

LiDAR-Event Stereo Fusion with Hallucinations

Luca Bartolomei , Matteo Poggi , Andrea Conti , and
Stefano Mattoccia 

University of Bologna, Italy
<https://eventvpstereo.github.io/>

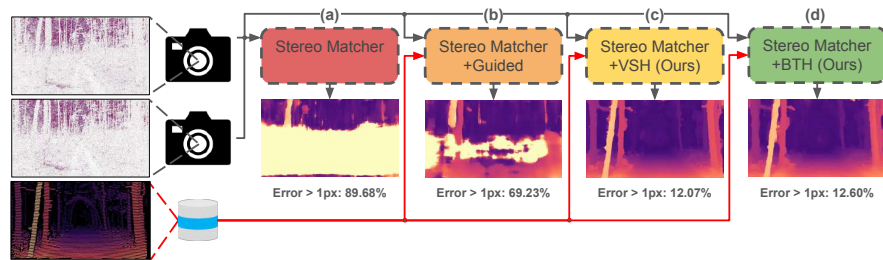


Fig. 1: LiDAR-Event Stereo Fusion with Hallucinations. In the absence of motion or brightness changes, sparse event streams lead stereo models to catastrophic failures (a). A LiDAR sensor can be used with existing strategies [49] to soften this problem, yet with limited impact (b), whereas our proposals are superior (c,d).

Abstract. Event stereo matching is an emerging technique to estimate depth from neuromorphic cameras; however, events are unlikely to trigger in the absence of motion or the presence of large, untextured regions, making the correspondence problem extremely challenging. Purposely, we propose integrating a stereo event camera with a fixed-frequency active sensor – e.g., a LiDAR – collecting sparse depth measurements, overcoming the aforementioned limitations. Such depth hints are used by hallucinating – i.e., inserting fictitious events – the stacks or raw input streams, compensating for the lack of information in the absence of brightness changes. Our techniques are general, can be adapted to any structured representation to stack events and outperform state-of-the-art fusion methods applied to event-based stereo.

1 Introduction

Depth estimation is a fundamental task with many applications ranging from robotics, 3D reconstruction, and augmented/virtual reality to autonomous vehicles. Accurate, prompt, and high-resolution depth information is crucial for most of these tasks, but obtaining it remains an open challenge. Among the many possibilities, depth-from-stereo is one of the longest-standing approaches

to deal with it, with a large literature of deep architecture [51] proposed in the last decade for processing rectified color images.

Event cameras [18] (or neuromorphic cameras) are recently emerging as an alternative to overcome the limitations of traditional imaging devices, such as their low dynamic range or the motion blur caused by fast movements. Unlike their traditional counterparts, event cameras do not capture frames at synchronous intervals. Instead, they mimic the dynamic nature of human vision by reporting pixel intensity changes, which can have *positive* or *negative* polarities, as soon as they happen. This peculiarity endows them with unparalleled features – notably microsecond temporal resolution, and an exceptionally high dynamic range – making them perfectly suited for applications where fast motion and challenging light conditions are persistent issues (*e.g.* autonomous driving). The events streams are often encoded in $W \times H \times C$ tensors, thus being fully compatible with the CNNs used for classical stereo [43], capable of estimating dense disparity maps driven by data, despite the sparse nature of events.

However, as the events trigger only with brightness changes any derived data is *semi-dense* and uninformative, for instance, when facing large untextured regions or in the absence of any motion – *e.g.*, as in the example in Fig. 1. This makes the downstream stereo network struggle to match events across left and right cameras, as shown in Fig. 1 (a). According to the RGB stereo literature, fusing color information with sparse depth measurements from an active sensor [4, 12, 49, 82] (*e.g.*, a LiDAR) considerably softens the weaknesses of passive depth sensing, despite the much lower resolution at which depth points are provided. We argue that such a strategy would counter the aforementioned issues even if applied to the event stereo paradigm, yet with a notable nodus caused by the fixed rate at which depth sensors work – usually, 10Hz for LiDARs – being in contrast with the asynchronous acquisition rate of event cameras. This would cause to either i) use depth points only when available, harming the accuracy of most fusion strategies known from the classical stereo literature [12, 49, 82], or ii) limiting processing to the LiDAR pace, nullifying one of the greatest strength of event cameras – *i.e.*, microseconds resolution. Nonetheless, this track on event stereo/active sensors fusion has remained unexplored so far.

In this paper, starting from the RGB literature [4, 12, 49, 82], we embark on a comprehensive investigation into the fusion of event-based stereo with sparse depth hints from active sensors. Inspired by [4], which projects distinctive color patterns on the images consistently with measured depth, we design a hallucination mechanism to generate fictitious events over time to densify the stream collected by the event cameras. Purposely, we propose two different strategies, respectively consisting of i) creating distinctive patterns directly at the stack level, *i.e.* a **Virtual Stack Hallucination** (VSH), just before the deep network processing, or ii) generating raw events directly in the stream, starting from the time instant t_d for which we aim to estimate a disparity map and performing **Back-in-Time Hallucination** (BTH). Both strategies, despite the different constraints – VSH requires explicit access to the stacked representation, whereas BTH does not – dramatically improve the accuracy of pure event-based

stereo systems, overcoming some of their harshest limitations as shown in Fig. 1 (c,d). Furthermore, despite depth sensors having a fixed acquisition rate that is in contrast with the asynchronous capture rate of event cameras, VSH and BTH can leverage depth measurements not synchronized with t_d (thus collected at $t_z < t_d$) with marginal drops in accuracy compared to the case of perfectly synchronized depth and event sensors ($t_z = t_d$). This strategy allows for exploiting both VSH and BTH while preserving the microsecond resolution peculiar of event cameras. Exhaustive experiments support the following claims:

- We prove that LiDAR-stereo fusion frameworks can effectively be adapted to the event stereo domain
- Our VSH and BTH frameworks are general and work effectively with any structured representation among the eight we surveyed
- Our strategies outperform existing alternatives inherited from RGB stereo literature on DSEC [21] and M3ED [9] datasets
- VSH and BTH can exploit even outdated LiDAR data to increase the event stream distinctiveness and ease matching, preserving the microsecond resolution of event cameras and eliminating the need for synchronous processing dictated by the constant framerate of the depth sensor

2 Related Work

Stereo Matching on color images. It is a longstanding open problem, with a large body of literature spanning from traditional approaches grounded on handcrafted features and priors [5, 24, 31, 36, 62, 68, 75, 76, 78] to contemporary deep learning approaches that brought significant improvements over previous methods, starting with [79]. Nowadays, the most effective solutions have emerged as end-to-end deep stereo networks [51], replacing the whole stereo pipeline with a deep neural network architecture through 2D and 3D architectures. The former, inspired by the U-Net model [53], adopts an encoder-decoder design [37, 42, 45, 50, 54, 59, 63, 64, 74, 77]. In contrast, the latter constructs a feature cost volume from image pair features and estimates the disparity map through 3D convolutions at the cost of substantially higher memory and runtime demands [10, 11, 13, 16, 23, 27, 28, 57, 70, 73, 80]. A recent trend in this field [34, 38, 65, 71, 83, 84] introduced innovative deep stereo networks that embrace an iterative refinement paradigm or use Vision Transformers [22, 35].

Stereo Matching with event cameras. This topic attracted significant attention due to the unique advantages of event sensors over traditional frame-based cameras. Similarly to conventional stereo matching, the first approaches focused on developing traditional algorithms by building structured representations, such as voxel grids [56], matched through handcrafted similarity functions [30, 56, 60, 85]. However, pseudo-images lose the high temporal resolution of the stream: to face this problem, [8, 52] handle events without an intermediate representation using an event-to-event matching approach, where for each reference event, a set of possible matches is given. Camuñas-Mesa *et al.* [7]

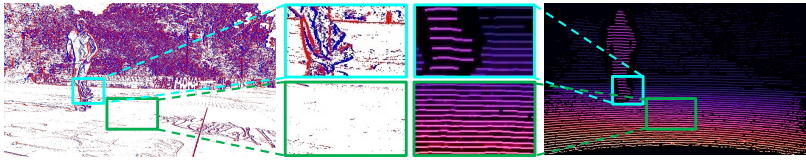


Fig. 2: Event cameras vs LiDARs – strengths and weaknesses. Event cameras provide rich cues at object boundaries where LiDARs cannot (cyan), yet LiDARs can measure depth where the lack of texture makes event cameras uninformative (green).

add filters to exploit orientation cues and increase matching distinctiveness. Instead, [47] revisited the cooperative network from [41]. Neural networks also showed promising results on event stereo matching with models directly processing raw events or using structured representation. The former are often inspired by [41] and typically employ Spiking Neural Networks (SNN) [1, 15, 44]. The latter adopts data-driven Convolutional Neural Networks (CNNs) to infer dense depth maps [43, 66, 67]. A detailed review of different event-based stereo techniques can be found in [17].

Sensor fusion for stereo. Recent research has delved into the fusion of color-cameras stereo vision with active sensors, starting with handcrafted algorithms: Badino *et al.* [2] integrated LiDAR data directly into the stereo algorithm using dynamic programming, Gandhi *et al.* [19] proposed an efficient seed-growing algorithm to fuse time-of-flight (ToF) depth data with stereo pairs, while Marin *et al.* [40] and Poggi *et al.* [48] exploited confidence measures. Eventually, contemporary approaches integrated depth from sensors with modern stereo networks, either by concatenating them to images as input [12, 46, 69, 81] or by using them to guide the cost optimization process by modulating existing cost volumes [25, 49, 69, 82]. More recently, Bartolomei *et al.* [4] followed a different path with Virtual Pattern Projection (VPP). Although LiDAR sensors and event cameras have been deployed together for some applications [6, 14, 20, 33, 55, 58, 61], this paper represents the first attempt at combining LiDAR with an event stereo framework. We argue that the two modalities are complementary, as shown in Fig. 2 – e.g., the lack of texture and motion makes an event camera uninformative, whereas this does not affect LiDAR systems.

3 Preliminaries: Event-based Deep Stereo

Event cameras measure brightness changes as an asynchronous stream of events. Accordingly, an event $e_k = (x_k, y_k, p_k, t_k)$ is triggered at time t_k if the intensity sensed by pixel (x_k, y_k) on the $W \times H$ sensor grid changes and surpasses a specific contrast threshold. Depending on the sign of this change, it will have polarity $p_k \in \{-1, 1\}$. Since this unstructured flow is not suitable for standard CNNs – as those proposed in the classical stereo literature [51] – converting it into $W \times H \times C$ structured representations is necessary if we are interested in obtaining a dense disparity map [21]. Purposely, given a timestamp t_d at which we want to estimate

a disparity map, events are sampled backward in time from the stream, either based on a time interval (SBT) or a maximum number of events (SBN), and *stacked* according to various strategies – among them:

Histogram [39]. Events of the two polarities are counted into per-pixel histograms, yielding a $W \times H \times 2$ stack.

Voxel grid [86]. The timelapse from which events are sampled is split into B uniform bins: polarities are accumulated in each bin of a $W \times H \times B$ stack.

Mixed-Density Event Stack (MDES) [43]. Similar to the voxel grid strategy, the timelapse is split into bins covering $1, \frac{1}{2}, \frac{1}{4}, \dots, \frac{1}{2^{N-2}}, \frac{1}{2^{N-1}}$ of the total interval. The latest event in each bin is kept, yielding a $W \times H \times N$ binary stack.

Concentrated stack [43]. A shallow CNN is trained to process a pre-computed stack (*e.g.*, an MDES) and aggregate it to a $W \times H \times 1$ data structure.

Time-Ordered Recent Event (TORE) [3]. It stores event timestamps into Q per-pixel queues for each polarity, yielding a $W \times H \times 2Q$ stack.

Time Surface [32]. A surface is derived from the timestamp distributions of the two polarities. S values are sampled for each, yielding a $W \times H \times 2S$ stack.

ERGO-12 [87]. An optimized representation of 12 channels, each built according to different strategies from the previous. It yields a $W \times H \times 12$ stack.

Tencode [26]. A color image representation in which R and B channels encode positive and negative polarities, with G encoding the timestamp relative to the total timelapse. It produces an RGB image, *i.e.* a $W \times H \times 3$ stack.

We can broadly classify stereo frameworks using these representations into three categories: i) *white boxes*, for which we have full access to the implementation of both the stereo backbone and the stacked event construction; ii) *gray boxes*, in case we do not have access to the stereo backbone; iii) *black boxes*, when the stacked event representation is not accessible neither.

4 Proposed Method

According to the sensor fusion literature for conventional cameras, the main strategies for combining stereo images with sparse depth measurements from active sensors consist of i) concatenating the two modalities and processing them as joint inputs with a stereo network [12, 46, 69, 81], ii) modulating the internal cost volume computed by the backbone itself [25, 49, 69, 82] or, more recently, iii) projecting distinctive patterns on images according to depth hints [4].

We follow the latter path, since it is more effective and flexible than the alternatives – which can indeed be applied to *white box* frameworks only. For this purpose, we design two alternative strategies suited even for gray and black box frameworks, respectively, as depicted in Fig. 3.

4.1 Virtual Stack Hallucination – VSH

Given left and right stacks $\mathcal{S}_L, \mathcal{S}_R$ of size $W \times H \times C$ and a set Z of depth measurements $z(x, y)$ by a sensor, we perform a Virtual Stack Hallucination (VSH), by augmenting each channel $c \in C$, to increase the distinctiveness of local patterns

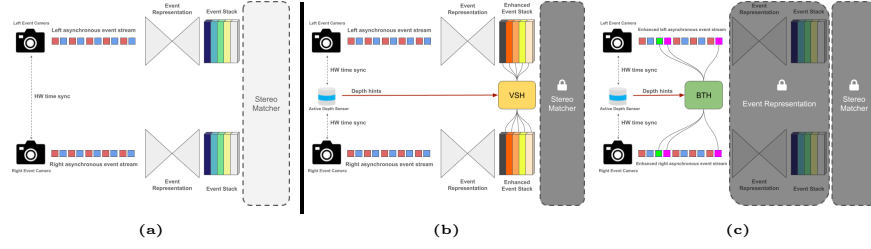


Fig. 3: Overview of a generic event-based stereo network and our hallucination strategies. State-of-the-art event-stereo frameworks (a) pre-process raw events to obtain event stacks fed to a deep network. In case the stacks are accessible, we define the model as a *gray box*, otherwise as a *black box*. In the former case (b), we can hallucinate patterns directly on it (VSH). When dealing with a black box (c), we can hallucinate raw events that will be processed to obtain the stacks (BTH).

and thus ease matching. This is carried out by injecting the same virtual stack $\mathcal{A}(x, y, x', c)$ into $\mathcal{S}_L, \mathcal{S}_R$ respectively at coordinates (x, y) and (x', y) .

$$\begin{aligned} \mathcal{S}_L(x, y, c) &\leftarrow \mathcal{A}(x, y, x', c) \\ \mathcal{S}_R(x', y, c) &\leftarrow \mathcal{A}(x, y, x', c) \end{aligned} \quad (1)$$

with x' obtained as $x - d(x, y)$, with disparity $d(x, y)$ triangulated back from depth $z(x, y)$ as $\frac{bf}{z(x, y)}$, according to the baseline and focal lengths b, f of the stereo system. We deploy a generalized version of the random pattern operator \mathcal{A} proposed in [4], agnostic to the stacked representation:

$$\mathcal{A}(x, y, x', c) \sim \mathcal{U}(\mathcal{S}^-, \mathcal{S}^+) \quad (2)$$

with \mathcal{S}^- and \mathcal{S}^+ the minimum and maximum values appearing across stacks $\mathcal{S}_L, \mathcal{S}_R$ and \mathcal{U} a uniform random distribution. Following [4], the pattern can either cover a single pixel or a local window. This strategy alone is sufficient already to ensure distinctiveness and to dramatically ease matching across stacks, even more than with color images [4], since acting on semi-dense structures – *i.e.*, stacks are uninformative in the absence of events. It also ensures a straightforward application of the same principles used on RGB images, *e.g.*, to combine the original content (color) with the virtual projection (pattern) employing alpha blending [4]. Nevertheless, we argue that acting at this level i) requires direct access to the stacks, *i.e.*, a gray-box deep event-stereo network, and ii) might be sub-optimal as stacks encode only part of the information from streams.

4.2 Back-in-Time Hallucination – BTH

A higher distinctiveness to ease correspondence can be induced by hallucinating patterns directly in the continuous events domain. Specifically, we act in the so-called *event history*: given a timestamp t_d at which we want to estimate

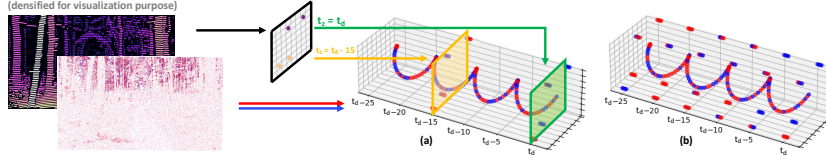


Fig. 4: Overview of Back-in-Time Hallucination (BTH). To estimate disparity at t_d , if LiDAR data is available – e.g., at timestamp $t_z = t_d$ (green) or $t_z = t_d - 15$ (yellow) – we can naïvely inject events of random polarities at the same timestamp t_z (a). More advanced injection strategies can be used – e.g. by hallucinating multiple events, starting from t_d , back-in-time at regular intervals (b).

disparity, raw events are sampled from the left and right streams starting from t_d and going backward, according to either SBN or SBT stacking approaches, to obtain a pair of event histories $\mathcal{E}_L = \{e_k^L\}_{k=1}^N$ and $\mathcal{E}_R = \{e_k^R\}_{k=1}^M$, where e_k^L, e_k^R are the k -th left and right events. Events in the history are sorted according to their timestamp – i.e., inequality $t_k \leq t_{k+1}$ holds for every two adjacent e_k, e_{k+1} .

At this point, we intervene to hallucinate novel events: given a depth measurement $z(\hat{x}, \hat{y})$, triangulated back into disparity $d(\hat{x}, \hat{y})$, we inject a pair of fictitious events $\hat{e}^L = (\hat{x}, \hat{y}, \hat{p}, \hat{t})$ and $\hat{e}^R = (\hat{x}', \hat{y}, \hat{p}, \hat{t})$ respectively inside \mathcal{E}_L and \mathcal{E}_R , producing $\hat{\mathcal{E}}_L = \{\hat{e}_1^L, \dots, \hat{e}_N^L\}$ and $\hat{\mathcal{E}}_R = \{\hat{e}_1^R, \dots, \hat{e}_M^R\}$. By construction, \hat{e}^L and \hat{e}^R adhere to i) the time ordering constraint, ii) the geometry constraint $\hat{x}' = \hat{x} - d(\hat{x}, \hat{y})$ and iii) a similarity constraint – i.e., \hat{p}, \hat{t} are the same for \hat{e}^L and \hat{e}^R . Fictitious polarity \hat{p} and fictitious timestamp \hat{t} are two degrees of freedom useful to ensure distinctiveness along the epipolar line and ease matching, according to which we can implement different strategies summarized in Fig. 4, and detailed in the remainder.

Single-timestamp injection. The simplest way to increase distinctiveness is to insert synchronized events at a fixed timestamp. Accordingly, for each depth measurement $d(\hat{x}, \hat{y})$, a total of $K_{\hat{x}, \hat{y}}$ pairs of fictitious events are inserted in $\mathcal{E}_L, \mathcal{E}_R$, having polarity \hat{p}_k randomly chosen from the discrete set $\{-1, 1\}$. Timestamp \hat{t} is fixed and can be, for instance, t_z at which the sensor infers depth, that can coincide with timestamp t_d at which we want to estimate disparity – e.g., $t_z = t_d = 0$ in the case depicted in Fig. 4 (a). Inspired by [4], events might be optionally hallucinated in patches rather than single pixels. However, as depth sensors usually work at a fixed acquisition frequency – e.g., 10Hz for LiDARs – sparse points might be unavailable at any specific timestamp. Nonetheless, since $\mathcal{E}_L, \mathcal{E}_R$ encode a time interval, we can hallucinate events even if derived from depth scans performed *in the past* – e.g., at $t_z < t_d$, – by placing them in the proper position inside $\mathcal{E}_L, \mathcal{E}_R$.

Repeated injection. The previous strategy does not exploit one of the main advantages of events over color images, i.e. the temporal dimension, at its best. Purposely, we design a more advanced hallucination strategy based on *repeated* naïve injections performed along the time interval sampled by $\mathcal{E}_L, \mathcal{E}_R$. As long as we are interested in recovering depth at t_d only, we can hallucinate as many

events as we want in the time interval *before* t – *i.e.*, for $t_z = t_d = 0$, over the entire interval as shown in Fig. 4 (b) – consistent with the depth measurements at t_d itself, which will increase the distinctiveness in the event histories and will ease the match by hinting the correct disparity. Inspired by the stacked representations introduced in Sec. 3, we can design a strategy for injecting multiple events along the stream. Accordingly, we define the *conservative* time range $[t^-, t^+]$ of the events histories $\mathcal{E}_L, \mathcal{E}_R$, with $t^- = \min \{t_0^L, t_0^R\}$ and $t^+ = \max \{t_N^L, t_M^R\}$ and divide it into B equal temporal bins. Then, inspired by MDES [43], we run B single-timestamp injections at $\hat{t}_b = \frac{2^b - 1}{2^b} (t^+ - t^-) + t^-$, with $b \in \{1, \dots, B\}$. Additionally, each depth measurement is used only once – *i.e.*, the number of fictitious events $K_{b,\hat{x},\hat{y}}$ in the b -th injection is set as $K_{b,\hat{x},\hat{y}} \leftarrow K_{\hat{x},\hat{y}} \delta(b, D_{\hat{x},\hat{y}})$ where $\delta(\cdot, \cdot)$ is the Kronecker delta and $D_{\hat{x},\hat{y}} \leftarrow \text{round}(X^U(B - 1) + 1)$ is a random slot assignment. We will show in our experiment how this simple strategy can improve the results of BTH, in particular increasing its robustness against misaligned LiDAR data – *i.e.*, measurements retrieved at a timestamp $t_z < t_d$.

5 Experiments

5.1 Implementation and Experimental Settings

We implement VSH and BTH in Python, using the Numba package.

General framework. We build our code base starting from SE-CFF [43] – state-of-the-art for event-based stereo – assuming the same stereo backbone as in their experiments, *i.e.* derived from AANet [72], and run SBN to generate the event history to be stacked. While we select a single architecture, we implement a variety of stacked representations: purposely, we implement a single instance of the stereo backbone for any stacked representations introduced in Sec. 3, taking the opportunity to evaluate their performance with the event stereo task. For Concentration representation, we use MDES as the prior stacking function following [43] and avoid considering future events during training. Furthermore, in this case, VSH is applied before the concentration network since it would interfere with gradient back-propagation during training – while this cannot occur with BTH. From [4], we adapt occlusion handling and hallucination on uniform/not uniform patches. We also implement alpha-blending, for VSH only – as it loses its purpose when acting on the raw streams. For all our methods, we inherit the same hyper-parameters from [4]; yet, we discard occluded points as the occlusion handling strategy for BTH since an equivalent strategy to deal with sparse event histories is not trivial. For VSH on Voxel Grids, we use the 5-th and 95-th percentile to calculate \mathcal{S}^- and \mathcal{S}^+ due to the frequent presence of extreme values in the stack. For BTH, we perform 12 injections (*i.e.*, $B = 12$).

Existing fusion methodologies. We compare our proposal with existing methods from the RGB stereo literature, consisting of i) modulating the cost volume built by the backbone – Guided Stereo Matching [49], ii) concatenating the sparse depth values to the inputs to the stereo network – *e.g.*, as done by LidarStereoNet [12], iii) a combination of both the previous strategies – in

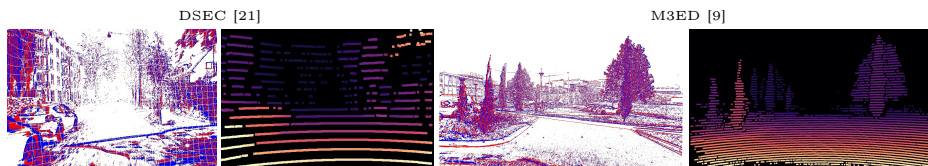


Fig. 5: Qualitative comparison – DSEC vs M3ED. DSEC features 640×480 event cameras and a 16-line LiDAR, M3ED has 1280×720 event cameras and a 64-line LiDAR. LiDAR scans have been dilated with a 7×7 kernel to ease visualization.

analogy to CCVNorm [69]. Any strategy is adapted to the same common stereo backbone [43] (see **supplementary material**). Running BTH and VSH adds respectively 10ms and 2-15ms (depending on representations) on the CPU.

Training protocol. Any model we train – either the original event stereo backbones or those implementing fusion strategies – runs for 25 epochs with a batch size of 4 and a maximum disparity set to 192. We use Adam [29] with beta (0.9, 0.999) and weight decay set to 10^{-4} . The learning rate starts at $5 \cdot 10^{-4}$ and decays with cosine annealing. We apply random crops and vertical flips to augment data during training.

5.2 Evaluation Datasets & Protocol

We introduce datasets and metrics used in our experiments.

DSEC [21]. An outdoor event stereo dataset, captured using wide-baseline (50 cm) stereo event cameras at 640×480 resolution. Ground-truth disparity is obtained by accumulating 16-line LiDAR scans, for a total of 26 384 maps organized into 41 sequences. We split them into train/test sets following [43]. From the training set, we retain a further *search* split for hyper-parameters tuning and ablation experiments. Sparse LiDAR measurements are obtained by aligning the raw scans with the ground-truth – both provided by the authors – by running a LiDAR inertial odometry pipeline followed by ICP registration (see the **supplementary material** for details).

M3ED [9]. This dataset provides 57 indoor/outdoor scenes collected with a compact multi-sensor block mounted on three different vehicles – *i.e.*, a car, a UAV, and a quadruped robot. A 64-line LiDAR generates semi-dense ground-truth depth, while the event stereo camera has a shorter baseline (12 cm) and a higher resolution (1280×720). We use 5 sequences from this dataset for evaluation purposes only – some of which contain several frames acquired with the cameras being static – to evaluate the generalization capacity of the models both to different domains and the density of the LiDAR sensor. Similarly to DSEC, we derived sparse LiDAR depth maps from the raw scans. Thanks to the SDK made available by the authors, we could derive LiDAR measurements aligned to any desired temporal offset according to linear interpolation of the ground-truth poses (see the **supplementary material** for details). This allows us to run ded-

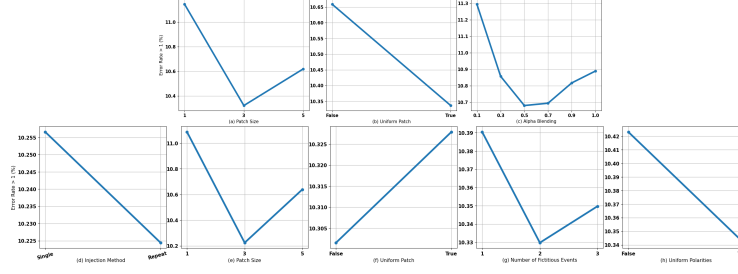


Fig. 6: Hyperparameters search. Results on DSEC search split. On top, we study the impact of (a) patch size, (b) uniform patches, and (c) alpha blending on VSH. At the bottom, we consider (d) single vs repeated injection, (e) patch size, (f) uniform patches, (g) number of fictitious events, and (h) uniform polarities on BTH.

icated experiments to assess the effect of time-misaligned depth measurements. Fig. 5 shows a qualitative comparison between the two datasets.

Evaluation Metrics. We compute the percentage of pixels with an error greater than 1 or 2 pixels (1PE, 2PE), and the mean absolute error (MAE). We highlight the **best** and **second best** methods per row on each metric.

5.3 Ablation Study

We ran hyper-parameters search and ablation experiments for VSH and BTH on the DSEC search split, reporting the 1PE error. We conducted these experiments using any representation listed in Sec. 3 – except for Concentration [43], which starts from pre-computed MDES stacks – and report the average results.

VSH. Fig. 6 (top) shows the impact of different hyper-parameters on VSH strategy. In (a), we can observe how VSH is improved by using 3×3 patches, while 5×5 cannot yield further benefits. Consequently, we select it as the default configuration from now on. In (b), we show that uniform patterns are more effective than random ones, and in (c) alpha equal to 0.5 works the best.

BTH. Fig. 6 (bottom) focuses on our second strategy. In (d) we show how repeated injection can improve the results; thus, we select it as the default configuration from now on. In the remainder, we will better appreciate how this setting is much more robust when dealing with misaligned LiDAR data. Next, (e) outlines how hallucinating events with 3×3 patches lead to the best results. Applying a uniform patch of events following (f) yields, again, better results. In (g), we tested different numbers $K_{\hat{x},\hat{y}}$ of injected fictitious events. Injecting more than one event is beneficial, yet saturating with two. Finally, (h) shows that using uniform polarities yields lower errors.

5.4 Experiments on DSEC

We now report experiments on the DSEC testing split, either when applying fusion strategies to pre-trained stereo models without retraining them or when training the networks from scratch to exploit LiDAR data.

Table 1: Results on DSEC [21] – pre-trained. We test the different stacked representations (rows) with several fusion strategies applied to pre-trained stereo backbones.

Stacked representation	Baseline			Guided [49]			VSH (ours)			BTH (ours)		
	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE
(A) Histogram [39]	16.21	4.73	0.74	16.07	4.68	0.73	13.71	4.20	0.69	13.32	3.92	0.66
(B) MDES [43]	15.32	4.40	0.70	15.13	4.34	0.70	12.94	3.52	0.63	12.61	3.50	0.62
(C) Concentration [43]	15.97	4.33	0.70	15.79	4.27	0.70	13.70	3.60	0.65	14.66	3.77	0.66
(D) Voxelgrid [86]	16.49	4.56	0.72	16.29	4.50	0.71	13.12	3.69	0.65	12.44	3.60	0.62
(E) TORE [3]	15.91	4.57	0.71	15.72	4.50	0.71	12.53	3.65	0.63	12.27	3.68	0.62
(F) Time Surface [32]	15.33	4.29	0.70	15.18	4.24	0.69	12.16	3.38	0.62	12.28	3.45	0.62
(G) ERGO-12 [87]	15.02	4.20	0.68	14.87	4.14	0.68	12.02	3.40	0.61	11.98	3.42	0.61
(H) Tencode [26]	14.46	4.17	0.68	14.29	4.11	0.67	12.12	3.37	0.61	11.86	3.45	0.61
Avg. Rank.	-	-	-	3.00	3.00	3.00	1.75	1.38	1.50	1.25	1.63	1.13

Table 2: Results on DSEC [21] – retrained. We test different stacked representations (rows) with several fusion strategies applied during training.

	Concat [12]			Guided+Concat [69]			Guided [49]			VSH (ours)			BTH (ours)		
	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE
(A)	12.57	3.37	0.62	12.81	3.41	0.63	15.57	4.58	0.72	9.90	3.26	0.53	10.91	3.41	0.59
(B)	12.37	3.17	0.61	12.40	3.25	0.61	14.66	4.36	0.70	9.31	3.01	0.51	9.62	3.01	0.54
(C)	12.38	3.41	0.63	12.74	3.44	0.66	15.15	4.45	0.71	9.70	3.04	0.53	9.66	2.98	0.55
(D)	12.23	3.18	0.60	11.90	3.10	0.60	14.52	4.21	0.68	10.16	3.20	0.56	9.68	2.90	0.54
(E)	12.99	3.33	0.62	12.62	3.25	0.61	16.00	4.56	0.73	9.91	3.05	0.53	9.83	2.98	0.54
(F)	12.18	3.09	0.61	12.47	3.17	0.61	14.40	4.21	0.68	9.47	2.90	0.52	9.58	2.92	0.54
(G)	12.43	3.14	0.61	12.82	3.19	0.62	13.85	3.97	0.66	9.25	2.88	0.50	9.37	2.87	0.54
(H)	11.95	3.08	0.60	11.75	3.10	0.60	14.72	4.21	0.69	9.39	3.00	0.52	9.59	2.97	0.55
	3.38	3.00	3.13	3.63	3.50	3.38	5.00	5.00	5.00	1.38	1.88	1.13	1.63	1.38	1.88

Pre-trained models. Tab. 1 reports, on each row, the results yielded by using a specific stacked representation. In the columns, we report the different fusion strategies involved in our experiments, starting with the baseline – *i.e.*, a stereo backbone processing events only. In the last row, we report the average ranking – for the three metrics – achieved by any fusion strategy over the eight representations. Starting from baseline models, we can notice how the different representations have an impact on the accuracy of the stereo backbone, with those modeling complex behaviors – *e.g.*, Time Surface [32] or ERGO-12 [87] – yielding up to 2% lower 1PE than simpler ones such as Histogram [39]. The Guided framework [49] can improve the results only moderately: this is caused by the very sparse measurements retrieved from the 16-line LiDAR sensor used in DSEC, as well as by the limited effect of the cost volume modulation in regions where events are not available for matching. Nonetheless, VSH and BTH consistently outperform Guided, always improving the baseline by 2-3% points on 1PE. In general, BTH achieves the best 1PE and MAE metrics in most cases; this strategy is the best when re-training the stereo backbone is not feasible.

Training from scratch. Tab. 2 reports the results obtained by training the stereo backbones from scratch to perform LiDAR-event stereo fusion. This allows either the deployment of strategies that process the LiDAR data directly as input [12, 69] or those not requiring it, *i.e.*, [49] and ours. Specifically, Concat [12] and Guided+Concat [69] strategy achieve results comparable to those by VSH and BTH observed before, thus outperforming Guided [49] which, on the contrary, cannot benefit much from the training process. When deploying our solutions during training, their effectiveness dramatically increases, often dropping 1PE

Table 3: Results on M3ED [9] – pre-trained. We test the different stacked representations (rows) with several fusion strategies applied to pre-trained stereo backbones.

Stacked representation	Baseline			Guided [49]			VSH (ours)			BTH (ours)		
	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE
(A) Histogram [39]	37.70	19.49	1.76	37.18	19.29	1.75	20.19	11.19	1.19	22.32	12.37	1.27
(B) MDES [43]	43.17	19.50	1.85	42.27	19.16	1.83	29.42	14.80	1.52	22.58	12.20	1.30
(C) Concentration [43]	45.78	20.84	1.82	45.06	20.57	1.80	33.63	16.19	1.53	25.22	12.68	1.28
(D) Voxelgrid [86]	37.33	17.66	1.70	36.64	17.38	1.68	20.40	11.41	1.22	20.94	11.72	1.23
(E) TORE [3]	41.70	19.09	1.81	41.00	18.78	1.80	28.25	14.01	1.47	21.91	12.34	1.30
(F) Time Surface [32]	38.58	18.52	1.72	37.91	18.23	1.70	24.89	13.34	1.37	22.60	12.77	1.31
(G) ERGO-12 [87]	36.33	17.81	1.66	35.61	17.50	1.64	22.53	12.33	1.26	20.41	11.69	1.21
(H) Tencode [26]	43.56	20.07	1.82	42.66	19.76	1.80	28.24	14.46	1.43	22.61	12.75	1.26
Avg. Rank.	-			3.00	3.00	3.00	1.75	1.75	1.75	1.25	1.25	1.25

Table 4: Results on M3ED [9] – retrained. We test different stacked representations (rows) with several fusion strategies applied during training.

	Concat [12]			Guided+Concat [69]			Guided [49]			VSH (ours)			BTH (ours)		
	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE	1PE	2PE	MAE
(A)	34.67	15.21	1.92	38.65	17.00	1.94	37.45	18.98	1.76	19.34	12.93	1.46	19.83	13.20	1.39
(B)	37.72	16.91	1.85	37.32	17.16	2.14	37.00	18.66	1.76	19.24	13.17	1.44	18.70	11.79	1.24
(C)	39.88	19.01	2.33	38.45	17.76	2.47	38.14	19.62	1.80	19.84	13.68	1.90	19.46	12.29	1.35
(D)	33.89	16.21	1.89	33.54	15.85	1.75	37.85	18.81	1.74	18.56	11.76	1.32	21.02	14.30	1.80
(E)	38.83	18.38	2.27	35.80	16.63	2.05	40.51	19.96	1.95	20.03	13.97	1.86	20.04	12.65	1.39
(F)	40.26	18.44	2.19	35.48	17.74	2.15	38.77	18.41	1.75	19.61	13.01	1.55	21.91	14.33	1.72
(G)	42.43	19.31	2.31	42.24	18.42	2.34	37.95	17.83	1.76	18.45	12.31	1.55	19.12	11.60	1.24
(H)	37.46	17.87	2.15	33.69	16.47	1.95	39.78	19.42	1.82	19.49	12.21	1.38	19.28	11.68	1.33
	4.38	4.00	4.50	3.63	3.38	4.38	4.00	4.63	2.75	1.38	1.63	1.88	1.63	1.38	1.50

error below 10%. VSH often yields the best 1PE and MAE overall, nominating it as the most effective – yet intrusive – among our solutions.

5.5 Experiments on M3ED

We test the effectiveness of BTH and alternative approaches on M3ED, using the backbones trained on DSEC **without** any fine-tuning on M3ED itself.

Pre-trained models. Tab. 3 collects the outcome of this experiment by applying Guided, VSH, and BTH to pre-trained models. Looking at the baselines, we can appreciate how M3ED is very challenging for models trained in a different domain, with 1PE errors higher than 30%. This is caused by both the domain shift and the higher resolution of the event cameras used. Even so, complex event representations – *e.g.*, ERGO-12 [87] – can better generalize. Guided confirms its limited impact, this time mainly because of the ineffectiveness of the cost volume modulation in the absence of any information from the events domain. On the contrary, we can appreciate even further the impact of VSH and BTH, almost halving the 1PE error. Specifically, BTH is the absolute winner with 6 out of 8 representations, and the best choice for pre-trained frameworks.

Training from scratch. Tab. 4 resumes the results obtained when training on DSEC the backbones implementing LiDAR-event stereo fusion strategies. The very different distribution of depth points observed across the two datasets – sourced respectively from 16 and 64-line LiDARs – yields mixed results for existing methods [12, 49, 69], with rare cases for which they fail to improve the baseline model (*e.g.*, Concat and Guided with Time Surface and ERGO-

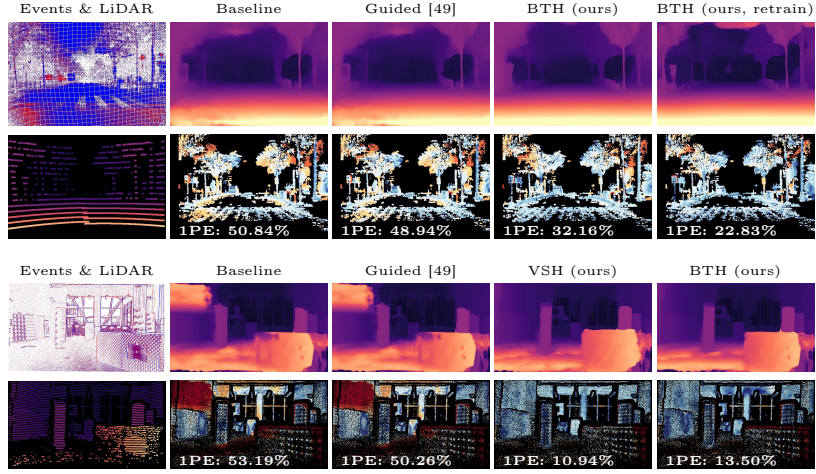


Fig. 7: Qualitative results. Results on DSEC *zurich_10_b* with Voxelgrid [86] (top) and M3ED *spot_indoor_obstacles* with Histogram [39] (bottom).

12, Guided+Concat with Histogram). On the contrary, backbones trained with VSH and BTH consistently improve over the baseline, often with larger gains compared to their use with pre-trained models. Overall, BTH is the best on 2PE and MAE, confirming it is better suited for robustness across domains and different LiDAR sensors.

Fig. 7 shows qualitative results. On DSEC (top), BTH dramatically improves results over the baseline and Guided, yet cannot fully recover some details in the scene except when retraining the stereo backbone. On M3ED (bottom), both VSH and BTH with pre-trained models reduce the error by $5\times$.

5.6 Experiments on M3ED – Time-misaligned LiDAR

We conclude by assessing the robustness of the considered strategies against the use of LiDAR not synchronized with the timestamp at which we wish to estimate disparity – occurring if we wish to maintain the microsecond resolution of the event cameras. Purposely, we extract raw LiDAR measurements collected 3, 13, 32, 61, and 100 ms in the past with the M3ED SDK.

Fig. 8 shows the trend of the 1PE metric achieved by Guided (red), VSH (yellow) and BTH (black and green) on pre-trained backbones. Not surprisingly, the error rates arise at the increase of the temporal distance: while this is less evident with Guided because of its limited impact, this becomes clear with VSH and BTH. Nonetheless, both can always retain a significant gain over the baseline model (blue) – i.e., the stereo backbone processing events only – even with the farthest possible misalignment with a 10Hz LiDAR (100ms). We can appreciate how BTH is often better than VSH (coherently with Tab. 3), yet only when repeated injections are performed (green). Indeed, using a single injection (black)

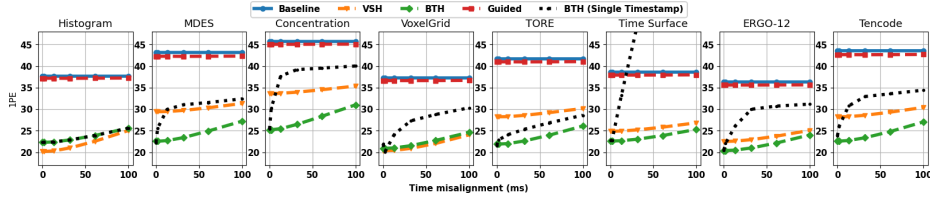


Fig. 8: Experiments with time-misaligned LiDAR on M3ED [9] – pre-trained. We measure the robustness of different fusion strategies against the use of out-of-sync LiDAR data, without retraining the stereo backbone.

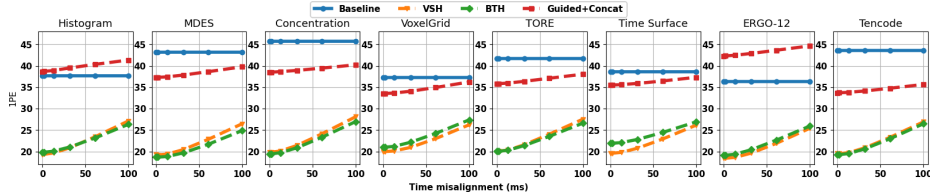


Fig. 9: Experiments with time-misaligned LiDAR on M3ED [9] – retrained. We measure the robustness of different fusion strategies against the use of out-of-sync LiDAR data when training the stereo backbone from scratch.

rapidly leads BTH to an accuracy drop when increasing the misalignment, except when using Histogram representation. Overall, BTH with ERGO-12 is the most robust solution. Fig. 9 shows the results achieved by VSH (yellow) and BTH (green) after retraining, against the best competitor according to average ranks in Tab. 4 – i.e., Guided+Concat (red). The impact of this latter is limited and sometimes fails to improve the baseline (see Histogram and ERGO-12). On the contrary, our solutions confirm their robustness and effectiveness even when dealing with time-misaligned LiDAR data.

6 Conclusion

This paper proposes a novel framework for implementing event stereo and LiDAR fusion. It works by hallucinating fictitious events either in the stacked representation processed by stereo backbones or the continuous streams sensed by event cameras, easing the matching process to the downstream stereo model estimating disparity. Our exhaustive experiments prove that our solutions, VSH and BTH, dramatically outperform alternative fusion strategies from the RGB stereo literature, retaining the microsecond resolution typical of event cameras despite the discrete frame rate of LiDARs and depth sensors in general.

Limitations. Despite the robustness shown with misaligned LiDAR data, a marginal drop in accuracy compared to the case of having LiDAR measurements at the very same timestamp at which we aim to infer disparity maps occurs. Future work will focus on studying new design mechanisms to deal with it.

Acknowledgement. This study was carried out within the MOST – Sustainable Mobility National Research Center and received funding from the European Union Next-GenerationEU – PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.4 – D.D. 1033 17/06/2022, CN00000023. This manuscript reflects only the authors’ views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support.

References

1. Andreopoulos, A., Kashyap, H.J., Nayak, T.K., Amir, A., Flickner, M.D.: A low power, high throughput, fully event-based stereo system. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7532–7542 (2018)
2. Badino, H., Huber, D.F., Kanade, T.: Integrating lidar into stereo for fast and improved disparity computation. 2011 International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission pp. 405–412 (2011)
3. Baldwin, R.W., Liu, R., Almatrafi, M., Asari, V., Hirakawa, K.: Time-ordered recent event (tore) volumes for event cameras. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(2), 2519–2532 (2022)
4. Bartolomei, L., Poggi, M., Tosi, F., Conti, A., Mattoccia, S.: Active stereo without pattern projector. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 18470–18482 (October 2023)
5. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE Transactions on pattern analysis and machine intelligence **23**(11), 1222–1239 (2001)
6. Brebion, V., Moreau, J., Davoine, F.: Learning to estimate two dense depths from lidar and event data. In: Scandinavian Conference on Image Analysis. pp. 517–533. Springer (2023)
7. Camuñas-Mesa, L.A., Serrano-Gotarredona, T., Ieng, S.H., Benosman, R.B., Linares-Barranco, B.: On the use of orientation filters for 3d reconstruction in event-driven stereo vision. Frontiers in neuroscience **8**, 48 (2014)
8. Carneiro, J., Ieng, S.H., Posch, C., Benosman, R.: Event-based 3d reconstruction from neuromorphic retinas. Neural Networks **45**, 27–38 (2013)
9. Chaney, K., Cladera, F., Wang, Z., Bisulco, A., Hsieh, M.A., Korpela, C., Kumar, V., Taylor, C.J., Daniilidis, K.: M3ed: Multi-robot, multi-sensor, multi-environment event dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 4015–4022 (June 2023)
10. Chang, J.R., Chen, Y.S.: Pyramid stereo matching network. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5410–5418 (2018)
11. Cheng, X., Wang, P., Yang, R.: Learning depth with convolutional spatial propagation network. IEEE transactions on pattern analysis and machine intelligence **42**(10), 2361–2379 (2019)
12. Cheng, X., Zhong, Y., Dai, Y., Ji, P., Li, H.: Noise-aware unsupervised deep lidar-stereo fusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6339–6348 (2019)

13. Cheng, X., Zhong, Y., Harandi, M., Dai, Y., Chang, X., Li, H., Drummond, T., Ge, Z.: Hierarchical neural architecture search for deep stereo matching. *Advances in Neural Information Processing Systems* **33** (2020)
14. Cui, M., Zhu, Y., Liu, Y., Liu, Y., Chen, G., Huang, K.: Dense depth-map estimation based on fusion of event camera and sparse lidar. *IEEE Transactions on Instrumentation and Measurement* **71**, 1–11 (2022). <https://doi.org/10.1109/TIM.2022.3144229>
15. Dikov, G., Firouzi, M., Röhrbein, F., Conradt, J., Richter, C.: Spiking cooperative stereo-matching at 2 ms latency with neuromorphic hardware. In: *Biomimetic and Biohybrid Systems: 6th International Conference, Living Machines 2017, Stanford, CA, USA, July 26–28, 2017, Proceedings 6*. pp. 119–137. Springer (2017)
16. Duggal, S., Wang, S., Ma, W.C., Hu, R., Urtasun, R.: Deeppruner: Learning efficient stereo matching via differentiable patchmatch. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4384–4393 (2019)
17. Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A.J., Conradt, J., Daniilidis, K., et al.: Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence* **44**(1), 154–180 (2020)
18. Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A.J., Conradt, J., Daniilidis, K., Scaramuzza, D.: Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(1), 154–180 (2022). <https://doi.org/10.1109/TPAMI.2020.3008413>
19. Gandhi, V., Čech, J., Horaud, R.: High-resolution depth maps based on tof-stereo fusion. In: *2012 IEEE International Conference on Robotics and Automation*. pp. 4742–4749. IEEE (2012)
20. Gao, L., Liang, Y., Yang, J., Wu, S., Wang, C., Chen, J., Kneip, L.: Vector: A versatile event-centric benchmark for multi-sensor slam. *IEEE Robotics and Automation Letters* **7**(3), 8217–8224 (2022)
21. Gehrig, M., Aarents, W., Gehrig, D., Scaramuzza, D.: Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters* (2021). <https://doi.org/10.1109/LRA.2021.3068942>
22. Guo, W., Li, Z., Yang, Y., Wang, Z., Taylor, R.H., Unberath, M., Yuille, A., Li, Y.: Context-enhanced stereo transformer. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2022)
23. Guo, X., Yang, K., Yang, W., Wang, X., Li, H.: Group-wise correlation stereo network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3273–3282 (2019)
24. Hirschmüller, H.: Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on pattern analysis and machine intelligence* **30**(2), 328–341 (2007)
25. Huang, Y.K., Liu, Y.C., Wu, T.H., Su, H.T., Chang, Y.C., Tsou, T.L., Wang, Y.A., Hsu, W.H.: S3: Learnable sparse signal superdensity for guided depth estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16706–16716 (2021)
26. Huang, Z., Sun, L., Zhao, C., Li, S., Su, S.: Eventpoint: Self-supervised interest point detection and description for event-based camera. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 5396–5405 (January 2023)
27. Kendall, A., Martirosyan, H., Dasgupta, S., Henry, P., Kennedy, R., Bachrach, A., Bry, A.: End-to-end learning of geometry and context for deep stereo regression. In: *The IEEE International Conference on Computer Vision (ICCV)* (Oct 2017)

28. Khamis, S., Fanello, S., Rhemann, C., Kowdle, A., Valentin, J., Izadi, S.: Stereonet: Guided hierarchical refinement for real-time edge-aware depth prediction. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 573–590 (2018)
29. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
30. Kogler, J., Sulzbachner, C., Humenberger, M., Eibensteiner, F.: Address-event based stereo vision with bio-inspired silicon retina imagers. Advances in theory and applications of stereo vision pp. 165–188 (2011)
31. Kolmogorov, V., Zabini, R.: What energy functions can be minimized via graph cuts? IEEE transactions on pattern analysis and machine intelligence **26**(2), 147–159 (2004)
32. Lagorce, X., Orchard, G., Galluppi, F., Shi, B.E., Benosman, R.B.: Hots: a hierarchy of event-based time-surfaces for pattern recognition. IEEE transactions on pattern analysis and machine intelligence **39**(7), 1346–1359 (2016)
33. Li, B., Meng, H., Zhu, Y., Song, R., Cui, M., Chen, G., Huang, K.: Enhancing 3-d lidar point clouds with event-based camera. IEEE Transactions on Instrumentation and Measurement **70**, 1–12 (2021)
34. Li, J., Wang, P., Xiong, P., Cai, T., Yan, Z., Yang, L., Liu, J., Fan, H., Liu, S.: Practical stereo matching via cascaded recurrent network with adaptive correlation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16263–16272 (2022)
35. Li, Z., Liu, X., Drenkow, N., Ding, A., Creighton, F.X., Taylor, R.H., Unberath, M.: Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6197–6206 (2021)
36. Liang, C.K., Cheng, C.C., Lai, Y.C., Chen, L.G., Chen, H.H.: Hardware-efficient belief propagation. IEEE Transactions on Circuits and Systems for Video Technology **21**(5), 525–537 (2011)
37. Liang, Z., Feng, Y., Guo, Y., Liu, H., Chen, W., Qiao, L., Zhou, L., Zhang, J.: Learning for disparity estimation through feature constancy. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
38. Lipson, L., Teed, Z., Deng, J.: Raft-stereo: Multilevel recurrent field transforms for stereo matching. In: International Conference on 3D Vision (3DV) (2021)
39. Maqueda, A.I., Loquercio, A., Gallego, G., García, N., Scaramuzza, D.: Event-based vision meets deep learning on steering prediction for self-driving cars. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5419–5427 (2018)
40. Marin, G., Zanuttigh, P., Mattoccia, S.: Reliable fusion of tof and stereo depth driven by confidence measures. In: European Conference on Computer Vision (ECCV). pp. 386–401 (2016)
41. Marr, D.C., Poggio, T.A.: Cooperative computation of stereo disparity. Science **194** 4262, 283–7 (1976)
42. Mayer, N., Ilg, E., Hausser, P., Fischer, P., Cremers, D., Dosovitskiy, A., Brox, T.: A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
43. Nam, Y., Mostafavi, M., Yoon, K.J., Choi, J.: Stereo depth from events cameras: Concentrate and focus on the future. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6114–6123 (2022)

44. Osswald, M., Ieng, S.H., Benosman, R., Indiveri, G.: A spiking neural network model of 3d perception for event-based neuromorphic stereo vision systems. *Scientific reports* **7**(1), 40703 (2017)
45. Pang, J., Sun, W., Ren, J.S., Yang, C., Yan, Q.: Cascade residual learning: A two-stage convolutional neural network for stereo matching. In: *The IEEE International Conference on Computer Vision (ICCV)* (Oct 2017)
46. Park, K., Kim, S., Sohn, K.: High-precision depth estimation with the 3d lidar and stereo fusion. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 2156–2163. IEEE (2018)
47. Piatkowska, E., Belbachir, A., Gelautz, M.: Asynchronous stereo vision for event-driven dynamic stereo sensor using an adaptive cooperative approach. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*. pp. 45–50 (2013)
48. Poggi, M., Agresti, G., Tosi, F., Zanuttigh, P., Mattoccia, S.: Confidence estimation for tof and stereo sensors and its application to depth data fusion. *IEEE Sensors Journal* **20**(3), 1411–1421 (2020). <https://doi.org/10.1109/JSEN.2019.2946591>
49. Poggi, M., Pallotti, D., Tosi, F., Mattoccia, S.: Guided stereo matching. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 979–988 (2019)
50. Poggi, M., Tosi, F.: Federated online adaptation for deep stereo. In: *CVPR* (2024)
51. Poggi, M., Tosi, F., Batsos, K., Mordohai, P., Mattoccia, S.: On the synergies between machine learning and binocular stereo for depth estimation from images: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(9), 5314–5334 (2022)
52. Rogister, P., Benosman, R., Ieng, S.H., Lichtsteiner, P., Delbruck, T.: Asynchronous event-based binocular stereo matching. *IEEE Transactions on Neural Networks and Learning Systems* **23**(2), 347–353 (2011)
53. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)
54. Saikia, T., Marrakchi, Y., Zela, A., Hutter, F., Brox, T.: Autodispnet: Improving disparity estimation with automl. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1812–1823 (2019)
55. Saucedo, M.A., Patel, A., Sawlekar, R., Saradagi, A., Kanellakis, C., Agha-Mohammadi, A.A., Nikolakopoulos, G.: Event camera and lidar based human tracking for adverse lighting conditions in subterranean environments. *IFAC-PapersOnLine* **56**(2), 9257–9262 (2023)
56. Schraml, S., Belbachir, A.N., Milosevic, N., Schön, P.: Dynamic stereo vision system for real-time tracking. In: *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. pp. 1409–1412. IEEE (2010)
57. Shen, Z., Dai, Y., Rao, Z.: Cfnet: Cascade and fused cost volume for robust stereo matching. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 13906–13915 (June 2021)
58. Song, R., Jiang, Z., Li, Y., Shan, Y., Huang, K.: Calibration of event-based camera and 3d lidar. In: *2018 WRC Symposium on Advanced Robotics and Automation (WRC SARA)*. pp. 289–295. IEEE (2018)
59. Song, X., Zhao, X., Hu, H., Fang, L.: Edgestereo: A context integrated residual pyramid network for stereo matching. In: *ACCV* (2018)

60. Sulzbachner, C., Zinner, C., Kogler, J.: An optimized silicon retina stereo matching algorithm using time-space correlation. In: CVPR 2011 WORKSHOPS. pp. 1–7. IEEE (2011)
61. Ta, K., Bruggemann, D., Brödermann, T., Sakaridis, C., Van Gool, L.: L2e: Lasers to events for 6-dof extrinsic calibration of lidars and event cameras. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 11425–11431. IEEE (2023)
62. Tanai, T., Matsushita, Y., Naemura, T.: Graph cut based continuous stereo matching using locally shared labels. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1613–1620 (2014)
63. Tankovich, V., Hane, C., Zhang, Y., Kowdle, A., Fanello, S., Bouaziz, S.: Hit-net: Hierarchical iterative tile refinement network for real-time stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14362–14372 (June 2021)
64. Tonioni, A., Tosi, F., Poggi, M., Mattoccia, S., Stefano, L.D.: Real-time self-adaptive deep stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
65. Tosi, F., Tonioni, A., De Gregorio, D., Poggi, M.: Nerf-supervised deep stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 855–866 (June 2023)
66. Tulyakov, S., Fleuret, F., Kiefel, M., Gehler, P., Hirsch, M.: Learning an event sequence embedding for dense event-based deep stereo. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1527–1537 (2019)
67. Uddin, S.N., Ahmed, S.H., Jung, Y.J.: Unsupervised deep event stereo for depth estimation. IEEE Transactions on Circuits and Systems for Video Technology **32**(11), 7489–7504 (2022)
68. Veksler, O.: Stereo correspondence by dynamic programming on a tree. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05). vol. 2, pp. 384–390. IEEE (2005)
69. Wang, T.H., Hu, H.N., Lin, C.H., Tsai, Y.H., Chiu, W.C., Sun, M.: 3d lidar and stereo fusion using stereo matching network with conditional cost volume normalization. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 5895–5902. IEEE (2019)
70. Wang, Y., Lai, Z., Huang, G., Wang, B.H., Van Der Maaten, L., Campbell, M., Weinberger, K.Q.: Anytime stereo image depth estimation on mobile devices. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 5893–5900 (2019)
71. Xu, G., Wang, X., Ding, X., Yang, X.: Iterative geometry encoding volume for stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21919–21928 (2023)
72. Xu, H., Zhang, J.: Aanet: Adaptive aggregation network for efficient stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1959–1968 (2020)
73. Yang, G., Manela, J., Happold, M., Ramanan, D.: Hierarchical deep stereo matching on high-resolution images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5515–5524 (2019)
74. Yang, G., Zhao, H., Shi, J., Deng, Z., Jia, J.: Segstereo: Exploiting semantic information for disparity estimation. In: ECCV. pp. 636–651 (2018)
75. Yang, Q., Wang, L., Ahuja, N.: A constant-space belief propagation algorithm for stereo matching. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 1458–1465. IEEE (2010)

76. Yang, Q., Wang, L., Yang, R., Stewénus, H., Nistér, D.: Stereo matching with color-weighted correlation, hierarchical belief propagation, and occlusion handling. *IEEE transactions on pattern analysis and machine intelligence* **31**(3), 492–504 (2008)
77. Yin, Z., Darrell, T., Yu, F.: Hierarchical discrete distribution decomposition for match density estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6044–6053 (2019)
78. Zabih, R., Woodfill, J.: Non-parametric local transforms for computing visual correspondence. In: *Third European Conference on Computer Vision (Vol. II)*. pp. 151–158. 3rd European Conference on Computer Vision (ECCV), Springer-Verlag New York, Inc., Secaucus, NJ, USA (1994)
79. Zbontar, J., LeCun, Y., et al.: Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.* **17**(1), 2287–2318 (2016)
80. Zhang, F., Prisacariu, V., Yang, R., Torr, P.H.: GA-Net: Guided aggregation net for end-to-end stereo matching. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
81. Zhang, J., Ramanagopal, M.S., Vasudevan, R., Johnson-Roberson, M.: Listereo: Generate dense depth maps from lidar and stereo imagery. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 7829–7836. IEEE (2020)
82. Zhang, Y., Zou, S., Liu, X., Huang, X., Wan, Y., Yao, Y.: Lidar-guided stereo matching with a spatial consistency constraint. *ISPRS Journal of Photogrammetry and Remote Sensing* **183**, 164–177 (2022)
83. Zhao, H., Zhou, H., Zhang, Y., Chen, J., Yang, Y., Zhao, Y.: High-frequency stereo matching network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1327–1336 (2023)
84. Zhao, H., Zhou, H., Zhang, Y., Zhao, Y., Yang, Y., Ouyang, T.: Eai-stereo: Error aware iterative network for stereo matching. In: *Proceedings of the Asian Conference on Computer Vision*. pp. 315–332 (2022)
85. Zhou, Y., Gallego, G., Rebecq, H., Kneip, L., Li, H., Scaramuzza, D.: Semi-dense 3d reconstruction with a stereo event camera. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 235–251 (2018)
86. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: Unsupervised event-based learning of optical flow, depth, and egomotion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 989–997 (2019)
87. Zubić, N., Gehrig, D., Gehrig, M., Scaramuzza, D.: From chaos comes order: Ordering event representations for object recognition and detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 12846–12856 (October 2023)