

Learnable Gated Temporal Shift Module for Deep Video Inpainting

Ya-Liang Chang*
yaliangchang@cmlab.csie.ntu.edu.tw

Zhe Yu Liu*
zhe2325138@cmlab.csie.ntu.edu.tw

Kuan-Ying Lee
r03922165@ntu.edu.tw

Winston Hsu
whsu@ntu.edu.tw

National Taiwan University,
Taipei, Taiwan

Abstract

How to efficiently utilize temporal information to recover videos in a consistent way is the main issue for video inpainting problems. Conventional 2D CNNs have achieved good performance on image inpainting but often lead to temporally inconsistent results where frames will flicker when applied to videos (see [video:Edge-Connect](#)); 3D CNNs can capture temporal information but are computationally intensive and hard to train. In this paper, we present a novel component termed Learnable Gated Temporal Shift Module (LGTSM) for video inpainting models that could effectively tackle arbitrary video masks without additional parameters from 3D convolutions. LGTSM is designed to let 2D convolutions make use of neighboring frames more efficiently, which is crucial for video inpainting. Specifically, in each layer, LGTSM learns to shift some channels to its temporal neighbors so that 2D convolutions could be enhanced to handle temporal information. Meanwhile, a gated convolution is applied to the layer to identify the masked areas that are poisoning for conventional convolutions. On the FaceForensics and Free-form Video Inpainting (FVI) dataset, our model achieves state-of-the-art results with simply 33% of parameters and inference time. The source code is available on <https://github.com/amjltc295/Free-Form-Video-Inpainting>.

1 Introduction

Free-form video inpainting, to fill in arbitrary missing regions in a video, is a very challenging task. It could be widely used for movie post-processing, damaged video recovery and video editing. For humans, it takes tremendous efforts to recover these missing areas like Fig. 1, while an autonomous method may complete it easily.

The key for free-form video inpainting is to model spatial-temporal features. That is, a model needs to capture the content of masked areas according to its surroundings and fill in these areas with related pixels. Traditional patch-based methods [1, 2] fill in these areas by finding similar patches from other parts of the videos. However, the searching algorithms




Figure 1: Our model takes videos with free-form masks (first row) and fills in the missing areas with proposed LGTSM to generate realistic completed results (second row) compared to the original videos (third row). It could be applied to video editing tasks such as video object removal, as shown in the first two columns. Best viewed in color and zoom-in. See corresponding videos in the following links: [object removal](#), [free-form masks](#), and [faces](#).

usually have high computational complexity and the missing area may not be found for complex objects or masks (see Fig. 4).

On the other hand, deep learning methods [2, 3, 15, 17, 28] could fill in unseen masked areas by the encoding and decoding process, based on the structures learned from the training data. Still, compared to the success in image inpainting, deep learning methods struggle to model these video features due to the additional temporal dimension. Using 3D convolution to model spatial-temporal features is the most intuitive way but it requires plenty parameters and is hard to train.

In this paper, we propose a novel component termed Learnable Gated Temporal Shift Module (LGTSM) to handle free-form video masks with 2D convolutions, motivated by the TSM originally for action recognition. Though inspired by TSM, we found that TSM perfectly suitable for free-form video inpainting as it cannot totally make use of neighboring frames from the beginning layers nor handle irregular masks, which makes us propose LGTSM. Specifically, in each layer, LGTSM learns to shift a part of feature channels in a frame to its neighboring frame, and then attends on masked/inpainted/unmasked areas by a gating convolutional filter. LGTSM enables 2D convolutions to process masked videos and generate state-of-the-art results as 3D convolutions with only 33% parameters and inference time.

Our paper makes the following contributions:

- We propose the Gated Temporal Shift Module that could recover videos with free-form masks in 2D convolutions by temporally shifting and gating features in each layer, which reduce the model size and computational time to 33% compared to 3D convolutions.
- Given that video inpainting requires more information from neighboring frames, we propose a novel Learnable Gated Temporal Shift Module that could learn the temporally shifting kernels and achieved state-of-the-art performance.
- We propose the TSMGAN loss which significantly improves the model performance for free-form video inpainting.

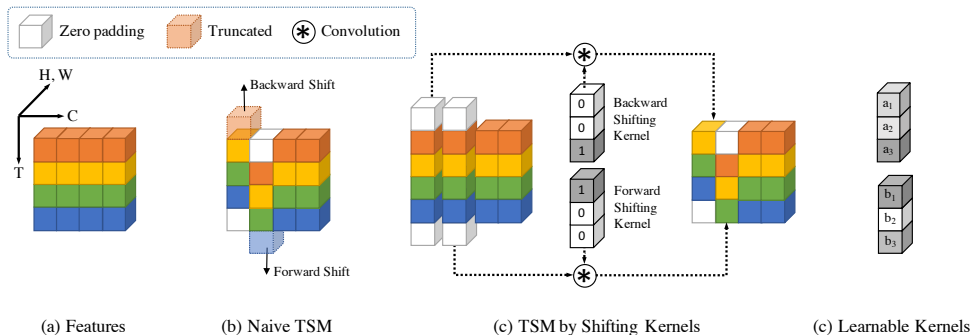


Figure 2: Explanation of the learnable shifting kernels in the proposed LGTSM. (a) Input features for the layer. We will do shifting operation on channel \times time dimensions. (b) Original TSM from [24]. (c) Equivalent TSM by temporal shifting kernels. (d) In the proposed LGTSM, temporal shifting kernels are also learnable and the size could be different.

2 Related Work

Image Inpainting. Recently, deep learning based methods have taken over the image inpainting task. Xie *et al.* [27] firstly apply convolutional neural networks (CNNs) on small-region image inpainting and denoising. Pathak *et al.* [49] then extended [27] to larger region with an encoder-decoder structure. Moreover, [49] adopted the generative adversarial network (GAN) [6], where a generator learning to create realistic images and a discriminator striving to tell fake ones are trained together to improve the image quality. Subsequently, [12, 28, 29] also developed new GAN architectures with different components for image inpainting.

Among these deep methods, Yu *et al.* [28] proposed gated convolutions for image inpainting that uses an additional gating convolution to learn the difference between masked, inpainted and unmasked areas in each layer. We integrate such gating mechanism to our model. Also, Nazeri *et al.* [17] developed a two-stage model that generates image edges first before recovering the whole images conditioned on edges. Their model achieved state-of-the-art results, and we set it as one of our baselines.

Video Inpainting. Generally, video inpainting could be viewed as an extension of image inpainting with temporal constraints (i.e content in different frames need to be consistent.) However, different from image inpainting, patch-based methods [4, 9, 18, 26] still play a role in video inpainting as more patches are available in videos. Among them, Huang *et al.* [9] jointly estimate optical flow and colors in the masked region to fix the moving camera problem and reached state-of-the-art results, so we also set it as one of our baselines.

Although patch-based methods have made great success in video inpainting, they are highly limited in computational time due to the search algorithms. In addition, the masked areas still need to be patchable; these methods do not work on complex objects such as faces. To address these problems, Wang *et al.* [25] proposed the first deep learning based method for video inpainting, with a two-stage CombCN that uses 3D convolutions to generate coarse but temporally consistent videos and then refines with 2D convolutions. Their model could learn to recover face videos so we also set it as a baseline.

Temporal Modeling. As most state-of-the-art deep video inpainting learning methods adopt the encoder-decoder structure, the key is to model the spatial-temporal structures in

videos. Over the past few years, a variety of deep learning architectures have been proposed to model video structures, especially for action recognition. These architectures include applying temporal pooling [10] or recurrent networks [60] on top of 2D convolutions to model temporal features structures, combining optical flows and RGB frames to make two-stream networks [22], directly using 3D convolutions [24] and variants of these models [8, 9]. For more details, we refer readers to [8].

Despite the great performance, many architectures for action recognition cannot be applied to video inpainting since 1) the input video is corrupted, so it is hard to derive optical flows or apply naive convolutions 2) these architectures only need an encoder, while for video inpainting requires a decoder to recover the missing areas.

Aside from architecture-level temporal modeling, there are also works that focus on the module-level [13, 14], which is more applicable for video inpainting. Lin *et al.* [14] proposed Temporal Shift Module that shifts part of feature channels in each frame to its neighboring frames so that 2D convolutions could handle temporal information. Similarly, Li *et al.* [13] developed Temporal Bilinear (TB) that applies factorized bilinear operation on features to model interactions between frames. Note that these models are for action recognition that all input frames are valid, while for video inpainting, many pixels are masked. Based on these ideas of integrating temporal information to 2D convolutions, we propose the Learnable Gated Temporal Shift Module for video inpainting.

3 Proposed Method

3.1 Learnable Gated Temporal Shift Module

Models with 3D convolutions could capture temporal information in an intuitive way but are hard to train due to the large number of parameters. To this end, we extend the residual Temporal Shift Module (TSM) [14], originally designed for action recognition, to video inpainting. TSM tackles temporal information in 2D convolutions. The input activation $F_{t,x,y}$ for each convolutional layer in video inpainting is in shape of (B, C, L, H, W) where B is the batch size, C is the channel number, L is the temporal length, H is the height and W is the width of the input activation. For each frame in L, TSM shifts a portion of channels to its previous and next frame before the convolutional operations, as shown in Fig. 2(b). These shifted channels contain features from other frames, so together with unshifted features, the original 2D convolutions could learn the temporal structures accordingly.

However, for free-form video inpainting, not every feature point is valid as many areas are masked. These masked areas are harmful to naive TSM as convolutions cannot tell the difference between valid and invalid feature points. To address this issue, we design the Gated Temporal Shift Module (GTSM) for free-form video inpainting (see Fig. 3). Specifically, in addition to the TSM module, a gating convolutional filter W_g is applied to input features $F_{t,x,y}$ to obtain a gating $Gating_{t,x,y}$. This gating will serve as a soft validity map to identify the masked/unmasked/inpainted areas for the output features $Features_{t,x,y}$ from the TSM module with original convolutional filter W_f .

Mathematically, GTSM could be expressed as:

$$Gating_{t,x,y} = \sum \sum W_g \cdot F_{t,x,y} \quad (1)$$

$$Features_{t,x,y} = \sum \sum W_f \cdot TSM(F_{t-1,x,y}, F_{t,x,y}, F_{t+1,x,y}) \quad (2)$$




Figure 3: Module design. We integrate (a) Residual TSM [14] and (b) gated convolution [28] to (c) Gated Temporal Shift Module (GTSM) and design learnable temporal shifting kernel (Fig. 2c) to make (d) the proposed Learnable Gated Temporal Shift Module (LGTSM).

$$Output_{t,x,y} = \sigma(Gating_{t,x,y})\phi(Features_{t,x,y}) \quad (3)$$

where σ is the sigmoid function that transforms gating to values between 0 (invalid) and 1 (valid), and ϕ is the activation function for the convolution. Note that the TSM module could be easily modified to online settings (without peeking future frames) for real-time applications.

Note that the temporal shifting operation in TSM is similar to applying forward/backward shifting kernels on the channel-temporal map, as shown in Fig. 2(c). In TSM, these kernels are fixed; a frame could only get features from its one-frame neighbors in each layer. Such fixed shifting kernels in TSM are insufficient to make use of further neighboring frames as temporal information could only be aggregated through deeper layers. Unlike action recognition, video inpainting models sorely need information from the beginning layers to capture the spatial-temporal structures so that deeper layers could recover the missing areas accordingly.

Therefore, we also propose the Learnable Gated Temporal Shift Module (LGTSM), where the temporal shifting kernels are also learnable and the kernel size could be larger (see Fig. 2(d)). With LGTSM, the model could learn to shift and scale features from specific temporal neighbors in each layer (or not). For example, the model could get features from more temporally further neighbors in the first few layers and remain unshifted in the deeper layers. This greatly enhances the model capability with very little cost.

In practice, the shifting operation only uses an additional buffer in size of 1/4 channels, so it has little cost in terms of computational time and run-time memory compared to traditional 2D convolutions. Note that the number of kernels could also be flexible (there are only two for TSM: forward and backward). Moreover, LGTSM could learn temporal information with very few extra parameters, and we found that it could achieve state-of-the-art results as the 3D convolutions with only 33% parameters and inference time.

3.2 Loss Functions

Design the combination of loss functions to train a video inpainting model is non-trivial due to the uncertainty of the free-form masks and the high complexity of video features. We use l_1 loss for low-level features, perceptual loss and style loss for image content, and propose TSMGAN loss for handle high-level features and enhance realism. In this section, we will introduce these loss functions to train our model.

Masked l_1 loss. The l_1 loss focuses on the pixel-level features, which is widely used for generative tasks on image [15, 16, 28, 29], and videos [4, 25]:

$$L_{l_1} = \mathbb{E}_{t,x,y}[|O_{t,x,y} - V_{t,x,y}|] \quad (4)$$

Perceptual loss and style loss. l_1 loss often leads to blurry results as it only focus on low-level features. To address this problem, we adopt the perceptual and style loss [6] to keep image contents. Similar loss functions could be found in many generative tasks such as image inpainting [15, 17] and super-resolution [11, 12].

The perceptual loss could be viewed as the l_1 loss in feature level:

$$L_{perc} = \sum_{t=1}^n \sum_{p=0}^{P-1} \frac{|\Psi_p^{O_t} - \Psi_p^{V_t}|}{N_{\Psi_p^{V_t}}} \quad (5)$$

where V_t is the input video, $\Psi_p^{V_t}$ is the activation from the p th selected layer of the pretrained network, and $N_{\Psi_p^{V_t}}$ is the number of elements in the p th layer. We choose layer $relu_{2_2}$, $relu_{3_3}$ and $relu_{4_3}$ from the VGG [23] network pre-trained on ImageNet [21].

Style loss is a variant of perceptual loss, with an auto-correlation (Gram matrix) applied to the features first:

$$L_{style} = \sum_{t=1}^n \sum_{p=0}^{P-1} \frac{1}{C_p C_p} \frac{|(\Psi^{O_{t,p}})^T (\Psi^{O_{t,p}}) - (\Psi^{V_{t,p}})^T (\Psi^{V_{t,p}})|}{C_p H_p W_p} \quad (6)$$

where $\Psi_p^{O_t}$ and $\Psi_p^{V_t}$ are both features from the pre-trained VGG network, as the ones in the perceptual loss 5.

TSMGAN loss. All the aforementioned loss functions are for image only, which do not take temporal consistency into consideration. Therefore, we develop TSMGAN to learn temporal consistency. We set up a generative adversarial network (GAN) with Gated Temporal Shift Module integrated on both the generator and discriminator as stated in 3.1.

The TSMGAN discriminator is composed of six 2D convolutional layers with TSM. Also, we apply the recently proposed spectral normalization [16] to both the generator and discriminator as [14] to enhance the training stability. The TSMGAN loss L_D for the discriminator to tell if the input video z is real or fake and L_G for the generator to fool the discriminator are defined as:

$$L_D = \mathbb{E}_{x \sim P_{data}(x)} [ReLU(1 + D(x))] + \mathbb{E}_{z \sim P_z(z)} [ReLU(1 - D(G(z)))] \quad (7)$$

$$L_G = -\mathbb{E}_{z \sim P_z(z)} [D(G(z))] \quad (8)$$

As [18], the kernel size is 5×5 , stride 2×2 and the shifting operation is applied to all 11 convolutional layers for the TSMGAN discriminator, so the receptive field of each output feature point includes the whole video. It could be viewed as several GANs on different feature points, and a local-global discriminator structure [19] is thus not needed.

Besides, the TSMGAN learns to classify real or fake for each spatial-temporal feature point from the last convolutional layer in the discriminator, which mostly consists of high-level features. Since the l_1 loss already focuses on low-level features, using TSMGAN could improve the model in an efficient way.

Overall loss. The overall loss function to train the model is defined as:

$$L_{total} = \lambda_{l_1} L_{l_1} + \lambda_{perc} L_{perc} + \lambda_{style} L_{style} + \lambda_G L_G \quad (9)$$

where λ_{l_1} , λ_{perc} , λ_{style} and λ_G are the weights for l_1 loss, perceptual loss, style loss and TSMGAN loss, respectively.

Dataset	Metric	Mask	TCCDS	EC	CombCN	3DGated	LGTSM
FaceFo-rensics	MSE↓	Curve	0.0031*	0.0022	0.0012	0.0008	0.0012
		Object	0.0096*	0.0074	0.0047	0.0048	0.0053
		BBox	0.0055	0.0019	0.0016	0.0018	0.0020
	LPIPS↓	Curve	0.0566*	0.0562	0.0483	0.0276	0.0352
		Object	0.1240*	0.0761	0.1353	0.0743	0.0770
		BBox	0.1260	0.0335	0.0708	0.0395	0.0432
	FID↓	Curve	1.281*	0.848	0.704	0.472	0.601
		Object	1.107*	0.946	0.913	0.766	0.782
		BBox	1.013	0.663	0.742	0.663	0.681
FVI	MSE↓	Curve	0.0220*	0.0048	0.0021	0.0025	0.0028
		Object	0.0110*	0.0075	0.0048	0.0056	0.0065
	LPIPS↓	Curve	0.2838*	0.1204	0.0795	0.0522	0.0569
		Object	0.2595*	0.1398	0.2054	0.1078	0.1086
	FID↓	Curve	2.1052*	1.0334	0.7669	0.6096	0.6436
		Object	1.2979*	1.0754	1.0908	0.9050	0.9323
Number of Parameters↓ (Generator/Discriminator)			-	20M/6M	16M/-	36M/18M	12M/6M
Inference FPS↑			0.05 [#]	55	120	23	80

Table 1: Quantitative comparison with baseline models on the FaceForensics and FVI dataset based on [2]. The results of FVI dataset are averaged of seven mask-to-frame ratios; detailed results of each ratio could be found in the supplementary materials. *TCCDS failed on some cases and the results are averaged of successful ones. [#]runs on CPU; others are on GPU.

3.3 Network Design

The model has a U-net like generator and a TSMGAN discriminator. The generator is composed of 11 convolutional layers with the proposed Gated Temporal Shift Module, including down-sampling, dilated and up-sampling ones. Similar structures are also adopted for state-of-the-art image inpainting models [15, 18]. Unlike U-net, there is no skip connection as there are many masked areas in the down-sampling layers. For down-sampling and up-sampling layers, we apply bilinear interpolation before convolutions.

4 Experimental Results

4.1 Setups

Datasets. To compare with the baselines [2, 9, 17, 25], we train and test our model on the FaceForensics [20] and Free-form Video Inpainting (FVI) [2] dataset. Both datasets are based on videos from YouTube, so they are close to real world scenarios. FaceForensics is composed of 1,004 videos with face, news-caster or newsprogram tags. There are only frontal faces cropped to 128×128 in the FaceForensics dataset, so it is rather simple for learning based models. Amongst, 150 videos are for evaluation while the rest are for training.

On the other hand, the FVI dataset contains 15,000 high-resolution videos with human activities, animals, natural scenes, etc. It also provides algorithms to generate free-form video masks for training. We re-size videos to 320×180 and split 100 videos for evaluation following the setup in [2]. Note the FVI dataset is considered more challenging as the videos are very diverse.




Figure 4: Visual comparison with the baselines on the FVI testing set with object-like masks. Best viewed in color and zoom-in. See video.

Training and Testing. Empirically we found that the model converges slower when directly trained as a whole. Thus, during training, we first pre-train the generator without the TSMGAN loss until convergence, and then fine-tune with the TSMGAN. We initialize the temporal shifting kernels in the LGTSM with the values that are equivalent to the original TSM. The pre-train stage takes about 1 day, while the fine-tune stage takes about 3 days on the FVI dataset, which is 3 times faster, reducing training time from 10 days to 3 days, than the model with 3D convolutions [2], demonstrating the merits of the proposed module. For other implementation details, please see the supplementary materials.

4.2 Quantitative Results

As [2], mean square error (MSE) and Learned Perceptual Image Patch Similarity (LPIPS) [31] are used to evaluate the image quality; the Fréchet Inception Distance (FID) [8] with pre-trained I3D [1] is used to evaluate video quality and temporal consistency.

We compare with state-of-the-art baselines with different strategies: the patch-based video inpainting method TCCDS by Huang *et al.* [9], two-stage deep image inpainting method Edge-Connect (EC) by Nazeri *et al.* [14], two-stage deep video inpainting method CombCN by Wang *et al.* [15], and one-stage deep video inpainting method with 3D gated convolutions (3DGated) by Chang *et al.* [2]. We train all learning based models on the FaceForensics and FVI datasets with free-form masks from [2] for fair comparison.

The averaged results of 7 ranges of mask-to-frame ratio from 0 - 10% to 60 - 70% on the FaceForensics and FVI testing are shown in Table 1. We could see that our model is on par with the state-of-the-art method 3DGated [2] in terms of perceptual distance (LPIPS and FID) and video quality (FID) with only 33% of parameters and inference time (note that the results are averaged; our model performs better for some mask-to-frame ratios). TCCDS failed on many cases since the masks are irregular and it cannot properly recover partially masked objects. Edge-connect (EC) performs better on FaceForensics dataset with bounding box masks because faces are all aligned and the generated edges could be stable. Still, it has serious temporal inconsistent problem under other circumstances (see Fig. 4). Although CombCN has lowest MSE scores, it could only generates blurry results for the

3D conv.	TSM	Learnable shifting	Gated conv.	GAN	LPIPS↓	FID ↓	Param. ↓
✓			✓	✓	0.1209	1.034	36M+18M
	✓			✓	0.2048	1.303	12M*+6M
	✓		✓		0.1660	1.198	12M
	✓		✓	✓	0.1256	1.091	12M+6M
	✓	✓	✓	✓	0.1213	1.039	12M+6M

Table 2: Ablation study on the FVI dataset with object-like masks. Number of parameters are shown as generator/discriminator. *We increase the channel of vanilla convolution to fairly compare with gated convolutions.

FVI dataset (see Fig. 4) and require more parameters for the generator than our model.

4.3 Qualitative Results

From Fig. 4 we could observe that the proposed model outperforms TCCDS (wrong patches), Edge-connect (temporally inconsistent) and CombCN (blurry), while almost the same as 3DGated. More visual comparison could be found in the supplementary materials.

4.4 Ablation Study

In order to validate the contribution of each component, we also conduct an ablation study on the FVI dataset (Table 2). We could observe that both the gated convolution and TSMGAN play important roles in our model; without them, the model will have a significant drop of performance. The proposed learnable shifting kernel further improves the performance with almost no additional parameters and achieve state-of-the-art results. Visual comparison be found in the supplementary materials and [videos](#).

5 Discussion and Future Work

Our model achieves state-of-the-art performance with LGTSM and 2D convolutions. However, the performance is still a bit lower than the 3D convolutions in [2]. How to design a module that could handle temporal information in a more efficient way is still a challenging future work.

Another future work is to extend the input videos to higher or arbitrary resolution. For now, deep learning models are limited to a fixed resolution, which is not sufficient for exquisite videos. Simply applying these models to different parts of a high-resolution video may cause in spatial inconsistency.

6 Conclusion

This paper presented a novel Learnable Gated Temporal Shift Module (LGTSM) for free-form video inpainting. LGTSM learns to shift some channels to its temporal neighbors in each frame and apply gating filter to attend on masked/inpainted/unmasked areas, which enables 2D convolutions to process temporal information and tackle poisoning masked areas at the same time. In addition, LGTSM is highly efficient, using only 33% of parameters and inference time compared to the state-of-the-art model with 3D convolutions. Experiments on the FaceForensics and FVI dataset suggest that the proposed model reach state-of-the-art performance in terms of evaluation metrics and visual results.

7 Acknowledgement

This work was supported in part by the Ministry of Science and Technology, Taiwan, under Grant MOST 108-2634-F-002-004. We also benefit from the NVIDIA grants and the DGX-1 AI Supercomputer. We are grateful to the National Center for High-performance Computing.

References

- [1] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [2] Ya-Liang Chang, Zhe Yu Liu, and Winston Hsu. Free-form video inpainting with 3d gated convolution and temporal patchgan. *arXiv preprint arXiv:1904.10247*, 2019.
- [3] Ya-Liang Chang, Zhe Yu Liu, and Winston Hsu. Vornet: Spatio-temporally consistent video inpainting for object removal. *arXiv preprint arXiv:1904.06726*, 2019.
- [4] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1933–1941, 2016.
- [5] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*, 2015.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [7] Miguel Granados, James Tompkin, K Kim, Oliver Grau, Jan Kautz, and Christian Theobalt. How not to be seen—Object removal from videos of crowded scenes. In *Computer Graphics Forum*, volume 31, pages 219–228. Wiley Online Library, 2012.
- [8] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017.
- [9] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. Temporally coherent completion of dynamic video. *ACM Transactions on Graphics (TOG)*, 35(6): 196, 2016.
- [10] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [11] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.

- [12] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [13] Yanghao Li, Sijie Song, Yuqi Li, and Jiaying Liu. Temporal bilinear networks for video action recognition. *arXiv preprint arXiv:1811.09974*, 2018.
- [14] Ji Lin, Chuang Gan, and Song Han. Temporal shift module for efficient video understanding. *arXiv preprint arXiv:1811.08383*, 2018.
- [15] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. *arXiv preprint arXiv:1804.07723*, 2018.
- [16] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [17] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edge-connect: Generative image inpainting with adversarial edge learning. 2019.
- [18] Alasdair Newson, Andrés Almansa, Matthieu Fradet, Yann Gousseau, and Patrick Pérez. Video inpainting of complex scenes. *SIAM Journal on Imaging Sciences*, 7(4):1993–2019, 2014.
- [19] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016.
- [20] Andreas Rösslner, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv preprint arXiv:1803.09179*, 2018.
- [21] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [22] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014.
- [23] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [24] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.

- [25] Chuan Wang, Haibin Huang, Xiaoguang Han, and Jue Wang. Video inpainting by jointly learning temporal structure and spatial details. In *Proceedings of the 33th AAAI Conference on Artificial Intelligence*, 2019.
- [26] Yonatan Wexler, Eli Shechtman, and Michal Irani. Space-time completion of video. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (3):463–476, 2007.
- [27] Junyuan Xie, Linli Xu, and Enhong Chen. Image denoising and inpainting with deep neural networks. In *Advances in neural information processing systems*, pages 341–349, 2012.
- [28] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. *arXiv preprint arXiv:1806.03589*, 2018.
- [29] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. *arXiv preprint*, 2018.
- [30] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4694–4702, 2015.
- [31] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *arXiv preprint*, 2018.