

TwoStreamVAN: Improving Motion Modeling in Video Generation

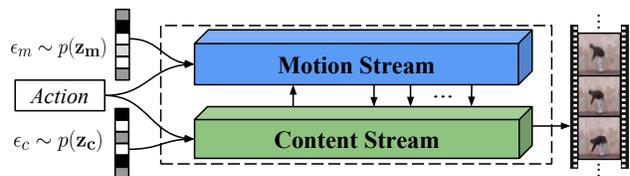
Ximeng Sun [†] Huijuan Xu ^{† ‡} Kate Saenko [†]
[†]Boston University [‡]University of California, Berkeley
[†]{sunxm, hxu, saenko}@bu.edu [‡]huijuan@berkeley.edu

Abstract

Video generation is an inherently challenging task, as it requires modeling realistic temporal dynamics as well as spatial content. Existing methods entangle the two intrinsically different tasks of motion and content creation in a single generator network, but this approach struggles to simultaneously generate plausible motion and content. To improve motion modeling in video generation task, we propose a two-stream model that disentangles motion generation from content generation, called a Two-Stream Variational Adversarial Network (TwoStreamVAN). Given an action label and a noise vector, our model is able to create clear and consistent motion, and thus yields photorealistic videos. The key idea is to progressively generate and fuse multi-scale motion with its corresponding spatial content. Our model significantly outperforms existing methods on the standard Weizmann Human Action, MUG Facial Expression and VoxCeleb datasets, as well as our new dataset of diverse human actions with challenging and complex motion. Our code is available at <https://github.com/sunxm2357/TwoStreamVAN/>.

1. Introduction

Despite great progress being made in generation of well-structured static images [47, 45] using methods such as GANs/VAEs ([6, 16, 30, 31, 33, 26, 21, 17, 18]), generation of pixel-level video has yet to achieve similarly impressive results [41, 34, 39, 15]. The challenge of video generation lies in the need to construct unstructured spatial content with both foreground objects and background, and simultaneously model natural motion at different scales and locations. Existing methods fail to generate convincing videos (see examples in Fig. 1b or the video in the supplementary material) mostly because of ineffective motion modeling, which in turn deteriorates content generation. A major problem with pixel-level video prediction [27, 36, 40] and generation methods is that they attempt to model both static content and dynamic motion in a single entangled generator, regardless of whether they disentangle the motion and content in the latent space or not.



(a) Overview of Our TwoStreamVAN



(b) TwoStreamVAN (left) v.s. MoCoGAN (right)

Figure 1: We propose Two-Stream Variational Adversarial Network (TwoStreamVAN) to improve motion modeling. Given an action class, it separately generates content and motion from the disentangled content and motion latent codes ϵ_c and ϵ_m and fuse them via a novel multi-scale fusion mechanism. Our model generates significantly better videos than MoCoGAN without separation motion generation.

We argue that separating motion and content modeling in the decoding phase is crucially important. On one hand, it removes the disturbance of motion during content generation and results in better content structures. On the other hand, separate motion modeling helps to produce action-relevant movement consistently across the entire sequence.

We propose a novel *Two-Stream Variational Adversarial Network (TwoStreamVAN)* that generates a video from an action label and a disentangled noise vector (Fig. 1a). Rather than over-estimating a single generator’s ability, we introduce two parallel generators that process content and motion separately and fuse them together to predict the next frame. Intuitively, motion is usually conditioned on the spa-

tial information within its local context window, e.g., different motions near the arms and legs when doing jumping jacks. Thus, we define the fusion of motion and content as a learned refinement of pixel values unique to each location. To precisely generate multi-scale motion, we conduct such refinement on intermediate content layers with several different resolutions. Moreover, we introduce motion mask to make the learning of motion only focus on the regions where it exists.

Considering the great success achieved in image generation, we expect that introducing image-level supervision will reduce the deterioration of content generation common in most video generation models. Previously, MoCoGAN [39] tried image-level supervision but made little progress, since the content generation was still affected by the motion modeling in a single shared generator. Notably, a key advantage of our disentangled two-stream generators is the ability to learn each stream separately and thus more accurately. We learn the image structure in its own generator fully via image-level supervision, which significantly enhances content generation performance. Well-trained content generation further benefits motion learning in video-level supervision.

We evaluate our approach on two standard video generation benchmarks, Weizmann Human Action [12] and MUG Facial Expression [1] as well as a large-scale speech video dataset – VoxCeleb [28]. To verify the advantage of our two-stream generator in modeling motion, we propose a large synthetic human action dataset, called *SynAction*, with challenging motion complexity (120 unique action models across 20 different non-rigid human actions, like running or squatting), using a library of video game actions, Mixamo [7]. To summarize, our contributions are:

- We propose a video generation model *TwoStreamVAN* as well as a more effective learning scheme, which disentangle motion and content in the generation phase.
- We design a multi-scale motion fusion mechanism and further improve motion modeling by conditioning on the spatial context;
- We create a large-scale synthetic video generation dataset available to the research community;
- We evaluate our model on four video datasets both quantitatively and qualitatively (with user studies), and demonstrate superior results to several strong baselines.

2. Related Work

Generative Models. The two main deep learning methods for image generation are VAEs [20, 14, 13, 44] and GANs [11, 30, 31, 6]. VAEs provide probabilistic descriptions of observations in latent spaces, but might generate blurry and unrealistic images in their vanilla form. GANs

propose an adversarial discriminator to encourage the generation of crisper images, but suffer from mode collapse [46] and unexpected bizarre artifacts. [23, 25] combine a VAE and a GAN and propose a Variational Adversarial Network (VAN) to learn an interpretable latent space as well as generate realistic images. In the light of the VAN’s success in image generation, we adopt it here for video generation.

Video Generation is a challenging video task, which maps random noise to the content and motion which form a plausible video sequence. VideoVAE [15] shows the VAE’s ability to produce video, proposing a structured latent space and an encoder-generator architecture to generate the video recurrently. Except for this work, the prevailing trend is to generate videos from noise utilizing the GAN training paradigm. VideoGAN [41] generates a sequence of foreground objects with a single shared static background. Instead of generating foreground and background separately, TGAN [34] decodes each frame from a unified spatio-temporal latent representation. MoCoGAN attempts to disentangle content and motion by sampling from separate latent spaces but uses a single generator to decode these two latent codes together. To overcome the ineffective motion modeling and the consequent content deterioration in the unified generation process, we further introduce disentangled content and motion generators to model spatial structures and temporal dynamics separately.

Video Prediction [9, 10, 8, 40, 2, 24, 36], a related yet different task, outputs additional frames for a partial input video by borrowing the content and extrapolating the motion from the given input video. The main difference of two tasks lies in the origins of the content and motion in the new output sequence: the generation samples them from noise, but the prediction borrows the content and extrapolates the motion from the input frames. Thus, they adopt different metrics, emphasis on the future reconstruction and the generation reality/diversity respectively. Despite of differences, they share some technical details in common. For instance, [40, 9] decompose the content and motion in the encode phase; [10] discusses different motion transformations as effective motion representations. In this paper, we mainly focus on Video Generation and show our superiority over these shared techniques in the pure generative context.

Multi-scale Motion Estimation and Prediction. To tackle multi-scale motion in actual videos, [3, 4, 37, 32] build a pyramid of real images, while [38, 43] constructs a pyramid of embedded feature maps. In our model, we take the second method. Traditionally motion between two adjacent frames is introduced by warping the dense optical flow features with the current frame [22, 19]. To avoid the high computational cost of optical flow, Xue *et al.* [43] generates multi-scale motion kernels from the difference map of adjacent frames. They learn generic motion kernels shared for the entire image which is hard to interpret, because motion

is usually conditioned on the local spatial context. In this paper, we produce motion kernels specifically for each spatial location and use these kernels to refine intermediate layers in content feature pyramid to model multi-scale motion. Furthermore, as [43] only predicts the next frame from the current, we design our model to address a more challenging task which generates the video sequence without receiving any visual clue as inputs.

3. Approach

We introduce a Two-Stream Variational Adversarial Network (Fig. 2), which generates a video given an input action label and a random noise vector. We define action-conditioned video generation as follows. Suppose we have K different action classes. In each class $k \in 1:K$, let C_k be the number of training videos. Let $V_k = \{v_{i,k}, \forall i \in 1:C_k\}$ be the set of all videos for class k and $v_{i,k} = \{x_1, x_2, \dots, x_T\}$ be a short video clip with T frames. The task is to define a function G which generates a plausible video $\hat{v} = \{\hat{x}_1, \dots, \hat{x}_T\}$ conditioned on the given class label k from a latent vector $\epsilon \in \mathbb{R}^{\mathcal{N}}$, i.e. $\hat{v} = G(k, \epsilon)$.

We disentangle the latent space ϵ into two independent codes: a content code $\epsilon_c \in \mathbb{R}^C$, and a motion code $\epsilon_m \in \mathbb{R}^{\mathcal{M}}$, with $\mathcal{N} = C + \mathcal{M}$. We also disentangle the video generation function G into two separate content and motion functions (G_c and G_m), in contrast to previous work that used a single generator [41, 34, 39], and we design a novel multi-scale fusion mechanism.

3.1. Two-Stream Generation

To learn generative functions G_c and G_m for content and motion modeling, we introduce two separate action-conditioned VAN streams with interactions at several stages. Each VAN stream contains an encoder, a generator and a discriminator, where the encoder and generator serve as the auto-encoder in the VAE, and the generator and discriminator comprise the GAN.

The **Content VAN Stream** consists of a Content Encoder E_c , a Content Generator G_c and an Image Discriminator D_I . After observing a single frame x , E_c generates the posterior content latent distribution $q(\mathbf{z}_c|x, k)$, which is close to its true prior distribution $p(\mathbf{z}_c|k)$. G_c decodes a content vector ϵ_c sampled from the content distribution into a frame \hat{x} . D_I discriminates real/generated frames to encourage G_c to generate realistic image patterns.

Similarly, the **Motion VAN Stream** consists of a Motion Encoder E_m , a Motion Generator G_m and a Video Discriminator D_V . Instead of encoding spatial content, in our approach E_m models the temporal dynamics in the difference map Δx between neighbor frames. It generates the posterior motion latent distribution $q(\mathbf{z}_m|\Delta x, k)$ which is close to its true prior distribution $p(\mathbf{z}_m|k)$. A convLSTM [42] ac-

cumulates the motion history and generates the current motion embedding e_m by receiving a sequence of ϵ_m sampled from the motion distributions at all previous time steps. G_m takes in ϵ_c and e_m to generate motion at different scales. We generate each video \hat{v} by fusing the generated motion with the corresponding content at T time steps (see Sec. 3.2). D_V discriminates the real/generated videos and additionally classifies their actions to encourage G_m to generate realistic motion for action k .

3.2. Multi-scale Motion Generation and Fusion

At pixel (a, b) , the motion usually happens within a local window between adjacent frames. Inspired by the spatial convolution for frame interpolation [29], we represent motion as a refinement of the current pixel value based on its local context and fuse such motion with content via spatially-adaptive convolution. Moreover, we propose a novel multi-scale fusion mechanism to overcome the drawbacks of their approach, namely: 1) ineffective modeling of multi-scale motion due to the single fusion step performed on the full-resolution image and 2) the high demand for memory due to the large convolution kernels used to represent the maximum possible motion.

To generate precise motion, G_m takes the disentangled content and motion embeddings (ϵ_c and e_m) as input. In G_m , motion at different scales is separated at the corresponding hidden layers: large motion comes from low resolution layers and small motion from high resolution layers. At each layer, G_m uses the current feature map to: 1) calculate the motion (pixel refinement) in the form of pixel-wise 2D kernels with size n ; 2) identify the regions containing actual movements in a motion mask; 3) generate the motion map for the following layer. To fuse motion with content, we convolve the generated 2D kernels with patches centered at their corresponding pixels respectively (i.e. perform spatially-adaptive convolution) only if the pixel is in the regions where motion is activated (see Fig. 3).

Specifically, suppose we separate the motion into S scales. For each scale s , let l_s be the resolution of its corresponding hidden layer, $w^s \in \mathbb{R}^{l_s \times l_s \times n^2}$ be the convolution kernels produced by G_m and $h_c^s \in \mathbb{R}^{l_s \times l_s \times d_s}$ be the corresponding content layer, where d_s is the content feature dimension. We perform the spatially-adaptive convolution in the following steps. First, for each location (a, b) , we recover a 2D convolution kernel $\mathcal{K}^s(a, b) \in \mathbb{R}^{n \times n}$ from its flattened form $w^s(a, b)$. Then, we convolve $\mathcal{K}^s(a, b)$ with the content local patch $\mathcal{P}^s(a, b) \in \mathbb{R}^{n \times n \times d_s}$ on h_c^s to produce an intermediate content representation $\tilde{h}_c^s(a, b)$ for the location (a, b) :

$$\tilde{h}_c^s(a, b) = \mathcal{K}^s(a, b) * \mathcal{P}^s(a, b). \quad (1)$$

In Sec 4.6, we show the effectiveness of our approach for capturing motion at different scales. Moreover, we handle larger motion by adding adaptive convolutions to

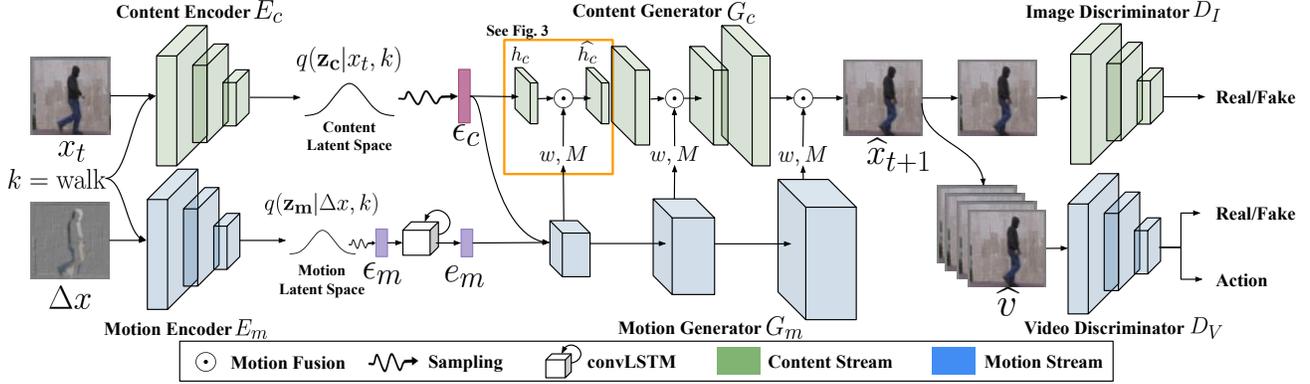


Figure 2: Our Two-Stream Variational Adversarial Network learns to generate the next frame \hat{x}_{t+1} from the current frame x_t . At training time, Content and Motion Encoders (E_c and E_m) generate latent distributions $q(\mathbf{z}_c|x_t, k)$ and $q(\mathbf{z}_m|\Delta x, k)$ by viewing x_t and the difference map Δx between x_{t+1} and x_t , respectively. From latent distributions, we sample the content latent vector ϵ_c and ϵ_m and employ a convLSTM to get e_m encoding the history of temporal dynamics. Content and Motion Generators (G_c and G_m) take in ϵ_c and e_m and then decode a content hidden layer h_c and motion kernels w along with a motion mask M at each scale. h_c , w and M are served as inputs to our defined multi-scale motion fusion (Fig. 3). Image and Video Discriminators (D_I and D_V) encourage the model to generate both realistic content and motion.

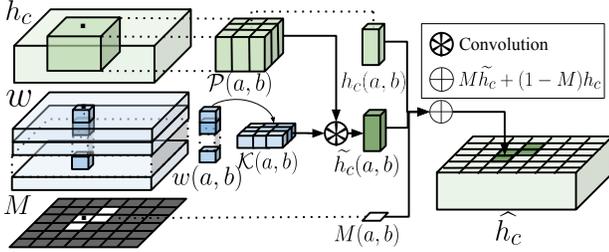


Figure 3: Motion fusion for location (a, b) on its local window $P(a, b)$, at scale s . Take kernel size $n = 3$ for example. h_c is the content hidden layer from G_c . w and M are convolutional kernels and the motion mask from G_m . We first recover the 2D kernel $\mathcal{K}(a, b)$ from its flattened form $w(a, b)$. Then we compute an intermediate content feature $\tilde{h}_c(a, b)$ by convolving $P(a, b)$ with $\mathcal{K}(a, b)$. Finally, we update h_c with \tilde{h}_c guided by M to get the new hidden layer \hat{h}_c .

smaller layers while [29] handles it by increasing the kernel size. Since the number of parameters to store and model is quadratic in the kernel size, we significantly reduce the memory usage and the model complexity by leveraging small kernels ($n = 3$ or 5) for all layers in our multi-scale fusion framework.

To focus G_m 's attention on learning the motion of scale s in the regions where it actually happens, we predict a motion mask $M^s \in \mathbb{R}^{l_s \times l_s}$ along with w^s , to identify such regions. Each entry $M^s(a, b)$ is in $[0, 1]$. We generate the new content map \hat{h}_c^s with fused motion from $h_c^s(a, b)$ and $\tilde{h}_c^s(a, b)$ guided by $M^s(a, b)$:

$$\hat{h}_c^s(a, b) = M^s(a, b)\tilde{h}_c^s(a, b) + (1 - M^s(a, b))h_c^s(a, b). \quad (2)$$

In addition to focusing the model's attention, motion mask helps preserve the pixel value in the refinement by simply deactivating $M^s(a, b)$. In contrast, [29] requires G_m to

learn a 2D kernel \mathcal{K} with a special pattern for pixel preservation. In Sec. 4.6, we show that our masks activate the correct areas at different scales and achieve better motion generation and background preservation.

3.3. Learning

Taking advantage of our disentangled content and motion streams, we propose an alternating dual-task learning scheme to learn each stream separately and thus more accurately and effectively. Specifically, the content stream is learned via image reconstruction, while the motion stream is learned via video reconstruction. We alternate training, such that each stream is trained while the other is fixed, and use both VAE and GAN losses to optimize each stream.

Content Learning. We learn our content stream solely by reconstructing the current frame x , without the disturbance of motion modeling, and reach compatible performance with image generation methods. We optimize the content VAE loss and image-level GAN loss. The VAE loss includes a reconstruction loss for the current frame and a KL-divergence between the prior content latent distribution $p(\mathbf{z}_c|k)$ and the posterior distribution $q(\mathbf{z}_c|x, k)$:

$$\min_{E_c, G_c} \mathcal{L}_{VAE-c} = \lambda_1 \mathcal{L}_2(\hat{x}, x) + \lambda_2 KL(q(\mathbf{z}_c|x, k) || p(\mathbf{z}_c|k)). \quad (3)$$

Image-level GAN training uses real frames x as positive examples, and treats the generated images x_p and \hat{x} sampled from $p(\mathbf{z}_c|k)$ and $q(\mathbf{z}_c|x, k)$ as negative examples following Larsen *et al.* [23]:

$$\max_{G_c} \min_{D_I} \mathcal{L}_{GAN-c} = \log [1 - D_I(x)] + \log [D_I(\hat{x})] + \log [D_I(x_p)]. \quad (4)$$

Motion Learning. We train the Motion Stream to generate the continuous motion for the next 10 frames after observing

the first frame. During training, the Motion Stream reconstructs the whole sequence recurrently. For every time step t , it reconstructs the current frame \hat{x}_t from $\{x_1, x_2, \dots, x_t\}$.

Similarly to content learning, we optimize the motion stream with a motion VAE loss and a video-level GAN loss. In addition to the video reconstruction loss and KL-divergence between the prior motion latent distribution $p(\mathbf{z}_m|k)$ and the posterior distribution $q(\mathbf{z}_m|\Delta x_t, k)$ where $\Delta x_t = x_t - x_{t-1}$, the VAE loss also contains a \mathcal{L}_2 loss between the refined content hidden layer of the previous frame $(\hat{h}_c^s)_{t-1}$ and the content layer of the current frame $(h_c^s)_t$. Thus, the total motion VAE loss is:

$$\begin{aligned} \min_{E_m, G_m} \mathcal{L}_{VAE-m} = & \lambda_3 \sum_{s,t} \mathcal{L}_2 \left((\hat{h}_c^s)_{t-1}, (h_c^s)_t \right) + \lambda_4 \mathcal{L}_2(\hat{v}, v) \\ & + \lambda_5 \sum_t KL(q(\mathbf{z}_m|\Delta x_t, k) || p(\mathbf{z}_m|k)). \end{aligned} \quad (5)$$

The video-level GAN loss contains the vanilla GAN loss \mathcal{L}_{GAN-m} and an auxiliary classification loss \mathcal{L}_{cls} [30]. The vanilla GAN takes real videos v as positive examples, and treats the generated videos v_p and \hat{v} sampled from $p(\mathbf{v}_c|k)$ and $q(\mathbf{v}_c|x, k)$ as negative examples. Apart from GAN training, both the generator and the discriminator minimize a classification loss on action labels to generate action-related motion:

$$\max_{G_m} \min_{D_V} \mathcal{L}_{GAN-m} = \log [1 - D_V(v)] + \log [D_V(\hat{v})] + \log [D_V(v_p)], \quad (6)$$

$$\min_{G_m, D_V} \mathcal{L}_{cls}(v) + \mathcal{L}_{cls}(\hat{v}) + \mathcal{L}_{cls}(v_p). \quad (7)$$

We provide implementation details (hyper-parameters, model architecture, etc.) in the supplementary material.

3.4. Generating a Video at Test Time

While learning relies on observing the ground truth, at test time, generating a video begins from sampling in the latent space. Given an action class k , we generate the first frame from a randomly sampled content vector $\epsilon_c \sim p(\mathbf{z}_c|k)$, and then the following frames from the content embedding of the last frame as well as the current motion embedding e_m computed by the convLSTM recurrently. To generate e_m at each time step, the convLSTM updates the accumulated motion history with an extra motion vector $\epsilon_m \sim p(\mathbf{z}_m|k)$ containing the current potential motion.

4. Experiments

4.1. SynAction Dataset

Existing synthetic datasets used in [36, 43, 39] only contain rigid motions (e.g. linear motion in any direction) which are not challenging for deep neural networks. To uncover the model’s ability to generate complex and action-related motion, we build a large-scale synthetic human action dataset, *SynAction Dataset*, with a powerful game engine *Unity*.

Dataset	SynAction	Moving MNIST	Shape Motion	Weizmann	MUG
Type	Synthetic	Synthetic	Synthetic	Real	Real
Videos	6000	10000	4000	90	882
Action Models	120	-	2	-	-
Action Classes	20	1	1	10	6
Scenes	150	10	1	9	52

Table 1: Compared to other datasets frequently used in video generation, our SynAction Dataset has a larger variation of content and motion and is thus more challenging.

The dataset contains 120 unique non-rigid human action models from the Mixamo motion library across 20 action classes. Every action is akin to real human actions but easy to distinguish from other actions. The dataset is further varied between 10 different actors and 5 different backgrounds. Table 1 shows that SynAction has more variation in content and motion than existing synthetic and standard real-world datasets for video generation.

We provide each video with four different annotations: actor identity, action class, background and viewpoint. In this paper, we only use the action class to generate videos.

4.2. Other Datasets

In addition to our proposed SynAction Dataset, we evaluate our model on three existing datasets: Weizmann Human Action [12], MUG Facial Expression [1] and VoxCeleb [28]. The Weizmann Human Action Dataset contains 90 videos of 9 actors performing 10 different actions. MUG Facial Expression Dataset contains 882 videos with 52 actors performing 6 different facial expressions. From the VoxCeleb Dataset, we form a training set containing 15184 videos of 186 people speaking.

With these three datasets, we cover a large range of motion, from large human actions (e.g. running, jumping) to subtle facial movements (e.g. happiness, disgust, speaking) as well as head movements (e.g. nodding, turning) and include both periodic and non-periodic motion.

4.3. Evaluation Metrics

Quantitative evaluation of generative models remains a challenging problem, and there is no consensus on the measurement which best evaluates the realism and diversity of the generated results. Thus, instead of just relying on a single measurement, we utilize four different metrics to examine both the realism and diversity of generated motion: Inception Score (IS) [35], Inter-Entropy $H(y)$ [15] and Intra-Entropy [15] $H(y|v)$, where v is the video for evaluation and y is the action predicted by a classifier. Because all these metrics utilize a pre-trained classifier for evaluation, we train a classifier separately on each dataset and show its performance by computing the same metrics on the test set, which only consists of real videos. We call these values the *Experimental Bound*. To make a fair comparison, we compute metrics on 10-frame video clips generated by each model.

Dataset	Weizmann (# action = 10)			MUG (# action = 6)			Syn-Action (# action = 20)		
Metric	H(y)	H(y v)	IS	H(y)	H(y v)	IS	H(y)	H(y v)	IS
MoCoGAN [39]	4.38	0.31	58.52	1.78	0.17	5.03	2.90	0.36	12.75
VideoVAE [15]	4.37	0.11	70.10	-	-	-	-	-	-
SGVAN	4.34	0.04	73.73	1.79	0.13	5.29	2.98	0.18	16.43
TwoStreamVAN (-M)	4.31	0.29	55.99	1.79	0.11	5.32	2.97	0.15	16.79
TwoStreamVAN	4.40	0.05	77.11	1.79	0.09	5.48	2.99	0.09	18.27
<i>Exp Bound</i>	4.50	0.01	88.94	1.79	0.01	5.91	3.00	0.01	19.85
<i>Math Bound</i>	4.50	0.00	90.00	1.79	0.00	6.00	3.00	0.00	20.00

Table 2: Quantitative Results on Weizmann, MUG and Syn-Action Datasets. For IS, the higher value is better; while for $H(y|v)$, the lower value is better. Compared with all baselines, our TwoStreamVAN model achieves the best results on most metrics.

4.4. Baselines

We compare against two existing methods to show our model’s superiority in generating videos of a given action: MoCoGAN¹ [39] and VideoVAE² [15], which are the current state-of-the-art.

We also design several ablated variants of TwoStreamVAN to examine key components of our model:

SGVAN adopts a single generator to generate a single frame from the disentangled content and motion vectors, keeping all parts of the model the same. This comparison evaluates the contribution of the parallel G_c and G_m .

TwoStreamVAN(-M) applies motion fusion to content hidden layers at multiple scales without the guidance of motion masks. This comparison helps us to examine the effectiveness of motion masks.

4.5. Results

Quantitative Results. We compute quantitative metrics on the results of all baselines and our TwoStreamVAN (see Table. 2) on Weizmann, MUG and SynAction Datasets. We train a normal action classifier on MUG and SynAction Datasets, and train a classifier to distinguish each actor-action pair on Weizmann to compare with VideoVAE.

TwoStreamVAN improves the Inception Scores of MoCoGAN by 32%, 9% and 43% on Weizmann, MUG and SynAction Datasets respectively. Despite that VideoVAE receives the first frame and the actor identity as additional inputs, our model still results in a better IS value on Weizmann. Meanwhile, TwoStreamVAN achieves both higher $H(y)$ and lower $H(y|v)$, indicating that it generates more diverse and more realistic videos than either MoCoGAN or VideoVAE.

Compared to our ablated models SGVAN and TwoStreamVAN(-M), our model pushes all metrics closer to their bounds. These results reveal that our full model benefits from key components in our design, namely the disentangled generators and the guidance of the motion mask in fusion.

¹We use categorical MoCoGAN implemented by its authors.

²Due to VideoVAE’s non-public implementation, we only compare with quantitative results on Weizmann Dataset reported in the paper.

Qualitative Results. We visualize videos generated by TwoStreamVAN and MoCoGAN. For Weizmann, MUG and SynAction datasets, we provide 4 generated videos from TwoStreamVAN and MoCoGAN of the given action (Fig. 4). Our TwoStreamVAN model succeeds in generating more accurate and fine-grained motion for different actions.

To evaluate the quality of content generation, we randomly sample 10 generated frames from TwoStreamVAN and MoCoGAN’s results respectively (Fig. 4). In comparison, TwoStreamVAN yields better content generation, with few severe distortions or bizarre artifacts across all three datasets, which demonstrates that TwoStreamVAN reduces the content deterioration significantly.

User Study on SynAction. To further test the ability of TwoStreamVAN and MoCoGAN to handle more diverse human action videos, we conduct user studies on SynAction via AMTurk [5]. In the pairwise comparison, **88% ± 0.45%** / **81% ± 0.71%** users think TwoStream generates better videos before/after knowing the ground truth action, showing that the superiority of our TwoStreamVAN to generate the clear and consistent motion and yield visually satisfying videos.

Experiment on large-scale VoxCeleb. To extensively test the ability of our model to generate real-world videos, we experiment on a new large-scale speech dataset VoxCeleb [28]. Compared to the standard Weizmann and MUG Datasets, VoxCeleb is more challenging due to the wide variety of content and motion in its huge amount of videos. Because VoxCeleb has no action label, we generate videos only from a noise vector. We compare to MoCoGAN via a user study on AMTurk and the qualitative visualization.

In the pairwise comparison, **82.55% ± 0.53%** users prefer TwoStreamVAN to MoCoGAN, evidence that our generated videos are more visually-pleasing. In the qualitative comparison (see Fig 5), TwoStreamVAN captures both lip and head movements, whereas MoCoGAN suffers from mode collapse and only generates videos without head movements. Our model precisely reproduces the subtle motion related to speaking, while MoCoGAN outputs exaggerated and unrealistic movements (e.g. teeth suddenly appear during speech). Moreover, TwoStreamVAN generates better-looking faces than MoCoGAN.

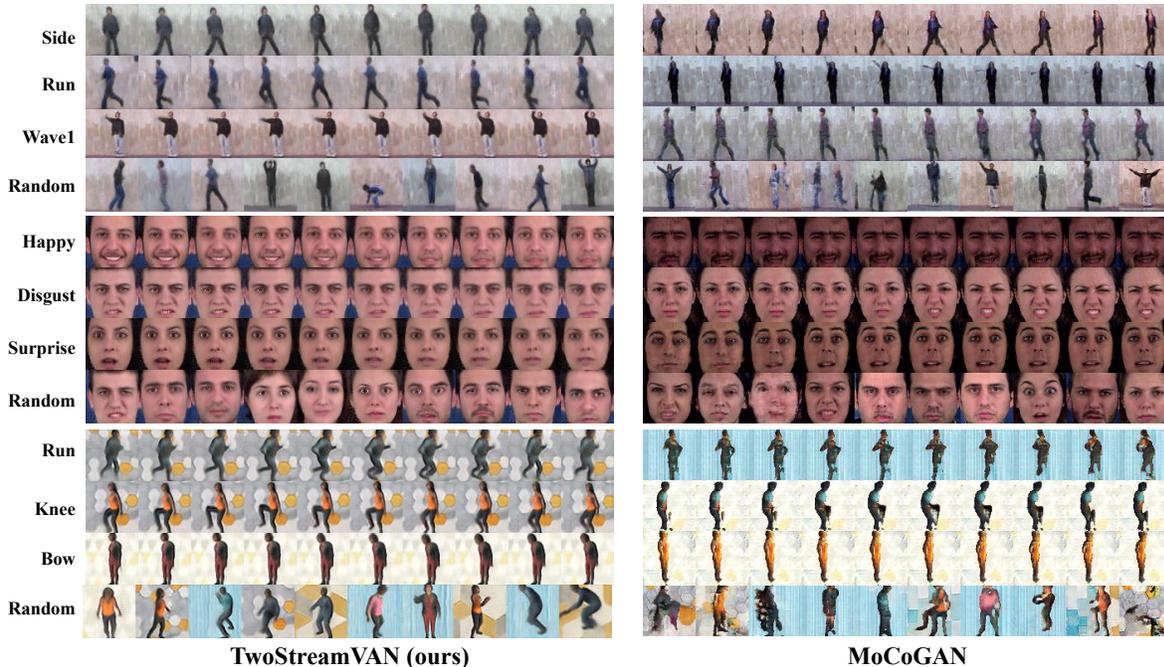


Figure 4: We provide 3 generated videos of TwoStreamVAN and MoCoGAN on Weizmann, MUG and our new SynAction datasets. Our model generates clear and consistent motion. We further randomly sample 10 generated frames for each model. TwoStreamVAN results in better content quality with fewer distortions in every single frame. More qualitative results are presented in the supplementary material.

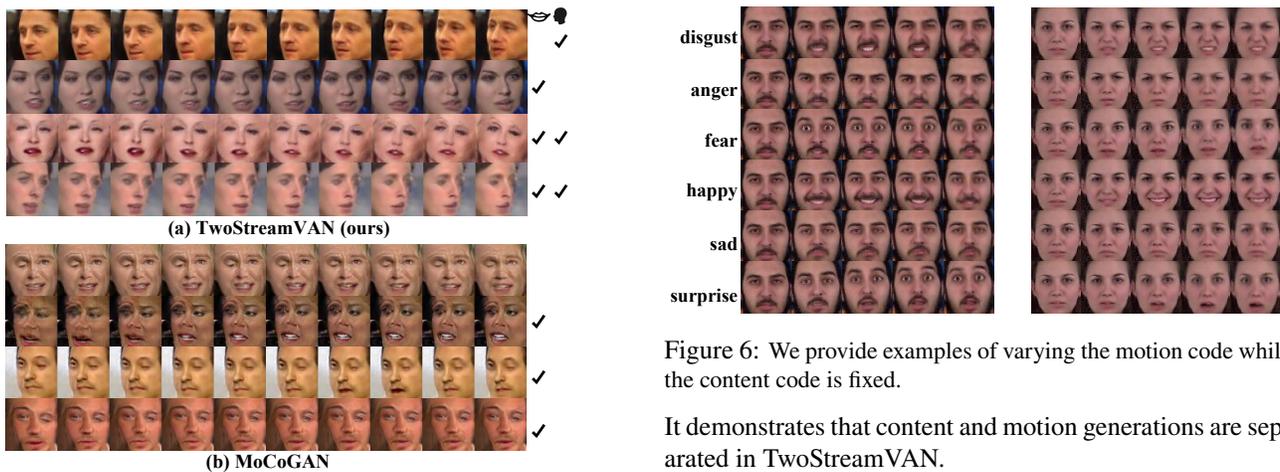


Figure 5: Qualitative comparison between TwoStreamVAN and MoCoGAN. TwoStreamVAN captures both lip and head movements, while MoCoGAN only generates non-moving heads. More examples are provided in the supplementary material.

4.6. Ablation Studies

Disentanglement Study To show that the disentanglement of content and motion really happens, we experiment on MUG dataset. With the fixed content code ϵ_c , we generate videos using different motion codes (see Fig 6). We classify actors and actions in the generated videos using pre-trained classifiers. 95.9% of our generated videos preserve the actor in the fixed content code and 95.6% of videos generate the correct action defined in the different motion code.



Figure 6: We provide examples of varying the motion code while the content code is fixed.

It demonstrates that content and motion generations are separated in TwoStreamVAN.

Multi-scale v.s. Single-scale Motion Fusion To examine the effectiveness of multi-scale motion fusion, we train four TwoStreamVAN models on the Weizmann Human Action Dataset, in which we apply fusions at 1, 2, 3 and 4 layers (scales) respectively. We add new fusion layers from the highest resolution to the lowest resolution. In each fusion, we implement motion kernels with a fixed size $n = 5$. Moreover, we train a model where we apply large motion kernels with $n = 17$ on the output image from the Content Stream to imitate the single-scale fusion in [29].

In Table. 3, it is not surprising that the performance drops when we reduce n from 17 to 5 with a single fusion on the image. As we increase fusion layers, IS value, $H(y)$ and

# layers	n	IS \uparrow	H(y) \uparrow	H(y v) \downarrow	H(y v) (run) \downarrow
1	17	75.89	4.38	0.05	0.104
1	5	74.95	4.39	0.07	0.166
2	5	74.57	4.38	0.07	0.096
3	5	77.83	4.40	0.05	0.062
4	5	77.11	4.40	0.05	0.059

Table 3: Quantitative results of models varying in the kernel size n and the number of scales to apply fusion. For single fusion, the performance drops by reducing the kernel size. After applying fusion at multiple scales, the performance recovers and finally outperforms the single fusion with large kernels.

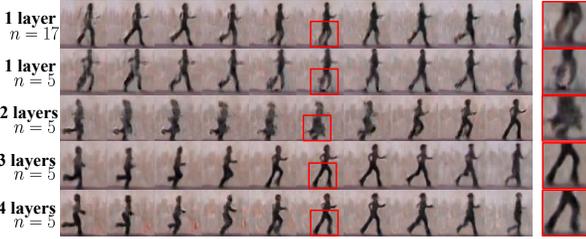


Figure 7: Qualitative comparison among similar sequences generated by single fusion with large kernels and multi-scale motion fusion with small kernels. When we apply fusion at 3 or 4 scales, our model generates the sharpest and clearest outlines of the runner’s leg among all models.

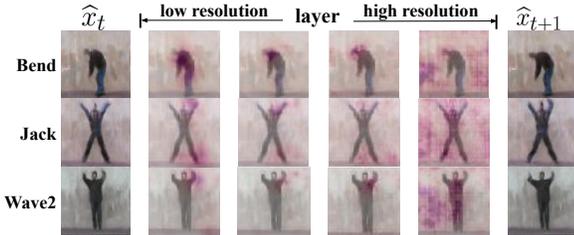


Figure 8: We overlay the motion mask at each scale with the current frame \hat{x}_t . Motion masks at lower resolutions are only activated on large-motion areas, while a large area in the motion mask at the highest resolution is activated to tackle small changes between neighboring frames.

$H(y|v)$ recover and finally outperform those of the model with a single large fusion step on the full resolution image.

To further analyze multi-scale fusion, we measure $H(y|v)$ of large motion, e.g. running (in Table. 3). The more layers such fusion is applied to, the lower $H(y|v)$ is, indicating that more realistic large motion is generated. We pick similar videos (Fig. 7) generated by different models and zoom in on the actor’s legs, where the largest motion happens. When we only apply fusion with kernel size $n = 5$ at the highest 2 resolutions, the model fails to tackle the large motion around the legs and generates blobs. After increasing the number of fusion layers, it finally generates an even sharper outline than the model using the single large fusion, showing the benefits of our multi-scale mechanism.

Visualization of Motion Masks We visualize the motion mask at each layer by overlaying it with the current frame \hat{x}_t



Figure 9: Qualitative Comparison of TwoStreamVAN(–M) and TwoStreamVAN. Motion masks improve generation in a few areas (background for Weizmann and SynAction and eyes for MUG).

(Fig. 8) to show it activates the correct area. For low resolutions, the mask is only activated at large-motion areas, e.g. torso for bending, arms and legs for jumping-jacks and arms for waving. For the highest resolution, the activation covers the background area to overcome some small changes, e.g. lighting changes or small camera movements. This correct activation focuses the model’s attention on learning the motion in these areas and boosts quantitative performance (see TwoStreamVAN (–M) v.s. TwoStreamVAN in Sec. 4.5).

We further visualize frames generated from these two models (Fig. 9). We observe that TwoStreamVAN(–M) does a worse job in small-motion areas. On Weizmann and SynAction, it messes up background patterns. On MUG, it generates unexpected brick patterns around eyes. Our full TwoStreamVAN does not suffer from these problems with the help of motion masks. This observation is consistent with our claim (in Sec. 3.2) that motion masks help to preserve static pixel values during the generation.

5. Conclusion

In this paper, we propose a novel Two-Stream Variational Adversarial Network to improve motion modeling in video generation. To generate motion efficiently, we decompose content and motion in the generation phase and fuse them via a novel multi-scale mechanism. Combined with the dual-task learning scheme, our disentangled generative network overcomes the common content deterioration and further benefits motion modeling. We propose a large-scale synthetic human action dataset *SynAction* to evaluate the motion modeling in video generation. Our model significantly outperforms the current state-of-the-art works across Weizmann, MUG, VoxCeleb and our SynAction datasets in quantitative and qualitative (user study) evaluations. In the future work, we hope to explore video generation controlled via textual description or interactive user manipulation.

6. Acknowledgements

This work is supported by DARPA LWLL Project. It reflects the opinions and conclusions of its authors, but not the funding agents.

References

- [1] N. Aifanti, C. Papachristou, and A. Delopoulos. The mug facial expression database. In *Image analysis for multimedia interactive services (WIAMIS), 2010 11th international workshop on*, pages 1–4. IEEE, 2010. 2, 5
- [2] M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine. Stochastic variational video prediction. *arXiv preprint arXiv:1710.11252*, 2017. 2
- [3] M. J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer vision and image understanding*, 63(1):75–104, 1996. 2
- [4] T. Brox, A. Bruhn, N. Papenbergh, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision*, pages 25–36. Springer, 2004. 2
- [5] M. Buhrmester, T. Kwang, and S. D. Gosling. Amazon’s mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science*, 6(1):3–5, 2011. 6
- [6] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016. 1, 2
- [7] S. Corazza and E. Gambaretto. Real time generation of animation-ready 3d character models, Feb. 25 2014. US Patent 8,659,596. 2
- [8] E. Denton and R. Fergus. Stochastic video generation with a learned prior. *arXiv preprint arXiv:1802.07687*, 2018. 2
- [9] E. L. Denton et al. Unsupervised learning of disentangled representations from video. In *Advances in neural information processing systems*, pages 4414–4423, 2017. 2
- [10] C. Finn, I. Goodfellow, and S. Levine. Unsupervised learning for physical interaction through video prediction. In *Advances in neural information processing systems*, pages 64–72, 2016. 2
- [11] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 2
- [12] L. Gorelick, M. Blank, E. Shechtman, M. Irani, and R. Basri. Actions as space-time shapes. *Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2247–2253, December 2007. 2, 5
- [13] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *International Conference on Machine Learning*, 2015. 2
- [14] I. Gulrajani, K. Kumar, F. Ahmed, A. A. Taiga, F. Visin, D. Vazquez, and A. Courville. Pixelvae: A latent variable model for natural images. *International Conference on Learning Representations*, 2017. 2
- [15] J. He, A. Lehrmann, J. Marino, G. Mori, and L. Sigal. Probabilistic video generation using holistic attribute control. *European Conference on Computer Vision*, 2018. 1, 2, 5, 6
- [16] S. Hong, D. Yang, J. Choi, and H. Lee. Inferring semantic layout for hierarchical text-to-image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7986–7994, 2018. 1
- [17] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 1
- [18] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*, 2018. 1
- [19] H. R. Kaviani and S. Shirani. Frame rate upconversion using optical flow and patch-based reconstruction. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(9):1581–1594, 2016. 2
- [20] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2014. 2
- [21] N. Kodali, J. Abernethy, J. Hays, and Z. Kira. On convergence and stability of gans. *arXiv preprint arXiv:1705.07215*, 2017. 1
- [22] R. Krishnamurthy, J. W. Woods, and P. Moulin. Frame interpolation and bidirectional prediction of video using compactly encoded optical-flow fields and label fields. *IEEE transactions on circuits and systems for video technology*, 9(5):713–726, 1999. 2
- [23] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther. Autoencoding beyond pixels using a learned similarity metric. *International Conference on Machine Learning*, 2016. 2, 4
- [24] A. X. Lee, R. Zhang, F. Ebert, P. Abbeel, C. Finn, and S. Levine. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018. 2
- [25] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey. Adversarial autoencoders. *International Conference on Learning Representations*, 2016. 2
- [26] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2017. 1
- [27] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *International Conference on Learning Representations*, 2016. 1
- [28] A. Nagrani, J. S. Chung, and A. Zisserman. Voxceleb: a large-scale speaker identification dataset. In *INTERSPEECH*, 2017. 2, 5, 6
- [29] S. Niklaus, L. Mai, and F. Liu. Video frame interpolation via adaptive convolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, page 3, 2017. 3, 4, 7
- [30] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2642–2651, 2017. 1, 2, 5
- [31] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *International Conference on Learning Representations*, 2016. 1, 2

- [32] A. Ranjan and M. J. Black. Optical flow estimation using a spatial pyramid network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, page 2. IEEE, 2017. 2
- [33] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text-to-image synthesis. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016. 1
- [34] M. Saito, E. Matsumoto, and S. Saito. Temporal generative adversarial nets with singular value clipping. In *IEEE International Conference on Computer Vision (ICCV)*, volume 2, page 5, 2017. 1, 2, 3
- [35] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2234–2242, 2016. 5
- [36] N. Srivastava, E. Mansimov, and R. Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852, 2015. 1, 2, 5
- [37] D. Sun, S. Roth, and M. J. Black. A quantitative analysis of current practices in optical flow estimation and the principles behind them. *International Journal of Computer Vision*, 106(2):115–137, 2014. 2
- [38] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8934–8943, 2018. 2
- [39] S. Tulyakov, M.-Y. Liu, X. Yang, and J. Kautz. Mocogan: Decomposing motion and content for video generation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 3, 5, 6
- [40] R. Villegas, J. Yang, S. Hong, X. Lin, and H. Lee. Decomposing motion and content for natural video sequence prediction. *ICLR*, 2017. 1, 2
- [41] C. Vondrick, H. Pirsaviash, and A. Torralba. Generating videos with scene dynamics. In *Advances In Neural Information Processing Systems*, pages 613–621, 2016. 1, 2, 3
- [42] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015. 3
- [43] T. Xue, J. Wu, K. Bouman, and B. Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2016. 2, 3, 5
- [44] X. Yan, J. Yang, K. Sohn, and H. Lee. Attribute2image: Conditional image generation from visual attributes. In *European Conference on Computer Vision*, pages 776–791. Springer, 2016. 2
- [45] F. Yu, A. Seff, Y. Zhang, S. Song, T. Funkhouser, and J. Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 1
- [46] Z. Zhou, H. Cai, S. Rong, Y. Song, K. Ren, W. Zhang, Y. Yu, and J. Wang. Activation maximization generative adversarial nets. *International Conference on Learning Representations*, 2018. 2
- [47] X. W. Ziwei Liu, Ping Luo and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015. 1