This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Efficient Flow-Guided Multi-frame De-fencing

Stavros Tsogkas

s Fengjia Zhang Allan Jepson Alex Levinshtein Samsung AI Center Toronto 101 College St., Toronto, ON, Canada, M5G 1L7

{stavros.t, f.zhang2, allan.jepson, alex.lev}@samsung.com

Abstract

Taking photographs "in-the-wild" is often hindered by fence obstructions that stand between the camera user and the scene of interest, and which are hard or impossible to avoid. De-fencing is the algorithmic process of automatically removing such obstructions from images, revealing the invisible parts of the scene. While this problem can be formulated as a combination of fence segmentation and image inpainting, this often leads to implausible hallucinations of the occluded regions. Existing multi-frame approaches rely on propagating information to a selected keyframe from its temporal neighbors, but they are often inefficient and struggle with alignment of severely obstructed images. In this work we draw inspiration from the video completion literature, and develop a simplified framework for multi-frame de-fencing that computes high quality flow maps directly from obstructed frames, and uses them to accurately align frames. Our primary focus is efficiency and practicality in a real world setting: the input to our algorithm is a short image burst (5 frames) – a data modality commonly available in modern smartphones- and the output is a single reconstructed keyframe, with the fence removed. Our approach leverages simple yet effective CNN modules, trained on carefully generated synthetic data, and outperforms more complicated alternatives real bursts, both quantitatively and qualitatively, while running real-time.

1. Introduction

Rapid improvements in both the camera hardware and image processing software have turned modern cell phones into powerful yet portable image and video recording devices. This has enabled and encouraged casual users to shoot photos without any time for special preparation, setup, or framing of the shot. On the flip side, photos and videos taken under these conditions rarely contain just the object(s) of interest, and are hindered by various obstructions that stand between the subject and the user.

One type of obstruction that is of special interest because



Figure 1: We train a simple and efficient model for **defencing**: removing fence obstructions from images, revealing the underlying scene of interest. Our method is fast and can accurately remove fences of varying size and appearance from *real* bursts, without online finetuning.

of its commonness is *fences*. Imagine, for instance, taking a photo of an animal through a zoo fence, or people playing basketball in a fenced-out outdoor court; these are just a few everyday scenes that are obstructed by fence structures that are either inconvenient or completely impossible to avoid. *De-fencing* uses computer vision algorithms to automatically remove such fence obstructions from images, as shown in Figure 1. De-fencing is a harder problem than it may initially seem. Fences have varying structure and appearance patterns, and come in different thicknesses and sizes. Furthermore, reconstruction of the background scene can become challenging because of low lighting and noise, or motion blur, caused by rapidly moving objects.

The first works that tackled de-fencing in a principled manner were by Liu et al. [32, 19]. They formulate the problem as the segmentation of a repeating foreground pattern that exhibits approximate translational symmetry, followed by inpainting to recover the occluded image regions. The approach of [32] is mainly limited by the fact that it uses a single image as input. Because of the opaque nature of the fence obstruction, the occluded parts of the scene must be hallucinated by the inpainting algorithm; [19] partially addresses this by using a photo taken from a different view, to reduce the number of pixels that must be hallucinated.

De-fencing can also be viewed as a special case of the more general problem of layer separation [21, 1, 5, 14], which models an image as a composition of individual lay-

ers, e.g., a foreground layer containing the obstruction and a background layer containing the scene of interest. Xue et al. [31] formulate generic obstruction removal as a layer separation problem, driven by motion parallax. Although their solution is generic and works well, it involves multiple time-consuming, hand-tuned optimization steps, and the use of hand-crafted motion and image priors. SOLD [14] is a deep learning re-incarnation of [31] that achieved stateof-the-art results on obstruction removal, and can also be adapted to remove fences from a multi-frame input burst. Unfortunately, SOLD depends on computationally expensive networks for flow computation and frame reconstruction, making it impractical for use in low-powered devices. It also often requires an online optimization step that takes \sim 3 minutes to produce acceptable results on real bursts, and even without this input-specific finetuning, it cannot be run in real time. Finally, since the background frames are the output of a reconstruction CNN module, they occasionally contain inconsistencies or artifacts.

Flow-guided completion methods [7, 29, 29, 13, 35] reduce such artifacts by computing inpainted flow maps between pairs of obstructed frames and using them to explicitly transfer pixel values to a reference frame from its temporal neighbors. Inpainting flows in the occluded area is easier than directly inpainting pixel values, so a less powerful network can be used, and the results tend to look more plausible because the pixel values are taken from real frames instead of being hallucinated by a generative network. These works do not make any assumptions regarding the shape and type of the occlusion, other the fact that it is completely opaque. However, the mask marking the occluded area is considered to be known, which is an unrealistic requirement for our purposes.

Here our goal is to develop a de-fencing algorithm that prioritizes efficiency and practicality. We develop a framework that enjoys the realism and modularity of flow-based video completion approaches, while being significantly simpler to train and deploy. Instead of videos, the input to our algorithm is shorter bursts of K = 5 frames, which are a very common photo modality in modern smartphones. The type of occlusion is known (fences), but we do not make any assumptions regarding its spatial extent or location; instead, we train a class-specific segmentation model to automatically detect fences in images. Computing flow maps of scenes occluded by fences, presents us with a new challenge, as standard optical flow networks fail under the presence of repeated patterns [9, 14]. To solve this problem, we train a segmentation-aware SPyNet [23] that can simultaneously compute and inpaint flow maps corresponding to the occluded background scene, ignoring foreground occlusions. Finally, to quantitatively evaluate the performance of our approach on real data, we collect a dataset of multiframe sequences and corresponding "pseudo-groundtruth"

for the reference frame using a alignment procedure, akin to [3]. To summarize:

- We design a CNN pipeline for multi-frame de-fencing that is simple, modular, efficient, and easy to train.
- Unlike flow-based works, which assume the occlusion is known, we estimate it automatically from the input.
- We train a segmentation-aware optical flow model that can reliably estimate flows corresponding to the background scene despite severe fence obstructions.
- Our method achieves state-of-the-art results on synthetic and real bursts, without requiring sequencespecific finetuning. As a result, it has significantly lower runtimes compared to alternatives.

2. Related Work

2.1. Image and Video De-fencing

Liu et al. [32] is perhaps the first work to formally introduce de-fencing in a computer vision context, as symmetrydriven automatic fence segmentation, followed by inpainting. [19] improves on this work by using online learning to aid lattice detection and segmentation, and by leveraging a second viewpoint to improve inpainting. Jonna et al. [11, 10] also improve fence segmentation, by complementing RGB with depth data.

[16, 34] extend de-fencing to video sequences of arbitrary number of frames. Mu et al [16] rely on motion parallax to separate foreground fence obstructions from background (although the definition of "fence" is quite loose), while Yi et al. [34] describe a bottom-up approach for video de-fencing that groups pixels in each frame using color and motion cues. Both methods rely on optimization techniques to refine their optical flow or frame inpainting results.

More recently, deep learning has been adopted for video defencing [9, 4]. Jonna et al. [9] use a pretrained classification CNN as a feature extractor and train a SVM classifier that distinguishes fence from non-fence patches. The authors reformulate an existing optical flow algorithm to make it occlusion-aware and recover the de-fenced image using FISTA optimization [2]. Du et al. [4] replace the CNN-SVM combination with a fully convolutional network (FCN) [15] and apply temporal refinement to the extracted segmentations by aggregating information from neighboring frames. Our approach shares a similar pipeline but simplifies both the segmentation extraction and occlusionaware flow computation steps, while being considerably faster, as we do not perform test-time optimization.

2.2. Layer Separation

A more generic formulation of the fence removal problem views an image as a composition of layers, each with its own alpha map (which can be semi transparent), and the goal is to separate the layers. In [5] the foregroundbackground layers are the output of two convolutional network, trained per image in an unsupervised fashion, and are recovered using a deep image prior [28]. A similar idea is used by Alayrac et al. [1] for video decomposition, but with supervised training. Other approaches for video decomposition use explicit motion information [31, 14]. Xue et al. [31] describe a computational method for decomposing a scene into an foreground obstruction layer and a background scene from multi-scale motion cues of a multi-frame sequence, in an unsupervised way. They solve an optimization problem that alternatingly finds the constituent layers and the respective motion fields that, when used to align the burst to a reference frame, can reconstruct the original frames with low error. A modern reincarnation of this approach is proposed by Liu et al. with SOLD [14]. SOLD follows a similar multi-scale approach, using a convolutional framework, both for layer reconstruction and motion estimation – the latter wth a pre-trained PWC-Net [25].

2.3. Flow-based Video Completion

Video completion is a related problem to multi-frame fence removal, the main difference being that the segmentation mask is assumed to be provided and the emphasis is typically on longer frame sequences. Our work is motivated by Xu et al. [29], who proposed the idea of first tackling the easier problem of flow inpainting, and then using the completed flows to propagate color values to a reference frame from its temporal neighbors. Since it not guaranteed that all occluded pixels are visible in some frame, a separate image inpainting step must be used to fill any remaining holes. Gao et al. [6] improve this approach by synthesizing sharp flow edges along the object boundaries and using nonlocal temporal neighborhoods for propagating pixels across frames. These works involve a series of individual, separately trained processing stages, some of which are handcrafted, inefficient, and can potentially compromise performance of subsequent stages. Li et al. [13] address this issue by proposing an end-to-end framework for flow-guided video completion. Their approach was developed concurrently to our own and shares some of its simplicity and efficiency advantages. However, their framework still operates under the assumption that the occlusion mask is provided.

3. Method

We begin with an overview of the problem we are solving, and establish the notation that we use throughout the paper. The input to our algorithm is a burst of K RGB frames, $\{I_i\}$, composed from an unknown background scene of interest $\{B_i\}$ and an unknown opaque foreground occlusion in the form of a fence $\{F_i\}$. Specifically,

$$I_i = S_i \cdot F_i + (1 - S_i) \cdot B_i, \tag{1}$$

where $S_i \in [0, 1]$ is a soft fence occlusion mask. Our goal is to train a model that removes the fence obstruction from $\{I_i\}$ and recovers a single keyframe background image B_k , where k is the keyframe index.

Instead of outputting the unobstructed frame directly, we break down the problem into the steps illustrated in Figure 2. We start by training a network that is applied individually on each frame in $\{I_i\}$, and outputs fence segmentation predictions $\{S_i\}$. The role of these segmentations is two-fold: i) they mark the occluded area that needs to be recovered; ii) they are used to condition a segmentationaware network that computes optical flows corresponding to the background scene only, directly from the occluded input $\{I_i\}$. With this network we extract flows $\{f_{kj}\}$ between the keyframe I_k and each other frame I_j in the sequence, and align the burst. Finally, we employ learned flow-guided image inpainting, to recover the parts of the keyframe that are occluded by the fence, yielding the final output \tilde{B}_k . In the following subsections we explain in detail each step.

3.1. Single-frame Fence Segmentation

Our fence segmentation model takes as input a single RGB frame, possibly containing a fence, and outputs a soft fence segmentation mask. Although this sounds like a relatively simple task, there are, in fact, multiple challenges. First, datasets of large size and with high quality annotations for fence segmentation are surprisingly scarce. The most appropriate for this task is probably the De-fencing dataset [4]. Fences in this dataset do not exhibit significant variance in terms of appearance, scale, or structure, so we rely on substantial data augmentation to train a network that is robust to different types of fences and environments.

More specifically, we apply different degrees of downscaling to the original image and its associated annotation, to effectively create fences at different scales (varying fence width/distance from the camera). To augment the limited scene variety in the De-fencing dataset, we also mask out fences, using the groundtruth segmentations, and overlay them on images from the DAVIS dataset [22]. Finally, we apply random horizontal flipping to the fence image and take randomly crop a 320×192 window for training.

The segmentation network itself is a U-net [24] backbone, with four encoder and four decoder blocks, that is trained from scratch on our augmented fence data using a binary cross entropy loss and the ADAM optimizer [12]. To obtain segmentation scores in the [0, 1] range, we apply a sigmoid in the output logits from the last U-net layer. Table 1 lists precision-recall and f-measure scores of our method at different thresholds. Even though we do not



Figure 2: Given a burst of K frames with fence obstructions as input, we reconstruct a single keyframe, after removing the fence. Our pipeline is composed of three distinct steps: a) initially, a fence mask is estimated on each input frame individually with a U-net fence segmentation model (Sec. 3.1); b) the estimated masks are used to condition a segmentation-aware optical flow SPyNet^m, which simultaneously computes and inpaints flows corresponding only to the background scene, ignoring the repeated fence occlusion patterns (Sec. 3.2); c) finally, an image inpainting module takes the estimated masks and flows, aligns the frames with respect to a selected keyframe and fills in the missing pixel values (Sec. 3.3).

Method	Precision	Recall	F-measure
Du et al. [4]	0.910	0.959	0.934
U-net (thresh=0.05)	0.908	0.958	0.931
U-net (thresh=0.1)	0.934	0.942	0.937
U-net (thresh=0.3)	0.969	0.899	0.932

Table 1: Segmentation results on the De-fencing test set [4].



Figure 3: Standard optical flow networks fail under repeated occlusion patterns. Our occlusion-aware SPyNet^m can reliably estimate the optical flow *of the background scene*, ignoring the foreground occlusion. For training, "ground-truth" flows are computed using a vanilla SPyNet on the original background frames.

use temporal information from multiple frames like [4], we achieve comparable performance.

3.2. Segmentation-aware Optical Flow Estimation

Optical flow computation is an integral step in many obstruction removal and video completion pipelines. The challenge is in how to align the background regions without being distracted by foreground occlusions. SOLD [14] uses a pretrained PWC-Net [25] to compute flows between all frame pairs in the burst, and uses frames warped to the keyframe to prime background reconstruction. Note that, in the case of fence removal, flows are only computed for the background layers, after removing the obstruction. The reason, according to the authors, is that the flow estimation network cannot handle the repetitive structures, and often predicts noisy results, which renders the alignment step unreliable. This is further aggravated by the fact that the weights of PWC-Net are frozen, so it cannot adapt to deal with potential errors in background layer reconstructions from the first levels in the coarse-to-fine SOLD architecture; this can consequently yield inaccurate flow estimates in subsequent levels, compounding errors. Finally, PWC-Net relies on cost volume computation, whose runtime does not scale favorably with input size, at least when using a publicly available implementation [17]. On the other hand, [29, 6] compute flow maps between obstructed pairs of frames using FlowNet [8]. One key difference in this scenario is that the obstruction does not follow a repetitive structure pattern, but is typically a large, compact area. This causes the flow maps to contain holes which are inpainted in a separate step.

In our work we drastically simplify flow estimation for obstructed scenes by utilizing the fence segmentation network described in Section 3.1. First, we replace PWC-Net with the faster, more lightweight, SPyNet [23] architecture¹. Second, we modify its first convolution layer to input both the fence segmentation masks, S_i , S_j , along with their corresponding input frames I_i , I_j . Our modified SPY^m architecture then estimates the mask-conditional flow map

$$f_{ij}^{m} = \mathbf{SPY}^{m}([I_i; S_i], [I_j; S_j]),$$
(2)

[·] denoting concatenation along the channel dimension. We use the original pretrained weights to initialize SPY^m, except for the modified part of the input layer, which we initialize randomly. During training, f_{ij}^m are computed between synthetically generated frames of fence images overlaid on clean background frames. Consequently, we use flow maps computed with the vanilla SPyNet on the *clean*

¹We use "SPY" in equations, for short.



Figure 4: **Frame inpainting module.** We use the predicted fence segmentations to mask out (\odot) the occluded areas in the input frames. We then use optical flow to warp (*W*) the masked frames and respective masks. Finally, the flows, the validity maps (see text), and the aligned frames and masks are concatenated ([·]) and passed as features to a CNN predicting the keyframe residual in the occluded regions.

background frames B_i , B_j as pseudo ground truth targets. We use an L_1 loss to finetune SPY^m:

$$\mathcal{L}_f = \frac{1}{2N} \sum_{\mathbf{x}} |\mathbf{SPY}(B_i, B_j)|_{\mathbf{x}} - f_{ij|\mathbf{x}}^m|, \qquad (3)$$

where N is the number of image pixels, x denotes the location at which we evaluate, and we average over 2N to account for the u, v flow channels.

Conditioning SPyNet on segmentation predictions allows us to denote parts of the scene corresponding to obstructions and ignore them while computing *background* flows, solving a fundamental problem faced by SOLD. This idea was previously explored in [4] but it involved a costly optimization process. Our approach is simple but robust to the presence of significant fence obstructions. Segmentation-aware flow estimation can be useful in a variety of practical settings where one wants to ignore parts of the scene as distractions or sources of noise. Figure 3 demonstrates the effectiveness of our approach by comparing outputs of the vanilla SPyNet and our segmentationaware SPyNet^m, on the same obstructed scene.

3.3. Flow-guided Multi-frame Fence Removal

The final component in our fence removal pipeline is a frame inpainting module, depicted in Figure 4. The frame inpainting module takes as input the sequence of obstructed frames $\{I_i\}$, forward flow maps f_{ki}^m , computed between the keyframe I_k and each other frame in the burst using the mask-conditional SPyNet^m (Section 3.2), and the fence segmentation masks $\{S_i\}$ computed using our single-frame segmentation model (Section 3.1). We first use $\{S_i\}$ to mask out the areas that correspond to fences in each frame, obtaining masked frames $I_i^m = I_i \odot S_i$. Then flows f_{ki}^m are used to warp all frames and their respective segmentations

with respect to the reference frame, giving rise to aligned masked frames $\tilde{I}_i^m = W(I_i^m, f_{ki}^m)$ and aligned fence masks $\tilde{S}_i = W(S_i, f_{ki}^m)$. We also compute binary masks $\{V_i\}$ that mark valid warped regions, and are "on" for all pixels that fall inside the image grid after warping.

 $f_{in} = [\{\bar{I}_i^m\}; \{S_i\}; \{V_i\}]$ is passed as input to a Residual Dense Network (RDN) [36] that is responsible for filling in the missing areas in the keyframe. We also add a skip connection between the masked keyframe I_k^m and the output of the RDN, so the latter only has to learn to fill in the missing areas instead of reconstructing the entire image. The inpainting module is trained in a supervised fashion using an L_1 loss and the clean background as the ground truth:

$$\mathcal{L}_{in} = \frac{1}{N} \sum_{\mathbf{x}} |B_{k|\mathbf{x}} - (I_k^m + \text{RDN}(f_{in}))|_{\mathbf{x}}|.$$
(4)

3.4. Implementation details

We implement our pipeline in Python 3 and Py-Torch [20]. For U-net, PWC-Net, SPyNet, and RDN, we use their publicly available third-party implementations [27, 18, 17, 33]. To facilitate our experiments, we have also re-implemented SOLD in PyTorch (SOLD_{pt}), following closely the original Tensorflow implementation [14]; we plan to make our re-implementation publicly available to allow for broader use by the community and replication of results. Unless otherwise mentioned, we train all our models for 1000 epochs, using a starting learning rate $l_r = 10^{-4}$, a weight decay rate $w_r = 4 \cdot 10^{-5}$, and the ADAM optimizer [12] with parameters $\alpha = 10^{-4}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$. All three models (fence segmentation, occlusion-aware flow estimation, frame inpainting) are trained independently.

4. Data for Training and Evaluation

We use two types of data in our experiments. The first type is synthetic multi-frame sequences, generated similarly to previous works [4, 14]. These are used predominantly for training and validation experiments, but a held-out test set is also used for evaluation. The second type is real bursts with fence obstructions, which include *uncontrolled* sequences, for which no ground truth clean frame is available, and *controlled* sequences, which come with a clean background scene (without the fence) as ground truth.

Synthetic bursts are generated by overlaying obstruction (foreground) layers on a clean scene (background). We source background scenes (which are also used as ground truth during training and evaluation) from Vimeo-90k [30], which consists of videos depicting every day activities in realistic settings, often including people and other objects. We specifically use the original test split of the dataset²,

²http://data.csail.mit.edu/tofu/testset/

which contains sequences of seven (7) frames. Training and validation splits are generated on the fly, but for our evaluation experiments we use a fixed test set of 100 bursts.

The foreground fence obstructions are sourced from the De-fencing dataset [4], which contains 545 training and 100 test images with fences, along with corresponding binary masks as ground truth for the fence segmentation. The sequences in this dataset have been collected in various outdoor conditions and have a variable frame count per scene. Since we have the ground truth fence masks, we can use them to mask out the fence from any given frame and overlay it on a clean background from Vimeo. To obtain a fence image burst of size K, we mask out the fence from a single frame and apply K random perspective distortions to it, to simulate the changes caused by slightly different viewpoints and motion. To increase the variability of fences and background scenes, we apply various forms of data augmentations before fusing them into a single frame; these are listed in detail in the supplemental material.

Real bursts. Because we want to develop a practical algorithm for fence removal, good performance under realistic motion, lighting, and obstruction patterns is of paramount importance. In previous works, performance on real sequences is –for the most part– evaluated qualitatively, since obtaining the ground truth background is far from trivial. Liu et al. [14] include only two sequences with fence-like obstructions, collected in a controlled environment, which is too small a dataset for a proper quantiative evaluation.

In this paper we construct a wider set of controlled sequences, specifically for quantitative evaluation. Rather than collecting toy scenes as in Liu et al. [14], we capture real world hand-held sequences with a fence and a corresponding background ground truth image without a fence. As we cannot physically remove a fence, we instead bring our camera to the fence and center it in one of the fence cells such that only the background is visible. To maintain a similar level of brightness and sharpness of the background in the input and ground truth images, we fix the exposure and focus of the camera on the background during capture. Due to camera motion and possible changes in illumination, the input keyframe and its respective ground truth may be misaligned or have color discrepancies. We align crops of the scene using standard feature-based RANSAC fitting of homographies, similar to [3] and correct color discrepancies using color histogram matching. We then filter out any misaligned crops using SSIM, PSNR, and human visual check, avoiding mostly homogeneous regions, to promote diversity in our dataset. Our final real burst dataset consists of $185\ 320 \times 192$ input bursts and corresponding ground truth keyframes. More details on dataset generation and image sampling are provided in supplemental material.

We compare our method and other works on synthetic and real bursts. For quantitative evaluations we use the test set of our synthetically generated fence-obstructed sequences, and our real bursts described in Section 4. For all baselines, we use the officially released model weights, with the exception of our SOLD reimplementation. We also purposedly omit the sequence-specific online optimization step of SOLD in our comparisons. Although online optimization improves performance, its runtime is quite slow (~ 3 minutes per burst), pushing it outside the scope of our work, which is centered around efficiency and practicality. For qualitative evaluations and visual comparison, we use real sequences from previous work and the data we collected.

5.1. Baselines

Single-frame baseline. We pass either ground truth or our (thresholded) U-net fence mask predictions as input to LaMa [26], a state-of-the-art CNN-based inpainting method, to create a single-frame de-fencing baseline. LaMa takes as input a (possibly obstructed) image and a binary mask and inpaints the area marked by the mask.

SOLD [14] primarily targets reflection removal, but it can be adapted to deal with opaque obstructions such as fences or raindrops on glass. We evaluate both the original Tensor-flow model (SOLD $_{tf}$) and our PyTorch reimplementation (SOLD $_{pt}$), with the latter trained on our synthetic data.

Flow-guided video completion operates in a setting that is different than ours in a few ways. First, the mask denoting the occluded area is known, and its shape is either rectangular or in the shape of an object in the video. Second, the number of frames in a typical input video sequence is $K \gg 5$. Lastly, the output is the entire inpainted video. Nevertheless, we can apply these methods for de-fencing in a relatively straightforward fashion by passing the fence segmentation as the occlusion mask, treating the burst as a (short) video sequence, and keeping the inpainted result for the reference frame only. In our experiments we compare against two recent flow-guided approaches, FGVC [6] and E2FGVI [13], using their publicly provided code.

5.2. Fence Removal on Synthetic and Real Data

Quantitative comparisons on synthetic and real bursts are shown in Table 2. We report performance in terms of the commonly used SSIM, PSNR, and LPIPS metrics. For LPIPS we use a VGG-16 backbone as the feature extractor. PSNR and SSIM can be computed as an aggregation of pixel-wise scores, so we use the fence masks (ground truth in the case of synthetic data, thresholded and binarized Unet predictions in the case of real data³) to dissect perfor-

^{5.} Experiments

vimeo_test_clean.zip

³To get better pseudo ground truth fence masks, we run U-net at multiple scales and compute the pixel-wise maximum across scales.



Figure 5: Qualitative de-fencing results on real sequences. We highlight areas of interest in red (shown as zoomed insets) and report PSNR inside the fence mask for all methods. Last example is from Xue et al. [31] and has no ground truth.

Synthetic data (Methods in blue use GT fence masks)						Real bursts (Methods in blue use <i>pseudo</i> -GT fence masks)								
Method	SSIM ↑		PSNR (dB) \uparrow		LPIPS \downarrow	SSIM ↑		PSNR (dB) \uparrow			LPIPS \downarrow			
Wiethou	in	out	total	in	out	total	(VGG)	in	out	total	in	out	total	(VGG)
$SOLD_{tf}[14]$.783	.970	.941	23.36	37.82	30.34	.111	.728	.911	.885	23.50	30.27	27.71	.132
SOLD _{pt} [14]	.893	.993	.977	28.14	45.70	35.82	.040	.813	.916	.902	26.41	30.90	29.71	.094
LaMa [26]	.788	.995	.964	24.97	51.38	33.10	.039	.480	.902	.845	19.95	29.96	26.25	.133
LaMa [26]*	.655	.955	.910	20.96	31.74	27.38	.089	.477	.867	.816	20.85	28.02	25.98	.132
FGVC [6]	.846	.943	.928	25.56	33.85	30.36	.068	.848	.910	.901	27.56	30.48	29.73	.095
FGVC [6]*	.784	.896	.879	22.73	27.80	26.05	.113	.856	.907	.900	27.49	29.90	29.36	.090
E2FGVI [13]	.918	.997	.985	30.79	55.87	38.89	.030	.571	.902	.856	19.69	29.88	25.87	.167
E2FGVI [13]*	.890	.984	.969	29.34	38.64	35.25	.044	.709	.901	.875	25.58	30.25	29.14	.117
Ours	.954	.999	.992	33.76	56.55	41.78	.015	.869	.917	.909	28.60	31.14	30.46	.080
Ours-fence _{gt}	.957	.999	.993	34.33	58.58	42.42	.012	.872	.918	.910	28.77	31.15	30.53	.078

Table 2: Results on synthetic test data (left) and our collected real bursts (right). Results in rows denoted with "*" are computed after thresholding the fence segmentations at t = 0.1, and dilating the binary mask 4 times with a 3×3 square.

mance in three different regions: a) inside the mask (in); b) outside the mask (out); and c) in the entire image (total). Performance outside the mask is high for all methods, since this part of the image is not occluded. Performance *inside* the mask is the most important criterion, since it quantifies the quality of reconstruction only in the occluded area. We outperform all single- and multi-frame baselines, based on all metrics, with FGVC [6] performing second best at the cost of much higher runtime (see Section 5.3). We would like to draw the reader's attention to the results of LaMa [26], in particular. LaMa is a state-of-the-art inpainting method, yet it achieves surprisingly low PSNR-in and SSIM-in scores. The reason becomes clear if one looks at Figure 5 (e.g., antennae structure in the first example): even

though LaMa produces perfectly plausible results under the occluded area, these are often very different than the actual background scene. These results tellingly demonstrate the advantage of using multiple frames for de-fencing. Figure 5 also illustrates that alternatives like SOLD and FGVC can yield blurry or completely scrambled results, likely due to issues with frame alignment. For more qualitative results see our supplemental material.

5.3. Runtime Analysis

Table 3 compares the total runtime of our pipeline with other approaches. In our case, timing includes all necessary processing steps: fence segmentation, optical flow computation, alignment, and frame inpainting. LaMa, FGVC, and

Method	LaMa	$SOLD_{pt}$	FGVC	E2FGVI	Ours
Runtime (s)	0.2	0.8	0.7	0.16	0.14 (0.08)

Table 3: Runtime comparison for a 5-frame burst (LaMa processes a single frame). Times for other methods *do not* include segmentation (fence segmentations are part of the SOLD output). We provide the runtime of our method without the segmentation step in parentheses, for comparison.

E2FGVI times do not include the time spent on fence segmentation, since these methods assume that the occlusion masks are precomputed. Our method is clocked at $\sim 7 fps$, for 320×192 inputs, which is $\sim 5 \times$ faster than the next best performing methods, SOLD and FGVC, and comparable to E2FGVI and LaMa. The runtime difference with the last two may not seem large in absolute terms, but it still amounts to a 12.5% and 30% lower runtime respectively, while our method *significantly* outperforms them in terms of reconstruction quality on real data. If we exclude the time spent on segmentation from our pipeline, the speedup becomes even more noticeable (50% and 60% respectively). All timings are performed on a workstation equipped with a Nvidia GTX 1080 Ti, with 12GB of GPU RAM. The detailed breakdown of timings for our pipeline is: i) segmentation (for a 5-frame burst): 0.06s; ii) flow estimation and alignment: 0.04s; iii) frame inpainting (RDN): 0.04s.

5.4. Ablation

Changing the frame inpainting module architecture allows us to trade-off reconstruction performance for efficiency. Replacing the RDN with a simple CNN consisting of 8 convolution + LeakyReLU layers decreases performance by 3.5 dB on synthetic test data but also reduces runtime by $\sim 30\%$, from 0.14s to 0.1s.

Does frame alignment matter? Inaccurate computation of the motion corresponding to the (occluded) background can lead to errors in frame alignment, impacting the quality of frame reconstruction. We experiment with the following two alternatives for computing flows: i) original SPyNet on obstructed frames; ii) SPyNet followed by masking the occluded areas and using Laplacian inpainting [29] to complete the missing flows (SPyNet^{inp}). We also consider the option of not aligning frames at all, and leaving our RDN inpainting network to learn how to complete the occluded areas in the reconstructed keyframe.

The importance of alignment and the quality of flows used become becomes clear when looking at the results in Table 4. Not aligning the input frames at all leads to a noticeable drop in keyframe reconstruction performance. Standard flow networks cannot handle the repeated fence obstruction patterns, and explicitly inpainting the flows under the occluded area does not help either since flow artifacts extend to non-fence areas as well, as shown in Fig-

Alignment		SSIM ↑		PS	LPIPS \downarrow		
Anginnent	in	out	total	in	out	total	(VGG)
None	.792	.996	.965	24.82	53.64	32.84	.028
SPyNet	.841	.997	.973	26.84	54.32	34.93	.047
SPyNet ^{inp}	.841	.997	.973	26.81	54.13	34.89	.048
SPyNet ^m	.954	.999	.992	33.76	56.55	41.78	.015

Table 4: Effect of frame alignment on keyframe reconstruction quality (results on our synthetic test data).



Figure 6: **Failure example.** Imperfect fence predictions (middle) compromise the quality of the inpainting (right).

ure 3. Our occlusion-aware SPyNet^m, on the other hand, can accurately estimate background flows, resulting in superior frame alignment and reconstruction quality.

5.5. Limitations and Failure Cases.

The main limitation of our approach is that the quality of the final reconstruction depends on the outputs of U-net and SPyNet^m in the two previous stages. Errors in fence segmentation affect segmentation-aware optical flow computation, potentially compromising frame alignment, which is crucial for good inpainting (see ablation in Section 5.4). In addition, fence segmentations are also used in masking out the occluded areas in $\{I_i\}$, The most deleterious mistakes occur when the fence occlusion is out of our training distribution, e.g., when the fence has an unusual shape/pattern, or when the contrast with the background is low. This is an issue that can be handled to a certain extent through better data augmentation during training or by having access to richer datasets with varied types of fences, as we show in the supplemental material. A failure example and its effect on frame inpainting is shown in Figure 6.

6. Conclusions

We have developed a simple, modular, and efficient pipeline for removing fence obstructions from a singe frame in a photo burst. Our algorithm enjoys the realism of flow-guided video completion methods, while addressing some of their practical limitations, such as complicated training and long runtimes. Our method runs at 7 fps for 5-frame 320×192 bursts, on a Nvidia GTX 1080 Ti, and is particularly effective on *real* data, outperforming other single- and multi-frame de-fencing baselines on a dataset of obstructed bursts we collected specifically for this problem.

References

- Jean-Baptiste Alayrac, Joao Carreira, and Andrew Zisserman. The visual centrifuge: Model-free layered video representations. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 2457– 2466, 2019.
- [2] Amir Beck and Marc Teboulle. A fast iterative shrinkagethresholding algorithm for linear inverse problems. SIAM journal on imaging sciences, 2(1):183–202, 2009.
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Deep burst super-resolution. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9209–9218, 2021.
- [4] Chen Du, Byeongkeun Kang, Zheng Xu, Ji Dai, and Truong Nguyen. Accurate and efficient video de-fencing using convolutional neural networks and temporal information. In 2018 IEEE International Conference on Multimedia and Expo (ICME), pages 1–6. IEEE, 2018.
- [5] Yosef Gandelsman, Assaf Shocher, and Michal Irani. " double-dip": Unsupervised image decomposition via coupled deep-image-priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11026–11035, 2019.
- [6] Chen Gao, Ayush Saraf, Jia-Bin Huang, and Johannes Kopf. Flow-edge guided video completion. In *European Confer*ence on Computer Vision, pages 713–729. Springer, 2020.
- [7] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. Temporally coherent completion of dynamic video. ACM Transactions on Graphics (TOG), 35(6):1–11, 2016.
- [8] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2462–2470, 2017.
- [9] Sankaraganesh Jonna, Krishna K Nakka, and Rajiv R Sahay. Deep learning based fence segmentation and removal from an image using a video sequence. In *European Conference* on Computer Vision, pages 836–851. Springer, 2016.
- [10] Sankaraganesh Jonna, Sukla Satapathy, and Rajiv R Sahay. Stereo image de-fencing using smartphones. In 2017 IEEE international conference on acoustics, speech and signal processing (ICASSP), pages 1792–1796. IEEE, 2017.
- [11] Sankaraganesh Jonna, Vikram S Voleti, Rajiv R Sahay, and Mohan S Kankanhalli. A multimodal approach for image defencing and depth inpainting. In 2015 Eighth International Conference on Advances in Pattern Recognition (ICAPR), pages 1–6. IEEE, 2015.
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [13] Zhen Li, Cheng-Ze Lu, Jianhua Qin, Chun-Le Guo, and Ming-Ming Cheng. Towards an end-to-end framework for flow-guided video inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17562–17571, 2022.

- [14] Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Learning to see through obstructions with layered decomposition. *IEEE Transactions* on Pattern Analysis and Machine Intelligence, 2021.
- [15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [16] Yadong Mu, Wei Liu, and Shuicheng Yan. Video de-fencing. IEEE Transactions on Circuits and Systems for Video Technology, 24(7):1111–1121, 2013.
- [17] Simon Niklaus. A reimplementation of PWC-Net using PyTorch. https://github.com/sniklaus/ pytorch-pwc, 2018.
- [18] Simon Niklaus. A reimplementation of SPyNet using PyTorch. https://github.com/sniklaus/ pytorch-spynet, 2018.
- [19] Minwoo Park, Kyle Brocklehurst, Robert T Collins, and Yanxi Liu. Image de-fencing revisited. In Asian Conference on Computer Vision, pages 422–434. Springer, 2010.
- [20] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [21] M Pawan Kumar, Philip HS Torr, and Andrew Zisserman. Learning layered motion segmentations of video. *International Journal of Computer Vision*, 76(3):301–319, 2008.
- [22] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Computer Vision and Pattern Recognition*, 2016.
- [23] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4161–4170, 2017.
- [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [25] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018.
- [26] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF Win*-

ter Conference on Applications of Computer Vision, pages 2149–2159, 2022.

- [27] Nikhil Tomar. Semantic-segmentation-architecture. https://github.com/nikhilroxtomar/ Semantic-Segmentation-Architecture, 2020.
- [28] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.
- [29] Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. Deep flow-guided video inpainting. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3723–3732, 2019.
- [30] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 127(8):1106–1125, 2019.
- [31] Tianfan Xue, Michael Rubinstein, Ce Liu, and William T Freeman. A computational approach for obstruction-free photography. ACM Transactions on Graphics (TOG), 34(4):1–11, 2015.
- [32] Liu Yanxi, Belkina Tamara, H Hays James, and Roberto Lublinerman. Image de-fencing. In *Computer Vision and Pattern Recognition*, 2008. CVPR 2008. IEEE Conference on, pages 1–8, 2008.
- [33] Jeffrey Yeo. Rdn. https://github.com/yjn870/ RDN-pytorch, 2019.
- [34] Renjiao Yi, Jue Wang, and Ping Tan. Automatic fence segmentation in videos of dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 705–713, 2016.
- [35] Kaidong Zhang, Jingjing Fu, and Dong Liu. Inertia-guided flow completion and style fusion for video inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5982–5991, 2022.
- [36] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018.