

Playable Video Generation

Willi Menapace

University of Trento

willi.menapace@unitn.it

Stéphane Lathuilière*

LTCI, Télécom Paris

Institut Polytechnique de Paris

stephane.lathuiliere@telecom-paris.fr

Sergey Tulyakov*

Snap Inc.

stulyakov@snap.com

Aliaksandr Siarohin

University of Trento

aliaksandr.siarohin@unitn.it

Elisa Ricci

University of Trento

Fondazione Bruno Kessler

e.ricci@unitn.it

Abstract

This paper introduces the unsupervised learning problem of playable video generation (PVG). In PVG, we aim at allowing a user to control the generated video by selecting a discrete action at every time step as when playing a video game. The difficulty of the task lies both in learning semantically consistent actions and in generating realistic videos conditioned on the user input. We propose a novel framework for PVG that is trained in a self-supervised manner on a large dataset of unlabelled videos. We employ an encoder-decoder architecture where the predicted action labels act as bottleneck. The network is constrained to learn a rich action space using, as main driving loss, a reconstruction loss on the generated video. We demonstrate the effectiveness of the proposed approach on several datasets with wide environment variety. Further details, code and examples are available on our project page willi-menapace.github.io/playable-video-generation-website.

1. Introduction

Humans at a very early age can identify key objects and how each object can interact with its environment. This ability is particularly notable when watching videos of sports or games. In tennis and football, for example, the skill is taken to the extreme. Spectators and sportscasters often argue which action or movement the player should have performed in the field. We can understand and anticipate actions in videos despite never being given an explicit list of plausible actions. We develop this skill in an unsupervised manner as we see actions live and on the screen. We can further analyze the technique with which an action is performed as well as the “amount” of action, i.e. how much to the left. Furthermore, we can reason about what

*The second and third authors contributed equally to the work.

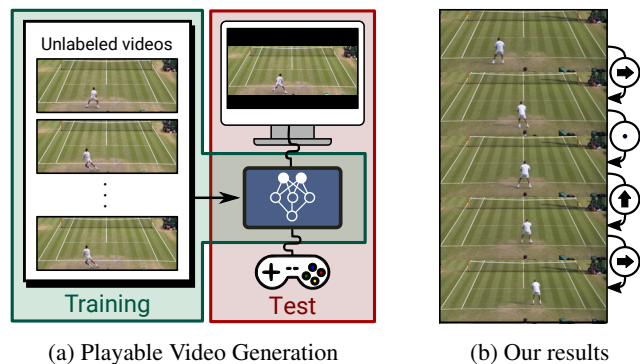


Figure 1: We introduce the task of playable video generation in an unsupervised setting (left). Given a set of unlabeled video sequences, a set of discrete actions are learned in order to condition video generation. At test time, using our method, named CADDY, the user can control the generated video on-the-fly providing action labels.

happens if the player took a different action and how this would change the video.

From this observation, we propose a new task, Playable Video Generation (PVG) illustrated in Fig 1a. In PVG, the goal is to learn a set of distinct actions from real-world video clips in an unsupervised manner (green block) in order to offer the user the possibility to interactively generate new videos (red block). As shown in Fig 1b, at test time, the user provides a discrete action label at every time step and can see live its impact on the generated video, similarly to video games. Introducing this novel problem paves the way toward methods that can automatically simulate real-world environments and provide a gaming-like experience.

PVG is related to the future frame prediction problem [13, 25, 28, 40, 42, 44] and in particular to methods that condition future frames on action labels [20, 30, 31]. Given one or few initial frames and the labels of the performed actions, such systems aim at predicting what happens next in

the video. For example, this framework can be used to imitate a desired video game using a neural network with a remarkable generation quality [20, 31]. However, at training time, these methods require videos with their corresponding frame-level action at every time step. Consequently, these methods are limited to video game environments [20, 31] or robotic data [30] and cannot be employed in real-world environments. As an alternative, the annotation effort required to address real-world videos can be reduced using a single action label to condition the whole video [45], but it limits interactivity since the user cannot control the generated video on-the-fly. Conversely, in PVG, the user can control the generation process by providing an action at every time-step after observing the last generated frame.

This paper addresses these limitations introducing a novel framework for PVG named *Clustering for Action Decomposition and DiscoverY* (CADDY). Our approach discovers a set of distinct actions after watching multiple videos through a clustering procedure blended with the generation process that, at inference time, outputs playable videos. We adopt an encoder-decoder architecture where a discrete bottleneck layer is employed to obtain a discrete representation of the transitions between consecutive frames. A reconstruction loss is used as main driving loss avoiding the need for neither action label supervision nor even the precise number of actions. A major difficulty in PVG is that discrete action labels cannot capture the stochasticity typical of real-world videos. To address this difficulty, we introduce an action network that estimates the action label posterior distribution by decomposing actions in a discrete label and a continuous component. While the discrete action label captures the main semantics of the performed action, the continuous component captures how the action is performed. At test time, the user provides only the discrete actions to interact with the generation process.

Finally, we experimentally demonstrate that our approach can learn consistent actions in varied environments. We conduct experiments on three different datasets including both real-world (*i.e.* tennis and robotics [9]) and synthetic (*i.e.* video game [2]). Our experiments show that CADDY generates high-quality videos while offering the user a better playability experience than existing future frame prediction methods.

2. Related Works

Video generation. Recent advances in deep generative models have led to impressive progress in video generation. A variety of formulations have been explored including Generative Adversarial Networks (GANs) [35, 42, 44], variational auto-encoders (VAEs) [1] and auto-regressive models [47]. While early works addressed the unconditional case [35, 44], many conditional video generation tasks have been addressed. The generated video can be con-

ditioned on specific type of information. Among numerous examples, we can cite video to video translation [46], image animation [37, 38] or pose-based generation [4].

Another typical example is the problem of future frame prediction that consists in generating a video conditioned on the first video frames. Early works employ deterministic predictive models [10, 28, 44] that cannot handle the stochasticity inherent to real-world videos. Several methods have been proposed to incorporate stochasticity in the model using VAE formulations [13, 25, 42], GANs [24], or normalizing flows [23]. Inspired by VAE-based approaches, we adopt a probabilistic formulation that is compatible with our encoder-decoder pipeline. Nevertheless, our playable video generation framework goes beyond video prediction methods and allows the user to control generation of future frames with discrete actions.

To offer better control over the generation process than video prediction methods, some approaches condition the generated video on an input action label [21, 45]. However, these approaches require action annotations for training and the generated video cannot be controlled on-the-fly since a single label is used to generate the whole video. To further increase the level of control, other approaches condition each frame in the generated video on a different action label [6, 20, 30, 31]. In particular, Kim *et al.* recently proposed *GameGAN* [20], a framework that visually imitates a desired game. During training, *GameGAN* ingests a large collection of videos and the corresponding keyboard actions pressed by the player. The network is trained to predict the next frame from past frames and keyboard actions. While these methods allow playable video generation similarly to our approach, the requirement for frame-level annotation limits their application to constrained environments such as video games. Conversely, our approach does not use action label supervision and can be employed in real-world environments. In the context of future frame predictions, Rybkin *et al.* [34] propose to infer latent actions. However, this method learns a continuous action space that is later mapped to discrete actions with action supervision.

Few works have focused on the generation of videos with controllable characters in real-world environments [14, 49]. However, these approaches heavily rely on prior knowledge specific to the environment, such as pre-trained full-body pose estimators. On the contrary, in PVG, the goal is to design a general framework that can be applied to varied environments without modifications.

Deep clustering and unsupervised learning. Unsupervised learning has attracted growing attention with the raise of self-supervised learning techniques. Recent works address jointly clustering and representation learning so that the network is trained to categorize the training data while learning deep representation. For instance, DeepCluster [3] alternates between k-mean clustering and a feature learn-

ing phase where cluster assignments are used as supervision. As an alternative to this iterative algorithm, other approaches [18, 29] propose the use of networks that directly output cluster assignment rather than feature representations. In these methods, the network is trained with an information-theoretic criterion formulated as a mutual information maximization problem. Differently from these approaches that focus only on image clustering, CADDY tackles the problem of learning actions in videos by blending clustering with a generation formulation.

Note that, recent works address the problem of unsupervised learning for videos [8, 27, 32, 39] but they are either limited to representation learning [8, 27] or perform clustering at the video-level [32, 39] while PVG requires action labels at every time step.

3. Method

3.1. Overall Pipeline

In this work, we propose a framework where a user can interactively control the video generation process by selecting an action at every time step among a set of K discrete actions. Our method, named CADDY, is trained on a dataset of unannotated videos. We only assume that the video sequences depict a single agent acting in an environment. No action labels are required.

Inspired by Reinforcement Learning (RL) literature [41], the object in the scene is modeled as an agent interacting with its environment by performing an action at every time step. Differently from RL, our goal is to jointly learn the action space, the state representation that describes the state of the agent and its environment, and a decoder that reconstructs the observations (*i.e.* frames) from the state.

CADDY is articulated into four main modules illustrated in Fig. 2: (i) an encoder employs a network E to extract frame representations; (ii) a temporal model estimates the label corresponding to the action performed in the current frame and predicts a state s_{t+1} that describes the environment at the next time step, after performing the detected action. The action label is predicted via a network A that receives the frame representations from the current and next frames. To predict the next frame environment state s_{t+1} , we employ a recurrent neural network R that we refer to as *Dynamics network*. (iii) a decoder module employs a network D to reconstruct each frame from the frame embedding predicted by the temporal model. (iv) the reconstructed frames are fed to the encoder to assess the quality of the estimated action labels.

The overall pipeline is trained in an end-to-end fashion using as main driving loss a reconstruction loss on the output frames. The key idea of our approach is that the action network A needs to predict consistent action labels in order to correctly estimate the next frame embeddings s_{t+1} ,

and then, accurately reconstruct the input frames. We now describe each network.

Video encoder. Let us consider a video sequence $\{x_t\}_{t=1}^T$ of length T . First, for every frame, we extract a feature representation f_t with $t \in \{1, \dots, T\}$ which we call *frame features*. The aim of this representation is to encode information about the appearance and the semantics of the input frame:

$$f_t = E(x_t). \quad (1)$$

Action network. Next, the action network A is used to infer the discrete action $a_t \in \{1..K\}$ performed between the frames x_t and x_{t+1} . The action label a_t characterizes the transition between the features f_t and its successor f_{t+1} . In a deterministic environment, the evolution of the environment is fully described by a_t . However, in real scenarios this assumption is rarely satisfied and agents acting in a complex environment may express a behavior that can only be partially described with discrete actions. To handle this, we propose to model the transition between two consecutive frame features as the combination of the discrete action a_t and a continuous component v_t referred to as *action variability embedding*, *i.e.*:

$$a_t, v_t = A(f_t, f_{t+1}). \quad (2)$$

While the discrete action a_t describes the action performed by the agent at a high level, the continuous component v_t describes the variability of each action and captures the possible non-determinism in the environment.

In order to handle this non-determinism, we propose a probabilistic auto-encoder formulation to train the action network. In this regard, details are given in Sec. 3.2.

Dynamics network. Given the features f_t , the action a_t , and the corresponding *action variability embedding* v_t , the dynamics network R is used to embed the dynamics of the environment observed in the input sequence. We model the dynamics network as a recurrent network based on convolutional LSTMs [36] with the *environment state* s_t :

$$s_{t+1} = R(s_t, f_t, a_t, v_t). \quad (3)$$

The induction in Eq. (3) is initialized with a parameter s_0 that we consider to be trainable. The action label a_t and the *action variability embedding* v_t are concatenated channel-wise with the input feature f_t in the dynamics network.

Decoder. We use a decoder network D to reconstruct the frames \hat{x}_{t+1} from the *environment state* s_{t+1} .

3.2. Probabilistic Action Network

As introduced earlier, the goal of the action network is to estimate the action label a_t and *action variability embedding* v_t . The action network first uses an *action state network* A_s to extract an *action embedding* e_t which represents the information in f_t that is related to the action currently performed. Following a probabilistic formulation, A_s

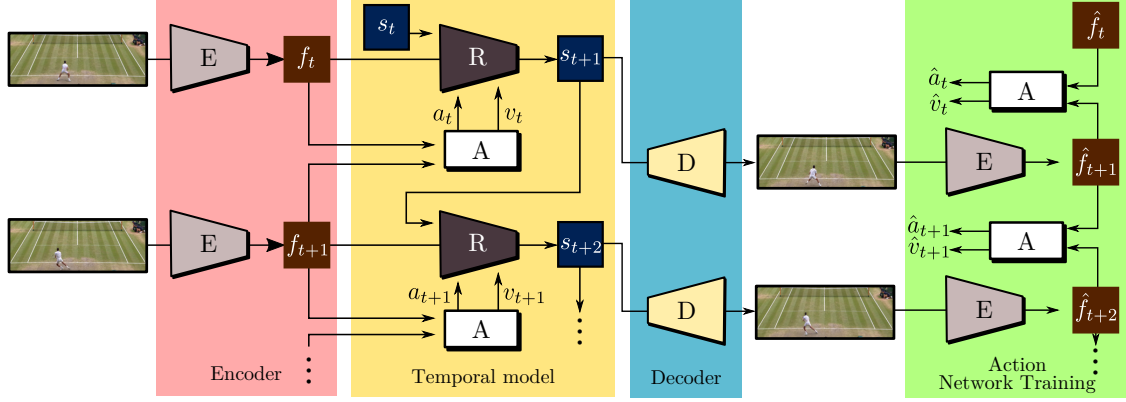


Figure 2: CADDY’s training procedure for unsupervised playable video generation. An encoder E extracts frame representations from the input sequence. A temporal model estimates the successive states using a recurrent dynamics network R and an action network A which predicts the action label corresponding to the current action performed in the input sequence. Finally, a decoder D reconstructs the input frames. The model is trained using reconstruction as the main driving loss.

predicts the posterior distribution of the *action embedding* modeled as Gaussian:

$$e_t \sim \mathcal{N}(\mu_{e_t}, \sigma_{e_t}^2) \quad (4)$$

$$\text{with } \mu_{e_t}, \sigma_{e_t}^2 = A_s(f_t). \quad (5)$$

Similarly, the distribution of the *action embedding* $e_{t+1} \sim \mathcal{N}(\mu_{e_{t+1}}, \sigma_{e_{t+1}}^2)$ is estimated from f_{t+1} . We can then combine the action states e_t and e_{t+1} to predict the performed action. More precisely, we propose to model the posterior distribution of the action (a_t, v_t) using the difference between the two *action embeddings* which still follows a Gaussian distribution. From the independence between random variables e_t and e_{t+1} we have:

$$d_t = e_{t+1} - e_t \sim \mathcal{N}(\mu_{d_t}, \sigma_{d_t}^2) \quad (6)$$

$$\text{with } \mu_{d_t} = \mu_{e_{t+1}} - \mu_{e_t} \text{ and } \sigma_{d_t}^2 = \sigma_{e_{t+1}}^2 + \sigma_{e_t}^2. \quad (7)$$

Finally, the *action direction* d_t is sampled according to (6) using the re-parametrization trick [22] and fed to a single-layer classifier to estimate the probability of each action $p_t = \{p_t^k\}_{k=1}^K \in [0, 1]^K$. Importantly, we implement the classification layer using a Gumbel-Softmax layer [17] to obtain the discrete label a_t from the probabilities p_t while preserving differentiability for Stochastic Gradient Descent (SGD).

Regarding the *action variability embedding* v_t , considering $v_t = d_t$ would not favour the model to learn meaningful action labels a_t since the changes in the environment could be directly encoded in d_t without using a_t . Thus, we propose to make v_t dependent on both d_t and a_t in a way that d_t cannot be recovered from v_t alone, enforcing the learning of action labels a_t . To this aim, we consider a set of K *action direction centroids* $\{c_k\}_{k=1}^K$ (one per action) that are defined as the expected *action directions* for each action. Practically, the cluster centroids are estimated using an exponential moving average over the action direction associated with each discrete action. We propose to define the

action variability embedding v_t as the difference between the observed action direction d_t and its assigned cluster centroid. Following a soft-assignment formulation, the *action variability embedding* is given by the expected difference with its action centroid:

$$v_t = \sum_{k=1}^K p_t^k (d_t - c_k). \quad (8)$$

In this way, v_t acts as a bottleneck for d_t since, assuming distinct centroids c_k , d_t cannot be completely encoded in v_t . To ensure distinct centroids, we introduce a loss that prevents their collapse to a single point (see Sec. 3.3).

3.3. Training objectives and test

CADDY is trained in an end-to-end fashion with a combination of objectives that aim at obtaining both high-quality output sequences and a discrete action space that captures the agent’s high-level actions.

Reconstruction losses. The main driving loss of our system is a frame reconstruction loss based on the perceptual loss of Johnson *et al.* [19]:

$$\mathcal{L}_{\text{rec}}^x = \frac{1}{T} \sum_{t=1}^T \sum_{j=1}^J \|N_j(x_t) - N_j(\hat{x}_t)\|_1, \quad (9)$$

where $N_j, j \in \{1, \dots, J\}$ indicates the j^{th} channel extracted from a specific VGG-19 layer and J is the number of feature channels in this layer. The reconstruction loss is averaged over different layers and resolutions [38]: D is equipped with multiple heads to output \hat{x}_t at different resolutions and x_t is down-sampled to form a pyramid. This loss is completed by a L_1 reconstruction loss (denoted as \mathcal{L}_1) between \hat{x}_t and x_t to improve convergence.

Additionally, we propose to assess reconstruction quality at the frame feature and action levels. The computation of these additional losses is represented in the green block of Fig. 2. The reconstructed frame \hat{x}_t is fed to E leading to

frame features \hat{f}_t . Then, the Euclidean distance between initial and reconstructed features is employed:

$$\mathcal{L}_{\text{rec}}^f = \frac{1}{T} \sum_{t=1}^T \|f_t - \hat{f}_t\|_2^2. \quad (10)$$

The motivation for this loss is that it enforces that the semantic information extracted from input frames x_t is also preserved in \hat{x}_t . To avoid trivial minimization of (10), we do not back-propagate the loss gradient through f_t .

Action losses. Regarding action understanding, we propose an information-theoretic objective that acts on the action probabilities. The reconstructed *frame features* \hat{f}_t and \hat{f}_{t+1} are fed to A that internally estimates the action probabilities \hat{p}_t . We then maximize the mutual information between the actions extracted from the input sequences and the corresponding ones extracted from the reconstructed sequences:

$$\max_{\theta} \mathcal{MI}(p_t, \hat{p}_t) = \max_{\theta} (\mathcal{H}(p_t) - \mathcal{H}(p_t|\hat{p}_t)). \quad (11)$$

Maximizing this objective function has two desirable effects. First, minimizing the conditional entropy term $\mathcal{H}(p_t|\hat{p}_t)$ imposes that the same action must be inferred from the input and the reconstructed sequence. Second, maximizing the entropy term $\mathcal{H}(p_t)$ enforces that the maximal number of actions must be discovered. This avoids trivial solutions where only a single action label a_t is constantly predicted and ensures that the K *action direction centroids* $\{c_k\}_{k=1}^K$ do not collapse to a single point.

To compute the mutual information in Eq. (11), we consider a mini-batch \mathcal{B} of action probabilities (p_t, \hat{p}_t) . We estimate the joint probability matrix $P_{ij} = P(a_t=i, \hat{a}_t=j)$, with $(i, j) \in \{1, \dots, K\}^2$, as follows:

$$P_{ij} = P(a_t=i, \hat{a}_t=j) = \frac{1}{\bar{\mathcal{B}}} \sum_{(p_t, \hat{p}_t) \in \mathcal{B}} p_t \cdot \hat{p}_t^\top, \quad (12)$$

where $\bar{\mathcal{B}}$ denotes the size of \mathcal{B} . We obtain $P_i = P(a_t=i)$ and $P_j = P(\hat{a}_t=j)$ by marginalization over P_{ij} . The mutual information loss term is then given by:

$$\mathcal{L}_{\text{act}} = -\mathcal{MI}(p_t, \hat{p}_t) = \sum_{t=1}^K \sum_{j=1}^K P_{ij} \ln \frac{P_{ij}}{P_i \cdot P_j}. \quad (13)$$

This action matching loss \mathcal{L}_{act} is completed by a loss that enforces that the *action variability embeddings* match. Considering the *action direction* distributions $d_t \sim \mathcal{N}(\mu_{d_t}, \sigma_{d_t}^2)$ and $\hat{d}_t \sim \mathcal{N}(\mu_{\hat{d}_t}, \sigma_{\hat{d}_t}^2)$ associated with each (p_t, \hat{p}_t) in \mathcal{B} , we measure the average Kullback–Leibler (KL) divergence between the distribution of \hat{d}_t and d_t :

$$\mathcal{L}_{\text{rec}}^a = \frac{1}{\bar{\mathcal{B}}} \sum_{(\hat{d}_t, d_t) \in \mathcal{B}} \mathcal{D}_{\text{KL}}(\mathcal{N}(\mu_{\hat{d}_t}, \sigma_{\hat{d}_t}^2) \parallel \mathcal{N}(\mu_{d_t}, \sigma_{d_t}^2)). \quad (14)$$

Assuming a Gaussian prior for d_t , we minimize the average KL divergence \mathcal{D}_{KL} between the distributions of d_t and a

Gaussian prior with unit variance $\mathcal{N}(0, I_n)$,

$$\mathcal{L}_{\text{KL}} = \frac{1}{\bar{\mathcal{B}}} \sum_{d_t \in \mathcal{B}} \mathcal{D}_{\text{KL}}(\mathcal{N}(\mu_{d_t}, \sigma_{d_t}^2) \parallel \mathcal{N}(0, I_n)). \quad (15)$$

Total loss. The total objective is given by:

$$\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_{\text{rec}}^x + \lambda_{\text{rec}}^f \mathcal{L}_{\text{rec}}^f + \lambda_{\text{rec}}^a \mathcal{L}_{\text{rec}}^a + \lambda_{\text{act}} \mathcal{L}_{\text{act}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}}, \quad (16)$$

where $\lambda_{\text{act}}, \lambda_{\text{rep}}, \lambda_{\text{KL}}, \lambda_{\text{rec}}^a$ are positive weighting parameters. These parameters are estimated on the training set: for every term, we iteratively experiment raising parameter values until we can observe that the first back-propagation steps lead to a decrease of the loss term.

Test time. At test time, CADDY receives as input the initial video frame and the user provides an action label at every time-step. We follow an auto-regressive approach where the estimated \hat{x}_{t+1} is used as input at the next time step (instead of x_{t+1}). We no longer employ the action network A and use the action labels a_t provided by the user. Regarding the *action variability embedding*, we employ $v_t=0$ that corresponds to the maximum of its posterior distribution.

Note that, this difference between training and test would affect the performance of the network that has never received generated images as inputs. To mitigate this problem, we propose a mixed training procedure. For the first T_f frames, we use the training procedure as described above where the *Dynamic Network R* receives in input the *frame features* f_t computed from original frames. Then, for $t > T_f$, R is fed with the *frame features* \hat{f}_t computed from the reconstructed frames \hat{x}_t . By mixing original and reconstructed frames, we obtain a network that can handle both real and reconstructed frames, and consequently, is less prone to the shift issue typical of auto-regressive methods.

4. Experiments

Datasets. We evaluate our method on three video datasets:

- *BAIR* robot pushing dataset [9]. We employ a version of the dataset in 256x256 resolution, composed of about 44K videos of 30 frames. Ground-truth robotic arm positions are available but are only used for evaluation purposes.
- *Atari Breakout* dataset. We collect a dataset using a Rainbow DQN agent [15] trained on the Atari Breakout video game environment. We collect 1407 sequences of about 32 frames with resolution 160x210 (358 for training, 546 sequences for validation and 503 for testing).
- *Tennis* dataset. We collect Youtube videos corresponding to two tennis matches from which we extract about 900 videos with resolution 256x96. (See *Sup. Mat.* for more details). To respect the single agent assumption, we consider only the lower part of the field.

Evaluation Protocol. We propose to evaluate both action and video generation qualities by comparing the models on the video reconstruction task. A test sequence is considered

and the action network is used to extract the sequence of learned discrete actions characterizing the input sequence. Starting from the initial frame, the extracted actions are used to reconstruct the remaining of the sequence. The evaluation is completed with a user-study that directly assesses the quality of the learned set of discrete actions. We adopt a large set of metrics.

Video quality metrics.

- *LPIPS* [50]: We report the average *LPIPS* computed on corresponding frames of input and reconstructed sequences.
- *FID* [16]: We report the average *FID* between the original and reconstructed frames.
- *FVD* [43]: We compute the *FVD* between the original and reconstructed videos.

Action-space metrics. We introduce two metrics that measure the quality of the action space. Both metrics use additional knowledge (*i.e.* ground-truth information or an externally trained detector) to measure motion consistency among frames where the same action is performed. Assuming two consecutive frames, we measure the displacement Δ of a reference point on the object of interest. On the *Tennis* dataset, we employ FasterRCNN [33] to detect the player and use the bounding box center as the reference point. On *Atari Breakout*, we employ a simple pixel matching search to detect the rectangular platform and use its center as the reference point. Finally, on *BAIR*, we use the ground-truth location information of the robotic arm. The key idea of the two action quality metrics is to assess whether the predicted action labels and the displacement Δ are consistent by trying to predict one from the other:

- Δ *Mean Squared Error* (Δ -*MSE*): This metric measures how Δ can be regressed from the action label. For each action, we estimate the average displacement Δ , which is the optimal estimator for Δ (in terms of MSE). We report the MSE to evaluate regression quality. To facilitate comparison among datasets, we normalize the MSE by the variance of Δ over the dataset.
- Δ -*based Action Accuracy* (Δ -*Acc*): This metric measures how the predicted action can be predicted from the displacement Δ . To this aim, we train a linear classifier and report the action-accuracy measured on the test set.

Action-conditioning metrics. Finally, we consider two metrics that measure how the action label conditions the generated video. Therefore, it evaluates both generation quality and the learned action space:

- *Average Detection Distance* (*ADD*). Similarly to [38], we report the Euclidean distance between the reference object keypoints in the original and reconstructed frames. We employ the same detector as for the action-quality metrics to estimate the reference points.
- *Missing Detection Rate* (*MDR*): We report the percentage of detections that are successful in the input sequence, but not in the reconstructed one. This represents the proportion

Variant	G.S	v_t	\mathcal{L}_{act}	LPIPS↓	FID↓	FVD↓	Δ -MSE↓	Δ -Acc↑
(i)				0.263	80.0	1300	69.7	51.2
(ii)	✓			0.209	42.3	571	64.8	37.9
(iii)	✓		✓	0.249	76.4	1130	92.7	24.1
(iv)	✓	✓		0.245	76.9	1130	93.7	27.6
CADDY (Full)	✓	✓	✓	0.202	35.9	423	54.8	69.0

Table 1: Ablation results on the *BAIR* dataset. G.S: use of Gumbel-Softmax, v_t ; use of the *action variability embedding* v_t ; \mathcal{L}_{act} : training with the mutual information loss \mathcal{L}_{act} . Δ -*MSE* and Δ -*Acc* in %.

of frames where the object of interest is missing.

4.1. Experimental analysis of CADDY

Ablation study. In this section, we study the impact of three key elements: Gumbel-Softmax sampling, the action loss \mathcal{L}_{act} and *action variability embeddings* v_t . We produce the following variants of our method: (i) uses none of these components, (ii) introduces G.S., (iii) employs G.S. and \mathcal{L}_{act} , (iv) uses G.S. and v_t . The *BAIR* dataset is used because ground-truth displacement Δ are available.

The results are presented in Tab. 1. In (i), when no component is used, the lack of G.S. sampling makes the model exploit p_t at training time to encode information about the next frame, learning a continuous rather than a discrete action representation. At test time, when discrete actions are used, this mismatch leads to poor performance. When G.S. sampling is introduced in (ii), the model learns discrete action representations with a Δ -*MSE* of 64.8%. However, the model lacks a mechanism to resolve the ambiguity of which variation of the discrete action should be performed. This results in difficulties in optimizing the reconstruction objective and leads to reduced video quality. In (iii), when both G.S. sampling and \mathcal{L}_{act} are present, the optimization process favors \mathcal{L}_{act} rather than the reconstruction objective, resulting in degraded performance. Lastly, in (iv), when both G.S. sampling and v_t are used without \mathcal{L}_{act} , the model uses v_t at training time to encode complete information about the next frame and a discrete action space is not learned. At test time, when $v_t=0$, performance of the model is poor. Overall, this ablation study confirms the positive impact of G.S sampling, the *action variability embeddings*, and our mutual-information loss on the performance.

Visualization. Fig. 3 shows the object motion distribution corresponding to each action learned by CADDY. On *BAIR*, the model learns actions related to movements on the x (1, 4), y (2, 3) and z (6) axes as well as a no-movement action (7). On *Atari Breakout*, the model learns actions corresponding to the three possible movements, as well as the inertia of the platform and the physics of the ball and blocks. Finally, on *Tennis*, the model learns actions corresponding to forward (2) and backward (3) movement, lateral movement (4, 7), no-movement (6) and hitting (1, 5).

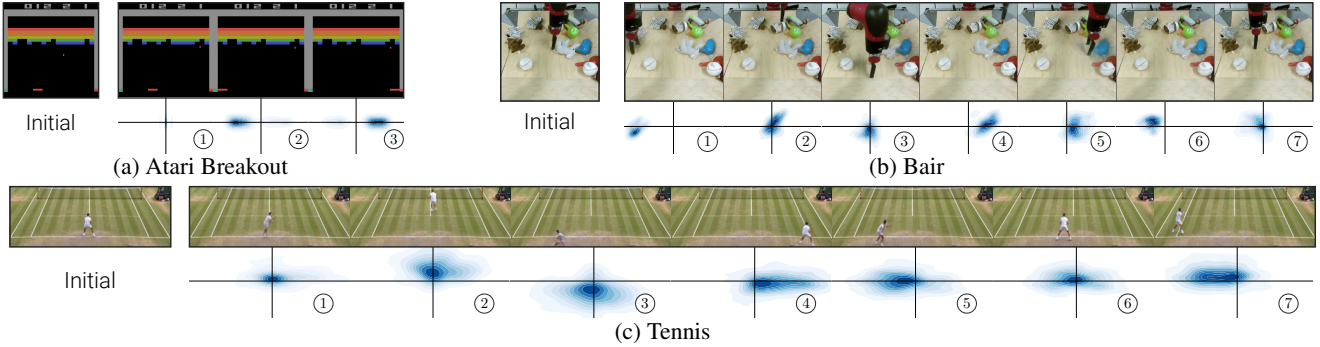


Figure 3: Qualitative evaluation of the action space learned by our model on the three datasets. We consider an initial frame and, for every action, produce a sequence by repeatedly selecting that action and show the final frame. The bottom row shows, for each action, the distribution of the displacement Δ associated with the object of interest.

4.2. Comparison with Previous Works

Baseline selection. Since we present the first method for unsupervised PVG, we compare the proposed model against a selection of baselines, adapting existing methods to this new settings. We consider two criteria in the selection of the baselines: the architecture should be adaptable to the new setting without deep modifications, and the source code must be available to include these modifications. Comparing existing video prediction methods in terms of FVD on a 64x64 version of the *BAIR* dataset (quantitative comparison available in *Sup. Mat*) shows that SAVP [25] is the best performing method with the code publicly available. Only the methods described in [7, 26, 47] are marginally outperforming SAVP but their code is not available.

Furthermore, SAVP [25] is widely used and features an architecture prone to be adapted to the current setting. During training, SAVP learns an encoder E that encodes information relative to the transition between successive frames which is used by the generator in the synthesis of the next frame. After training the SAVP model, we cluster the latent space learned by the encoder network on the training sequences using K-means. This procedure produces a set of K centroids that we use as action labels. During evaluation, we compute the latent representations for the input sequence and assign them the index of the nearest centroid.

As a second baseline, we choose MoCoGAN [42]. This method possesses favorable characteristics for adaptation: it separates the content from the motion space and it employs an InfoGAN [5] loss to learn discrete video categories. We consider 6-frame short videos and assume a constant action over them. We employ an action discriminator that predicts the action performed in the input video sequence. This predicted action is used as an auxiliary input to the generator and the InfoGAN loss is used to learn the action space.

Finally, we also include SRVP [13] in our comparison. SRVP is modified similarly to SAVP [25] to handle PVG. However, due to the poor empirical results and the high training costs, we consider this baseline only on the *BAIR*

dataset. Several other works have been considered, but not adopted [21, 48] since they would require too important modifications to handle the PVG problem.

SAVP [25], MoCoGAN [42] and SRVP [13] are designed for video generation at resolutions lower than our datasets. Therefore, we generate videos at the resolution of 64x80, 64x64 and 128x48 respectively for the *Atari Breakout*, *BAIR* and *Tennis* datasets, and then upscale to full resolution for evaluation. In addition, we create two other baselines, referred to as SAVP+ and MoCoGAN+, obtained by increasing the capacity of the respective networks to generate videos at full resolution. Due to the high memory requirements of SRVP [13], it was not feasible to produce a version operating at full resolution.

Quantitative evaluation. We show the results for the evaluation on the *BAIR* dataset in Tab. 2. According to Δ -MSE, MoCoGAN, SAVP and our model learn an action space correlated with the movement of the robot (Δ -MSE < 100), but our method surpasses these baselines by 26.1% in Δ -MSE and by 24.2% in Δ -Acc, showing greater capacity in learning a set of discrete actions. Finally, SRVP clearly underperforms all the other methods in term of reconstruction (LPIPS), video quality (FID and FVD). Moreover, SRVP predicts a single action class for the whole test set. This degenerated behavior results in a 100% Δ -ACC and a poor Δ -MSE. Because of these poor results and its very high computational requirement, SRVP is not shown in the following.

Tab. 3 reports the results on the *Atari Breakout* dataset. Our method outperforms the baselines in both action and video quality metrics. In particular, our model replicates the movement of the platform with an average ADD of 7.29 pixels and an MDR of 2.70%, indicating that the learned set of actions are consistent with the movement of the user-controlled *Atari Breakout* platform. Note that MoCoGAN obtains a lower MDR but a much higher ADD showing that it generates correctly the platform but at the wrong position.

The evaluation results on the *Tennis* dataset are reported in Tab. 4. With the exception of FVD, where the result is

	LPIPS ↓	FID ↓	FVD ↓	Δ -MSE ↓	Δ -Acc ↑
MoCoGAN [42]	0.466	198	1380	88.8	20.7
MoCoGAN+	0.201	66.1	849	98.4	22.9
SAVP [25]	0.433	220	1720	80.9	41.4
SAVP+	0.154	27.2	303	82.0	44.8
SRVP [13]	0.491	224	3540	(100)	(100)
CADDY (Ours)	0.202	35.9	423	54.8	69.0

Table 2: Comparison with baselines on the *BAIR* dataset. Δ -MSE, Δ -Acc and MDR in %, ADD in pixels. Parentheses indicate degenerated cases resulting in uninformative metrics (see details in the text).

	LPIPS ↓	FID ↓	FVD ↓	Δ -MSE ↓	Δ -Acc ↑	(ADD, MDR) ↓
MoCoGAN [42]	0.234	99.9	447	99.6	81.9	(46.0, 0.795)
MoCoGAN+	65.8e-3	10.4	103	103	57.5	(54.6, 17.4)
SAVP [25]	0.239	98.4	487	103	58.1	(24.7, 21.0)
SAVP+	39.3e-3	4.84	84.4	104	85.6	(15.8, 51.5)
CADDY (Ours)	7.66e-3	0.716	5.94	82.7	91.6	7.29 , 2.70

Table 3: Comparison with baselines on the *Atari Breakout* dataset. Δ -MSE, Δ -Acc and MDR in %, ADD in pixels.

	LPIPS ↓	FID ↓	FVD ↓	Δ -MSE ↓	Δ -Acc ↑	(ADD, MDR) ↓
MoCoGAN [42]	0.266	132	3400	101	26.4	(28.5, 20.2)
MoCoGAN+	0.166	56.8	1410	103	28.3	(48.2, 27.0)
SAVP [25]	0.245	156	3270	112	19.6	(10.7, 19.7)
SAVP+	0.104	25.2	223	116	33.1	(13.4, 19.2)
CADDY (Ours)	0.102	13.7	239	72.2	45.5	8.85 , 1.01

Table 4: Comparison with baselines on the *Tennis* dataset. Δ -MSE, Δ -Acc and MDR in %, ADD in pixels.




Figure 4: Reconstruction results on the *BAIR* (left) and *Tennis* (right) datasets. We zoom in for better visualization.

close to SAVP+, our method obtains the best performance in all the metrics. A Δ -MSE of 72.2% shows that actions with a consistent associated movement are learned. On the other hand, the Δ -MSE scores of the other baselines show that the movements associated with each action do not present a consistent meaning. Moreover, our method features the lowest ADD and a significantly lower MDR of 1.01% which indicate that our method consistently generates the player and moves it accurately on the field.

Qualitative evaluation In Fig. 4, we show examples of re-

	Agreement	Diversity	Other votes
MoCoGAN [42]	-3.15e-3	1.58	1.80%
MoCoGAN+	-2.84e-3	1.51	28.0%
SAVP [25]	0.0718	1.69	7.14%
SAVP+	-1.97e-3	1.80	5.40%
CADDY (Ours)	0.469	1.65	1.61%

Table 5: User study results on the *Tennis* dataset.

constructed sequences on the *BAIR* and *Tennis* datasets. Our method achieves a precise placement of the object of interest with respect to the input sequence. In the *Sup. Mat.* we show additional qualitative results as well as sample videos generated interactively.

User evaluation. To complete our evaluation, we perform a user study on the *Tennis* dataset. We sample 23 random frames that we use as initial frames. For each of them, we generate K continuations of the sequences, one for each action, each produced by repeatedly selecting the corresponding action. For each sequence, users are asked to select the performed action among a predefined set (*Left, Right, Forward, Backward, Hit the ball* or *Stay*). An additional option *Other* is provided in case the user cannot recognize any action. We measure user agreement using the Fleiss’ kappa measure [11] that is commonly used to evaluate agreement between categorical ratings [12]. In addition, to validate that all the actions can be generated (*i.e.* detecting mode collapse), we compute the diversity of the action space, expressed as the entropy of the user-selected actions. Results are shown in Tab. 5. While all methods capture actions with high diversity, our method shows a higher agreement, indicating that users consistently associate the same action label to each learned action. Looking at the details of the votes for our method (see *Sup. Mat.*), we observe that most disagreements are due to cases where the player hits the ball while moving. In contrast, the other methods do not obtain high user agreement, with actions assuming different meanings, depending on the particular frame used as context.

5. Conclusions

In this work, we propose the unsupervised learning problem of playable video generation. We introduce CADDY, a self-supervised method based on an encoder-decoder architecture that uses predicted action labels as a bottleneck. We evaluate our method on three varied datasets and show state-of-the-art performance. Our experiments show that we can learn a rich set of actions that offer the user a gaming-like experience to control the generated video. As future work, we plan to extend our method to multi-agent environments.

6. Acknowledgements

This project was supported by the EU H2020 AI4MEDIA No. 951911 project and by the Caritro Deep Learning Lab of the ProM Facility of Rovereto.

References

- [1] Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H. Campbell, and Sergey Levine. Stochastic variational video prediction. In *International Conference on Learning Representations (ICLR)*, 2018. [2](#)
- [2] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *CoRR*, abs/1606.01540, 2016. [2](#)
- [3] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Computer Vision – European Conference of Computer Vision (ECCV)*, pages 139–156, 2018. [2](#)
- [4] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 5933–5942, 2019. [2](#)
- [5] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, volume 29, pages 2172–2180, 2016. [7](#)
- [6] Silvia Chiappa, Sébastien Racanière, Daan Wierstra, and Shakir Mohamed. Recurrent environment simulators. *CoRR*, abs/1704.02254, 2017. [2](#)
- [7] Aidan Clark, Jeff Donahue, and Karen Simonyan. Efficient video generation on complex datasets. *CoRR*, abs/1907.06571, 2019. [7](#)
- [8] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal cycle-consistency learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1801–1810, 2019. [3](#)
- [9] Frederik Ebert, Chelsea Finn, Alex X. Lee, and S. Levine. Self-supervised visual planning with temporal skip connections. In *CoRL*, 2017. [2](#), [5](#)
- [10] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction. In *Advances in Neural Information Processing Systems (NIPS)*, volume 29, pages 64–72, 2016. [2](#)
- [11] Joseph L Fleiss. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378, 1971. [8](#)
- [12] Joseph L Fleiss, Bruce Levin, Myunghee Cho Paik, et al. The measurement of interrater agreement. *Statistical methods for rates and proportions*, 2(212-236):22–23, 1981. [8](#)
- [13] Jean-Yves Franceschi, Edouard Delasalles, Mickael Chen, Sylvain Lamprier, and P. Gallinari. Stochastic latent residual video prediction. *ArXiv*, abs/2002.09219, 2020. [1](#), [2](#), [7](#), [8](#)
- [14] Oran Gafni, Lior Wolf, and Yaniv Taigman. Vid2game: Controllable characters extracted from real-world videos. In *International Conference on Learning Representations (ICLR)*, 2020. [2](#)
- [15] Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Gheshlaghi Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *AAAI*, pages 3215–3222, 2018. [5](#)
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems (NIPS)*, volume 30, pages 6626–6637, 2017. [6](#)
- [17] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. 2017. [4](#)
- [18] Xu Ji, Andrea Vedaldi, and João F. Henriques. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 9864–9873, 2019. [3](#)
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision – European Conference of Computer Vision (ECCV)*, 2016. [4](#)
- [20] Seung Wook Kim, Yuhao Zhou, Jonah Philion, Antonio Torralba, and Sanja Fidler. Learning to Simulate Dynamic Environments with GameGAN. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020. [1](#), [2](#)
- [21] Yunji Kim, Seonghyeon Nam, In Cho, and Seon Joo Kim. Unsupervised keypoint learning for guiding class-conditional video prediction. In *Advances in Neural Information Processing Systems (NIPS)*. 2019. [2](#), [7](#)
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *stat*, 1050:10, 2014. [4](#)
- [23] Manoj Kumar, Mohammad Babaeizadeh, Dumitru Erhan, Chelsea Finn, Sergey Levine, Laurent Dinh, and Durk Kingma. Videoflow: A conditional flow-based model for stochastic video generation. In *International Conference on Learning Representations*, 2019. [2](#)
- [24] Yong-Hoon Kwon and Min-Gyu Park. Predicting future frames using retrospective cycle gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1811–1820, 2019. [2](#)
- [25] Alex X. Lee, Richard Zhang, Frederik Ebert, P. Abbeel, Chelsea Finn, and S. Levine. Stochastic adversarial video prediction. *ArXiv*, abs/1804.01523, 2018. [1](#), [2](#), [7](#), [8](#)
- [26] Pauline Luc, A. Clark, S. Dieleman, D. Casas, Yotam Doron, Albin Cassirer, and K. Simonyan. Transformation-based adversarial video prediction on large-scale data. *ArXiv*, abs/2003.04035, 2020. [7](#)
- [27] Zelun Luo, Boya Peng, De-An Huang, Alexandre Alahi, and Li Fei-Fei. Unsupervised learning of long-term motion dynamics for videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2203–2212, 2017. [3](#)
- [28] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. [1](#), [2](#)
- [29] Willi Menapace, Stéphane Lathuilière, and Elisa Ricci. Learning to cluster under domain shift. In *Computer Vision – European Conference of Computer Vision (ECCV)*, pages 736–752, Cham, 2020. [3](#)
- [30] Manuel Serra Nunes, Atabak Dehban, Plinio Moreno, and José Santos-Victor. Action-conditioned benchmarking of

- robotic video prediction models: a comparative study. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8316–8322. IEEE, 2020. 1, 2
- [31] 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 (NIPS)*, pages 2863–2871, 2015. 1, 2
- [32] B. Peng, J. Lei, H. Fu, C. Zhang, T. Chua, and X. Li. Unsupervised video action clustering via motion-scene interaction constraint. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(1):131–144, 2020. 3
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, volume 28, pages 91–99, 2015. 6
- [34] Oleh Rybkin, Karl Pertsch, Konstantinos G Derpanis, Kostas Daniilidis, and Andrew Jaegle. Learning what you can do before doing anything. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [35] Masaki Saito, Shunta Saito, Masanori Koyama, and Sotuke Kobayashi. Train sparsely, generate densely: Memory-efficient unsupervised training of high-resolution temporal gan. *International Journal of Computer Vision (IJCV)*. 2
- [36] Xingjian SHI, Zhourong 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 (NIPS)*, pages 802–810. 2015. 3
- [37] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [38] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Conference on Neural Information Processing Systems (NeurIPS)*, December 2019. 2, 4, 6
- [39] Khurram Soomro and Mubarak Shah. Unsupervised action discovery and localization in videos. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 696–705, 2017. 3
- [40] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research (JMLR)*, 15(1):1929–1958, 2014. 1
- [41] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018. 3
- [42] 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 (CVPR)*, June 2018. 1, 2, 7, 8
- [43] 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. 6
- [44] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating the future by watching unlabeled video. *arXiv preprint arXiv:1504.08023*, 2, 2015. 1, 2
- [45] Yaohui Wand, Piotr Bilinski, Francois Bremond, and Antitza Dantcheva. Imaginator: Conditional spatio-temporal gan for video generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, March 2020. 2
- [46] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1144–1156, 2018. 2
- [47] Dirk Weissenborn, Oscar Täckström, and Jakob Uszkoreit. Scaling autoregressive video models. In *International Conference on Learning Representations (ICLR)*, 2020. 2, 7
- [48] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3964–3973, 2018. 7
- [49] Haotian Zhang, Cristobal Sciutto, Maneesh Agrawala, and Kayvon Fatahalian. Vid2player: Controllable video sprites that behave and appear like professional tennis players, 2020. 2
- [50] 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 (CVPR)*, 2018. 6