Fast-Vid2Vid: Spatial-Temporal Compression for Video-to-Video Synthesis

Long Zhuo¹[®], Guangcong Wang²[®], Shikai Li³[®], Wayne Wu^{1,3}[®], and Ziwei Liu^{2⊠}[®]

¹ Shanghai AI Laboratory
² S-Lab, Nanyang Technological University
³ SenseTime Research
zhuolong@pjlab.org.cn {guangcong.wang, ziwei.liu}@ntu.edu.sg
lishikai@sensetime.com wuwenyan0503@gmail.com



Fig. 1. Fast-Vid2Vid. Our proposed Fast-Vid2Vid accelerates the video-to-video synthesis and generates photo-realistic videos more efficiently compared to the original Vid2Vid model.

Abstract. Video-to-Video synthesis (Vid2Vid) has achieved remarkable results in generating a photo-realistic video from a sequence of semantic maps. However, this pipeline suffers from high computational cost and long inference latency, which largely depends on two essential factors: 1) network architecture parameters, 2) sequential data stream. Recently, the parameters of image-based generative models have been significantly compressed via more efficient network architectures. Existing methods mainly focus on slimming network architectures and ignore the size of the sequential data stream. Moreover, due to the lack of temporal coherence, image-based compression is not sufficient for the compression of

the video task. In this paper, we present a spatial-temporal compression framework, **Fast-Vid2Vid**, which focuses on data aspects of generative models. It makes the first attempt at time dimension to reduce computational resources and accelerate inference. Specifically, we compress the input data stream spatially and reduce the temporal redundancy. After the proposed spatial-temporal knowledge distillation, our model can synthesize key-frames using the low-resolution data stream. Finally, Fast-Vid2Vid interpolates intermediate frames by motion compensation with slight latency. On standard benchmarks, Fast-Vid2Vid achieves around real-time performance as 20 FPS and saves around $8 \times$ computational cost on a single V100 GPU. Code and models are publicly available¹.

Keywords: Video-to-Video Synthesis, GAN Compression

1 Introduction

Video-to-video synthesis (Vid2Vid) [42] targets at synthesizing a photo-realistic video given a sequence of semantic maps as input. A wide range of applications are derived from this task, such as face-talking video generation (Sketch2Face) [42,41], driving video generation (Segmentation2City) [42,41] and human pose transferring generation (Pose2Body) [4,25,51]. With the advance of Generative Adversarial Networks (GANs) [14], Vid2Vid models [42,41] have made significant progress in video quality. However, these approaches need largescale computational resources to yield the results, and they are computationally prohibitive and environmentally unfriendly. For example, the standard Vid2Vid [42] consumes 2066 G MACs for generating each frame, which is $500 \times$ more than ResNet-50 [16]. Recent studies demonstrate that lots of recognition compression approaches have been successfully extended to image-based GAN compression methods [1,6,10,24,29,27]. Can we directly employ these existing image-based GAN compression methods to achieve promising Vid2Vid compression models?

In the literature, image-based GAN compression methods can be roughly categorized into three groups, including knowledge distillation [1, 6, 24, 29, 2], network pruning [29,40], and neural architecture search (NAS) [27,13,10,12,28]. They focus on obtaining a compact network by cutting the network architecture parameters of the original network. However, the input data, another factor that significantly affects the inference speed of a deep neural network, has been ignored by the existing GAN compression methods. Moreover, since they are image-based synthesis tasks, they do not consider redundant temporal information hidden in neighbor frames of a video. Therefore, directly applying image-based compression models to Vid2Vid synthesis is difficult to achieve the desired results. In this work, we aim to compress the input data stream while maintaining the well-designed network parameters, and generate the photo-realistic results for Vid2Vid synthesis. Furthermore, we make an initial attempt at removing temporal redundancy to accelerate Vid2Vid model.

 $\mathbf{2}$

¹ Project page: https://fast-vid2vid.github.io/

Code and models: https://github.com/fast-vid2vid/fast-vid2vid

There are three critical challenges for Vid2Vid compression. First, the typical Vid2Vid model [42] consists of several encoder-decoders to capture both spatial and temporal features. It is difficult to reduce the parameters from such a complicated structure due to the intricate connections between these encoders and decoders. Second, it is a challenge to compress the input data stream and achieve decent performance for GAN generation since the perceptual fields of GAN are much more erratic than image recognition. Third, transferring knowledge from a teacher model to a student model temporally is challenging to align with the spatial knowledge distillation, as the temporal knowledge is implicitly hidden within adjacent frames and is more difficult to capture than the spatial knowledge.

To address the above issues, in this paper, we propose a novel spatialtemporal compression framework for Vid2Vid synthesis, named **Fast-Vid2Vid**. As shown in Fig. 2, we reduce the computational resources by only compressing the input data stream through Motion-Aware Inference (MAI) without destroying the well-designed and complicated network parameters of the original Vid2Vid model, which addresses challenge 1. For challenges 2 and 3, we propose a Spatial-Temporal Knowledge Distillation method (STKD) that transfers spatial and temporal knowledge from the original model to the student network using compressed input data. Our goal is to transfer the knowledge from large-size synthesized videos to small-size synthesized ones to make GAN robust enough to gain promising visual performance when the input data is compressed.

We first train a spatially low-demand generator by taking low-resolution sequences as input but generating the full-resolution sequences. We perform Spatial Knowledge Distillation (Spatial KD) that transfers the spatial knowledge from the original generator to the spatially low-demand generator to obtain highresolution information. Furthermore, we train a part-time generator by uniformly sampling video frames from sequences. We perform Temporal-aware Knowledge Distillation (Temporal KD) and distill the temporal knowledge of the original generator to the part-time student generator to obtain full-time motion information by introducing two losses, i.e., local temporal knowledge distillation loss and global temporal knowledge distillation loss. This design aims to capture the implicit knowledge in the time dimension.

To summarize, to the best of our knowledge, we make the first attempt to tackle the Vid2Vid compression problem at data aspects. On a single V100 GPU, Fast-Vid2Vid achieves 18.56 FPS ($6.1 \times$ acceleration) with $8.1 \times$ less computational cost on Sketch2Face, 24.77 FPS ($5.8 \times$ acceleration) with $8.3 \times$ less computational cost on Segmentation2City, and 21.39 FPS ($7.1 \times$ acceleration) with $9.3 \times$ less computational cost on Pose2Body. The main contributions of this paper are concluded as three-fold:

- We present Fast-Vid2Vid, an sequential data stream compression method in spatial and temporal dimensions to greatly accelerate the Vid2Vid model.
- We introduce Spatial KD that transfers knowledge from a teacher model with high-resolution input data to a student model with low-resolution input data to learn high-resolution information.

- 4 L. Zhuo et al.
- We propose Temporal KD to distill knowledge from a full-time teacher model to a part-time student model. A new temporal knowledge distillation loss globally is further presented to capture the time-series correlation.

2 Related Work

Video-to-Video Synthesis. Video-to-video synthesis (Vid2Vid) is a computer vision task that generates a photo-realistic sequence using the corresponding semantic sequence. Based on high-resolution image-based synthesis [43], Wang *et al.* [42] developed a standard Vid2Vid synthesis model by introducing temporal coherence. Few-shot Vid2Vid model [41] further extended to a few-shot training version. Recently, Vid2Vid has been successfully extended to a wide range of video generation tasks, including video super-resolution [35,7,46], video in-painting [52,47], image-to-video synthesis [36,37] and human pose-to-body synthesis [4,11,51,25]. Most of these methods exploited temporal information to improve the performance of generated videos. However, they do not focus on Vid2Vid synthesis compression but on better visual performance.

Model Compression. Model compression aims at reducing superfluous parameters of deep neural networks to accelerate inference. In the computer vision task, lots of model pruning approaches [15,26,22,30,17,50,40] have greatly cut the weights of neural networks. Commonly, compression methods [23,17,26,18] reduced the unnecessary channels with low activations. GAN compression has been proved by [49] that it is far more difficult than normal CNN compression. A content-aware approach [29] was proposed to use salient regions to identify specific redundancy for GAN pruning. Wang *et al.* [40] reduced the redundant weights by NAS. Notably, the mentioned methods simplify the network structure but ignore the input information. Furthermore, these approaches do not consider the essential temporal coherence for video-based GAN, and thus yield sub-optimal results for compressing Vid2Vid models.

Knowledge Distillation. Knowledge Distillation aims to make a student network imitate its teacher. Hinton et al. [20] proposed an effective framework for model distillation in classification. Knowledge distillation has been widely used in recognition models [5,6,31,32,48]. Recently, lots of response-based knowledge distillation methods [1,6,10,2] were proposed for image-based GAN compression. For example, Jin *et al.* [24], developed distillation techniques from [27], used a global kernel alignment module to gain more potential information. Liu et al. [29] utilized a salient mask to guide the knowledge distillation process based on the norm. These methods only address image-based knowledge distillation, and thus only spatial knowledge is exploited, and they do not consider movements. Different from these spatial-aware knowledge distillation, we consider both spatial information and temporal information into knowledge distillation, which tailors for Vid2Vid model compression. Most recently, Feng et al. [9] have presented a resolution-aware knowledge distillation method that ignores the network parameters and compresses the input information for image recognition. In our work, it is the first time to introduce the input data compression for GAN synthesis.



Fig. 2. The pipeline of our Fast-Vid2Vid. It maintains the same amount of parameters as the original generator but compresses the input data in space and time dimensions.

3 Fast-Vid2Vid

3.1 A Revisit of GAN Compression

The function of a deep neural network (DNN) can be written as f(X) = W * X, where W denotes the parameters of the networks, * represents the operation of DNN and X denotes the input data. Two essential factors accounting for computational cost are the parameters and the input data. Existing GAN compression methods [1,6,10,24,29,27] have intended to cut the computational cost by reducing the parameters of network structures. However, the network structures of GAN for specific tasks are carefully designed and it would cause poor visual results if the network parameters are cut arbitrarily. Another way to reduce computational cost is compressing the input data. In this work, we seek for compressing the input data instead of the parameters of well-designed networks. To the best of our knowledge, there is little literature working on compressing data for GAN compression.

3.2 Overview of Fast-Vid2Vid

The typical Vid2Vid framework [42] takes a sequence of semantic maps $\{X\}_0^T \in \mathbb{R}^{T \times H \times W}$ with T frames and the initial real frames as input and predicts a photo-realistic video sequence $\{Y\}_0^T \in \mathbb{R}^{T \times H \times W}$. H and W denotes the height and weight of each frame. The vanilla inference of the Vid2Vid model (the full-time teacher generator) that utilizes a full-size sequential input data stream is a consecutive process that synthesizes a video sequence frame by frame. Considering both image synthesis and temporal coherence, a Vid2Vid model often contains several encoder-decoders to capture spatial-temporal cues, which are computationally prohibitive. In this paper, we propose a Fast-Vid2Vid compression framework, an input data compression method, to reduce the computational resources of the Vid2Vid framework in both space and time dimensions.

Fig. 2 illustrates the overview of the proposed method. Fast-Vid2Vid first replaces the resBlock of the original Vid2Vid generator [42] with decomposed



Fig. 3. Left: The proposed Spatial Knowledge Distillation (Spatial KD). Right: The proposed temporal loss for temporal global knowledge distillation.

convolutional block [21] to obtain a modern network architecture, which is similar to [27]. During knowledge distillation, we train a student generator using the compressed data and distill knowledge from the teacher generator by our proposed spatial-temporal knowledge distillation method (STKD). STKD, including spatial knowledge distillation (Spatial KD) and temporal knowledge distillation (Temporal KD), performs **spatial resolution compression** and **temporal sequential data compression**. After STKD, a part-time generator cooperating with motion compensation synthesizes a full-size sequence by **motion-aware inference (MAI)**.

3.3 Spatial Resolution Compression for Vid2Vid

To reduce the spatial input data, a straightforward way [9] is to predict the low-resolution results and re-size them into the full-resolution by a distortion algorithm. However, in our preliminary experiments, this method leads to severe artifacts since the distortion algorithm lacks high-frequency information and loses the important textures. Therefore, we make an adaptive change for Vid2Vid synthesis. We replace the downsampling layers with the normal convolution layers to generate the high-resolution results. The modified generator only takes the low-resolution semantic sequence $\{X\}_0^{'T} \in \mathbb{R}^{T \times h \times w}$ as the input, where $h \times w = \frac{1}{(2^d)^2} H \times W$, and d is the number of the modified downsampling layers. d is set to 1. In this way, we obtain a spatially low-demand generator.

Next, the spatially low-demand generator is required to learn high-frequency representation from the full-time teacher generator. We present a spatial knowledge distillation method (Spatial KD) to model high-frequency knowledge from the teacher net. Specifically, as shown in the left figure of Fig. 3, Spatial KD shrinks the margin between the low-resolution domain and the high-resolution domain to improve the performance of the student network. Spatial KD implicitly transfers spatial knowledge from the teacher net to the student net. Particularly, Spatial KD applies a knowledge distillation loss to mimic the visual features of the teacher net, and the spatial knowledge distillation loss L_{SKD} can be written as:

$$L_{SKD} = \frac{1}{T} \sum_{t=0}^{T} [MSE(Y_t, Y'_t) + L_{per}(Y_t, Y'_t)], \qquad (1)$$



Fig. 4. The proposed temporal-aware knowledge distillation (Temporal KD).

where t means the current timestamp, T is the total timestamps of the sequences, Y is the output sequence of the teacher net and Y' is the predicted sequence of the spatially low-demand generator. MSE represents a mean squared error between two frames, and L_{per} denotes a perceptual loss [42].

3.4 Temporal Sequential Data Compression for Vid2Vid

Each video sequence consists of dense video frames, which brings an enormous burden on computational devices. How to efficiently synthesize dense frames is a difficult yet important issue for lightweight Vid2Vid models.

In Section 3.3, we obtain a spatially low-demand generator. To ease the burden of generating dense frames for each video, we re-train the spatially lowdemand generator on sparse video sequences, which are uniformly sampled from dense video sequences. The sampling interval is randomly selected in each training iteration. The original Vid2Vid generator is regarded as a full-time teacher generator and the re-trained spatially low-demand generator is regarded as a part-time student generator. To distill the temporal knowledge from the fulltime teacher generator to the part-time student generator, we propose a local temporal-aware knowledge distillation method and a global temporal-aware knowledge distillation method for temporal distillation, as shown in Fig. 4.

Both the full-time teacher generator and the part-time student generator take the previous p-1 synthesized frames $\{Y\}_1^{p-1}$ and p semantic maps $\{X\}_1^p$ as input and generate the next frame. The previous frames are used to capture the temporal coherence of the sequences and generate more coherent video frames. The generation process of the full-time teacher generator is the consecutive generation. More generally, each generation iteration of the full-time teacher generator can be formulated as follows:

$$Y_k = f_{FG}(\{X\}_{k-p}^k, \{Y\}_{k-p}^{k-1}), \tag{2}$$

where Y_k denotes the predicted current generative frame of the full-time teacher generator. f_{FG} denotes the generation function of the full-time teacher generator.

 ${X}_{k-p}^k$ denotes p+1 frames of semantic maps, and ${Y}_{k-p}^{k-1}$ denotes the previous p generated frames.

Different from the full-time teacher generator whose uniform sampling interval is 1, the uniform sampling interval of the part-time student generator is g, where 1 < g < T. g is a random number and randomly selected in each training iteration. Similarly, the frame generation of the part-time student generator can be formulated as follows:

$$Y_k^* = f_{PG}(f_R^d(\{X_t\}_{k-p}^{*k}), \{Y\}_{k-p}^{*k-1}), \tag{3}$$

where Y_k^* denotes the predicted current generative frame of the part-time student generator. f_{PG} denotes the generation function of the part-time student generator, f_R^d denotes the function of reducing the resolution into $\frac{1}{(2^d)^2}$, $\{X\}_{k-p}^{*k}$ includes p frames of the sparse semantic sequences and $\{Y\}_{k-p}^{*k-1}$ is the previous frames of the synthesized sparse sequences. To better illustrate Temporal KD, we set p = 1 in Fig. 4.

Because the full-time teacher generator is trained on sequences with dense frames and is learned to generate dense coherent frames, the full-time teacher generator cannot directly skip partial frames to generate sparse frames, leading to expensive computational cost. The part-time student generator generate sparse frames and interpolate intermediate frames with the slight computational cost. However, since the part-time student generator is trained on sequences with sampled sparse frames, the low sample rate could be two times less than temporal motion frequency and thus leads to aliasing according to Nyquist–Shannon sampling theorem. Our preliminary experiments also show that the large changes between two non-adjacent frames cause remarkable inter-frame incoherence and generate a bad result.

Local Temporal-aware Knowledge Distillation. We first introduce the local temporal-aware knowledge distillation. Our goal is to distill the knowledge from the full-time generator to the part-time student generator to reduce aliasing. A straightforward idea is to align the outputs of the full-time generator and the outputs of the part-time student generator and reduce the distances between the corresponding synthesized frame pairs. The loss function of local temporal-aware knowledge distillation loss L_{LTKD} is given by

$$L_{LTKD} = \frac{1}{T} \sum_{t=0}^{T} [MSE(Y_t^*, Y_t) + L_{per}(Y_t^*, Y_t)], \qquad (4)$$

where Y' denotes the resulting frame of the spatially low-demand generator and Y denotes the resulting frame of the teacher net. MSE represents a mean squared error between two frames. L_{per} denotes a perceptual loss [42].

Global Temporal-aware Knowledge Distillation. Local temporal-aware knowledge distillation allows the part-time student generator to imitate the local motion of the full-time teacher generator. However, it does not consider the global semantic consistency. Therefore, we further propose a global temporal-aware knowledge distillation to distill global temporal coherence.

It is observed that the current frame synthesis deeply relies on the results of the previous synthesis. This indicates that the generated current frame implicitly contains information from the previous frames. The global temporal-aware knowledge distillation exploits the generated non-adjacent frames generated by the full-time teacher generator to distill the hidden information of temporal coherence. The frames generated by the full-time teacher generator at the same timestamps as the non-adjacent frames by the part-time student generator are extracted and concatenated into a predicted sequence $\{Y^*\}$ (Seq-s). Similarly, the results of the full-time teacher generator are concatenated into a sequence $\{Y^S\}$ (Seq-t) in time order, as shown in the right figure of Fig. 3. We introduce a global temporal loss to minimize the distance between $\{Y^*\}$ and $\{Y^S\}$, namely I3D-loss. The global temporal-aware knowledge distillation employs a well-trained I3D [3] model, a well-known video recognition model, to extract the time-series features of neighbor frames of $\{Y^S\}$ and $\{Y^S\}$. The global temporalaware knowledge distillation loss is given by

$$L_{GTKD} = MSE(f_{I3D}(\{Y^*\}\}), f_{I3D}(\{Y^S\})),$$
(5)

where MSE calculates the mean squared error between two feature vectors, f_{I3D} is the function of the pre-trained I3D model. Finally, we obtain a temporal knowledge distillation loss, as written as,

$$L_{TKD} = \alpha L_{LTKD} + \beta L_{GTKD}, \tag{6}$$

where α and β control the importance of local and global losses. We set $\alpha = 2$ and $\beta = 15$ to enhance the global temporal coherence.

Full Objective Function. We integrate local temporal-aware and global temporalaware knowledge distillation into a unified optimization framework, which enables the part-time student generator to imitate both global and local motions of the full-time teacher generator. The full objective function is given by

$$L_{KD} = \sigma L_{STD} + \gamma L_{TKD},\tag{7}$$

where σ and γ control the weights of spatial and temporal KD losses, respectively. In particular, σ is set as 1 and γ is set as 2 to learn more knowledge of the temporal features from the teacher net.

3.5 Semantic-driven Motion Compensation for Interpolation Synthesis

The temporal compression further greatly reduces the computational cost compared with the original Vid2Vid generator. However, the part-time student generator can only synthesize sparse frames $\{Y'\}$. To compensate for this problem, we use a fast motion compensation method [33], a zero-parameter algorithm, to complete the sequence. Motion compensation enables the synthesis of the inter frames between key-frames. As the adjacent frames are with slight changes,

the final results remain a reliable visual performance by reducing the temporal redundancy. During inference, another question is which frames should be synthesized by the part-time student generator as sparse frames and how to determine these key-frames without sufficient photo-realistic frames. In this paper, we surprisingly find that we can distinguish key-frames $\{X'_k | k \in K\}$, where Kis a set containing the numbers of key-frame, from semantic maps $\{X'_0\}^T$. With the key semantic maps, the part-time student generator generates sparse frames $\{Y'_k | k \in K\}$ and finally, we interpolate other inter frames to a full-size result sequence $\{Y\} \in \mathbb{R}^{T \times H \times W}$ by the fast motion compensation method.

4 Experiments

4.1 Experimental Setup

Models. We conduct our experiments using Vid2Vid[42] model. The original Vid2Vid model uses a coarse-to-fine generator for high-resolution output. To simplify the compression problem, we only retrain the first-scale Vid2Vid model based on the official repository². We also evaluate three compression methods for image synthesis models, including NAS compression [27], CA compression [29] and CAT [24]. Since Vid2Vid is a 2D-based generation framework, these three methods can be easily transferred into this task with adjustments.

Datasets. We evaluate our proposed compression method on several datasets. We pre-process the following datasets as the same as the settings of Vid2Vid, including Face Video dataset [34], Cityscape Video dataset [8] and Youtube Dancing Video dataset [42]. Since we only apply the first scale of the Vid2Vid model, we re-size the datasets for convenience. Face Video dataset is re-sized into 512×512 , Cityscape Video dataset is re-size into 256×512 , and Youtube Dancing Video dataset is re-sized to 384×512 .

Evaluation Metrics. We apply three metrics for quantitative evaluation, including FID [19], FVD [39] and pose error(PE) [41]. The lower score of the three metrics represents better performance.

Key-frame Selection. We first calculate the residual maps between the adjacent frames, and sum up each map to draw smooth statistical curves using sliding windows. Thus, the peak ones of the curves represent the local maximum of the difference between two adjacent frames and are used as keyframes. Note that our keyframe selection only consumes about 0.5 milliseconds for processing a video of 30 frames.

Motion Compensation. Motion compensation is to predict a video frame given the previous frames and future frames in video compression. We adopt the overlapped block motion compensation (OBMC) [33] and the enhanced predictive zonal search (EPZS) method [38] to generate the non-keyframes by "FFMPEG" toolbox. EPZS and OBMC consumes about 0.0008146G MACs for a 512×512 frame, which is much less than our generative model part.

² https://github.com/NVIDIA/vid2vid

Task	$ _{ $ Method	m.MACs(G)	FPS	Metric		
				$ \mathrm{FID}(\downarrow) $	FVD (\downarrow)	$\mathrm{PE}(\downarrow)$
Sketch2Face	original	2066	2.77	34.17	6.74	_
	NAS	303	10.50	33.64	6.86	
	CA	290	11.07	35.82	6.95	
	CAT	294	10.89	33.90	6.43	
	Ours	282	18.56	29.02	5.79	_
Segmentation2City	original	1254	4.27	_	2.76	_
	NAS	277	12.44	-	2.99	
	CA	187	15.48	_	3.98	
	CAT	178	13.29	_	3.44	
	Ours	132	24.77	_	2.33	
Pose2Body	original	1769	3.01	_	12.31	2.60
	NAS	280	12.57		12.53	3.28
	CA	253	13.92	-	15.89	4.85
	CAT	257	12.48	_	15.75	4.18
	Ours	191	21.39	_	10.03	2.18

 Table 1. Quantitative results of Fast-Vid2Vid. m.MACs represents the mean of the total MACs of the calculation resources for a sequence of video.

4.2 Quantitative Results

We compare our method with the state-of-the-art GAN compression methods, NAS [27], CA [29] and CAT [24] on three benchmark datasets to validate the effectiveness of our approach. For a fair comparison, we perform a decent pruning rate by removing around 60% channels of the Vid2Vid model in CA and CAT, and use NAS to find out the best network configuration with similar mMACs.

The experimental results are shown in Table 1. We can see that given the lower computational budget, our method achieves the best FID and FVD on three datasets. Specifically, other GAN compression methods have lower performance than the full-size model while our method outperforms the original model. Other compression methods speed up the original model by simply cutting the network structures, but they ignore the temporal coherence. Meanwhile, the original Vid2Vid model significantly accumulates losses during inference. Our proposed motion-aware inference accumulates losses since it only generates several frames of the sequence. Such results show the advantage of our spatial-temporal aware compression methods.

4.3 Ablation Study

We adopt Face Video as the benchmark dataset for our ablation study.

Effectiveness of Temporal KD Loss. Based on the spatially low-demand generator mentioned before, we analyze the temporal knowledge distillation for the Vid2Vid model. We set 6 different distillation loss schemes as: (1) w/o TKD: the spatially low-demand generator was retrained on the dataset; (2) TKD-local: the spatially low-demand generator is transferred the only local knowledge from the teacher net; (3) TKD-global: the spatially low-demand generator is transferred the only global knowledge from the teacher net; (4) MMD: the spatially lowdemand generator is transferred the knowledge using MMD-loss [9]. (5) LSTM:



Fig. 5. Ablation Study for Fast Vid2Vid. Left:Temporal Loss ablation study for temporal loss. Right:The trade-off experiments for the windows of key-frames selector.



Fig. 6. The comparison of the results of different knowledge distillation.

the spatially low-demand generator is transferred the knowledge based on LSTM regulation [45]. (6) TKD: the spatially low-demand generator is transferred both local and global knowledge from the teacher net;

As shown in the left figure in Fig. 5, the local knowledge distillation loss improves the performance of the model without KD. The temporal KD loss globally further improves the performance of the common local KD loss, especially in FVD. Our proposed KD loss outperforms MMD-loss and LSTM-based KD loss. It indicates that the temporal KD loss effectively enhances the similarity of distribution on the video recognition network between the videos generated by the teacher network and the student network. We also provide the comparison results in Fig. 6 and our STKD generates more realistic frame than others.

Effectiveness of Spatial KD Loss. We conduct an ablative study for Spatial KD on the Sketch2Face benchmark. The spatial compression methods are used together with our proposed Temporal KD to perform Vid2Vid compression. Table 2 shows that our proposed Spatial KD performs better than other image compression methods. Our Spatial KD does not destroy network structures of the original GAN while other methods tweak the sophisticated parameters.

Effectiveness of Windows for Key-frame Selection. We investigate the sliding windows to select the key-frames to find out the best trade-off between the sliding windows and the performance. The larger sliding windows mean that there are fewer key-frames selected and thus use less computational resources. Interestingly, as shown in the right figure of Fig. 5, it rises significantly in FVD when increasing the sliding windows, and achieves the best performance when the

Method	$\mathrm{MACs}(\mathrm{G})$	FPS	FID	FVD
CA	331	17.00	36.65	6.76
CAT	310	18.02	35.64	6.85
NAS	344	16.78	32.41	6.71
Spatial KD	$\boldsymbol{282}$	18.56	29.02	5.79

Table 2. Ablation study for spatial compression with the proposed Temporal KD.

 Table 3. Ablation study for the interpolation.
 Table 4. Ablation study for the time gap.

	Gap	FID FVD
Interpolation FID FVD	Fixed time gap	32.21 6.85
Linear Interpolation 31.55 6.45	Random time gap	31.43 6.22
Motion Compensation 29.02 5.79	Key-frames	29.02 5.79

sliding windows are three in three tasks. It indicates that the part-time student generator needs enough independent motion to maintain decent performance. **Effectiveness of Interpolation.** We apply two common interpolation methods for completing the video, namely linear interpolation and motion compensation. We also conduct ablative studies on interpolation methods. As shown in Table 3, motion compensation outperforms the simple linear interpolation. Therefore, we apply motion compensation as our zero-parameters interpolation method.

Effectiveness of Time Gap. We also study the ways of selected frames for the generation of the part-time student generator. We use the same numbers of the selected frames. Specifically, the fixed time gap strategy selects the frames in fixed time intervals, the random time gap strategy chooses the frames in random intervals, and the key-frame strategy selects the key-frames as the frames to be synthesized by the part-time student generator. As shown in Table 4, the keyframe strategy outperforms other strategies since the key-frames of a sequence consist of all essential motions and texture.

4.4 Qualitative Results

We illustrate the output sequences of the mentioned methods in Fig. 7. The generators in Vid2Vid synthesis would accumulate the visual losses sequentially since the current frame relies on the previously generated frames. Compared with other GAN compression methods, our proposed method generates more realistic results. We have provided more details in supplementary materials.

5 Conclusion

In this paper, we present Fast-Vid2Vid to accelerate Vid2Vid synthesis. We propose a spatial-temporal compression framework to accelerate the inference by compressing the sequential input data stream but maintaining the parameters of



Fig. 7. Qualitative results compared with the advanced GAN compression methods in the task of Sketch2Face, Segmentation2City and Pose2Body.

the network. In the space dimension, we distill knowledge from the full-resolution domain to the low-resolution domain. In the time dimension, we use temporalaware knowledge distillation for local and global knowledge to obtain a part-time generator. Finally, the part-time generator is used for motion-aware inference where it only generates the key-frames of the sequence and interpolates the middle frames by motion compensation. By reducing the resolution of the input data and extracting the key-frames of the data stream, Fast-Vid2Vid saves the computational resources significantly.

Discussion. We discuss some future directions for this work. Recently, sequencein and sequence-out methods, like transformer, are challenging for model compression. On the contrary, our Fast-Vid2Vid accelerates Vid2Vid by optimizing a part-time student generator (via Temporal KD) and a lower-resolution spatial generator (via Spatial KD), which is versatile for various networks. When combined with seq-in and seq-out transformers like visTR [44], Fast-Vid2Vid first synthesizes a partial video by a part-time transformer-based student generator (via fully parallel computation) and then recovers the full video by motion compensation.

Acknowledgements. This work is supported by NTU NAP, MOE AcRF Tier 1 (2021-T1-001-088), and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- 1. Aguinaldo, A., Chiang, P.Y., Gain, A., Patil, A., Pearson, K., Feizi, S.: Compressing gans using knowledge distillation. arXiv preprint arXiv:1902.00159 (2019) 2, 4, 5
- Belousov, S.: Mobilestylegan: A lightweight convolutional neural network for highfidelity image synthesis. arXiv preprint arXiv:2104.04767 (2021) 2, 4
- Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017) 9
- Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5933–5942 (2019) 2, 4
- Chen, G., Choi, W., Yu, X., Han, T., Chandraker, M.: Learning efficient object detection models with knowledge distillation. Advances in neural information processing systems 30 (2017) 4
- Chen, H., Wang, Y., Shu, H., Wen, C., Xu, C., Shi, B., Xu, C., Xu, C.: Distilling portable generative adversarial networks for image translation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 3585–3592 (2020) 2, 4, 5
- Chu, M., Xie, Y., Mayer, J., Leal-Taixé, L., Thuerey, N.: Learning temporal coherence via self-supervision for gan-based video generation. ACM Transactions on Graphics (TOG) 39(4), 75–1 (2020) 4
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016) 10
- Feng, Z., Lai, J., Xie, X.: Resolution-aware knowledge distillation for efficient inference. IEEE Transactions on Image Processing 30, 6985–6996 (2021) 4, 6, 11
- Fu, Y., Chen, W., Wang, H., Li, H., Lin, Y., Wang, Z.: Autogan-distiller: Searching to compress generative adversarial networks. arXiv preprint arXiv:2006.08198 (2020) 2, 4, 5
- Gafni, O., Wolf, L., Taigman, Y.: Vid2game: Controllable characters extracted from real-world videos. arXiv preprint arXiv:1904.08379 (2019) 4
- Gao, C., Chen, Y., Liu, S., Tan, Z., Yan, S.: Adversarialnas: Adversarial neural architecture search for gans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5680–5689 (2020) 2
- Gong, X., Chang, S., Jiang, Y., Wang, Z.: Autogan: Neural architecture search for generative adversarial networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3224–3234 (2019) 2
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems. vol. 27 (2014) 2
- Han, S., Pool, J., Tran, J., Dally, W.J.: Learning both weights and connections for efficient neural networks. arXiv preprint arXiv:1506.02626 (2015) 4
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) 2
- 17. He, Y., Kang, G., Dong, X., Fu, Y., Yang, Y.: Soft filter pruning for accelerating deep convolutional neural networks. arXiv preprint arXiv:1808.06866 (2018) 4

- 16 L. Zhuo et al.
- He, Y., Zhang, X., Sun, J.: Channel pruning for accelerating very deep neural networks. In: Proceedings of the IEEE international conference on computer vision. pp. 1389–1397 (2017) 4
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017) 10
- Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) 4
- Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for mobilenetv3. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1314–1324 (2019) 6
- Hu, H., Peng, R., Tai, Y., Tang, C., Trimming, N.: A data-driven neuron pruning approach towards efficient deep architectures. arXiv preprint arXiv:1607.03250 46 (2016) 4
- Hu, H., Peng, R., Tai, Y.W., Tang, C.K.: Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. arXiv preprint arXiv:1607.03250 (2016) 4
- Jin, Q., Ren, J., Woodford, O.J., Wang, J., Yuan, G., Wang, Y., Tulyakov, S.: Teachers do more than teach: Compressing image-to-image models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13600–13611 (2021) 2, 4, 5, 10, 11
- Kappel, M., Golyanik, V., Elgharib, M., Henningson, J.O., Seidel, H.P., Castillo, S., Theobalt, C., Magnor, M.: High-fidelity neural human motion transfer from monocular video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1541–1550 (2021) 2, 4
- 26. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710 (2016) 4
- Li, M., Lin, J., Ding, Y., Liu, Z., Zhu, J.Y., Han, S.: Gan compression: Efficient architectures for interactive conditional gans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5284–5294 (2020) 2, 4, 5, 6, 10, 11
- Lin, J., Zhang, R., Ganz, F., Han, S., Zhu, J.Y.: Anycost gans for interactive image synthesis and editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14986–14996 (2021) 2
- Liu, Y., Shu, Z., Li, Y., Lin, Z., Perazzi, F., Kung, S.Y.: Content-aware gan compression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12156–12166 (2021) 2, 4, 5, 10, 11
- Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., Zhang, C.: Learning efficient convolutional networks through network slimming. In: Proceedings of the IEEE international conference on computer vision. pp. 2736–2744 (2017) 4
- Lopez-Paz, D., Bottou, L., Schölkopf, B., Vapnik, V.: Unifying distillation and privileged information. arXiv preprint arXiv:1511.03643 (2015) 4
- Luo, P., Zhu, Z., Liu, Z., Wang, X., Tang, X.: Face model compression by distilling knowledge from neurons. In: Thirtieth AAAI conference on artificial intelligence (2016) 4
- Orchard, M.T., Sullivan, G.J.: Overlapped block motion compensation: An estimation-theoretic approach. IEEE Transactions on Image Processing (1994) 9, 10

- Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Nießner, M.: Faceforensics: A large-scale video dataset for forgery detection in human faces. arXiv preprint arXiv:1803.09179 (2018) 10
- Sajjadi, M.S., Vemulapalli, R., Brown, M.: Frame-recurrent video super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6626–6634 (2018) 4
- Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: First order motion model for image animation. Advances in Neural Information Processing Systems 32, 7137–7147 (2019) 4
- Siarohin, A., Woodford, O.J., Ren, J., Chai, M., Tulyakov, S.: Motion representations for articulated animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13653–13662 (2021) 4
- Tourapis, A.M.: Enhanced predictive zonal search for single and multiple frame motion estimation. In: Visual Communications and Image Processing 2002. vol. 4671, pp. 1069–1079. SPIE (2002) 10
- Unterthiner, T., van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Towards accurate generative models of video: A new metric & challenges. arXiv preprint arXiv:1812.01717 (2018) 10
- 40. Wang, H., Gui, S., Yang, H., Liu, J., Wang, Z.: Gan slimming: All-in-one gan compression by a unified optimization framework. In: European Conference on Computer Vision. pp. 54–73. Springer (2020) 2, 4
- Wang, T.C., Liu, M.Y., Tao, A., Liu, G., Kautz, J., Catanzaro, B.: Few-shot videoto-video synthesis. arXiv preprint arXiv:1910.12713 (2019) 2, 4, 10
- Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. arXiv preprint arXiv:1808.06601 (2018) 2, 3, 4, 5, 7, 8, 10
- 43. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: Highresolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018) 4
- 44. Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-end video instance segmentation with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8741–8750 (2021) 14
- Xiao, Z., Fu, X., Huang, J., Cheng, Z., Xiong, Z.: Space-time distillation for video super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2113–2122 (2021) 12
- Xu, G., Xu, J., Li, Z., Wang, L., Sun, X., Cheng, M.M.: Temporal modulation network for controllable space-time video super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6388– 6397 (2021) 4
- Xu, R., Li, X., Zhou, B., Loy, C.C.: Deep flow-guided video inpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3723–3732 (2019) 4
- 48. Yim, J., Joo, D., Bae, J., Kim, J.: A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4133–4141 (2017) 4
- Yu, C., Pool, J.: Self-supervised gan compression. arXiv preprint arXiv:2007.01491 (2020) 4

- 18 L. Zhuo et al.
- Zhang, T., Ye, S., Zhang, K., Tang, J., Wen, W., Fardad, M., Wang, Y.: A systematic dnn weight pruning framework using alternating direction method of multipliers. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 184–199 (2018) 4
- Zhou, Y., Wang, Z., Fang, C., Bui, T., Berg, T.: Dance dance generation: Motion transfer for internet videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019) 2, 4
- Zou, X., Yang, L., Liu, D., Lee, Y.J.: Progressive temporal feature alignment network for video inpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16448–16457 (2021) 4