

Improved Conditional VRNNs for Video Prediction

Lluís Castrejon*
 Mila, Université de Montréal
 lluis.castrejon@gmail.com

Nicolas Ballas
 Facebook AI Research
 nballas@fb.com

Aaron Courville
 CIFAR, Mila, Université de Montréal
 aaron.courville@umontreal.ca

Abstract

Predicting future frames for a video sequence is a challenging generative modeling task. Promising approaches include probabilistic latent variable models such as the Variational Auto-Encoder. While VAEs can handle uncertainty and model multiple possible future outcomes, they have a tendency to produce blurry predictions. In this work we argue that this is a sign of underfitting. To address this issue, we propose to increase the expressiveness of the latent distributions and to use higher capacity likelihood models. Our approach relies on a hierarchy of latent variables, which defines a family of flexible prior and posterior distributions in order to better model the probability of future sequences. We validate our proposal through a series of ablation experiments and compare our approach to current state-of-the-art latent variable models. Our method performs favorably under several metrics in three different datasets.

1. Introduction

We investigate the task of video prediction, a particular instantiation of self-supervision [7, 9] where generative models learn to predict future frames in a video. Training such models does not require any annotated data, yet the models need to capture a notion of the complex dynamics of real-world phenomena (such as physical interactions) to generate coherent sequences.

Uncertainty is an inherent difficulty associated with video prediction, as many future outcomes are plausible for a given sequence of observations [1, 5]. Predictions from deterministic models rapidly degrade over time as uncertainty grows, converging to an average of the possible future outcomes [34]. To address this issue, probabilistic latent variable models such as Variational Auto-Encoders (VAEs) [19, 31], and more specifically Variational Recurrent Neural Networks (VRNNs) [3], have been proposed for video prediction [1, 5]. These models define a prior distri-

*Work done while part-time at Facebook AI Research.

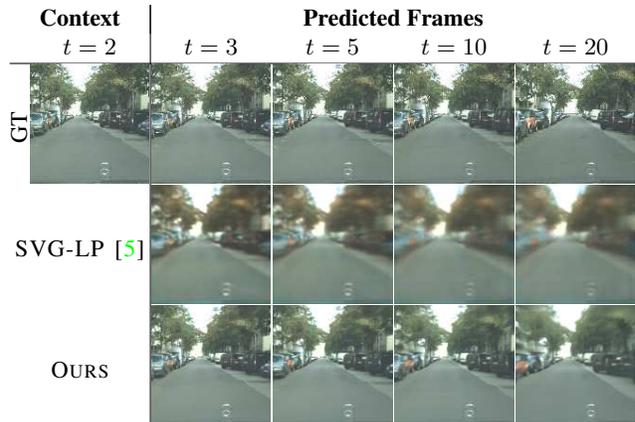


Figure 1: **Can generative models predict the future?** We propose an improved VAE model for video prediction. Our model uses hierarchical latents and a higher capacity likelihood network to improve upon previous VAE approaches, generating more visually appealing samples that remain coherent for longer temporal horizons.

bution over a set of latent variables, allowing different samples from these latents to capture multiple outcomes.

It has been empirically observed that VAE and VRNN-based models produce blurry predictions [21, 22]. This tendency is usually attributed to the use of a similarity metric in pixel space [21, 26] such as Mean Squared Error (corresponding to a log-likelihood loss under a fully factorized Gaussian distribution). This has led to alternative models such as VAE-GAN [21, 22], which extends the traditional VAE objective with an adversarial loss in order to obtain more visually compelling generations.

In addition, the lack of expressive latent distributions has been shown to lead to poor model fitting [13]. Training VAEs involves defining an approximate posterior distribution over the latent variables which models their probability after the generated data has been observed. If the approximate posterior is too constrained, it will not be able to match the true posterior distribution and this will prevent the model from accurately fitting the training data. On the other hand, the prior distribution over the latent variables can be interpreted as a model of uncertainty.

The decoder or likelihood network needs to transform latent samples into data observations covering all plausible outcomes. Given a simple prior, this transformation can be very complex and require high capacity networks. We hypothesize that the reduced expressiveness of current VRNN models is limiting the quality of their predictions and investigate two main directions to improve video prediction models. First, we propose to scale the capacity of the likelihood network. We empirically demonstrate that by using a high capacity decoder we can ease the latent modeling problem and better fit the data.

Second, we introduce more flexible posterior and prior distributions [32]. Current video prediction models usually rely on one shallow level of latent variables and the prior and approximate posterior are parameterized using diagonal Gaussian distributions [1]. We extend the VRNN formulation by proposing a hierarchical variant that uses multiple levels of latents per timestep.

Models leveraging a hierarchy of latents are known to be hard to optimize as they are required to backpropagate through a stack of stochastic latent variables, usually resulting in models that only make use of a small subset of the latents [19, 25, 32]. We mitigate this problem by using a warmup regime for the KL loss [33] and a dense connectivity pattern [14, 24] between the input and latent variables. Specifically, each stochastic latent variable is connected to the input and to all subsequent stochastic levels in the hierarchy. Our empirical findings confirm that only with these techniques our model is able to take advantage of different layers in a latent hierarchy.

We validate our hierarchical VRNN in three datasets with varying levels of future uncertainty and realism: Stochastic Moving MNIST [5], the BAIR Push Dataset [8] and Cityscapes [4]. When compared to current state of the art models [5, 22], our approach performs favorably under several metrics. In particular for the BAIR Push Dataset, our hierarchical-VRNN shows an improvement of 44% in Video Fréchet Distance (FVD) [36] and 9.8% in term of LPIPS score [43] over SVG-LP [5], the previous best VAE-based model. It also achieves a similar FVD than the SAVP VAE-GAN model [22], while showing a 11.2% improvement in terms of LPIPS over this baseline.

2. Related Work

Initial video prediction approaches relied on deterministic models. Ranzato et al. [29] divided frames into patches and predicted their evolution in time given previous neighboring patches. In [34] Srivastava et al. used LSTM networks on pretrained image embeddings to predict the future. Similarly, Oh et al. [27] used LSTMs on CNN representations to predict frames in Atari games when given the player actions.

ConvLSTMs [42] adapt the LSTM equations to spatial

feature maps by replacing matrix multiplications with convolutions. They were originally used for precipitation nowcasting and are commonly used for video prediction. PredNet [23] and ContextVP [2] are deterministic models using LSTMs and other recurrent architectures that have been used for video prediction of car mounted scenes.

Other works have proposed to disentangle the motion and context of the frames to generate [37, 35, 6]. They assume that a scene can be decomposed as multiple objects, which allows them to use a fixed representation for the background. Our approach does not follow this modeling assumption and instead tries to capture the uncertainty in the future.

Autoregressive models [16, 30] approximate the full joint data distribution $p(x_1, x_2, \dots, x_N)$ over pixels, which allows them to capture complex pixel dependencies but at the expense of making their inference mechanism slow and not scalable to high resolutions. Latent variable models using GANs [10] were proposed in [39, 38, 35]. Training pure GAN video models is still an open research direction: training is unstable and most models require auxiliary losses. Concurrent work [20] has proposed a fully invertible model for video.

A more successful approach so far has been based on VAE [19, 31]/VRNN [3] models. SV2P [1] proposed to capture sequence uncertainty in a single set of latent variables kept fixed for each predicted sequence. SVG [5] adopted the VRNN formulation [3], introducing per-step latent variables (SVG-FP) and a variant with a learned prior (SVG-LP), which makes the prior at a certain timestep a function of previous frames. In recent work, SAVP [22] proposed to use the VAE-GAN [21] framework for video, a hybrid model that offers a trade-off between VAEs and GANs. Our model extends the VRNN formulation by introducing a hierarchy of latents to better approximate the data likelihood.

There are multiple works addressing hierarchical VAEs for non-sequential data [28, 25, 33, 18]. While hierarchical VAEs can model more flexible latent distributions, training them is usually difficult due to the multiple layers of conditional latents [32], resulting in most latents being unused. Ladder Variational Autoencoders [33] proposed a series of techniques to partially alleviate this issue. IAF [18] used a similar architecture to Ladder VAEs and extended it with a novel normalizing flow. Recent work [24] has trained very deep hierarchical models that produce visually compelling samples. We extend hierarchical latent variable models to sequential data and apply them to video prediction.

3. Preliminaries

We follow previous work in video prediction [5]. Given D context frames $\mathbf{c} = (c_1, c_2, \dots, c_D)$ and the T following future frames $\mathbf{x} = (x_1, x_2, \dots, x_T)$, our goal is to learn a

generative model that maximizes the probability $p(\mathbf{x}|\mathbf{c})$.

VRNN follows the VAE formalism and introduces a set of latent variables $\mathbf{z} = (z_1, z_2, \dots, z_T)$ to capture the variations in the observed variables \mathbf{x} at each timestep t . It defines a *likelihood* model $p(\mathbf{x}|\mathbf{z}, \mathbf{c}) = \prod_{t=1}^T p(x_t|\mathbf{z}_{\leq t}, \mathbf{x}_{<t}, \mathbf{c})$ and a *prior* distribution $p(\mathbf{z}|\mathbf{c}) = \prod_{t=1}^T p(z_t|\mathbf{z}_{<t}, \mathbf{x}_{<t}, \mathbf{c})$ which are parametrized in an autoregressive manner; *i.e.* at each timestep observed and latent variables are conditioned on the past latent samples and observed frames. VRNN therefore uses a *learned prior* [3, 5]. Taking into account the temporal structure of the data, the probability $p(\mathbf{x}, \mathbf{z}|\mathbf{c})$ is factorized as

$$p(\mathbf{x}, \mathbf{z}|\mathbf{c}) = \prod_{t=1}^T p(x_t|\mathbf{z}_{\leq t}, \mathbf{x}_{<t}, \mathbf{c})p(z_t|\mathbf{z}_{<t}, \mathbf{x}_{<t}, \mathbf{c}). \quad (1)$$

Computing $p(\mathbf{x}|\mathbf{c})$ requires marginalizing over the latent variables \mathbf{z} , which is computationally intractable. Instead, VRNN relies on Variation Inference [15] and defines an *amortized approximate posterior* $q(\mathbf{z}|\mathbf{x}, \mathbf{c}) = \prod_{t=1}^T q(z_t|\mathbf{z}_{<t}, \mathbf{x}_{\leq t}, \mathbf{c})$ that approximates the *true posterior* distribution $p(\mathbf{z}|\mathbf{x}, \mathbf{c})$. We then can derive the evidence lower bound (ELBO), a lower bound to the marginal log-likelihood $p(\mathbf{x}|\mathbf{c})$:

$$\begin{aligned} \log p(\mathbf{x}|\mathbf{c}) &\geq \sum_{t=1}^T \mathbb{E}_{q(\mathbf{z}_{\leq t}|\mathbf{x}_{\leq t}, \mathbf{c})} \log p(x_t|\mathbf{z}_{\leq t}, \mathbf{x}_{<t}, \mathbf{c}) \\ &\quad - D_{KL}(q(z_t|\mathbf{z}_{<t}, \mathbf{x}_{\leq t}, \mathbf{c})||p(z_t|\mathbf{z}_{<t}, \mathbf{x}_{<t}, \mathbf{c})) \end{aligned} \quad (2)$$

VRNN can be optimized to fit the training data by maximizing the ELBO using stochastic backpropagation and the reparameterization trick [19, 31].

4. Hierarchical VRNN

We now introduce a hierarchical version of the VRNN model. At each timestep, we consider L levels of latent variables $\mathbf{z}_t = (z_t^1, \dots, z_t^L)$. We then further factorize the latent *prior* as

$$p(\mathbf{z}_t|\mathbf{z}_{<t}, \mathbf{x}_{<t}, \mathbf{c}) = \prod_{l=1}^L p(z_t^l|\mathbf{z}_t^{<l}, \mathbf{z}_{<t}^1, \mathbf{x}_{<t}, \mathbf{c}). \quad (3)$$

The sampling process of the latent variable z_t^l now depends on the latent variables from previous time steps $\mathbf{z}_{<t}^1$ for that level and on the latent variables of the previous levels at the current timestep $\mathbf{z}_t^{<l}$. Similarly, we can write the *approximate posterior* as:

$$q(\mathbf{z}_t|\mathbf{z}_{<t}, \mathbf{x}_{\leq t}, \mathbf{c}) = \prod_{l=1}^L q(z_t^l|\mathbf{z}_t^{<l}, \mathbf{z}_{<t}^1, \mathbf{x}_{\leq t}, \mathbf{c}). \quad (4)$$

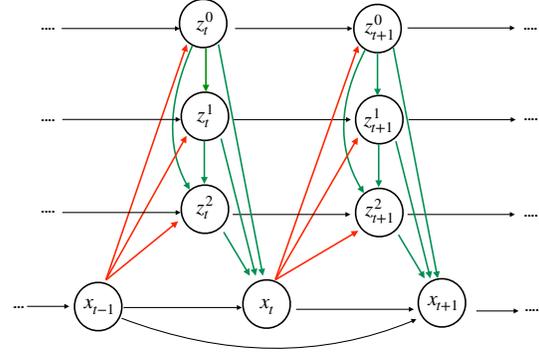


Figure 2: **Graphical model for the learned prior with the dense latent connectivity pattern.** Arrows in **red** show the connections from the input at the previous timestep to current latent variables. Arrows in **green** highlight skip connections between latent variables and connections to outputs. Arrows in **black** indicate recurrent temporal connections. We empirically observe that this dense-connectivity pattern eases the training of latent hierarchies.

Using eq. 3 and eq.4, we can rewrite the ELBO as

$$\begin{aligned} \log p(\mathbf{x}|\mathbf{c}) &\geq \sum_{t=1}^T [\mathbb{E}_{q(\mathbf{z}_t|\mathbf{z}_{<t}, \mathbf{x}_{\leq t}, \mathbf{c})} \log p(x_t|\mathbf{z}_{\leq t}, \mathbf{x}_{<t}, \mathbf{c}) \\ &\quad - \sum_{l=1}^L D_{KL}(q(z_t^l|\mathbf{z}_t^{<l}, \mathbf{z}_{<t}^1, \mathbf{x}_{\leq t}, \mathbf{c})||p(z_t^l|\mathbf{z}_t^{<l}, \mathbf{z}_{<t}^1, \mathbf{x}_{<t}, \mathbf{c}))]. \end{aligned} \quad (5)$$

Refer to the Appendix for a full derivation of the ELBO.

4.1. Dense Latent Connectivity

Training a hierarchy of latent variables is known to be challenging as it requires to backpropagate through multiple stochastic layers. Usually this results in models that only use one specific level of the hierarchy [19, 25, 32]. To ease the optimization we use a dense connectivity pattern between latent levels both for the prior and the approximate posterior, following [14, 24].

Fig 2 illustrates the dense connection of the learned prior (refer to the Appendix for the approximate posterior). For each latent level, the prior and posterior are implemented using recurrent neural networks which take as input a deterministic transformation of x_{t-1} (red arrows in Fig 2), and to all the latent variables from the previous levels (green arrows in in Fig 2). In addition, each latent variable has a direct connection to the output variables x_t .

4.2. Model Parametrization

We now describe an instantiation of the VRNN model that we will use in the experiments, illustrated in Fig. 3. First we compute features for each context frame and use them to initialize the hidden state of the

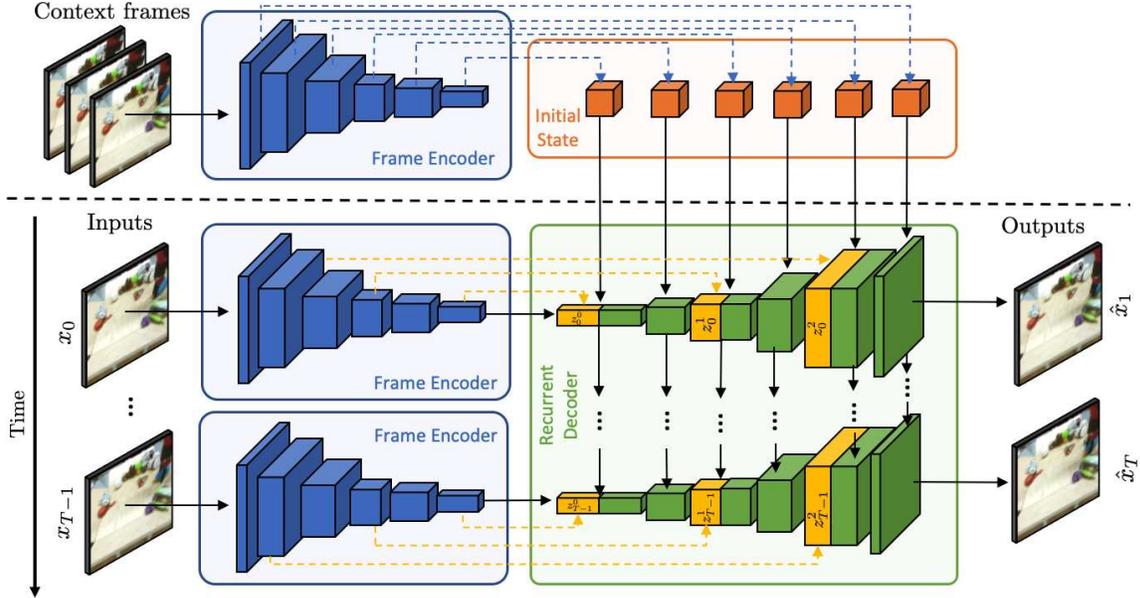


Figure 3: **Model Parametrization.** Our model uses a CNN to encode frames individually. The representation of the context frames is used to initialize the states of the prior, posterior and likelihood networks, all of which use recurrent networks. At each timestep, the decoder receives an encoding of the previous frame, a set of latent variables (either from the prior or the posterior) and its previous hidden state and predicts the next frame in the sequence.

prior/posterior/decoder networks, all of which have recurrent components. At a given timestep t , the model takes as input the latent variable samples $\mathbf{z}_t = (z_t^0, \dots, z_t^L)$ with the embedding of the previously generated frame x_{t-1} and outputs the next frame \hat{x}_t . During training we draw latent samples from the approximate posterior distribution $q(\mathbf{z}_t | \mathbf{z}_{<t}, \mathbf{x}_{<t}, \mathbf{c})$ and maximize the ELBO. To generate unseen sequences, we sample \mathbf{z}_t from the learned prior $p(\mathbf{z}_t | \mathbf{z}_{<t}, \mathbf{x}_{<t}, \mathbf{c})$. Note that since we have multiple levels of conditional latents we use ancestral sampling to generate \mathbf{z}_t , *i.e.* we first sample from the top level of the hierarchy and we then sequentially sample the lower levels conditioning on the sampled values of the previous layers in the hierarchy.

Frame Encoder We use a 2D CNN with ResNet [12] blocks and max-pooling layers to represent input frames.

Prior/Approximate Posterior We parametrize both the prior and the posterior as a hierarchy of diagonal Normal distributions $\mathcal{N}(\mu, \sigma)$, where the parameters μ and σ are recurrent functions of samples from i) previous levels in the hierarchy and ii) the frame encoder features. Each level in the hierarchy operates at a different spatial resolution, with the top level features operating at a 1×1 resolution, *i.e.* not having a spatial topology. At a given timestep t , the parameters for a specific latent level z_t^l are given by a ConvLSTM that consumes i) a previous hidden state, ii) samples from the previous levels in the hierarchy $z_t^{<l}$, iii) the feature map of a frame with the same spatial resolution as the ConvLSTM. For the prior network, the input frame embedding

corresponds to the previously generated frame x_{t-1} , while for the posterior the input comes from the frame to generate x_t .

Likelihood/Frame Decoder At each timestep t , the decoder takes a representation of the previously generated frame x_{t-1} and the samples $\mathbf{z}_t = (z_t^1, \dots, z_t^L)$ and generates x_t according to $p(x_t | \mathbf{z}_t, \mathbf{x}_{<t}, \mathbf{c})$. The decoder consists of ConvLSTMs interleaved with transposed convolutions that upscale the feature maps back to the input resolution.

Initial State The initial states of our prior, posterior and decoder/likelihood models are functions of the context. We use small CNNs to initialize each of the ConvLSTMs layers used in the VAE components.

5. Experiments

All our models are trained with Adam [17] and a batch size of $b = 128$ on Nvidia DGX-1s. We use a learning rate warmup [11] starting with an initial learning rate $\lambda = 2e-5$ that is linearly increased at each timestep until reaching $\lambda = 1.6e-4$ in 5 epochs. We use $\beta_1 = 0.5$ and $\beta_2 = 0.9$ and weight decay $\delta = 1e-4$. We train the autoregressive components of our models using teacher forcing [41].

Our models are also trained using beta warmup [33], which consists in gradually increasing the weight of the KL divergence in the ELBO, turning the model from an unregularized Autoencoder into a VAE progressively. VAEs trained with beta warmup usually encode more information in the latent variables. Refer to the Appendix for a complete description of our models.

MODEL	PARAMETERS	TRAIN/TEST ELBO(↑)
1-ConvLSTM	62.22M	-3237/-3826
3-ConvLSTM	86.64M	-1948/-2355
6-ConvLSTM	93.81M	-1279.21/-1731.31
+ higher capacity	194.15M	-1113.31/-1589.72

Table 1: **Ablation - Likelihood** We compare models with different number of recurrent layers for the likelihood network. We observe that the model performance increases monotonically as we add more ConvLSTMs. We further increase the size of the recurrent hidden states for the 6-ConvLSTM model (+ higher capacity variant), also leading to a better data fit. These results suggest that current video prediction models might underfit the data because of reduced likelihood capacity.

5.1. Ablation Study

We first investigate the importance of each VRNN component, namely the likelihood, the prior and the posterior. We focus on the BAIR Push dataset [8] with 64x64 color sequences of a robotic arm interacting with children toys in a sandbox. Similarly to previous works [22], we use trajectories 256 to 511 as our test set and the rest for training, resulting in the 43264 train and 256 test sequences. At training we randomly subsample 12 frames from each train sequence, use the first 2 frames as the context, and learn to predict the remaining 10 frames. To evaluate the different model variations, we report the training objective (ELBO) obtained for the training set and the test set.

5.1.1 Scaling the Likelihood Model

We assess the importance of the likelihood model $p(x_t | \mathbf{z}_{\leq t}, \mathbf{x}_{< t}, \mathbf{c})$. For this purpose, we build a VRNN with a single level of latent variables and modify the number of ConvLSTM layers in the decoder. Our aim is to investigate whether increasing the capacity of the latents to observations mapping results in better predictions.

In this experiment, our baseline likelihood model has one LSTM at 1x1 spatial resolution. We then gradually replace convolutional layers in the decoder with ConvLSTM layers, which increases the amount of information that can be carried from previous timesteps and, by extension, also increases the overall likelihood model capacity. We compare to a model with 3 ConvLSTM layers at resolutions 1x1, 4x4 and 8x8 and a model with 6 ConvLSTM layers at 1x1, 4x4, 8x8, 16x16, 32x32 and 64x64. Additionally, we also increase the size of the ConvLSTM layers for the model with 6 layers as another way of adding capacity.

Results can be found in Fig 1. We observe that, as a general trend, both the training and test ELBO decrease as we increase the model capacity, which suggests that current video prediction models might operate in an underfitting regime and benefit from higher capacity decoders.

MODEL	PARAMETERS	TRAIN/TEST ELBO (↑)
1	166.55M	-1141.85/-1536.93
1-8	220.60M	-989.39/-1313.02
1-8-32	230.74M	-883.10/-1162.24
1-8-16-32	245.19M	-956.63/-1256.22
Naive Training	224.18M	-1127.33/-1440.58
BW	224.18M	-1101.39/-1440.62
Dense	230.74M	-1182.60/-1547.05
BW and Dense	230.74M	-883.10/-1162.24

Table 2: **Ablation - Hierarchy of Latents Top half:** We compare a VRNN baseline with a single level of latents with no spatial topology (1), a model with two levels of latents at resolutions 1x1 and 8x8 (1-8), our full model with three levels of latents at 1x1, 8x8 and 32x32 (1-8-32), and a model with 4 levels of latents (1-8-16-32). Adding more levels of latents leads to a better fit, with reduced ELBOs. However, adding too many levels of latents reduces the performance due to the difficulties in training hierarchical latent variable models. *Bottom half:* To highlight the difficulties in training hierarchies of latents, we investigate the effects of using beta warmup (BW) [33] and having a dense connectivity (Dense) between latents when training the 1-8-32 model. Without these techniques the hierarchy of latents does not improve performance.

5.1.2 More Flexible Prior and Posterior

We now investigate the importance of having more flexible prior and approximate posterior distributions and augment the 6-ConvLSTM VRNN model with a hierarchy of latent variables. For all models, we fix the frame encoder and likelihood model¹ and change the networks that estimate the learned prior $p(\mathbf{z}_t | \mathbf{z}_{< t}, \mathbf{x}_{< t}, \mathbf{c})$, and the approximate posterior $q(\mathbf{z}_t | \mathbf{z}_{< t}, \mathbf{x}_{< t}, \mathbf{c})$ over the latent variables. All these models use a dense connectivity pattern and beta warmup.

We compare a VRNN baseline with a single level of latents with no spatial topology, with a model with two levels of latents at resolutions 1x1 and 8x8 (1-8), three levels of latents at 1x1, 8x8 and 32x32 (1-8-32), and four levels of latents (1-8-16-32) in the top half of Table 2. All models are trained with beta warmup and dense latent connectivity. We observe that in general adding more levels of latents with higher resolution reduces the train and test ELBOs, supporting the hypothesis that a more flexible prior and posterior leads to a better data fit. However, we observe diminishing returns past 3 levels, as our 1-8-16-32 model does not outperform the 3 layers model. We attribute this to the difficulties in training deep hierarchies of latents, which remains a challenging optimization problem.

To further highlight the difficulties in training hierarchies of latents, we investigate the importance of using beta warmup [33] and having a dense connectivity between la-

¹To add the multiple levels of latents we need to modify the decoder and slightly increase the number of parameters. However, most (> 85%) of the added capacity goes towards the prior and posterior networks.

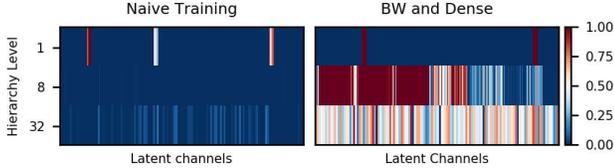


Figure 4: **Average normalized KL per latent channel.** We visualize the mean normalized KL for each latent channel for models from Table 2. Without beta warmup and dense connectivity the hierarchy of latents is underutilized, with most information being encoded in a few latents of the top level. In contrast, the same model with these techniques utilizes all latent levels.

tents. The results of this experiment can be found in the bottom half of Table 2. We observe that these techniques are required to make our 1-8-32 model make use of the hierarchy of latents and improve upon the single level model.

This is analyzed in more detail in Fig 4, where we visualize the KL between the prior and the posterior distributions for the test sequences of the BAIR Push dataset for the 1-8-32 model and the variant without warmup or dense connectivity (Naive training). We consider a channel to be active if its average KL is higher than 0.01 following [24], and consider that a unit with a KL higher than 0.15 is maximally activated. We observe that without these techniques the model only uses a few latents of the top level in the hierarchy. However, when using beta warmup and a dense connectivity most of the latents are active across levels.

5.2. Comparisons to Previous Approaches

Next, we compare our single latent level VRNN (Ours w/o Hier) and our 3-level hierarchical VRNN (Ours w/ Hier) to previous approaches on three datasets.

5.2.1 Evaluation and Metrics

Defining evaluation metrics for video prediction is an open research question. We want models to predict *realistic* sequences that *cover* all possible outcomes. Unfortunately, we are not aware of any metric that reflects all these aspects.

To measure *coverage* we adopt the evaluation protocol in [5, 22]. For each ground truth test sequence, we sample N predictions from the model which are conditioned on the test sequence initial frames. Then we find the sample that best matches the ground truth sequence according to a given metric and report that metric value. Some common metric choices are Mean-Square Error (MSE), Structural Similarity (SSIM) [40] or Peak Signal-to-Noise Ratio (PSNR). In practice, these metrics have been shown to not correlate well with human judgement as they tend to prefer blurry predictions over sharper but imperfect generations [43, 22, 36]. LPIPS [43], on the other hand, is a perceptual metric that compares CNN features and has better correlation to human judgment. For this evaluation we

MODEL	FVD (\downarrow)	LPIPS (\downarrow)	SSIM (\uparrow)
SVG-LP [5]	90.81	0.153 ± 0.03	0.668 ± 0.04
Ours w/o HIER	63.81	0.102 ± 0.04	0.763 ± 0.09
Ours w/ HIER	57.17	0.103 ± 0.03	0.760 ± 0.08

Table 3: **Stochastic Moving MNIST.** We compute the FVD metric between samples from different models and test sequences as well as the average LPIPS and SSIM of the best sample for each test sequence. Our models outperform the SVG-LP baseline on all metrics by a significant margin. While our model with hierarchical latent variables obtains a better FVD score, both variants obtain comparable results in this relatively simple dataset.

generate $N = 100$ samples following previous work and use SSIM and LPIPS as metrics. We have empirically observed that the metrics have small variance when using 100 samples. We report the metric average over the test set.

We also use the recently proposed Fréchet Video Distance (FVD), to measure sample *realism*. FVD uses features from a 3D CNN and has also been shown to correlate well with human perception [36]. FVD compares populations of samples to assert whether they were both generated by the same distribution (it does not compare pairs of ground truth/generated frames directly). We form the ground truth population by using all the test sequences with their context. For the predicted population we randomly sample one video out of the N generated for each test sequence. We repeat this process 5 times and report the mean of the FVD scores obtained, which are similar across samplings.

5.2.2 Stochastic Moving MNIST [5]

Stochastic Moving MNIST is a synthetic dataset proposed in [5] which consists of sequences of MNIST digits moving over a black background and bouncing off the frame borders. As opposed to the original Moving MNIST dataset [34] with deterministic motion, Stochastic Moving MNIST has uncertain digit trajectories - the digits bounce off the border with a random new trajectory. We train two variants of our model and compare to the SVG-LP baseline [5], for which we use a pretrained model from the official codebase. All models are trained using 5 frames of context and 10 future frames to predict. To evaluate the models, we follow the procedure in [5] described in section 5.2.1.

We report the results of the experiment in Table 3. We observe that both versions of our model (with/out the latent hierarchy) outperform the SVG-LP baseline by a significant margin on all metrics. Note that LPIPS and FVD might not be suited to this dataset as they use features from CNNs trained on real world images, but we report them for completeness. Visually, our samples (see Appendix) depict the digits with reduced degradation over time. There are small differences between the two versions of our model, suggesting that the extra expressiveness of the hierarchical model

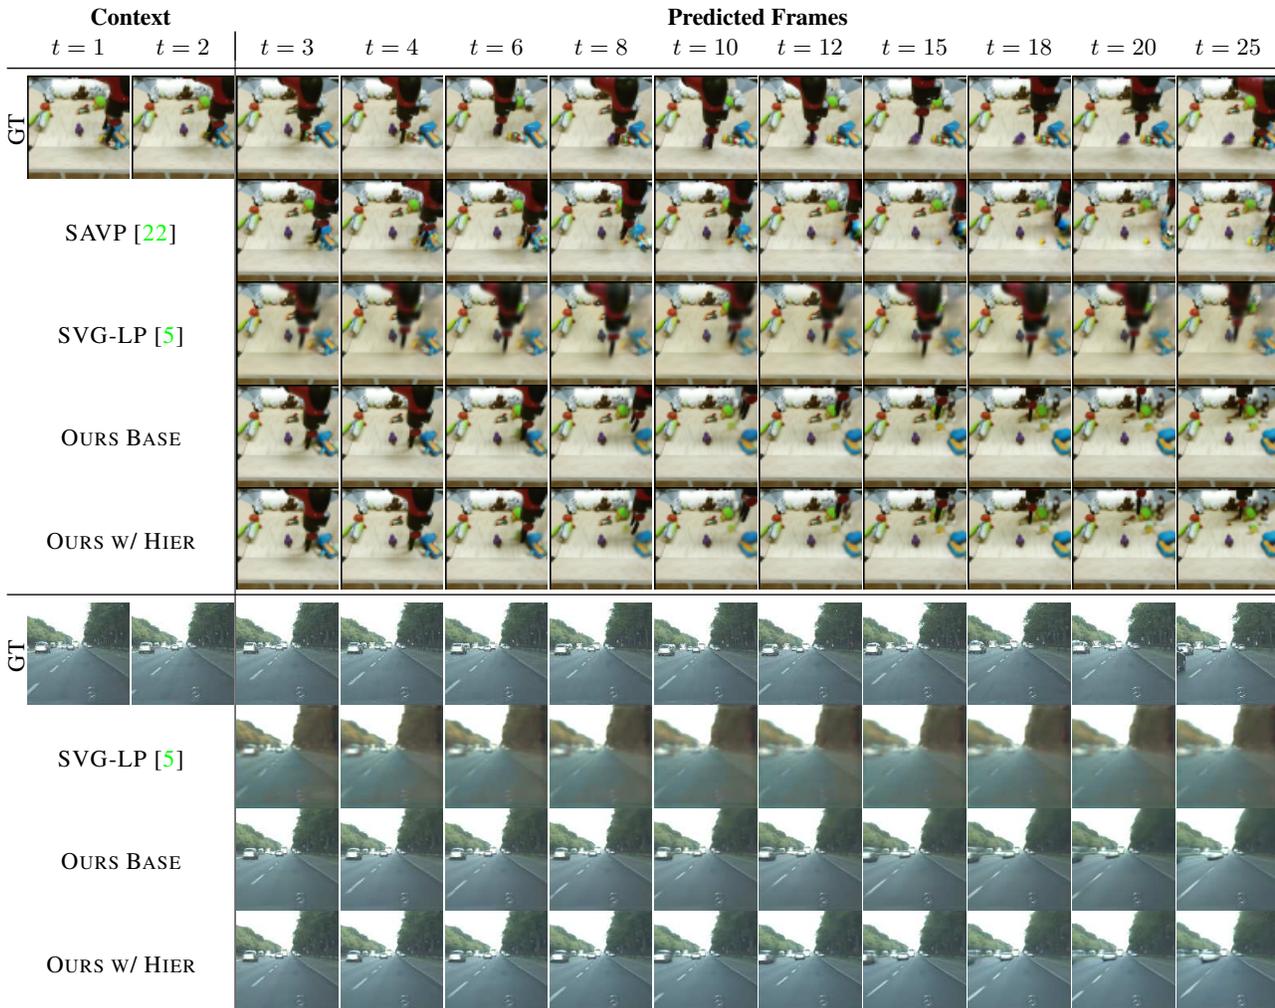


Figure 5: **Selected Samples for BAIR Push and Cityscapes.** We show a sequence for BAIR Push and Cityscapes and random generations from our model and baselines. On BAIR Push we observe that the SAVP predictions are crisp but sometimes depict inconsistent arm-object interactions. SVG-LP produces blurry predictions in uncertain areas such as occluded parts of the background or those showing object interactions. Our model generates plausible interactions with reduced blurriness relatively to SVG-LP. On Cityscapes, the SVG-LP baseline is unable to model any motion. Our model, using a hierarchy of latents, generates more visually compelling predictions. More samples can be found in the Appendix.

is not necessary in this synthetic dataset.

5.2.3 BAIR Push

We compare our VRNN models to SVG-LP [5] and SAVP [22]. We use their official implementations and pre-trained models to reproduce their results. We use the experimental setup of previous works [5, 22], using 2 context frames and generating 28 frames.

Results can be found on Fig 6. When the robotic arm is interacting with an object, SVG-LP tends to generate blurry predictions characterized by a high FVD score. SAVP exhibits a lower FVD as it produces more realistically looking predictions. However, SAVP does not have a better cov-

erage of the ground truth sequences compared to SVG-LP as measured by LPIPS and SSIM. By inspecting the SAVP samples we notice that the SAVP generations tend to be sharper but sometimes they exhibit temporal inconsistencies or implausible interactions (see Fig 5).

Our models obtain better scores than SVG-LP, the previous best VAE model. This highlights the importance of having a high-capacity likelihood model. In addition, our hierarchical VRNN further improves both the FVD and LPIPS metrics, suggesting that the hierarchy of latents helps modeling the data. In particular, it shows an improvement of 44% in terms of FVD and 9.8% in terms of LPIPS over SVG-LP. It also achieves a similar FVD to that of SAVP, while outperforming it in terms of LPIPS by 11.2%.

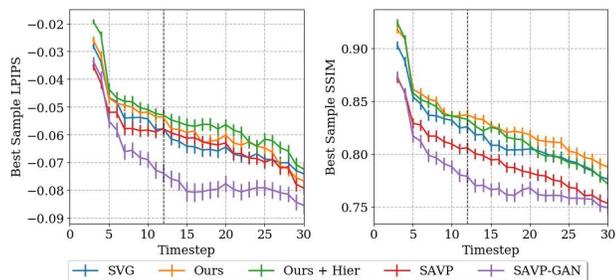


Figure 6: **BAIR Push - Results.** *Left:* We show the evolution in time of the Average LPIPS and SSIM of the best predicted sample per test sequence. *Right:* We report the Average FVD, SSIM and LPIPS of the best sample for each test sequence. Compared to SVG-LP, both our model with a single level of latents and the hierarchical models improve all metrics. Compared to SAVP, we obtain better LPIPS and SSIM. Our model with a single level of latents performs better in SSIM but worse on perceptual metrics. When adding the hierarchy of latents, our model matches the FVD of SAVP and improves the LPIPS, indicating samples of similar visual quality and better coverage of the ground-truth sequences.

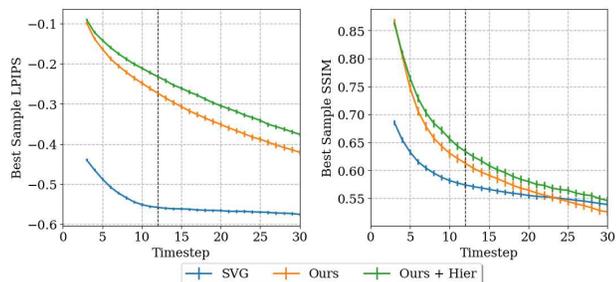


Figure 7: **Cityscapes - Quantitative Results** We report FVD, SSIM and LPIPS scores on Cityscapes at 128x128 resolution for the SVG-LP [5] baseline and two variants of our model. Increasing the capacity of the likelihood model brings an improvement in all metrics over the SVG baseline. When adding a hierarchy of latents we observe further improvements, validating its usefulness. Even though SVG matches our models in SSIM at later timesteps, this does not correlate well with human judgement, as the generated SVG samples show more blurriness (see Fig. 5).

5.2.4 Cityscapes

The Cityscapes dataset contains sequences recorded from a car driving around multiple cities under different conditions. Cityscapes is a challenging dataset - while contiguous frames are locally similar, uncertainty grows significantly with time. Compared to previous datasets, the backgrounds in Cityscapes do not stay constant across time.

We consider sequences with 30 frames from the training set cities for a total of 1877 train sequences and randomly select 256 test sequences. We use 2 context and 10 prediction frames to train the models. At test time we predict 28 frames following the BAIR Push setup. We preprocess the videos by taking a 1024x1024 center crop and resizing them to 128x128 pixels. For evaluating the models we use the standard setup where we generate 100 samples per test sequence and report FVD, SSIM and LPIPS metrics. Since none of the baselines from previous experiments are trained on Cityscapes, we use the official SVG implementation (that defines models with 128x128 inputs) and train a SVG-LP model. We train all models for 100 epochs.

Results can be found in Fig. 7. SVG-LP has difficulties

generating motion, usually predicting a static image similar to the last context frame. In contrast, our model without a hierarchy of latents is able to model the changing scene. When adding hierarchical latents our model is able to capture more fine-grained details, and as a result, it produces more visually appealing samples with a boost in all metrics. We note that the SSIM scores for SVG-LP match those of our models at later timesteps in the prediction, however this does not translate to better samples as can be seen in Fig. 5 or in the Appendix. This further suggests that SSIM is not a suitable metric to evaluate video prediction models.

6. Conclusions

We propose a hierarchical VRNN for video prediction featuring an improved likelihood model and a hierarchy of latents. Our approach compares favorably to current state of the art models in terms of the FVD, LPIPS and SSIM metrics, producing visually appealing and coherent samples. Our results show that current video prediction models benefit from increased capacity and opens the door to further gains with flexible higher capacity generative models.

References

- [1] Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H Campbell, and Sergey Levine. Stochastic variational video prediction. *arXiv preprint arXiv:1710.11252*, 2017. [1](#), [2](#)
- [2] Wonmin Byeon, Qin Wang, Rupesh Kumar Srivastava, and Petros Koumoutsakos. Contextvp: Fully context-aware video prediction. In *The European Conference on Computer Vision (ECCV)*, September 2018. [2](#)
- [3] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pages 2980–2988, 2015. [1](#), [2](#), [3](#), [11](#)
- [4] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. [2](#)
- [5] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *International Conference on Machine Learning*, pages 1182–1191, 2018. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [14](#), [15](#), [16](#), [17](#), [18](#), [19](#), [20](#), [21](#), [22](#), [23](#)
- [6] Emily L Denton et al. Unsupervised learning of disentangled representations from video. In *Advances in Neural Information Processing Systems*, pages 4414–4423, 2017. [2](#)
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [1](#)
- [8] Frederik Ebert, Chelsea Finn, Alex X Lee, and Sergey Levine. Self-supervised visual planning with temporal skip connections. *arXiv preprint arXiv:1710.05268*, 2017. [2](#), [5](#)
- [9] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. [1](#)
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. [2](#)
- [11] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. [4](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [4](#)
- [13] Matthew D Hoffman and Matthew J Johnson. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *Workshop in Advances in Approximate Bayesian Inference, NIPS*, 2016. [1](#)
- [14] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. [2](#), [3](#)
- [15] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999. [3](#)
- [16] Nal Kalchbrenner, Aäron van den Oord, Karen Simonyan, Ivo Danihelka, Oriol Vinyals, Alex Graves, and Koray Kavukcuoglu. Video pixel networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1771–1779. JMLR. org, 2017. [2](#)
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [4](#)
- [18] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016. [2](#)
- [19] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [1](#), [2](#), [3](#)
- [20] Manoj Kumar, Mohammad Babaeizadeh, Dumitru Erhan, Chelsea Finn, Sergey Levine, Laurent Dinh, and Durk Kingma. Videoflow: A flow-based generative model for video. *arXiv preprint arXiv:1903.01434*, 2019. [2](#)
- [21] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. *arXiv preprint arXiv:1512.09300*, 2015. [1](#), [2](#)
- [22] Alex X Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#), [17](#), [18](#), [19](#), [20](#)
- [23] William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. *ICLR*, 2016. [2](#)
- [24] Lars Maaløe, Marco Fraccaro, Valentin Liévin, and Ole Winther. Biva: A very deep hierarchy of latent variables for generative modeling. *arXiv preprint arXiv:1902.02102*, 2019. [2](#), [3](#), [6](#)
- [25] Lars Maaløe, Casper Kaae Sønderby, Søren Kaae Sønderby, and Ole Winther. Auxiliary deep generative models. *arXiv preprint arXiv:1602.05473*, 2016. [2](#), [3](#)
- [26] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. [1](#)
- [27] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L Lewis, and Satinder Singh. Action-conditional video prediction using deep networks in atari games. In *Advances in neural information processing systems*, pages 2863–2871, 2015. [2](#)
- [28] Rajesh Ranganath, Dustin Tran, and David Blei. Hierarchical variational models. In *International Conference on Machine Learning*, pages 324–333, 2016. [2](#)
- [29] MarcAurelio Ranzato, Arthur Szlam, Joan Bruna, Michael Mathieu, Ronan Collobert, and Sumit Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv preprint arXiv:1412.6604*, 2014. [2](#)

- [30] Scott Reed, Aäron van den Oord, Nal Kalchbrenner, Sergio Gómez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov, and Nando de Freitas. Parallel multiscale autoregressive density estimation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2912–2921. JMLR. org, 2017. 2
- [31] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014. 1, 2, 3
- [32] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. How to train deep variational autoencoders and probabilistic ladder networks. In *33rd International Conference on Machine Learning (ICML 2016) International Conference on Machine Learning*, 2016. 2, 3
- [33] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder variational autoencoders. In *Advances in neural information processing systems*, pages 3738–3746, 2016. 2, 4, 5
- [34] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852, 2015. 1, 2, 6
- [35] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 2
- [36] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 2, 6
- [37] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. *arXiv preprint arXiv:1706.08033*, 2017. 2
- [38] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 98–106, 2016. 2
- [39] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *Advances In Neural Information Processing Systems*, pages 613–621, 2016. 2
- [40] Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [41] Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989. 4
- [42] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015. 2
- [43] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 2, 6