

000  
001           **Flow-edge Guided Video Completion**  
002           **Supplementary Material**  
003  
004

005           Anonymous ECCV submission  
006  
007           Paper ID 1715  
008  
009

## 010           Overview

011           In this supplementary document, we provide additional implementation details  
012           and results to complement the main manuscript.  
013

- 014           1. We summarize the complete pipeline of our algorithm in pseudo-code in  
015           Algorithm 1 and visualize the steps in [algorithm\\_illustration.mp4](#).  
016           2. We show the runtime analysis and profiling to analyze the speed of our  
017           algorithm.  
018           3. We describe the training details for the edge completion network.  
019           4. We demonstrate the ability of our method to handle high-resolution (up to  
020           2K) videos.  
021           5. We present additional visual examples of the ablation study, highlighting the  
022           effectiveness of our design choices.  
023           6. We provide detailed per-sequence results in terms of PSNR, SSIM, and LPIPS  
024           on the DAVIS dataset.  
025           7. We show extensive video comparison between our method and the state-of-  
026           the-arts. Please find them in [index.html](#).
- 027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044

## 045 1 Algorithm

046  
047 We show our method pipeline in [algorithm\\_illustration.mp4](#). We summarize our  
048 complete pipeline in Algorithm 1. Our pipeline consists of three main components:  
049 **flow prediction**, **edge-guided flow completion**, and **video completion**. We will  
050 release the pre-trained flow-edge completion model as well as the source code to  
051 facilitate future research.  
052

---

**053 Algorithm 1:** Summary of our video completion algorithm.

---

```

054 1 Input: Color frames  $I_1 \dots I_n$ , mask frames  $M_1 \dots M_n$ .
055 2 Output: Completed frames  $I_1 \dots I_n$  (updated in place).
056 3 for every frame  $i \in 1 \dots n$  do
057   4 | Compute local and non-local optical flow  $F_{i \rightarrow j}$ ,
058   5 |  $j \in \{i - 1, i + 1, 1, \lceil n/2 \rceil, n\}$  (Equations 1 and 2).
059 6 for each computed flow field  $F_{i \rightarrow j}$  do
060   7 | Extract flow edges  $E_{i \rightarrow j}$  using Canny edge detector [1].
061   8 | Complete flow edges  $\tilde{E}_{i \rightarrow j}$  using EdgeConnect [3] edge model.
062   9 | Complete flow  $\tilde{F}_{i \rightarrow j}$  with edge guidance (Equation 3).
063 10 | Compute flow error  $\tilde{D}_{i \rightarrow j}$  (Equation 4).
064 11 while any missing pixels exist in  $M_1 \dots M_n$  do
065   12 | for every frame  $i \in 1 \dots n$  do
066     13 | | Obtain temporal neighbors through propagation.
067     14 | | Fuse gradient images  $\tilde{G}_{x,i}$  and  $\tilde{G}_{y,i}$  (Equation 7).
068     15 | | Reconstruct color image  $\tilde{I}_i$  (Equation 8).
069     16 | | Update mask  $M_i(p) = 0$ , where  $|N(p)| \geq 1$ .
070     17 | Select frame  $\tilde{I}_f$  with most remaining missing pixels.
071     18 | Complete  $\tilde{I}_f$  with DeepFill [5].
072     19 | Set  $M_f = 1$  (all pixels in this frame).
073     20 | Set  $I = \tilde{I}$ .
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
```

---

## 2 Runtime analysis and profiling

085 Following [2], we also show the detailed running time analysis of our method in  
086 Table 1. We report the time for each component of our method on the “CAMEL”  
087 video sequence under the object removal setting. The resolution is  $960 \times 512$ .  
088 There are 10721523 pixels being removed, which is 9.1% of the total pixels. Our  
089 method runs at 7.2 frames per minute.

090  
091 Table 1: **Running time analysis.** We report the running time for each compo-  
092 nent of our method on the “CAMEL” video sequence under the object removal  
093 setting. The resolution is  $960 \times 512$ .

	Component	Time (min.)
Flow completion	Flow prediction	1.20
	Edge extraction and completion	0.45
	Edge-guided flow completion	4.20
Video completion	Temporal propagation	4.31
	Spatial inpainting	0.10
	Poisson blending	2.29
Total		12.55

### 102 3 Training details

103 The only trainable component in our method is the flow edge completion network.  
104 We build our flow edge completion network upon the publicly available official im-  
105 plementation of EdgeConnect [3] edge model<sup>1</sup>. We load weights pretrained on the  
106 Places2 dataset [6], and then finetune on 60 sequences in DAVIS 2017-test-dev  
107 and 2017-test-challenge for three epochs.

108 Starting from the predicted flow between adjacent frames  $i$  and  $j$ ,  $\mathbf{F}_{i \rightarrow j}$ , we  
109 first calculate the flow magnitude image. We use the Canny edge detector [1]  
110 to extract a flow edge map  $\mathbf{E}_{i \rightarrow j}$  from the flow magnitude image. We use the  
111 following parameters for the Canny edge detector [1]: Gaussian  $\sigma = 1$ , low  
112 threshold 0.1, high threshold 0.2. We randomly choose a mask from NVIDIA  
113 Irregular Mask Dataset testing split and resize it to  $256 \times 256$ .<sup>2</sup> We crop the flow  
114 edge map  $\mathbf{E}_{i \rightarrow j}$  and the corresponding flow magnitude images to  $256 \times 256$ , and  
115 corrupt them with the mask. The input to the flow edge completion network is  
116 the mask (Figure 1a), the corrupted flow edge map (Figure 1b) and the corrupted  
117 flow magnitude image (Figure 1c). We train the network to complete the flow  
118 edge using batches of 8 randomly cropped  $256 \times 256$  patches.

119 Note that our edge completion network does *not* receive any additional  
120 information regarding the stationary mask with a uniform grid of  $5 \times 4$  square  
121 blocks during training.

### 122 4 High resolution videos

123 Our proposed method is memory efficient. Our method handles videos with  
124 up to 4K resolution, while other methods fail due to excessive GPU memory  
125 requirements. Due to the file size limit, we show five evenly selected frames of  
126 the object removal result on the 2K “HIKE” sequence in the HIKE\_2K folder.

133 <sup>1</sup> <https://github.com/knazeri/edge-connect>

134 <sup>2</sup> [https://www.dropbox.com/s/01dfayns9s0kevy/test\\_mask.zip?dl=0](https://www.dropbox.com/s/01dfayns9s0kevy/test_mask.zip?dl=0)

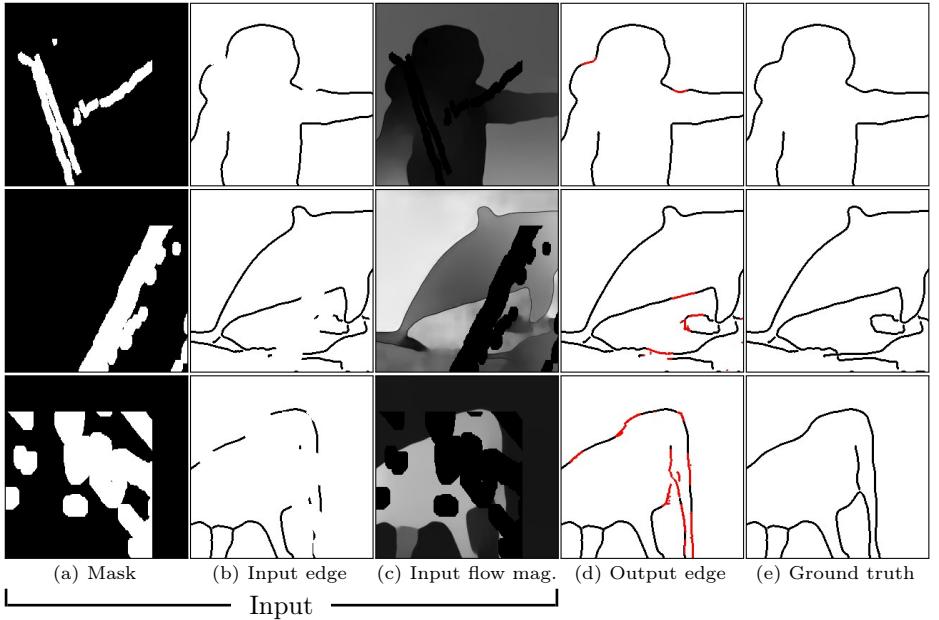


Fig. 1: **Flow edge completion.** Our flow edge completion network takes the mask, the corrupted flow edge map and the corrupted flow magnitude image as input, and complete the flow edge.

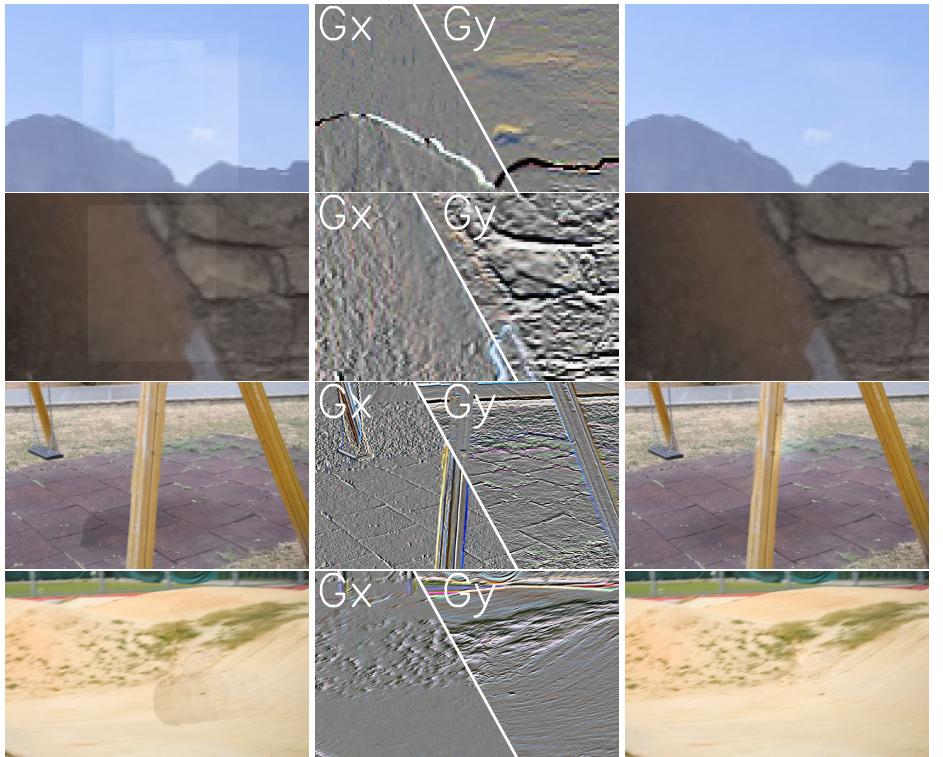
## 5 Additional visual examples of the ablation study

In this section, we show additional visual examples of the ablation study to highlight the effectiveness of our design choices.

**Gradient domain processing.** We compare the proposed gradient propagation process with color propagation (used in [2,4]). Figure 2 shows a visual comparison. When filling the missing region with directly propagated colors, the result contains visible seams due to color differences in different source frames (Figure 2a). Our method operates in the gradient domain and does not suffer from such artifacts (Figure 2c).

**Non-local temporal neighbors.** We study the effectiveness of the non-local temporal neighbors. Figure 3 shows such an example. Using non-local neighbors enables us to transfer the correct contents from temporally distant frames.

**Edge-guided flow completion.** We evaluate the performance of completing the flow field with different methods. In Figure 4, we show examples of flow completion results using diffusion, a trained flow completion network [4], and our proposed edge-guided flow completion. The diffusion-based method maximizes smoothness in the flow field everywhere and thus cannot create sharp motion boundaries. The learning-based flow completion network [4] fails to predict a



203 (a) Color propagation    (b) Propagated x/y gradient    (c) Reconstruction  
 204 Fig. 2: **Gradient domain reconstruction.** Previous methods operate directly  
 205 in the color domain, which results in visible seams in the completed video (a).  
 206 We propagate in the gradient domain (b), and reconstruct the results via Poisson  
 207 reconstruction (c).

210 smooth flow field and sharp flow edges. In contrast, the proposed edge-guided  
 211 flow completion fills the missing region with a piecewise-smooth flow and no  
 212 visible seams along the hole boundary.

## 214 6 Per-sequence results on the DAVIS dataset

215 We show the detailed per-sequence results in terms of PSNR, SSIM, and LPIPS  
 216 under stationary screen-space masks setting in Figure 5. Our proposed method  
 217 improves the performance over state-of-the-art methods for most of the video  
 218 sequences, under all three metrics.



Without non-local neighbors

With non-local neighbors

Fig. 3: **Non-local temporal neighbor ablation.** Video completion results *with* and *without* non-local temporal neighbors. The result without non-local neighbors (left) does not recover well from the lack of well-propagated content.

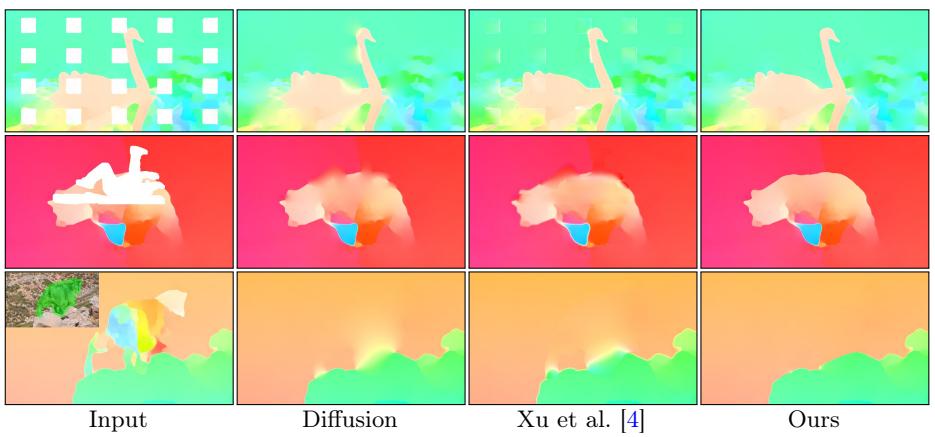


Fig. 4: **Flow completion.** Comparing different methods for flow completion. Our method has better ability to retain the piecewise-smooth nature of flow fields (sharp motion boundaries, smooth everywhere else) than the other two methods.

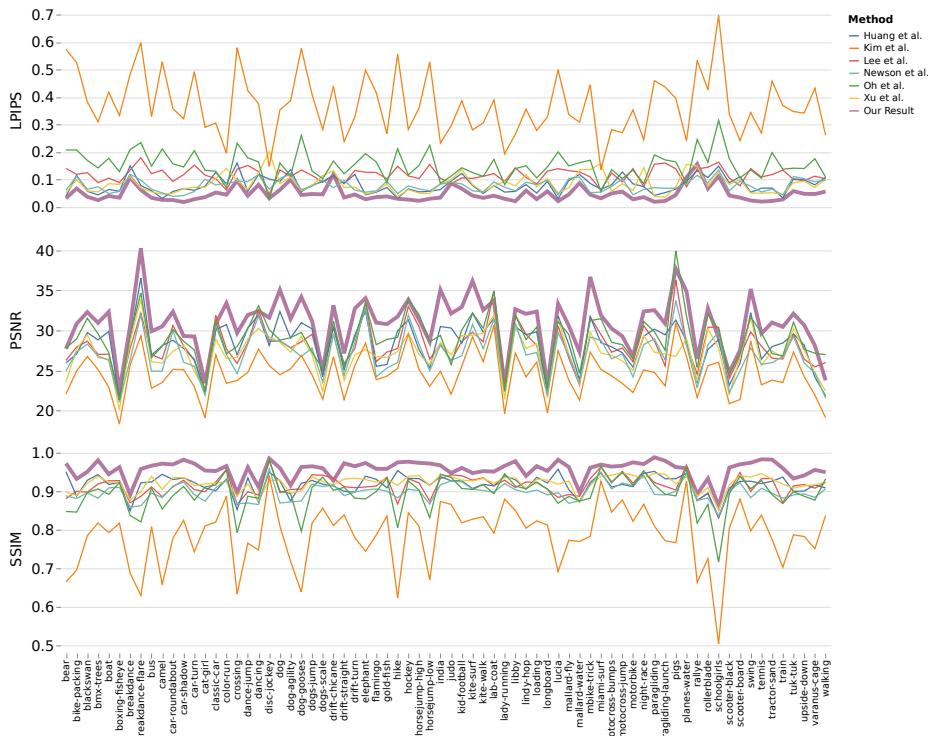


Fig. 5: Per-sequence PSNR, SSIM and LPIPS on DAVIS under the stationary masks inpainting setting.

## 315 References

- 316 1. Canny, J.: A computational approach to edge detection. *IEEE Trans. Pattern Anal.*  
317 *Mach. Intell.* pp. 679–698 (1986)
- 318 2. Huang, J.B., Kang, S.B., Ahuja, N., Kopf, J.: Temporally coherent completion of  
319 dynamic video. *ACM Transactions on Graphics (TOG)* (2016)
- 320 3. Nazeri, K., Ng, E., Joseph, T., Qureshi, F., Ebrahimi, M.: Edgeconnect: Generative  
321 image inpainting with adversarial edge learning. In: *ICCVW* (2019)
- 322 4. Xu, R., Li, X., Zhou, B., Loy, C.C.: Deep flow-guided video inpainting. In: *CVPR*  
323 (2019)
- 324 5. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting  
325 with contextual attention. In: *CVPR* (2018)
- 326 6. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million  
327 image database for scene recognition. *IEEE transactions on pattern analysis and  
328 machine intelligence* **40**(6), 1452–1464 (2017)