Seeing Behind Dynamic Occlusions with Event Cameras

Rong Zou   Manasi Muglikar   Nico Messikommer   Davide Scaramuzza
Robotics and Perception Group, University of Zurich, Switzerland

Abstract

Unwanted camera occlusions, such as debris, dust, rain-drops, and snow, can severely degrade the performance of computer-vision systems. Dynamic occlusions are particularly challenging because of the continuously changing pattern. Existing occlusion-removal methods currently use synthetic aperture imaging or image inpainting. However, they face issues with dynamic occlusions as these require multiple viewpoints or user-generated masks to hallucinate the background intensity. We propose a novel approach to reconstruct the background from a single viewpoint in the presence of dynamic occlusions. Our solution relies for the first time on the combination of a traditional camera with an event camera. When an occlusion moves across a background image, it causes intensity changes that trigger events. These events provide additional information on the relative intensity changes between foreground and background at a high temporal resolution, enabling a truer reconstruction of the background content. We present the first large-scale dataset consisting of synchronized images and event sequences to evaluate our approach. We show that our method outperforms image inpainting methods by 3dB in terms of PSNR on our dataset.

1. Introduction

The majority of computer vision algorithms operate under the assumption that the scene being analyzed is visible in its entirety and without any obstructions. However, occlusions resulting from non-ideal environmental factors, such as dirt from road debris, insects, raindrops, snow, and others, have the potential to significantly reduce the quality of captured images. The dynamic nature of these occlusions, affects the downstream applications such as autonomous navigation, object detection, tracking, 3D reconstruction, etc. Dynamic occlusions pose an interesting challenge for scene understanding, especially for robotics and autonomous driving, where it is crucial to know the occluded scene for safe navigation in unknown environments.

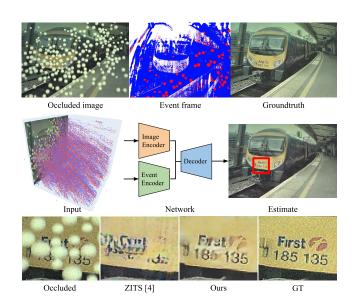The literature, however, has focused more on static oc-

Figure 1: (Top) The background image (right) occluded by moving particles observed with a traditional camera (left) and event camera (center). (Middle) The proposed method for background reconstruction in the presence of moving objects uses the continuous event stream and the occluded image to learn to reconstruct the background image. (Bottom) Our approach outperforms the image inpainting baseline and accurately reconstructs the background intensity.

clusions and many different approaches have been proposed that use multiple viewpoints (synthetic aperture imaging) [21, 44], optical diffraction cloaking [35], or image inpainting [2, 9, 18, 19, 38, 43, 20, 4, 12]. The assumption of static occlusion often renders the above approaches infeasible in the presence of dynamic obstructions [49, 39]. Image inpainting approaches are often trained on large-scale datasets to predict intensity in areas specified by a pre-defined mask. However, these methods hallucinate the scene since they are trained to favour more aesthetically pleasing images over the true scene content as seen in Fig. 1. While there have been several advancements in these approaches with the advent of deep learning, these methods fail to reconstruct true scene content in the presence of dynamic occlusions. Synthetic Aperture Imaging (SAI) captures the target scene from multiple viewpoints to create a virtual camera with a

large-aperture lens [47, 49]. This makes it easier to blur the foreground occlusion and refocus the image on the background scene. These methods assume the occlusions are at the same depth and stationary throughout the capture process, which can be up to a few milliseconds, rendering this method infeasible for fast-moving occlusions.

In this paper, we propose a novel approach to estimate the background image in the presence of dynamic occlusions using a single viewpoint. Our solution complements a traditional camera with an event camera. Event cameras [6] are bio-inspired sensors that asynchronously measure changes in brightness with microsecond resolution. When an occlusion crosses over a background image, it causes changes in intensity, thus triggering events as seen in Fig. 1. These events give additional information on the relative intensity changes between the foreground and background at a high temporal resolution, which leads to a more precise reconstruction of the background, compared to image inpainting methods which only use unoccluded image statistics. Since this is the first time an event camera is used to tackle dynamic occlusion removal, we collect a large-scale dataset containing 233 challenging scenes with synchronized events, occluded images, and ground truth unoccluded images.

The classical event generation model [16] can be used to reconstruct the background image intensity, given contrast threshold and occlusion intensity. However, due to inherent sensor noise and instability of the contrast threshold, this approach results in poor performance at higher occlusion densities. We, therefore, propose to use a supervised data-driven approach that learns to reconstruct the unoccluded image, implicitly learning the sensor parameters and noise characteristics. The network predicts the background intensity using only a single occluded image and events. We show that our method relying on a single frame and events leads to a performance improvement of $3.3\,\mathrm{dB}$ on our synthetic dataset and $2.8\,\mathrm{dB}$ on our real dataset over the image-inpainting baselines in terms of PSNR. Our results show that our method is capable of accurately reconstructing scenes in the presence of dynamic occlusions and has the potential to increase perceptual robustness in unknown environments. To summarize, our contributions are:

- A novel solution to the problem of background image reconstruction in the presence of dynamic occlusions using a single viewpoint. We are the first to tackle this problem using an event camera in addition to a standard camera.

- A large-scale dataset recorded in the real world containing 233 challenging scenes for background reconstruction with synchronized events, occluded frames, and groundtruth unoccluded images.

2. Related works

2.1. Fame-based methods

Image inpainting is the task closest related to our tackled problem of removing dynamic occlusions in a single camera view. In image inpainting, the goal is to reconstruct missing pixel intensities in areas specified by a pre-defined mask. Traditional image inpainting methods rely on the statistics of the unmasked image areas to directly infer the missing regions [2, 1] or copy textures from database images [3, 9]. Because of the ability to account for the context in masked images, current research on image inpainting mostly focuses on deep generative approaches [10, 18, 19, 25, 28, 50]. To provide global contextual information early on in the network, LaMa[38] proposes to perform convolutions in the Fourier domain. In [7], a two-stream network fuses texture and structure information for a more accurate image generation. To generalize to different scenes, MISF [13] proposes to apply predictive filtering on the feature level to reconstruct semantic information and on the image level to recover details. More recently, transformers were also applied to the task of image inpainting [43, 20, 4, 12] and achieve state-of-the-art performance. In PUT [20], an encoder-decoder is applied to reduce the information loss suffered by tokenizing and dekonizing the tokens for the transformer backbone. Instead of directly predicting the pixel intensities, ZITS [4] uses a transformer model to restore the structures with low resolution in form of edges, which are then further processed by a CNN. Generally, since the only information for reconstructing image patches is given by the structure and context of the unmasked regions, image inpainting methods aim to reconstruct plausible and aesthetically pleasing images without the focus on reconstructing the real scene content. Additionally, the application of image inpainting methods for occlusion removal requires the estimation of the occlusion mask. Therefore, image inpainting methods tackle a different problem setting and are not designed for the task of removing dynamic occlusions in a single-camera view.

Instead of considering only one image, a sequence of images taken at different positions can be leveraged to remove reflections and static occlusions [21] or contaminant artifacts on lenses [14]. In the field of synthetic aperture imaging (SAI), different camera views are used to estimate the incoming lights from the background onto a synthetic camera view. By considering a large number of images captured at different positions, SAI methods can focus on different planes in the scene, effectively removing the static occlusions by blurring them out [47, 42, 41, 22, 49, 29]. Furthermore, various imaging setups with hardware modifications are developed. In DeOccNet [44], multiple views of a light field camera are used in an encoder-decoder fashion to remove the static occlusion. Capturing the same scene from

multiple viewpoints constrains the occlusion to be stationary during the capture process, which can range from a few milliseconds to seconds depending on the camera speed and exposure time. Our problem, therefore, becomes more challenging to solve with SAI, as these methods are designed to remove static occlusions by considering multiple camera views. Another interesting work [35] proposes to learn a diffractive optical element, which can be inserted in front of the lens to focus better on the objects further away from the camera, which effectively removes thin occluders close to the camera.

2.2. Event-based methods

It was shown that event cameras are capable of generating image sequences in high-speed and high-dynamic range scenes in which standard cameras fail [32, 33, 24, 52]. Moreover, event cameras are also successfully applied in combination with frame cameras for different imaging tasks. Specifically, the high temporal resolution of event cameras is used for filling in missing information between two frames to create high frame-rate videos [40, 46, 54] or to deblur images [36, 26], whereas the high-dynamic range property is leveraged in HDR imaging [8, 46, 23]. Recently, event cameras achieved impressive results for removing static occlusions using a moving camera, which is termed event-based synthetic aperture [53, 51]. In [15], frames are additionally considered to reconstruct the background information behind a static, dense occlusion. In contrast to our proposed method for dynamic occlusion removal, the stated event-based synthetic aperture approaches tackle a different task of removing static occlusions, which are captured at different positions using a moving camera.

3. Method

Problem formulation Let us assume an event-based dynamic occlusion setting, where we are given an occluded frame I_0 at timestep 0 and an event sequence recorded between timestep 0 and t. We aim to reconstruct the background image from event sequences by integrating information from events and the occluded image.

Basic considerations Event-cameras are novel, bio-inspired sensors that asynchronously measure *changes* (i.e., temporal contrast) in illumination at every pixel, at the time they occur [16, 37, 5, 30]. When a particle occludes a background pixel at position x_k at time t, the intensity changes from $L(\mathbf{x_k}, t_k - \Delta t_k)$ to $L(\mathbf{x_k}, t_k)$. In particular, an event camera generates an event $e_k = (\mathbf{x_k}, t_k, p_k)$ at time t_k when the difference of logarithmic brightness at the same pixel $\mathbf{x}_k = (x_k, y_k)^\top$ reaches a predefined threshold C:

$$L(\mathbf{x_k}, t_k) - L(\mathbf{x_k}, t_k - \Delta t_k) = p_k C, \qquad (1)$$

where $p_k \in \{-1, +1\}$ is the sign (or polarity) of the brightness change, and Δt_k is the time since the last event at the pixel $\mathbf{x}_k$. The result is a sparse sequence of events that are asynchronously triggered by illumination changes. To reconstruct the original intensity, we use the above event generation model as follows:

$$L(\mathbf{x_k}, t_k) = L(\mathbf{x_k}, t_k - \Delta t_k) + \sum p_k C, \qquad (2)$$

Using the above equation, the intensity of the occluded pixels can be estimated under the condition that the contrast threshold C and the position of the occlusion are known. Thus, a model-based method requires segmenting the occluded pixels and estimating the contrast threshold.

We design such a method as a baseline for comparison and refer to it as *Accumulation Method*. The designed method assumes that the occluded pixels have similar intensities and uses this heuristic to segment the occluded pixels from the background. However, this method breaks down in the presence of multiple overlapping occlusions, which leads to an occluded area in the image with different intensities. For the estimation of the unknown contrast threshold C, as is commonly done, a reasonable value is defined for the event accumulation.

3.1. Network Architecture

We design a learning-based method to tackle these deficiencies by implicitly learning the contrast threshold as well as detecting the occluded areas using the input frame and events. The proposed network takes as input a single occluded image I_0 and N event representations $E_{0\to\tau}, ..., E_{(N-1)\tau\to t}$, which are constructed in time intervals of τ based on the events recorded immediately after the image between timestep 0 and t. The proposed fully convolutional network leverages ConvLSTM layers [48] to integrate the events and follows a U-Net structure [34] to predict a residual image $\tilde{I}_0$. The residual image is added to the occluded input image to obtain the final occlusion-free intensity image $\hat{I}_0$, which should contain the true background information at the occluded areas of the input image. The overview of our network is visualized in Fig. 2 (a).

As a first step, we integrate the event information in the *Event Accumulation Module* (EAM) to retrieve accumulated intensity changes since the events capture instantaneous changes caused by the movement of the occlusions. The Event Accumulation Module should focus on the intensity changes relevant to the background information, i.e., it should select the intensity information at the timesteps at which the background behind the occlusions becomes visible. To achieve this, we integrate the event information at multiple timesteps and let the network select the intensity change corresponding to the unoccluded background. This also helps the network to ignore events triggered by overlapping occlusions. We make use of the accumulation and
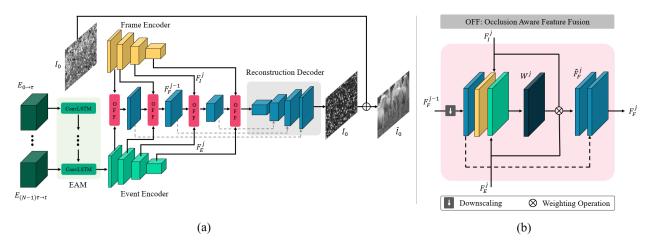
Figure 2: As shown in (a), our proposed method first integrates N event representations $E_{0 \to \tau}, ..., E_{(N-1)\tau \to t}$ using ConvL-STM layers in our *Event Accumulation Module* (EAM). The final hidden state of the ConvLSTM layers is processed with the multi-scale *Event Encoder*. Similarly, the occluded input image I_0 is processed with a *Frame Encoder*. At each scale j, the *Occlusion-aware Feature Fusion* (OFF) module (b) decides in spatial and channel dimensions which frame features F_I^j and event features F_E^j to fuse. The fused features F_F^j are reconstructed to a residual image $\tilde{I}_0$ using the *Reconstruction Decoder*. The occlusion-free image $\hat{I}_0$ is obtained by adding the residual image to the input image I_0.

gating mechanism of ConvLSTM layers to model the event integration mechanism, as well as the timestep selection. Specifically, we apply two ConvLSTM layers in sequence to recurrently process the event representation starting with an initial zero-valued hidden and cell states. The output of the Event Accumulation Module represents the last hidden state of the last ConvLSTM layer at the timestamp of the final event representation.

The integrated event features are then further encoded using a multi-scale *Event Encoder* consisting of a convolution layer and a batch normalization layer with ReLU activation per scale to generate event features. In a similar fashion, the occluded image is processed using a corresponding *Frame Encoder*. In total there are j=6 scales. Due to the differential principle of event cameras, we noticed that event features are highly beneficial for the reconstruction of textures, e.g., edges, but struggle to reconstruct uniform image areas. In contrast, standard frames contain absolute intensities and thus are better suited for filling in uniform areas since, in the simplest case of constant pixel values, a border value copying suffices. Because of this duality, we process event and frame features separately and employ a spatial as well as channel-wise weighting between them at each encoding scale j. The spatial weighting is done in the *Occlusion-aware Feature Fusion* (OFF) module, which computes gating weights W^j based on the event features F_E^j and frame features F_I^j at the considered scale j as well as the fused features from the previous scale F_F^{j-1}, which are downscaled by a strided convolution layer. Specifically, the event and frame features are fused by weighting them

with W^j and $1 - W^j$ and summing them up.

$$\hat{F}_F^j = (1 - W^j)F_I^j + W^j F_E^j, \tag{3}$$

Together with the fused feature of the previous scale, the weighted sum of the event and frame features is given as input to convolution layers outputting the fused features F_F^j. The *Occlusion-aware Feature Fusion* module is visualized in Fig. 2 (b).

Finally, the fused features at each scale are used as skip connections inside a *Reconstruction Decoder*. For upsampling, we use a bilinear interpolation followed by a convolution layer. The output of the *Reconstruction Decoder* is the residual image $\tilde{I}_0$, which is added to the input image I_0 and given to a sigmoid function resulting in the final occlusion-free image $\hat{I}_0$. For more details about the network structure, we refer to the supplementary.

4. Dataset

To the best of our knowledge, there exists no dataset aimed at tackling the problem of dynamic occlusion removal using an event camera. Therefore, we generated a new synthetic dataset containing frames and events to compare our approach against existing baselines in controlled conditions. Furthermore, to validate our approach in the real world, we record a novel dataset with real events and images. Both datasets lay the foundation for future research in dynamic occlusion removal using event cameras.
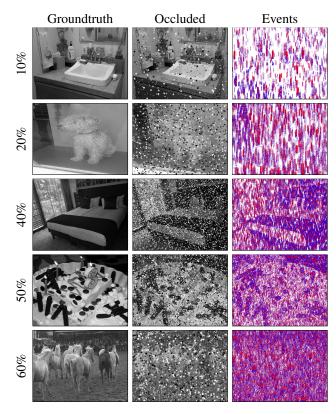
Figure 3: Sequence samples of different coverages from our synthetic dataset showing the occlusion-free ground truth frame (left), the occluded frame (middle), and the events accumulated into an event frame (right).

4.1. Synthetic Dataset

For the synthetic dataset, we simulate 320 training and 160 test sequences, each containing a sequence of occluded images, synthetic events, and an occlusion-free ground truth image. The background images are sampled from the MS-COCO dataset [17] and cropped to a resolution of $384{\times}512$. Each sequence consists of an image covered with circular-shaped particles which smoothly move across the image. To introduce more variability in the data generation process, we randomize the size, intensity, and velocity of the occlusion particles. This is done to simulate different occlusion densities from 10% to 60% occlusion ratio as seen in Fig. 3. Since the generation of events requires high-frame-rate videos, we compute the continuous-time trajectory of the particles over the duration of the sequences and render the occluded frames at a high framerate. Events are generated from these rendered images using ESIM [31]. Finally, for each sequence, we create synchronized event sequences, the occluded image, groundtruth occlusion mask, and groundtruth background image. These groundtruth occlusion masks are only used as an input to the image inpainting baselines. We consider grayscale images instead
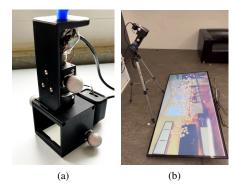

Figure 4: For the recording of our real-world dataset, we use a beamsplitter setup consisting of an event and a frame camera (a) to record particles moving on top of a screen (b). The experimental setup consists of the TV screen displaying a background image with the occluding particles rolling on the screen.

of RGB images because events from the event camera only provide relative intensity changes but lack color information. Therefore, the reconstruction of RGB images using events introduces additional challenges.

4.2. Real-world Dataset

We build a dynamic occlusion dataset where the event streams are captured by a Prophesee Gen4 camera [5], with a resolution of 1280×720, and the images with a FLIR Blackfly S camera with a resolution of $4000{\times}3000$. The images are captured at a framerate of 15 fps and with an exposure time of $7\,\mathrm{ms}$. The cameras are hardware synchronized and mounted in a beam splitter setup (Fig. 4 (a)), which contains a mirror that splits the incoming light to the event and frame camera, ensuring alignment between events and frames. Similar to the synthetic dataset, background images are sampled from the MSCOCO dataset [17] and projected on a TV screen, which is recorded with the beam splitter setup pointed towards the screen. The occlusions are created by rolling Styrofoam balls over the TV screen as shown in Fig. 4 (b). For each sequence, the ground truth is collected by recording the image projected on the screen without any occlusions. We sample 3 frames from each sequence corresponding to low, medium, and high occlusion densities as shown in Fig. 5. Overall, the real-world dataset consists of 154 training sequences and 79 test sequences.

5. Experiments

Implementation details Our models are implemented in Pytorch [27] and trained from scratch with random weights on our synthetic dataset. For the experiments on our real dataset, we finetune our pre-trained network on the real training data. We use Adam optimizer[11] with a learning rate of $1\mathrm{e}{-}3$ and batch size of 4 and train for 1250 epochs
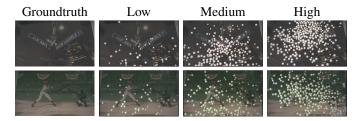
Figure 5: Our real-world dataset contains samples with low, medium, and high occlusion density.

Method	Input	PSNR ↑	SSIM ↑	MAE ↓
MAT [12]	I	30.7620	0.9217	0.0107
MISF [13]	I	31.1884	0.9229	0.0101
PUT [20]	I	26.9858	0.8608	0.0187
ZITS [4]	I	31.2971	0.9328	0.0100
EF-SAI [15]	I+E	26.7557	0.8688	0.0257
E2VID [32]	E	19.2278	0.6086	0.0508
Ours (Acc. Method)	E	20.3444	0.6955	0.0373
Ours (Learning)	I+E	**34.6203**	**0.9536**	**0.0085**

Table 1: Reconstruction performance on our synthetic dataset. 'I' and 'E' stand for image and events, respectively.

Method	Input	PSNR ↑					
Coverage		10%	20%	30%	40%	50%	60%
MAT [12]	I	35.4937	31.4344	31.895	29.8665	27.7057	28.1765
MISF [13]	I	35.6868	31.8720	31.7689	30.5100	28.3776	28.9148
PUT [20]	I	32.7979	28.6378	27.0988	26.1156	23.7555	23.5093
ZITS [4]	I	35.6619	31.9159	32.3340	30.4923	28.3828	28.9959
EF-SAI [15]	I+E	29.2994	27.9049	27.7354	26.2447	24.7267	24.6231
E2VID [32]	E	22.0048	19.8868	20.2063	18.1452	17.5353	17.5884
Ours (Acc. Method)	E	28.0390	22.4906	20.6897	17.9743	17.0200	15.8527
Ours (Learning)	I+E	**40.1286**	**36.3622**	**35.8476**	**33.2527**	**31.0083**	**31.1224**

Table 2: Reconstruction performance on our synthetic dataset in terms of PSNR divided according to the occlusion density of the test samples.

on the synthetic and real dataset. We supervise our network with the $L1$ loss computed on the predicted and ground truth occlusion-free images. To measure the quality of the reconstructed image, we use structural similarity index measure (SSIM) [45], peak signal-to-noise ratio (PSNR), and mean absolute error (MAE).

Baselines We evaluate our approach against four state-of-the-art image-inpainting solutions: MAT [12], MISF [13], PUT [20], and ZITS [4]. These baselines are trained on challenging large-scale datasets with over eight million images. Thus, for our experiments, we use the weights provided by the authors. Since dynamic occlusion removal from a single viewpoint using an event camera is a novel task, there currently do not exist any event-based approaches that tackle this problem directly. The closest baseline uses events and frames to reconstruct the background image from multiple viewpoints (EF-SAI) [15]. EF-SAI [15] uses a refocus module with events to blur the foreground and focus on the target depth plane of the background. We adapt this baseline for our task by providing it with a single image and events from only a single viewpoint and train this network on our dataset from scratch. Additionally, we also compare against events-to-image reconstruction methods [32]. This method is adapted by combining the reconstructed event images with the occluded images using the groundtruth mask. Note that all of the evaluated baselines were originally not designed for this particular task but are the closest related methods applicable to our task.

5.1. Results on Synthetic Dataset

To compare our method against the baselines in controlled conditions, we evaluate all methods on our synthetic dataset. Table 1 summarizes the quantitative results. Our method outperforms the best image inpainting baseline by $3\,\mathrm{dB}$ in terms of PSNR. Fig. 6 shows qualitative comparisons between different methods. Image inpainting methods tend to hallucinate the background, resulting in visually more pleasing information rather than the true scene. An example of this can be seen in Fig. 6 third row, where the inpainting method is unable to reconstruct the characters on

the bus, whereas our method is able to better preserve this information. We also outperform the event-based synthetic aperture imaging baseline EF-SAI [15] and the event-to-image reconstruction baseline E2VID [32]. Although the input to our method and EF-SAI consists of events and a single image, the EF-SAI baseline was designed for multi-view reconstruction. The simple event accumulation baseline is one of the lowest performing baselines even with the knowledge of the correct contrast threshold, as the occlusions are too complex for the basic event generation model to capture the intensity changes as discussed in the Section 3. We also analyze the effect of occlusion density on the performance of all the methods and summarize them in Table 2. As expected, increasing the occlusion density decreases the performance of all the approaches. However, at higher occlusion densities, the image inpainting methods drastically degrade in performance as they tend to hallucinate occluded areas. In contrast, our method uses events that provide continuous intensity changes, which results in a better performance at higher occlusion densities. We provide more qualitative results of other baselines in the supplementary material.

Ablation Study To study the effect of the EAM and OFF modules in our network, we report the network performance by removing the specific modules. For computational reasons, we performed the ablation of the OFF module without
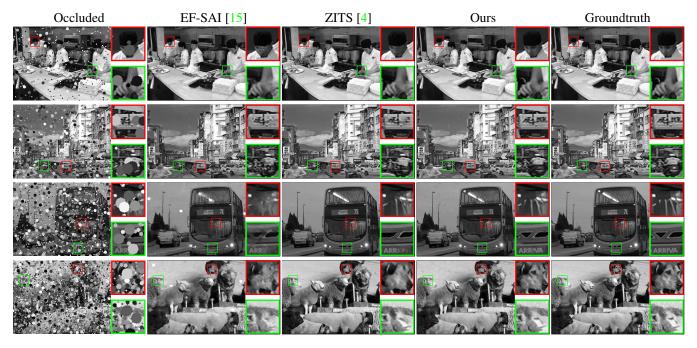
Figure 6: Images showing the occluded input frame, the reconstructed frame of the best-performing frame-based and event-based baselines and our method, and the ground truth frame for our synthetic dataset.

Method	Input	PSNR ↑	SSIM ↑	MAE ↓
Ours (w/o EAM & OFF)	I+E	31.0776	0.9212	0.0139
Ours (w/o EAM)	I+E	32.7652	0.9425	0.0102
Ours (Full)	I+E	**34.6203**	**0.9536**	**0.0085**

Table 3: The reconstruction performance on the synthetic dataset of our network obtained by removing adaptively the introduced modules.

including the EAM module, which reduces significantly the training time. As shown in Table 3, adding the OFF module to the network improves the performance with respect to all metrics. This improvement can be explained by the ability of the OFF module to better localize the occlusions and adaptively fuse image and event features throughout the network scales. Including the EAM module leads to an even higher performance increase of 1.9 dB in PSNR. One possible reason for the improvement over the network without recurrent event encoding is that the EAM module is better at handling multiple overlapping occlusions by selecting events relevant to the true background and ignoring the redundant event information. For more ablation studies showing the effect of the event and image features on the textured and uniform image regions, we refer to the supplementary material.

5.2. Results on Real Dataset

We also evaluate our method on our real dataset collected with a custom build beamsplitter setup. Unlike the synthetic dataset, we do not have an occlusion mask available for the sequences. We, therefore, approximate the mask by subtracting the occluded frame from the groundtruth frame and applying thresholding. This mask is used as an input for the image inpainting methods and event image reconstruction baselines. The quantitative results are shown in Table 4. Our results on the real dataset follow a similar trend to the synthetic dataset. In Fig. 7, we show the qualitative results for different sequences. While the image inpainting method results in visually clean images, the hallucination artifacts can be clearly seen in the small image patches, e.g., the fingers around the camera are missing in the sample visualized in the first row or the watch in the sample shown in the second row. Our method, on the other hand, can reconstruct the background information accurately.

6. Conclusion

We introduce a novel event-based approach for background image reconstruction in the presence of dynamic occlusions. It leverages the complementary nature of event camera and frames to reconstruct true scene information instead of hallucinating occluded areas as done by image inpainting approaches. Specifically, our proposed data-driven approach reconstructs the background image using only one
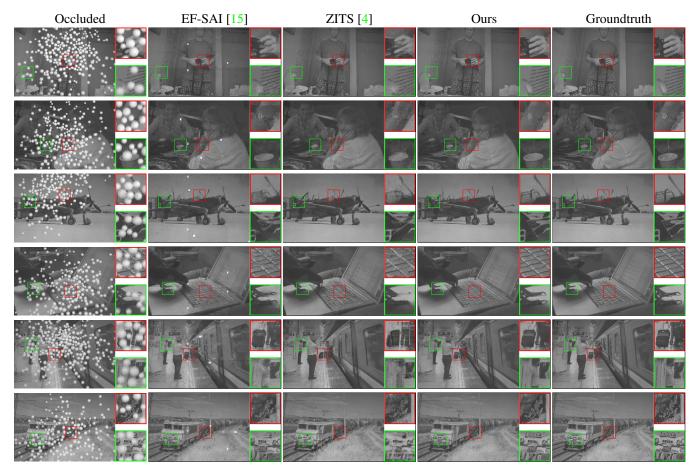
Figure 7: Images showing the occluded input frame, the reconstructed frame of the best-performing frame-based and event-based baselines and our method, and the ground truth frame for our real-world dataset.

Method	Input	PSNR ↑	SSIM ↑	MAE ↓
MAT [12]	I	26.7451	0.5670	0.0285
MISF [13]	I	29.0281	0.6951	0.0222
PUT [20]	I	18.9135	0.2726	0.0801
ZITS [4]	I	30.2675	0.7542	0.0190
EF-SAI [15]	I+E	30.2390	0.8037	0.0194
E2VID [32]	E	16.5642	0.3203	0.1066
Ours ((Acc. Method)	E	19.5110	0.5685	0.0606
Ours (Learning)	I+E	**33.0701**	**0.8173**	**0.0166**

Table 4: Reconstruction performance on our real-world dataset.

groundtruth images. Our method achieves an improvement of 3dB in PSNR over state-of-the-art frame-based and event-based methods on both synthetic and real datasets. We will release our synthetic and recorded dataset representing the first datasets for background image reconstruction using events and images in the presence of dynamic occlusions. We believe that our proposed method and dataset lay the foundation for future research.

occluded image and events. The high temporal resolution of the events provides our method additional information on the relative intensity changes between the foreground and background, making it robust to dense occlusions. To evaluate our approach, we present the first large-scale dataset recorded in the real world containing over 230 challenging scenes with synchronized events, occluded images, and

Seeing Behind Dynamic Occlusions with Event Cameras
—Supplementary Material—

Rong Zou   Manasi Muglikar   Nico Messikommer   Davide Scaramuzza
Robotics and Perception Group, University of Zurich, Switzerland

7. Network Structure

In this section, we provide details of the structure of our network. As explained in the main paper, our network consists of the *Frame Encoder*, the *Event Encoder*, the *Event Accumulation Module* (EAM), the *Occlusion-aware Feature Fusion* (OFF) module and the *Reconstruction Decoder*.

The architecture of the Frame Encoder and the Event Encoder is shown in Table 5. Both encoders have the same structure except at the first two scales, where the frame encoder is provided with a grayscale image of dimension (h, w, 1) and the event encoder receives the output of the EAM of dimension (h, w, 32). The architecture of the EAM module is shown in Fig. 8.

We provide the details of the OFF module in Table 6. The OFF module computes gating weights W^j at each scale j based on the current-scale event features F_E^j and frame features F_I^j as well as the previous-scale fused features F_F^{j-1} downscaled to $\bar{F}_F^{j-1}$. The fused feature F_F^j is then obtained based on the current-scale event and frame features F_E^j and F_I^j, the gating weights W^j, and the downsampled previous-scale fused feature $\bar{F}_F^{j-1}$. The OFF module can be described with the following equations:

$$\tilde{F}_F^j = G(F_E^j, F_I^j) \tag{4}$$

$$W^j = G(\tilde{F}_F^j, \bar{F}_F^{j-1}) \tag{5}$$

$$\hat{F}_F^j = (1 - W^j)F_I^j + W^j F_E^j \tag{6}$$

$$F_F^j = G(\hat{F}_F^j, \bar{F}_F^{j-1}) \tag{7}$$

where G indicates concatenation followed by convolution operations.

Finally, the structure of the *Reconstruction Decoder* is shown in Table 7 which uses the fused features at each scale as skip connections to predict the residual image.

8. Detailed Ablation Results

Here we present a more detailed ablation study on the network architecture.

	Frame Encoder		**Event Encoder**	
Scale j	Layer	Out Size (h^j, w^j, c^j)	Layer	Out Size (h^j, w^j, c^j)
0	input image	(h, w, 1)	–	–
1	conv(5,1,2)	(h, w, 32)	–	–
2	conv(5,2,2)	(h/2, w/2, 64)	conv(5,2,2)	(h/2, w/2, 64)
3	conv(5,2,2)	(h/4, w/4, 128)	conv(5,2,2)	(h/4, w/4, 128)
4	conv(5,2,2)	(h/8, w/8, 256)	conv(5,2,2)	(h/8, w/8, 256)
5	conv(5,2,2)	(h/16, w/16, 512)	conv(5,2,2)	(h/16, w/16, 512)
6	conv(3,1,1)	(h/16, w/16, 512)	conv(3,1,1)	(h/16, w/16, 512)

Table 5: The architecture of the *Frame Encoder* and *Event Encoder* in our proposed method, in which (h^j, w^j, c^j) denotes the height, width, and number of channels of the output of the layer corresponding to scale j, and conv(k, s, p) denotes a convolution block with kernel size k, stride s and padding p. The input layer corresponds to $j = 0$, and the input size used in our paper is h= $h^0 = 384$, w= $w^0 = 512$.

Occlusion-aware Feature Fusion (OFF)			
Input	Input Size	Layer	Output
F_E^j & F_I^j	(h^j, w^j, c^j)	conv(5,1,2) $\times$ 2	$\tilde{F}_F^j$
F_F^{j-1}	$(h^{j-1}, w^{j-1}, c^{j-1})$	conv(5,s^j,2)	$\bar{F}_F^{j-1}$
$\bar{F}_F^{j-1}$ & $\tilde{F}_F^j$	(h^j, w^j, c^j)	conv(5,1,2) $\times$ 2	W^j
F_E^j, F_I^j & W^j	(h^j, w^j, c^j)	Feature Selection	$\hat{F}_F^j$
$\hat{F}_F^j$ & $\bar{F}_F^{j-1}$	(h^j, w^j, c^j)	conv(5,1,2) $\times$ 2	F_F^j

Table 6: The structure of the *Occlusion-aware Feature Fusion* module in the proposed method.

Reconstruction Decoder			
Scale	Input Size	Layer	Output Size
6	(h/16, w/16, 512+512)	up(2)-conv(5,1,2)	(h/8, w/8, 512)
5	(h/8, w/8, 512+256)	up(2)-conv(5,1,2)	(h/4, w/4, 256)
4	(h/4, w/4, 256+128)	up(2)-conv(5,1,2)	(h/2, w/2, 128)
3	(h/2, w/2, 128+64)	up(2)-conv(5,1,2)	(h, w, 64)
2	(h, w, 64+32)	conv(3,1,1)	(h, w, 32)
1	(h, w, 32)	conv(1,1,0)	(h, w, 1)

Table 7: The structure of the *Reconstruction Decoder*, in which 'up(2)' denotes updampling by 2.
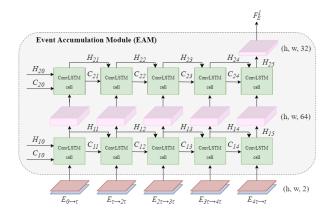
Figure 8: The EAM in our proposed method which integrates $N = 5$ event representations $E_{0 \to \tau}, ..., E_{(N-1)\tau \to t}$ using two ConvLSTM layers. The final hidden state of the second ConvLSTM layers formulates the first-scale event feature F_E^1 in the multi-scale *Event Encoder*.

Method	Used Feature Scale Image —— Event		PSNR ↑	SSIM ↑	MAE ↓
Shared F & E Encoder	All scale	All scale	27.8297	0.8642	0.0217
Indep. F Encoder	All scale	1st scale	31.6567	0.9297	0.0133
Indep. E Encoder	1st scale	All scale	31.3298	0.9263	0.0125
Indep. F & E Encoder Simple Fusion	All scale	All scale	31.0776	0.9212	0.0139
Indep. F & E Encoder with OFF	All scale	All scale	**32.7652**	**0.9425**	**0.0102**

Table 8: The effects of event and frame encoders and feature fusion methods on reconstruction performance evaluated on the synthetic dataset

8.1. Importance of Event and Frame Encoders

We study four different strategies to combine the event and frame information and train each corresponding network on our synthetic dataset. The first strategy combines events and frame by concatenating them at the input stage and uses a single encoder-decoder structure to predict the reconstructed image, termed as *Shared F & E Encoder*. The second strategy, termed *Independent Frame Encoder*, uses an independent frame encoder, which computes the frame features without considering the event information. A second encoder takes as input the events and fuses for each subsequent scale the frame features from the independent frame encoder with the feature from the previous scale. As the third strategy, we switch the inputs of the second strategy, i.e., we use an independent event encoder and a second encoder for fusing the event features with shallow frame features (referred to as *Independent Event Encoder*). Since we want to focus on the effect of processing the features for each sensing modality separately, we use two independent encoders for events and frame respectively and fuse features at each scale using a simple convolution layer. This is referred to as *Independent Event and Frame Encoder Simple Fusion*. Lastly, to show the effect of a more sophisticated fusion, we also show the reconstructions using our OFF module to perform the fusion of the event and frame feature at each scale (*Independent Event and Frame Encoder with OFF*).

We show the qualitative comparison of these approaches in Fig. 9 and provide quantitative evaluations in Table 8. The *Shared F & E Encoder* results in the worst performance in both the textured and untextured regions of the image. The *Independent Frame Encoder* results in an improved performance in uniform areas, whereas textured regions are poorly reconstructed, as can be observed in the third column of Fig. 9. The uniform texture of the red patch is captured perfectly by this network, indicating the importance of image features for uniform areas. This can be explained by the fact that images contain absolute intensity information of surrounding spatial regions, which provides more information to the network for filling in the uniform patches. However, the edges are not well preserved, which results in bleeding edges, as seen in the edge of the letter 'D' in the first row of Fig. 9. The *Independent Event Encoder*, in comparison to the *Independent Frame Encoder*, shows better ability in reconstructing textured regions, which indicates the importance of event features for reconstructing textured areas. However, it struggles to remove occlusions in uniform regions,

Both the second and the third strategy performs significantly better than the first one, highlighting the importance of having separate encoders for frames and events to enable high-quality reconstruction of both textured and untextured regions. However, a vanilla fusion of *Independent Event and Frame Encoder Simple Fusion* does not take full advantage of the strengths of the two respective modalities and therefore results in lower performance, as can be seen in the fourth row in Table 8. In contrast, the *Independent Event and Frame Encoder with OFF* achieves higher performance than all the above-mentioned strategies. Especially, it can better reconstruct textured regions, e.g., the edge of the letter 'D' is recovered, as visualized in the first row of Fig. 9. We believe that the differential nature of event cameras enables the preservation of high-frequency structures such as edges.

In summary, the proposed model using independent encoders for frame and events with a sophisticated fusion shows the best overall performance.

8.2. Importance of EAM and OFF Modules

In this section, we study the impact of *Event Accumulation module* (EAM) and *Occlusion-aware Feature Fusion* (OFF) on network performance. The quantitative performance is shown in Table 9 and the qualitative results can be seen in Fig. 10. The base model without EAM and OFF results in the worst performance of $31dB$ PSNR, and adding
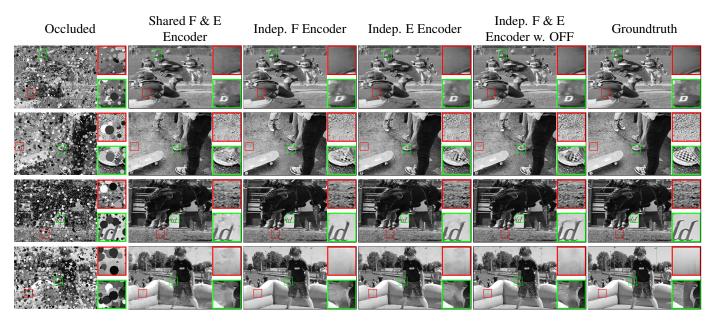
Figure 9: Images showing the importance of independent image and event encoders, conducted without EAM. A uniform region and a region containing interesting textures are shown in the red and green square boxes, respectively.
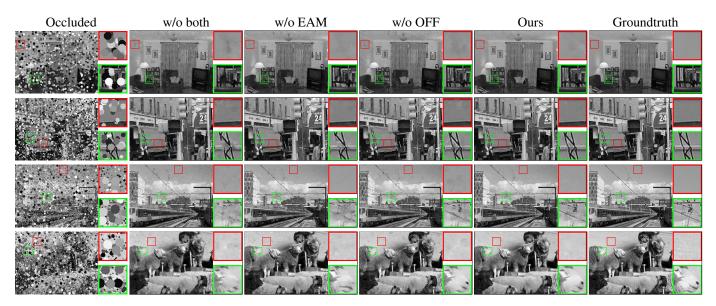

Figure 10: Images showing the importance of EAM and OFF modules. A uniform region and a region containing interesting textures are shown in the red and green square boxes, respectively.

EAM improves it by $1.8dB$ in terms of PSNR. This gain is achieved because the EAM is designed to select event data that is relevant to the true background and filter out redundant information. As can be seen in Fig. 10, the network with EAM (4^{th} column) shows better results in both uniform and textured regions compared to the base model. Also, in contrast to the two methods without EAM, the true background information is recovered to a significantly better extent in textured areas, even where there are multiple overlapping occlusions.

The performance of the network gains a substantial improvement from the OFF module as well, with a $1.7dB$ increase in PSNR. As illustrated in Fig. 10, the OFF module also improves reconstruction quality in both uniform and textured areas when compared with the base model. In addition, compared to the two models without OFF, the one with OFF is significantly better at reconstructing uniform or non-textured regions.

Method	Input	PSNR ↑	SSIM ↑	MAE↓
w/o EAM & OFF	I+E	31.0776	0.9212	0.0139
w/o EAM	I+E	32.7652	0.9425	0.0102
w/o OFF	I+E	32.8892	0.9378	0.0119
Full	I+E	**34.6203**	**0.9536**	**0.0085**

Table 9: The reconstruction performance on the synthetic dataset of our network obtained by removing adaptively the introduced modules.

Combining EAM and OFF, the full model achieves an increase of more than $3.5dB$ PSNR over the base model. Qualitatively, the full model performs better than all other models in both types of regions, implying that the two modules benefit from each other.

9. Additional Qualitative Results

We show the qualitative comparisons between our method and all the baselines on our synthetic and real-world datasets in Fig. 11 and Fig. 12 respectively. Qualitatively, our method outperforms other baselines for both datasets. We re-state our conclusions that image inpainting baselines tend to hallucinate the missing areas, resulting in clean but inaccurate image reconstructions. In comparison, our method is better at preserving the details of the original image, as can be seen in the highlighted patches.

10. Inference Time

In this section, we compare the computational complexity of the methods in terms of inference time. Note that while it is not feasible to compute the precise inference time, as GPU loads vary over time due to other processes running simultaneously, we approximate this number by averaging the inference time over the entire test set. The runtime is computed for a batch size of 1 using a Quadro RTX 8000 GPU for the learning-based baselines and Intel(R) Core(TM) i7-3720QM CPU @ 2.60GHz for our model-based baseline. The comparison for all the baselines is presented in Table 10. Note that only our model-based baseline is run on CPU and therefore marked with a star. Our method has an inference time of 0.18 sec, slightly slower than the fastest image inpainting baseline MAT [12]. The best performing image-inpainting baseline ZITS [4], on the other hand, has an inference time of 1.94 sec.

Method	MAT [12]	MISF [13]	PUT [20]	ZITS [4]	EF-SAI [15]	E2VID [32]	Ours (Acc. Method)*	Ours (Learning)
Input	I	I	I	I	I+E	E	E	I+E
Infer. time	0.15	0.16	>60	1.94	3.13	0.01	0.001	0.18

Table 10: Comparison of inference time (sec).

References

[1] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE Transactions on Image Processing*, 10(8):1200–1211, 2001. 2

[2] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, page 417–424, 2000. 1, 2

[3] A. Criminisi, P. Perez, and K. Toyama. Object removal by exemplar-based inpainting. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–II, 2003. 2

[4] Qiaole Dong, Chenjie Cao, and Yanwei Fu. Incremental transformer structure enhanced image inpainting with masking positional encoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 6, 7, 8, 12, 13, 14

[5] Thomas Finateu, Atsumi Niwa, Daniel Matolin, Koya Tsuchimoto, Andrea Mascheroni, Etienne Reynaud, Pooria Mostafalu, Frederick Brady, Ludovic Chotard, Florian LeGoff, Hirotsugu Takahashi, Hayato Wakabayashi, Yusuke Oike, and Christoph Posch. A 1280x720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86μm pixels, 1.066geps readout, programmable event-rate controller and compressive data-formatting pipeline. In *IEEE Intl. Solid-State Circuits Conf. (ISSCC)*, 2020. 3, 5

[6] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. 2

[7] Xiefan Guo, Hongyu Yang, and Di Huang. Image inpainting via conditional texture and structure dual generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14134–14143, October 2021. 2

[8] Jin Han, Chu Zhou, Peiqi Duan, Yehui Tang, Chang Xu, Chao Xu, Tiejun Huang, and Boxin Shi. Neuromorphic camera guided high dynamic range imaging. In *CVPR*, 2020. 3

[9] James Hays and Alexei A Efros. Scene completion using millions of photographs. *ACM Transactions on Graphics (SIGGRAPH 2007)*, 26(3), 2007. 1, 2

[10] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and Locally Consistent Image Completion. *ACM Transactions on Graphics (Proc. of SIGGRAPH)*, 36(4), 2017. 2

[11] Diederik P. Kingma and Jimmy L. Ba. Adam: A method for stochastic optimization. *Int. Conf. Learn. Representations (ICLR)*, 2015. 5

[12] Wenbo Li, Zhe Lin, Kun Zhou, Lu Qi, Yi Wang, and Jiaya Jia. Mat: Mask-aware transformer for large hole image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 6, 8, 12, 13, 14

Figure 11: Images showing the occluded input frame, the reconstructed frame of all methods, and the ground truth frame for our synthetic dataset.

[13] Xiaoguang Li, Qing Guo, Di Lin, Ping Li, Wei Feng, and Song Wnag. Misf: Multi-level interactive siamese filtering for high-fidelity image inpainting. *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022. 2, 6, 8, 12, 13, 14

[14] Xiaoyu Li, Bo Zhang, Jing Liao, and Pedro V. Sander. Let's see clearly: Contaminant artifact removal for moving cam-

eras. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1991–2000, 2021. 2

[15] Wei Liao, Xiang Zhang, Lei Yu, Shijie Lin, Wen Yang, and Ning Qiao. Synthetic aperture imaging with events and frames. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages
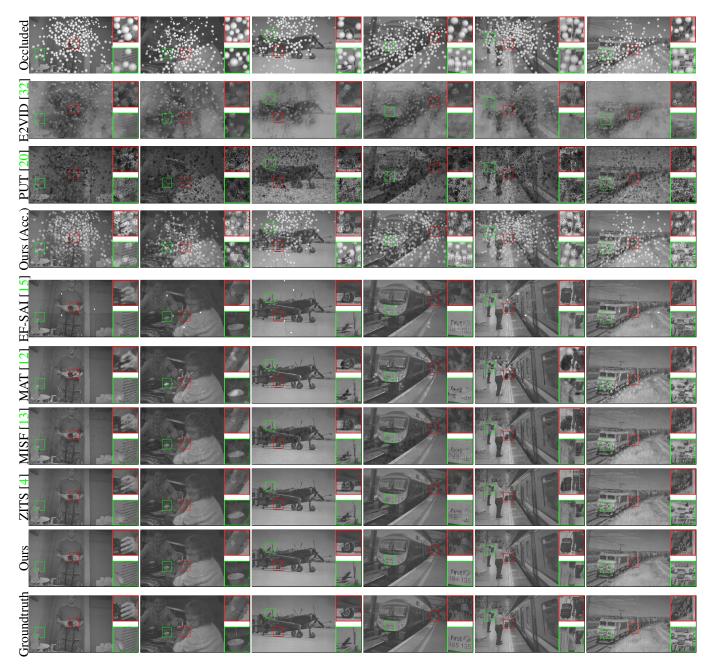
Figure 12: Images showing the occluded input frame, the reconstructed frame of all methods, and the ground truth frame for our real-world dataset.

17735–17744, June 2022. 3, 6, 7, 8, 12, 13, 14

[16] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128×128 120 dB 15 μs latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circuits*, 43(2):566–576, 2008. 2, 3

[17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 740–755, 2014. 5

[18] Guilin Liu, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *The European Conference on Computer Vision (ECCV)*, 2018. 1, 2

[19] Hongyu Liu, Bin Jiang, Yi Xiao, and Chao Yang. Coherent semantic attention for image inpainting. In *IEEE International Conference on Computer Vision (ICCV)*, July 2019. 1,

2

[20] Qiankun Liu, Zhentao Tan, Dongdong Chen, Qi Chu, Xiyang Dai, Yinpeng Chen, Mengchen Liu, Lu Yuan, and Nenghai Yu. Reduce information loss in transformers for pluralistic image inpainting. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022. 1, 2, 6, 8, 12, 13, 14

[21] Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Learning to see through obstructions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1, 2

[22] W. Matusik, A. Hornung, S. Avidan, and T. Basha. Structure and motion from scene registration. In *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1426–1433. IEEE Computer Society, 2012. 2

[23] Nico Messikommer*, Stam Georgoulis*, Daniel Gehrig, Stepan Tulyakov, Julius Erbach, Alfredo Bochicchio, Yuanyou Li, and Davide Scaramuzza. Multi-bracket high dynamic range imaging with event cameras. *IEEE Conf. Comput. Vis. Pattern Recog. Workshops (CVPRW)*, 2022. 3

[24] S.M. Mostafavi I., Lin Wang, Yo-Sung Ho, and Kuk-Jin Yoon Yoon. Event-based high dynamic range image and very high frame rate video generation using conditional generative adversarial networks. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019. 3

[25] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edgeconnect: Structure guided image inpainting using edge prediction. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. 2

[26] Liyuan Pan, Cedric Scheerlinck, Xin Yu, Richard Hartley, Miaomiao Liu, and Yuchao Dai. Bringing a blurry frame alive at high frame-rate with an event camera. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019. 3

[27] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Workshops*, 2017. 5

[28] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2536–2544, 2016. 2

[29] Zhao Pei, Yanning Zhang, Xida Chen, and Yee-Hong Yang. Synthetic aperture imaging using pixel labeling via energy minimization. *Pattern Recognition*, 46(1):174–187, 2013. 2

[30] Christoph Posch, Daniel Matolin, and Rainer Wohlgenannt. A QVGA 143 dB dynamic range frame-free PWM image sensor with lossless pixel-level video compression and time-domain CDS. *IEEE J. Solid-State Circuits*, 46(1):259–275, Jan. 2011. 3

[31] Henri Rebecq, Daniel Gehrig, and Davide Scaramuzza. ESIM: an open event camera simulator. In *Conf. on Robotics Learning (CoRL)*, 2018. 5

[32] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. Events-to-video: Bringing modern computer vision to event cameras. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019. 3, 6, 8, 12, 13, 14

[33] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019. 3

[34] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2015. 3

[35] Zheng Shi, Yuval Bahat, Seung-Hwan Baek, Qiang Fu, Hadi Amata, Xiao Li, Praneeth Chakravarthula, Wolfgang Heidrich, and Felix Heide. Seeing through obstructions with diffractive cloaking. *ACM Trans. Graph.*, 41(4), 2022. 1, 3

[36] Lin Songnan, Zhang Jiawei, Pan Jinshan, Jiang Zhe, Zou Dongqing, Wang Yongtian, Chen Jing, and S Ren Jimmy. Learning event-driven video deblurring and interpolation. *Eur. Conf. Comput. Vis. (ECCV)*, 2020. 3

[37] Yunjae Suh, Seungnam Choi, Masamichi Ito, Jeongseok Kim, Youngho Lee, Jongseok Seo, Heejae Jung, Dong-Hee Yeo, Seol Namgung, Jongwoo Bong, Jun seok Kim, Paul K. J. Park, Joonseok Kim, Hyunsurk Ryu, and Yongin Park. A 1280x960 Dynamic Vision Sensor with a 4.95-μm pixel pitch and motion artifact minimization. In *IEEE Int. Symp. Circuits Syst. (ISCAS)*, 2020. 3

[38] 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. *IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, 2022. 1, 2

[39] Zinovi Tauber, Ze-Nian Li, and Mark S. Drew. Review and preview: Disocclusion by inpainting for image-based rendering. *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)*, 37(4):527–540, July 2007. 1

[40] Stepan Tulyakov, Alfredo Bochicchio, Daniel Gehrig, Stamatios Georgoulis, Yuanyou Li, and Davide Scaramuzza. Time lens++: Event-based frame interpolation with parametric non-linear flow and multi-scale fusion. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3

[41] V. Vaish, M. Levoy, R. Szeliski, C.L. Zitnick, and Sing Bing Kang. Reconstructing occluded surfaces using synthetic apertures: Stereo, focus and robust measures. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2331–2338, 2006. 2

[42] Vaibhav Vaish, Bennett Wilburn, Neel Joshi, and Marc Levoy. Using plane + parallax for calibrating dense camera arrays. *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, 1:I–I, 2004. 2

[43] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. *Int. Conf. Comput. Vis. (ICCV)*, 2021. 1, 2

[44] Yingqian Wang, Tianhao Wu, Jungang Yang, Longguang Wang, Wei An, and Yulan Guo. DeOccNet: Learning to see through foreground occlusions in light fields. In *Winter Conference on Applications of Computer Vision (WACV)*, Mar 2020. 1, 2

[45] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, Apr. 2004. 6

[46] Ziwei Wang, Yonhon Ng, Cedric Scheerlinck, and Robert Mahony. An asynchronous kalman filter for hybrid event cameras. In *ICCV*, pages 448–457, October 2021. 3

[47] Bennett Wilburn, Neel Joshi, Vaibhav Vaish, Eino-Ville Talvala, Emilio Antunez, Adam Barth, Andrew Adams, Mark Horowitz, and Marc Levoy. High performance imaging using large camera arrays. *ACM Transactions on Graphics*, 24(3):765–776, July 2005. 2

[48] SHI Xingjian, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *NeurIPS*, pages 802–810, 2015. 3

[49] Tao Yang, Yanning Zhang, Jingyi Yu, Jing Li, Wenguang Ma, Xiaomin Tong, Rui Yu, and Lingyan Ran. All-in-focus synthetic aperture imaging. In *Computer Vision – ECCV 2014*. Springer International Publishing, 2014. 1, 2

[50] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018. 2

[51] Lei Yu, Xiang Zhang, Wei Liao, Wen Yang, and Gui-Song Xia. Learning to see through with events. In *IEEE TPAMI*, 2022. 3

[52] Song Zhang, Yu Zhang, Zhe Jiang, Dongqing Zou, Jimmy Ren, and Bin Zhou. Learning to see in the dark with events. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 666–682, Cham, 2020. Springer International Publishing. 3

[53] Xiang Zhang, Wei Liao, Lei Yu, Wen Yang, and Gui-Song Xia. Event-based synthetic aperture imaging with a hybrid network. In *CVPR*, 2021. 3

[54] Yu Zhiyang, Zhang Yu, Liu Deyuan, Zou Dongqing, Chen Xijun, Liu Yebin, and S Ren Jimmy. Training weakly supervised video frame interpolation with events. *Int. Conf. Comput. Vis. (ICCV)*, 2021. 3