

Efficient Video Prediction via Sparsely Conditioned Flow Matching

Aram Davtyan*, Sepehr Sameni*, Paolo Favaro

Computer Vision Group, Institute of Computer Science, University of Bern, Switzerland

{aram.davtyan, sepehr.sameni, paolo.favaro}@unibe.ch

Abstract

We introduce a novel generative model for video prediction based on latent flow matching, an efficient alternative to diffusion-based models. In contrast to prior work, we keep the high costs of modeling the past during training and inference at bay by conditioning only on a small random set of past frames at each integration step of the image generation process. Moreover, to enable the generation of high-resolution videos and to speed up the training, we work in the latent space of a pretrained VQGAN. Finally, we propose to approximate the initial condition of the flow ODE with the previous noisy frame. This allows to reduce the number of integration steps and hence, speed up the sampling at inference time. We call our model Random frame conditioned flow Integration for VidEo pRediction, or, in short, RIVER. We show that RIVER achieves superior or on par performance compared to prior work on common video prediction benchmarks, while requiring an order of magnitude fewer computational resources. Project website: <https://araachie.github.io/river>.

1. Introduction

Video prediction, *i.e.*, the task of predicting future frames given past ones, is a fundamental component of an agent that needs to interact with an environment [5]. This capability enables planning and advanced reasoning, especially when other agents are in the scene [21, 20, 74]. More in general, however, a video prediction model that can generalize to new unseen scenarios needs to implicitly understand the scene, *i.e.*, detect and classify objects, learn how each object moves and interacts, estimate the 3D shape and location of the objects, model the laws of physics of the environment, and so on. In addition to naturally leading to rich and powerful representations of videos, this task does not require any labeling and is thus an excellent candidate for learning from readily available unannotated datasets.

While the literature in video prediction is by now rela-

*Equal contribution.

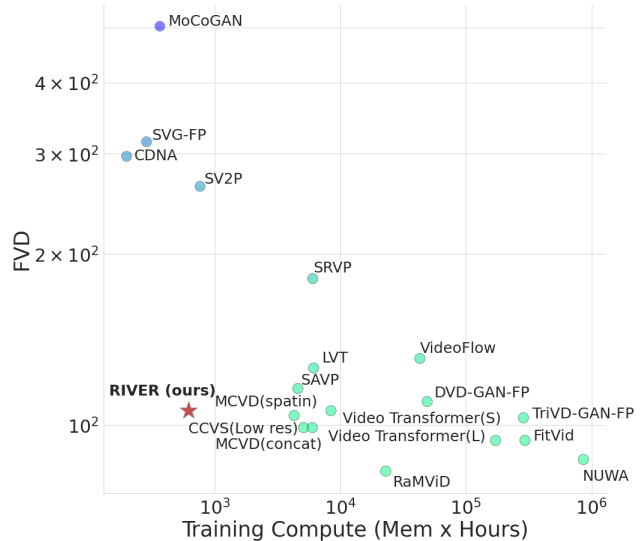


Figure 1. RIVER achieves an ideal trade-off between quality of generated videos and compute needed to train the model. This makes research on video models more easily scalable.

tively rich [13, 5, 39], the quality of the predicted frames has been achieving realistic levels only recently [67, 33, 27, 76]. This has been mostly due to the exceptional complexity of this task and the difficulty of training models that can generalize well to unseen (but in-domain) data.

To address these challenges, we propose a novel training procedure for video prediction that is computationally efficient and delivers high quality frame prediction. One of the key challenges of synthesizing realistic predicted frames is to ensure the temporal consistency of the generated sequence. To this aim, conditioning on as many past frames as possible is a desirable requirement. In fact, with only two past frames it is possible to predict only constant motions at test time, and for general complex motions, such as object interactions (*e.g.*, a ball bouncing off a cube in CLEVRER [77]), many more frames are needed. However, conditioning on many past frames comes either at the sacrifice of the video quality or at a high computational cost, as shown in Figure 1. In the literature, we see two main approaches to address these issues: 1) models that take a fixed

large temporal window of past frames as input and 2) models that compress all the past into a state, such as recurrent neural networks (RNNs) [4, 13, 5]. Fixed window models require considerable memory and computations both during training and at inference time. Although methods such as Flexible Diffusion [27] can gain considerable performance by choosing carefully non contiguous past frames, their computational cost still remains demanding. RNNs also require considerable memory and computations resources at training time, as they always need to feed a sequence from the beginning to learn how to predict the next frame. Moreover, training these models is typically challenging due to the vanishing gradients.

In the recent growing field of diffusion models for image generation, the 3DiM method [70] introduces the idea of a sparse conditioning on randomly chosen scene views during the diffusion process, and showed impressive results in novel view synthesis. In our approach, we adapt this idea to the case of video prediction by also conditioning the generation of the next frame on a randomly chosen sparse set of past frames during the diffusion process. In practice, this is an effective remedy to limit the computational complexity of the model at both training and test time, because the conditioning at each training step is limited to a small set of past frames, but the frame prediction at test time can incorporate an arbitrary number of past frames (also efficiently). To further speed up the training and generation of videos at test time, we compress videos via VQGAN autoencoding [19] and work in the latent space. This design choice has been shown to work well in the case of image generation [54] and to enable the efficient generation of images at high resolution. Unlike other methods that employ a temporal VQGAN [75, 26], we adopt a per-frame VQGAN approach to minimize the training cost. Additionally, we incorporate a refinement network (more details in section 3.3) to improve the frame quality and to correct any temporal inconsistencies between pairs of frames during post-processing. We gain another significant performance boost both in terms of better convergence at training time and in terms of better image quality generation, by adapting Flow Matching [42] to video prediction. The key insight is that it is possible to build an explicit mapping of a noise instance to an image sample, and diffusion models are the result of a specific choice of such mapping, which has been shown experimentally to be sub-optimal [42]. Finally, to make our method more efficient at inference time, we introduce a *warm-start sampling*. In the case of video prediction, the content changes slowly over time. Thus, we expect that a very good guess for the next frame is the current frame itself. Therefore, we propose to speed up the integration of the flow to generate the next frame by starting from a noisy current frame rather than from zero-mean Gaussian noise. We call our method Random frame conditioned flow In-

tegration for Video Prediction (RIVER). We demonstrate RIVER on common video prediction benchmarks and show that it performs on par or better than state of the art methods, while being much more efficient to train. We also show that RIVER can be used for video generation and interpolation and can predict non-trivial long-term object interactions. In summary, our contributions are the design of a video prediction model that

1. Extends flow matching to video prediction;
2. Is efficient to train and to use at test time;
3. Can be conditioned on arbitrarily many past frames;
4. Is efficient at test time (via warm-start sampling).

2. Prior work

Conventional Methods. Video prediction models are used to generate realistic future frames of a video sequence based on past frames. Up until recently, most video prediction models relied on Recurrent Neural Networks (RNNs) in the bottleneck of a convolutional autoencoder [5]. Training such models is known to be challenging, but so is handling the stochastic nature of generative tasks. Most approaches in this domain benefit from variational methods [36]. These methods usually use a recurrent network [32] conditioned on a global [4] or a per-frame [13, 5] latent variable. To model longer sequences, hierarchical variational models have been proposed [72, 40]. Another approach that scales to long sequences is keypoint-based video prediction [49, 35, 23], where they cast the problem into first keypoint dynamics prediction (with a variational method) and then pixel prediction. The usage of such methods in complex datasets with non homogeneous keypoints is yet to be seen. To overcome the blurry results that variational methods have been known for, SAVP [39] combined both an adversarial loss [25] and a VAE [36]. Instead, Mathieu et al. [47] used a multiscale network for video prediction. Many methods deal with motion and content separately [60, 66, 41, 14]. The fundamental problem with using GANs for video prediction is to ensure long-term temporal consistency [3]. This issue is tackled in recent works [12, 45], but at a huge computational cost.

Transformers and Quantized Latents. Following the success of large language models [8], autoregressive transformers [64] emerged in the video synthesis domain and are replacing RNNs. However, because of the attention mechanism, transformers incur high computational cost that scales quadratically with the number of inputs. In order to scale these methods to long and higher-resolution videos, an established method is to predict vector-quantized codes [63] (usually obtained with VQGAN [19]) either per frame [38, 52, 26, 57] or a set of frames [75], instead of pixels.

Diffusion Methods. Following the impressive results of score-based diffusion models [59] on image generation [15, 53], several researchers extended these models for either video generation [31] or video prediction [67, 33, 27, 28, 76]. Even though unconditional models can be used to approximate conditional distributions (as in video prediction) [31], it has been shown that directly modeling the conditional distribution yields a better performance [61]. MCVD [67] uses masking to train a single model capable of generating past, future, or intermediate frames. Masking allows this model to generate longer sequences by applying a moving window, even though it was only trained with a fixed number of frames. However, a common shortcoming of all these methods is that conditioning on past frames increases the number of input frames, and thus also the computational cost of training.

Conditions for generative models are usually formulated as a fixed window of previous frames. However, FDM [27] uses a per-frame UNet [55] and attention to take a variable number of frames as input. Each input frame can be set as a conditioning input or as a prediction target, and hence this model can be conditioned on frames arbitrarily far in the past and can even predict multiple frames at the same time. More recently 3DiM [70] introduced the idea of “implicit” conditioning in the task of 3D multi-view reconstruction. The idea is that at each step of the image generation in the diffusion process, the denoising network is conditioned only on a random view, instead of all the views. In this way the conditioning on multiple frames can be distributed over the denoising steps. In RIVER we extend this idea to the case of videos, where instead of views we use past frames.

Recently, [42] introduced conditional flow matching. They showed that it generalizes diffusion models, and it achieves faster training convergence and better results than other denoising diffusion models. Thus, we adopt this framework in our generative model. There are many ways to accelerate or improve vanilla diffusion models [37, 34, 16]. Among all, CCDF [11] is the most relevant to our warm-start sampling scheme. CCDF starts the backward denoising process from some time t other than T (which is the final time of the forward diffusion process) using an initial guess of the final output (e.g., a low-resolution sample). In RIVER, we develop an analogous technique within the formulation of flow matching for video generation.

3. Method

Let $\mathbf{x} = \{x^1, \dots, x^m\}$, where $x^i \in \mathbb{R}^{3 \times H \times W}$, be a video consisting of m RGB images. The task of video prediction is to forecast the upcoming n frames of a video given the first k frames, where $m = n + k$. Thus, it requires mod-

elling the following distribution:

$$p(x^{k+1}, \dots, x^{k+n} | x^1, \dots, x^k) = \prod_{i=1}^n p(x^{k+i} | x^1, \dots, x^{k+i-1}). \quad (1)$$

The decomposition in eq. (1) suggests an autoregressive sampling of the future frames. However, explicitly conditioning the next frame on all the past frames is computationally and memory-wise demanding. In order to overcome this issue, prior work suggests to use a recurrently updated memory variable [68, 50, 65, 10] or to restrict the conditioning window to a fixed number of frames [71, 67, 33, 76]. We instead propose to model each one-step predictive conditional distribution as a denoising probability density path that starts from a standard normal distribution. Moreover, rather than using score-based diffusion models [59] to fit those paths, we choose flow matching [42], a simpler method to train generative models. We further leverage the iterative nature of sampling from the learned flow and use a single random conditioning frame from the past at each iteration. This results in a simple and efficient training. An idea similar to ours was first introduced in [70] for novel view synthesis in 3D applications. In this paper, however, we made some design choices to adapt it to videos.

3.1. Latent Image Compression

Although we could operate directly on the pixels of the frames x^i , we introduce a compression step that reduces the dimensionality of the data samples and thus the overall numerical complexity of our approach. Given a dataset of videos D , we train a VQGAN [19] on single frames from that dataset. The VQGAN consists of an encoder \mathcal{E} and a decoder \mathcal{D} and allows to learn a perceptually rich latent codebook through a vector quantization bottleneck and an adversarial reconstruction loss [63]. A trained VQGAN is then used to compress the images to much lower resolution feature maps. That is, $z = \mathcal{E}(x) \in \mathbb{R}^{c \times \frac{H}{f} \times \frac{W}{f}}$, where $x \in \mathbb{R}^{3 \times H \times W}$. Commonly used values for c are 4 or 8 and for f are 8 or 16, which means that a 256×256 image can be downsampled to up to a 16×16 grid. Following [54], we let the decoder \mathcal{D} absorb the quantization layer and work in the pre-quantized latent space. Further in the paper, when referring to video frames we always assume that they are encoded in the latent space of a pretrained VQGAN.

3.2. Flow Matching

Flow matching was introduced in [42] as a simpler albeit more general and more efficient alternative to diffusion models [29]. A similar framework incorporating straight flows has also been proposed independently in [43, 2]. We assume that we are given samples from an unknown data distribution $q(z)$. In our case, the data sample z is the

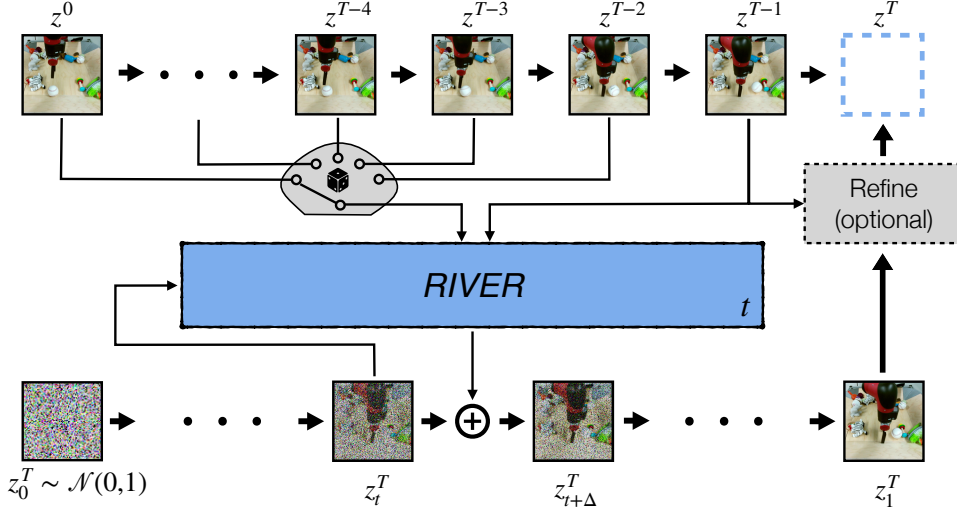


Figure 2. Inference with RIVER. In order to generate the next frame z^T (top-right), we sample an initial estimate from the standard normal distribution z_0^T (bottom-left) and integrate the ODE (2) by querying our model at each step with a random conditioning frame from the past z^c and previous frame z^{T-1} (top). We omitted the encoding/decoding for simplicity.

encoding of a video frame x via VQGAN. The aim of flow matching is to learn a temporal vector field $v_t(z) : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, with $t \in [0, 1]$, such that the following ordinary differential equation (ODE)

$$\dot{\phi}_t(z) = v_t(\phi_t(z)) \quad (2)$$

$$\phi_0(z) = z \quad (3)$$

defines a flow $\phi_t(z) : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ that pushes $p_0(z) = \mathcal{N}(z | 0, 1)$ towards some distribution $p_1(z) \approx q(z)$ along some probability density path $p_t(z)$. That is, $p_t = [\phi_t]_* p_0$, where $[\cdot]_*$ denotes the push-forward operation. If one were given a predefined probability density path $p_t(z)$ and the corresponding vector field $u_t(z)$, then one could parameterize $v_t(z)$ with a neural network and solve

$$\min_{v_t} \mathbb{E}_{t, p_t(z)} \|v_t(z) - u_t(z)\|^2. \quad (4)$$

However, this would be unfeasible in the general case, because typically we do not have access to $u_t(z)$. Lipman *et al.* [42] suggest to instead define a conditional flow $p_t(z | z_1)$ and the corresponding conditional vector field $u_t(z | z_1)$ per sample z_1 in the dataset and solve

$$\min_{v_t} \mathbb{E}_{t, p_t(z | z_1), q(z_1)} \|v_t(z) - u_t(z | z_1)\|^2. \quad (5)$$

This formulation enjoys two remarkable properties: 1) all the quantities can be defined explicitly; 2) Lipman *et al.* [42] show that solving eq. (5) is guaranteed to converge to the same result as in eq. (4). The conditional flow can be explicitly defined such that all intermediate distributions are Gaussian. Moreover, Lipman *et al.* [42] show that a linear transformation of the Gaussians' parameters yields

the best results in terms of convergence and sample quality. They define $p_t(z | z_1) = \mathcal{N}(z | \mu_t(z_1), \sigma_t^2(z_1))$, with $\mu_t(x) = tx_1$ and $\sigma_t(x) = 1 - (1 - \sigma_{\min})t$. With these choices, the corresponding target vector field is given by

$$u_t(z | z_1) = \frac{z_1 - (1 - \sigma_{\min})z}{1 - (1 - \sigma_{\min})t}. \quad (6)$$

Sampling from the learned model can be obtained by first sampling $z_0 \sim \mathcal{N}(z | 0, 1)$ and then numerically solving eq. (2) for $z_1 = \phi_1(z_0)$.

3.3. Video Prediction

We introduce the main steps to train and use RIVER. First, as described in sec. 3.1 we use a per-frame perceptual autoencoder to reduce the dimensionality of data. Since the encoding is per-frame and thus the reconstruction error could be temporally inconsistent, we improve the quality of a generated video by also introducing an optional small autoregressive refinement step in the decoding network. Second, we train a denoising model via flow matching in the space of encoded frames with our distributed conditioning. Moreover, we accelerate the video generation by introducing a warm-start sampling procedure.

Training. We adapt Flow Matching [42] to video prediction by letting the learned vector field v_t condition on the past context frames. Furthermore, we randomize the conditioning at each denoising step to only 2 frames. This results in a very simple training procedure, which is described in Algorithm 1. Given a training video $\mathbf{z} = \{z^1, \dots, z^m\}$ (pre-encoded with VQGAN), we randomly sample a target frame z^T and a random (diffusion) timestep $t \sim U[0, 1]$. We can then draw a sample from the conditional probability

Algorithm 1 Video Flow Matching with RIVER

Input: dataset of videos D , number of iterations N
for i in range(1, N) **do**
 Sample a video \mathbf{x} from the dataset D
 Encode it with a pre-trained VQGAN to obtain z
 Choose a random target frame $z^\tau, \tau \in \{3, \dots, |\mathbf{x}|\}$
 Sample a timestamp $t \sim U[0, 1]$
 Sample a noisy observation $z \sim p_t(z | z^\tau)$
 Calculate $u_t(z | z^\tau)$
 Sample a condition frame $z^c, c \in \{1, \dots, \tau - 2\}$
 Update the parameters θ of v_t via gradient descent

$$\nabla_{\theta} \|v_t(z | z^{\tau-1}, z^c, \tau - c; \theta) - u_t(z | z^\tau)\|^2 \quad (7)$$

end for

distribution $z \sim p_t(z | z^\tau)$ and calculate the target vector field $u_t(z | z^\tau)$ using eq. (6). We then sample another index c uniformly from $\{1, \dots, \tau - 2\}$ and use z^c , which we call *context frame*, together with $z^{\tau-1}$, which we call *reference frame*, as the two conditioning frames. Later, we show that the use of the reference is crucial for the network to learn the scene motion, since one context frame carries very little information about such motion. The vector field regressor v_t is then trained to minimize the following objective

$$\mathcal{L}_{\text{FM}}(\theta) = \|v_t(z | z^{\tau-1}, z^c, \tau - c; \theta) - u_t(z | z^\tau)\|^2, \quad (8)$$

where θ are the parameters of the model. Note that at no point during the training the whole video sequence must be stored or processed. Moreover, the generation of frames is never needed, which further simplifies the training process.

Inference. At inference time, in order to generate the T -th frame, we start from sampling an initial estimate z_0^T from the standard normal distribution (see Figure 2). We then use an ODE solver to integrate the learned vector field along the time interval $[0, 1]$. At each integration step, the ODE solver queries the network for $v_t(z_t^T | z^{T-1}, z^c, T - c)$, where $c \sim U\{1, \dots, T - 2\}$. In the simplest case, the Euler step of the ODE integration takes the form

$$z_{t_{i+1}}^T = z_{t_i}^T + \frac{1}{N} v_{t_i}(z_{t_i}^T | z^{T-1}, z^{c_i}, T - c_i), \quad (9)$$

$$c_i \sim U\{1, \dots, T - 2\}, \quad (10)$$

$$z_{t_0}^T \sim \mathcal{N}(z | 0, 1), \quad (11)$$

$$t_i = \frac{i}{N}, i \in \{0, \dots, N - 1\}, \quad (12)$$

where N is the number of integration steps. We then use z_1^T as an estimate of z^T .

Refinement. When using a per-frame VQGAN [19], the autoencoded videos may not always be temporally consistent. To address this issue without incurring a significant




Figure 3. Higher values of s for warm-start sampling lead to faster sampling, but worse FVD. Interestingly, $s = 0.1$ acts like the truncation trick [46, 7] and slightly improves the FVD.

computational cost, we optionally utilize a refinement network that operates in the pixel space. This deep convolutional network, based on the architecture of RCAN [79], is trained using the previous frame and the decoded next frame to refine the second frame. We train the model using an L_2 and a perceptual loss by refining 16 consecutive frames independently and then by feeding all frames to a perceptual network (I3D [9] in our case). We train the refinement network separately after training the autoencoder.

Sampling Speed. A common issue of models based on denoising processes is the sampling speed, as the same denoising network is queried multiple times along the denoising path in order to generate an image. This is even more apparent for the video domain, where the generation speed scales with the number of frames to generate. Some video diffusion models [27, 67] overcome this issue by sampling multiple frames at a time. However, the price they have to pay is the inability to generate arbitrarily long videos. We instead leverage the temporal smoothness of videos, that is, the fact that subsequent frames in a video do not differ much. This allows us to use a noisy previous frame as the initial condition of the ODE instead of pure noise. More precisely, instead of starting the integration from $z_0 \sim \mathcal{N}(z | 0, 1)$, we start at $z'_s \sim p_s(z | z^{T-1})$, where $1 - s$ is the speed up factor. We call this technique *warm-start sampling*. Intuitively, larger s results in a lower variability in the future frames. Moreover, we found that starting closer to the end of the integration path reduces the magnitude of the motion in the generated videos, since the model is required to sample closer to the previous frame. Therefore, there is a tradeoff between the sampling speed and the quality of the samples. We further emphasize this tradeoff by computing the FVD [62] of the generated videos depending on the speed up factor $1 - s$ (see Figure 3).




Figure 4. Video prediction on the *KTH* dataset. In order to predict the future frames, the model conditions on the first 10 (context) frames. Of this sequence, only the last context frame is shown. By definition, a proper stochastic predictive model generates realistic predictions of future frames that do not necessarily match the GT data.

Method	FVD↓	PSNR↑	SSIM↑
<i>10→30</i>			
SRVP [22]	222	29.7	0.87
SLAMP [1]	228	29.4	0.87
MCVD [67]	323	27.5	0.84
RIVER (ours)	180	30.4	0.86
<i>10→40</i>			
MCVD [67]	276.7	26.4	0.81
GridKeypoints [24]	144.2	27.1	0.84
RIVER (ours)	170.5	29.0	0.82

Table 1. *KTH* dataset evaluation. The evaluation protocol is to predict the next 30/40 frames given the first 10 frames.

3.4. Implementation

A commonly leveraged architecture for flow matching and diffusion models is UNet [55]. However, we found that training UNet could be time demanding. Instead, we propose to model $v_t(z | z^{\tau-1}, z^c, \tau - c; \theta)$ with the recently introduced U-ViT [6]. U-ViT follows the standard ViT [17] architecture and adds several long skip-connections, like in UNet. This design choice allows U-ViT to achieve on par or better results than UNet on image generation benchmarks with score-based diffusion models.

The inputs to the network are HW/f^2 tokens constructed by concatenating z , $z^{\tau-1}$ and z^c in feature axis as well as one additional time embedding token t that makes the network time-dependent. We additionally add spatial position encodings to the image tokens and augment $z^{\tau-1}$ and z^c

with an encoded relative distance $\tau - c$ to let the network know how far in the past the condition is. That is, the overall input to the network is of size $[HW/f^2 + 1, 3 \times d]$, where the first dimension refers to the number of tokens, while the second refers to the number of channels. For further details, see the supplementary material.

4. Experiments

In section 4.1, we report our results on several video prediction benchmarks. We evaluate our method using standard metrics, such as FVD [62], PSNR and SSIM [69]. We additionally show in section 4.2 that our model is able to perform visual planning. Video generation is demonstrated in section 4.3. Note that if not explicitly specified, we use the model without the refinement stage and with $s = 0$ in warm-start sampling. For additional results and training details, see the supplementary material.

4.1. Conditional Video Prediction

We test our method on 2 datasets. First, to assess the ability of RIVER to generate structured human motion, we test it on the **KTH** dataset [56]. KTH is a dataset containing 6 different human actions performed by 25 subjects in different scenarios. We follow the standard evaluation protocol predicting 30/40 future frames conditioned on the first 10 at a 64×64 pixel resolution. The results are reported in Table 1. We show that RIVER achieves state of the art prediction quality compared to prior methods that do not use domain-specific help. For instance, [24] models the motion

of the keypoints, which works well for human-centric data, but does not apply to general video generation. Figure 4 shows qualitative results.

Additionally, in Table 2 we evaluate the capability of RIVER to model complex interactions on BAIR [18], which is a dataset containing around 44K clips of a robot arm pushing toys on a flat square table. For BAIR, we generate and refine 15 future frames conditioned on one initial frame at a 64×64 pixel resolution. Due to the high stochasticity of motion in the BAIR dataset, the standard evaluation protocol in [5] is to calculate the metrics by comparing 100×256 samples to 256 random test videos (*i.e.*, 100 generated videos for each test video, by starting from the same initial frame as the test example). Additionally, we report the compute (memory in GB and hours) needed to train the models. RIVER reaches a tradeoff between the FVD and the compute and generates smooth realistic videos while requiring much less computational effort (see also Figure 1). In addition, we calculate the FVD vs the autoencoded test set, as we find that FVD (like FID [51]) can be affected even by different interpolation techniques. This way we eliminate the influence of potential autoencoding artifacts on the metrics in order to assess the consistency of the motion only. In fact, there is an improvement of about 30% in the FVD. Furthermore, although the standard benchmark on BAIR uses 64×64 pixels resolution, with the help of the perceptual compression, we are able to generate higher-resolution videos under the same training costs. See Figure 5 for qualitative results on the BAIR dataset at 256×256 resolution. Finally, we would like to point out that we observed DDPM fail to converge on BAIR, which further justifies our choice of flow matching (see also the supplementary material).

4.2. Visual Planning

One way to show the ability of the model to learn the dynamics of the environment is to do planning [21, 20, 74]. With a small change to the training of our model, RIVER is able to infill the video frames given the source and the target images. The only change to be done to the model is to remove the reference frame and to let two condition frames be sampled from both the future frames and the past ones. At inference time, given the source and the target frames, our model sequentially infills the frames between those. We show in Figure 6 some qualitative results of video interpolation on the CLEVRER [77] dataset, which is a dataset containing 10K training clips capturing a synthetic scene with multiple objects interacting with each other through collisions. It is a dataset suitable for planning, as it allows to show the ability of the method to model the dynamics of the separate objects and their interactions. We test our model at the 128×128 pixels resolution. Note how the model has learned the interactions between the objects and is able to manipulate the objects in order to achieve the given goals.

Method	FVD↓	Mem (GB)	Hours
TriVD-GAN-FP [44]	103.0	1024	280
Video Transformer [71] (L)	94.0	512	336
CCVS [38] (low res)	99.0	128	40
CCVS [38] (high res)	80.0	-	-
LVT [52] ($n_c = 4$)	125.8	128	48
FitVid [5]	93.6	1024	288
MaskViT [26]	93.7	-	-
MCVD [67] (concat)	98.8	77	78
MCVD [67] (spatin)	103.8	86	50
NÜWA [73]	86.9	2560	336
RaMViD [33]	84.2	320	72
VDM [31]	66.9	-	-
RIVER <i>w/ refine</i>	106.1	25	25
RIVER <i>w/o refine</i>	145.8	-	-
RIVER <i>w/o refine vs ae GT</i>	73.5	-	-

Table 2. BAIR dataset evaluation. We follow the standard evaluation protocol, which is to predict 15 future frames given 1 initial frame. The common way to compute the FVD is to compare 100×256 generated sequences to 256 randomly sampled test videos. Additionally, we report the numbers of the network without the refinement stage versus the original ground truth (RIVER *w/o refine*) and the autoencoded ground truth (RIVER *w/o refine vs ae GT*) to highlight the influence of the VQGAN’s artifacts on the assessment of the motion consistency.

Method	FVD↓	PSNR↑
w/ reference	94.38	30.53
w/o reference	217.13	26.95

Table 3. Ablations on the use of the reference frame. We generate 14 frames given 2 initial ones and the metrics are calculated on 256 test videos with 1 sample per video and 10 integration steps per frame. All models are trained for 80K iterations.

4.3. Video Generation

RIVER can be easily adapted to support *video generation*. Inspired by the classifier-free guidance [30] we train a single model to both generate (the first frame of a video) and predict the next frames by simply feeding noise instead of the condition frames 10% of the times during training. Then, during inference we generate the first frame and then predict the rest of the video given the first frame. Figure 7 shows our results for video generation on CLEVRER [77] (FVD = 23.63). Other methods [48, 78, 58] have difficulties in modeling the motions and interactions of objects. For videos and qualitative comparisons, visit our website¹.

4.4. Ablations

In this section, we ablate several design choices in order to illustrate their impact on the performance of RIVER.

First, we ablate the importance of using the reference

¹<https://araachie.github.io/river>




Figure 5. Video prediction on the *BAIR* dataset. The model predicts future frames conditioned on a single initial frame. Thanks to VQGAN, RIVER can be used to generate high resolution videos.

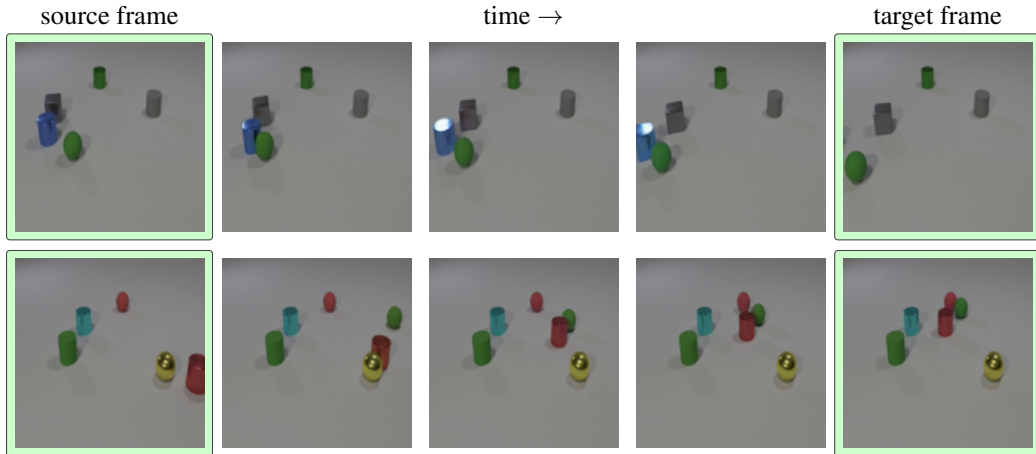


Figure 6. Visual planning with RIVER on the *CLEVRER* dataset. Given the source and the target frames, RIVER infills the frames inbetween. Note how the model manipulates the objects by forcing them to interact in order to achieve the goal. In some cases this even requires introducing new objects into the scene.

frame in the condition. In [70], where the stochastic conditioning was first introduced, only one view from the memory was used at each denoising step for generating a novel view. However, conditioning on one frame from the past does not work for video prediction, since one frame does not contain any information about pre-existing motion. We train a model, where we remove the reference frame from the condition and compare its performance to the full model. For this ablation we test RIVER on the *CLEVRER* [77] dataset. We found that without the reference frame in the condition the model is confused about the direction of the motion, which results in jumping objects (see Figure 8). For the quantitative results, check Table 3.

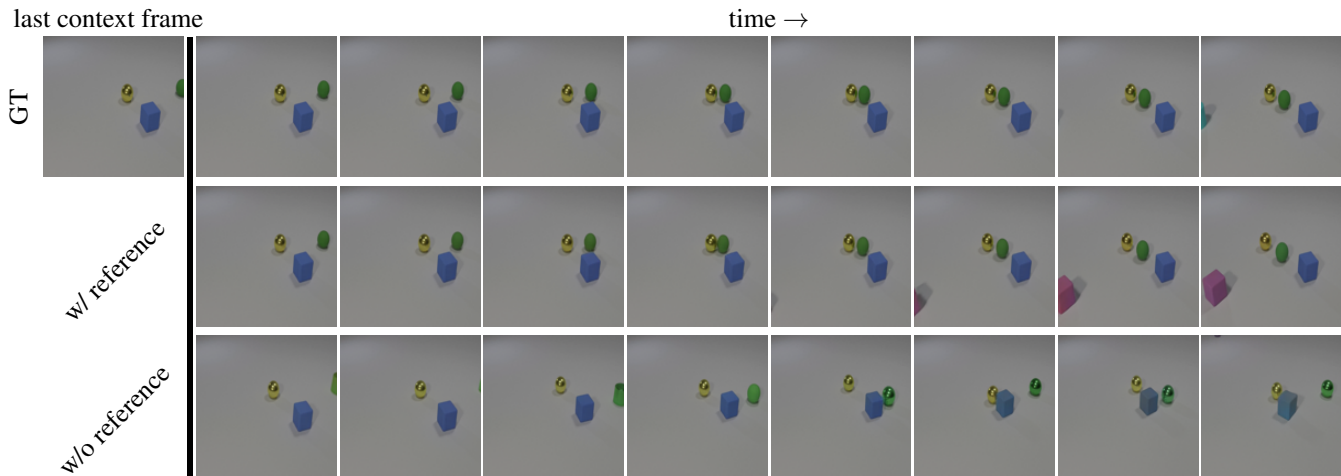
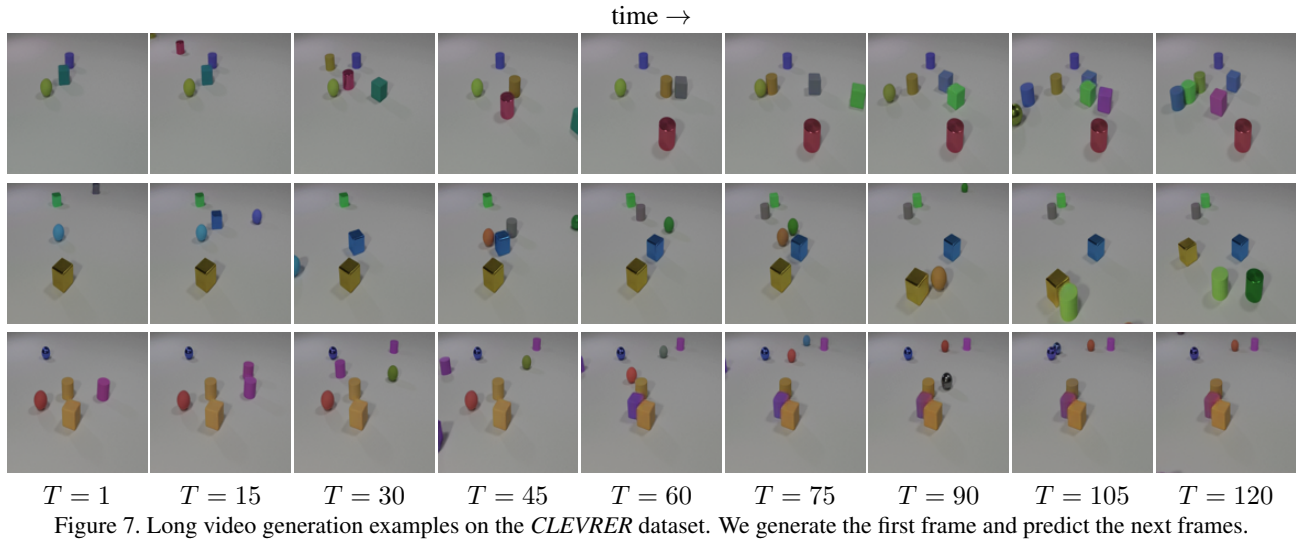
Given a model trained so that the context frames are sampled from the whole past of a sequence, at inference time we ablate the size of the past window used for the context frames to better understand the impact of the history on the video generation performance. In this ablation, we uniformly sample the context frames from $\{\tau - 1 - k, \dots, \tau - 2\}$ for $k = 2, 4, 6, 8$, and show which past frames better support RIVER’s predictions. For this experiment we use our trained model on the *BAIR* [18]

Context	BAIR / PSNR \uparrow	KTH / PSNR \uparrow
2 frames	25.64	28.53
4 frames	25.94	29.07
6 frames	26.00	30.17
8 frames	25.28	29.40

Table 4. Ablations on the context size. Using a pretrained model on *BAIR* [18] and *KTH* [56] we observe a trade-off wrt the number of conditioning frames. We believe that datasets with more challenging scenes and dynamics may require more context frames.

and *KTH* [56] datasets. Since there are occlusions in *BAIR*, we suspect that having more context can help to predict the future frames more accurately. Having more context frames also helps to predict a smoother motion for humans in *KTH*. Table 4 shows that there is a trade-off in context size and although having more context can be useful, on simple datasets having only a few frames is better to solve the prediction task.

Finally we show in Figure 3 that warm-start sampling can be used to generate samples faster (with fewer integration steps) but with a cost on quality. Interestingly we ob-



served that a small speed up factor actually helps the sampling and despite having fewer integration steps leads to better performance. We suspect that this effect is similar to the truncation trick [46, 7] in GANs. Notice however, that compared to other diffusion-based video generation approaches, RIVER conditions only on 2 past frames for a single neural function evaluation (NFE). Hence, a single NFE is generally less expensive. For instance, it takes 9.97 seconds for RIVER to generate 16 frames video, while RaMViD [33] requires 40.47 seconds with a vanilla scheduler on a single Nvidia GeForce RTX 3090 GPU (on BAIR with 64×64 resolution). For more results, see the supplementary material.

5. Conclusion

In this paper we have introduced RIVER, a novel training procedure and a model for video prediction that are based

on the recently proposed Flow Matching for image synthesis. We have adapted the latter to videos and incorporated conditioning on an arbitrarily large past frames window through randomly sampling a new context frame at each integration step of the learned flow. Moreover, working in the latent space of a pretrained VQGAN enabled the generation of high-resolution videos. All these have resulted in a simple and effective training procedure, which we hope future works on video synthesis will largely benefit from. We have tested RIVER on several video datasets and found that it is not only able to predict high-quality videos, but is also flexible enough to be trained to perform other tasks, such as visual planning and video generation.

Acknowledgements. This work was supported by grant 188690 of the Swiss National Science Foundation.

References

- [1] Adil Kaan Akan, Erkut Erdem, Aykut Erdem, and Fatma Guney. Slamp: Stochastic latent appearance and motion prediction. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14708–14717, 2021. [6](#)
- [2] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022. [3](#)
- [3] Nuha Aldausari, Arcot Sowmya, Nadine Marcus, and Gelareh Mohammadi. Video generative adversarial networks: A review. *ACM Computing Surveys (CSUR)*, 55:1–25, 2020. [2](#)
- [4] Mohammad Babaeizadeh, Chelsea Finn, D. Erhan, Roy H. Campbell, and Sergey Levine. Stochastic variational video prediction. *ArXiv*, abs/1710.11252, 2018. [2](#)
- [5] Mohammad Babaeizadeh, Mohammad Taghi Saffar, Suraj Nair, Sergey Levine, Chelsea Finn, and Dumitru Erhan. Fitvid: Overfitting in pixel-level video prediction. *arXiv preprint arXiv:2106.13195*, 2021. [1](#), [2](#), [7](#)
- [6] Fan Bao, Chongxuan Li, Yue Cao, and Jun Zhu. All are worth words: a vit backbone for score-based diffusion models. *arXiv preprint arXiv:2209.12152*, 2022. [6](#)
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *ArXiv*, abs/1809.11096, 2019. [5](#), [9](#)
- [8] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020. [2](#)
- [9] João Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733, 2017. [5](#)
- [10] Lluís Castrejon, Nicolas Ballas, and Aaron Courville. Improved conditional vrnn for video prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7608–7617, 2019. [3](#)
- [11] Hyungjin Chung, Byeongsu Sim, and Jong-Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12403–12412, 2022. [3](#)
- [12] Aidan Clark, Jeff Donahue, and Karen Simonyan. Adversarial video generation on complex datasets. *arXiv: Computer Vision and Pattern Recognition*, 2019. [2](#)
- [13] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *International conference on machine learning*, pages 1174–1183. PMLR, 2018. [1](#), [2](#)
- [14] Emily L. Denton and Vighnesh Birodkar. Unsupervised learning of disentangled representations from video. *ArXiv*, abs/1705.10915, 2017. [2](#)
- [15] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. *ArXiv*, abs/2105.05233, 2021. [3](#)
- [16] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Genie: Higher-order denoising diffusion solvers. *ArXiv*, abs/2210.05475, 2022. [3](#)
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929, 2021. [6](#)
- [18] Frederik Ebert, Chelsea Finn, Alex X. Lee, and Sergey Levine. Self-supervised visual planning with temporal skip connections. In *CoRL*, 2017. [7](#), [8](#)
- [19] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. [2](#), [3](#), [5](#)
- [20] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction. *Advances in neural information processing systems*, 29, 2016. [1](#), [7](#)
- [21] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2786–2793. IEEE, 2017. [1](#), [7](#)
- [22] Jean-Yves Franceschi, Edouard Delasalles, Mickaël Chen, Sylvain Lamprier, and Patrick Gallinari. Stochastic latent residual video prediction. In *ICML*, 2020. [6](#)
- [23] Xiaojie Gao, Yueming Jin, Qi Dou, Chi-Wing Fu, and Pheng-Ann Heng. Accurate grid keypoint learning for efficient video prediction. *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5908–5915, 2021. [2](#)
- [24] Xiaojie Gao, Yueming Jin, Qi Dou, Chi-Wing Fu, and Pheng-Ann Heng. Accurate grid keypoint learning for efficient video prediction. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5908–5915. IEEE, 2021. [6](#)
- [25] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. [2](#)
- [26] Agrim Gupta, Stephen Tian, Yunzhi Zhang, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. Maskvit: Masked visual pre-training for video prediction. *arXiv preprint arXiv:2206.11894*, 2022. [2](#), [7](#)
- [27] William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weillbach, and Frank Wood. Flexible diffusion modeling of long videos. *arXiv preprint arXiv:2205.11495*, 2022. [1](#), [2](#), [3](#), [5](#)
- [28] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. [3](#)
- [29] Jonathan Ho, Ajay Jain, and P. Abbeel. Denoising diffusion probabilistic models. *ArXiv*, abs/2006.11239, 2020. [3](#)

- [30] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 7
- [31] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022. 3, 7
- [32] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, nov 1997. 2
- [33] Tobias Höppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen, and Andrea Dittadi. Diffusion models for video prediction and infilling. *arXiv preprint arXiv:2206.07696*, 2022. 1, 3, 7, 9
- [34] Alexia Jolicœur-Martineau, Ke Li, Remi Piche-Taillefer, Tal Kachman, and Ioannis Mitliagkas. Gotta go fast when generating data with score-based models. *ArXiv*, abs/2105.14080, 2021. 3
- [35] Yunji Kim, Seonghyeon Nam, I. Cho, and Seon Joo Kim. Unsupervised keypoint learning for guiding class-conditional video prediction. *ArXiv*, abs/1910.02027, 2019. 2
- [36] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2014. 2
- [37] Zhifeng Kong and Wei Ping. On fast sampling of diffusion probabilistic models. *ArXiv*, abs/2106.00132, 2021. 3
- [38] Guillaume Le Moing, Jean Ponce, and Cordelia Schmid. Ccvs: Context-aware controllable video synthesis. *Advances in Neural Information Processing Systems*, 34:14042–14055, 2021. 2, 7
- [39] 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
- [40] Wonkwang Lee, Whie Jung, Han Zhang, Ting Chen, Jing Yu Koh, Thomas E. Huang, Hyungsuk Yoon, Honglak Lee, and Seunghoon Hong. Revisiting hierarchical approach for persistent long-term video prediction. *ArXiv*, abs/2104.06697, 2021. 2
- [41] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P. Xing. Dual motion gan for future-flow embedded video prediction. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1762–1770, 2017. 2
- [42] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 2, 3, 4
- [43] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 3
- [44] Pauline Luc, Aidan Clark, Sander Dieleman, Diego de Las Casas, Yotam Doron, Albin Cassirer, and Karen Simonyan. Transformation-based adversarial video prediction on large-scale data. *arXiv preprint arXiv:2003.04035*, 2020. 7
- [45] Pauline Luc, Aidan Clark, Sander Dieleman, Diego de Las Casas, Yotam Doron, Albin Cassirer, and Karen Simonyan. Transformation-based adversarial video prediction on large-scale data. *ArXiv*, abs/2003.04035, 2020. 2
- [46] Marco Marchesi. Megapixel size image creation using generative adversarial networks. *ArXiv*, abs/1706.00082, 2017. 5, 9
- [47] Michaël Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *CoRR*, abs/1511.05440, 2016. 2
- [48] Kangfu Mei and Vishal M. Patel. Vidm: Video implicit diffusion models. *ArXiv*, abs/2212.00235, 2022. 7
- [49] Matthias Minderer, Chen Sun, Ruben Villegas, Forrester Cole, Kevin P. Murphy, and Honglak Lee. Unsupervised learning of object structure and dynamics from videos. *ArXiv*, abs/1906.07889, 2019. 2
- [50] Marc Oliu, Javier Selva, and Sergio Escalera. Folded recurrent neural networks for future video prediction. In *ECCV*, 2018. 3
- [51] Gaurav Parmar, Richard Zhang, and Junyan Zhu. On aliased resizing and surprising subtleties in gan evaluation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11400–11410, 2022. 7
- [52] Ruslan Rakhimov, Denis Volkhonskiy, Alexey Artemov, Denis Zorin, and Evgeny Burnaev. Latent video transformer. *arXiv preprint arXiv:2006.10704*, 2020. 2, 7
- [53] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv*, abs/2204.06125, 2022. 3
- [54] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2, 3
- [55] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *ArXiv*, abs/1505.04597, 2015. 3, 6
- [56] Christian Schuldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: a local svm approach. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 3, pages 32–36. IEEE, 2004. 6, 8
- [57] Younggyo Seo, Kimin Lee, Fangchen Liu, Stephen James, and P. Abbeel. Harp: Autoregressive latent video prediction with high-fidelity image generator. *ArXiv*, abs/2209.07143, 2022. 2
- [58] Ivan Skorokhodov, S. Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3616–3626, 2021. 7
- [59] Yang Song, Jascha Narain Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *ArXiv*, abs/2011.13456, 2021. 3
- [60] Ximeng Sun, Huijuan Xu, and Kate Saenko. A two-stream variational adversarial network for video generation. *ArXiv*, abs/1812.01037, 2018. 2
- [61] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csd: Conditional score-based diffusion models for probabilistic time series imputation. In *NeurIPS*, 2021. 3
- [62] 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. 5, 6

- [63] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *NIPS*, 2017. 2, 3
- [64] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 2
- [65] Ruben Villegas, Arkanath Pathak, Harini Kannan, D. Erhan, Quoc V. Le, and Honglak Lee. High fidelity video prediction with large stochastic recurrent neural networks. *ArXiv*, abs/1911.01655, 2019. 3
- [66] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. *ArXiv*, abs/1706.08033, 2017. 2
- [67] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. Mcvd: Masked conditional video diffusion for prediction, generation, and interpolation. *arXiv preprint arXiv:2205.09853*, 2022. 1, 3, 5, 6, 7
- [68] Yunbo Wang, Zhifeng Gao, Mingsheng Long, Jianmin Wang, and S Yu Philip. Predrnn++: Towards a resolution of the deep-in-time dilemma in spatiotemporal predictive learning. In *International Conference on Machine Learning*, pages 5123–5132. PMLR, 2018. 3
- [69] Zhou Wang, Alan Conrad Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13:600–612, 2004. 6
- [70] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. *arXiv preprint arXiv:2210.04628*, 2022. 2, 3, 8
- [71] Dirk Weissenborn, Oscar Täckström, and Jakob Uszkoreit. Scaling autoregressive video models. *arXiv preprint arXiv:1906.02634*, 2019. 3, 7
- [72] Nevan Wichers, Ruben Villegas, D. Erhan, and Honglak Lee. Hierarchical long-term video prediction without supervision. *ArXiv*, abs/1806.04768, 2018. 2
- [73] Chenfei Wu, Jian Liang, Lei Ji, F. Yang, Yuejian Fang, Daxin Jiang, and Nan Duan. Nüwa: Visual synthesis pre-training for neural visual world creation. In *ECCV*, 2022. 7
- [74] Annie Xie, Dylan Losey, Ryan Tolsma, Chelsea Finn, and Dorsa Sadigh. Learning latent representations to influence multi-agent interaction. In *Conference on robot learning*, pages 575–588. PMLR, 2021. 1, 7
- [75] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 2
- [76] Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. Diffusion probabilistic modeling for video generation. *arXiv preprint arXiv:2203.09481*, 2022. 1, 3
- [77] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B. Tenenbaum. Clevrer: Collision events for video representation and reasoning. *ArXiv*, abs/1910.01442, 2020. 1, 7, 8
- [78] Sihyun Yu, Jihoon Tack, Sangwoo Mo, Hyunsu Kim, Junho Kim, Jung-Woo Ha, and Jinwoo Shin. Generating videos with dynamics-aware implicit generative adversarial networks. *ArXiv*, abs/2202.10571, 2022. 7
- [79] Yulun Zhang, Kungpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Raymond Fu. Image super-resolution using very deep residual channel attention networks. In *European Conference on Computer Vision*, 2018. 5