

FROM HERE TO THERE: VIDEO INBETWEENING USING DIRECT 3D CONVOLUTIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

We consider the problem of generating plausible and diverse video sequences, when we are only given a start and an end frame. This task is also known as *inbetweening*, and it belongs to the broader area of stochastic video generation, which is generally approached by means of recurrent neural networks (RNN). In this paper, we propose instead a fully convolutional model to generate video sequences directly in the pixel domain. We first obtain a latent video representation using a stochastic fusion mechanism that learns how to incorporate information from the start and end frames. Our model learns to produce such latent representation by progressively increasing the temporal resolution, and then decode in the spatiotemporal domain using 3D convolutions. The model is trained end-to-end by minimizing an adversarial loss. Experiments on several widely-used benchmark datasets show that it is able to generate meaningful and diverse in-between video sequences, according to both quantitative and qualitative evaluations.

1 INTRODUCTION

Imagine if we could teach an intelligent system to automatically turn comic books into animations. Being able to do so would undoubtedly revolutionize the animation industry. Although such an immensely labor-saving capability is still beyond the current state-of-the-art, advances in computer vision and machine learning are making it an increasingly more tangible goal. Situated at the heart of this challenge is video *inbetweening*, that is, the process of creating intermediate frames between two given key frames.

Recent development in artificial neural network architectures (Simonyan & Zisserman, 2015; Szegedy et al., 2015; He et al., 2016) and the emergence of generative adversarial networks (GAN) (Goodfellow et al., 2014) have led to rapid advancement in image and video synthesis (Aigner & Körner, 2018; Tulyakov et al., 2017). At the same time, the problem of inbetweening has received much less attention. The majority of the existing works focus on two different tasks: i) unconditional video generation, where the model learns the input data distribution during training and generates new plausible videos without receiving further input (Srivastava et al., 2015; Finn et al., 2016; Lotter et al., 2016); and ii) video prediction, where the model is given a certain number of past frames and it learns to predict how the video evolves thereafter (Vondrick et al., 2016; Saito et al., 2017; Tulyakov et al., 2017; Denton & Fergus, 2018).

In most cases, the generative process is modeled as a recurrent neural network (RNN) using either long-short term memory (LSTM) cells (Hochreiter & Schmidhuber, 1997) or gated recurrent units (GRU) (Cho et al., 2014). Indeed, it is generally assumed that some form of a recurrent model is necessary to capture long-term dependencies, when the goal is to generate videos over a length that cannot be handled by pure frame-interpolation methods based on optical flow. In this paper, we show that it is in fact possible to address the problem of video inbetweening using a stateless, fully convolutional model. A major advantage of this approach is its simplicity. The absence of recurrent components implies shorter gradient paths, hence allowing for deeper networks and more stable training. The model is also more easily parallelizable, due to the lack of sequential states. Moreover, in a convolutional model, it is straightforward to enforce temporal consistency with the start and end frames given as inputs. Motivated by these observations, we make the following contributions in this paper:

- We propose a **fully convolutional model** to address the task of video inbetweening. The proposed model consists of three main components: i) a 2D-convolutional image encoder, which maps the input key frames to a latent space; ii) a 3D-convolutional latent representation generator, which learns how to incorporate the information contained in the input frames with progressively increasing temporal resolution; and iii) a video generator, which uses transposed 3D-convolutions to decode the latent representation into video frames.
- Our key finding is that separating the **generation of the latent representation** from video decoding is of crucial importance to successfully address video inbetweening. Indeed, attempting to generate the final video directly from the encoded representations of the start and end frames tends to perform poorly, as further demonstrated in Section 4. To this end, we carefully design the latent representation generator to stochastically fuse the key frame representations and progressively increase the temporal resolution of the generated video.
- We carried out extensive experiments on several widely used benchmark datasets, and demonstrate that the model is able to produce **realistic** video sequences, considering key frames that are well over a half second apart from each other. In addition, we show that it is possible to generate **diverse** sequences given the same start and end frames, by simply varying the input noise vector driving the generative process.

The rest of the paper is organized as follows: We review the outstanding literature related to our work in Section 2. Section 3 describes our proposed model in details. Experimental results, both quantitative and qualitative, are presented in Section 4, followed by our conclusions in Section 5.

2 RELATED WORK

Recent advances based on deep networks have led to tremendous progress in three areas related to the current work: i) video prediction, ii) video generation and iii) video interpolation.

Video prediction: Video prediction addresses the problem of producing future frames given one (or more) past frames of a video sequence. The methods that belong to this group are deterministic, in the sense that always produce the same output for the same input and they are trained to minimize the L2 loss between the ground truth and the predicted future frames.

Most of the early works in this area adopted recurrent neural networks to model the temporal dynamics of video sequences. In Srivastava et al. (2015) a LSTM encoder-decoder framework is used to learn video representations of image patches. The work in Finn et al. (2016) extends the prediction to video frames rather than patches, training a convolutional LSTM. The underlying idea is to compute the next frame by first predicting the motions of either individual pixels or image segments and then merge these predictions via masking. A multi-layer LSTM is also used in Lotter et al. (2016), progressively refining the prediction error. Some methods do not use recurrent networks to address the problem of video prediction. For example, a 3D convolutional neural network is adopted in Mathieu et al. (2016). An adversarial loss is used in addition to the L2 loss to ensure that the predicted frames look realistic. More recently, Aigner & Körner (2018) proposed a similar approach, though in this case layers are added progressively to increase the image resolution during training (Karras et al., 2017).

All the aforementioned methods aim at predicting the future frames in the pixel domain directly. An alternative approach is to first estimate local and global transformations (e.g., affine warping and local filters), and then apply them to each frame to predict the next, by locally warping the image content accordingly (De Brabandere et al., 2016; Chen et al., 2017a; van Amersfoort et al., 2017).

Video generation: Video generation differs from video prediction in that it aims at modelling future frames in a probabilistic manner, so as to generate diverse and plausible video sequences. To this end, methods based on generative adversarial networks (GAN) and variational autoencoder networks (VAN) are being currently explored in the literature.

In Vondrick et al. (2016) a GAN architecture is proposed, which consists of two generators (to produce, respectively, foreground and static background pixels), and a discriminator to distinguish between real and generated video sequences. While Vondrick et al. (2016) generates the whole output video sequence from a single latent vector, in Saito et al. (2017) a temporal generator is first used to produce a sequence of latent vectors that captures the temporal dynamics. Subsequently an

image generator produces the output images from the latent vectors. Both the generators and the discriminator are based on CNNs. The model is also able to generate video sequences conditionally on an input label, as well as interpolating between frames by first linearly interpolating the temporal latent vectors.

To address mode collapse in GANs, Denton & Fergus (2018) proposes to use a variational approach. Each frame is recursively generated combining the previous frame encoding with a latent vector. This is fed to a LSTM, whose output goes through a decoder. Similarly to this, Babaeizadeh et al. (2018) samples a latent vector, which is then used as conditioning for the deterministic frame prediction network in Finn et al. (2016). A variational approach is used to learn how to sample the latent vector, conditional on the past frames. Other methods do not attempt to predict the pixels of the future frame directly. Conversely, a variational autoencoder is trained to generate plausible differences between consecutive frames (Xue et al., 2016), or motion trajectories (Walker et al., 2016). Recently, Lee et al. (2018) proposed to use a loss function that combines a variational loss (to produce diverse videos) (Denton & Fergus, 2018) with an adversarial loss (to generate realistic frames) (Saito et al., 2017).

Video sequences can be modelled as two distinct components: content and motion. In Tulyakov et al. (2017) the latent vector from which the video is generated is divided in two parts: content and motion. This leads to improved quality of the generated sequences when compared with previous approaches (Vondrick et al., 2016; Saito et al., 2017). A similar idea is explored in Villegas et al. (2017a), where two encoders, one for motion and one for content, are used to produce hidden representations that are then decoded to a video sequence. Also Sun et al. (2018) explicitly separates motion and content in two streams, which are generated by means of a variational network and then fused to produce the predicted sequence. An adversarial loss is then used to improve the realism of the generated videos.

All of the aforementioned methods are able to predict or generate just a few video frames into the future. Long-term video prediction has been originally addressed in Oh et al. (2015) with the goal of predicting up to 100 future frames of an Atari game. The current frame is encoded using a CNN or LSTM, transformed conditionally on the player action, and decoded into the next frame. More recently, Villegas et al. (2017b) addressed a similar problem, but for the case of real-world video sequences. The key idea is to first estimate high-level structures from past frames (e.g., human poses). Then, a LSTM is used to predict a sequence of future structures, which are decoded to future frames. One shortcoming of Villegas et al. (2017b) is that it requires ground truth landmarks as supervision. This is addressed in Wichers et al. (2018), which proposes a fully unsupervised method that learns to predict a high-level encoding into the future. Then, a decoder with access to the first frame generates the future frames from the predicted high-level encoding.

Video interpolation: Video interpolation is used to increase the temporal resolution of the input video sequence. This is addressed with different approaches: optical flow based interpolation (Ilg et al., 2017; Liu et al., 2017), phase-based interpolation (Meyer et al., 2018), and pixels motion transformation (Niklaus et al., 2017; Jiang et al., 2018). These methods typically target temporal super-resolution and the frame rate of the input sequence is often already sufficiently high. Interpolating frames becomes more difficult when the temporal distance between consecutive frames increases. Long-term video interpolation received far less attention in the past literature. Deterministic approaches have been explored using either block-based motion estimation/compensation (Ascenso et al., 2005), or convolutional LSTM models (Kim et al., 2018). Our work is closer to those using generative approaches. In Chen et al. (2017b) two convolutional encoders are used to generate hidden representations of both the first and last frame, which are then fed to a decoder to reconstruct all frames in between. A variational approach is presented in Xu et al. (2018). A multi-layer convolutional LSTM is used to interpolate frames given a set of extended reference frames, with the goal of increasing the temporal resolution from 2 fps to 16 fps. In our experiments, we compare our method with those in Niklaus et al. (2017), Jiang et al. (2018), and Xu et al. (2018).

3 MODEL

The proposed model receives three inputs: a start frame x_s , an end frame x_e , and a Gaussian noise vector $\mathbf{u} \in \mathbb{R}^D$. The output of the model is a video $(x_s, \hat{x}_1, \dots, \hat{x}_{T-2}, x_e)$, where different se-

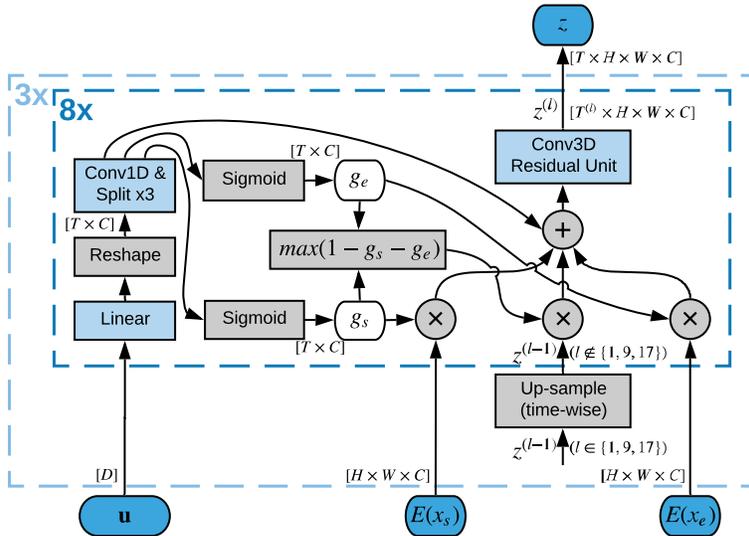


Figure 1: Layout of the model used to generate the latent video representation z . The inputs are the encoded representations of the start and end frames $E(x_s)$ and $E(x_e)$, together with a noise vector \mathbf{u} .

quences of plausible in-between frames $(\hat{x}_i)_{i=1}^{N-2}$ are generated by feeding different instantiations of the noise vector \mathbf{u} . In the rest of this paper, we set $T = 16$ and $D = 128$.

The model consists of three components: an image encoder, a latent representation generator and a video generator. In addition, a video discriminator and an image discriminator are added so that the whole model can be trained using adversarial learning (Goodfellow et al., 2014) to produce realistic video sequences.

3.1 IMAGE ENCODER

The image encoder $E(x)$ receives as input a video frame of size $H_0 \times W_0$ and produces a feature map of shape $H \times W \times C$, where C is the number of channels. The encoder architecture consists of six layers, alternating between 4×4 convolutions with stride-2 down-sampling and regular 3×3 convolutions, followed by a final layer to condense the feature map to the target depth C . This results in spatial dimensions $H = H_0/8$ and $W = W_0/8$. We set $C = 64$ in all our experiments.

3.2 LATENT REPRESENTATION GENERATOR

The latent representation generator $G_Z(\cdot)$ receives as input $E(x_s)$, $E(x_e)$ and \mathbf{u} , and produces an output tensor of shape $T \times H \times W \times C$. Its main function is to gradually fill in the video content between the start and end frames, working directly in the latent space defined by the image encoder.

The model architecture is composed of a series of L residual blocks (He et al., 2016), each consisting of 3D convolutions and stochastic fusion with the encoded representations of x_s and x_e . This way, each block progressively learns a transformation that improves the video content generated by the previous block. The generic l -th block is represented by the inner rectangle in Figure 1. Note that the lengths of the intermediate representations can differ from the final video length T , due to the use of a coarse-to-fine scheme in the time dimension. To simplify the notation, we defer its description to the end of this section and omit the implied temporal up-sampling from the equations.

Let $T^{(l)}$ denote the representation length within block l . First, we produce a layer-specific noise tensor of shape $T^{(l)} \times C$ by applying a linear transformation to the input noise vector \mathbf{u} :

$$\mathbf{u}^{(l)} = A^{(l)}\mathbf{u} + \mathbf{b}^{(l)}, \quad (1)$$

where $A^{(l)} \in \mathbb{R}^{T^{(l)}C \times D}$ and $\mathbf{b}^{(l)} \in \mathbb{R}^{T^{(l)}C}$, and reshaping the result into a $T^{(l)} \times C$ tensor $u^{(l)}$. This is used to drive two stochastic “gating” functions for the start and end frames, respectively:

$$g_s^{(l)} = \sigma(u^{(l)} * k_s^{(l)} + b_s^{(l)}), \quad (2)$$

$$g_e^{(l)} = \sigma(u^{(l)} * k_e^{(l)} + b_e^{(l)}), \quad (3)$$

where $*$ denotes convolution along the time dimension, k_s, k_e are kernels of width 3 and depth C , and $\sigma(\cdot)$ is the sigmoid activation function. The gating functions are used to progressively fuse the encoded representations of the start and end frames with the intermediate output of the previous layer $z^{(l-1)}$, as described by the following equation:

$$z_{\text{in}}^{(l)} = g_s^{(l)} \cdot E(x_s) + g_e^{(l)} \cdot E(x_e) + \max(0, 1 - g_s^{(l)} - g_e^{(l)}) \cdot z^{(l-1)} + n^{(l)}, \quad (4)$$

where $n^{(l)} \triangleq u^{(l)} * k_n^{(l)} + b_n^{(l)}$ denotes an additional learned stochastic component added to stimulate diversity in the generative process. Note that $z_{\text{in}}^{(l)}$ has shape $T^{(l)} \times H \times W \times C$. Therefore, to compute the component-wise multiplication \cdot , $E(x_s)$ and $E(x_e)$ (each of shape $S \times S \times C$) are broadcast (i.e., replicated uniformly) $T^{(l)}$ times along the time dimension, while $g_s^{(l)}, g_e^{(l)}$ and $n^{(l)}$ (each of shape $T^{(l)} \times C$) are broadcast $H \times W$ times over the spatial dimensions. The idea of the fusion step is similar to that of StyleGAN (Zhang et al., 2018), albeit with different construction and purposes. Finally, the fused input is convolved spatially and temporally with $3 \times 3 \times 3$ kernels $k_1^{(l)}$ and $k_2^{(l)}$ in a residual unit (He et al., 2016):

$$z^{(l)} = h(z^{(l-1)} + h(z_{\text{in}}^{(l)} * k_1^{(l)} + b_1^{(l)}) * k_2^{(l)} + b_2^{(l)}), \quad (5)$$

where h is the leaky ReLU (Maas et al., 2013) activation function (with parameter $\alpha = 0.2$). Hence Equation 1–5 collectively define the stochastic transformation $G^{(l)}(\cdot)$ from $z^{(l-1)}$ to $z^{(l)}$ given $E(x_s)$ and $E(x_e)$, with A, k, b being its learnable parameters. The generation of the overall latent video representation $z \in \mathbb{R}^{T \times S \times S \times C}$ can be expressed as:

$$\begin{aligned} z \triangleq z^{(L)} &= G_Z(E(x_s), E(x_e), \mathbf{u}) \\ &= G^{(L)} \circ \dots \circ G^{(1)}(z^{(0)}; E(x_s), E(x_e), \mathbf{u}). \end{aligned} \quad (6)$$

Coarse-to-fine generation: For computational efficiency, we adopt a coarse-to-fine scheme in the time dimension, represented by the outer dashed rectangle in Figure 1. More specifically we double the length of $z^{(l)}$ every $L/3$ generator blocks, i.e., $z^{(1)}, \dots, z^{(L/3)}$ have length $T/4 = 4$, $z^{(L/3+1)}, \dots, z^{(2L/3)}$ have $T/2 = 8$, and $z^{(2L/3+1)}, \dots, z^{(L)}$ have the full temporal resolution $T = 16$. We initialize $z^{(0)}$ to $(E(x_s), E(x_e))$ (which becomes $(E(x_s), E(x_s), E(x_e), E(x_e))$ after the first up-sampling) and set $L = 24$, resulting in 8 blocks per granularity level.

3.3 VIDEO GENERATOR

The video generator G_V produces the output video sequence $(x_s, \hat{x}_1, \hat{x}_2, \dots, x_e) = G_V(z)$ from the latent video representation z using spatially transposed 3D convolutions. The generator architecture alternates between $3 \times 3 \times 3$ regular convolutions and transposed $3 \times 4 \times 4$ convolutions with a stride of $(1, 2, 2)$, hence applying only spatial (but not temporal) up-sampling. Note that it actually generates all T frames including the “reconstructed” start frame \hat{x}_0 and end frame \hat{x}_{T-1} , though they are not used and are always replaced by the real x_s and x_e in the output.

3.4 LOSS FUNCTIONS

We train our model end-to-end by minimizing an adversarial loss function. To this end, we train two discriminators: a 3D convolutional video discriminator D_V and a 2D convolutional image discriminator D_I , following the approach of Tulyakov et al. (2017). The video discriminator has a similar architecture to Tulyakov et al. (2017), except that in our case we produce a single output for the entire video rather than for its sub-volumes (“patches”). For the image discriminator, we use a Resnet-based architecture (He et al., 2016) instead of the DCGAN-based architecture (Radford et al., 2016) used in Tulyakov et al. (2017).

Let $X = (x_s, x_1, \dots, x_{T-2}, x_e)$ denote a real video and $\hat{X} = (x_s, \hat{x}_1, \dots, \hat{x}_{T-2}, x_e)$ denote the corresponding generated video conditioned on x_s and x_e . Adopting the non-saturating log-loss, training amounts to optimizing the following adversarial objectives:

$$\min_{D_V} : \mathcal{L}(D_V) = \mathbb{E}_{(X, \hat{X})} \left[-\log D_V(X) - \log(1 - D_V(\hat{X})) \right] \quad (7)$$

$$\min_{D_I} : \mathcal{L}(D_I) = \mathbb{E}_{(X, \hat{X})} \left[\frac{1}{T-2} \sum_{i=1}^{T-2} [-\log D_I(x_i) - \log(1 - D_I(\hat{x}_i))] \right] \quad (8)$$

$$\min_{G=\{E, G_Z, G_V\}} : \mathcal{L}(G) = \mathbb{E}_{(X, \hat{X})} \left[-\log D_V(\hat{X}) - \frac{1}{T-2} \sum_{i=1}^{T-2} \log D_I(\hat{x}_i) \right] \quad (9)$$

During optimization we replace the average over the $T - 2$ intermediate frames with a single uniformly sampled frame to save computation, as is done in Tulyakov et al. (2017). This does not change the convergence properties of stochastic gradient descent, since the two quantities have the same expectation.

We regularize the discriminators by penalizing the derivatives of the pre-sigmoid logits with respect to their input videos and images, as is proposed in Roth et al. (2017) to improve GAN stability and prevent mode collapse. In our case, instead of the adaptive scheme of Roth et al. (2017), we opt for a constant coefficient of 0.1 for the gradient magnitude, which we found to be more reliable in our experiments. We use batch normalization (Ioffe & Szegedy, 2015) on all 2D and 3D convolutional layers in the generator and layer normalization (Ba et al., 2016) in the discriminators. 1D convolutions and fully-connected layers are not normalized. Architectural details of the encoder, decoder, and discriminators are further provided in Appendix A.

4 EXPERIMENTS

We evaluated our approach on three well-known public datasets: BAIR robot pushing (Ebert et al., 2017), KTH Action Database (Schuldt et al., 2004), and UCF101 Action Recognition Data Set (Soomro et al., 2012). All video frames were down-sampled and cropped to 64×64 , and subsequences of 16 frames were used in all the experiments, that is, 14 intermediate frames are generated. The videos in KTH and UCF101 datasets are 25 fps, translating to key frames 0.6 seconds apart. The frame rate of BAIR videos is not provided, though visually it appears to be much lower, hence longer time in between key frames. For all the datasets, we adopted the conventional train/test splits practised in the literature. A validation set held out from the training set was used for model checkpoint selection. More details on the exact splits are provided in Appendix B. We did not use any dataset-specific tuning of hyper-parameters, architectural choices, or training schemes.

4.1 METRICS AND METHODOLOGY

Our main objective is to generate plausible transition sequences with characteristics similar to real videos, rather than predicting the exact content of the original sequence from which the key frames were extracted. Therefore we use the recently proposed Fréchet video distance (FVD) (Unterthiner et al., 2018) as our primary evaluation metrics. The FVD is equivalent to the Fréchet Inception distance (FID) (Heusel et al., 2017) widely used for evaluating image generative models, but revisited in a way that it can be applied to evaluate videos, by adopting a deep neural network architecture that computes video embeddings taking the temporal dimension explicitly into account. The FVD is a more suitable metrics for evaluating video inbetweening than the widely used structural similarity index (SSIM) (Wang et al., 2004). The latter is suitable when evaluating prediction tasks, as it compares each synthesized frame with the original reference at the pixel level. Conversely, FVD compares the distributions of generated and ground-truth videos in an embedding space, thus measuring whether the synthesized video belongs to the distribution of realistic videos. Since the FVD was only recently proposed, we also report the SSIM to be able to compare with the previous literature.

During testing, we ran the model 100 times for each pair of key frames, feeding different instances of the noise vector \mathbf{u} to generate different sequences consistent with the given key frames, and computed the FVD for each of these stochastic generations. This entire procedure was repeated 10

Table 1: We report the mean FVD for both the full model and two baselines, averaged over all 10 training runs with 100 stochastic generations each run, and the corresponding 95% confidence intervals. A lower value of the FVD corresponds to higher quality of the generated videos.

	BAIR	KTH	UCF101
Full model	152 [144, 160]	153 [148, 158]	424 [411, 438]
- w/o fusion	175 [166, 184]	171 [163, 180]	463 [453, 474]
- Naïve	702 [551, 895]	346.1 [328, 361]	1101 [1070, 1130]

Table 2: Diversity measured by the average pairwise cosine distance in the FVD embedding space, over 100 stochastic generations. A higher value corresponds to more diverse videos. The mean of the 10 training runs is reported, together with its 95%-confidence interval.

	BAIR	KTH	UCF101
Full model	0.071 [0.065, 0.076]	0.013 [0.010, 0.016]	0.131 [0.122, 0.139]
- w/o fusion	0.051 [0.043, 0.059]	0.006 [0.004, 0.008]	0.121 [0.112, 0.129]

times for each model variant and dataset to account for the randomness in training. We report the mean over all training runs and stochastic generations as well as the confidence intervals obtained by means of the bootstrap method (Efron & Tibshirani, 1993).

For training we used the ADAM (Kingma & Ba, 2015) optimizer with $\beta_1 = 0.5$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and ran it on batches of 32 samples with a conservative learning rate of 5×10^{-5} for 500,000 steps. A checkpoint was saved every 5000 steps, resulting in 100 checkpoints. Training took around 5 days on a single Nvidia Tesla V100 GPU. The checkpoint for evaluation was selected to be the one with the lowest FVD on the validation set.

4.2 RESULTS

To assess the impact of the stochastic fusion mechanism as well the importance of having a separate latent video representation generator component, we compare the full model with baselines in which the corresponding components are omitted.

- **Baseline without fusion:** The gating functions (Equation 2 and 3) are omitted and Equation 4 reduces to $z_{\text{in}}^{(l)} = z^{(l-1)} + n^{(l)}$.
- **Naïve:** The entire latent video representation generator described in Section 3.2 is omitted. Instead, decoding with transposed 3D convolution is performed directly on the (stacked) start/end frame encoded representations $z^{(0)} = (E(x_1), E(x_N))$ (which has dimensionality $2 \times 8 \times 8$), using a stride of 2 in both spatial and temporal dimensions when up-scaling, to eventually produce 16 64×64 frames. To maintain stochasticity in the generative process, a spatially uniform noise map is generated by sampling a Gaussian vector \mathbf{u} , applying a (learned) linear transform, and adding the result in the latent space before decoding.

The results in Table 1 shows that the dedicated latent video representation generator is indispensable, as the naïve baseline performs rather poorly. Moreover, stochastic fusion improves the quality of video generation. Note that the differences are statistically significant at 95% confidence level across all three datasets.

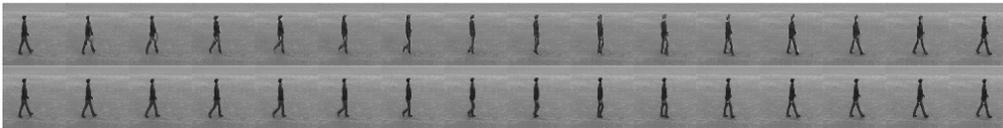
To illustrate the generated videos, Figure 2 shows some exemplary outputs of our full model. The generated sequence is not expected (or even desired) to reproduce the ground truth, but only needs to be similar in style and consistent with the given start and end frames. The samples show that the model does well in this area.

For stochastic generation, good models should produce samples that are not only high-quality but also diverse. Following the approach of Lee et al. (2018), we measure diversity by means of the average pairwise cosine distance (i.e., $1 - \text{cosine similarity}$) in the FVD embedding space among samples generated from the same start/end frames. The results Table 2 shows that incorporating fusion increases sample diversity and the difference is statistically significant.

BAIR:



KTH:



UCF101:



Figure 2: Examples of videos generated with the proposed model. For each of the three datasets, the top row represents the generated video sequences, the bottom row the original video from which the key frames are sampled.



Figure 3: Output diversity illustrated by taking the average of 100 generated videos conditioned on the same start and end frames.

A qualitative illustration of the diversity in the generated videos is further illustrated in Figure 3, where we take the average of 100 generated videos conditioned on the same start and end frames. If the robot arm has a very diverse set of trajectories, we should expect to see it “diffuse” into the background due to averaging. Indeed this is the case, especially near the middle of the sequence.

Finally we computed the average SSIM for our method for each dataset in order to compare our results with those previously reported in the literature, before the FVD metrics was introduced. The results are shown in Table 3 alongside several existing methods that are capable of video inbetweening, ranging from RNN-based video generation (Xu et al., 2018) to optical flow-based interpolation (Niklaus et al., 2017; Jiang et al., 2018)¹. Note that the competing methods generate 7 frames

¹The numbers for these methods are cited directly from Xu et al. (2018).

Table 3: Average SSIM of our model using direct 3D convolution and alternative methods based on RNN, such as SDVI (Xu et al., 2018), or optical flow, such as SepConv (Niklaus et al., 2017) and SuperSloMo (Jiang et al., 2018). Higher is better. Note the difference in setup: our model spans a time base twice as long as the others. The SSIM for each test example is computed on the best sequence out of 100 stochastic generations, as per standard practice (Babaeizadeh et al., 2018; Denton & Fergus, 2018; Lee et al., 2018; Xu et al., 2018). We report the mean and the 95%-confidence interval for our model over 10 training runs.

	BAIR	KTH	UCF101
14 in-between frames			
3D-Conv (ours)	0.836 [0.832, 0.839]	0.733 [0.729, 0.737]	0.686 [0.680, 0.693]
7 in-between frames			
SDVI, full	0.880	0.901	0.598
SDVI, cond. 2 frames	0.852	0.831	—
SepConv	0.877	0.904	0.443
SuperSloMo	—	0.893	0.471

and are conditioned on potentially multiple frames before and after. In contrast our model generates 14 frames, i.e., over a time base twice as long, and it is conditioned on only one frame before and after. Consequently, the SSIM figures are not directly comparable. However it is interesting to see that on UCF101, the most challenging dataset among the three, our model attains higher SSIM than all the other methods despite having to generate much longer sequences. This demonstrates the potential of the direct convolution approach to outperform existing methods, especially on difficult tasks. It is also worth noting from Table 3 that purely optical flow-based interpolation methods achieve essentially the same level of SSIM as the sophisticated RNN-based SDVI on BAIR and KTH, which suggests either that a 7-frame time base is insufficient in length to truly test video inbetweening models or that the SSIM is not an ideal metric for this task.

5 CONCLUSION

We presented a method for video inbetweening using only direct 3D convolutions. Despite having no recurrent components, our model produces good performance on most widely-used benchmark datasets. The key to success for this approach is a dedicated component that learns a latent video representation, decoupled from the final video decoding phase. A stochastic gating mechanism is used to progressively fuse the information of the given key frames. The rather surprising fact that video inbetweening can be achieved over such a long time base without sophisticated recurrent models may provide a useful alternative perspective for future research on video generation.

REFERENCES

- Sandra Aigner and Marco Körner. FutureGAN: Anticipating the Future Frames of Video Sequences using Spatio-Temporal 3d Convolutions in Progressively Growing GANs. Technical report, ArXiv, 2018. URL <http://arxiv.org/abs/1810.01325>.
- João Ascenso, Catarina Brites, and Fernando Pereira. Improving frame interpolation with spatial motion smoothing for pixel domain distributed video coding. In *EURASIP Conference on Speech and Image Processing, Multimedia Communications and Services*, Smolenice, Slovak Republic, 2005. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.329.6290&rep=rep1&type=pdf>.
- Lei Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. Technical report, arXiv, 2016. URL <http://arxiv.org/abs/1607.06450>.
- Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H. Campbell, and Sergey Levine. Stochastic Variational Video Prediction. In *International Conference on Learning Representations (ICLR)*, 2018. URL <http://arxiv.org/abs/1710.11252>.
- Baoyang Chen, Wenmin Wang, Jinzhuo Wang, and Xiongtao Chen. Video Imagination from a Single Image with Transformation Generation. In *Proceedings of the on Thematic Workshops of ACM Multimedia*, Mountain View, California, USA, 2017a. doi: 10.1145/3126686.3126737. URL <http://arxiv.org/abs/1706.04124>.
- Xiongtao Chen, Wenmin Wang, Jinzhuo Wang, Weimian Li, and Baoyang Chen. Long-Term Video Interpolation with Bidirectional Predictive Network. In *IEEE Visual Communications and Image Processing (VCIP)*, 2017b. URL <https://arxiv.org/pdf/1706.03947.pdf>.
- Kyunghyun Cho, Bart van Merriënboer, Çalar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- Bert De Brabandere, Xu Jia, Tinne Tuytelaars, and Luc Van Gool. Dynamic Filter Networks. In *Neural Information Processing Systems (NIPS)*, 2016. URL <http://arxiv.org/abs/1605.09673>.
- Emily Denton and Rob Fergus. Stochastic Video Generation with a Learned Prior. In *International Conference on Machine Learning (ICML)*, 2018. URL <http://arxiv.org/abs/1802.07687>.

- Frederik Ebert, Chelsea Finn, Alex Lee, and Sergey Levine. Self-supervised visual planning with temporal skip connections. In *Conference on Robot Learning (CoRL)*, 2017.
- Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Number 57 in Monographs on Statistics and Applied Probability. Chapman & Hall/CRC, Boca Raton, Florida, USA, 1993.
- Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised Learning for Physical Interaction through Video Prediction. In *Neural Information Processing Systems (NIPS)*, 2016. URL <http://arxiv.org/abs/1605.07157>.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Neural Information Processing Systems (NIPS)*, pp. 2672–2680, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- 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 *Neural Information Processing Systems (NIPS)*, 2017.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <http://dx.doi.org/10.1162/neco.1997.9.8.1735>.
- Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. URL <http://arxiv.org/abs/1612.01925>.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, 2015.
- Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super SloMo: High Quality Estimation of Multiple Intermediate Frames for Video Interpolation. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. URL <http://arxiv.org/abs/1712.00080>.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations (ICLR)*, 2017. URL <http://arxiv.org/abs/1710.10196>.
- Nayoung Kim, Jung Kyung Lee, Chae Hwa Yoo, Seunghyun Cho, and Je-won Kang. Video Generation and Synthesis Network for Long-term Video Interpolation. In *APSIPA Annual Summit and Conference*, 2018. ISBN 9789881476852.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- Alex X. Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic Adversarial Video Prediction. Technical report, ArXiv, 2018. URL <http://arxiv.org/abs/1804.01523>.
- Ziwei Liu, Raymond A. Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *International Conference on Computer Vision (ICCV)*, 2017.
- William Lotter, Gabriel Kreiman, and David Cox. Deep Predictive Coding Networks for Video Prediction and Unsupervised Learning. Technical report, ArXiv, 2016. URL <http://arxiv.org/abs/1605.08104>.
- Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.

- Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. In *International Conference on Learning Representations (ICLR)*, 2016. URL <http://arxiv.org/abs/1511.05440>.
- Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, and Christopher Schroers. PhaseNet for Video Frame Interpolation. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. URL <http://arxiv.org/abs/1804.00884>.
- Simon Niklaus, Long Mai, and Feng Liu. Video Frame Interpolation via Adaptive Separable Convolution. In *International Conference on Computer Vision (ICCV)*, 2017. URL <http://arxiv.org/abs/1708.01692>.
- Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard Lewis, and Satinder Singh. Action-Conditional Video Prediction using Deep Networks in Atari Games. In *Neural Information Processing Systems (NIPS)*, 2015. URL <http://arxiv.org/abs/1507.08750>.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations ICLR*, 2016.
- Kevin Roth, Aurélien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. In *Neural Information Processing Systems (NIPS)*, 2017.
- Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal Generative Adversarial Nets with Singular Value Clipping. In *International Conference on Computer Vision*, 2017. URL <http://arxiv.org/abs/1611.06624>.
- Christian Schuldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: A local svm approach. In *International Conference on Pattern Recognition (ICPR)*, 2004.
- Karen. Simonyan and Andrew. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human action classes from videos in the wild. Technical Report 12-01, UCF-CRCV, November 2012.
- Nitish Srivastava, Elman Mansimov, and Ruslan Salakhutdinov. Unsupervised Learning of Video Representations using LSTMs. In *International Conference on Machine Learning (ICML)*, 2015. URL <http://arxiv.org/abs/1502.04681>.
- Ximeng Sun, Huijuan Xu, and Kate Saenko. A Two-Stream Variational Adversarial Network for Video Generation. Technical report, ArXiv, 2018. URL <http://arxiv.org/abs/1812.01037>.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. URL <http://arxiv.org/abs/1409.4842>.
- Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. MoCoGAN: Decomposing Motion and Content for Video Generation. In *Computer Vision and Pattern Recognition Conference (CVPR)*, 2017. URL <http://arxiv.org/abs/1707.04993>.
- Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *ArXiv*, abs/1812.01717, 2018. URL <http://arxiv.org/abs/1812.01717>.
- Joost van Amersfoort, Anitha Kannan, Marc’ Aurelio Ranzato, Arthur Szlam, Du Tran, and Soumith Chintala. Transformation-Based Models of Video Sequences. Technical report, ArXiv, 2017. URL <http://arxiv.org/abs/1701.08435>.

- Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing Motion and Content for Natural Video Sequence Prediction. In *International Conference on Learning Representations (ICLR)*, 2017a. URL <http://arxiv.org/abs/1706.08033>.
- Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, and Honglak Lee. Learning to Generate Long-term Future via Hierarchical Prediction. In *International Conference on Machine Learning (ICML)*, 2017b. URL <http://arxiv.org/abs/1704.05831>.
- Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating Videos with Scene Dynamics. In *Neural Information Processing Systems (NIPS)*, 2016. URL <http://arxiv.org/abs/1609.02612>.
- Jacob Walker, Carl Doersch, Abhinav Gupta, and Martial Hebert. An Uncertain Future: Forecasting from Static Images using Variational Autoencoders. In *European Conference on Computer Vision*, 2016. URL <http://arxiv.org/abs/1606.07873>.
- Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.
- Nevan Wichers, Ruben Villegas, Dumitru Erhan, and Honglak Lee. Hierarchical Long-term Video Prediction without Supervision. In *International Conference on Machine Learning (ICML)*, 2018. URL <http://arxiv.org/abs/1806.04768>.
- Qiangeng Xu, Hanwang Zhang, Weiyue Wang, Peter N. Belhumeur, and Ulrich Neumann. Stochastic Dynamics for Video Infilling. Technical report, ArXiv, 2018. URL <http://arxiv.org/abs/1809.00263>.
- Tianfan Xue, Jiajun Wu, Katherine L. Bouman, and William T. Freeman. Visual Dynamics: Probabilistic Future Frame Synthesis via Cross Convolutional Networks. In *Neural Information Processing Systems (NIPS)*, 2016. URL <http://arxiv.org/abs/1607.02586>.
- Rui Zhang, Sheng Tang, Yu Li, Junbo Guo, Yongdong Zhang, Jintao Li, and Shuicheng Yan. Style separation and synthesis via generative adversarial networks. In *ACM International Conference on Multimedia*, pp. 183–191, 2018. ISBN 978-1-4503-5665-7. doi: 10.1145/3240508.3240524. URL <http://doi.acm.org/10.1145/3240508.3240524>.

A NETWORK ARCHITECTURE

Architectural configurations of the image encoder, the video decoder, and the discriminators. Notation: C = number of channels, K = kernel size, S = stride, P = size of padding (by zero), (H, W) = frame height and width. Inputs and outputs have 3 channels (RGB) if the videos are colored or 1 channel if they are gray-scale. The batch dimension is omitted. Shape broadcasting is assumed wherever necessary. All generator components uses batch normalization followed by Leaky ReLU activation at the end of each layer, while discriminators use layer normalization. Note that regular convolution and transposed convolution are equivalent when stride is 1 (i.e., not up/down-sampling).

Image Encoder

Input:	Image x , $H_0 \times W_0 \times \{1, 3\}$
L1:	Conv2D(x , C=64, K=4, S=2, P=1)
L2:	Conv2D(L1, C=64, K=3, S=1, P=1)
L3:	Conv2D(L2, C=128, K=4, S=2, P=1)
L4:	Conv2D(L3, C=128, K=3, S=1, P=1)
L5:	Conv2D(L4, C=256, K=4, S=2, P=1)
L6:	Conv2D(L5, C=256, K=3, S=1, P=1)
L7:	Conv2D(L6, C=64, K=3, S=1, P=1)
Output:	Feature map $E(x)=L7$, $H_0/8 \times W_0/8 \times 64$

Video Generator

Input:	Feature map z , $16 \times H_0/8 \times W_0/8 \times 64$
L1:	TransposedConv3D(z , C=256, K=3, S=1, P=1)
L2:	TransposedConv3D(L1, C=256, K=3, S=1, P=1)
L3:	TransposedConv3D(L2, C=128, K=(3,4,4), S=(1,2,2), P=1)
L4:	TransposedConv3D(L3, C=128, K=3, S=1, P=1)
L5:	TransposedConv3D(L4, C=64, K=(3,4,4), S=(1,2,2), P=1)
L6:	TransposedConv3D(L5, C=64, K=3, S=1, P=1)
L7:	TransposedConv3D(L6, C={1, 3}, K=(3,4,4), S=(1,2,2), P=1)
Output:	Video $G(z)=L7$, $16 \times H_0 \times W_0 \times \{1, 3\}$

Video Discriminator (MoCoGAN-style)

Input:	Video X , $16 \times H_0 \times W_0 \times \{1, 3\}$
L1:	Conv3D(X , C=64, K=4, S=(1,2,2), P=(0,1,1))
L2:	Conv3D(L1, C=128, K=4, S=(1,2,2), P=(0,1,1))
L3:	Conv3D(L2, C=256, K=4, S=(1,2,2), P=(0,1,1))
L4:	Conv3D(L3, C=512, K=4, S=(1,2,2), P=(0,1,1))
L5:	Sigmoid(Linear(Flatten(L4), C=1))
Output:	Scalar $D_V(i)=L5 \in (0, 1)$

Image Discriminator (Resnet-based)

Notation: Shortcut(\cdot , C) = Conv2D(AvgPool(\cdot , K=2, S=2, P=0), C, K=1, S=1, P=0)

Input:	Image x , $H_0 \times W_0 \times \{1, 3\}$
L1:	Conv2D(x , C={1,3}, K=3, S=1, P=1)
L2:	Conv2D(L1, C=64, K=4, S=2, P=1) + Shortcut(x , C=64)
L3:	Conv2D(L2, C=64, K=3, S=1, P=1)
L4:	Conv2D(L3, C=128, K=4, S=2, P=1) + Shortcut(x , C=128)
L5:	Conv2D(L4, C=128, K=3, S=1, P=1)
L6:	Conv2D(L5, C=256, K=4, S=2, P=1) + Shortcut(x , C=256)
L7:	Conv2D(L6, C=256, K=3, S=1, P=1)
L8:	Conv2D(L7, C=512, K=4, S=2, P=1) + Shortcut(x , C=512)
L9:	Sigmoid(Linear(Flatten(L8), C=1))
Output:	Scalar $D_I(i)=L9 \in (0, 1)$

B DATASETS

Three well-known datasets are used: BAIR, KTH, and UCF101.

B.1 PREPROCESSING

All videos are center-cropped to square-sized frames and down-sampled to 64×64 . While BAIR and UCF101 are colored (3-channel RGB), KTH is in reality black and white (even though the raw videos come in a colored format). We treat KTH as 1-channel gray-scale. For computing the FVD, we simply duplicate the KTH videos channel-wise, since the pre-trained network requires 3-channel inputs.

BAIR and UCF101 contains many short videos, and we use a random (16-frame) sub-sequence of each video for training and the center (time-wise) one for evaluation (FVD and SSIM). Since KTH contains much fewer but longer videos, we divide each video up into sub-sequences and use all of them.

B.2 TRAIN/VALIDATION/TEST SPLITTING

BAIR:

- Test: Sequence 0–255.
- Validation: Sequence 256–2559.
- Train: All the rest.

KTH:

- Test: Person 17–25
- Validation: Person 16
- Train: Person 1–15

UCF101:

- Test: `testlist01.txt` from the "Action Recognition" train/test split).
- Validation: Randomly sampled 5% of `trainlist01.txt` (from the same zip file as above).
- Train: The other 95% of `trainlist01.txt`.

C ADDITIONAL RESULTS

To illustrate how the latent representation generator progressively transforms the feature map from layer to layer, we show what the generated video would look like if we were to connect the final video generator to one of intermediate layers of the latent representation generator *after* training the full model. Figure 4 shows the results for the last 8 layers (i.e., $l = 17, \dots, 24$) on some example. It is interesting to see that larger, more prominent features, such as the robot arm, become visible earlier in the model, while finer details, such as the cluttered background objects, tend to emerge in later stages.

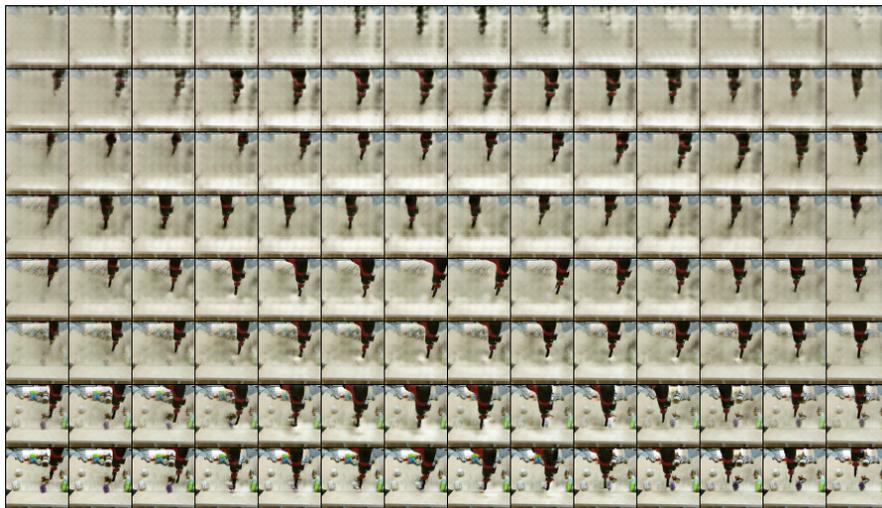


Figure 4: Sample output from intermediate representations. Each row corresponds to connecting the final video generator to one of the last 8 latent representation generator layers, from layer 17 (top) to 24 (bottom), the last of which is the actual output of the full model. Only the 14 in-between frames are shown.