

# Compositional Video Synthesis with Action Graphs

Amir Bar<sup>\*1</sup> Roi Herzig<sup>\*1</sup> Xiaolong Wang<sup>23</sup> Anna Rohrbach<sup>3</sup> Gal Chechik<sup>45</sup>  
Trevor Darrell<sup>3</sup> Amir Globerson<sup>1</sup>

## Abstract

Videos of actions are complex signals containing rich compositional structure in space and time. Current video generation methods lack the ability to condition the generation on multiple coordinated and potentially simultaneous timed actions. To address this challenge, we propose to represent the actions in a graph structure called Action Graph and present the new “Action Graph To Video” synthesis task. Our generative model for this task (AG2Vid) disentangles motion and appearance features, and by incorporating a scheduling mechanism for actions facilitates a timely and coordinated video generation. We train and evaluate AG2Vid on the CATER and Something-Something V2 datasets, and show that the resulting videos have better visual quality and semantic consistency compared to baselines. Finally, our model demonstrates zero-shot abilities by synthesizing novel compositions of the learned actions.<sup>1</sup>

## 1. Introduction

Learning to generate visual content is a fundamental task in computer vision and graphics, with numerous applications from sim-to-real training of autonomous agents, to creating visuals for games and movies. While the quality of generating still images has leaped forward recently (Brock et al., 2019; Karras et al., 2020), generating videos remains in its infancy. Generating actions and interactions between objects is perhaps the most challenging aspect of conditional video generation. When constructing a complex scene with multiple agents (e.g. a driving scene with several cars or a basketball game with a set of players) it would be important to have *control* over multiple agents and their actions as well

<sup>\*</sup>Equal contribution <sup>1</sup>The Blavatnik School of Computer Science, Tel Aviv University <sup>2</sup>UC San Diego <sup>3</sup>UC Berkeley <sup>4</sup>NVIDIA Research <sup>5</sup>Bar-Ilan University. Correspondence to: Amir Bar <amir.bar@cs.tau.ac.il>.

<sup>1</sup>See the project page for code and pretrained models:

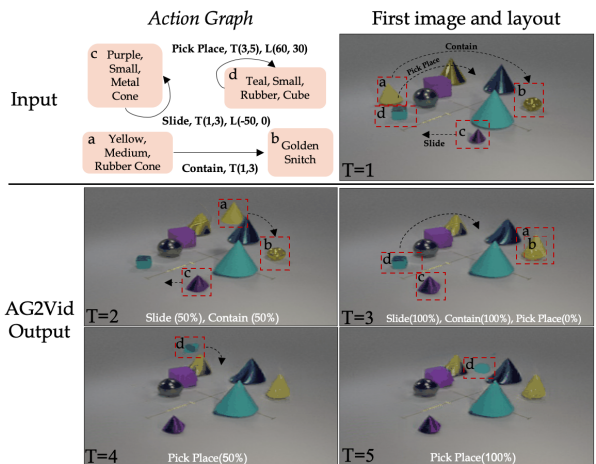


Figure 1. We propose a new task called *Action Graph to Video*. To represent input actions, we use a graph structure called *Action Graph*, where nodes are objects and edges are actions with start/end frames specified by  $T(s, e)$  and optional attributes (such as object destination coordinates) specified by  $L(x, y)$ , where  $(x, y) \in \mathbb{R}^2$  is a spatial offset. Together with input first frame and scene layout, our goal is to synthesize a video that follows the input Action Graph instructions. We include an example over CATER (Girdhar & Ramanan, 2020) dataset. The annotated object-action trajectories are only shown for illustration purposes.

as to *coordinate* them w.r.t. each other (e.g. to generate a Pick 'n' Roll move in basketball). Actions create long-range spatio-temporal dependencies between people and objects they interact with. For example, when a player passes a ball, the entire movement sequence of all the entities (thrower, ball, receiver) must be coordinated and carefully timed. This paper specifically addresses this challenge, the task of generating coordinated and timed actions, as an important step towards generating videos of complex scenes.

Current approaches for conditional video generation are not well suited to condition the generation on multiple coordinated (potentially simultaneous) timed actions. One line of work, *Class Conditional Video Generation* (Tulyakov et al., 2018; Nawhal et al., 2020a) relies on a single action (Tulyakov et al., 2018) and a potential object class (Nawhal et al., 2020a). However, this approach cannot deal with variable number of objects, multiple simultaneous

actions and does not control the timing of the actions. *Future Video Prediction* (Watters et al., 2017; Ye et al., 2019), generates future frames based on an initial input frame, but one frame will not suffice for generation of multiple coordinated actions. In another line of work, *Video-to-Video Translation* (Wang et al., 2018a; 2019), the goal is to translate a sequence of semantic masks into an output video. However, segmentation maps contain only object class information, and thus do not explicitly capture the action information. As Wang et al. (2018a) noted for the case of generating car turns, semantic-maps do not have enough information to distinguish car turns from a normally driving car, which leads to mistakes in video-translation. Finally, *Text-to-Video* (Li et al., 2018; Gupta et al., 2018) work is promising because language can describe complex actions. However, in applications that require a precise description of a scene, natural language is not ideal due to ambiguities (MacDonald et al., 1994), user subjectivity (Wiebe et al., 2004) and difficulty to infer specific timings, if such information is missing. Hence, we address this problem with a more structured approach.

To generate videos based on actions, we propose to represent actions in an object-centric graph structure we call an “Action Graph” (AG) and define the new task of “Action Graph to Video” generation. An AG is a graph structure aimed at representing coordinated and timed actions. Its nodes represent objects and edges represent actions annotated with their start and end time (Figure 1). AGs are an intuitive representation for describing timed actions and would be a natural way to provide precise inputs to generative models. A key advantage of AGs is their ability to describe the dynamics of objects and actions precisely in a scene. In the “Action Graph to Video” task, the input is an initial frame of the video with the ground-truth layout consisting of the objects classes and bounding boxes, and an AG.

There are multiple unique challenges in generating a video from an AG that cannot be addressed using current generation methods. First, each action in the graph unfolds over time, so the model needs to “keep track” of the progress of actions rather than just condition on previous frames as commonly done. Second, AGs may contain multiple concurrent actions and the generation process needs to combine them in a realistic way. Third, one has to design a training loss that captures the spatio-temporal video structure to ensure that the semantics of the AG is accurately captured.

To address these challenges, our AG2Vid model uses three levels of abstraction. First, an action scheduling mechanism we call “Clocked Edges” tracks the progress of actions over time. Based on this, a graph neural network (Kipf & Welling, 2016) operates on the AG with the Clocked Edges and predicts a sequence of scene layouts. Finally, pixels are generated conditioned on the predicted layouts. We apply our AG2Vid model to CATER (Girdhar & Ramanan, 2020)

and Something-Something V2 (Goyal et al., 2017) datasets and show that our approach results in videos that are more realistic and semantically compliant with the input AG than the baselines.

Last, we evaluate the AG2Vid model in a zero-shot setting where AGs involve novel compositions of the learned actions. The structured nature of our approach enables it to successfully generalize to new “composite” actions, as confirmed by the human raters. We also provide qualitative results of our model trained on Something-Something V2 in a challenging setting where AGs include multiple actions, while at training time each AG has a single action.

## 2. Related work

Video generation is challenging because videos often contain long range dependencies. Recent generation approaches (Vondrick et al., 2016; Denton & Fergus, 2018; Lee et al., 2018; Babaeizadeh et al., 2018; Villegas et al., 2019; Kumar et al., 2020) extended the framework of unconditional image generation to video, based on a latent representation. For example, MoCoGAN (Tulyakov et al., 2018) disentangles the latent space representations of motion and content to generate a sequence of frames using RNNs; TGAN (Saito et al., 2017) generates each frame in a video separately while also using a temporal generator to model dynamics across the frames. Here, we tackle a different problem by aiming to generate videos that comply with Action Graphs (AGs).

Class conditional action generation allows the generation of video based on a single action (Tulyakov et al., 2018; Nawhal et al., 2020b). A recent method, HOI-GAN (HG) (Nawhal et al., 2020b), was proposed for the generation task of a single action and object. Specifically, HG addresses the zero-shot setting, and the model is tested on action and object compositions which are first presented at test time. Our focus is on generation of multiple simultaneous actions over time, performed by multiple objects. Our approach directly addresses this challenge via the AG representation and the notion of clocked edges.

Other conditional video generation methods have attracted considerable interest recently, with focus on two main tasks: video prediction (Mathieu et al., 2015; Battaglia et al., 2016; Walker et al., 2016; Watters et al., 2017; Kipf et al., 2018; Ye et al., 2019) and video-to-video translation (Wang et al., 2019; Chan et al., 2019; Siarohin et al., 2019; Kim et al., 2019; Mallya et al., 2020). In video prediction, the goal is to generate future video frames conditioned on few initial frames. For example, it was proposed to train predictors with GANs (Goodfellow et al., 2014) to predict future pixels (Mathieu et al., 2015). However, directly predicting pixels is challenging (Walker et al., 2016). Instead of

pixels, researchers explored object-centric graphs and performed prediction on these (Battaglia et al., 2016; Luc et al., 2018; Ye et al., 2019). While inspired by object-centric representations, our method is different from these works as our generation is goal-oriented and guided by an AG. The video-to-video translation task was proposed by Bansal et al. (2018); Wang et al. (2018a), where a natural video is generated from a different video based on frame-wise semantic segmentation annotations or keypoints. However, densely labeling pixels for each frame is expensive, and might not even be necessary. Motivated by this, researchers have sought to perform generation conditioned on more accessible signals including audio or text (Song et al., 2018; Fried et al., 2019; Ginosar et al., 2019). Here, we propose to synthesize videos conditioned on an AG, which is cheaper to obtain compared to semantic segmentation and is a more structured representation compared to natural audio or text.

Various methods have been proposed to generate videos based on input text (Marwah et al., 2017; Pan et al., 2017; Li et al., 2018). Most recent methods typically used very short captions which do not contain complex descriptions of actions. For example, Li et al. (2018) used video-caption pairs from YouTube, where typical captions are "playing hockey" or "flying a kite". Gupta et al. (2018) proposed the Flinstones animated dataset and introduced the CRAFT model for text-to-video generation. While the CRAFT model relies on text-to-video retrieval, our approach works in an end-to-end manner and aims to accurately synthesize the given input actions.

Scene Graphs (SG) (Johnson et al., 2015; 2018) are a structured representation that models scenes, where objects are nodes and relations are edges. SGs have been widely used in various tasks including image retrieval (Johnson et al., 2015; Schuster et al., 2015), relationship modeling (Krishna et al., 2018; Schroeder et al., 2019; Raboh et al., 2020), SG prediction (Xu et al., 2017; Newell & Deng, 2017; Zellers et al., 2018; Herzig et al., 2018), and image captioning (Xu et al., 2019). SGs have also been applied to image generation (Johnson et al., 2018; Deng et al., 2018; Herzig et al., 2020), where the goal is to generate a natural image corresponding to the input SG. Recently, Ji et al. (2019) presented Action Genome, a video dataset where videos of actions from Charades (Sigurdsson et al., 2016) are also annotated by SGs. More generally, spatio-temporal graphs have been explored in the field of action recognition (Jain et al., 2016; Sun et al., 2018; Wang & Gupta, 2018; Yan et al., 2018; Girdhar et al., 2019; Herzig et al., 2019; Materzynska et al., 2020). For example, a space-time region graph was proposed by Wang & Gupta (2018) where object regions are taken as nodes and a GCN (Kipf & Welling, 2016) is applied to perform reasoning across objects for classifying actions. Recently, it was also shown (Ji et al., 2019; Yi et al., 2019; Girdhar & Ramanan, 2020) that a key obstacle in action

recognition is the ability to capture the long-range dependencies and compositionality of actions. While inspired by these approaches, we focus on generating videos which is a very different challenge.

### 3. Action Graphs

Our goal in this work is to build a model for synthesizing videos that contain a specified set of timed actions. A key component in this effort is developing a semantic representation to describe the actions performed by different objects in the scene. Towards this end, we propose to use a graph-structure we call *Action Graph* (AG). In an AG, nodes correspond to objects, and edges correspond to actions that the objects participate in. Objects and actions are labeled with semantic categories, while actions are also annotated with their start and end times. Formally, an AG is a tuple  $(\mathcal{C}, \mathcal{A}, \mathcal{O}, \mathcal{E})$  defined as follows:

**A vocabulary of object categories  $\mathcal{C}$ :** Object categories can include attributes, e.g., "Blue Cylinder" or "Large Box".

**A vocabulary of actions  $\mathcal{A}$ :** Some actions, such as "Slide" may also have attributes (e.g., the destination coordinates).

**Object nodes  $\mathcal{O}$ :** A set  $\mathcal{O} \in \mathcal{C}^n$  of  $n$  objects.

**Action edges  $\mathcal{E}$ :** Actions are represented as labeled directed edges between object nodes. Each edge is annotated with an action category and with the time period during which the action is performed. Formally, each edge is of the form  $(i, a, j, t_s, t_e)$ , where  $i, j \in \{1, \dots, n\}$  are object instances,  $a \in \mathcal{A}$  is an action and  $t_s, t_e \in \mathbb{N}$  are action start and end time. Thus, this edge implies that object  $i$  (which has category  $o_i$ ) performs an action  $a$  over object  $j$ , and that this action takes place between times  $t_s$  and  $t_e$ . We note that an AG edge can directly model actions over a single object and a pair of objects. For example, "Swap the positions of objects  $i$  and  $j$  between time 0 and 9" is an action over two objects corresponding to the edge  $(i, swap, j, 0, 9)$ . Some actions, such as "Rotate", involve only one object and will therefore be specified as self-loops.

### 4. Action Graph to Video via Clocked Edges

We now turn to the key challenge of this paper: transforming an AG into a video. Naturally, this transformation will be learned from data. The generation problem is defined as follows: we wish to build a generator that takes as input an AG and outputs a video. We will also allow conditioning on the first video frame and layout, so we can preserve the visual attributes of the given objects.<sup>2</sup>

<sup>2</sup>Using the first frame and layout can be avoided by using a SG2Image model (Johnson et al., 2018; Ashual & Wolf, 2019; Herzig et al., 2020) for generating the first frame.

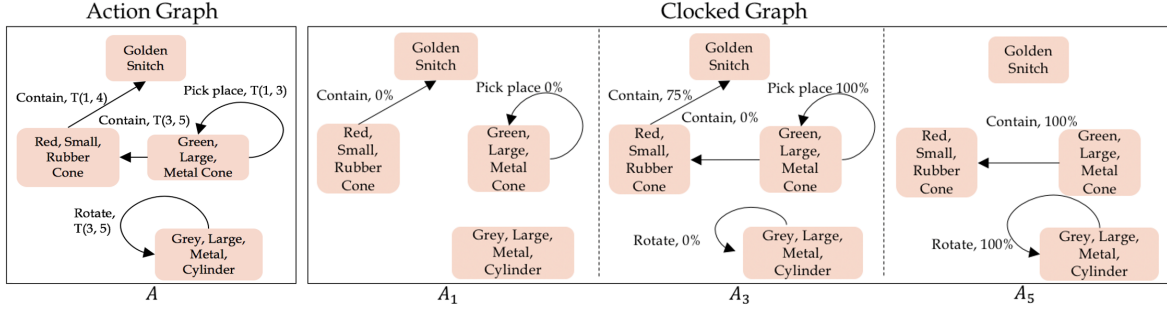


Figure 2. Example of a partial Action Graph execution schedule in different time-steps.

There are multiple unique challenges in generating a video from an AG. First, each action in the graph unfolds over time, so the model needs to “keep track” of the progress of actions. Second, AGs may contain multiple concurrent actions and the generation process needs to combine them in a realistic way. Third, one has to design a training loss that captures the spatio-temporal video structure to ensure that the semantics of the AG is accurately captured.

**Clocked Edges.** As discussed above, we need a mechanism for monitoring the progress of action execution during the video generation. A natural approach is to add a “clock” for each action, for keeping track of action progress. See Figure 2 for an illustration. Formally, we keep a clocked version of the graph  $A$  where each edge is augmented with a temporal state. Let  $e = (i, a, j, t_s, t_e) \in E$  be an edge in the AG  $A$ . We define the progress of  $e$  at time  $t$  to be  $r_t = \frac{t-t_s}{t_e-t_s}$ , and clip  $r_t$  to  $[0, 1]$ . Thus, if  $r_t = 0$  the action has not started yet, if  $r_t \in (0, 1)$  it is currently being executed, and if  $r_t = 1$  it has completed. We then create an augmented version of the edge  $e$  at time  $t$ , denoted  $e_t = (i, a, j, t_s, t_e, r_t)$ . We define  $A_t = \{e_t | e \in A\}$  to be the AG at time  $t$ . To summarize the above, we take the original graph  $A$  and turn it into a sequence of AGs  $A_0, \dots, A_T$ , where  $T$  is the last time-step. Each action edge in the graph now has a unique clock for its execution. This facilitates a timely execution and coordination between actions.

#### 4.1. The AG2Vid Model

Next, we describe our proposed AG-to-video model (AG2Vid). Figure 3 provides a high-level illustration of the model architecture. The idea of the generation process is that the AG is first used to produce intermediate layouts, and then these layouts are used to produce frame pixels. We let  $\ell_t = (x_t, y_t, w_t, h_t, z_t)$  denote the set of predicted layouts for the  $n$  objects in the video at time  $t$ . The values  $x_t, y_t, w_t, h_t \in [0, 1]^n$  are the bounding box coordinates for all objects, and  $z_t$  are the object descriptors, aimed to represent objects given the actions they participate in, as well as other actions in the scene that might affect them. Let  $v_t$  denote the generated frame at time  $t$ , and

$p(v_2, \dots, v_T, \ell_2, \dots, \ell_T | A, v_1, \ell_1)$  denote the generating distribution of the frames and layouts given the AG and the first frame  $v_1$  and scene layout  $\ell_1$ .

We assume that the generation of the frame and layout directly depends only on the current AG and the recently generated frames and layouts. Specifically, this corresponds to the following form for  $p$ :

$$p(v_2, \dots, v_T, \ell_2, \dots, \ell_T | A, v_1, \ell_1) = \prod_{t=2}^T p(\ell_t | A_t, \ell_{t-1}) p(v_t | v_{t-1}, \ell_t, \ell_{t-1}). \quad (1)$$

Conditioning over  $A_t$  provides both short and long term information which makes it possible to predict future layouts. For example, it can be inferred from  $A_3$  in Figure 2 that the “Grey Large Metal Cylinder” is scheduled to “Rotate” and thus is expected to stay in the same spatial location for the rest of the video.

We refer to the distribution  $p(\ell_t | \cdot)$  as *The Layout Generating Function (LGF)* and to  $p(v_t | \cdot)$  as *The Frame Generating Function (FGF)*. Next, we describe how we model these distributions as functions.

**The Layout Generating Function (LGF).** At time  $t$  we want to use the previous layout  $\ell_{t-1}$  and current AG  $A_t$  to predict the current layout  $\ell_t$ . The rationale is that  $A_t$  captures the current state of the actions and can thus “propagate”  $\ell_{t-1}$  to the next layout. This prediction requires integrating information from different objects as well as the progress of the actions given by the edges of  $A_t$ . Thus, a natural architecture for this task is a *Graph Convolutional Network (GCN)* that operates on the graph  $A_t$  whose nodes are “enriched” with the layouts  $\ell_t$ . Formally, we construct a new graph of the same structure as  $A_t$ , with new features on nodes and edges. The graph nodes features are comprised of the previous object locations defined in  $\ell_{t-1}$  and the embeddings of the object categories  $O$ . The features on the edges are the embedding of action  $a$  and the progress of the action,  $r_t$ , taken from the edge  $(i, a, j, t_s, t_e, r_t)$  of  $A_t$ . Then, node and edge features are repeatedly re-estimated for  $K$  steps using a GCN. The resulting activations at timestep

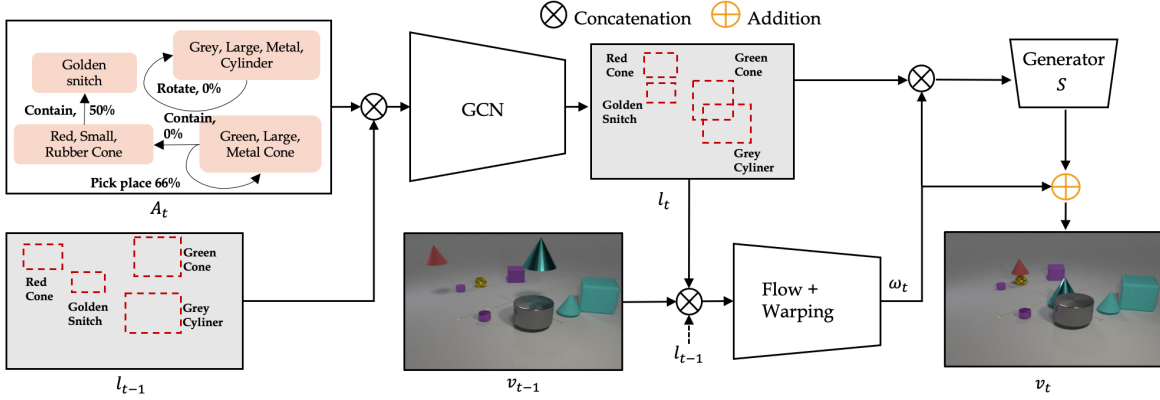


Figure 3. Our AG2Vid Model. The AG  $A_t$  describes the execution stage of each action at time  $t$ . Together with the previous layout  $l_{t-1}$ , it is used to generate the next layout  $l_t$  which has object representations that are enriched with actions information from  $A_t$ . Then,  $l_t, l_{t-1}, v_{t-1}$  are used to generate the next frame.

$t$  are  $z_t \in \mathbb{R}^{O \times D}$  which we use as the new object descriptors. Each such object descriptor is enriched with the action information. An MLP is then applied to it to produce the new box coordinates, which together form the layout  $l_t$ .<sup>3</sup>

**The Frame Generating Function (FGF).** After obtaining the layout  $l_t$  which contains the updated object representations  $z_t$ , we use it along with  $v_{t-1}$  and  $l_{t-1}$  to predict the next frame. The idea is that  $l_t, l_{t-1}$  characterize how objects should move,  $z_t, z_{t-1}$  capture the object features enriched with action information, and  $v_{t-1}$  shows their last appearance. Combining all of them should allow us to generate the next frame accurately. As the first step, we transform the layouts  $l_t, l_{t-1}$  to feature maps  $m_{t-1}, m_t \in \mathbb{R}^{H \times W \times D}$  which have a spatial representation similar to images. To transform a layout  $l$  to a feature map  $m$ , we follow a similar process as in (Johnson et al., 2018). We construct a feature map per object  $\hat{m} \in \mathbb{R}^{O \times H \times W \times D}$ , which is initialized with zeros, and for every object assign its embedding from  $z$  to its location from  $l$ . Finally,  $m$  is obtained by summing over the objects feature maps  $\hat{m}$ . These feature maps provide a coarse object motion. To obtain more fine-grained motion, we estimate how pixels in the image will move from time  $t-1$  to  $t$  using optical flow. We compute  $f_t = F(v_{t-1}, m_{t-1}, m_t)$ , where  $F$  is a flow prediction network similar to (Ilg et al., 2017). The idea is that given the previous frame and two consecutive feature maps, we should be able to predict the flow.  $F$  is trained using an auxiliary loss and does not require additional supervision (see Section 4.2). Given the flow  $f_t$  and previous frame  $v_{t-1}$ , a natural estimate of the next frame is to use a warping function (Zhou et al., 2016)  $w_t = W(f_t, v_{t-1})$ . Finally we refine  $w_t$  via a network  $S(m_t, w_t)$  that provides an additive correction resulting in the final frame prediction:  $v_t = w_t + S(m_t, w_t)$ , where the  $S$  network is the SPADE generator (Park et al., 2019).

<sup>3</sup>For more details refer to Sec.1 in the Supplemental material.

## 4.2. Losses and Training

We use ground truth frames  $v_t^{GT}$  and layouts  $l_t^{GT}$  for training,<sup>4</sup> and rely on the following losses:

**Layout Prediction Loss  $\mathcal{L}_\ell$ .** This loss is defined as  $\mathcal{L}_\ell = \|l_t - l_t^{GT}\|_1$ , an  $L_1$  loss between ground-truth bounding boxes  $l_t^{GT}$  and predicted boxes  $l_t$ .

**Pixel Action Discriminator Loss  $\mathcal{L}_A$ .** We employ a GAN loss that uses a discriminator to distinguish the generated frames  $v_t$  from the GT frames  $v_t^{GT}$  conditioned on  $A_t$  and  $l_t^{GT}$ . The purpose of this discriminator is to ensure that each generated frame matches a real image and respective actions  $A_t$  and layout  $l_t^{GT}$ . Let  $D_A$  be a discriminator with output in  $(0, 1)$ . First, we follow a similar process as described in the previous section, to obtain a feature map  $m_t$ . For this, we construct a new graph based on  $A_t$  and  $l_t^{GT}$ , and apply a GCN to obtain object embeddings  $z_t$ . Then,  $m_t$  is constructed using  $l_t^{GT}$  and  $z_t$ . The input frame and feature map are then channel-wise concatenated and fed into a multi-scale PatchGAN discriminator (Wang et al., 2018b). The loss is then the GAN loss (e.g., see Isola et al., 2017):

$$\mathcal{L}_A = \max_{D_A} \mathbb{E}_{GT} [\log D_A(A_t, v_t^{GT}, l_t^{GT})] + \mathbb{E}_p [\log(1 - D_A(A_t, v_t, l_t^{GT}))], \quad (2)$$

where  $GT/p$  correspond to sampling frames from the real/generated videos, respectively. Optimization is done in the standard way of alternating gradient ascent on  $D_A$  parameters and descent on generator parameters.

**Flow Loss  $\mathcal{L}_f$ .** This loss measures the error between the warps of the previous frame and the ground truth of the next frame  $v_t^{GT}$ :  $\mathcal{L}_f = \frac{1}{T-1} \sum_{t=1}^{T-1} \|w_t - v_t\|_1$ , where  $w_t = W(f_t, v_{t-1})$ , as defined in Section 4.1. The loss was proposed by Zhou et al. (2016); Wang et al. (2018a).

**Perceptual and Feature Matching Loss  $\mathcal{L}_P$ .** We use

<sup>4</sup>With slight abuse of notation, here we ignore the latent object descriptor part of  $l_t$ .

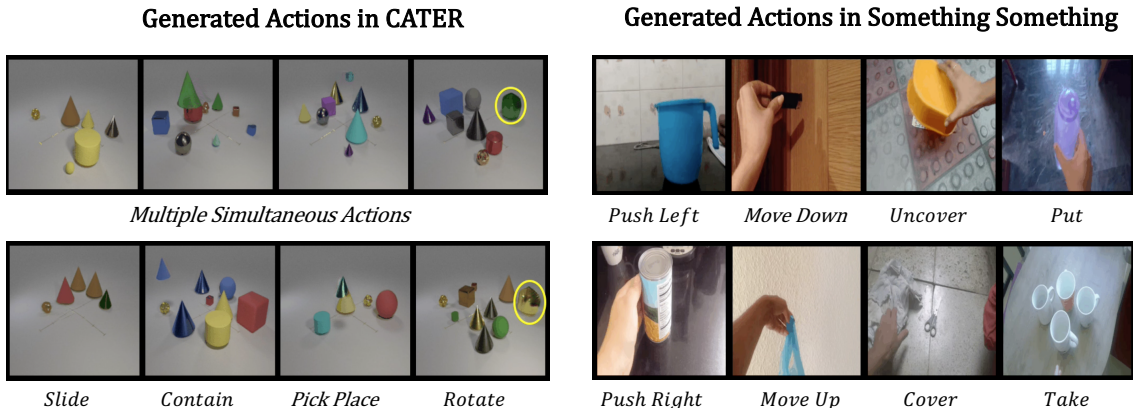


Figure 4. Qualitative examples of generation on CATER and Something-Something V2. AG2Vid generated videos of four and eight standard actions on CATER and Something-Something V2, respectively. For CATER we also used AGs with multiple simultaneous actions, and the generated actions indeed correspond to those (verified manually). For more examples please refer to Figure 9 and Figure 10 in the Supplemental. **Click on the image to play the video clip in a browser.**

these losses as proposed in pix2pixHD (Wang et al., 2018b; Larsen et al., 2016) and other previous works.

The overall optimization problem is to minimize the weighted sum of the losses:

$$\min_{\theta} \max_{D_A} \mathcal{L}_A(D_A) + \lambda_\ell \mathcal{L}_\ell + \lambda_f \mathcal{L}_f + \lambda_P \mathcal{L}_P, \quad (3)$$

where  $\theta$  are the trainable parameters of the generative model.

## 5. Experiments and Results

We evaluate our AG2Vid model on two datasets: *CATER* (Girdhar & Ramanan, 2020) and *Something-Something V2* (Goyal et al., 2017) (henceforth SmthV2). For each dataset, we learn an AG2Vid model with a given set of actions. We then evaluate the visual quality of the generated videos and measure how they semantically comply with the input actions. Finally, we estimate the generalization of the AG2Vid model to novel composed actions.

**Datasets.** We use two datasets: **(1) CATER** is a synthetic video dataset originally created for action recognition and reasoning. Each video contains multiple objects performing actions. The dataset contains bounding-box annotations for all objects, as well as the list of attributes per object (color, shape, material and size), the labels of the actions and their time. Actions include: “Rotate”, “Contain”, “Pick Place” and “Slide”. For “Pick Place” and “Slide” we include the action destination coordinates  $(x, y)$  as an attribute. We use these actions to create Action Graphs for training and evaluation. We employ the standard CATER training partition (3,849 videos) and split the validation into 30% val (495 videos) and use the rest for testing (1,156 videos).

**(2) Something-Something V2** contains real videos of humans performing basic actions. Here we use the eight most frequent actions: “Push Left”, “Push Right”, “Move Down”,

AG2Vid vs. Baseline	Semantic Accuracy		Visual Quality	
	CATER	SmthV2	CATER	SmthV2
CVP (Ye et al., 2019)	<b>85.7</b>	<b>90.6</b>	<b>76.2</b>	<b>93.8</b>
HG (Nawhal et al., 2020a)	—	<b>84.6</b>	—	<b>88.5</b>
V2V (Wang et al., 2018a)	<b>68.8</b>	<b>84.4</b>	<b>68.8</b>	<b>96.9</b>
RNN	<b>56.0</b>	<b>80.6</b>	<b>52.0</b>	<b>77.8</b>
AG2Vid-GTL	48.6	46.2	42.9	50.0

Table 1. **Human evaluation** of action generation with respect to Semantic Accuracy and Visual Quality. Raters have to select the better image between the AG2Vid and a baseline generation method. Each number means that AG2Vid was selected as better for  $X\%$  of the presented pairs. Image resolution is  $256 \times 256$ .

“Move Up”, “Cover”, “Uncover”, “Put” and “Take”. Every videos has a single action and up to three different objects, coming from 122 different object classes, including the hand which is performing the action. We use the bounding box annotations of the objects from Materzynska et al. (2020). We set the first frame where the objects are present as the action’s start time and the last one as the end time.

**Implementation Details.** The GCN model uses  $K = 3$  hidden layers and an embedding layer of 128 units for each object and action. For optimization we use ADAM (Kingma & Ba, 2014) with  $lr = 1e - 4$  and  $(\beta_1, \beta_2) = (0.5, 0.99)$ . Models were trained with a batch size of 2 which was the maximal batch size to fit on a single NVIDIA V100 GPU. For loss weights we use  $\lambda_B = \lambda_F = \lambda_P = 10$  and  $\lambda_A = 1$ . We use videos of 8 FPS and 6 FPS for CATER and SmthV2 and evaluate on videos consisting of 16 frames which correspond to spans of 2.7 and 2 seconds accordingly.

**Performance Metrics.** The AG2Vid outputs are quantitatively evaluated as follows. **a) Visual Quality:** It is common in video generation to evaluate the visual quality of videos regardless of the semantic content. To evaluate visual quality, we use the Learned Perceptual Image Patch Similar-

## Compositional Video Synthesis with Action Graphs

Loss	Inception $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$	
	SmthV2	SmthV2	CATER	SmthV2
Flow	1.59 $\pm$ 0.02	107.26 $\pm$ 1.46	0.14 $\pm$ 0.01	0.70 $\pm$ 0.06
+Percept.	2.21 $\pm$ 0.07	71.70 $\pm$ 1.46	0.08 $\pm$ 0.03	0.29 $\pm$ 0.07
+Act Disc.	<b>3.02 <math>\pm</math> 0.11</b>	<b>66.70 <math>\pm</math> 1.29</b>	<b>0.07 <math>\pm</math> 0.02</b>	<b>0.25 <math>\pm</math> 0.08</b>

Table 2. Ablation experiment for losses of the frame generation. Losses are added one by one.

ity (LPIPS) by Zhang et al. (2018) (lower is better) over predicted and GT videos. For the SmthV2 dataset, since videos contain single actions we can report the Inception Score (IS) (Salimans et al., 2016) (higher is better) and Fréchet Inception Distance (FID) (Heusel et al., 2017) (lower is better) using a TSM (Lin et al., 2019) model, pretrained on SmthV2. For CATER, we avoid reporting FID and IS as these scores rely on models pre-trained for single-action recognition, and CATER videos contain multiple simultaneous actions. Finally, to assess the visual quality, we present human annotators with our model and a baseline generated results pairs and ask them to pick the video with the better quality. **b) Semantic Accuracy:** The key goal of AG2Vid is to generate videos that accurately depict the specified actions. To evaluate this, we ask human annotators to select which of the two given video generation models provides a better depiction of actions shown in the real video. The protocol is similar to the visual quality evaluation above. We also evaluate action timing as discussed below.

**Compared Methods.** Generating videos based on AG input is a new task. There are no off-the-shelf models that can be used for direct comparison with our approach, since no existing models take AG as input and output a video. To provide fair evaluation, we compare with two types of baselines. (1) Existing baseline models that share some functionality with AG2Vid. (2) Variants of the AG2Vid model that shed light on its design choices. Each baseline serves to evaluate specific aspects of the model, as described next. **Baselines:** (1) **HOI-GAN (HG) (Nawhal et al., 2020b)** generates videos given a *single* action-object pair, an initial frame and a layout. It can be viewed as operating on a two-node AG without timing information. We compare HG to AG2Vid on the SmthV2 dataset because it contains exactly such action graphs. HG is not applicable to CATER since the videos there contain multiple action-object pairs. (2) **CVP (Ye et al., 2019)** uses the first image and layout as input for future frame prediction without access to action information. CVP allows us to assess the visual quality of the AG2Vid videos. However, it is not expected that CVP captures the semantics of the action-graph, unless the first frame and action are highly correlated (e.g., a hand at the top-left corner always moves down). (3) **V2V (Wang et al., 2018a)**: This baseline uses a state-of-the-art Vid2Vid model to generate videos from *ground-truth* layouts. Since it uses ground-truth layouts it provides an upper bound on Vid2Vid performance for this task. We note that Vid2Vid cannot

Methods	Inception $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$	
	SmthV2	SmthV2	CATER	SmthV2
64x64				
Real Videos	3.90 $\pm$ 0.120	0.00 $\pm$ 0.00	0.00 $\pm$ 0.00	0.00 $\pm$ 0.00
HG (Nawhal et al., 2020a)	1.66 $\pm$ 0.03	35.18 $\pm$ 3.60	-	0.33 $\pm$ 0.08
<b>AG2Vid (Ours)</b>	<b>2.51 <math>\pm</math> 0.08</b>	<b>26.05 <math>\pm</math> 0.73</b>	<b>0.04 <math>\pm</math> 0.01</b>	<b>0.13 <math>\pm</math> 0.01</b>
256x256				
Real Videos	7.58 $\pm$ 0.20	0.00 $\pm$ 0.00	0.00 $\pm$ 0.00	0.00 $\pm$ 0.00
CVP (Ye et al., 2019)	1.92 $\pm$ 0.03	67.77 $\pm$ 1.43	0.24 $\pm$ 0.04	0.55 $\pm$ 0.08
RNN	1.99 $\pm$ 0.05	74.17 $\pm$ 1.54	0.14 $\pm$ 0.05	0.26 $\pm$ 0.08
V2V (Wang et al., 2018a)	2.22 $\pm$ 0.07	67.51 $\pm$ 1.42	0.11 $\pm$ 0.02	0.29 $\pm$ 0.09
<b>AG2Vid (Ours)</b>	<b>3.02 <math>\pm</math> 0.11</b>	<b>66.70 <math>\pm</math> 1.29</b>	<b>0.07 <math>\pm</math> 0.02</b>	<b>0.25 <math>\pm</math> 0.08</b>
<b>AG2Vid-GTL (Ours)</b>	<b>3.52 <math>\pm</math> 0.14</b>	<b>65.04 <math>\pm</math> 1.25</b>	<b>0.06 <math>\pm</math> 0.02</b>	<b>0.22 <math>\pm</math> 0.09</b>

Table 3. Visual quality metrics of conditional video-generation methods in CATER and SmthV2. All methods use resolution 256  $\times$  256 except for HG, which only supports 64  $\times$  64.

use the AG, and thus it is not provided as input. **AG2Vid variants:** (4) **RNN:** This AG2Vid variant replaces the LGF GCN with an RNN that processes the AGs edges sequentially. The frame generation part is the same as in AG2Vid. The motivation behind this baseline is to compare the design choice of the GCN to a model that processes edges sequentially (RNN). (5) **AG2Vid-GTL:** An AG2Vid model that uses GT layouts at inference time. It allows us to see if using the GT layouts for all frames improves an overall AG2Vid video quality and semantics.

**Layout Generation Ablation Results.** We compare the choice of GCN to the RNN baseline, as an alternative to the implementation of the AG2Vid LGF. To evaluate, we use the mIOU (mean Intersection Over Union) between the predicted and GT bounding boxes. The results confirm the advantage of GCN over RNN with a mIOU of 93.09 vs. 75.71 over CATER and 51.32 vs. 41.28 over SmthV2. We include the full results in Sec. 4.2 in the Supplementary.

**Loss Ablation Results.** Table 2 reports ablations over the losses used in FGF, confirming that each loss, including the pixel actions discriminator loss, contributes to the overall visual quality on CATER and SmthV2.

**Semantic and Visual Quality.** We include human evaluation results in Table 1, which indicate that AG2Vid is more semantically accurate and has better visual quality than the baselines. Comparing to AG2Vid-GTL, it is slightly worse, as expected. We also provide qualitative AG2Vid generation examples in Figure 4, and comparison to baselines in Figure 5. For more results and a per-action analysis of the correctness of the generated actions see the Suppl. Section 4.5. Table 3 evaluates visual quality using several automatic metrics, with similar takeaways as in Table 1.

**Action Timings.** To evaluate to what extent the AG2Vid model can control the timing of actions, we generated pairs of AGs where in the first AG an action is coordinated to appear earlier than in the second AG, keeping everything else constant. We then use the AG2Vid model to create corresponding videos and ask MTurk annotators to choose the video where the action appears first. In 89.45% of the cases, the annotators confirm the intended result. For more information see Section 4.4 in the Supplementary.




Figure 5. Comparison to the baseline methods. The top row is from *CATER*, the bottom row is from *Something-Something V2*. [Click on the image to play the video clip in a browser.](#)

**Composing New Actions.** Finally, we validate the ability of our AG2Vid model to generalize at test time to zero-shot compositions of actions. Here, we manually define four new “composite” actions. We are using learned atomic actions to generate new combinations that did not appear in the training data (either by having the same object perform multiple actions at the same time, or having multiple objects perform coordinated actions), see Figure 6. Specifically, in *CATER*, we created the action “Swap” based on “Pick-Place” and “Slide”, and action “Huddle” based on multiple “Contain” actions. In *SmthV2* we composed the “Push-Left” and “Move-Down” to form the “Left-Down” action, and “Push-Right” with “Move-Up” to form the “Right-Up” action. For each generated video, raters were asked to choose the correct action class from the list of all possible actions. The average class recall for *CATER* and *SmthV2* is 96.65 and 87.5 respectively. See Supplementary for more details.

To further demonstrate that AG2Vid can generalize to AGs with more objects and actions than seen in training, we include additional examples. In Figure 7, the AG contains two simultaneous actions that were never observed together during the training. In Figure 8, the AG contains three consecutive actions, a scenario also not seen during training. These examples demonstrate the ability of our AG2Vid model to generalize to new and more complex scenarios.

## 6. Discussion

In this work, we aim to generate videos of multiple coordinated and timed actions. To represent the actions, we use an Action Graph, and propose the “Action Graph to Video” task. Based on the AG representation, our AG2Vid model can synthesize complex videos of new action compositions and control the timing of actions. The AG2Vid model uses three levels of abstraction: scheduling actions using the “Clocked Edges” mechanism, coarse layout prediction, and fine-grained pixels level prediction. Despite outperform-

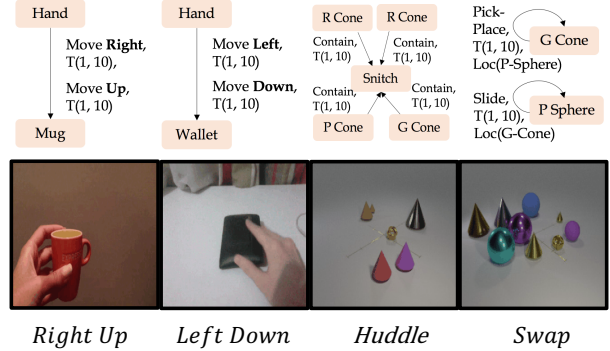


Figure 6. Zero-shot “composite” actions in *Something-Something V2* and *CATER*. For example, the “Swap” action is composed by performing the “Pick-Place” and “Slide” actions simultaneously. [Click on the image to play the video clip in a browser.](#)




Figure 7. AG2Vid generation example of a video with four objects and two simultaneous actions. The hands (a,d) were not part of the original image and were manually added. Here, we illustrate the ability of our model to work on AGs that are much more complex than the typical *Something-Something V2* examples.




Figure 8. AG2Vid generation example of a video with three consecutive actions. Here, we illustrate the ability of our model to work on AGs with multiple actions, while trained on single actions. [Click on the image to play the video clips in a browser.](#)

ing other methods, AG2Vid still fails in certain cases, e.g., when synthesizing occluded or overlapping objects. We believe more fine-grained intermediate representations like segmentation masks can be incorporated to alleviate this difficulty. Furthermore, while it was shown that AG2Vid can synthesize actions that have rigid motion, synthesizing actions with non-rigid motion might require conditioning on random noise, or by explicitly defining action attributes like speed or style. We leave such extensions for future work.



## Acknowledgements

We would like to thank Lior Bracha for her help running MTurk experiments, and to Haggai Maron and Yuval Atzmon for helpful feedback and discussions. This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant ERC HOLI 819080). Trevor Darrell’s group was supported in part by DoD including DARPA’s XAI, LwLL, and/or SemaFor programs, as well as BAIR’s industrial alliance programs. Gal Chechik’s group was supported by the Israel Science Foundation (ISF 737/2018), and by an equipment grant to Gal Chechik and Bar-Ilan University from the Israel Science Foundation (ISF 2332/18). This work was completed in partial fulfillment for the Ph.D degree of Amir Bar.

## References

- Ashual, O. and Wolf, L. Specifying object attributes and relations in interactive scene generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4561–4569, 2019.
- Babaeizadeh, M., Finn, C., Erhan, D., Campbell, R. H., and Levine, S. Stochastic variational video prediction. In *International Conference on Learning Representations*, 2018.
- Bansal, A., Ma, S., Ramanan, D., and Sheikh, Y. Recycle-gan: Unsupervised video retargeting. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 119–135, 2018.
- Battaglia, P., Pascanu, R., Lai, M., Rezende, D. J., et al. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, pp. 4502–4510, 2016.
- Brock, A., Donahue, J., and Simonyan, K. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- Chan, C., Ginosar, S., Zhou, T., and Efros, A. A. Everybody dance now. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5933–5942, 2019.
- Deng, Z., Chen, J., Fu, Y., and Mori, G. Probabilistic neural programmed networks for scene generation. In *Advances in Neural Information Processing Systems*, pp. 4028–4038, 2018.
- Denton, E. and Fergus, R. Stochastic video generation with a learned prior. In *ICML*, pp. 1174–1183, 2018.
- Fried, O., Tewari, A., Zollhöfer, M., Finkelstein, A., Shechtman, E., Goldman, D. B., Genova, K., Jin, Z., Theobalt, C., and Agrawala, M. Text-based editing of talking-head video. *ACM Transactions on Graphics (TOG)*, 38(4): 1–14, 2019.
- Ginosar, S., Bar, A., Kohavi, G., Chan, C., Owens, A., and Malik, J. Learning individual styles of conversational gesture. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3497–3506, 2019.
- Girdhar, R. and Ramanan, D. CATER: A diagnostic dataset for Compositional Actions and Temporal Reasoning. In *ICLR*, 2020.
- Girdhar, R., Carreira, J., Doersch, C., and Zisserman, A. Video action transformer network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 244–253, 2019.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Goyal, R., Kahou, S. E., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haanel, V., Freund, I., Yianilos, P., Mueller-Freitag, M., et al. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, pp. 5, 2017.
- Gupta, T., Schwenk, D., Farhadi, A., Hoiem, D., and Kembhavi, A. Imagine this! scripts to compositions to videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 598–613, 2018.
- Herzig, R., Raboh, M., Chechik, G., Berant, J., and Globerson, A. Mapping images to scene graphs with permutation-invariant structured prediction. In *Advances in Neural Information Processing Systems (NIPS)*, 2018.
- Herzig, R., Levi, E., Xu, H., Gao, H., Brosh, E., Wang, X., Globerson, A., and Darrell, T. Spatio-temporal action graph networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 0–0, 2019.
- Herzig, R., Bar, A., Xu, H., Chechik, G., Darrell, T., and Globerson, A. Learning canonical representations for scene graph to image generation. In *European Conference on Computer Vision*, 2020.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 6626–6637. Curran Associates, Inc., 2017.

- Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., and Brox, T. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul 2017. URL <http://lmb.informatik.uni-freiburg.de/Publications/2017/IMKDB17>.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
- Jain, A., Zamir, A. R., Savarese, S., and Saxena, A. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5308–5317, 2016.
- Ji, J., Krishna, R., Fei-Fei, L., and Niebles, J. C. Action genome: Actions as composition of spatio-temporal scene graphs. *arXiv preprint arXiv:1912.06992*, 2019.
- Johnson, J., Krishna, R., Stark, M., Li, L.-J., Shamma, D., Bernstein, M., and Fei-Fei, L. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3668–3678, 2015.
- Johnson, J., Gupta, A., and Fei-Fei, L. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1219–1228, 2018.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. Analyzing and improving the image quality of stylegan. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Kim, D., Woo, S., Lee, J.-Y., and Kweon, I. S. Deep video inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5792–5801, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. *arXiv preprint arXiv:1802.04687*, 2018.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Krishna, R., Chami, I., Bernstein, M. S., and Fei-Fei, L. Referring relationships. *ECCV*, 2018.
- Kumar, M., Babaeizadeh, M., Erhan, D., Finn, C., Levine, S., Dinh, L., and Kingma, D. Videoflow: A conditional flow-based model for stochastic video generation. In *International Conference on Learning Representations*, 2020.
- Larsen, A. B. L., Sønderby, S. K., Larochelle, H., and Winther, O. Autoencoding beyond pixels using a learned similarity metric. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 1558–1566, 2016.
- Lee, A. X., Zhang, R., Ebert, F., Abbeel, P., Finn, C., and Levine, S. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018.
- Li, Y., Min, M. R., Shen, D., Carlson, D. E., and Carin, L. Video generation from text. In *AAAI*, pp. 5, 2018.
- Lin, J., Gan, C., and Han, S. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- Luc, P., Couprie, C., Lecun, Y., and Verbeek, J. Predicting Future Instance Segmentation by Forecasting Convolutional Features. In *ECCV*, pp. 593–608, 2018.
- MacDonald, M. C., Pearlmutter, N. J., and Seidenberg, M. S. The lexical nature of syntactic ambiguity resolution. *Psychological review*, 101(4):676, 1994.
- Mallya, A., Wang, T.-C., Sapra, K., and Liu, M.-Y. World-consistent video-to-video synthesis. In *European Conference on Computer Vision (ECCV)*, 2020.
- Marwah, T., Mittal, G., and Balasubramanian, V. N. Attentive semantic video generation using captions. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1426–1434, 2017.
- Materzynska, J., Xiao, T., Herzig, R., Xu, H., Wang, X., and Darrell, T. Something-else: Compositional action recognition with spatial-temporal interaction networks. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- Mathieu, M., Couprie, C., and LeCun, Y. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015.
- Nawhal, M., Zhai, M., Lehrmann, A., Sigal, L., and Mori, G. Generating videos of zero-shot compositions of actions and objects. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020a.
- Nawhal, M., Zhai, M., Lehrmann, A., Sigal, L., and Mori, G. Generating videos of zero-shot compositions of actions and objects. In *European Conference on Computer Vision (ECCV)*, 2020b.

- Newell, A. and Deng, J. Pixels to graphs by associative embedding. In *Advances in Neural Information Processing Systems 30 (to appear)*, pp. 1172–1180. Curran Associates, Inc., 2017.
- Pan, Y., Qiu, Z., Yao, T., Li, H., and Mei, T. To create what you tell: Generating videos from captions. In *Proceedings of the 25th ACM international conference on Multimedia*, pp. 1789–1798, 2017.
- Park, T., Liu, M.-Y., Wang, T.-C., and Zhu, J.-Y. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2337–2346, 2019.
- Raboh, M., Herzig, R., Chechik, G., Berant, J., and Globerson, A. Differentiable scene graphs. In *Winter Conf. on App. of Comput. Vision*, 2020.
- Saito, M., Matsumoto, E., and Saito, S. Temporal generative adversarial nets with singular value clipping. In *ICCV*, 2017.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., and Chen, X. Improved techniques for training gans. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 2234–2242. Curran Associates, Inc., 2016.
- Schroeder, B., Tripathi, S., and Tang, H. Triplet-aware scene graph embeddings. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- Schuster, S., Krishna, R., Chang, A., Fei-Fei, L., and Manning, C. D. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *Proceedings of the fourth workshop on vision and language*, pp. 70–80, 2015.
- Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., and Sebe, N. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2377–2386, 2019.
- Sigurdsson, G. A., Varol, G., Wang, X., Farhadi, A., Laptev, I., and Gupta, A. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision*, pp. 510–526. Springer, 2016.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Song, Y., Zhu, J., Li, D., Wang, X., and Qi, H. Talking face generation by conditional recurrent adversarial network. *arXiv preprint arXiv:1804.04786*, 2018.
- Sun, C., Shrivastava, A., Vondrick, C., Murphy, K., Sukthankar, R., and Schmid, C. Actor-centric relation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 318–334, 2018.
- Tulyakov, S., Liu, M.-Y., Yang, X., and Kautz, J. MoCoGAN: Decomposing motion and content for video generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1526–1535, 2018.
- Villegas, R., Pathak, A., Kannan, H., Erhan, D., Le, Q. V., and Lee, H. High fidelity video prediction with large stochastic recurrent neural networks. In *Advances in Neural Information Processing Systems 32*, pp. 81–91. Curran Associates, Inc., 2019.
- Vondrick, C., Pirsivash, H., and Torralba, A. Generating videos with scene dynamics. In *Advances in neural information processing systems (NIPS)*, 2016.
- Walker, J., Doersch, C., Gupta, A., and Hebert, M. An uncertain future: Forecasting from static images using variational autoencoders. In *European Conference on Computer Vision*, pp. 835–851. Springer, 2016.
- Wang, T.-C., Liu, M.-Y., Zhu, J.-Y., Liu, G., Tao, A., Kautz, J., and Catanzaro, B. Video-to-video synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018a.
- Wang, T.-C., Liu, M.-Y., Zhu, J.-Y., Tao, A., Kautz, J., and Catanzaro, B. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8798–8807, 2018b.
- Wang, T.-C., Liu, M.-Y., Tao, A., Liu, G., Kautz, J., and Catanzaro, B. Few-shot video-to-video synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Wang, X. and Gupta, A. Videos as space-time region graphs. In *ECCV*, 2018.
- Watters, N., Zoran, D., Weber, T., Battaglia, P., Pascanu, R., and Tacchetti, A. Visual interaction networks: Learning a physics simulator from video. In *Advances in neural information processing systems*, pp. 4539–4547, 2017.
- Wiebe, J., Wilson, T., Bruce, R., Bell, M., and Martin, M. Learning subjective language. *Computational linguistics*, 30(3):277–308, 2004.
- Xu, D., Zhu, Y., Choy, C. B., and Fei-Fei, L. Scene Graph Generation by Iterative Message Passing. In *Proc. Conf. Comput. Vision Pattern Recognition*, pp. 3097–3106, 2017.

- Xu, N., Liu, A.-A., Liu, J., Nie, W., and Su, Y. Scene graph captioner: Image captioning based on structural visual representation. *Journal of Visual Communication and Image Representation*, pp. 477–485, 2019.
- Yan, S., Xiong, Y., and Lin, D. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- Ye, Y., Singh, M., Gupta, A., and Tulsiani, S. Compositional video prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 10353–10362, 2019.
- Yi, K., Gan, C., Li, Y., Kohli, P., Wu, J., Torralba, A., and Tenenbaum, J. B. Clevrer: Collision events for video representation and reasoning. *arXiv preprint arXiv:1910.01442*, 2019.
- Zellers, R., Yatskar, M., Thomson, S., and Choi, Y. Neural motifs: Scene graph parsing with global context. In *Conference on Computer Vision and Pattern Recognition*, 2018.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- Zhou, T., Tulsiani, S., Sun, W., Malik, J., and Efros, A. A. View synthesis by appearance flow. In *ECCV*, 2016.