be optimal to associate it with diverse video contents. In the future, we will consider to train the RDN with more frames and extend our model to generate videos with a larger size and a longer length in a cascade manner. Furthermore, we plan to expand the annotated videos in details for better text-driven video generation.

References

1. Bin, H., Weihai, C., Xingming, W., Chun-Liang, L.: High-quality face image SR using conditional generative adversarial networks. arXiv preprint arXiv:1707.00737 (2017)
2. Dosovitskiy, A., Tobias Springenberg, J., Brox, T.: Learning to generate chairs with convolutional neural networks. In: CVPR (2015)
3. Goodfellow, I., et al.: Generative adversarial nets. In: NeurIPS (2014)
4. Gregor, K., Danihelka, I., Graves, A., Rezende, D., Wierstra, D.: Draw: a recurrent neural network for image generation. In: ICML (2015)
5. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: CVPR (2006)
6. Kalchbrenner, N., et al.: Video pixel networks. arXiv:1610.00527 (2016)
7. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv:1312.6114 (2013)
8. Kiros, R., et al.: Skip-thought vectors. In: NeurIPS (2015)
9. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv:1411.1784 (2014)
10. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 527–544. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_32
11. Nemirovski, A., Yudin, D.: On Cezari's convergence of the steepest descent method for approximating saddle point of convex-concave functions. In: Soviet Math. Dokl (1978)
12. van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al.: Conditional image generation with PixelCNN decoders. In: NeurIPS (2016)
13. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis. In: ICML (2016)

14. Reed, S., van den Oord, A., Kalchbrenner, N., Bapst, V., Botvinick, M., deFreitas, N.: Generating interpretable images with controllable structure. Technical report (2016)
15. Reed, S.E., Zhang, Y., Zhang, Y., Lee, H.: Deep visual analogy-making. In: NeurIPS (2015)
16. Saito, M., Matsumoto, E., Saito, S.: Temporal generative adversarial nets with singular value clipping. In: ICCV (2017)
17. Soomro, K., Zamir, A.R., Shah, M.: UCF101: a dataset of 101 human actions classes from videos in the wild. arXiv:1212.0402 (2012)
18. Tan, W.R., Chan, C.S., Aguirre, H., Tanaka, K.: Improved ArtGAN for conditional synthesis of natural image and artwork. arXiv preprint arXiv:1708.09533 (2017)
19. Tieleman, T.: Optimizing neural networks that generate images. Ph.D. thesis (2014)
20. Tulyakov, S., Liu, M.Y., Yang, X., Kautz, J.: MoCoGAN: decomposing motion and content for video generation. In: CVPR (2018)
21. Van DenOord, A., Kalchbrenner, N., Kavukcuoglu, K.: Pixel recurrent neural networks. In: ICML (2016)
22. Vondrick, C., Pirsiavash, H., Torralba, A.: Generating videos with scene dynamics. In: NeurIPS (2016)
23. Wang, L., Xiong, Y., Wang, Z., Qiao, Y.: Towards good practices for very deep two-stream ConvNets. arXiv:1507.02159 (2015)
24. Wang, X., Gupta, A.: Generative image modeling using style and structure adversarial networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9908, pp. 318–335. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_20
25. Wu, H., Zheng, S., Zhang, J., Huang, K.: GP-GAN: towards realistic high-resolution image blending. arXiv preprint arXiv:1703.07195 (2017)
26. Yang, J., Reed, S.E., Yang, M.H., Lee, H.: Weakly-supervised disentangling with recurrent transformations for 3D view synthesis. In: NeurIPS (2015)
27. Zhang, H., et al.: StackGAN: text to photo-realistic image synthesis with stacked generative adversarial networks. In: ICCV (2017)
28. Zhou, Y., Berg, T.L.: Learning temporal transformations from time-lapse videos. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9912, pp. 262–277. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_16