

Free-form Video Inpainting with 3D Gated Convolution and Temporal PatchGAN

Ya-Liang Chang*

Zhe Yu Liu*

Kuan-Ying Lee

Winston Hsu

National Taiwan University, Taipei, Taiwan

{yaliangchang, zhe2325138}@cmlab.csie.ntu.edu.tw, {r03922165, whsu}@ntu.edu.tw

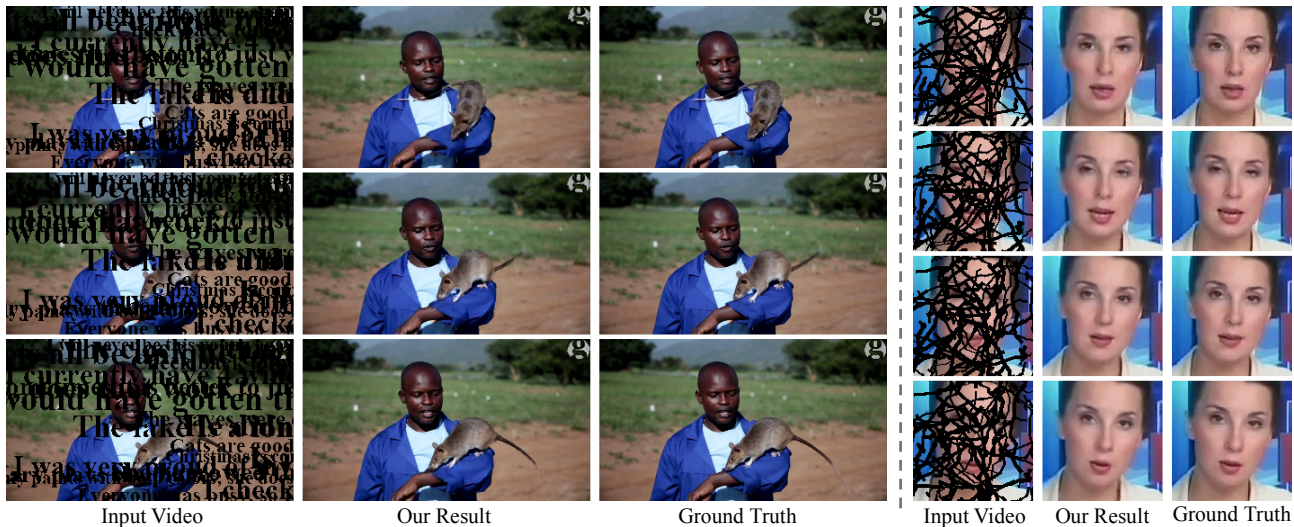


Figure 1. Our model takes incomplete videos with free-form masks (e.g. random text on the left and curves on the right) as inputs, and generates the completed videos as outputs. We propose using 3D gated convolutions to deal with the uncertainty of free-form masks and a combination of designed loss functions to enhance temporal consistency. Best viewed in [videos](#).

Abstract

Free-form video inpainting is a very challenging task that could be widely used for video editing such as text removal (see Fig. 1). Existing patch-based methods could not handle non-repetitive structures such as faces, while directly applying image-based inpainting models to videos will result in temporal inconsistency (see [videos](#)). In this paper, we introduce a deep learning based free-form video inpainting model, with proposed 3D gated convolutions to tackle the uncertainty of free-form masks and a novel Temporal PatchGAN loss to enhance temporal consistency. In addition, we collect videos and design a free-form mask generation algorithm to build the free-form video inpainting (FVI) dataset for training and evaluation of video inpainting models. We demonstrate the benefits of these components and experiments on both the FaceForensics and our

FVI dataset suggest that our method is superior to existing ones. Related source code, full-resolution result videos and the FVI dataset could be found on [Github](#).

1. Introduction

Video inpainting, to recover missing parts in a video, is a very challenging task that remains unsolved. It is a very practical and crucial problem and solving this problem could be beneficial for movie post-production and general video editing. Among them, free-form video inpainting is the most difficult and unconstrained problem because the missing area could be of arbitrary shape (see Fig. 1). In this paper, we propose a novel model to tackle the free-form video inpainting task, and both the quantitative and qualitative evaluations show our model can generate state-of-the-art results with high video quality.

There are many methods proposed for the video inpainting problem, such as patch-based algorithms [11, 13, 22,

*The two authors contributed equally to this paper.

30], which aim to find the most similar patch from the unmasked parts of the video to fill in the masked region. However, patch-based models often fail to recover complex objects that could not be seen or found easily in the unmasked parts of the video. The nearest neighbor search algorithms in patch-based methods also may not work when the ratio of covered area by free-form masks to that of uncovered area is high (see Table 2).

Aside from patch-based methods, many deep learning based models have made tremendous progress on free-form image inpainting. Nevertheless, simply applying image inpainting model to videos tends to cause twisted or flickering results that are temporally inconsistent (See [Edge-Connect](#)).

We extend the work of free-form image inpainting to videos by developing a novel architecture that enhances temporal consistency. The model is learning based, so it could model the data distribution based on the training videos and recover the masked regions. It could even recover objects that are mostly occluded in the video such as the face in Fig. 1, which is impossible for patch-based methods. Besides, our method fully utilizes the temporal information in videos, so the flickering problem of image inpainting is mitigated.

Specifically, we observe that input videos contain many masked voxels that are potentially harmful to vanilla convolutions, we design a generator with 3D gated convolutional layers that could properly handle the masked video by learning the difference between the unmasked region, filled in region and masked region in each layer and attend on proper features correspondingly. In addition, different from image inpainting, video inpainting has to be temporally coherent, so we propose a novel Temporal PatchGAN discriminator that penalizes high-frequency spatial-temporal features and enhances the temporal consistency through the combination of different losses. We also design a new algorithm to generate diverse free-form video masks, and collect a new video dataset based on existing videos that could be used to train and evaluate learning-based video inpainting models.

Our contributions could be summarized as follows:

- We extend the work of image inpainting and propose the first learning-based model for free-form video inpainting and achieve state-of-the-art results qualitatively and quantitatively on the FaceForensics and our dataset.
- We introduce a novel Temporal PatchGAN (T-PatchGAN) discriminator to enhance the temporal consistency and video quality. It could also be extended to other video generation tasks such as video object removal or video super-resolution.
- We design a new algorithm to generate free-form masks. We design and evaluate several types of masks with different mask-to-frame ratios.

- We collect the free-form video inpainting (FVI) dataset, the first dataset to our knowledge for training and evaluation of free-form video inpainting methods, including 1940 videos from the YouTube-VOS [32] dataset and 12600 videos from the YouTube-BoundingBoxes [24] dataset.

2. Related Work

Image Inpainting. Image inpainting, to recover the damaged or missing region in a picture, is firstly introduced in [4]. Many approaches have been proposed to solve the image inpainting task, including diffusion-based [3, 4] and patch-based [2, 5, 8] ones. In general, these methods performs well on simple structure but often fails to generate complex objects or recover large missing area.

Over the past few years, deep learning based methods have made tremendous progress on image inpainting. Xie *et al.* [31] is the first to introduce convolutional neural networks (CNNs) for image inpainting and denoising on small regions. Subsequently, Pathak *et al.* [23] extended image inpainting to larger region with an encoder-decoder structure and used generative adversarial network (GAN) [10] where a generator that strives to create genuine images and a discriminator learns to recognize fake ones are jointly trained to improve the blurry issue caused by the l_2 loss. Yu *et al.* [35] further proposed a contextual attention layer with local and global WGANs to achieve better results.

Free-form Image Inpainting. Recently, image inpainting with irregular holes (free-form masks) caught more attention because it is closer to the real case. Yan *et al.* [33] designed a special shift-connection layer in the U-Net architecture; Lui *et al.* [19] proposed the partial convolution; Yu *et al.* [34] developed the gated convolution with spectral-normalized discriminator to improve free-form image inpainting. Besides from these works, Nazeri *et al.* [21] proposed a two-stage adversarial model EdgeConnect, where the edge generator firstly hallucinates edges of the missing region, and the image completion network generates the final output image using hallucinated edges as a priori. Nazeri *et al.* provided a pretrained model that reaches state-of-the-art, which we set as a baseline in our work.

Although state-of-the-art image inpainting models could recover missing regions in a picture in a reasonable manner, extending these models to videos will cause serious temporal inconsistency as each inpainted frame is different (see [videos](#)).

Video Inpainting. Traditionally, patch-based methods [11, 13, 22, 30] are used for video inpainting. Wexler *et al.* [30] considered video inpainting as a global optimization problem, where all missing regions could be filled in with

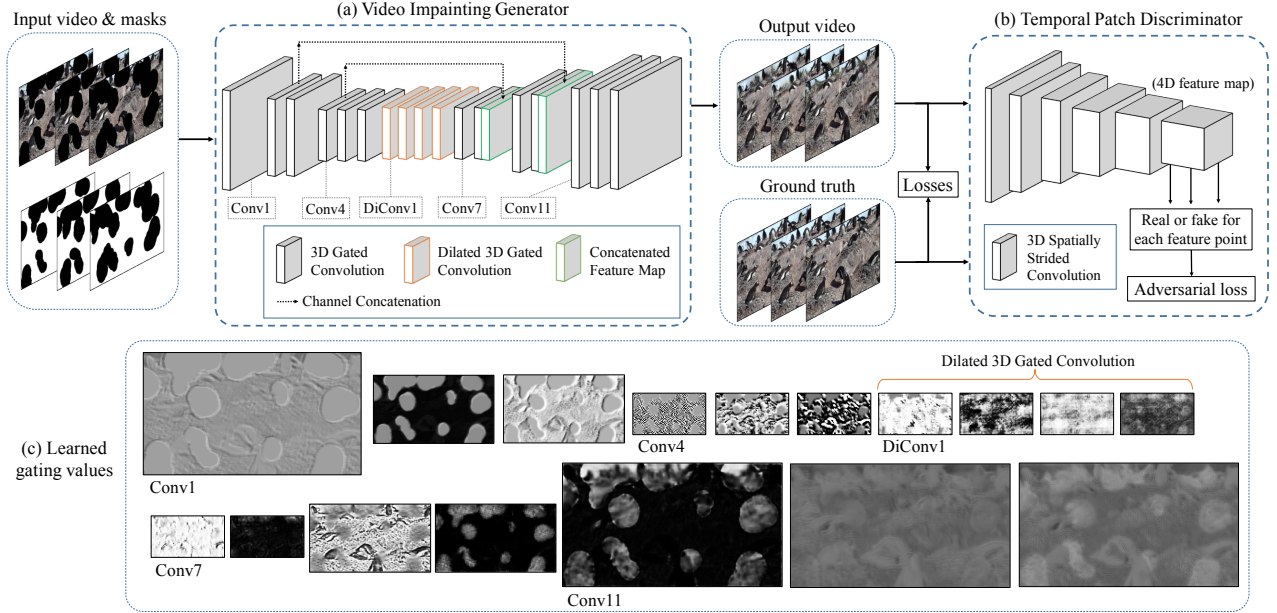


Figure 2. Model architecture and learned gated value visualization. Our model is composed of (a) a video inpainting generator with 3D gated convolutional layers that fully utilizes information for neighboring frames to handle irregular video masks and (b) a Temporal PatchGAN (T-PatchGAN) discriminator that focuses on different spatial-temporal features to enhance output quality. (c) The visualization of learned gating values $\sigma(Gating_{t,x,y})$. The 3D gated convolution will attend on the masked area and gradually fill in the missing feature points. Note that all gating values are extracted from the first channel of each layer without manual picking.

similar patches from the unmasked parts. Afterwards, Newson *et al.* [22] further improved the search algorithm, integrating texture features and initialization scheme. Lastly, Huang *et al.* [13] tackled the moving camera problem by jointly estimating optical flow and colors in the missing regions, and we also take it as our baseline.

State-of-the-art patch-based methods could generate plausible videos under certain conditions, but the computation time of these methods is too high for real-time applications. In addition, patch-based models are limited to repetitive patterns or appeared objects and not feasible to complex structures and large/long-lasting occlusions. The proposed model is learning based that could solve both problems by modeling the distribution of real videos and generate realistic results only by forward inference, without searching.

To solve issues in patch based methods, Wang *et al.* [28] proposed the first deep learning based method CombCN for video inpainting, and we also set it as a baseline. It is a two-stage model with a 3D convolutional network for temporal consistency, followed by a 2D completion network to improve video quality. Although their model could be applied to some random holes during validation/testing stage, it is rather limited so we do not consider it as a free-form video inpainting method. Besides, their model only uses traditional convolution and is trained with the l_1 loss, so the results tend to be blurry in complex scenes [7]. Our model is single-stage, uses gated convolution to attend on valid

features, and integrates perceptual and temporal generative adversarial loss to generate clear and plausible videos for irregular moving masks.

3. Proposed Method

The proposed model (Fig. 2) consists of a generator network G with 3D gated convolution to inpaint videos, and a Temporal PatchGAN discriminator D with several losses.

3.1. Video Inpainting Generator

We extend the single-stage UNet-like network used for image inpainting [19] to video inpainting and integrate the gated convolutional layers in [34] (see Fig. 2 (a)). During training, we combine ground truth video frames $\{V_t \mid t = 1 \dots n\}$ and masks $\{M_t \mid t = 1 \dots n\}$ into masked input video $\{I_t \mid t = 1 \dots n\}$. The model will inpaint the masked region and generate the output video frames $\{O_t \mid t = 1 \dots n\}$.

3.2. Spatial-temporally Aware 3D Gated Conv.

In vanilla convolutional layers, all pixels are treated as valid ones, which makes sense for tasks with real images/videos as inputs, such as object detection or action recognition. However, for inpainting problems, masked regions are filled with black pixels, so input features for convolutional layers include invalid pixels (shallow layers)

or synthesized pixels (deep layers), which should not be treated exactly as normal ones.

To address this problem, we propose the 3D gated convolutions extended from [34] to replace vanilla convolutions in our generator. The 3D convolutions utilizes information from neighboring frames, while the gated convolutions attends on the irregular masked areas; together, 3D gated convolutions could properly handle the uncertainty of free-form video masks. Specifically, for each convolutional layer, an additional gating convolutional filter W_g is applied to input features $F_{t,x,y}$ to obtain a gating $Gating_{t,x,y}$, which is used as an attention map on the output features $Features_{t,x,y}$ from original vanilla convolutional filter W_f according to the validity (see Fig. 2(c)). The t, x, y are the spatial-temporal coordinates of the video. It could be expressed as:

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

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

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

where σ is the sigmoid function to transform gating to values between 0 (invalid) and 1 (valid), and ϕ is the original activation function (e.g. LeakyReLU).

3.3. Loss Functions

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

$$L_{total} = \lambda_{l_1} L_{l_1} + \lambda_{l_{1_{masked}}} L_{l_{1_{masked}}} + \lambda_{perc} L_{perc} + \lambda_{style} L_{style} + \lambda_G L_G \quad (4)$$

where λ_{l_1} , $\lambda_{l_{1_{masked}}}$, λ_{perc} , λ_{style} and λ_G are the weights for l_1 loss, masked l_1 loss, perceptual loss, style loss and Temporal PatchGAN loss, respectively.

Masked l_1 loss. The l_1 loss focuses on the pixel-level features. Since the unmasked area will be pasted onto the final output video, we separate the l_1 loss for all videos:

$$L_{l_1} = \mathbb{E}_{t,x,y}[|O_{t,x,y} - V_{t,x,y}|] \quad (5)$$

and the l_1 loss for the masked region:

$$L_{l_{1_{masked}}} = \mathbb{E}_{t,x,y}[M_{t,x,y}|O_{t,x,y} - V_{t,x,y}|] \quad (6)$$

Perceptual loss. Perceptual loss is firstly proposed in [9] to keep image contents for style transfer, and is now widely used for image inpainting [19, 21] and super-resolution [15, 17] to mitigate the blurriness caused by the l_1 loss. The perceptual loss computes the l_1 loss in feature level:

$$L_{perc} = \sum_{t=1}^n \sum_{p=0}^{P-1} \frac{|\Psi_p^{O_t} - \Psi_p^{V_t}|}{N_{\Psi_p^{V_t}}} \quad (7)$$

where $\Psi_p^{V_t}$ denotes the activation from the p th selected layer of the pretrained network given the input V_t , and $N_{\Psi_p^{V_t}}$ is the number of elements in the p th layer. We use layer $relu_{2,2}$, $relu_{3,3}$ and $relu_{4,3}$ from the VGG [27] network pre-trained on ImageNet [26].

Style loss. We also include the style loss, which is introduced in [9] to keep the image style for style transfer, and also used in image inpainting [19, 21]. Style loss is similar to perceptual loss, except that an auto-correlation (Gram matrix) is firstly applied to the features:

$$L_{style} = \sum_{t=1}^n \sum_{p=0}^{P-1} \frac{1}{C_p C_p} \frac{|(\Psi^{O_{t_p}})^T(\Psi^{O_{t_p}}) - (\Psi_p^{V_t})^T(\Psi_p^{V_t})|}{C_p H_p W_p} \quad (8)$$

where $\Psi^{O_{t_p}}$ and $\Psi^{V_{t_p}}$ are both VGG features in the shape of (H_p, W_p, C_p) as the ones in perceptual loss 7.

Temporal PatchGAN loss. For the free-form video inpainting problem, masks could be anywhere in a video, so we should consider global and local features in each frame, and the temporal consistency of these features. A naive idea will be applying a loss function for each of the three aspects respectively. However, empirically we found that it is hard to balance the weights of these loss functions, especially when some of them are GAN losses (adding GAN loss is a very common strategy to make image inpainting results more realistic [21, 35, 34]).

Yu *et al.* proposed an efficient SN-PatchGAN [34], which applies GAN loss on feature maps of the discriminator to replace the use of global and local GAN and tackle the problem that masks could be anywhere and of any form. Although their work tackles the balancing issue between GAN losses and solves free-form image inpainting problem, it does not consider temporal consistency, a pivotal factor for high-quality video inpainting. Inspired by their work, we further integrate the temporal dimension and design a novel Temporal PatchGAN (T-PatchGAN) discriminator that focuses on different spatial-temporal features to fully utilize all the global and local image features and temporal information together.

Our T-PatchGAN discriminator is composed of 6 3D convolutional layers with kernel size $3 \times 5 \times 5$ and stride $1 \times 2 \times 2$. The recently proposed spectral normalization [20] is applied to both the generator and discriminator, similar to [21] to enhance training stability. In addition, we use the hinge loss as the objective function as to discriminate if the input video is real or fake:

$$L_D = \mathbb{E}_{x \sim P_{data}(x)}[ReLU(1 + D(x))] + \mathbb{E}_{z \sim P_z(z)}[ReLU(1 - D(G(z)))] \quad (9)$$

$$L_G = -\mathbb{E}_{z \sim P_z(z)}[D(G(z))] \quad (10)$$

where G is the video inpainting network that takes input video z and D is the T-PatchGAN discriminator.

Note that we use kernel size $3 \times 5 \times 5$ for each layer in the discriminator, so the receptive field of each output feature covers the whole videos, and a global discriminator like [35] is not needed. The T-PatchGAN learns to classify each

spatial-temporal patch as real or fake, which restricts it to focus on high-frequency features because it only penalizes at the scale of patches. As the l_1 loss already focus on low-frequency features, using T-PatchGAN could improve the output video quality in an efficient way.

3.4. Free-form Video Masks Generation

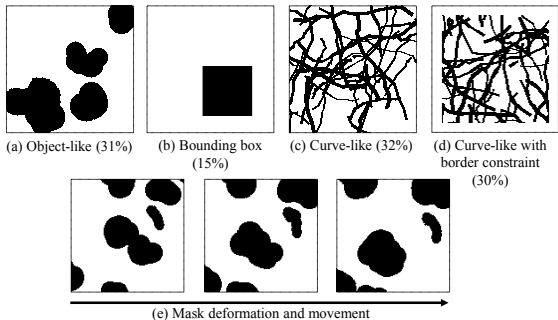


Figure 3. Masks generated by our algorithm with different mask types and mask-to-image ratios. The components in a mask video may move and deform independently as shown in (e).

Training data is extremely important for learning based methods, and the generation of our input mask videos is non-trivial as it should consider different scenarios to be “free-form”. There is no existing database or algorithm to generate such free-form video masks, so we develop an video mask generation algorithm based on the image one by [34] (see Algorithm 1 in the supplementary materials).

The image mask generation in [34] uses several strokes drawing on a blank image to represent a mask. Each stroke is composed of a ordered set of control points, which is determined by the trace of a head point initialized at a random position and repeatedly moving to a nearby position.

Additionally, for free-form video masks, we introduce the concept of motion: strokes may move and deform over time (see Fig. 3(e)). Stroke deformation is achieved by randomly moving each control point of a stroke with a certain probability. For the movement, the concepts of velocity and acceleration are applied on the strokes. The initial speeds of strokes are sampled from a normal distribution centered on 0 since most objects in videos do not have large speeds.

As stated in [19], many methods [14, 19] have degraded performance when masks cover the border. Therefore, aside from motion simulation, we also take such border constraints into consideration. That is, we generate both masks that either cover or do not cover the edges of the frame (see Fig. 3. Masks without border constraint are more difficult since convolutional filters will have no valid pixels as inputs in the masked border areas.

Moreover, we consider three different types of masks: long thin curve-like and round thick object-like masks generated with different hyper-parameters in our mask generation algorithm along with bounding-box masks, as shown in

Fig. 3. The curve-like masks are considered easier as most masked areas are close to valid pixels (unmasked area), while object-like and bounding-box masks are challenging since it is hard to inpaint large invalid voxels.

Totally 28,000 free-form videos with mask-to-frame ratio from 0 - 10% to 60 - 70% are generated for training. And for each mask type, 100 videos are generated for testing.

4. Experimental Results

4.1. Datasets

FaceForensics. We compare with [28] on the FaceForensics dataset [25], which contains 1004 face videos from YouTube and the YouTube-8m dataset [1] with face, newscaster or newsprogram tags. The videos are cropped into 128×128 with the face in the middle during the data preparation stage, following the setting in [28]. Among them, 150 videos are used for evaluation and the rest are used for training. The FaceForensics dataset is rather easy for learning-based methods since the data is less diverse.

Free-form video inpainting (FVI) dataset To test on more practical cases, we collect videos from the YouTube-VOS [32] dataset and the YouTube-BoundingBoxes dataset [24]. The former has about 2000 videos with 94 categories of object segmentation in 6 frame per second (FPS) and the latter has about 380,000 videos with 23 kinds of object bounding boxes in 30 FPS. We choose videos with resolution higher than 640×480 and manually filter out videos with shot transitions. We set 100 videos from the YouTube-VOS as testing set, while the training set includes about 15,000 videos. Together with the 28,000 free-form mask videos, we build the FVI dataset, the first dataset for free-form video inpainting, to the best of our knowledge.

Our FVI dataset is very challenging for the video inpainting task due to the high diversity, including different kinds of objects, animals and human activities. All videos are from YouTube, closer to real-world scenario. Moreover, the provided object segmentation and bounding boxes could be used to test video object removal.

For the experiments, we only use 1940 videos from the training set as we do not witness a significant improvement for our model using the full training set. During data pre-processing stage, we resize videos to 384×216 and randomly crop them to 320×180 with random horizontal flip.

4.2. Evaluation metrics

We use mean square error (MSE) and Learned Perceptual Image Patch Similarity (LPIPS) [36] to evaluate the image quality. Furthermore, to evaluate the video quality and temporal consistency, we also calculate the Frchet Inception Distance (FID) [12] with I3D [6] pre-trained video recognition CNN as Vid2vid [29]. See the supplementary materials for the details.

4.3. Quantitative Results

We evaluate our model on the FaceForensics and FVI testing set with free-form masks in 7 ranges of mask-to-frame ratio from 0 - 10% to 60 - 70% (higher mask-to-frame ratio makes the task more difficult, see Fig. 4). The state-of-the-art patch-based video inpainting method TCCDS by Huang *et al.* [13], image inpainting method Edge-Connect(EC) by Nazeri *et al.* [21] and learning based video inpainting method CombCN by Wang *et al.* [28] are set for comparison. We train Nazeri *et al.*'s model and Wang *et al.*'s model on our dataset. Note that Wang *et al.*'s model is originally trained on bounding boxes, but we train it with our free-form mask for fair comparison.

From Table 1 we could see that the FaceForensics dataset is easy for learning-based models but not for the patch-based method TCCDS [13], because face features are non-repetitive and hence cannot be recovered with nearby patches. Yet, the overall structure of faces are learnable and thus learning-based methods are favorable. Compared with the two deep learning based methods, our model has superior performance on curve-like and object-like masks since it fully utilizes information of neighboring frames to recover the missing areas by 3D convolutions and the proposed T-PatchGAN loss. As for bounding-box masks, our model outperforms CombCN while on par with EC. Note our model is only trained on FVI. Hence, for fair comparison, we train EC from scratch on FVI without having it pre-trained on Celeb-A as stated in the original paper.

On the other hand, Table 2 shows that our FVI dataset is more challenging for learning-based methods for its high diversity. It is rather difficult for learning-based models to capture the distribution of the highly diverse masks, while patch-based methods like TCCDS could easily find realistic enough patches to fill in the mask given the mask is not large. Nonetheless, we could notice that for some masks, the nearest neighbor search in TCCDS fails to find candidates when most patches are covered by the mask. Note that CombCN is only trained with the l_1 loss, so while it reports a lower MSE, its results are actually blurry, bearing high perceptual distance to ground truths (high LPIPS). Our method generates clear results (low LPIPS and FID) and demonstrates high temporal consistency (low FID), which is of crucial importance in video inpainting task.

4.4. Qualitative Results

We also demonstrate the visual comparison in Fig. 5 with the corresponding video link. As mentioned in 4.3, CombCN's [28] outputs are blurry due to the l_1 loss, TCCDS [13] may paste wrong patches, and Edge-Connect [21] will have flickering results (best viewed in videos). Our model could generate reasonable frames with high temporal consistency.

In addition, our trained model can be easily applied on

object removal, as shown in Fig. 7. More visual comparisons could be found in the supplementary materials.

	Mask Type	TCCDS	EC	CombCN	3DGated (Ours)
MSE↓	Curve	0.0031*	0.0022	0.0012	0.0008
	Object	0.0096*	0.0074	0.0047	0.0048
	BBox	0.0055	0.0019	0.0016	0.0018
LPIPS↓	Curve	0.0566*	0.0562	0.0483	0.0276
	Object	0.1340*	0.0761	0.1353	0.0743
	BBox	0.1260	0.0335	0.0708	0.0395
FID↓	Curve	1.281*	0.848	0.704	0.472
	Object	1.107*	0.946	0.913	0.766
	BBox	1.013	0.663	0.742	0.663

Table 1. Quantitative results on the FaceForensics testing set with masks without border. Our model has superior performance for the curve-like and object-like masks. *TCCDS fails on some masks; the results are average of the successful cases.

	Mask Type	TCCDS	EC#	CombCN	3DGated (Ours)
MSE	Curve	0.0219*	0.0047	0.0021	0.0024
	Object	0.0110*	0.0079	0.0049	0.0056
LPI	Curve	0.2838*	0.1204	0.0794	0.0521
	Object	0.2001*	0.1420	0.2054	0.1078
FID	Curve	2.105*	1.033	0.766	0.609
	Object	1.287*	1.083	1.091	0.905

Table 2. Quantitative comparison on the FVI testing set without border. The results are the average of different mask-to-frame ratios (see the supplementary materials for original data). Our model outperforms the baselines for perceptual distance (LPIPS) and temporal consistency (FID). CombCN has better MSE but their results are blurry (see Fig. 5). #Pretrained on Places2 [37].




Figure 4. Effect of mask size on LPIPS and FID on the FVI test set with curve-like masks. Larger masks are harder for all methods.

4.5. User Study

Aside from qualitative comparison, we also conduct a human subjective study to evaluate our method. During the study, we display a pair of result videos (ours against baselines or ground truth, in random sequence), and ask subjects to choose the more realistic and consistent one. The mask video is shown meanwhile for reference. For each




Figure 5. Visual comparison with the baselines. TCCDS: pasting wrong patches; Edge-Connect: inconsistent between frames; CombCN: blurry. See [videos](#)

mask type (object-like and curve-like) and mask-to-hole ratio (0-10% to 60-70%), we randomly select 20 video pairs to compare and each video pair is presented to 5 subjects.

Results from 150 participants are shown in Fig. 6. Our model outperforms the baselines in both object-/curve-like masks for all mask-to-frame ratios. In addition, when compared with ground truth, our method still has 23% preference on average, which indicates subjects could not tell our results and the original videos apart in 23% cases.

4.6. Ablation Study

We conduct an ablation study to evaluate the contribution of each proposed component. From Table 3 we could see




Figure 6. User preference on the FVI testing set (ours versus baselines and ground truth). 50% means that the two methods are equally good. Our model outperforms the baselines in both the object-like and curve-like masks for all mask-to-frame ratios. When compared with ground truth (GT), our method could still have about 23% preference for curve-like masks.

that 3D convolution and T-PatchGAN are both crucial because the two components provide a great amount of temporal information by 3D convolutions. Corresponding video comparisons could be found on [YouTube](#).

3D conv.	Gated conv.	T-Patch GAN	LPIPS↓	FID↓
	✓	✓	0.1769	1.243
✓		✓	0.1321	1.121
✓	✓		0.1716	1.201
✓	✓	✓	0.1209	1.034

Table 3. Ablation study on the FVI dataset with object-like masks. We can see that all components are important. We set up all models with about the same number of parameters (i.e., increase the channel number for 2D convolution and vanilla convolution) to exclude the gain from additional parameters.

4.7. Extension to Video Super-Resolution

Our model could be extended to video super-resolution, interpolation or prediction by using proper masks. For video super-resolution, given a low-resolution video with width W , height H , length L and up-sampling factor K , we construct the input mask video in shape $(W \times K, H \times K, L)$ where each pixel (x, y, t) is masked if x or y is not a multiple of K . For frame interpolation, masks could be added between frames. In Fig. 8 and Table 4, we compare our method with well-known super-resolution methods SRResNet and SRGAN in [17]. Our model could generate plausible high-resolution videos with low perceptual distance. See [result videos](#).

	Bicubic	SRResNet	SRGAN	Ours
MSE↓	0.0089	0.0044	0.0074	0.0076
LPIPS↓	0.5141	0.3582	0.1785	0.1631
FID↓	1.502	1.083	1.035	1.096

Table 4. Quantitative comparison for spatial super-resolution on the VOR testing set for 4x up-sample. We can see that our model could reach low perceptual quality.




Figure 7. Our model could be easily extended to video object removal. See [videos](#).

5. Discussion and Future Work

Our model fails when the testing video is very different from the training data as most learning-based methods do. In addition, when the masked area is too thick, our model fails to generate natural results. Still, compared with the baselines, our model performs better under the two conditions (see [videos](#)).

Besides, compared to 2D convolutions, 3D convolutions require more parameters that could lead to higher redundancy. Nonetheless, our model is single-stage, feed-forward and does not depend on optical flows, so the inference speed is fast, and the performance gain is significant. A potential solution to reduce the number of parameters is to integrate the Temporal Shift Module [18] so that 2D convolutions could deal with temporal information.

Also, we found that we could reach a similar performance to that of gated convolutions by simply increasing the number of channels in the ablation study. This may im-

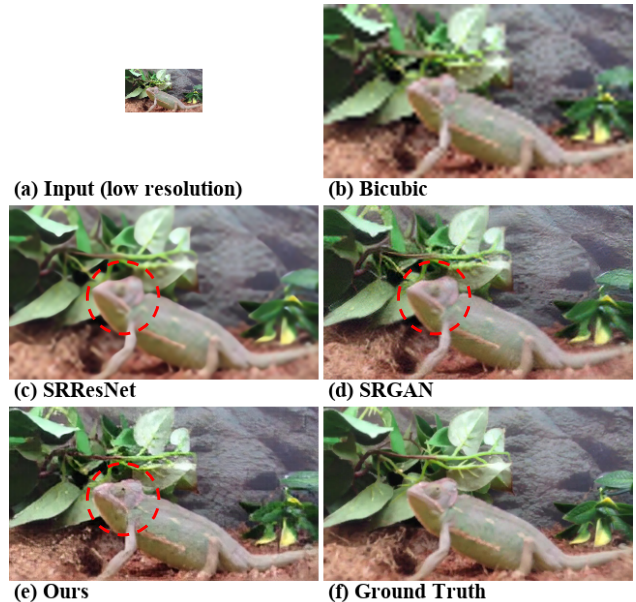


Figure 8. Two examples of spatial super-resolution with 4x up-sampling. Compared with the two baselines, our model could recover the eyes of the chameleon. See corresponding [videos](#).

ply that our model still underfits the FVI dataset, or gated convolutions have less impacts for video inpainting compared to image inpainting. Potential future works would be to compare and improve convolutional layers for free-form video inpainting or combine with optical flows [16].

Another future work is to integrate the user guided inputs as [34, 21] by training the model with edge images of video frames as additional inputs. The model could generate more plausible results given the object shape information from edges. Additionally, during inference, users could draw lines to change the edge images to manipulate the output videos, which is useful for video editing.

6. Conclusion

In this paper, we proposed the first learning based free-form video inpainting network to our knowledge, using 3D gated convolution and a novel GAN loss. We demonstrate the power of 3D gated convolution and temporal PatchGAN to enhance video quality and temporal consistency in the video inpainting task. Our system could also be extended to video object removal, video super-resolution or video interpolation. Both the quantitative and qualitative results indicate that our model achieve state-of-the-art 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] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016.
- [2] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. In *ACM Transactions on Graphics (ToG)*, volume 28, page 24. ACM, 2009.
- [3] M. Bertalmio, A. L. Bertozzi, and G. Sapiro. Navier-stokes, fluid dynamics, and image and video inpainting. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–I. IEEE, 2001.
- [4] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 417–424. ACM Press/Addison-Wesley Publishing Co., 2000.
- [5] R. Bornard, E. Lecan, L. Laborelli, and J.-H. Chenot. Missing data correction in still images and image sequences. In *Proceedings of the tenth ACM international conference on Multimedia*, pages 355–361. ACM, 2002.
- [6] J. Carreira and A. 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.
- [7] Y.-L. Chang, Z. Y. Liu, and W. Hsu. Vornet: Spatio-temporally consistent video inpainting for object removal. *arXiv preprint arXiv:1904.06726*, 2019.
- [8] I. Drori, D. Cohen-Or, and H. Yeshurun. Fragment-based image completion. In *ACM Transactions on graphics (TOG)*, volume 22, pages 303–312. ACM, 2003.
- [9] L. A. Gatys, A. S. Ecker, and M. Bethge. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*, 2015.
- [10] 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.
- [11] M. Granados, J. Tompkin, K. Kim, O. Grau, J. Kautz, and C. 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.
- [12] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. 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.
- [13] J.-B. Huang, S. B. Kang, N. Ahuja, and J. Kopf. Temporally coherent completion of dynamic video. *ACM Transactions on Graphics (TOG)*, 35(6):196, 2016.
- [14] S. Iizuka, E. Simo-Serra, and H. Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):107, 2017.
- [15] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [16] D. Kim, S. Woo, J.-Y. Lee, and I. S. Kweon. Deep video inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5792–5801, 2019.
- [17] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. 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.
- [18] J. Lin, C. Gan, and S. Han. Temporal shift module for efficient video understanding. *arXiv preprint arXiv:1811.08383*, 2018.
- [19] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro. Image inpainting for irregular holes using partial convolutions. *arXiv preprint arXiv:1804.07723*, 2018.
- [20] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [21] K. Nazeri, E. Ng, T. Joseph, F. Qureshi, and M. Ebrahimi. Edgeconnect: Generative image inpainting with adversarial edge learning. 2019.
- [22] A. Newson, A. Almansa, M. Fradet, Y. Gousseau, and P. Pérez. Video inpainting of complex scenes. *SIAM Journal on Imaging Sciences*, 7(4):1993–2019, 2014.
- [23] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016.
- [24] E. Real, J. Shlens, S. Mazzocchi, X. Pan, and V. Vanhoucke. Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5296–5305, 2017.
- [25] A. Rössler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv preprint arXiv:1803.09179*, 2018.
- [26] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [28] C. Wang, H. Huang, X. Han, and J. Wang. Video inpainting by jointly learning temporal structure and spatial details. In *Proceedings of the 33th AAAI Conference on Artificial Intelligence*, 2019.
- [29] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, G. Liu, A. Tao, J. Kautz, and B. Catanzaro. Video-to-video synthesis. *arXiv preprint arXiv:1808.06601*, 2018.
- [30] Y. Wexler, E. Shechtman, and M. Irani. Space-time completion of video. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (3):463–476, 2007.

- [31] J. Xie, L. Xu, and E. Chen. Image denoising and inpainting with deep neural networks. In *Advances in neural information processing systems*, pages 341–349, 2012.
- [32] N. Xu, L. Yang, Y. Fan, J. Yang, D. Yue, Y. Liang, B. Price, S. Cohen, and T. Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 585–601, 2018.
- [33] Z. Yan, X. Li, M. Li, W. Zuo, and S. Shan. Shift-net: Image inpainting via deep feature rearrangement. *arXiv preprint arXiv:1801.09392*, 2018.
- [34] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang. Free-form image inpainting with gated convolution. *arXiv preprint arXiv:1806.03589*, 2018.
- [35] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang. Generative image inpainting with contextual attention. *arXiv preprint*, 2018.
- [36] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. *arXiv preprint*, 2018.
- [37] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.