

SNR Enhancement in Event-Based Grayscale Image Reconstruction via Pixel Sensitivity Characterization

Adrián Leonardo Delucchi¹; Johannes Baßler¹; Paul Bäcker²; Michael Heizmann¹ 

¹Karlsruhe Institute of Technology, Institute of Industrial Information Technology (IIIT), Karlsruhe

²Fraunhofer Institute of Optronics, System Technologies and Image Exploitation (IOSB), Karlsruhe

Abstract This paper presents a robust method for enhancing the signal-to-noise ratio (SNR) in grayscale images reconstructed via the `Ev2Gray` method from event-based vision sensors. Reconstructed images often exhibit significant noise originating from the sensor's inherent fixed-pattern noise (FPN). We propose a two-phase technique involving the empirical extraction of a pixel sensitivity matrix followed by an algorithmic normalization. By adopting a logarithmic sensor response model, the multiplicative noise is transformed into a linear additive problem, allowing for computationally efficient correction. The contributions of this work are two-fold: First, we introduce a practical and straightforward method to quantify the FPN of a sensor without the requirement of laboratory equipment. Second, we demonstrate that using the knowledge of this quantified FPN, we can significantly improve image quality.

The proposed dynamic parameter estimation enables the algorithm to adapt to varying conditions, such as temperature and gain fluctuations, without requiring manual recalibration. Experimental results demonstrate a substantial reduction in structural correlation between the FPN and the reconstructed images, decreasing from $r = 0.257$ to -0.063 and $r = 0.138$ to -0.009 , respectively, while maintaining the global mean intensity. This approach effectively suppresses persistent sensor-level artifacts across diverse operating scenarios, providing a critical advancement for high-quality semantic and structural scene interpretation in event-based vision.

Keywords: Event-based vision, Image denoising, Fixed-pattern noise (FPN), Grayscale image reconstruction, Logarithmic sensor model

1 Introduction

Event-based vision is a novel change-based image acquisition method offering improved dynamic range, reduced data redundancy and latency relative to conventional frame-based vision [1]. With the help of the `Ev2Gray` method, grayscale images can be reconstructed from static scenes, using only information from an event-based camera [2]. This robust technique enables semantic and structural scene interpretation exclusively through event-based vision information [3]. However, reconstructed images exhibit pronounced sensor-specific noise. We propose a technique to decrease the noise in the reconstructed images through fixed-pattern noise (FPN) characterization of the sensor. This is done by stacking multiple recordings of the same homogeneous scene under identical conditions, which are then processed with the method `Ev2Gray`. By averaging these reconstructed images, statistical noise factors, such as shot noise or temporal dark-current noise, are diminished while the FPN is isolated. This allows us to determine a sensitivity matrix of the sensor, which can then be used to compensate and reduce the noise in newly reconstructed images. The proposed method not only improves the quality of reconstructed images, but quantifies the FPN of a sensor in an accessible way without expensive laboratory equipment, which can be later used for further processing.

2 Related Work

Event-based vision cameras exhibit several distinct noise sources. The dominant physical noise mechanisms include: photon-shot noise, leakage and dark-current noise, and FPN [4]. Most denoising approaches rely on the empirical observation that true events cluster in space and time, whereas noise events appear isolated [1].

2.1 Classical Methods for Filtering Noise

Early methods of noise reduction in event streams rely on filtering to remove the events caused by noise. A common method for this is the Background Activity Filter (BAF) (or Nearest-Neighbour-Filter). This filter removes events that have no other events occurring within their spatial neighborhood and a short temporal window. The underlying idea is that real brightness changes in the scene trigger multiple events in neighboring pixels and within close temporal proximity, whereas noise events often remain isolated [5]. Although the BAF is already efficient and can be implemented in hardware, it was further improved by Guo and Delbruck [6].

In addition to neighborhood filters, spatio-temporal smoothing methods have been proposed. Local event sequences can be evaluated by fitting a plane in space-time – if an event does not lie on a surface that can be explained by motion, it is classified as noise [7].

Statistical criteria have also been used for filtering out noise. A threshold for the event density within a small space-time window can be defined and events whose local density falls below this threshold are eliminated [8]. This removes random isolated events in a coarse filtering step, followed by fine filtering of hot-pixel activity.

Purely neighborhood based filters, while simple and fast, can reach their limits in the presence of complex noise patterns. Nevertheless, BAF and related methods remain the foundation of real-time noise suppression in event-based cameras to this day [8].

2.2 Learning-based Approaches

In recent years, machine-learning methods have been developed to detect noise events more reliably. Baldwin *et al.* (2020) proposed an approach in which an Event Probability Mask (EPM) is first computed, resulting in the probability for each pixel that a detected event is genuine. These probabilities are estimated from auxiliary information such as simultaneously recorded intensity images and IMU data. Using such labels, a CNN for noise removal could be trained (EDnCNN). This approach not only treats background activity like the filtering methods, but also false negatives, temporal jitter and random fluctuations in the event rate [9]. Further improvements were done by Fang *et al.* with AEDNet and Jiang *et al.* with EDFormer, which achieve state-of-the-art noise suppression through learned feature processing [10, 11].

While these networks provide high accuracy in separating signals from noise events, they are often computationally expensive. This creates a bottleneck, given the high data rates inherent to event-based vision cameras. Ruan *et al.* (2025) propose EDMamba, a framework based on state-space models. EDMamba encodes events as 4D point clouds and employs efficient spatio-temporal projections to remove noise events, allowing it to achieve state-of-the-art accuracy with significantly lower resource requirements [12].

2.3 Denoising in Image Reconstruction

The reconstruction of grayscale images using event-based vision data is significantly complicated by noise. For this reason, many reconstruction methods directly incorporate mechanisms for noise compensation.

Bardow *et al.* (2016) combined optical flow and intensity estimation and used a robust smoothing term to suppress noise and improve image contrast. This reduces image noise, but can cause a loss in detail [13]. New deep-learning approaches for event-to-video reconstruction (e.g. E2VID by Rebecq *et al.* (2019)) have learned how to handle noise during the training process, likely because simulated noise events were included during training. It has been observed that these approaches behave more

robustly in comparison to optical-flow-based methods [14]. Prefiltering methods, which include spatial smoothing via a time surface of temporal smoothing to remove isolated noise, are often applied to further increase the quality of the data [14].

To the best of our knowledge, no method currently exists to compensate for one of the biggest noise factors in event-based vision cameras: the FPN. This paper presents a robust, fast and efficient approach to compensate for the FPN in grayscale image reconstruction, significantly enhancing the quality of the resulting images.

3 Method Overview

The proposed method in this paper aims to improve the quality of reconstructed grayscale images from event-based vision data using Ev2Gray [3] by reducing the noise present in these images and thus increasing the signal-to-noise ratio (SNR). This is achieved by determining the particular FPN of the sensor and compensating for it in reconstructed images. The approach is structured into two main phases: the empirical extraction of the FPN and the subsequent correction of reconstructed images. In this work we treat the Ev2Gray method as a virtual sensor and only analyze its input and output.

3.1 Ev2Gray Method Description

Ev2Gray is a novel method for reconstructing pseudo-grayscale images from pure event-based vision data. By leveraging a moving opaque stripe to impose controlled occlusion of the scene, every pixel observes a uniform intensity anchor during the sweep. This approach enables the recovery of absolute brightness structures and fine details in static scenes without requiring external sensors or pre-trained models. The method effectively bridges the gap between asynchronous event streams and static scene interpretation, and further provides a computationally simple pathway for applying conventional machine vision algorithms developed for intensity images to event-based vision use cases. For further details about the Ev2Gray method the reader is referred to Baßler *et al.* (2025, 2026) [2, 3].

3.2 Determining the FPN

To quantify the FPN inherent to the used sensor, it must be isolated from transient noise sources. This is achieved by exploiting the different temporal characteristics of the noise components. While FPN is time-invariant due to the physical properties of the sensor (considering all parameters remain identical, e.g. temperature), all other noise sources (e.g., photon-shot noise, dark-current noise) are stochastic and time-variant.

Although these stochastic sources exhibit spatial variance in any single frame, their zero-mean distribution allows for effective suppression through temporal integration. By averaging over multiple

recordings, the variance of the stochastic spatial deviations is reduced, causing them to converge towards their expectation value. Consequently, the temporally averaged image $\bar{I}$ remains primarily affected by the time-invariant FPN, which is reinforced as the dominant spatial non-uniformity.

The practical implementation involves the following procedure: a homogeneous, static, evenly lit scene is recorded and reconstructed using the method `Ev2Gray` [3]. This is repeated N times while maintaining static scene conditions and ensuring that all camera parameters and `Ev2Gray` settings remain identical to guarantee experimental consistency. The resulting reconstructed grayscale images are then pixel-wise stacked to create a new grayscale image, which contains the average gray value for each pixel computed across all N images. This temporally averaged image is referred to as $\bar{I}$ and serves as the basis for deriving the sensitivity matrix S_{meas} , which provides an approximation of the relative sensitivity of each sensor pixel.

This approximation is achieved by determining the pixel-wise deviation of $\bar{I}$ from a reference value μ_{ref} . While a global average of $\bar{I}$ could serve as a theoretical baseline, achieving perfect homogeneity across the entire scene is difficult. Therefore, a moving average filter is used to determine μ_{ref} instead, resulting in a bounded positive value due to the nature of the output of `Ev2Gray`. This allows the reference value to adapt to local intensity variations while maintaining an ideal uniform response. We define the sensitivity matrix S_{meas} as:

$$S_{\text{meas}}(x, y) := \frac{\bar{I}(x, y)}{\mu_{\text{ref}}(x, y)} \quad (1)$$

The matrix S_{meas} represents the FPN of the sensor by stating the approximated empirically inferred sensitivity of each pixel. Each entry of S_{meas} expresses the deviation for that specific pixel, where values greater than 1 indicate increased sensitivity and values smaller than 1 indicate decreased sensitivity relative to the sensor's local average.

3.3 Algorithmic Correction

Once the FPN of the used event-based vision sensor is quantified with the sensitivity matrix S_{meas} , grayscale images which were recorded with the same sensor – thereby exhibiting the same FPN characteristics – can be corrected.

3.3.1 Modelling the Intensity Measurement

A logarithmic sensor response model is chosen due to the nature of event-based vision sensors [1]. The FPN present in these is predominantly caused by spatial variations in pixel sensitivity and contrast threshold. Additional additive offsets and background activity may contribute under low-signal conditions, but typically represent a secondary effect. For this reason, the mathematical model of

the sensor response used in this work will only take into consideration the FPN.

While noise in conventional frame-based imaging is often modeled as an additive component, event-based vision sensors require a different approach due to their inherent response to relative contrast changes rather than absolute intensity. In this context, spatial variations in pixel sensitivity act as a gain error. This has its cause in the fact that a change in the sensitivity of an event-based pixel produces an error which scales with the radiance. Under low-light conditions the absolute error is small, whereas in high-intensity scenarios, the absolute error becomes large [15].

Furthermore, the FPN is assumed to be of exponential nature to facilitate the calculation of the scaling factor α . This factor accounts for the varying magnitudes of FPN intensity resulting from the inconsistent normalization inherent in `Ev2Gray`.

We assume that the reconstructed pixel intensity I_{meas} is a function of I_{ideal} , the ideal scene radiance without FPN, modulated by S_{ideal} , the true spatial sensitivity pattern representing the exact FPN inherent to the sensor used, according to the following relationship:

$$I_{\text{meas}} := I_{\text{ideal}} \cdot \exp(\alpha \cdot S_{\text{ideal}}) \quad (2)$$

where α denotes a scalar weighting factor.

3.3.2 Logarithmic Linearization

To facilitate a computationally efficient correction, the multiplicative model is linearized by transforming the intensity values into the logarithmic domain. To ensure numerical stability and avoid undefined values at $I_{\text{meas}} = 0$, a small regularization constant $\epsilon = 10^{-9}$ is introduced:

$$\ln(I_{\text{meas}} + \epsilon) = \ln(I_{\text{ideal}}) + \alpha \cdot S_{\text{ideal}} \quad (3)$$

Defining $L := \ln(I)$, the relationship can be expressed as follows:

$$L_{\text{meas}} := L_{\text{ideal}} + \alpha \cdot S_{\text{ideal}} \quad (4)$$

3.3.3 Dynamic Parameter Estimation

The goal is to determine L_{ideal} by estimating the unknown coefficient α , as this would yield the FPN-free image. Since the intensity of FPN fluctuates based on environmental factors such as sensor temperature or gain settings, the scaling factor α must be estimated dynamically for each individual image [16]. Starting from the linear model defined in Eq. (4), we treat the estimation of α as an ordinary least-squares (OLS) problem and aim to minimize the residual difference between the observed signal (L_{meas}) and the noise model prediction ($\alpha \cdot S_{\text{ideal}}$). By centering both the measured log-intensity L_{meas} and the sensitivity matrix S_{ideal} around their respective global spatial means, $\bar{L}_{\text{meas}}$ and $\bar{S}_{\text{ideal}}$, we

isolate the structural fluctuations and eliminate the need for an explicit intercept term, focusing solely on the structural correlation.

The optimal scaling factor α is then found by minimizing the residual variance between the observed noise and our model:

$$J = \min_{\alpha} \sum_{\substack{x \in [0, W] \\ y \in [0, H]}} [(L_{\text{meas}}(x, y) - \bar{L}_{\text{meas}}) - \alpha(S_{\text{ideal}}(x, y) - \bar{S}_{\text{ideal}})]^2 \quad (5)$$

After centering L_{meas} and S_{ideal} by subtracting their respective spatial means, the estimation of α reduces to a least-squares regression without an intercept. The corresponding minimizer is:

$$\alpha = \frac{\text{cov}(L_{\text{meas}}, S_{\text{ideal}})}{\text{var}(S_{\text{ideal}})} \quad (6)$$

Hence, α captures the degree to which the measured log-intensity is structurally aligned with the FPN pattern and adapts the correction strength to the current image. A larger value of α indicates that the spatial deviations in L_{meas} more strongly follow the pattern encoded in S_{ideal} , and therefore that a stronger FPN correction is required.

3.3.4 Image Reconstruction and Validation

Once the optimal α is determined, the noise-compensated image in the log-domain can be calculated by solving Eq. (4) for L_{ideal} .

Because S_{meas} only approximates the true sensitivity matrix S_{ideal} , the corrected result is likewise an approximation of L_{ideal} , which we define as $L_{\text{corrected}}$:

$$L_{\text{corrected}} := L_{\text{meas}} - \alpha \cdot S_{\text{meas}} \quad (7)$$

$L_{\text{corrected}}$ represents an estimate for the log-image without FPN. The final image $I_{\text{corrected}}$ is reconstructed using the exponential function:

$$I_{\text{corrected}} := \exp(L_{\text{corrected}}) - \epsilon \quad (8)$$

To preserve numerical consistency for standard display pipelines, we apply min-max scaling to map the floating-point values onto the interval $[0, 255]$, followed by a conversion to 8-bit unsigned integer format.

4 Results

In the following chapter, the results obtained using the presented method are discussed. To evaluate the performance, a Prophesee EVK4 camera was utilized. Although the evaluation is limited to a single sensor type, FPN is a well-documented characteristic of event-based vision, and therefore this method remains applicable to other event-based architectures [1].

4.1 Quantitative Performance Analysis

To determine the sensitivity matrix S_{meas} , a kernel size of 250 was used for the moving average filter (see Section 3.2). The efficacy of the logarithmic FPN correction was evaluated by quantifying the structural correlation between the image and the sensitivity matrix S_{meas} . Ideally, a successful correction yields a correlation coefficient near zero, indicating a complete decoupling of sensor-induced patterns from the signal.

4.1.1 Test Case *Schloss Karlsruhe*

For the test case *Schloss Karlsruhe*, the algorithm estimated an optimal scaling factor $\alpha = 0.0033$.



Figure 1: Example *Schloss Karlsruhe*. Left: uncorrected image I_{meas} . Right: FPN corrected image $I_{\text{corrected}}$.

The test case image is presented in Fig. 1, where the left half displays the uncorrected reconstruction I_{meas} and the right half shows the result after FPN correction $I_{\text{corrected}}$.

As summarized in Table 1, the high initial correlation with the FPN structures ($r_{\text{uncorrected}} = 0.257$) – representing the visible FPN – was reduced to a minimal level ($r_{\text{corrected}} = -0.063$). The remaining residual correlation is attributed to the fact that the measured sensitivity matrix S_{meas} serves as an approximation of the underlying ideal sensitivity S_{ideal} , limited by experimental uncertainties. Nevertheless, this result indicates a successful suppression of the dominant noise patterns. Furthermore, the global mean intensity remained largely unchanged, confirming that the correction does not introduce unwanted brightness shifts.

Metric	Uncorrected	Corrected
α	–	0.0033
r	0.257	–0.063
$\bar{I}$	147.871	147.103

Table 1: Statistics of test case *Schloss Karlsruhe*

4.1.2 Test Case *Grayscale Chart*

For the test case *Grayscale Chart*, the algorithm estimated an optimal scaling factor $\alpha = 0.0019$. The test case image is presented in Fig. 2, displaying the uncorrected reconstruction I_{meas} in the left half, while the FPN corrected reconstruction $I_{\text{corrected}}$ is visible as the right half of the image.

As summarized in Table 2, the notable initial correlation between the image and the FPN structures ($r_{\text{uncorrected}} = 0.134$) was reduced to a negligible level ($r_{\text{corrected}} = -0.009$). Consistent with the aforementioned limitations of S_{meas} as an approx-

imation, this minor residual indicates an effective suppression of the dominant noise patterns. Furthermore, the global mean intensity also remained largely unchanged.

Metric	Uncorrected	Corrected
α	–	0.0019
r	0.134	–0.009
$\bar{I}$	139.505	139.167

Table 2: Statistics of test case *Grayscale Chart*

4.2 Visual Validation and Signal Quality

A qualitative comparison of the reconstructed grayscale images before and after FPN correction shows quite clearly that in the uncorrected state (Fig. 1 left half, Fig. 2 left half) pronounced high-frequency noise artifacts are visible, especially in homogeneous regions of the scene.

Following the application of the proposed method, these artifacts are severely reduced, resulting in a visually smoother intensity distribution with higher global SNR (Fig. 1 right half, Fig. 2 right half).

Unlike a Gaussian filter, our approach does not function as a simple low-pass filter. While it effectively suppresses noise, it preserves fine structural details and high-frequency components – such as the windows and fence in Fig. 1 – which would otherwise be lost to smoothing.

4.3 Correlation between Scene Illumination and Scaling Factor

To evaluate the adaptability of the proposed method, the relationship between the scene illumination E and the estimated coefficient α was evaluated. This investigation is of interest to determine whether the sensor’s FPN scaling follows a




Figure 2: Example *Grayscale Chart*. Left: uncorrected image I_{meas} . Right: FPN corrected image $I_{\text{corrected}}$.




Figure 3: Correlation between scene illumination E and estimated scaling factor α . The scatter plot shows individual recordings, while the red line tracks the mean value per illumination level.

predictable, deterministic trend relative to illumination intensity, which could potentially allow for a static calibration bypass. Measurements of the same homogeneous, static, evenly lit scene were made, solely varying the scene illumination in 21 steps from 2 lux to 530 lux, capturing five recordings per illumination level.

Fig. 3 illustrates the distribution: blue dots represent individual scaling factors, while red points indicate the mean scaling factor for each illumination value. The Pearson correlation factor for this distribution is 0.2884, indicating a weak correlation between the scaling factor α and the lux value E of the scene. The high variance observed in the scaling factor suggests that while scene illumination influences the presence of FPN, it represents a secondary effect. Other factors appear to have a more significant impact on the scaling factor's distribution.

5 Conclusion

This paper presented a robust method for enhancing the SNR in event-based grayscale image reconstruction through characterization and compensation for the FPN of the sensor. By adopting a logarithmic sensor response model, we successfully transformed a multiplicative noise problem into an additive one. The proposed dynamic parameter estimation allows the algorithm to adapt to varying environmental conditions – such as gain and temperature fluctuations – without manual recalibration or additional measurements of the scene.

Experimental results on reconstructed images demonstrate that the method achieves a substantial reduction in structural correlation, decreasing from $r_{\text{uncorrected}} = 0.257$ to $r_{\text{corrected}} = -0.063$ for the test case *Schloss Karlsruhe* and from $r_{\text{uncorrected}} = 0.138$ to $r_{\text{corrected}} = -0.009$ for the test case *Grayscale Chart*. Furthermore, the stability of the

global mean intensity confirms that the correction preserves the original scene luminance. The presented method provides a critical step forward in making event-based vision more viable for high-quality semantic and structural scene interpretation by effectively suppressing persistent sensor-level artifacts.

6 Outlook

The findings in this work provide a foundation for several promising research avenues. First, while this study serves as a proof-of-concept using a single event-based vision sensor, future research will involve multi-camera setups to demonstrate the generalizability of the proposed method across different hardware architectures. Furthermore, the sensitivity matrix S_{meas} offers a unique diagnostic tool: by comparing matrices from various sensors, inherent manufacturing differences and hardware-specific noise characteristics can be quantified. This cross-sensor analysis will be crucial for developing universal denoising standards in event-based vision. Second, the direct integration of the presented FPN correction into the *Ev2Gray* reconstruction provides an immediate refinement of the method, yielding higher-fidelity images natively. This integration leads to a more streamlined and automated reconstruction pipeline.

Third, incorporating the sensitivity matrix S_{meas} directly into the event processing pipeline offers significant potential. By weighting or filtering individual events in the stream based on their specific sensitivity, it may be possible to suppress FPN in real-time. This approach could significantly enhance motion compensation and enable more efficient bulk sorting of event data, as events can be prioritized based on their reliability (weighted by S_{meas}). Such “event-level” denoising would enhance the data quality for all subsequent higher-level algorithms, such as object detection or tracking.

Finally, the core principles of this compensation method could be expanded to conventional frame-based CMOS cameras. Given that pixel-wise sensitivity variations are a universal phenomenon in CMOS technology, adapting this robust, empirical calibration approach could enhance the dynamic range and image quality of standard vision systems.

References

- [1] Guillermo Gallego, Tobi Delbrück, et al. "Event-based Vision: A Survey". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44.1 (2022), pp. 154–180. DOI: 10.1109/TPAMI.2020.3008413. URL: <https://ieeexplore.ieee.org/document/9138762>.
- [2] J. Baßler and M. Heizmann. "Ev2Gray: generating intensity images from event streams". In: *Proc. Sensor and Measurement Science International (SMSI)*. Nürnberg, 2025, pp. 117–118. DOI: 10.5162/SMSI2025/B6.4.
- [3] J. Baßler, A. L. Delucchi, et al. "Ev2Gray: event-only intensity imaging – occlusion-referenced log-intensity reconstruction". In: *tm – Technisches Messen* (2026). DOI: 10.1515/teme-2025-0124.
- [4] Daxin Gu, Jia Li, et al. "Reliable Event Generation With Invertible Conditional Normalizing Flow". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46.2 (2024), pp. 927–943. DOI: 10.1109/TPAMI.2023.3326538.
- [5] Vandana Padala, Arindam Basu, and Garrick Orchard. "A Noise Filtering Algorithm for Event-Based Asynchronous Change Detection Image Sensors on TrueNorth and Its Implementation on TrueNorth". In: *Frontiers in Neuroscience* Volume 12 - 2018 (2018). ISSN: 1662-453X. DOI: 10.3389/fnins.2018.00118. URL: <https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2018.00118>.
- [6] Shasha Guo and Tobi Delbrück. "Low Cost and Latency Event Camera Background Activity Denoising". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45.1 (2023), pp. 785–795. DOI: 10.1109/TPAMI.2022.3152999.
- [7] Tobi Delbrück. "Frame-free dynamic digital vision". In: *Proceedings of the International Symposium on Secure-Life Electronics, Advanced Electronics for Quality Life and Society* (Mar. 2008). DOI: 10.5167/uzh-17620.
- [8] Yang Feng, Hengyi Lv, et al. "Event Density Based Denoising Method for Dynamic Vision Sensor". In: *Applied Sciences* 10 (Mar. 2020), p. 2024. DOI: 10.3390/app10062024.
- [9] R. Wes Baldwin, Mohammed Almatrafi, et al. "Event Probability Mask (EPM) and Event Denoising Convolutional Neural Network (EDnCNN) for Neuromorphic Cameras". In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020, pp. 1698–1707. DOI: 10.1109/CVPR42600.2020.00177.
- [10] Huachen Fang, Wu Jinjian, et al. "AEDNet: Asynchronous Event Denoising with Spatial-Temporal Correlation among Irregular Data". In: Oct. 2022, pp. 1427–1435. DOI: 10.1145/3503161.3548048.
- [11] Bin Jiang, Bo Xiong, et al. "EDformer: Transformer-Based Event Denoising Across Varied Noise Levels". In: Nov. 2024, pp. 200–216. ISBN: 978-3-031-73382-6. DOI: 10.1007/978-3-031-73383-3_12.
- [12] Ciyu Ruan, Zihang Gong, et al. *EDmamba: A Simple yet Effective Event Denoising Method with State Space Model*. May 2025. DOI: 10.48550/arXiv.2505.05391.
- [13] Hanme Kim, Stefan Leutenegger, and Andrew Davison. "Real-Time 3D Reconstruction and 6-DoF Tracking with an Event Camera". In: vol. 9910. Oct. 2016, pp. 349–364. ISBN: 978-3-319-46465-7. DOI: 10.1007/978-3-319-46466-4_21.
- [14] Shintaro Shiba, Yoshimitsu Aoki, and Guillermo Gallego. *Simultaneous Motion And Noise Estimation with Event Cameras*. 2025. DOI: 10.48550/arXiv.2504.04029. arXiv: 2504.04029 [cs.CV]. URL: <https://arxiv.org/abs/2504.04029>.
- [15] Ziwei Wang, Yonhon Ng, et al. *Event Camera Calibration of Per-pixel Biased Contrast Threshold*. 2020. DOI: 10.48550/arxiv.2012.09378. arXiv: 2012.09378 [cs.CV]. URL: <https://arxiv.org/abs/2012.09378>.
- [16] Christian Brandli, Raphael Berner, et al. "A 240×180 130 dB 3 μ s Latency Global Shutter Spatiotemporal Vision Sensor". In: *IEEE Journal of Solid-State Circuits* 49.10 (2014), pp. 2333–2341. DOI: 10.1109/JSSC.2014.2342715.