

Article

Event Collapse in Contrast Maximization Frameworks

Shintaro Shiba ^{1,2,*} , Yoshimitsu Aoki ¹ and Guillermo Gallego ^{2,3} 

¹ Department of Electronics and Electrical Engineering, Faculty of Science and Technology, Keio University, 3-14-1, Kohoku-ku, Yokohama 223-8522, Kanagawa, Japan; aoki@elec.keio.ac.jp

² Department of Electrical Engineering and Computer Science, Technische Universität Berlin, 10587 Berlin, Germany; guillermo.gallego@tu-berlin.de

³ Einstein Center Digital Future and Science of Intelligence Excellence Cluster, 10117 Berlin, Germany

* Correspondence: sshiba@keio.jp

Abstract: Contrast maximization (CMax) is a framework that provides state-of-the-art results on several event-based computer vision tasks, such as ego-motion or optical flow estimation. However, it may suffer from a problem called event collapse, which is an undesired solution where events are warped into too few pixels. As prior works have largely ignored the issue or proposed workarounds, it is imperative to analyze this phenomenon in detail. Our work demonstrates event collapse in its simplest form and proposes collapse metrics by using first principles of space–time deformation based on differential geometry and physics. We experimentally show on publicly available datasets that the proposed metrics mitigate event collapse and do not harm well-posed warps. To the best of our knowledge, regularizers based on the proposed metrics are the only effective solution against event collapse in the experimental settings considered, compared with other methods. We hope that this work inspires further research to tackle more complex warp models.

Keywords: computer vision; intelligent sensors; robotics; event-based camera; contrast maximization; optical flow; motion estimation



Citation: Shiba, S.; Aoki, Y.; Gallego, G. Event Collapse in Contrast Maximization Frameworks. *Sensors* **2022**, *22*, 5190. <https://doi.org/10.3390/s22145190>

Academic Editor: Jing Tian

Received: 9 June 2022

Accepted: 7 July 2022

Published: 11 July 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Event cameras [1–3] offer potential advantages over standard cameras to tackle difficult scenarios (high speed, high dynamic range, low power). However, new algorithms are needed to deal with the unconventional type of data they produce (per-pixel asynchronous brightness changes, called events) and unlock their advantages [4]. Contrast maximization (CMax) is an event processing framework that provides state-of-the-art results on several tasks, such as rotational motion estimation [5,6], feature flow estimation and tracking [7–11], ego-motion estimation [12–14], 3D reconstruction [12,15], optical flow estimation [16–19], motion segmentation [20–24], guided filtering [25], and image reconstruction [26].

The main idea of CMax and similar event alignment frameworks [27,28] is to find the motion and/or scene parameters that align corresponding events (i.e., events that are triggered by the same scene edge), thus achieving motion compensation. The framework simultaneously estimates the motion parameters and the correspondences between events (data association). However, in some cases CMax optimization converges to an undesired solution where events accumulate into too few pixels, a phenomenon called event collapse (Figure 1). Because CMax is at the heart of many state-of-the-art event-based motion estimation methods, it is important to understand the above limitation and propose ways to overcome it. Prior works have largely ignored the issue or proposed workarounds without analyzing the phenomenon in detail. A more thorough discussion of the phenomenon is overdue, which is the goal of this work.

Contrary to the expectation that event collapse occurs when the event transformation becomes sufficiently complex [16,27], we show that it may occur even in the simplest case of one degree-of-freedom (DOF) motion. Drawing inspiration from differential geometry

and electrostatics, we propose principled metrics to quantify event collapse and discourage it by incorporating penalty terms in the event alignment objective function. Although event collapse depends on many factors, our strategy aims at modifying the objective's landscape to improve the well-posedness of the problem and be able to use well-known, standard optimization algorithms.

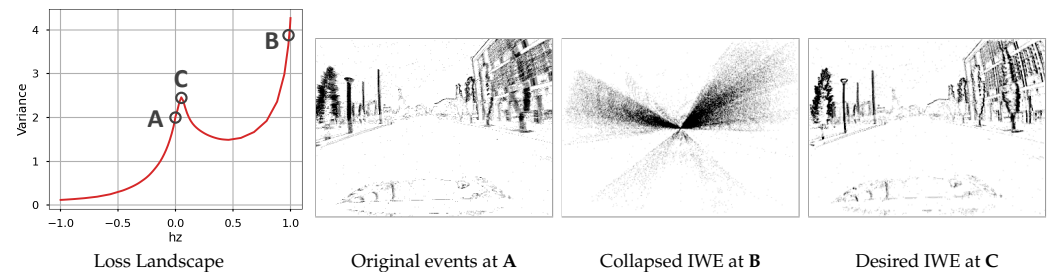


Figure 1. *Event Collapse.* **Left:** Landscape of the image variance loss as a function of the warp parameter h_z . **Right:** The IWEs at the different h_z marked in the landscape. (A) Original events (identity warp), accumulated over a small Δt (polarity is not used). (B) Image of warped events (IWE) showing event collapse due to maximization of the objective function. (C) Desired IWE solution using our proposed regularizer: sharper than (A) while avoiding event collapse (C).

In summary, our contributions are:

1. A study of the event collapse phenomenon in regard to event warping and objective functions (Sections 3.3 and 4).
2. Two principled metrics of event collapse (one based on flow divergence and one based on area-element deformations) and their use as regularizers to mitigate the above-mentioned phenomenon (Sections 3.4 to 3.6).
3. Experiments on publicly available datasets that demonstrate, in comparison with other strategies, the effectiveness of the proposed regularizers (Section 4).

To the best of our knowledge, this is the first work that focuses on the paramount phenomenon of event collapse, which may arise in state-of-the-art event-alignment methods. Our experiments show that the proposed metrics mitigate event collapse while they do not harm well-posed warps.

2. Related Work

2.1. Contrast Maximization

Our study is based on the CMax framework for event alignment (Figure 2, bottom branch). The CMax framework is an iterative method with two main steps per iteration: transforming events and computing an objective function from such events. Assuming constant illumination, events are triggered by moving edges, and the goal is to find the transformation/warping parameters θ (e.g., motion and scene) that achieve motion compensation (i.e., alignment of events triggered at different times and pixels), hence revealing the edge structure that caused the events. Standard optimization algorithms (gradient ascent, sampling, etc.) can be used to maximize the event-alignment objective. Upon convergence, the method provides the best transformation parameters and the transformed events, i.e., motion-compensated image of warped events (IWE).

The first step of the CMax framework transforms events according to a motion or deformation model defined by the task at hand. For instance, camera rotational motion estimation [5,29] often assumes constant angular velocity ($\theta \equiv \omega$) during short time spans, hence events are transformed following 3-DOF motion curves defined on the image plane by candidate values of ω . Feature tracking may assume constant image velocity $\theta \equiv \mathbf{v}$ (2-DOF) [7,30], hence events are transformed following straight lines.

In the second step of CMax, several event-alignment objectives have been proposed to measure the goodness of fit between the events and the model [10,13], establishing connections between visual contrast, sharpness, and depth-from-focus. Finally, the choice

of iterative optimization algorithm also plays a big role in finding the desired motion-compensation parameters. First-order methods, such as non-linear conjugate gradient (CG), are a popular choice, trading off accuracy and speed [12,21,22]. Exhaustive search, sampling, or branch-and-bound strategies may be affordable for low-dimensional (DOF) search spaces [14,29]. As will be presented (Section 3), our proposal consists of modifying the second step by means of a regularizer (Figure 2, top branch).

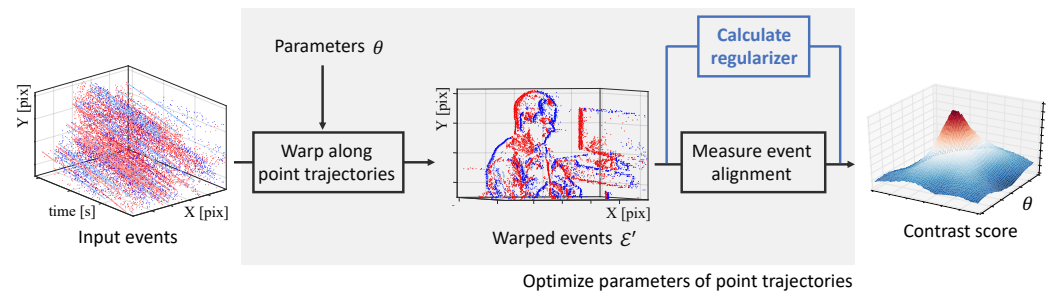


Figure 2. Proposed modification of the contrast maximization (CMax) framework in [12,13] to also account for the degree of regularity (collapsing behavior) of the warp. Events are colored in red/blue according to their polarity. Reprinted/adapted with permission from Ref. [13], 2019, Gallego et al.

2.2. Event Collapse

In which estimation problems does event collapse appear? At first look, it may appear that event collapse occurs when the number of DOFs in the warp becomes large enough, i.e., for complex motions. Event collapse has been reported in homographic motions (8 DOFs) [27,31] and in dense optical flow estimation [16], where an artificial neural network (ANN) predicts a flow field with $2N_p$ DOFs (N_p pixels), whereas it does not occur in feature flow (2 DOFs) or rotational motion flow (3 DOFs). However, a more careful analysis reveals that this is not the entire story because event collapse may occur even in the case of 1 DOF, as we show.

How did previous works tackle event collapse? Previous works have tackled the issue in several ways, such as: (i) initializing the parameters sufficiently close to the desired solution (in the basin of attraction of the local optimum) [12]; (ii) reformulating the problem, changing the parameter space to reduce the number of DOFs and increase the well-posedness of the problem [14,31]; (iii) providing additional data, such as depth [27], thus changing the problem from motion estimation given only events to motion estimation given events and additional sensor data; (iv) whitening the warped events before computing the objective [27]; and (v) redesigning the objective function and possibly adding a strong classical regularizer (e.g., Charbonnier loss) [10,16]. Many of the above mitigation strategies are task-specific because it may not always be possible to consider additional data or reparametrize the estimation problem. Our goal is to approach the issue without the need for additional data or changing the parameter space, and to show how previous objective functions and newly regularized ones handle event collapse.

3. Method

Let us present our approach to measure and mitigate event collapse. First, we revise how event cameras work (Section 3.1) and the CMax framework (Section 3.2), which was informally introduced in Section 2.1. Then, Section 3.3 builds our intuition on event collapse by analyzing a simple example. Section 3.4 presents our proposed metrics for event collapse, based on 1-DOF and 2-DOF warps. Section 3.5 specifies them for higher DOFs, and Section 3.6 presents the regularized objective function.

3.1. How Event Cameras Work

Event cameras, such as the Dynamic Vision Sensor (DVS) [2,3,32], are bio-inspired sensors that capture pixel-wise intensity changes, called events, instead of intensity images.

An event $e_k \doteq (\mathbf{x}_k, t_k, p_k)$ is triggered as soon as the logarithmic intensity L at a pixel exceeds a contrast sensitivity $C > 0$,

$$L(\mathbf{x}_k, t_k) - L(\mathbf{x}_k, t_k - \Delta t_k) = p_k C, \quad (1)$$

where $\mathbf{x}_k \doteq (x_k, y_k)^\top$, t_k (with μs resolution) and polarity $p_k \in \{+1, -1\}$ are the spatio-temporal coordinates and sign of the intensity change, respectively, and $t_k - \Delta t_k$ is the time of the previous event at the same pixel \mathbf{x}_k . Hence, each pixel has its own sampling rate, which depends on the visual input.

3.2. Mathematical Description of the CMax Framework

The CMax framework [12] transforms events in a set $\mathcal{E} = \{e_k\}_{k=1}^{N_e}$ geometrically

$$e_k \doteq (\mathbf{x}_k, t_k, p_k) \xrightarrow{\mathbf{W}} e'_k \doteq (\mathbf{x}'_k, t_{\text{ref}}, p_k), \quad (2)$$

according to a motion model \mathbf{W} , producing a set of warped events $\mathcal{E}' = \{e'_k\}_{k=1}^{N_e}$. The warp $\mathbf{x}'_k = \mathbf{W}(\mathbf{x}_k, t_k; \boldsymbol{\theta})$ transports each event along the point trajectory that passes through it (Figure 2, left), until t_{ref} is reached. The point trajectories are parametrized by $\boldsymbol{\theta}$, which contains the motion and/or scene unknowns. Then, an objective function [10,13] measures the alignment of the warped events \mathcal{E}' . Many objective functions are given in terms of the count of events along the point trajectories, which is called the image of warped events (IWE):

$$I(\mathbf{x}; \boldsymbol{\theta}) \doteq \sum_{k=1}^{N_e} b_k \delta(\mathbf{x} - \mathbf{x}'_k(\boldsymbol{\theta})). \quad (3)$$

Each IWE pixel \mathbf{x} sums the values of the warped events \mathbf{x}'_k that fall within it: $b_k = p_k$ if polarity is used or $b_k = 1$ if polarity is not used. The Dirac delta δ is in practice replaced by a smooth approximation [33], such as a Gaussian, $\delta(\mathbf{x} - \boldsymbol{\mu}) \approx \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \epsilon^2)$ with $\epsilon = 1$ pixel. A popular objective function $G(\boldsymbol{\theta})$ is the visual contrast of the IWE (3), given by the variance

$$G(\boldsymbol{\theta}) \equiv \text{Var}(I(\mathbf{x}; \boldsymbol{\theta})) \doteq \frac{1}{|\Omega|} \int_{\Omega} (I(\mathbf{x}; \boldsymbol{\theta}) - \mu_I)^2 d\mathbf{x}, \quad (4)$$

with mean $\mu_I \doteq \frac{1}{|\Omega|} \int_{\Omega} I(\mathbf{x}; \boldsymbol{\theta}) d\mathbf{x}$ and image domain Ω . Hence, the alignment of the transformed events \mathcal{E}' (i.e., the candidate “corresponding events”, triggered by the same scene edge) is measured by the strength of the edges of the IWE. Finally, an optimization algorithm iterates the above steps until the best parameters are found:

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} G(\boldsymbol{\theta}). \quad (5)$$

3.3. Simplest Example of Event Collapse: 1 DOF

To analyze event collapse in the simplest case, let us consider an approximation to a translational motion of the camera along its optical axis Z (1-DOF warp). In theory, translational motions also require the knowledge of the scene depth. Here, inspired by the 4-DOF in-plane warp in [20] that approximates a 6-DOF camera motion, we consider a simplified warp that does not require knowledge of the scene depth. In terms of data, let us consider events from one of the driving sequences of the standard MVSEC dataset [34] (Figure 1).

For further simplicity, let us normalize the timestamps of \mathcal{E} to the unit interval $t \in [t_1, t_{N_e}] \mapsto \tilde{t} \in [0, 1]$, and assume a coordinate frame at the center of the image plane, then the warp \mathbf{W} is given by

$$\mathbf{x}'_k = (1 - \tilde{t}_k h_z) \mathbf{x}_k, \quad (6)$$

where $\boldsymbol{\theta} \equiv h_z$. Hence, events are transformed along the radial direction from the image center, acting as a virtual focus of expansion (FOE) (cf. the true FOE is given by the data).

Letting the scaling factor in (6) be $s_k \doteq 1 - \tilde{t}_k h_z$, we observe the following: (i) s_k cannot be negative since it would imply that at least one event has flipped the side on which it lies with respect to the image center; (ii) if $s_k > 1$ the warped event gets away from the image center (“expansion” or “zoom-in”); and (iii) if $s_k \in [0, 1)$ the warped event gets closer to the image center (“contraction” or “zoom-out”). The equivalent conditions in terms of h_z are: (i) $h_z < 1$, (ii) $h_z < 0$ is an expansion, and (iii) $0 < h_z < 1$ is a contraction.

Intuitively, event collapse occurs if the contraction is large ($0 < s_k \ll 1$) (see Figures 1C and 3a). This phenomenon is not specific of the image variance; other objective functions lead to the same result. As we see, the objective function has a local maximum at the desired motion parameters (Figure 1B). The optimization over the entire parameter space converges to a global optimum that explains the event collapse.




Figure 3. Point trajectories (streamlines) defined on $x - y - t$ image space by various warps. (a) Zoom in/out warp from image center (1 DOF). (b) Constant image velocity warp (2 DOF). (c) Rotational warp around X axis (3 DOF).

Discussion

The above example shows that event collapse is enabled (or disabled) by the type of warp. If the warp does not enable event collapse (contraction or accumulation of flow vectors cannot happen due to the geometric properties of the warp), as in the case of feature flow (2 DOF) [7,30] (Figure 3b) or rotational motion flow (3 DOF) [5,29] (Figure 3c), then the optimization problem is well posed and multiple objective functions can be designed to achieve event alignment [10,13]. However, the disadvantage is that the type of warps that satisfy this condition may not be rich enough to describe complex scene motions.

On the other hand, if the warp allows for event collapse, more complex scenarios can be described by such a broader class of motion hypotheses, but the optimization framework designed for non-event-collapsing scenarios (where the local maximum is assumed to be the global maximum) may not hold anymore. Optimizing the objective function may lead to an undesired solution with a larger value than the desired one. This depends on multiple elements: the landscape of the objective function (which depends on the data, the warp parametrization, and the shape of the objective function), and the initialization and search strategy of the optimization algorithm used to explore such a landscape. The challenge in this situation is to overcome the issue of multiple local maxima and make the problem better posed. Our approach consists of characterizing event collapse via novel metrics and including them in the objective function as weak constraints (penalties) to yield a better landscape.

3.4. Proposed Regularizers

3.4.1. Divergence of the Event Transformation Flow

Inspired by physics, we may think of the flow vectors given by the event transformation $\mathcal{E} \mapsto \mathcal{E}'$ as an electrostatic field, whose sources and sinks correspond to the location of electric charges (Figure 4). Sources and sinks are mathematically described by the divergence operator $\nabla \cdot$. Therefore, the divergence of the flow field is a natural choice to characterize event collapse.




Figure 4. Divergence of different vector fields, $\nabla \cdot \mathbf{v} = \partial_x v_x + \partial_y v_y$. From left to right: contraction (“sink”, leading to event collapse), expansion (“source”), and incompressible fields. Image adapted from [khanacademy.org](https://www.khanacademy.org) (accessed on 6 July 2022).

The warp \mathbf{W} is defined over the space-time coordinates of the events, hence its time derivative defines a flow field over space-time:

$$\mathbf{f} \doteq \frac{\partial \mathbf{W}(\mathbf{x}, t; \boldsymbol{\theta})}{\partial t}. \quad (7)$$

For the warp in (6), we obtain $\mathbf{f} = -h_z \mathbf{x}$, which gives $\nabla \cdot \mathbf{f} = -h_z \nabla \cdot \mathbf{x} = -2h_z$. Hence, (6) defines a constant divergence flow, and imposing a penalty on the degree of concentration of the flow field accounts to directly penalizing the value of the parameter h_z .

Computing the divergence at each event gives the set

$$\mathcal{D}(\mathcal{E}, \boldsymbol{\theta}) \doteq \{\nabla \cdot \mathbf{f}_k\}_{k=1}^{N_e}, \quad (8)$$

from which we can compute statistical scores (mean, median, min, etc.):

$$R_D(\mathcal{E}, \boldsymbol{\theta}) \doteq \frac{1}{N_e} \sum_{k=1}^{N_e} \nabla \cdot \mathbf{f}_k. \quad (\text{mean}) \quad (9)$$

To have a 2D visual representation (“feature map”) of collapse, we build an image (like the IWE) by taking some statistic of the values $\nabla \cdot \mathbf{f}_k$ that warp to each pixel, such as the “average divergence per pixel”:

$$\text{DIWE}(\mathbf{x}; \mathcal{E}, \boldsymbol{\theta}) \doteq \frac{1}{N_e(\mathbf{x})} \sum_k (\nabla \cdot \mathbf{f}_k) \delta(\mathbf{x} - \mathbf{x}'_k), \quad (10)$$

where $N_e(\mathbf{x}) \doteq \sum_k \delta(\mathbf{x} - \mathbf{x}'_k)$ is the number of warped events at pixel \mathbf{x} (the IWE). Then we aggregate further into a score, such as the mean:

$$R_{\text{DIWE}}(\mathcal{E}, \boldsymbol{\theta}) \doteq \frac{1}{|\Omega|} \int_{\Omega} \text{DIWE}(\mathbf{x}; \mathcal{E}, \boldsymbol{\theta}) d\mathbf{x}. \quad (11)$$

In practice we focus on the collapsing part by computing a trimmed mean: the mean of the DIWE pixels smaller than a margin α (-0.2 in the experiments). Such a margin does not penalize small, admissible deformations.

3.4.2. Area-Based Deformation of the Event Transformation

In addition to vector calculus, we may also use tools from differential geometry to characterize event collapse. Building on [12], the point trajectories define the streamlines of the transformation flow, and we may measure how they concentrate or disperse based on how the area element deforms along them. That is, we consider a small area element $dA = dx dy$ attached to each point along the trajectory and measure how much it deforms when transported to the reference time: $dA' = |\det(\mathbf{J})| dA$, with the Jacobian

$$\mathbf{J}(\mathbf{x}, t; \boldsymbol{\theta}) \doteq \frac{\partial \mathbf{W}(\mathbf{x}, t; \boldsymbol{\theta})}{\partial \mathbf{x}} \quad (12)$$

(see Figure 5). The determinant of the Jacobian is the amplification factor: $|\det(J)| > 1$ if the area expands, and $|\det(J)| < 1$ if the area shrinks.




Figure 5. Area deformation of various warps. An area of $dA \text{ pix}^2$ at (\mathbf{x}_k, t_k) and is warped to t_{ref} , giving an area $dA' = |\det(J_k)|dA \text{ pix}^2$ at $(\mathbf{x}'_k, t_{\text{ref}})$, where $J_k \equiv J(e_k) \equiv J(\mathbf{x}_k, t_k; \theta)$ (see (12)). From left to right, increasing area amplification factor $|\det(J)| \in [0, \infty)$.

For the warp in (6), we have the Jacobian $J = (1 - \tilde{t}h_z)\text{Id}$, and so $\det(J) = (1 - \tilde{t}h_z)^2$. Interestingly, the area deformation around event e_k , $J(e_k) \equiv J(\mathbf{x}_k, t_k; \theta)$, is directly related to the scaling factor s_k : $\det(J(e_k)) = s_k^2$.

Computing the amplification factors at each event gives the set

$$\mathcal{A}(\mathcal{E}, \theta) \doteq \{|\det(J(e_k))|\}_{k=1}^{N_e} \quad (13)$$

from which we can compute statistical scores. For example,

$$R_A(\mathcal{E}, \theta) \doteq \frac{1}{N_e} \sum_{k=1}^{N_e} |\det(J(e_k))| \quad (\text{mean}) \quad (14)$$

gives an average score: $R_A > 1$ for expansion, and $R_A < 1$ for contraction.

We build a deformation map (or image of warped areas (IWA)) by taking some statistic of the values $|\det(J(e_k))|$ that warp to each pixel, such as the “average amplification per pixel”:

$$\text{IWA}(\mathbf{x}) \doteq 1 + \frac{1}{N_e(\mathbf{x})} \sum_{k=1}^{N_e} (|\det(J(e_k))| - 1) \delta(\mathbf{x} - \mathbf{x}'_k). \quad (15)$$

This assumes that if no events warp to a pixel \mathbf{x}_p , then $N_e(\mathbf{x}_p) = 0$, and there is no deformation ($\text{IWA}(\mathbf{x}_p) = 1$). Then, we summarize the deformation map into a score, such as the mean:

$$R_{\text{IWA}}(\mathcal{E}, \theta) \doteq \frac{1}{|\Omega|} \int_{\Omega} \text{IWA}(\mathbf{x}; \mathcal{E}, \theta) d\mathbf{x}. \quad (16)$$

To concentrate on the collapsing part, we compute a trimmed mean: the mean of the IWA pixels smaller than a margin α (0.8 in the experiments). The margin approves small, admissible deformations.

3.5. Higher DOF Warp Models

3.5.1. Feature Flow

Event-based feature tracking is often described by the warp $\mathbf{W}(\mathbf{x}, t; \theta) = \mathbf{x} + (t - t_{\text{ref}})\theta$, which assumes constant image velocity θ (2 DOFs) over short time intervals. As expected, the flow for this warp coincides with the image velocity, $\mathbf{f} = \theta$, which is independent of the space-time coordinates (\mathbf{x}, t) . Hence, the flow is incompressible ($\nabla \cdot \mathbf{f} = 0$): the streamlines given by the feature flow do not concentrate or disperse; they are parallel. Regarding the area deformation, the Jacobian $J = \partial(\mathbf{x} + (t - t_{\text{ref}})\theta)/\partial\mathbf{x} = \text{Id}$ is the identity matrix. Hence $|\det(J)| = 1$, that is, translations on the image plane do not change the area of the pixels around a point.

In-plane translation warps, such as the above 2-DOF warp, are well-posed and serve as reference to design the regularizers that measure event collapse. It is sensible for well-designed regularizers to penalize warps whose characteristics deviate from those of the reference warp: zero divergence and unit area amplification factor.

3.5.2. Rotational Motion

As the previous sections show, the proposed metrics designed for the zoom in/out warp produce the expected characterization of the 2-DOF feature flow (zero divergence and unit area amplification), which is a well-posed warp. Hence, if they were added as penalties into the objective function they would not modify the energy landscape. We now consider their influence on rotational motions, which are also well-posed warps. In particular, we consider the problem of estimating the angular velocity of a predominantly rotating event camera by means of CMax, which is a popular research topic [5,14,27–29]. By using calibrated and homogeneous coordinates, the warp is given by

$$\mathbf{x}^{h'} \sim \mathbf{R}(t\boldsymbol{\omega}) \mathbf{x}^h, \quad (17)$$

where $\boldsymbol{\theta} \equiv \boldsymbol{\omega} = (\omega_1, \omega_2, \omega_3)^\top$ is the angular velocity, $t \in [0, \Delta t]$, and \mathbf{R} is parametrized by using exponential coordinates (Rodrigues rotation formula [35,36]).

Divergence: It is well known that the flow is $\mathbf{f} = B(\mathbf{x}) \boldsymbol{\omega}$, where $B(\mathbf{x})$ is the rotational part of the feature sensitivity matrix [37]. Hence

$$\nabla \cdot \mathbf{f} = 3(x\omega_2 - y\omega_1). \quad (18)$$

Area element: Letting \mathbf{r}_3^\top be the third row of \mathbf{R} , and using (32)–(34) in [38],

$$\det(\mathbf{J}) = (\mathbf{r}_3^\top \mathbf{x}^h)^{-3}. \quad (19)$$

Rotations around the Z axis clearly present no deformation, regardless of the amount of rotation, and this is captured by the proposed metrics because: (i) the divergence is zero, thus the flow is incompressible, and (ii) $\det(\mathbf{J}) = 1$ since $\mathbf{r}_3 = (0, 0, 1)^\top$ and $\mathbf{x}^h = (x, y, 1)^\top$. For other, arbitrary rotations, there are deformations, but these are mild if the rotation angle $\Delta t \|\boldsymbol{\omega}\|$ is small.

3.5.3. Planar Motion

Planar motion is the term used to describe the motion of a ground robot that can translate and rotate freely on a flat ground. If such a robot is equipped with a camera pointing upwards or downwards, the resulting motion induced on the image plane, parallel to the ground plane, is an isometry (Euclidean transformation). This motion model is a subset of the parametric ones in [12], and it has been used for CMax in [14,27]. For short time intervals, planar motion may be parametrized by 3 DOFs: linear velocity (2 DOFs) and angular velocity (1 DOF). As the divergence and area metrics show in the Appendix A, planar motion is a well-posed warp. The resulting motion curves on the image plane do not lead to event collapse.

3.5.4. Similarity Transformation

The 1-DOF zoom in/out warp in Section 3.3 is a particular case of the 4-DOF warp in [20], which is an in-plane approximation to the motion induced by a freely moving camera. The same idea of combining translation, rotation, and scaling for CMax is expressed by the similarity transformation in [27]. Both 4-DOF warps enable event collapse because they allow for zoom-out motion curves. Formulas justifying it are given in the Appendix A.

3.6. Augmented Objective Function

We propose to augment previous objective functions (e.g., (5)) with penalties obtained from the metrics developed above for event collapse:

$$\theta^* = \operatorname{argmin}_{\theta} J(\theta) = \operatorname{argmin}_{\theta} (-G(\theta) + \lambda R(\theta)). \quad (20)$$

We may interpret $G(\theta)$ (e.g., contrast or focus score [13]) as the data fidelity term and $R(\theta)$ as the regularizer, or, in Bayesian terms, the likelihood and the prior, respectively.

4. Experiments

We evaluate our method on publicly available datasets, whose details are described in Section 4.1. First, Section 4.2 shows that the proposed regularizers mitigate the overfitting issue on warps that enable collapse. For this purpose we use driving datasets (MVSEC [34], DSEC [39]). Next, Section 4.3 shows that the regularizers do not harm well-posed warps. To this end, we use the ECD dataset [40]. Finally, Section 4.4 conducts a sensitivity analysis of the regularizers.

4.1. Evaluation Datasets and Metrics

4.1.1. Datasets

The MVSEC dataset [34] is a widely used dataset for various vision tasks, such as optical flow estimation [16,18,19,41,42]. Its sequences are recorded on a drone (indoors) or on a car (outdoors), and comprise events, grayscale frames and IMU data from an mDAVIS346 [43] (346×260 pixels), as well as camera poses and LiDAR data. Ground truth optical flow is computed as the motion field [44], given the camera velocity and the depth of the scene (from the LiDAR). We select several excerpts from the *outdoor_day1* sequence with a forward motion. This motion is reasonably well approximated by collapse-enabled warps such as (6). In total, we evaluate 3.2 million events spanning 10 s.

The DSEC dataset [39] is a more recent driving dataset with a higher resolution event camera (Prophesee Gen3, 640×480 pixels). Ground truth optical flow is also computed as the motion field using the scene depth from a LiDAR [41]. We evaluate on the *zurich_city_11* sequence, using in total 380 million events spanning 40 s.

The ECD dataset [40] is the de facto standard to assess event camera ego-motion [5,8,28,45–48]. Each sequence provides events, frames, a calibration file, and IMU data (at 1kHz) from a DAVIS240C camera [49] (240×180 pixels), as well as ground-truth camera poses from a motion-capture system (at 200Hz). For rotational motion estimation (3DOF), we use the natural-looking *boxes_rotation* and *dynamic_rotation* sequences. We evaluate 43 million events (10 s) of the box sequence, and 15 million events (11 s) of the dynamic sequence.

The driving datasets (MVSEC, DSEC) and the selected sequences in the ECD dataset have different type of motions: forward (which enables event collapse) vs. rotational (which does not suffer from event collapse). Each sequence serves a different test purpose, as discussed in the next sections.

4.1.2. Metrics

The metrics used to assess optical flow accuracy (MVSEC and DSEC datasets) are the average endpoint error (AEE) and the percentage of pixels with AEE greater than N pixels (denoted by “NPE”, for $N = \{3, 10, 20\}$). Both are measured over pixels with valid ground-truth values. We also use the FWL metric [50] to assess event alignment by means of the IWE sharpness (the FWL is the IWE variance relative to that of the identity warp).

Following previous works [13,27,28], rotational motion accuracy is assessed as the RMS error of angular velocity estimation. Angular velocity ω is assumed to be constant over a window of events, estimated and compared with the ground truth at the midpoint of the window. Additionally, we use the FWL metric to gauge event alignment [50].

The event time windows are as follows: the events in the time spanned by $dt = 4$ frames in MVSEC (standard in [16,18,41]), 500k events for DSEC, and 30k events for ECD [28]. The

regularizer weights for divergence (λ_{div}) and deformation (λ_{def}) are as follows: $\lambda_{\text{div}} = 2$ and $\lambda_{\text{def}} = 5$ for MVSEC, $\lambda_{\text{div}} = 50$ and $\lambda_{\text{def}} = 100$ for DSEC, and $\lambda_{\text{div}} = 5$ and $\lambda_{\text{def}} = 10$ for ECD experiments.

4.2. Effect of the Regularizers on Collapse-Enabled Warps

Tables 1 and 2 report the results on the MVSEC and DSEC benchmarks, respectively, by using two different loss functions G : the IWE variance (4) and the squared magnitude of the IWE gradient, abbreviated “Gradient Magnitude” [13]. For MVSEC, we report the accuracy within the time interval of $dt = 4$ grayscale frame (at ≈ 45 Hz). The optimization algorithm is the Tree-Structured Parzen Estimator (TPE) sampler [51] for both experiments, with a number of sampling points equal to 300 (1 DOF) and 600 (4 DOF). The tables quantitatively capture the collapse phenomenon suffered by the original CMax framework [12] and the whitening technique [27]. Their high FWL values indicate that contrast is maximized; however, the AEE and NPE values are exceedingly high (e.g., >80 pixels, $20PE > 80\%$), indicating that the estimated flow is unrealistic.

Table 1. Results of MVSEC dataset [44].

	Variance					Gradient Magnitude					
	AEE ↓	3PE ↓	10PE ↓	20PE ↓	FWL ↑	AEE ↓	3PE ↓	10PE ↓	20PE ↓	FWL ↑	
Ground truth flow	–	–	–	–	1.05	–	–	–	–	1.05	
Identity warp	4.85	60.59	10.38	0.31	1.00	4.85	60.59	10.38	0.31	1.00	
1 DOF	No regularizer	89.34	97.30	95.42	92.39	1.90	85.77	93.96	86.24	83.45	1.87
	Whitening [27]	89.58	97.18	96.77	93.76	1.90	81.10	90.86	89.04	86.20	1.85
	Divergence (Ours)	4.00	46.02	2.77	0.05	1.12	2.87	32.68	2.52	0.03	1.17
	Deformation (Ours)	4.47	52.60	5.16	0.13	1.08	3.97	48.79	3.21	0.07	1.09
	Div. + Def. (Ours)	3.30	33.09	2.61	0.48	1.20	2.85	32.34	2.44	0.03	1.17
4 DOF [20]	No regularizer	90.22	90.22	96.94	93.86	2.05	91.26	99.49	95.06	91.46	2.01
	Whitening [27]	90.82	99.11	98.04	95.04	2.04	88.38	98.87	92.41	88.66	2.00
	Divergence (Ours)	7.25	81.75	18.53	0.69	1.09	5.37	66.18	10.81	0.28	1.14
	Deformation (Ours)	8.13	87.46	18.53	1.09	1.03	5.25	64.79	13.18	0.37	1.15
	Div. + Def. (Ours)	5.14	65.61	10.75	0.38	1.16	5.41	66.01	13.19	0.54	1.14

Table 2. Results of DSEC dataset [39].

	Variance					Gradient Magnitude					
	AEE ↓	3PE ↓	10PE ↓	20PE ↓	FWL ↑	AEE ↓	3PE ↓	10PE ↓	20PE ↓	FWL ↑	
Ground truth flow	–	–	–	–	1.09	–	–	–	–	1.09	
Identity warp	5.84	60.45	16.65	3.40	1.00	5.84	60.45	16.65	3.40	1.00	
1 DOF	No regularizer	156.13	99.88	99.33	98.18	2.58	156.08	99.93	99.40	98.11	2.58
	Whitening [27]	156.18	99.95	99.51	98.26	2.58	156.82	99.88	99.38	98.33	2.58
	Divergence (Ours)	12.49	69.86	20.78	6.66	1.43	5.47	63.48	14.66	1.35	1.34
	Deformation (Ours)	9.01	68.96	18.86	4.77	1.40	5.79	64.02	16.11	2.75	1.36
	Div. + Def. (Ours)	6.06	68.48	17.08	2.27	1.36	5.53	64.09	15.06	1.37	1.35
4 DOF [20]	No regularizer	157.54	99.97	99.64	98.67	2.64	157.34	99.94	99.53	98.44	2.62
	Whitening [27]	157.73	99.97	99.66	98.71	2.60	156.12	99.91	99.26	97.93	2.61
	Divergence (Ours)	14.35	90.84	41.62	10.82	1.35	10.43	91.38	41.63	9.43	1.21
	Deformation (Ours)	15.12	94.96	62.59	22.62	1.25	10.01	90.15	39.45	8.67	1.25
	Div. + Def. (Ours)	10.06	90.65	40.61	8.58	1.26	10.39	91.02	41.81	9.40	1.23

By contrast, our regularizers (Divergence and Deformation rows) work well to mitigate the collapse, as observed in smaller AEE and NPE values. Compared with the values of no regularizer or whitening [27], our regularizers achieve more than 90% improvement for AEE on average. The AEE values are high for optical flow standards (4–8 pix in MVSEC vs. 0.5–1 pixel [16], or 10–20 pix in DSEC vs. 2–5 pix [41]); however, this is due to the fact that the warps used have very few DOFs (≤ 4) compared to the considerably higher DOFs ($2N_p$) of optical flow estimation algorithms. The same reason explains the high 3PE values (standard in [52]): using an end-point error threshold of 3 pix to consider that the flow is correctly estimated does not convey the intended goal of inlier/outlier classification for the low-DOF warps used. This is the reason why Tables 1 and 2 also report 10PE, 20PE metrics,

and the values for the identity warp (zero flow). As expected, for the range of AEE values in the tables, the 10PE and 20PE figures demonstrate the large difference between methods suffering from collapse (20PE > 80%) and those that do not (20PE < 1.1% for MVSEC and <22.6% for DSEC).

The FWL values of our regularizers are moderately high (≥ 1), indicating that event alignment is better than that of the identity warp. However, because the FWL depends on the number of events [50], it is not easy to establish a global threshold to classify each method as suffering from collapse or not. The AEE, 10PE, and 20PE are better for such a classification.

Tables 1 and 2 also include the results of the use of both regularizers simultaneously (“Div. + Def.”). The results improve across all sequences if the data fidelity term is given by the variance loss, whereas they remain approximately the same for the gradient magnitude loss. Regardless of the choice of the proposed regularizer, the results in these tables clearly show the effectiveness of our proposal, i.e., the large improvements compared with prior works (rows “No regularizer” and [27]).

The collapse results are more visible in Figure 6, where we used the variance loss. Without a regularizer, the events collapse in the MVSEC and DSEC sequences. Our regularizers successfully mitigate overfitting, having a remarkable impact on the estimated motion.

4.3. Effect of the Regularizers on Well-Posed Warps

Table 3 shows the results on the ECD dataset for a well-posed warp (3-DOF rotational motion, in the benchmark). We use the variance loss and the Adam optimizer [53] with 100 iterations. All values in the table (RMS error and FWL, with and without regularization, are very similar, indicating that: (i) our regularizers do not affect the motion estimation algorithm, and (ii) results without regularization are good due to the well-posed warp. This is qualitatively shown in the bottom part of Figure 6. The fluctuations of the divergence and deformation values away from those of the identity warp (0 and 1, respectively) are at least one order of magnitude smaller than the collapse-enabled warps (e.g., 0.2 vs. 2).

Table 3. Results on ECD dataset [40].

	boxes_rot		dynamic_rot	
	RMS ↓	FWL ↑	RMS ↓	FWL ↑
Ground truth pose	–	1.559	–	1.414
No regularizer	8.858	1.562	4.823	1.420
Divergence (Ours)	9.237	1.558	4.826	1.420
Deformation (Ours)	8.664	1.561	4.822	1.420

4.4. Sensitivity Analysis

The landscapes of loss functions as well as sensitivity analysis of λ are shown in Figure 7, for the MVSEC experiments. Without regularizer ($\lambda = 0$), all objective functions tested (variance, gradient magnitude, and average timestamp [16]) suffer from event collapse, which is the undesired global minimum of (20). Reaching the desired local optimum depends on the optimizing algorithm and its initialization (e.g., starting gradient descent close enough to the local optimum). Our regularizers (divergence and deformation) change the landscape: the previously undesired global minimum becomes local, and the desired minimum becomes the new global one as λ increases.

Specifically, the larger the weight λ , the smaller the effect of the undesired minimum (at $h_z = 1$). However, this is true only within some reasonable range: a too large λ discards the data-fidelity part G in (20), which is unwanted because it would remove the desired local optimum (near $h_z \approx 0$). Minimizing (20) with only the regularizer is not sensible.

Observe that for completeness, we include the average timestamp loss in the last column. However, this loss also suffers from an undesired optimum in the expansion region ($h_z \approx -1$). Our regularizers could be modified to also remove this undesired

optimum, but investigating this particular loss, which was proposed as an alternative to the original contrast loss, is outside the scope of this work.

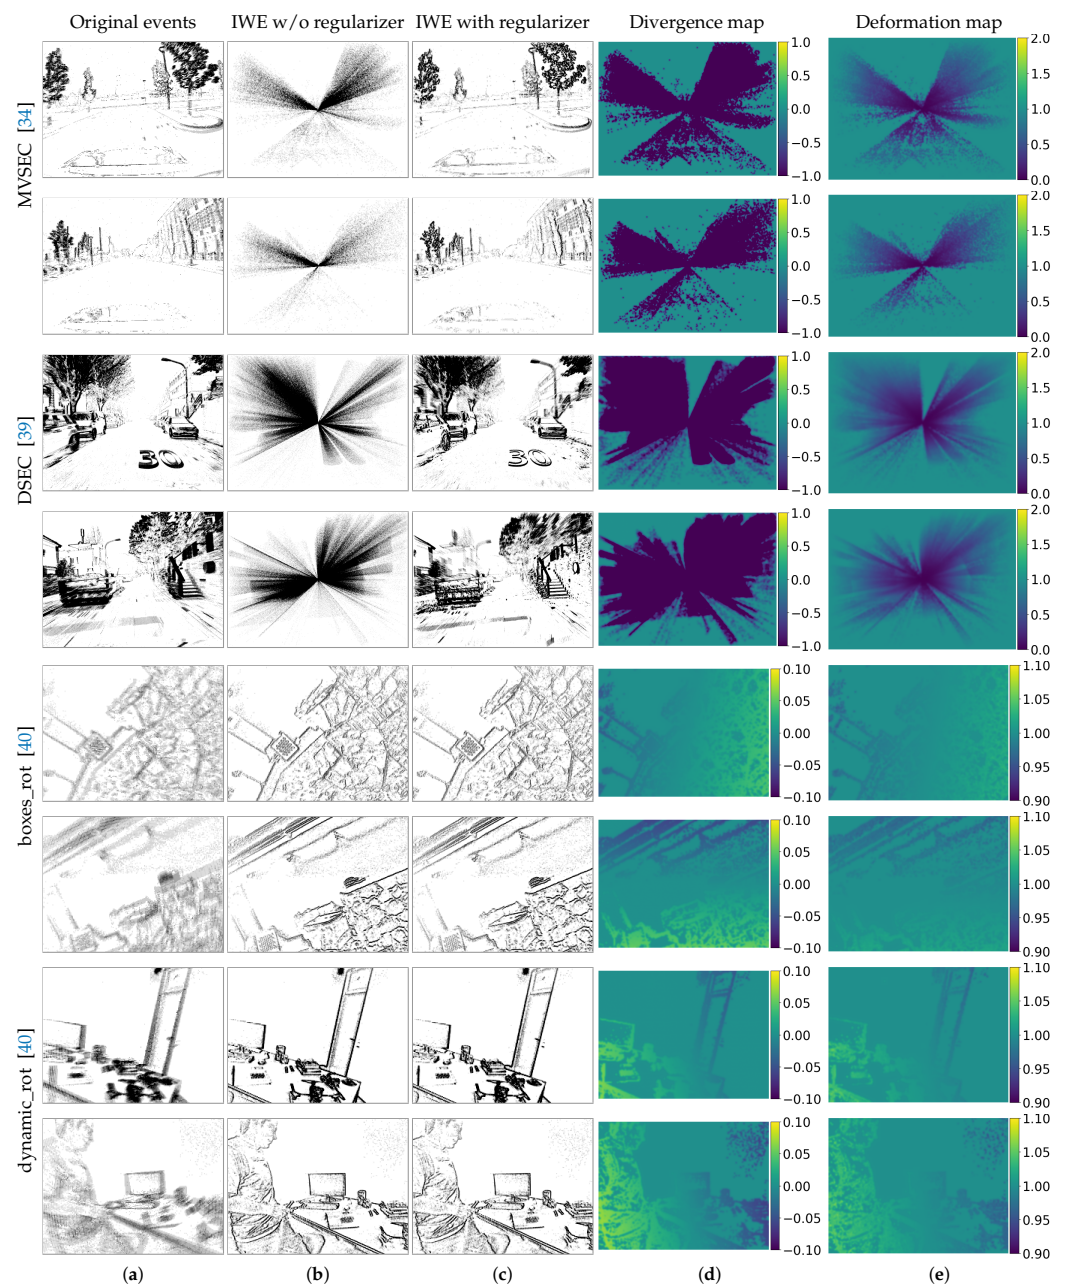


Figure 6. Proposed regularizers and collapse analysis. The scene motion is approximated by 1-DOF warp (zoom in/out) for MVSEC [34] and DSEC [39] sequences, and 3-DOF warp (rotation) for boxes and dynamic ECD sequences [40]. (a) Original events. (b) Best warp without regularization. Event collapse happens for 1-DOF warp. (c) Best warp with regularization. (d) Divergence map ((10) is zero-based). (e) Deformation map ((15), centered at 1). Our regularizers successfully penalize event collapse and do not damage non-collapsing scenarios.

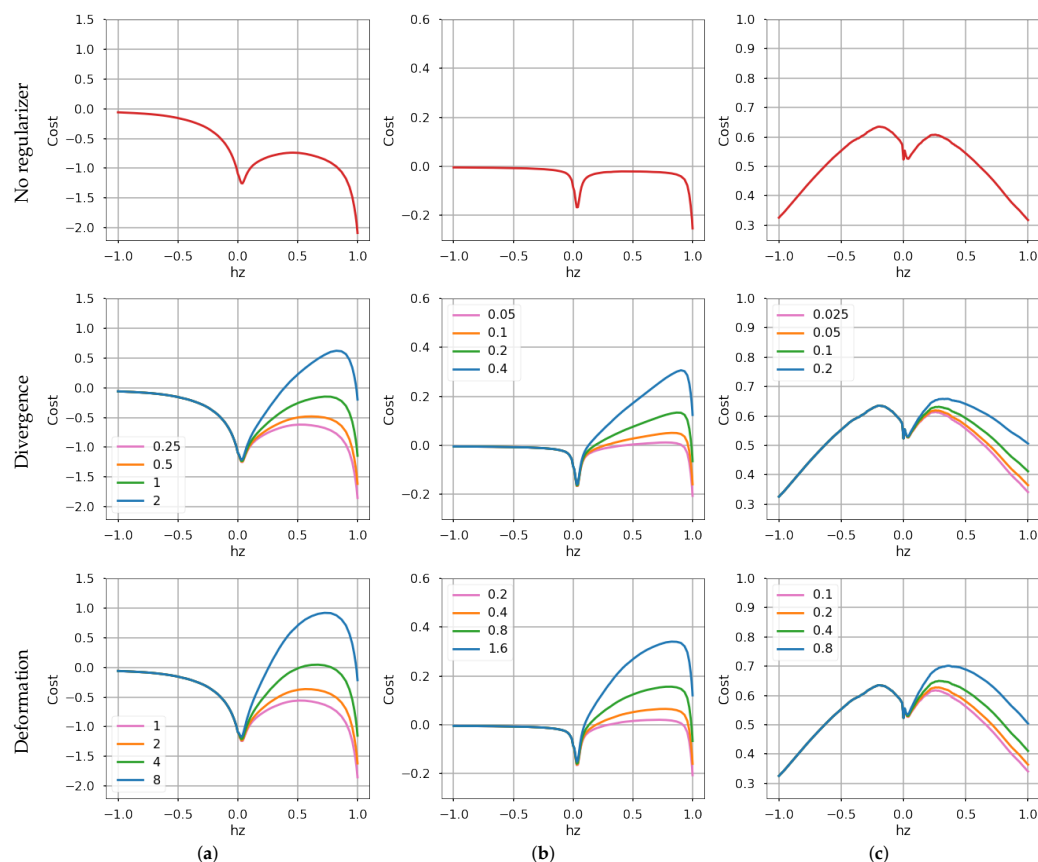


Figure 7. Cost function landscapes over the warp parameter h_z for: (a) Image variance [12], (b) gradient magnitude [13], and (c) mean square of average timestamp [16]. Data from MVSEC [34] with dominant forward motion. The legend weights denote λ in (20).

4.5. Computational Complexity

Computing the regularizer(s) requires more computation than the non-regularized objective. However, complexity is linear with the number of events and the number of pixels, which is an advantage, and the warped events are reutilized to compute the DIWE or IWA. Hence, the runtime is less than doubled (warping is the dominant runtime term [13] and is computed only once). The computational complexity of our regularized CMax framework is $O(N_e + N_p)$, the same as that of the non-regularized one.

4.6. Application to Motion Segmentation

Although most of the results on standard datasets comprise stationary scenes, we have also provided results on a dynamic scene (from dataset [40]). Because the time spanned by each set of events processed is small, the scene motion is also small (even for complicated objects like the person in the bottom row of Figure 6), hence often a single warp fits the scene reasonably well. In some scenarios, a single warp may not be enough to fit the event data because there are distinctive motions in the scene of equal importance. Our proposed regularizers can be extended to such more complex scene motions. To this end, we demonstrate it with an example in Figure 8.

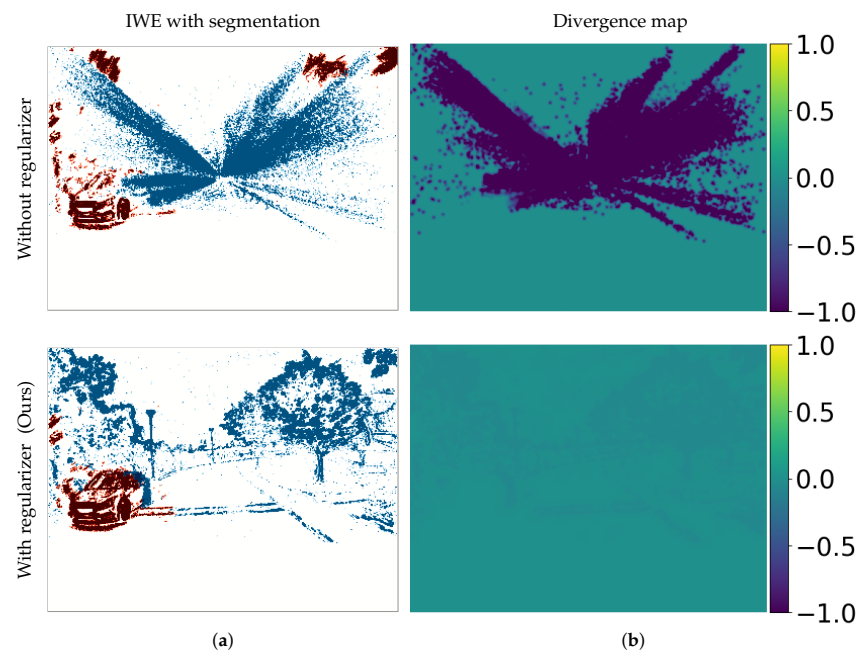


Figure 8. Application to Motion Segmentation. (a) Output IWE, whose colors (red and blue) represent different clusters of events (segmented according to motion). (b) Divergence map. The range of divergence values is larger in the presence of event collapse than in its absence. Our regularizer (divergence in this example) mitigates the event collapse for this complex motion, even with an independently moving object (IMO) in the scene.

Specifically, we use the MVSEC dataset, in a clip where the scene consists of two motions: the ego-motion (forward motion of the recording vehicle) and the motion of a car driving in the opposite direction in a nearby lane (an independently moving object—IMO). We model the scene by using the combination of two warps. Intuitively, the 1-DOF warp (6) describes the ego-motion, while the feature flow (2 DOF) describes the IMO. Then, we apply the contrast maximization approach (augmented with our regularizing terms) and the expectation-maximization scheme in [21] to segment the scene, to determine which events belong to each motion. The results in Figure 8 clearly show the effectiveness of our regularizer, even for such a commonplace and complex scene. Without regularizers, (i) event collapse appears in the ego-motion cluster of events and (ii) a considerable portion of the events that correspond to ego-motion are assigned to the second cluster (2-DOF warp), thus causing a segmentation failure. Our regularization approach mitigates event collapse (bottom row of Figure 8) and provides the correct segmentation: the 1-DOF warp fits the ego-motion and the feature flow (2-DOF warp) fits the IMO.

5. Conclusions

We have analyzed the event collapse phenomenon of the CMax framework and proposed collapse metrics using first principles of space-time deformation, inspired by differential geometry and physics. Our experimental results on publicly available datasets demonstrate that the proposed divergence and area-based metrics mitigate the phenomenon for collapse-enabled warps and do not harm well-posed warps. To the best of our knowledge, our regularizers are the only effective solution compared to the unregularized CMax framework and whitening. Our regularizers achieve, on average, more than 90% improvement on optical flow endpoint error calculation (AEE) on collapse-enabled warps.

This is the first work that focuses on the paramount phenomenon of event collapse. No prior work has analyzed this phenomenon in such detail or proposed new regularizers without additional data or reparameterizing the search space [14,16,27]. As we analyzed various warps from 1 DOF to 4 DOFs, we hope that the ideas presented here inspire further research to tackle more complex warp models. Our work shows how the divergence and

area-based deformation can be computed for warps given by analytical formulas. For more complex warps, like those used in dense optical flow estimation [16,18], the divergence or area-based deformation could be approximated by using finite difference formulas.

Author Contributions: Conceptualization, S.S. and G.G.; methodology, G.G.; software, S.S.; validation, S.S.; formal analysis, S.S. and G.G.; investigation, S.S. and G.G.; resources, Y.A.; data curation, S.S.; writing—original draft preparation, S.S. and G.G.; writing—review and editing, S.S., Y.A. and G.G.; visualization, S.S. and G.G.; supervision, Y.A. and G.G.; project administration, S.S.; funding acquisition, S.S., Y.A. and G.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the German Academic Exchange Service (DAAD), Research Grant-Bi-nationally Supervised Doctoral Degrees/Cotutelle, 2021/22 (57552338). Ref. no.: 91803781.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data available in publicly accessible repositories. The data presented in this study are openly available in reference number [34,39,40].

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Warp Models, Jacobians and Flow Divergence

Appendix A.1. Planar Motion — Euclidean Transformation on the Image Plane, $SE(2)$

If the point trajectories of an isometry are $\mathbf{x}(t)$, the warp is given by [27]

$$\begin{pmatrix} \mathbf{x}'_k \\ 1 \end{pmatrix} \sim \begin{pmatrix} \mathbf{R}(t_k \omega_Z) & t_k \mathbf{v} \\ \mathbf{0}^\top & 1 \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{x}_k \\ 1 \end{pmatrix}, \quad (\text{A1})$$

where \mathbf{v}, ω_Z comprise the 3 DOFs of a translation and an in-plane rotation. The in-plane rotation is

$$\mathbf{R}(\phi) = \begin{pmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{pmatrix}. \quad (\text{A2})$$

Since

$$\begin{pmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{0}^\top & 1 \end{pmatrix}^{-1} = \begin{pmatrix} \mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{b} \\ \mathbf{0}^\top & 1 \end{pmatrix} \quad (\text{A3})$$

and $\mathbf{R}^{-1}(\phi) = \mathbf{R}(-\phi)$, we have

$$\begin{pmatrix} \mathbf{x}'_k \\ 1 \end{pmatrix} \sim \begin{pmatrix} \mathbf{R}(-t_k \omega_Z) & -\mathbf{R}(-t_k \omega_Z)(t_k \mathbf{v}) \\ \mathbf{0}^\top & 1 \end{pmatrix} \begin{pmatrix} \mathbf{x}_k \\ 1 \end{pmatrix}. \quad (\text{A4})$$

Hence, in Euclidean coordinates the warp is

$$\mathbf{x}'_k = \mathbf{R}(-t_k \omega_Z)(\mathbf{x}_k - t_k \mathbf{v}). \quad (\text{A5})$$

The Jacobian and its determinant are:

$$\mathbf{J}_k = \frac{\partial \mathbf{x}'_k}{\partial \mathbf{x}_k} = \mathbf{R}(-t_k \omega_Z), \quad (\text{A6})$$

$$\det(\mathbf{J}_k) = 1. \quad (\text{A7})$$

The flow corresponding to (A5) is:

$$\mathbf{f} = \frac{\partial \mathbf{x}'}{\partial t} = \mathbf{R}^\top \left(\frac{\pi}{2} + t \omega_Z \right) (\mathbf{x} - t \mathbf{v}) \omega_Z - \mathbf{R}(-t \omega_Z) \mathbf{v}, \quad (\text{A8})$$

whose divergence is

$$\nabla \cdot \mathbf{f} = -2\omega_z \sin(t\omega_z). \quad (\text{A9})$$

Hence, for small angles $|t\omega_z| \ll 1$, the divergence of the flow vanishes.

In short, this warp has the same determinant and approximate zero divergence as the 2-DOF feature flow warp (Section 3.5.1), which is well-behaved. Note, however, that the trajectories are not straight in space-time.

Appendix A.2. 3-DOF Camera Rotation, $SO(3)$

Using calibrated and homogeneous coordinates, the warp is given by [5,12]

$$\mathbf{x}_k^{h'} \sim \mathbf{R}(t_k \boldsymbol{\omega}) \mathbf{x}_k^h, \quad (\text{A10})$$

where $\boldsymbol{\theta} = \boldsymbol{\omega} = (\omega_1, \omega_2, \omega_3)^\top$ is the angular velocity, and \mathbf{R} (3×3 rotation matrix in space) is parametrized using exponential coordinates (Rodrigues rotation formula [35,36]).

By the chain rule, the Jacobian is:

$$\mathbf{J}_k = \frac{\partial \mathbf{x}'_k}{\partial \mathbf{x}_k} = \frac{\partial \mathbf{x}'_k}{\partial \mathbf{x}_k^{h'}} \frac{\partial \mathbf{x}_k^{h'}}{\partial \mathbf{x}_k^h} \frac{\partial \mathbf{x}_k^h}{\partial \mathbf{x}_k} = \frac{1}{(\mathbf{x}_k^{h'})_3} \begin{pmatrix} 1 & 0 & -x'_k \\ 0 & 1 & -y'_k \end{pmatrix} \mathbf{R}(t_k \boldsymbol{\omega}) \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}. \quad (\text{A11})$$

Letting $\mathbf{r}_{3,k}^\top$ be the third row of $\mathbf{R}(t_k \boldsymbol{\omega})$, and using (32)–(34) in [38], gives

$$\det(\mathbf{J}_k) = (\mathbf{r}_{3,k}^\top \mathbf{x}_k^h)^{-3}. \quad (\text{A12})$$

Connection between Divergence and Deformation Maps

If the rotation angle $t_k \|\boldsymbol{\omega}\|$ is small, using the first two terms of the exponential map we approximate $\mathbf{R}(t_k \boldsymbol{\omega}) \approx \text{Id} + (t_k \boldsymbol{\omega})^\wedge$, where the hat operator $^\wedge$ in $SO(3)$ represents the cross product matrix [54]. Then, $\mathbf{r}_{3,k}^\top \mathbf{x}_k^h \approx (-t_k \omega_2, t_k \omega_1, 1)^\top (x_k, y_k, 1) = 1 + (y_k \omega_1 - x_k \omega_2) t_k$. Substituting this expression into (A12) and using the first two terms in Taylor's expansion around $z = 0$ of $(1+z)^{-3} \approx 1 - 3z + 6z^2$ (convergent for $|z| < 1$) gives $\det(\mathbf{J}_k) \approx 1 + 3(x_k \omega_2 - y_k \omega_1) t_k$. Notably, the divergence (18) and the approximate amplification factor depend linearly on $3(x_k \omega_2 - y_k \omega_1)$. This resemblance is seen in the divergence and deformation maps of the bottom rows in Figure 6 (ECD dataset).

Appendix A.3. 4-DOF In-Plane Camera Motion Approximation

The warp presented in [20],

$$\mathbf{x}'_k = \mathbf{x}_k - t_k (\mathbf{v} + (h_z + 1) \mathbf{R}(\phi) \mathbf{x}_k - \mathbf{x}_k) \quad (\text{A13})$$

has 4 DOFs: $\boldsymbol{\theta} = (\mathbf{v}, \phi, h_z)^\top$. The Jacobian and its determinant are:

$$\mathbf{J}_k = \frac{\partial \mathbf{x}'_k}{\partial \mathbf{x}_k} = (1 + t_k) \text{Id} - (h_z + 1) t_k \mathbf{R}(\phi), \quad (\text{A14})$$

$$\det(\mathbf{J}_k) = (1 + t_k)^2 - 2(1 + t_k) t_k (h_z + 1) \cos \phi + t_k^2 (h_z + 1)^2. \quad (\text{A15})$$

The flow corresponding to (A13) is given by

$$\mathbf{f} = \frac{\partial \mathbf{x}'_k}{\partial t} = -(\mathbf{v} + (h_z + 1) \mathbf{R}(\phi) \mathbf{x} - \mathbf{x}), \quad (\text{A16})$$

whose divergence is:

$$\nabla \cdot \mathbf{f} = -(h_z + 1) \nabla \cdot (\mathbf{R}(\phi) \mathbf{x}) + \nabla \cdot \mathbf{x} \quad (\text{A17})$$

$$= 2 - 2(h_z + 1) \cos(\phi). \quad (\text{A18})$$

As particular cases of this warp, one can identify:

- 1-DOF Zoom in/out ($\mathbf{v} = \mathbf{0}, \phi = 0$). $\mathbf{x}'_k = (1 - t_k h_z) \mathbf{x}_k$.
- 2-DOF translation ($\phi = 0, h_z = 0$). $\mathbf{x}'_k = \mathbf{x}_k - t_k \mathbf{v}$.
- 1-DOF "rotation" ($\mathbf{v} = \mathbf{0}, h_z = 0$). $\mathbf{x}'_k = \mathbf{x}_k - t_k (\mathbf{R}(\phi) \mathbf{x}_k - \mathbf{x}_k)$.

Using a couple of approximations of the exponential map in $SO(2)$, we obtain

$$\mathbf{x}'_k = \mathbf{x}_k - t_k (\mathbf{R}(\phi) - \text{Id}) \mathbf{x}_k \quad (\text{A19})$$

$$\approx \mathbf{x}_k - t_k \phi \wedge \mathbf{x}_k \quad \text{if } \phi \text{ is small} \quad (\text{A20})$$

$$= (\text{Id} + (-t_k \phi) \wedge) \mathbf{x}_k \quad (\text{A21})$$

$$\approx \mathbf{R}(-t_k \phi) \mathbf{x}_k \quad \text{if } t_k \phi \text{ is small.} \quad (\text{A22})$$

Hence, ϕ plays the role of a small angular velocity ω_Z around the camera's optical axis Z , i.e., in-plane rotation.

- 3-DOF planar motion ("isometry") ($h_z = 0$). Using the previous result, the warp splits into translational and rotational components:

$$\mathbf{x}'_k = \mathbf{x}_k - t_k (\mathbf{v} + \mathbf{R}(\phi) \mathbf{x}_k - \mathbf{x}_k) \quad (\text{A23})$$

$$\stackrel{(\text{A22})}{\approx} -t_k \mathbf{v} + \mathbf{R}(-t_k \phi) \mathbf{x}_k. \quad (\text{A24})$$

Appendix A.4. 4-DOF Similarity Transformation on the Image Plane, $\text{Sim}(2)$

Another 4-DOF warp is proposed in [27]. Its DOFs are the linear, angular and scaling velocities on the image plane: $\theta = (\mathbf{v}, \omega_Z, s)^\top$.

Letting $\beta_k = 1 + t_k s$, the warp is:

$$\begin{pmatrix} \mathbf{x}'_k \\ 1 \end{pmatrix} \sim \begin{pmatrix} \beta_k \mathbf{R}(t_k \omega_Z) & t_k \mathbf{v} \\ \mathbf{0}^\top & 1 \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{x}_k \\ 1 \end{pmatrix}. \quad (\text{A25})$$

Using (A3) gives

$$\begin{pmatrix} \mathbf{x}'_k \\ 1 \end{pmatrix} \sim \begin{pmatrix} \beta_k^{-1} \mathbf{R}(-t_k \omega_Z) & -\beta_k^{-1} \mathbf{R}(-t_k \omega_Z) (t_k \mathbf{v}) \\ \mathbf{0}^\top & 1 \end{pmatrix} \begin{pmatrix} \mathbf{x}_k \\ 1 \end{pmatrix}. \quad (\text{A26})$$

Hence, in Euclidean coordinates the warp is

$$\mathbf{x}'_k = \beta_k^{-1} \mathbf{R}(-t_k \omega_Z) (\mathbf{x}_k - t_k \mathbf{v}). \quad (\text{A27})$$

The Jacobian and its determinant are:

$$\mathbf{J}_k = \frac{\partial \mathbf{x}'_k}{\partial \mathbf{x}_k} = \beta_k^{-1} \mathbf{R}(-t_k \omega_Z), \quad (\text{A28})$$

$$\det(\mathbf{J}_k) = \beta_k^{-2} = \frac{1}{(1 + t_k s)^2}. \quad (\text{A29})$$

The following result will be useful to simplify equations. For a 2D rotation $\mathbf{R}(\phi(t))$, it holds that:

$$\frac{\partial \mathbf{R}(\phi(t))}{\partial t} = -\mathbf{R}^\top \left(\frac{\pi}{2} - \phi \right) \frac{\partial \phi}{\partial t}. \quad (\text{A30})$$

To compute the flow of (A27), there are three time-dependent factors. Hence, applying the product rule we obtain three terms, and substituting (A30) (with $\phi = -t \omega_Z$) gives:

$$\mathbf{f}_k = \left(\frac{\partial \beta_k^{-1}}{\partial t_k} \mathbf{R}(-t_k \omega_Z) + \beta_k^{-1} \omega_Z \mathbf{R}^\top \left(\frac{\pi}{2} + t_k \omega_Z \right) \right) (\mathbf{x}_k - t_k \mathbf{v}) - \beta_k^{-1} \mathbf{R}(-t_k \omega_Z) \mathbf{v}, \quad (\text{A31})$$

where, by the chain rule,

$$\frac{\partial \beta_k^{-1}}{\partial t_k} = -\beta_k^{-2} \frac{\partial \beta_k}{\partial t_k} = -\beta_k^{-2} s = -\frac{s}{(1+t_k s)^2}. \quad (\text{A32})$$

Hence, the divergence of the flow is:

$$\nabla \cdot \mathbf{f}_k = \frac{\partial \beta_k^{-1}}{\partial t_k} \nabla \cdot (\mathbf{R}(-t_k \omega_Z) \mathbf{x}_k) + \beta_k^{-1} \omega_Z \nabla \cdot (\mathbf{R}^\top \left(\frac{\pi}{2} + t_k \omega_Z \right) \mathbf{x}_k) \quad (\text{A33})$$

$$= \frac{\partial \beta_k^{-1}}{\partial t_k} 2 \cos(t_k \omega_Z) + \beta_k^{-1} \omega_Z 2 \sin(-t_k \omega_Z) \quad (\text{A34})$$

The formulas for $SE(2)$ are obtained from the above ones with $s = 0$ (i.e., $\beta_k = 1$).

References

1. Delbruck, T. Frame-free dynamic digital vision. In Proceedings of the International Symposium on Secure-Life Electronics, Tokyo, Japan, 6–7 March 2008; pp. 21–26. [\[CrossRef\]](#)
2. Suh, Y.; Choi, S.; Ito, M.; Kim, J.; Lee, Y.; Seo, J.; Jung, H.; Yeo, D.H.; Namgung, S.; Bong, J.; et al. A 1280x960 Dynamic Vision Sensor with a 4.95- μm Pixel Pitch and Motion Artifact Minimization. In Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS), Seville, Spain, 12–14 October 2020. [\[CrossRef\]](#)
3. Finateu, T.; Niwa, A.; Matolin, D.; Tsuchimoto, K.; Mascheroni, A.; Reynaud, E.; Mostafalu, P.; Brady, F.; Chotard, L.; LeGoff, F.; et al. A 1280x720 Back-Illuminated Stacked Temporal Contrast Event-Based Vision Sensor with 4.86 μm Pixels, 1.066GEPs Readout, Programmable Event-Rate Controller and Compressive Data-Formatting Pipeline. In Proceedings of the IEEE International Solid-State Circuits Conference-(ISSCC), San Francisco, CA, USA, 16–20 February 2020; pp. 112–114. [\[CrossRef\]](#)
4. Gallego, G.; Delbruck, T.; Orchard, G.; Bartolozzi, C.; Taba, B.; Censi, A.; Leutenegger, S.; Davison, A.; Conradt, J.; Daniilidis, K.; et al. Event-based Vision: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, *44*, 154–180. [\[CrossRef\]](#) [\[PubMed\]](#)
5. Gallego, G.; Scaramuzza, D. Accurate Angular Velocity Estimation with an Event Camera. *IEEE Robot. Autom. Lett.* **2017**, *2*, 632–639. [\[CrossRef\]](#)
6. Kim, H.; Kim, H.J. Real-Time Rotational Motion Estimation with Contrast Maximization Over Globally Aligned Events. *IEEE Robot. Autom. Lett.* **2021**, *6*, 6016–6023. [\[CrossRef\]](#)
7. Zhu, A.Z.; Atanasov, N.; Daniilidis, K. Event-Based Feature Tracking with Probabilistic Data Association. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Singapore, 29 May–3 June 2017; pp. 4465–4470. [\[CrossRef\]](#)
8. Zhu, A.Z.; Atanasov, N.; Daniilidis, K. Event-based Visual Inertial Odometry. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 5816–5824. [\[CrossRef\]](#)
9. Seok, H.; Lim, J. Robust Feature Tracking in DVS Event Stream using Bezier Mapping. In Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV), Snowmass, CO, USA, 1–5 March 2020; pp. 1647–1656. [\[CrossRef\]](#)
10. Stoffregen, T.; Kleeman, L. Event Cameras, Contrast Maximization and Reward Functions: An Analysis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 12292–12300. [\[CrossRef\]](#)
11. Dardelet, L.; Benosman, R.; Ieng, S.H. An Event-by-Event Feature Detection and Tracking Invariant to Motion Direction and Velocity. *TechRxiv* **2021**. [\[CrossRef\]](#)
12. Gallego, G.; Rebecq, H.; Scaramuzza, D. A Unifying Contrast Maximization Framework for Event Cameras, with Applications to Motion, Depth, and Optical Flow Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 3867–3876. [\[CrossRef\]](#)
13. Gallego, G.; Gehrig, M.; Scaramuzza, D. Focus Is All You Need: Loss Functions For Event-based Vision. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 12272–12281. [\[CrossRef\]](#)
14. Peng, X.; Gao, L.; Wang, Y.; Kneip, L. Globally-Optimal Contrast Maximisation for Event Cameras. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *44*, 3479–3495. [\[CrossRef\]](#)
15. Rebecq, H.; Gallego, G.; Mueggler, E.; Scaramuzza, D. EMVS: Event-based Multi-View Stereo—3D Reconstruction with an Event Camera in Real-Time. *Int. J. Comput. Vis.* **2018**, *126*, 1394–1414. [\[CrossRef\]](#)
16. 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 (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 989–997. [\[CrossRef\]](#)
17. Paredes-Valles, F.; Scheper, K.Y.W.; de Croon, G.C.H.E. Unsupervised Learning of a Hierarchical Spiking Neural Network for Optical Flow Estimation: From Events to Global Motion Perception. *IEEE Trans. Pattern Anal. Mach. Intell.* **2019**, *42*, 2051–2064. [\[CrossRef\]](#)

18. Hagenaaers, J.J.; Paredes-Valles, F.; de Croon, G.C.H.E. Self-Supervised Learning of Event-Based Optical Flow with Spiking Neural Networks. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Virtual-only Conference, 7–10 December 2021; Volume 34, pp. 7167–7179.
19. Shiba, S.; Aoki, Y.; Gallego, G. Secrets of Event-based Optical Flow. In Proceedings of the European Conference on Computer Vision (ECCV), Tel-Aviv, Israel, 23–27 October 2022.
20. Mitrokhin, A.; Fermüller, C.; Parameshwara, C.; Aloimonos, Y. Event-based Moving Object Detection and Tracking. In Proceedings of the IEEE/RISJ International Conference on Intelligent Robots and Systems (IROS), Madrid, Spain, 1–5 October 2018; pp. 1–9. [\[CrossRef\]](#)
21. Stoffregen, T.; Gallego, G.; Drummond, T.; Kleeman, L.; Scaramuzza, D. Event-Based Motion Segmentation by Motion Compensation. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Seoul, Korea, 27 October–2 November 2019; pp. 7243–7252. [\[CrossRef\]](#)
22. Zhou, Y.; Gallego, G.; Lu, X.; Liu, S.; Shen, S. Event-based Motion Segmentation with Spatio-Temporal Graph Cuts. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, 1–13. [\[CrossRef\]](#)
23. Parameshwara, C.M.; Sanket, N.J.; Singh, C.D.; Fermüller, C.; Aloimonos, Y. 0-MMS: Zero-shot multi-motion segmentation with a monocular event camera. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May–5 June 2021. [\[CrossRef\]](#)
24. Lu, X.; Zhou, Y.; Shen, S. Event-based Motion Segmentation by Cascaded Two-Level Multi-Model Fitting. In Proceedings of the IEEE/RISJ International Conference on Intelligent Robots and Systems (IROS), Prague, Czech Republic, 27 September–1 October 2021; pp. 4445–4452. [\[CrossRef\]](#)
25. Duan, P.; Wang, Z.; Shi, B.; Cossairt, O.; Huang, T.; Katsaggelos, A. Guided Event Filtering: Synergy between Intensity Images and Neuromorphic Events for High Performance Imaging. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**. [\[CrossRef\]](#)
26. Zhang, Z.; Yezzi, A.; Gallego, G. Image Reconstruction from Events. Why learn it? *arXiv* **2021**, arXiv:2112.06242.
27. Nunes, U.M.; Demiris, Y. Robust Event-based Vision Model Estimation by Dispersion Minimisation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**. [\[CrossRef\]](#) [\[PubMed\]](#)
28. Gu, C.; Learned-Miller, E.; Sheldon, D.; Gallego, G.; Bideau, P. The Spatio-Temporal Poisson Point Process: A Simple Model for the Alignment of Event Camera Data. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 10–17 October 2021; pp. 13495–13504. [\[CrossRef\]](#)
29. Liu, D.; Parra, A.; Chin, T.J. Globally Optimal Contrast Maximisation for Event-Based Motion Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 13–19 June 2020; pp. 6348–6357. [\[CrossRef\]](#)
30. Stoffregen, T.; Kleeman, L. Simultaneous Optical Flow and Segmentation (SOFAS) using Dynamic Vision Sensor. In Proceedings of the Australasian Conference on Robotics and Automation (ACRA), Sydney, Australia, 11–13 December 2017.
31. Ozawa, T.; Sekikawa, Y.; Saito, H. Accuracy and Speed Improvement of Event Camera Motion Estimation Using a Bird's-Eye View Transformation. *Sensors* **2022**, *22*, 773. [\[CrossRef\]](#)
32. Lichtsteiner, P.; Posch, C.; Delbruck, T. A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circuits* **2008**, *43*, 566–576. [\[CrossRef\]](#)
33. Ng, M.; Er, Z.M.; Soh, G.S.; Foong, S. Aggregation Functions For Simultaneous Attitude And Image Estimation with Event Cameras At High Angular Rates. *IEEE Robot. Autom. Lett.* **2022**, *7*, 4384–4391. [\[CrossRef\]](#)
34. Zhu, A.Z.; Thakur, D.; Ozaslan, T.; Pfrommer, B.; Kumar, V.; Daniilidis, K. The Multivehicle Stereo Event Camera Dataset: An Event Camera Dataset for 3D Perception. *IEEE Robot. Autom. Lett.* **2018**, *3*, 2032–2039. [\[CrossRef\]](#)
35. Murray, R.M.; Li, Z.; Sastry, S. *A Mathematical Introduction to Robotic Manipulation*; CRC Press: Boca Raton, FL, USA, 1994.
36. Gallego, G.; Yezzi, A. A Compact Formula for the Derivative of a 3-D Rotation in Exponential Coordinates. *J. Math. Imaging Vis.* **2014**, *51*, 378–384. [\[CrossRef\]](#)
37. Corke, P. *Robotics, Vision and Control: Fundamental Algorithms in MATLAB*; Springer Tracts in Advanced Robotics; Springer: Berlin/Heidelberg, Germany, 2017. [\[CrossRef\]](#)
38. Gallego, G.; Yezzi, A.; Fedele, F.; Benetazzo, A. A Variational Stereo Method for the Three-Dimensional Reconstruction of Ocean Waves. *IEEE Trans. Geosci. Remote Sens.* **2011**, *49*, 4445–4457. [\[CrossRef\]](#)
39. Gehrig, M.; Aarents, W.; Gehrig, D.; Scaramuzza, D. DSEC: A Stereo Event Camera Dataset for Driving Scenarios. *IEEE Robot. Autom. Lett.* **2021**, *6*, 4947–4954. [\[CrossRef\]](#)
40. Mueggler, E.; Rebecq, H.; Gallego, G.; Delbruck, T.; Scaramuzza, D. The Event-Camera Dataset and Simulator: Event-based Data for Pose Estimation, Visual Odometry, and SLAM. *Int. J. Robot. Res.* **2017**, *36*, 142–149. [\[CrossRef\]](#)
41. Gehrig, M.; Millhäusler, M.; Gehrig, D.; Scaramuzza, D. E-RAFT: Dense Optical Flow from Event Cameras. In Proceedings of the International Conference on 3D Vision (3DV), London, UK, 1–3 December 2021. [\[CrossRef\]](#)
42. Nagata, J.; Sekikawa, Y.; Aoki, Y. Optical Flow Estimation by Matching Time Surface with Event-Based Cameras. *Sensors* **2021**, *21*, 1150. [\[CrossRef\]](#)
43. Taverni, G.; Moeys, D.P.; Li, C.; Cavaco, C.; Motsnyi, V.; Bello, D.S.S.; Delbruck, T. Front and Back Illuminated Dynamic and Active Pixel Vision Sensors Comparison. *IEEE Trans. Circuits Syst. II* **2018**, *65*, 677–681. [\[CrossRef\]](#)
44. Zhu, A.Z.; Yuan, L.; Chaney, K.; Daniilidis, K. EV-FlowNet: Self-Supervised Optical Flow Estimation for Event-based Cameras. In Proceedings of the Robotics: Science and Systems (RSS), Pittsburgh, PA, USA, 26–30 June 2018. [\[CrossRef\]](#)

45. Rosinol Vidal, A.; Rebecq, H.; Horstschaef, T.; Scaramuzza, D. Ultimate SLAM? Combining Events, Images, and IMU for Robust Visual SLAM in HDR and High Speed Scenarios. *IEEE Robot. Autom. Lett.* **2018**, *3*, 994–1001. [[CrossRef](#)]
46. Rebecq, H.; Horstschäfer, T.; Gallego, G.; Scaramuzza, D. EVO: A Geometric Approach to Event-based 6-DOF Parallel Tracking and Mapping in Real-Time. *IEEE Robot. Autom. Lett.* **2017**, *2*, 593–600. [[CrossRef](#)]
47. Mueggler, E.; Gallego, G.; Rebecq, H.; Scaramuzza, D. Continuous-Time Visual-Inertial Odometry for Event Cameras. *IEEE Trans. Robot.* **2018**, *34*, 1425–1440. [[CrossRef](#)]
48. Zhou, Y.; Gallego, G.; Shen, S. Event-based Stereo Visual Odometry. *IEEE Trans. Robot.* **2021**, *37*, 1433–1450. [[CrossRef](#)]
49. Brandli, C.; Berner, R.; Yang, M.; Liu, S.C.; Delbruck, T. A 240×180 130dB $3\mu\text{s}$ Latency Global Shutter Spatiotemporal Vision Sensor. *IEEE J. Solid-State Circuits* **2014**, *49*, 2333–2341. [[CrossRef](#)]
50. Stoffregen, T.; Scheerlinck, C.; Scaramuzza, D.; Drummond, T.; Barnes, N.; Kleeman, L.; Mahony, R. Reducing the Sim-to-Real Gap for Event Cameras. In Proceedings of the European Conference on Computer Vision (ECCV), Glasgow, UK, 23–28 August 2020.
51. Bergstra, J.; Bardenet, R.; Bengio, Y.; Kégl, B. Algorithms for Hyper-Parameter Optimization. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Granada, Spain, 12–15 December 2011; Volume 24.
52. Geiger, A.; Lenz, P.; Stiller, C.; Urtasun, R. Vision meets robotics: The KITTI dataset. *Int. J. Robot. Res.* **2013**, *32*, 1231–1237. [[CrossRef](#)]
53. Kingma, D.P.; Ba, J.L. Adam: A Method for Stochastic Optimization. In Proceedings of the International Conference on Learning Representations (ICLR), San Diego, CA, USA, 7–9 May 2015.
54. Barfoot, T.D. *State Estimation for Robotics—A Matrix Lie Group Approach*; Cambridge University Press: Cambridge, UK, 2015.