

---

# Hierarchical Long-term Video Prediction without Supervision

---

Nevan Wichers<sup>1</sup> Ruben Villegas<sup>2</sup> Dumitru Erhan<sup>1</sup> Honglak Lee<sup>1</sup>

## Abstract

Much of recent research has been devoted to video prediction and generation, yet most of the previous works have demonstrated only limited success in generating videos on short-term horizons. The hierarchical video prediction method by Villegas et al. (2017b) is an example of a state-of-the-art method for long-term video prediction, but their method is limited because it requires ground truth annotation of high-level structures (e.g., human joint landmarks) at training time. Our network encodes the input frame, predicts a high-level encoding into the future, and then a decoder with access to the first frame produces the predicted image from the predicted encoding. The decoder also produces a mask that outlines the predicted foreground object (e.g., person) as a by-product. Unlike Villegas et al. (2017b), we develop a novel training method that jointly trains the encoder, the predictor, and the decoder together without high-level supervision; we further improve upon this by using an adversarial loss in the feature space to train the predictor. Our method can predict about 20 seconds into the future and provides better results compared to Denton and Fergus (2018) and Finn et al. (2016) on the Human 3.6M dataset.

## 1. Introduction

Building a model that is able to predict the future states of an environment from raw high-dimensional sensory data (e.g., video) has recently emerged as an important research problem in machine learning and computer vision. Models that are able to accurately predict the future can play a vital role in developing intelligent agents that interact with their environment (Jayaraman and Grauman, 2015; 2016; Finn et al., 2016).

---

<sup>1</sup>Google Brain, Mountain View, CA, USA. <sup>2</sup>Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA.. Correspondence to: Nevan Wichers <wichersn@google.com>.

Popular video prediction approaches focus on recursively observing the generated frames to make predictions farther into the future (Oh et al., 2015; Mathieu et al., 2016; Goroshin et al., 2015; Srivastava et al., 2015; Ranzato et al., 2014; Finn et al., 2016; Villegas et al., 2017a; Lotter et al., 2017). In order to make reasonable long-term frame predictions in natural videos, these approaches need to automatically identify the dynamics of the main factors of variation changing through time, while also being highly robust to pixel-level noise. However, it is common for the previously mentioned methods to generate quality predictions for the first few steps, but then the prediction dramatically degrades until all of the video context is lost or the predicted motion becomes static.

A hierarchical method makes predictions in a high-level information hierarchy (e.g., landmarks) and then decodes the predicted future in high-level back into low-level pixel space. The advantage of predicting the future in high-level space first is that the predictions degrade less quickly compared to predictions made solely in pixel space. The method by Villegas et al. (2017b) is an example of a hierarchical model; however, it requires ground truth human landmark annotations during training time. In this work, we explore ways to generate videos using a hierarchical model *without* requiring ground truth landmarks or other high-level structure annotations during training. In a similar fashion to Villegas et al. (2017b), the proposed network predicts the pixels of future video frames given the first few frames. Specifically, our network never observes any of the predicted frames, and the predicted future frames are driven solely by the high-level space predictions.

The contributions of our work are summarized below:

- An unsupervised approach for discovering high-level features necessary for long-term future prediction.
- A joint training strategy for generating high-level features from low-level features and low-level features from high-level features simultaneously.
- Use of adversarial training in feature space for improved high-level feature discovery and generation.
- Long-term pixel-level video prediction for about 20 seconds into the future for the Human 3.6M dataset.

## 2. Related Work

**Patch-level prediction** The video prediction problem was initially studied at the patch level containing synthetic motions (Sutskever et al., 2009; Michalski et al., 2014; Mitelman et al., 2014). Srivastava et al. (2015) and Ranzato et al. (2014) followed up by proposing methods that can handle prediction in natural videos. However, predicting patches encounters the well-known aperture problem that causes blockiness as prediction advances in time.

**Frame-level prediction on realistic videos.** More recently, the video prediction problem has been formulated at the full frame level using convolutional encoder/decoder networks as the main component. Finn et al. (2016) proposed a network that can perform next frame video prediction by explicitly predicting pixel movement. For each pixel in the previous frame, the network outputs a distribution over locations that a pixel is predicted to move. The possible movement a pixel can make are then averaged to obtain the final prediction. The network is trained end-to-end to minimize L2 loss. Mathieu et al. (2016) proposed adversarial training with multiscale convolutional networks to generate sharper pixel-level predictions in comparison to the conventional L2 loss. Villegas et al. (2017b) proposed a network that decomposes motion and content in video prediction and obtained more accurate results over Mathieu et al. (2016). Lotter et al. (2017) proposed a deep predictive coding network in which each layer learns to predict the lower-level difference between the future frame and current frame. As an alternative approach to convolutional encoder-decoder networks, Kalchbrenner et al. (2017) proposed an autoregressive generation scheme for improved prediction performance. In a concurrent work, Babaeizadeh et al. (2018) and Denton and Fergus (2018) proposed stochastic video prediction method based on recurrent variational autoencoders. Despite these efforts, long-term prediction on high-resolution natural videos beyond approximately 20 frames has been known to be very challenging.

**Long-term prediction.** Oh et al. (2015) proposed an action conditional convolutional encoder-decoder architecture that demonstrated high-quality long-term prediction performance on video games (e.g., Atari games), but it has not been applied to real-world video prediction. Villegas et al. (2017b) proposed a long-term prediction method using a hierarchical approach, but it requires the ground truth landmarks as supervision. Our work proposes several techniques to address this limitation.

## 3. Background

The hierarchical video prediction model in Villegas et al. (2017b) relieves the blurring problem observed in previous prediction approaches by modeling the video dynamics in high-level feature space. This approach enables the pre-

diction of many frames into the future. The hierarchical prediction model is described below.

To predict the image at timestep  $t$ , the following procedure is used: First, the high-level features  $p_t \in \mathbb{R}^l$  — in this case human pose landmarks — are estimated from the first  $C$  context frames. Next, an LSTM is used to predict the future landmark states  $\hat{p}_t \in \mathbb{R}^l$  given the landmarks estimated from the context frames as follows:

$$\begin{cases} [\hat{p}_t, H_t] = LSTM(p_{t-1}, H_{t-1}) & \text{if } t \leq C, \\ [\hat{p}_t, H_t] = LSTM(\hat{p}_{t-1}, H_{t-1}) & \text{if } t > C, \end{cases}$$

where  $H_t \in \mathbb{R}^h$  is the hidden state of the LSTM at timestep  $t$ . Note that the predicted  $\hat{p}_t$  after  $C$  timesteps is used to generate the video frames. Additionally, they remove the auto-regressive connections that feed  $\hat{p}_{t-1}$  back into LSTM making the prediction only depend on  $H_{t-1}$ . In our formulation, however, the prediction depends on both  $\hat{p}_{t-1}$  and  $H_{t-1}$ , but  $\hat{p}_{t-1}$  is not a vector of landmarks.

Once all  $\hat{p}_t$  are obtained, the visual analogy network (VAN) (Reed et al., 2015) generates the corresponding image at time  $t$ . VAN identifies the change between  $g(p_C)$  and  $g(\hat{p}_t)$ , where  $g(\cdot)$  is a fixed function that takes in landmarks and converts them into Gaussian heatmaps. Next, it applies the identified difference to image  $I_C$  to generate image  $I_t$ . The VAN does this by mapping images to a space where analogies can be represented by additions and subtractions. Therefore, the image at timestep  $t$  is computed by

$$\begin{aligned} \hat{I}_t &= VAN(p_C, \hat{p}_t, I_C) = \\ &f_{dec}(f_{pose}(g(\hat{p}_t)) - f_{pose}(g(p_C)) + f_{img}(I_C)). \end{aligned}$$

In contrast to Villegas et al. (2017b), our method does not require landmarks  $p_t$ , and therefore the dependence on the fixed function  $g(\cdot)$  is removed. Our method automatically discovers the features needed as input to the VAN for generating frame at time  $t$ . These features locate the object moving through time, and help our network focus on generating the moving object pixels in future frames. In the following section, we describe our method and training variations for unsupervised future frame prediction.

## 4. Method

### 4.1. Network Architecture

Our method uses a network architecture similar to Villegas et al. (2017b). However, our predictor LSTM and VAN do not require landmark annotations and can be trained jointly. In our model, the predictor LSTM is defined by

$$\begin{cases} [\hat{e}_t, H_t] = LSTM(e_{t-1}, H_{t-1}) & \text{if } t \leq C \\ [\hat{e}_t, H_t] = LSTM(\hat{e}_{t-1}, H_{t-1}) & \text{if } t > C, \end{cases} \quad (1)$$

where  $e_{t-1} \in \mathbb{R}^d$  is a general feature vector computed from an input image  $I_t$  by an encoder network, and  $\hat{e}_t \in \mathbb{R}^d$  is

the feature vector predicted by the LSTM. To compute the frame at time  $t$ , we use a variation of the *deep* version of the image analogy formulation from Reed et al. (2015). In contrast to Villegas et al. (2017b), we use the first frame in the input video to compute the future frames via image analogy. Therefore, the frame at time  $t$  is computed by

$$\bar{I}_t, M_t = \text{VAN}(e_1, \hat{e}_t, I_1) = f_{dec}(f_{enc}(\hat{e}_t) + T(f_{img}(I_1), f_{enc}(e_1), f_{enc}(\hat{e}_t))), \quad (2)$$

$$\hat{I}_t = \bar{I}_t \odot M_t + (1 - M_t) \odot I_1, \quad (3)$$

where  $f_{enc} : \mathbb{R}^d \rightarrow \mathbb{R}^{s \times s \times m}$  is a convolutional network that maps a feature vector into a feature tensor,  $f_{img} : \mathbb{R}^{h \times w \times c} \rightarrow \mathbb{R}^{s \times s \times m}$  is a convolutional network that maps an input image into a feature tensor,  $f_{dec} : \mathbb{R}^{s \times s \times m} \rightarrow \mathbb{R}^{h \times w \times c}$  is a deconvolutional network that maps a feature tensor into an image, and  $T(\cdot, \cdot, \cdot)$  is defined as follows:

$$T(x, y, z) = f_{analogy}([f_{diff}(x - y), z]), \quad (4)$$

where  $f_{diff} : \mathbb{R}^{s \times s \times m} \rightarrow \mathbb{R}^{s \times s \times m}$  computes a feature tensor from the difference between  $x$  and  $y$ ,  $[\cdot, \cdot]$  denotes a concatenation along the depth dimension of the input tensors, and  $f_{analogy} : \mathbb{R}^{s \times s \times 2m} \rightarrow \mathbb{R}^{s \times s \times m}$  computes the analogy feature tensor to be added to  $f_{enc}(\hat{e}_t)$ . Finally,  $M_t$  is a gating mechanism that enables our network to identify the moving objects in the video frames. In Equation 3, our network chooses pixels from the input frame that can simply be copied into the predicted frame, and pixels that need to be generated are chosen from  $\bar{I}_t$ . In Section 5, we show that the selected areas resemble the structure of moving objects in the input and the predicted frames.

## 4.2. Training Objective

These networks can be trained in multiple ways. In Villegas et al. (2017b), the predictor LSTM and VAN are trained separately using ground truth landmarks. In this work, we explore alternative ways of training these networks in the absence of ground truth annotations of high-level structures.

### 4.2.1. END-TO-END PREDICTION

One simple baseline method is to simply connect the VAN and the predictor LSTM together, and train them end-to-end (E2E). Our full network is optimized to minimize the L2 loss between the predicted image and the ground truth by:

$$\min\left(\sum_{t=1}^T L_2(\hat{I}_t, I_t)\right).$$

Figure 1 illustrates a diagram of this training scheme. Although a straightforward objective function is optimized, minimizing the L2 loss directly on the image outputs from previous observations tends to produce blurry predictions. This phenomenon has also been observed in several previous works (Mathieu et al., 2016; Villegas et al., 2017b;a).

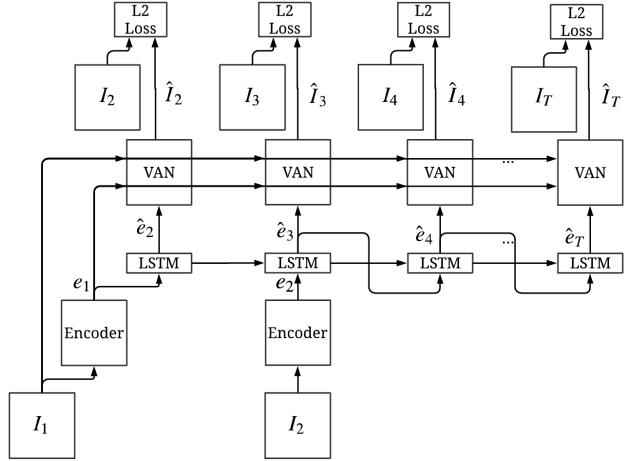


Figure 1. The E2E method. The first few frames are encoded and fed into the predictor as context. The predictor predicts the subsequent encodings, which the VAN uses to produce the pixel-level predictions. The average of the losses is minimized. This is the configuration of every method at inference time, even if the predictor and VAN are trained separately.

### 4.2.2. ENCODER PREDICTOR WITH ANALOGY MAKING

An alternative way to train our network is to constrain the features predicted by LSTM to be close to the outputs of the feature encoder (i.e.  $\hat{e}_t \approx e_t$ ). Simultaneously, the feature encoder outputs can be trained to be useful for analogy making. To accomplish this, we optimize the following objective function:

$$\min\left(\sum_{t=1}^T L_2(\hat{I}_t, I_t) + \alpha L_2(\hat{e}_t, e_t)\right), \quad (5)$$

where  $\hat{I}_t = \text{VAN}(e_1, e_t, I_1)$ ,  $e_t$  and  $e_1$  are both outputs of the feature encoder computed from the image at time  $t$  and the first image in the video, and  $\alpha$  is a balancing hyperparameter that controls the importance between predicting  $\hat{e}_t$  that is close to  $e_t$  and learning an encoding  $e_t$  that is good enough for image analogy.  $\alpha$  is used to prevent the predictor and encoder from both outputting the zero feature vector.

Figure 2 illustrates the flow of information by which the encoder and predictor are trained together with blue arrows, and the flow of information by which the VAN and encoder are trained together with red arrows. Separate gradient descent procedures (or optimizers, in TensorFlow parlance) could be used to minimize  $L_2(\hat{I}_t, I_t)$  and  $L_2(\hat{e}_t, e_t)$ , but we found that minimizing the sum is more accurate in our experiments. With this method, the predictor will generate the encoder outputs in future time steps, and the VAN will use the encoder output to produce the frame. The advantage of this training scheme is that the VAN learns to sharply predict the pixels since it is trained given the encoding from the ground truth frame. The predictor learns to approximate the ground truth high-level features from the encoder. There-

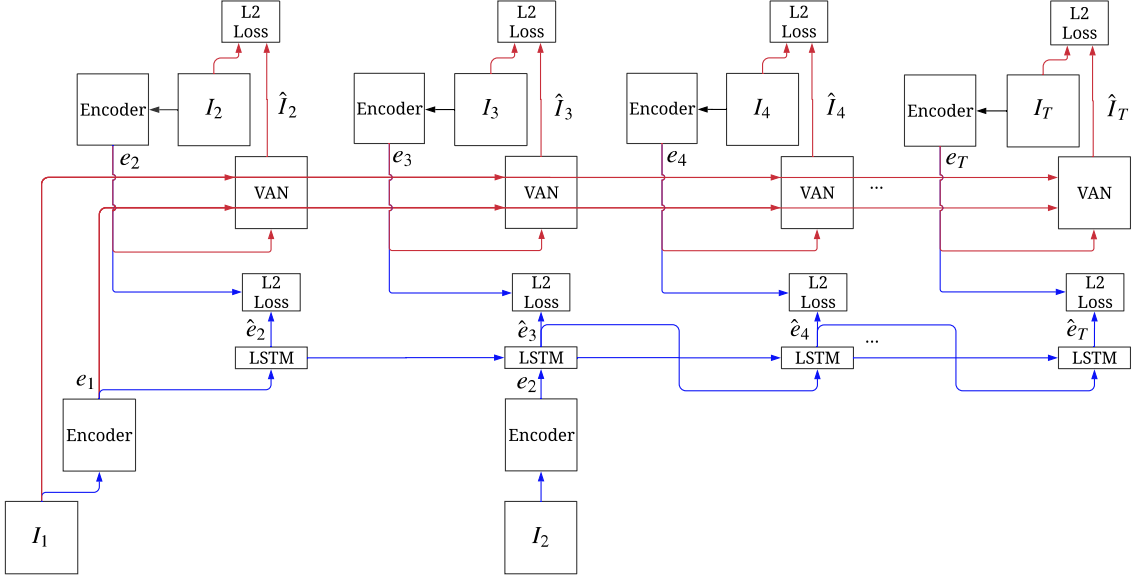


Figure 2. Blue lines represent the segment of the EPVA method in which the encoder and predictor are trained together. The encoder is trained to produce an encoding that is easy to predict, and the predictor is trained to predict that encoding into the future. Red lines represent the segment of the EPVA method in which the encoder and the VAN are trained together. The encoder is trained to produce an encoding that is informative to the VAN, while the VAN is trained to output the image given the encoding. The average of the losses in the diagram is minimized. This part of the method is similar to an autoencoder. Our method code is available at <https://bit.ly/2HqiHqx>

fore, at inference time the VAN knows how to decode the high-level structure features resulting in better predictions. Note that the encoder outputs  $e_t$  are given to VAN as input during training; however, the predictor outputs  $\hat{e}_t$  are given during testing. We refer to this method as EPVA.

The EPVA method works most accurately when experimented with  $\alpha$  starting small, around  $1e-7$ , and gradually increased to around 0.1 during training. As a result, the encoder will first be optimized to produce an informative encoding, then gradually optimized to make that encoding easy to predict by the predictor.

#### 4.2.3. EPVA WITH ADVERSARIAL LOSS IN PREDICTOR

A disadvantage of the EPVA training scheme alone is that the predictor is trained to minimize the L2 loss with respect to the encoder outputs. The L2 loss is notoriously known for the “blurriness effect,” and it causes our predictor LSTM to output blurry predictions in encoding space.

One solution to this problem is to use an adversarial loss (Goodfellow et al., 2014) between the predictor and encoder. We use an LSTM discriminator network, which takes a sequence of encodings and produces a score that indicates whether the encodings came from the predictor or the encoder network. We train the discriminator to minimize the improved Wasserstein loss (Gulrajani et al., 2017).

$$\min\left(\sum_{t=1}^T (D(\hat{e}) - D(e) + \lambda(\|\nabla_{\hat{e}} D(\hat{e})\|_2 - 1)^2)\right). \quad (6)$$

Here,  $e$  and  $\hat{e}$  are the sequence of inferred and predicted encodings respectively. We train both the encoder and the predictor, so we use a loss which takes both the encoder and predictor outputs into account. Therefore, we use the negative of the discriminator loss to optimize the generator.

$$\min\left(-\sum_{t=1}^T (D(\hat{e}) - D(e))\right) \quad (7)$$

We also still optimize the l2 loss between the predictor and encoder, weighted by a scale factor. This ensures the predictions will be accurate given the context frame. We also feed a Gaussian noise variable into the predictor in order to generate different results given the same input sequence. We found that the noise helps generate more complex predictions in practice.

In addition to passing the predictor or encoder output to the discriminator, we also pass the output of the VAN encoder, given the predictor or encoder output. This trains the predictor and encoder to encourage the VAN to produce similar quality images. This is achieved by substituting  $[f_{enc}(e), e]$  for  $e$  and  $[f_{enc}(\hat{e}), \hat{e}]$  for  $\hat{e}$  in the equations above, where  $f_{enc}$  is the VAN encoder. The encoder and VAN are trained together in the same way as previously discussed.

## 5. Experiments

We evaluated our methods on two datasets: the Human 3.6M dataset (Ionescu et al., 2014; 2011), and a toy dataset based

Table 1. Crowd-sourced human preference evaluation on the moving shapes dataset.

Method	Shape has correct color	Shape has wrong color	Shape disappeared
EPVA	96.9%	3.1%	0%
CDNA Baseline	24.6%	5.7%	69.7%

Figure 3. A visual comparison of the EPVA method and CDNA from Finn et al. (2016) as the baseline. This is a representative example of the quality of predictions from both methods. For videos please visit <https://bit.ly/2kS8r16>.

on videos of bouncing shapes. More sample videos and code to reproduce our results are available at our project website <https://bit.ly/2kS8r16>.

### 5.1. Long-term Prediction on a Toy Dataset

We train our method on a toy task with known factors of variation. We used a dataset with a generated shape that bounces around the image and changes size deterministically. We trained our EPVA method and the CDNA method from Finn et al. (2016) to predict 16 frames, given the first three frames as context. Both methods are evaluated on predicting approximately 1000 frames. We added noise to the LSTM states of the predictor network during training to help predict accurate motion further into the future. Results from a held out test set are described in the following.

After visually analyzing the results of both methods, we found that when the CDNA fails, the shape disappears entirely. In contrast, when the EPVA method fails, the shape changes color. See Figure 3 for sample predictions. For quantitative evaluation, we used a script to measure whether a shape was present from frames 1012 to 1022 and if that shape has the appropriate color. Table 1 shows the results averaged over 1000 runs. The CDNA method predicts a shape with the correct color about 25% of the time, and the EPVA method predicts a shape with the correct color about 97% of the time. The EPVA method sometimes fails by predicting the shape in the same location from frame to frame, but this is rare as the reader can confirm by examining the randomly sampled predictions on our project website. It is unrealistic to expect the methods to predict the location of the shape accurately in frame 1000 since small errors propagate in each prediction step.

### 5.2. Long-term Prediction on Human3.6M

In these experiments, we use subjects 1, 5, 6, 7, and 8 for training, and subject 9 for validation. Subject 11 results are reported in this paper for testing. We use 64 by 64 images, and subsample the dataset to 6.25 frames per second. We train the methods to predict 32 frames and the results in this paper show predictions over 126 frames. Each method is given the first five frames as context. In these images, the model predicts about 20 seconds into the future starting with 0.8 seconds of context. We use an encoding dimension of 64 for variations of our method on this dataset. The encoder in the EPVA method is initialized with the VGG network (Simonyan and Zisserman, 2015) pretrained on Imagenet (Deng et al., 2009). To speed up the convergence of the EPVA ADVERSARIAL method, we start training from a pretrained EPVA model.

We compare our method to the CDNA method in Finn et al. (2016) and the SVG-LP method in Denton and Fergus (2018). We trained each method with the same number of frames and context frames as ours. For Denton and Fergus (2018), we performed grid search on the  $\beta$  and learning rate to find the best configuration for this experiment, as well as, used a network as large as we could fit in the GPU. For Finn et al. (2016), we performed grid search on the learning rate. The method in Denton and Fergus (2018) can predict multiple futures, so we generate 5 futures for each context sequence, and compare against the one that most closely matches the ground truth in terms of SSIM. We find that this produces slightly better results than taking random predictions. Note that this protocol provides an unfair advantage to their method.

Figure 5 shows comparison to the baselines, and different variations of our method are compared in Figure 6. In Fig-

Table 2. Crowd-sourced human preference evaluation on the Human3.6M dataset.

Comparison	Ours is better	Same	Baseline is better
EPVA 1-127 vs Finn et al. (2016) 1-127	46.4%	40.7%	12.9%
EPVA ADV. 1-127 vs Finn et al. (2016) 1-127	73.9%	13.2%	12.9%
EPVA ADV. 63-127 vs Finn et al. (2016) 1-63	67.2%	17.5%	15.3%
EPVA ADV. 5-127 vs Denton and Fergus (2018) 5-127	58.2%	24.0%	17.8%

ure 5, we also show the discovered foreground motion segmentation mask from our method. This mask clearly shows that the feature embeddings from our encoder and predictor encode the rough location and outline of the moving human.

From visually analyzing the results, we found that the E2E and CDNA methods usually blur out very quickly. The EPVA method produces accurate predictions further into the future, but the figure sometimes disappears. The human predictions from the EPVA ADVERSARIAL method disappear less often and usually reappear in a later time step.

The CDNA (Finn et al., 2016) and the E2E methods produce blurry images because they are trained to minimize L2 loss directly. In the EPVA method, the predictor and VAN are trained separately. This prevents the VAN from learning to produce blurry images when the predictor is not confident. The predictions will be sharp as long as the predictor network outputs a valid encoding. The EPVA ADVERSARIAL method makes the predictor network more likely to produce a valid encoding since the discriminator is trained to produce valid predictions. We also observe that there is more movement in the EPVA ADVERSARIAL method.

### 5.2.1. PERSON DETECTOR EVALUATION

We propose to compare the methods quantitatively by considering whether the generated videos contain a recognizable person. To do this in an automated fashion, we ran a MobileNet (Howard et al., 2017) object detection model pre-trained on the MS-COCO (Lin et al., 2014) dataset for each of the generated frames. We record the confidence of the detector that a person (one of the MS-COCO labels) is in the image. We call this the “person score” (with value ranges from 0 to 1, with a higher score corresponding to a higher confidence level). The human detector achieves approximately an accuracy of 0.4 on the ground truth data. The results on each frame averaged over 1000 runs are shown in Figure 4. The EPVA ADVERSARIAL method stays relatively constant over the different frames. For longer-term predictions, the evaluation shows that the EPVA ADVERSARIAL method is significantly better than the baselines.

### 5.2.2. HUMAN EVALUATION

We also use a service similar to Mechanical Turk to collect comparisons of 1,000 generated videos from Finn et al. (2016) and Denton and Fergus (2018) to different variations of our method. The task presents videos generated by the

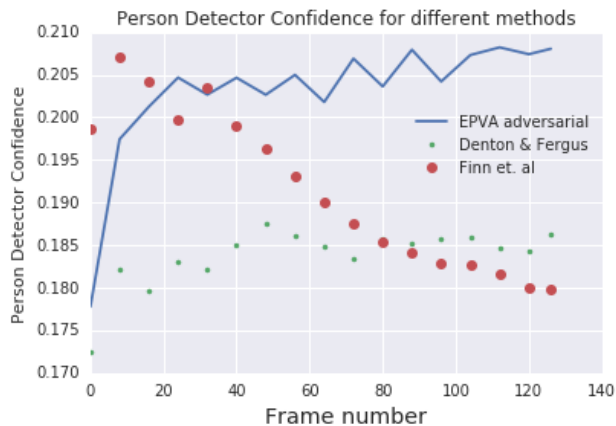


Figure 4. Confidence of the person detector that a person is recognized in the predicted frame (“person score”).

two methods side by side to human raters and asks them to confirm whether one of the videos is more realistic. The instructions tell raters to look for realistic motion, as well as a realistic person image. To evaluate the quality of the long-term predictions from the EPVA ADVERSARIAL method, we compare frames 64 to 127 of the EPVA ADVERSARIAL method to frames 1 to 63 of Finn et al. (2016). We evaluate frames 5-127 of Denton and Fergus (2018) against 5-127 of ours since their method isn’t designed to produce good results for the context frames.

The summary results are shown in Table 2. From these results, we conclude the following: the EPVA method generates significantly better long-term predictions than Finn et al. (2016). Further, the EPVA ADVERSARIAL method is a dramatic improvement over the EPVA method. The EPVA ADVERSARIAL method is capable of high-quality long-term predictions, as shown by frames 64 to 127 (seconds 10 to 20) of the EPVA ADVERSARIAL method being rated higher than frames 1-63 of Finn et al. (2016). The EPVA ADVERSARIAL is also significantly better than Denton and Fergus (2018) even after choosing the best out of 5 predictions after comparing with the ground truth in terms of SSIM.

### 5.2.3. POSE REGRESSION FROM LEARNED FEATURES

We perform experiments using the learned encoder features for human pose regression. We compare against a baseline based on features computed using the VGG network (Simonyan and Zisserman, 2015) trained for object recognition. The features are used as input to a 2-layer MLP, and trained to output human pose landmarks. The MLP trained with

### Long-term Video Prediction without Supervision




Figure 5. Comparison of the generated videos from EPVA with the ADVERSARIAL LOSS (ours), CDNA (Finn et al., 2016), and SVG-LP (Denton and Fergus, 2018). We let each method predict 127 frames and show the time steps indicated on top of the figure. The person completely disappears in all the predictions generated using Finn et al. (2016). For the SVG-LP method (Denton and Fergus, 2018), the person either stops moving or almost vanishes into the background. The EPVA with ADVERSARIAL LOSS method produces sharp predictions in comparison to the baselines. Additionally, we show the discovered foreground motion segmentation mask that allows our network to delete the human in the input frame (static mask in the top example) and generate the human in the future frames (moving mask in the top example). Please refer to our project website for video results: <https://bit.ly/2kS8r16>.

## Long-term Video Prediction without Supervision

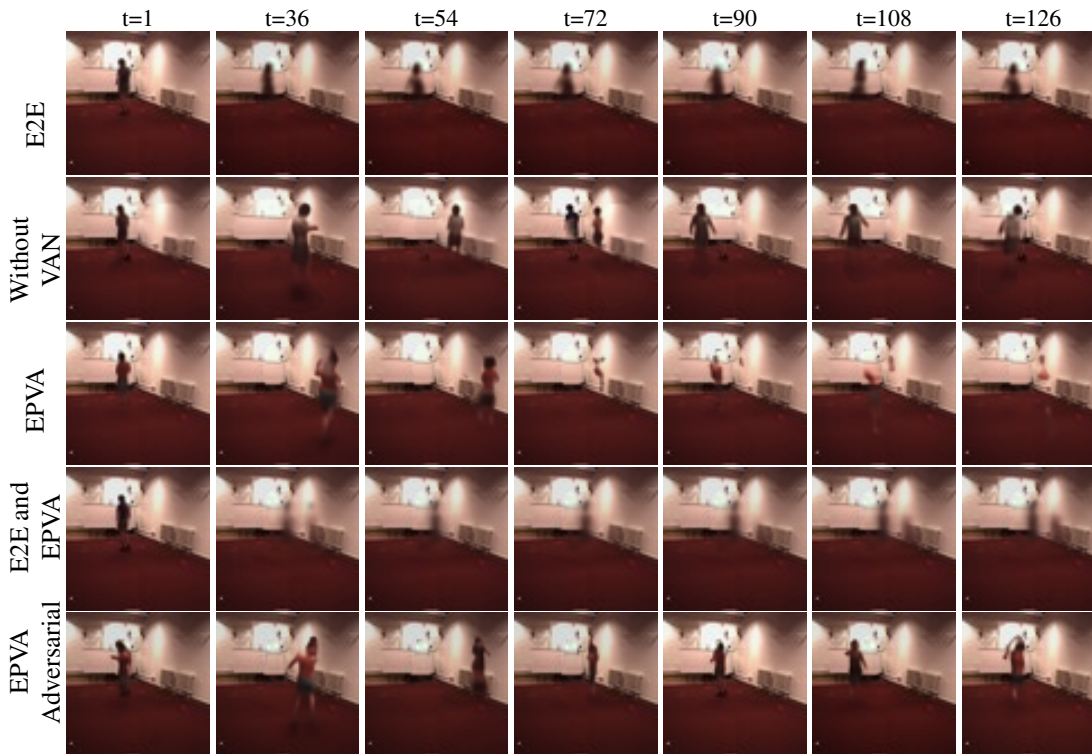


Figure 6. Ablative study illustration. We present comparisons between different variations of our architecture: E2E, loss without VAN, EPVA, combined E2E and EPVA loss, and our best model configuration (EPVA ADVERSARIAL). See our project website for videos.

our features achieves an error of 0.0687 against an error of 0.0758 from the baseline features. This is a relative improvement of approximately 9%. This along with the generated masks shows the usefulness of our discovered features.

### 5.3. Ablation Studies

We perform the following experiments to test different variations of the network and training. We hypothesize that using a VAN improves the quality of the predictions. To test this, we train a version of the network with the VAN replaced by a decoder network that only had access to the encoding and not the first observed frame.

In this method, as well as the methods with the VAN, the decoder outputs a mask that controls whether to use its own output, or the pixels of the first frame. Thus, the decoder will have to set the mask values to not use the pixels from the first frame that correspond to the image of the person. Without the VAN, the network is often unable to set the mask values to completely remove the human from the first frame when predicting frames beyond 32. This is because the network is not always given access to the first frame, so it has to represent both foreground and background information in the prediction, which degrades over time. Refer to Figure 6 for comparison.

We also tried to use a hybrid objective that combines E2E and EPVA losses, but the videos generated from this method are more blurry than the videos from the EPVA method.

These are called E2E and EPVA in Figure 6. Finally, we also trained and evaluated the EPVA method with 10 frames of context instead of 5. We found that this didn't improve the long-term prediction results.

## 6. Conclusion

We presented hierarchical long-term video prediction approaches that do not require ground truth high-level structure annotations. The proposed EPVA method has the limitation of the predictions occasionally disappearing, but it generates sharper images for a longer period of time compared to Finn et al. (2016), and the E2E method. By applying adversarial loss in the higher-level feature space, our EPVA ADVERSARIAL method generates more realistic predictions compared to all of the presented baselines including Finn et al. (2016) and Denton and Fergus (2018). This result suggests that it is beneficial to apply an adversarial loss in the higher-level feature space. For future work, applying other techniques in feature space such as the variational method described in Babaeizadeh et al. (2018) could enable our network to generate multiple future trajectories.

**Acknowledgments.** We thank colleagues at Google Brain and anonymous reviewers for their constructive feedback and suggestions about this work. We also thank Emily Denton for providing her code available for comparison. R. Villegas was supported by Rackham Merit Fellowship.



## References

- M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine. Stochastic variational video prediction. In *ICLR*, 2018.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- E. Denton and R. Fergus. Stochastic video generation with a learned prior. In *ICML*, 2018.
- C. Finn, I. Goodfellow, and S. Levine. Unsupervised learning for physical interaction through video prediction. In *NIPS*, 2016.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014.
- R. Goroshin, M. Mathieu, and Y. LeCun. Learning to linearize under uncertainty. In *NIPS*. 2015.
- I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein GANs. In *NIPS*, 2017.
- A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint:1704.04861*, 2017.
- C. Ionescu, F. Li, and C. Sminchisescu. Latent structured models for human pose estimation. In *ICCV*, 2011.
- C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014.
- D. Jayaraman and K. Grauman. Learning image representations tied to ego-motion. In *ICCV*. 2015.
- D. Jayaraman and K. Grauman. Look-ahead before you leap: end-to-end active recognition by forecasting the effect of motion. In *ECCV*, 2016.
- N. Kalchbrenner, A. v. d. Oord, K. Simonyan, I. Danihelka, O. Vinyals, A. Graves, and K. Kavukcuoglu. Video pixel networks. In *ICML*, 2017.
- T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *ECCV*, 2014.
- W. Lotter, G. Kreiman, and D. Cox. Deep predictive coding networks for video prediction and unsupervised learning. In *ICLR*. 2017.
- M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. In *ICLR*. 2016.
- V. Michalski, R. Memisevic, and K. Konda. Modeling deep temporal dependencies with recurrent “grammar cells”. In *NIPS*, 2014.
- R. Mittelman, B. Kuipers, S. Savarese, and H. Lee. Structured recurrent temporal restricted Boltzmann machines. In *ICML*. 2014.
- J. Oh, X. Guo, H. Lee, R. L. Lewis, and S. Singh. Action-conditional video prediction using deep networks in Atari games. In *NIPS*. 2015.
- M. Ranzato, A. Szlam, J. Bruna, M. Mathieu, R. Collobert, and S. Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv preprint:1412.6604*, 2014.
- S. E. Reed, Y. Zhang, Y. Zhang, and H. Lee. Deep visual analogy-making. In *NIPS*. 2015.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- N. Srivastava, E. Mansimov, and R. Salakhudinov. Unsupervised learning of video representations using LSTMs. In *ICML*. 2015.
- I. Sutskever, G. E. Hinton, and G. W. Taylor. The recurrent temporal restricted Boltzmann machine. In *NIPS*. 2009.
- R. Villegas, J. Yang, S. Hong, X. Lin, and H. Lee. Decomposing motion and content for natural video sequence prediction. In *ICLR*. 2017a.
- R. Villegas, J. Yang, Y. Zou, S. Sohn, X. Lin, and H. Lee. Learning to generate long-term future via hierarchical prediction. In *ICML*, 2017b.