

Enhancing Visual Place Recognition via Fast and Slow Adaptive Biasing in Event Cameras

Gokul B. Nair

Michael Milford

Tobias Fischer

Abstract—Event cameras are increasingly popular in robotics due to beneficial features such as low latency, energy efficiency, and high dynamic range. Nevertheless, their downstream task performance is greatly influenced by the optimization of bias parameters. These parameters, for instance, regulate the necessary change in light intensity to trigger an event, which in turn depends on factors such as the environment lighting and camera motion. This paper introduces feedback control algorithms that automatically tune the bias parameters through two interacting methods: 1) An immediate, on-the-fly *fast* adaptation of the refractory period, which sets the minimum interval between consecutive events, and 2) if the event rate exceeds the specified bounds even after changing the refractory period repeatedly, the controller adapts the pixel bandwidth and event thresholds, which stabilizes after a short period of noise events across all pixels (*slow* adaptation). Our evaluation focuses on the visual place recognition task, where incoming query images are compared to a given reference database. We conducted comprehensive evaluations of our algorithms’ adaptive feedback control in real-time. To do so, we collected the QCR-Fast-and-Slow dataset that contains DAVIS346 event camera streams from 366 repeated traversals of a Scout Mini robot navigating through a 100 meter long indoor lab setting (totaling over 35km distance traveled) in varying brightness conditions with ground truth location information. Our proposed feedback controllers result in superior performance when compared to the standard bias settings and prior feedback control methods. Our findings also detail the impact of bias adjustments on task performance and feature ablation studies on the fast and slow adaptation mechanisms.

I. INTRODUCTION

Event cameras distinguish themselves from traditional cameras by their unique operational principle. Unlike traditional cameras, which capture images at set intervals, event cameras generate data when individual pixels detect changes in brightness, leading to OFF or ON events when the change falls below or exceeds a certain threshold. The resulting output is an event stream that includes the (x, y) pixel coordinates, a high-precision timestamp t , and the polarity p indicating the event type (OFF/ON). This sparse and asynchronous data capture method grants event cameras high temporal resolution, low latencies, and reduced power consumption; qualities that could be advantageous for autonomous systems, particularly in tasks like motion segmentation, pose estimation, visual

This work received funding from Intel Labs to TF and MM, and from ARC Laureate Fellowship FL210100156 to MM. The authors acknowledge continued support from the Queensland University of Technology (QUT) through the Centre for Robotics and the Research Engineering Facility (REF) team for providing engineering support, expertise and research infrastructure.

The authors are with the Centre for Robotics, Faculty of Engineering, Queensland University of Technology, Brisbane, QLD Australia 4000 gokulbnr@gmail.com

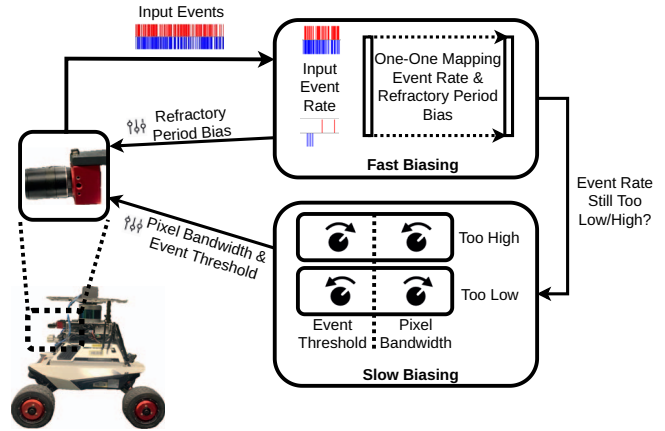


Fig. 1. **Closed-loop biasing of event cameras.** We perform feedback control on event camera biases to obtain improved performance for the visual place recognition downstream task. We achieve this by varying refractory period, pixel bandwidth, and event threshold parameters in real-time. As illustrated, we modulate these device parameters to monitor and maintain event-rate from an onboard event-camera mounted on a mobile robot.

odometry, simultaneous localization and mapping (SLAM) and neuromorphic control [1].

While traditional camera performance has greatly benefited from extensive optimization efforts over the years, such as adjustments to converter gain, noise filtering, and exposure time, this paper focuses on the development of similar automated feedback control strategies for event cameras. Given that it is often impractical to manually adjust an event camera’s bias parameters during robotic missions, this study presents a case for the necessity of intelligent, automated tuning based on the event rate – a metric that counts the number of events within a specified timeframe – to significantly enhance performance in downstream tasks (Figure 1).

We demonstrate our proposed feedback controller on the Visual Place Recognition (VPR) task [2], [3]. VPR benefits the long-term autonomy of mobile robots as a core component in localization and mapping algorithms, aiding loop-closure in SLAM and global re-localization in kidnapped robot scenarios. Event-based VPR has seen an increased interest over the past few years [4]–[11], but an in-depth investigation of the impact of the bias parameters on VPR performance is lacking until now. Our contributions are summarized as follows:

- 1) The introduction of a novel closed-loop feedback controller that, for the first time, seamlessly integrates continuous ‘fast’ adjustments of the refractory period with ‘slow’ periodic modifications of bias parameters that influence the camera’s sensitivity.

- 2) An in-depth analysis of bias parameter tuning for VPR, showcasing the significant performance improvements offered by our fast-and-slow feedback controller over both the standard bias settings and previous feedback control methods [12].
- 3) We collect and make publicly available the QCR-Fast-and-Slow-Dataset which is used to compare and evaluate our fast-slow feedback controller with prior approaches. The dataset contains 366 traverses of an indoor lab environment (100 meters) across high, medium, and low brightness conditions spanning various combinations of bias parameters. The dataset contains the event streams, conventional image data, and ground truth robot poses.

To support ongoing research in this area, we release both the algorithm and the QCR-Fast-and-Slow Dataset: <https://github.com/gokulbnr/fast-slow-biased-event-vpr>

II. RELATED WORKS

This section first provides an overview of typical approaches to process an event stream in Section II-A. We then introduce event-based VPR frameworks in Section II-B. We also note that [1] provides an excellent overview of event-based vision, with a focus on applications and algorithms that most benefit from the unique properties of event cameras.

A. Event stream representations and processing

As single events do not carry sufficient information for downstream tasks, they are often processed in batches [1]. Many early works relied on temporal collections of events that are processed into image-like event frames, where each pixel contains information about e.g. the count of events within a time window [13], [14]. Such batching enables the processing of these event frames with standard frame-based computer vision algorithms. The time window can be variable, with some approaches introducing adaptive time-slices of events based on slice event counts to perform feedback control [15]. Similarly, [16] couples event rates to global scene dynamics to adaptively determine the decay of events.

When processing event streams in real-time, the high temporal resolution requires algorithms that can process events rapidly. It is desirable to keep the event rate within certain bounds, as higher event-rates can lead to processing delays. [17] maintains desirable event-rates through various event-dropping strategies at the hardware level. [18] tackles process overflow due to high event rates by modulating event package sizes and controlled discarding of events depending on expectations of process times and hardware limitations. [19] uses a gain parameter to control the absolute delay of the system.

Tackling process overload through controlled discarding of events leads to task-relevant data loss. [20] adapts ON/OFF event thresholds in the device hardware to monitor and modulate event rates. The popular jAER software [21] exposes bias parameters along with underlying coarse and fine bias settings [22]. This opens up pathways to software-triggered adaptive biasing approaches, which couples various

information sources to infer environmental conditions and tune the device on the fly.

The most relevant prior work is [12]. They perform feedback control to regulate event-rates and noise by changing event threshold, pixel bandwidth, and refractory period. Section IV-D contains a detailed explanation of the various feedback controllers introduced in [12]. [23] evaluates how changes in device biases affect the output of a Prophesee event camera, using sharpness to assess event stream quality. [24] characterizes DVS performances to a set of quantifiable parameters and evaluates DVS responses based on parameters such as threshold, bandwidth, refractory period and scene illumination.

B. Event-based Visual Place Recognition

Visual Place Recognition (VPR) involves recognizing if current sensory data correspond to a previously visited location. This plays a crucial role in localization and navigation capabilities of mobile robotic systems [3]. Typically, VPR pipelines infer geo-locations of incoming query images by matching them with a set of geo-stamped reference images of known places. This addresses the revisiting problem and aids in loop-closures, making VPR a useful component of SLAM systems [25].

Consequently, VPR is an interesting downstream task for event-based sensing, and well suited for our experiments validating the effectiveness of our fast-and-slow feedback controller. In the remainder of this section, we provide an overview of prior works on event-based VPR.

[6] combines temporal windows of varying lengths through parallelized ensembles to obtain a high-performing VPR system at the cost of high computational requirements. [7] uses the popular Vector of Locally Aggregated Descriptors (VLAD) pooling mechanism to obtain image descriptors after reconstructing event edges. This idea has been taken further in [5] which introduced the first end-to-end deep-learning event-based VPR framework. [4] introduced a computationally efficient technique that relies on just a small subset of informative pixels. Another interesting approach is the combination of conventional frames with event streams for improved performance [11].

Moving towards more bio-inspired approaches, [9] developed a brain-inspired multimodal hybrid neural network that seeks to imitate the neurological processes of the brain for more efficient robot place recognition. Their model demonstrates the synergy between different sensory modalities and computational paradigms, offering insights into the development of advanced robotic systems. This approach is similar to [8] who also used a spiking neural network for processing the event stream. Similarly, [10] harnessed the power of neuromorphic computing, presenting a neuromorphic sequence learning model that uses event cameras to navigate through vegetative environments.

There are relatively few datasets publicly available that are suited for event-based VPR. The Brisbane-Event-Dataset [6] contains 6 traverses with varying time-of-day of an 8km long route captured with a DAVIS346 camera mounted on the

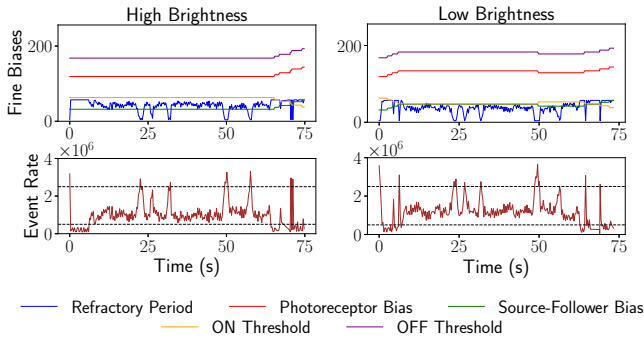


Fig. 2. **Dynamic bias adjustments for differing brightness conditions.** The top two graphs display the fine-tuning of bias parameters over time, while the bottom graphs show the corresponding event rates. The refractory period (blue) is constantly adjusted, while the other bias parameters are only adjusted if the refractory period changes alone do not suffice. On the left, under high brightness conditions, the biases are adjusted less frequently, resulting in a relatively stable event rate. Conversely, on the right, for low brightness conditions, there is more frequent fine-tuning of biases, which correlates with the more variable event rate observed. This illustrates how bias parameter management is critical for adapting to different luminosity levels and maintaining consistent event detection rates.

windshield of a car. It contains pseudo ground-truth from a GPS device. The QCR-Event-Dataset [4] contains 16 traverses of a 160m long route captured with a DAVIS346 camera mounted on a Clearpath Jackal robot, whereby the mounting position was varied between forward-, side-, or downward-facing. The ground truth for a subset of these traverses was manually annotated. The DDD20 dataset [26] is an enormous dataset containing recordings of over 4000km of driving (51h). It was mainly recorded for the steering prediction task, but a subset of the recordings could be used for the visual place recognition task. To the best of our knowledge, our QCR-Fast-and-Slow dataset is the largest indoor event-based VPR dataset to date (recordings totaling over 35km).

III. METHODOLOGY

We outline the foundational principles of event cameras for completeness in Section III-A. Our fast-and-slow feedback controller consists of two components that interact with each other (Figures 1 and 2): a fast controller that changes the refractory period aiming to keep the event rate within specified bounds (Section III-B), and a slow controller that changes the pixel bandwidth and event threshold if these event rate bounds are repeatedly exceeded (Section III-C).

A. Preliminaries: Event Camera Principles

Event cameras, or Dynamic Vision Sensors (DVS), capture the visual world differently from conventional cameras. They produce asynchronous events at the pixel level, triggered by changes in brightness that surpass a predefined threshold [27]. Each event $\mathbf{e} = (x, y, t, p) \in \mathcal{E}$ comprises the pixel coordinates $(x, y) \in \mathbb{N}^2$, a high-resolution timestamp $t \in \mathbb{R}^+$, and a polarity $p \in \{-1, 1\}$ indicating the direction of brightness change [1].

Ideally, event cameras would only generate events upon changes in the scene. However, in practice, they also produce noise events due to factors like junction leakage and parasitic




Fig. 3. **Refractory period.** Left: A short refractory period (black lines indicate the refractory period) allows more frequent firing of events (as indicated by the red and blue lines in the bottom plots), resulting in a signal that closely follows the input stimulus. Right: An extended refractory period limits the firing rate, leading to fewer events and a signal that may omit rapid fluctuations of the input stimulus.

photocurrents in bright conditions, and shot noise in low light [28], [29]. These unwanted events contribute to background activity, diminishing the signal-to-noise ratio [30]–[32].

Post-processing techniques exist to mitigate background noise [32], [33], but on-chip bias current generators also play a crucial role in adapting the camera to varying lighting conditions [22], [34], [35]. As detailed in the next section, our methodology automates the adjustment of these bias parameters, akin to tuning converter gain, exposure time, and ISO in traditional cameras.

B. Fast changes in refractory period

The refractory period in an event camera influences the duration before a pixel can generate another event, inversely affecting the event rate. Our control strategy uses this relationship to modulate the event rate in real time, targeting an optimal event rate range $[R^-, R^+]$ for enhanced place recognition performance. This is analogous to the refractory period in a biological nerve cell [36], [37]. As illustrated in Figure 3, higher refractory period settings lead to lower event rates in device output.

Let $\mathcal{E}(t, \tau)$ be the set of events that occurred within the time window $[t - \tau, t]$, where t' is the timestamp of the event:

$$\mathcal{E}(t, \tau) = \{\mathbf{e}(x, y, t', p) \in \mathcal{E} \mid t - \tau \leq t' < t\}. \quad (1)$$

Then, the event rate $R(t, \tau)$ (measured in events per second) at time t over the duration τ is defined as:

$$R(t, \tau) = \frac{|\mathcal{E}(t, \tau)|}{\tau}. \quad (2)$$

We define a linear mapping $B_r = f(R)$ from the event rate $R \in [R^-, R^+]$ to the refractory period $B_r \in [B_r^-, B_r^+]$. For the current event rate R , we obtain the refractory period bias as:

$$B_r = B_r^+ - \left(\frac{B_r^+ - B_r^-}{R^+ - R^-} \right) \cdot (R - R^-). \quad (3)$$

Further, we ensure the bias settings remain within the desired bounds $[B_r^-, B_r^+]$, resulting in the final bias setting \hat{B}_r :

$$\hat{B}_r = \begin{cases} B_r^+ & \text{if } B_r > B_r^+ \\ B_r^- & \text{if } B_r < B_r^- \\ B_r & \text{otherwise.} \end{cases} \quad (4)$$




Fig. 4. **Pixel bandwidth.** Left: A high bandwidth allows high frequency signal to pass through, which leads to higher event rates. Right: A low bandwidth filters high frequency components of the signal, leading to a low frequency signal and consequently lower event rate.

C. Slow changes in pixel bandwidth and event threshold

The pixel bandwidth determines the maximum permitted frequency of input signal, filtering away all higher frequency components [36]. Figure 4 illustrates how higher bandwidth increases the output event rate. Event thresholds concern with the minimum change in light intensity required to generate event spikes. Figure 5 illustrates how higher thresholds lower the output event rate.

We vary the pixel bandwidth and event thresholds to maintain event rates within desirable range $R \in [R^-, R^+]$. However, the instances of these changes have to be monitored because changing pixel bandwidth or event thresholds lead to momentarily disruptive event bursts in the output. Hence slow and simultaneous fixed-step adaptations ΔB_p and ΔB_e are performed in pixel bandwidth B_p and event thresholds B_e respectively if the event rate remains outside the desired bounds for the previous N consecutive time samples despite fast changes in the refractory period (Section III-B). Specifically, at time t :

$$B_p = \begin{cases} B_p + \Delta B_p, & \text{if } R_i < R_i^- \forall i \in (t - N, \dots, t] \\ B_p - \Delta B_p, & \text{if } R_i > R_i^+ \forall i \in (t - N, \dots, t], \end{cases} \quad (5)$$

and similarly for B_e noting a change in signs:

$$B_e = \begin{cases} B_e - \Delta B_e, & \text{if } R_i < R_i^- \forall i \in (t - N, \dots, t] \\ B_e + \Delta B_e, & \text{if } R_i > R_i^+ \forall i \in (t - N, \dots, t]. \end{cases} \quad (6)$$

IV. EXPERIMENTAL SETUP

We first describe the robotic platform and event camera processing pipeline in Section IV-A. Implementation details are provided in Section IV-B, and the underlying place recognition technique is briefly summarised in Section IV-C. Finally, the baseline techniques are described in Section IV-D.

A. Robot platform

We mount an iniVation DAVIS346 event camera [38] in the front of an AgileX Scout Mini mobile robot such that the camera faces in the direction of motion (Figure 6, left). We ensure that the camera focus remains constant in all our experiments, and use manual changes in camera aperture to simulate various environment brightness conditions (Figure 6, right). While no deliberate changes in the light source




Fig. 5. **Event threshold.** Left: A low event threshold setting (indicated by dashed horizontal lines) results in a high event rate (frequent events in the bottom plot), capturing most changes in the input signal. Right: A high threshold setting results in a lower event rate (much fewer events, as indicated by the sparse red and blue lines), only responding to larger variations in the input signal.

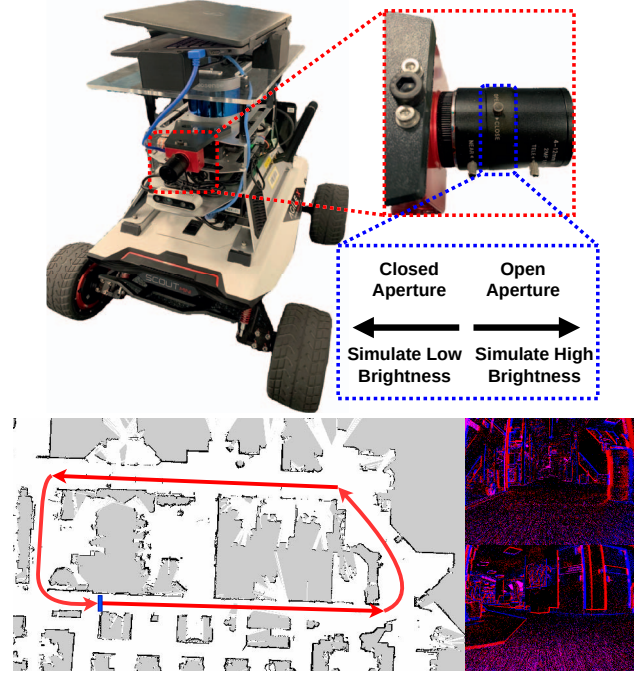


Fig. 6. **(Top) Robot system setup.** Left: Scout Mini platform equipped with various sensors, including a LiDAR to obtain pseudo ground truth and the DAVIS346 event camera mounted forward-facing. Right: Close-up of the DAVIS346 event camera, where the aperture setting can be manually adjusted. Altering the aperture simulates different brightness conditions: closing it simulates low-light environments, while opening it simulates high-brightness scenarios. **(Bottom) Map from the robot and frames from the event-data.** Left: The blue colored segment marks the starting position and red arrows indicate the traversal path. Right: Frames of accumulated event-data taken from one of the traversals.

positions were introduced, they might have occurred naturally as the dataset was collected over a prolonged time. Although Active Pixel Sensor (APS) grayscale images and Inertial Measurement Unit (IMU) data were also recorded along with the event data, they were not used in the experiments.

Pseudo ground truth robot positions in a fixed world frame are obtained using the SLAM Toolbox ROS package [39], which runs a full SLAM pipeline that takes the raw odometry from the wheel-encoders and LiDAR data as input to provide accurate ground truth position information in a global coordinate frame in a pre-mapped environment.

B. Closed-loop bias controller

Events are streamed over UDP for real-time computation of event rates. Our feedback bias controller uses event rates as feedback to control pixel bandwidth, event thresholds and refractory period in an online manner. Following up on Equation 1, the event rates are calculated and updated over time intervals of $\tau = 300ms$.

1) *Bias current initializations (in Amperes)*: Initial pixel bandwidth is set through photoreceptor bias $B_{pr} = 44.16pA$ and source-follower bias $B_{sf} = 1.48pA$. Initial event thresholds are set as ON threshold $B_{on} = 762.89nA$ and OFF threshold $B_{off} = 498.62pA$.

2) *Fast changes in refractory period*: As explained in Section III-B, we maintain a desired range of event rate $R \in [R^-, R^+]$ where $R^- = 5 \times 10^5 Hz$ and $R^+ = 2.5 \times 10^6 Hz$. We limit variations in refractory period to $B_r \in [759.37pA, 10.25nA]$.

3) *Slow changes in pixel bandwidth and event threshold*: Following the explanation in section III-C, we maintain the same desired range for event rate $R \in [R^-, R^+]$ as in Section IV-B.2. Slow changes are performed if $R \notin [R^-, R^+]$ for the last $N = 5$ consecutive time steps. Variations in pixel bandwidth $\Delta B_p = (\Delta B_{pr}, \Delta B_{sf})$ is performed as fixed step updates in photoreceptor bias $\Delta B_{pr} = 1.85pA$ and source-follower bias $\Delta B_{sf} = 0.23pA$. Similarly, variations in event threshold $\Delta B_e = (\Delta B_{on}, \Delta B_{off})$ is performed through fixed step updates in ON threshold $\Delta B_{on} = 60.55nA$ and OFF threshold $\Delta B_{off} = 14.84pA$.

C. Visual Place Recognition Pipeline

Event frames [13], [14] are generated by accumulating events into non-overlapping fixed-sized time-windows of $66ms$ duration. In brief, each pixel of the event frame contains the number of events that occurred within the time window. Post-processing involves filtering out event frames with event bursts and filtering out hot pixels, which produce noise in the form of high-frequency event spikes which are not caused by changes in light intensity [40]. Additionally, frames where the robot is stationary are discarded.

We use a previously proposed, simple visual place recognition pipeline to demonstrate our feedback controller [41]. To ensure robustness against appearance changes, the event frames are patch-normalized [42] with a patch-size of $8px$. Sum-of-absolute-differences (SAD) is performed between event frames to obtain distance matrices – the higher the SAD, the less similar the two event frames [42]. The event frame in the reference dataset that has the lowest SAD to the current query event frame is considered as the top match.

These top matches are compared with the ground truth robot positions in the map with a tolerance of $5m$. Recall R and precision P [3] are computed for each experiment as: $R = \frac{TP}{Q}$ and $P = \frac{TP}{TP+FP}$, where TP indicates the number of true positives, FP indicates the number of false positives and Q indicates the number of query images. The Recall@1 ($R@1$) metric is identical to the precision at 100% recall.

D. Baselines

1) *Default parameters*: The first baseline technique is to naively use the default biases that are used by jAER when it is started [21]. This corresponds to the use case where no manual and/or automated tuning is performed at all.

2) *Feedback control of refractory period to limit event rate*: The second baseline technique [12] is similar to our fast feedback controller and adapts the refractory period only. As a recap, a higher refractory periods reduces the event rate and vice versa. Specifically, the controller aims to keep event rate below a predefined upper bound by reducing the refractory period before returning to the system default. Note that this controller does not lead to spurious event bursts as only the refractory period is changed. Our *fast* adaptation method diverges from the above approach by tightly constraining the adaptation of the refractory period to a narrower range.

3) *Feedback control of pixel bandwidth to limit noise*: The output event rate tends to increase with an increase in pixel bandwidth. The noise event rate also increases with an increase in pixel bandwidth. This third baseline [12] limits the event noise rate in the output data by regulating the pixel bandwidth. Furthermore, the noise events are separated from the signal events using a background activity (BA) filter [40]. In contrast, our *slow* adaptation method performs simultaneous adjustments to pixel bandwidth and event threshold parameters on a case-by-case basis while monitoring the event rate.

4) *Feedback control of event threshold to bound noise*: The event threshold determines how much of a change is required in log input light intensity for an event spike to occur. The fourth baseline [12] exploits event thresholds to maintain event rates within desirable bounds. Unlike refractory period control, changes in the event threshold and pixel bandwidth lead to event bursts, making the device's output momentarily noisy. In contrast, similar to Section IV-D.3, our *slow* adaptation approach performs sparse and simultaneous changes in pixel bandwidth and event threshold parameters while monitoring the event rate.

V. RESULTS

Section V-A compares our closed-loop bias controller's performance against the baseline techniques, with the main results provided in Tables I and II. We also perform ablations studies on the optimal frequency of performing slow changes in Section V-B.1, and compare and analyse individual contributions of various decoupled components of our closed-loop bias controller pipeline in Section V-B.2.

A. Comparison to state-of-the-art

Tables I and II provide the Recall@1 metrics for varying query brightness conditions (low, medium, and high brightness) when the reference set was recorded under high brightness (Table I) or low brightness (Table II). Each experiment was performed with five repetitions. Generally speaking, in all conditions our method outperforms all baseline techniques or performs at least as well as the best performing baseline, and it performs the most consistently.

TABLE I

RECALL @ 1 COMPARISON FOR A REFERENCE SET RECORDED IN HIGH BRIGHTNESS CONDITIONS, WITH VARYING BRIGHTNESS CONDITIONS FOR THE QUERY SET. BASELINES INCLUDE DEFAULT PARAMETERS AND THREE BASELINE CONTROLLERS PROPOSED IN [12] (PxBw: CONTROL PIXEL BANDWIDTH, RfPr: CONTROL OF REFRACTORY PERIOD, PxTh: CONTROL PIXEL THRESHOLD; SEE SECTION IV-D). OUR METHOD CONSISTENTLY OUTPERFORMS OR PERFORMS AS WELL AS ALL OF THE BASELINE TECHNIQUES.

High Brightness Reference Set	Run #	Low Brightness (Maximum Variation)					Medium Brightness (Moderate Variation)					High Brightness (Minimum Variation)				
		Default	PxBw	RfPr	PxTh	Ours	Default	PxBw	RfPr	PxTh	Ours	Default	PxBw	RfPr	PxTh	Ours
		param.	[12]	[12]	[12]	Ours	param.	[12]	[12]	[12]	Ours	param.	[12]	[12]	[12]	Ours
	1	0.28	0.39	0.40	0.75	0.77	0.98	0.97	0.94	0.91	0.96	0.99	0.97	0.93	0.92	0.99
	2	0.47	0.71	0.62	0.86	0.79	0.88	0.98	0.98	0.93	0.96	0.95	0.95	0.97	0.95	0.94
	3	0.52	0.67	0.45	0.76	0.92	0.96	0.98	0.87	0.90	0.99	0.94	0.97	0.89	0.93	0.97
	4	0.44	0.33	0.49	0.84	0.90	0.98	0.95	0.85	0.91	0.92	0.97	0.96	0.85	0.86	0.93
	5	0.46	0.44	0.46	0.91	0.89	0.98	0.91	0.89	0.92	0.97	0.96	0.95	0.70	0.96	0.97
	Average	0.43	0.51	0.48	0.83	0.85	0.96	0.96	0.91	0.91	0.96	0.96	0.96	0.87	0.92	0.96
	Std. Dev.	0.09	0.17	0.08	0.07	0.07	0.04	0.03	0.05	0.01	0.03	0.02	0.01	0.10	0.04	0.02




Fig. 7. **Precision-Recall Curves for worst case performance (reference: high brightness).** The worst performing run from each method is chosen for plotting the PR curves. While our method showcases superior performance in comparison to all baselines when the query set is taken from low and medium brightness conditions, its performance remains on-par with the top three approaches when query set is from high brightness conditions.

TABLE II

RECALL @ 1 COMPARISON FOR A REFERENCE SET IN LOW BRIGHTNESS CONDITIONS, WITH VARYING BRIGHTNESS CONDITIONS FOR THE QUERY SET. SEE TABLE I FOR BASELINE TECHNIQUES.

Low Brightness Reference Set	Run #	Low Brightness (Minimum Variation)					Medium Brightness (Moderate Variation)					High Brightness (Maximum Variation)				
		Default	PxBw	RfPr	PxTh	Ours	Default	PxBw	RfPr	PxTh	Ours	Default	PxBw	RfPr	PxTh	Ours
		param.	[12]	[12]	[12]	Ours	param.	[12]	[12]	[12]	Ours	param.	[12]	[12]	[12]	Ours
	1	0.88	0.92	0.81	0.84	0.96	0.47	0.73	0.60	0.82	0.92	0.59	0.73	0.62	0.74	0.89
	2	0.88	0.99	0.83	0.91	0.92	0.73	0.97	0.81	0.60	0.91	0.73	0.94	0.80	0.56	0.97
	3	0.96	0.98	0.87	0.90	0.97	0.84	0.96	0.75	0.92	0.93	0.80	0.96	0.67	0.91	0.85
	4	0.85	0.94	0.72	0.94	0.97	0.61	0.92	0.47	0.80	0.91	0.65	0.92	0.59	0.78	0.88
	5	0.95	0.99	0.83	0.95	0.97	0.85	0.81	0.46	0.67	0.96	0.60	0.82	0.65	0.78	0.94
	Average	0.90	0.96	0.81	0.91	0.96	0.70	0.88	0.62	0.76	0.93	0.67	0.87	0.67	0.75	0.91
	Std. Dev.	0.05	0.03	0.05	0.04	0.02	0.16	0.11	0.16	0.13	0.02	0.09	0.10	0.08	0.13	0.05




Fig. 8. **Precision-Recall Curves for worst case performance (reference: low brightness).** The worst performing run from each method is chosen for plotting the PR curves. The PR curves demonstrate that our fast-and-slow bias controller performs better than all baseline techniques when considering the worst-case performance.

Specifically, in Table I (reference dataset: high brightness) one can observe that while the default bias parameters and the pixel bandwidth controller proposed in [12] perform as well as our technique in medium and high brightness

conditions (all $R@1=0.96$), they perform considerably worse in the challenging low query brightness conditions (absolute $R@1$ improvement of 42% and 34% respectively). This indicates that our proposed controller is more robust when

the appearance changes significantly (i.e. reference in high brightness and query in low brightness). We note that the pixel threshold controller consistently performs 2-5% worse than our proposed fast-and-slow controller. Another notable observation is that the lowest performance across all repetitions for our technique is almost always higher than the lowest performance of all baseline techniques, with the single exception being repetition 4 in the high brightness condition. This is further shown in Figure 7.

The same trend can be observed in Table II (reference dataset: low brightness), where the performance improvement over the second-best-performing pixel bandwidth controller [12] is an absolute 4% in the high brightness query condition, the query condition with the highest appearance change. Figure 8 demonstrates that the performance advantage in the worst-case run becomes even more pronounced, comparing a R@1 of 85% of our technique with 59% (default parameters), 73% (pixel bandwidth controller), 59% (refractory period controller) and 56% (pixel threshold controller).

B. Ablation Studies

1) *Frequency of slow changes:* The pixel bandwidth and event threshold bias settings are adjusted if the observed event rate remains outside desired range for the previous N times (Section III-C). Larger N lead to fewer number of slow changes made in a traversal, which is preferable since changes in pixel bandwidth and event threshold lead to event bursts. Filtering out events due to event bursts leads to a loss of useful data; and in the worst case, loss of localization. However, not tuning the pixel bandwidth and event threshold would also be destructive as those changes are required to keep the event rate within desired bounds (Figure 2).

We analyze the place recognition performance when running experiments with $N = \{2, 5, 7, 10\}$. Figure 9 (left) illustrates the average R@1 values obtained for three runs of each experiment. The most consistent performance is obtained when $N^* = 5$, including the highest performance in low-brightness query sets.

2) *Components:* We analyze individual components in our bias control pipeline to understand their contribution toward the overall pipeline performance. The variations are 1) to only initialize the event camera with the parameters outlined in Section IV-B.1 without changing the bias parameters during the run (constant), 2) only performing ‘fast’ changes in the refractory period, 3) only performing ‘slow’ adaptations in pixel bandwidth and event threshold, and 4) performing both slow and fast changes in refractory period.

Figure 9 (right) shows the R@1 performance from three iterations of each variation. When the query set is recorded in low brightness conditions, all variants significantly improve upon the constant baseline technique. In the medium and high brightness query conditions, the performance differences are relatively small; with our fast+slow controller achieving the highest performance in these conditions. Interestingly, only changing the refractory period (‘fast’) leads to a small performance decrease compared to the constant baseline.




Fig. 9. **Ablation studies.** **Left:** We investigate the frequency of slow changes. Generally, the performance differences are rather small, indicating good robustness of our technique to this hyperparameter. **Right:** Contributions of individual components of our online bias controller. The combination of slow+fast leads to the highest performance in all cases.

VI. DISCUSSION AND CONCLUSIONS

This paper has introduced a novel feedback control approach for event cameras, integrating fast-and-slow mechanisms to dynamically adjust bias parameters, thereby optimizing event stream processing for VPR tasks. Our methodology not only enhances the adaptability of event cameras to varying environmental conditions but also demonstrates that a feedback controller built specifically for the downstream task of VPR outperforms a general feedback controller in the same downstream task. Future work will extend this approach to longer trajectories and larger environmental variations, such as day-night transitions and diverse weather conditions. We also plan to integrate our feedback control system with Spiking Neural Networks (SNNs) [8]–[10], [43], [44]. SNNs, with their bio-inspired processing capabilities, could significantly benefit from the refined event streams produced by our method, potentially leading to more robust and efficient visual recognition systems.

Incorporating velocity signals from inertial measurement units (IMUs) to dynamically adjust the temporal windows in the VPR pipeline presents another pathway that could deal with the variability of motion in event-based vision. Unlike fixed temporal intervals, velocity-informed event frames could adjust the processing of event streams in real-time, better aligning with the movement dynamics of the scene or the observer. This concept could also extend to bias parameter adjustments, where velocity data could inform the tuning of event camera sensitivities, potentially improving the camera’s responsiveness to motion-induced event generation.

Another intriguing direction for future research is the incorporation of environmental light intensity measurements to refine the tuning of event camera biases. This method promises a finer, more context-sensitive calibration of camera settings, achieving an ideal equilibrium in event generation across various lighting scenarios. However, this strategy may necessitate a more sophisticated approach, as relying solely on raw light intensity might be overly simplistic. Incorporating some level of scene comprehension could offer a more nuanced optimization, potentially leading to an improvement in the signal-to-noise ratio and thereby enhancing the quality of the event stream for downstream tasks.

In conclusion, our fast-and-slow feedback controller represents a significant step forward in the practical application of event cameras for VPR and potentially other robotic vision

tasks. By addressing the dynamic nature of environmental conditions and the inherent limitations of current event camera technology, we have laid the groundwork for more adaptive, efficient, and effective event-based vision systems.

REFERENCES

- [1] G. Gallego, T. Delbrück, *et al.*, “Event-based vision: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 1, pp. 154–180, 2020.
- [2] S. Garg, T. Fischer, and M. Milford, “Where is your place, visual place recognition?” in *International Joint Conference on Artificial Intelligence*, 2021, pp. 4416–4425.
- [3] S. Schubert, P. Neubert, *et al.*, “Visual place recognition: A tutorial,” *IEEE Robotics & Automation Magazine*, 2023.
- [4] T. Fischer and M. Milford, “How many events do you need? Event-based visual place recognition using sparse but varying pixels,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 12 275–12 282, 2022.
- [5] D. Kong, Z. Fang, *et al.*, “Event-VPR: End-to-end weakly supervised deep network architecture for visual place recognition using event-based vision sensor,” *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–18, 2022.
- [6] T. Fischer and M. Milford, “Event-based visual place recognition with ensembles of temporal windows,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6924–6931, 2020.
- [7] A. J. Lee and A. Kim, “EventVLAD: Visual place recognition with reconstructed edges from event cameras,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 2247–2252.
- [8] H. Lee and H. Hwang, “Ev-reconnet: Visual place recognition using event camera with spiking neural networks,” *IEEE Sensors Journal*, vol. 23, no. 17, pp. 20 390–20 399, 2023.
- [9] F. Yu, Y. Wu, *et al.*, “Brain-inspired multimodal hybrid neural network for robot place recognition,” *Science Robotics*, vol. 8, no. 78, 2023.
- [10] L. Zhu, M. Mangan, and B. Webb, “Neuromorphic sequence learning with an event camera on routes through vegetation,” *Science Robotics*, vol. 8, no. 82, 2023.
- [11] K. Hou, D. Kong, *et al.*, “Fe-fusion-vpr: Attention-based multi-scale network architecture for visual place recognition by fusing frames and events,” *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3526–3533, 2023.
- [12] T. Delbruck, R. Graca, and M. Paluch, “Feedback control of event cameras,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1324–1332.
- [13] H. Rebecq, T. Horstschäfer, G. Gallego, and D. Scaramuzza, “Evo: A geometric approach to event-based 6-dof parallel tracking and mapping in real time,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 593–600, 2016.
- [14] A. I. Maqueda, A. Loquercio, *et al.*, “Event-based vision meets deep learning on steering prediction for self-driving cars,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5419–5427.
- [15] M. Liu and T. Delbruck, “Adaptive time-slice block-matching optical flow algorithm for dynamic vision sensors,” in *British Machine Vision Conference*, 2018.
- [16] U. M. Nunes, R. Benosman, and S.-H. Ieng, “Adaptive global decay process for event cameras,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9771–9780.
- [17] T. Finateau, A. Niwa, *et al.*, “5.10 a 1280× 720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86 μm pixels, 1.066 geps readout, programmable event-rate controller and compressive data-formatting pipeline,” in *IEEE International Solid-State Circuits Conference*, 2020, pp. 112–114.
- [18] R. Tapia, J. Martínez-de Dios, A. Gómez Eguíluz, and A. Ollero, “ASAP: adaptive transmission scheme for online processing of event-based algorithms,” *Autonomous Robots*, vol. 46, no. 8, pp. 879–892, 2022.
- [19] A. Glover, V. Vasco, and C. Bartolozzi, “A controlled-delay event camera framework for on-line robotics,” in *IEEE International Conference on Robotics and Automation*, 2018, pp. 2178–2183.
- [20] R. Berner, C. Brändli, and M. Zannoni, “Data rate control for event-based vision sensor,” U.S. Patent US10 715 750B2, August 30, 2022.
- [21] T. Delbruck *et al.*, “Frame-free dynamic digital vision,” in *International Symposium on Secure-Life Electronics, Advanced Electronics for Quality Life and Society*, 2008, pp. 21–26.
- [22] M. Yang, S.-C. Liu, C. Li, and T. Delbruck, “Addressable current reference array with 170db dynamic range,” in *IEEE International Symposium on Circuits and Systems*, 2012, pp. 3110–3113.
- [23] M. S. Dilmaghani, W. Shariff, *et al.*, “Control and evaluation of event cameras output sharpness via bias,” in *International Conference on Machine Vision*, 2023, pp. 455–462.
- [24] B. McReynolds, R. Graca, and T. Delbruck, “Experimental methods to predict dynamic vision sensor event camera performance,” *Optical Engineering*, vol. 61, no. 7, pp. 074 103–074 103, 2022.
- [25] S. P. Engelson and D. V. McDermott, “Error correction in mobile robot map learning,” in *IEEE International Conference on Robotics and Automation*, 1992, pp. 2555–2556.
- [26] Y. Hu, J. Binas, *et al.*, “DDD20 end-to-end event camera driving dataset: Fusing frames and events with deep learning for improved steering prediction,” in *IEEE International Conference on Intelligent Transportation Systems*, 2020.
- [27] S.-C. Liu, T. Delbruck, *et al.*, *Event-Based Neuromorphic Systems*. Wiley, 2015.
- [28] R. Graca and T. Delbruck, “Unraveling the paradox of intensity-dependent DVS pixel noise,” in *International Image Sensor Workshop*, 2021.
- [29] R. Graca, B. McReynolds, and T. Delbruck, “Optimal biasing and physical limits of DVS event noise,” in *International Image Sensor Workshop*, 2023.
- [30] G. Taverni, D. P. Moeys, *et al.*, “Front and back illuminated dynamic and active pixel vision sensors comparison,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 65, no. 5, pp. 677–681, 2018.
- [31] Y. Nozaki and T. Delbruck, “Temperature and parasitic photocurrent effects in dynamic vision sensors,” *IEEE Transactions on Electron Devices*, vol. 64, no. 8, pp. 3239–3245, 2017.
- [32] S. Guo and T. Delbruck, “Low cost and latency event camera background activity denoising,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 1, pp. 785–795, 2022.
- [33] B. McReynolds, R. Graca, and T. Delbruck, “Exploiting alternating DVS shot noise event pair statistics to reduce background activity,” in *International Image Sensor Workshop*, 2023.
- [34] T. Delbrück and A. V. Schaik, “Bias current generators with wide dynamic range,” *Analog Integrated Circuits and Signal Processing*, vol. 43, pp. 247–268, 2005.
- [35] T. Delbruck, R. Berner, P. Lichtsteiner, and C. Dualibe, “32-bit configurable bias current generator with sub-off-current capability,” in *IEEE International Symposium on Circuits and Systems*, 2010, pp. 1647–1650.
- [36] R. Graça, B. McReynolds, and T. Delbruck, “Shining light on the DVS pixel: A tutorial and discussion about biasing and optimization,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 4044–4052.
- [37] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- [38] R. Berner, C. Brandli, *et al.*, “A 240× 180 10mw 12us latency sparse-output vision sensor for mobile applications,” in *Symposium on VLSI Circuits*, 2013, pp. C186–C187.
- [39] S. Macenski and I. Jambrecic, “SLAM Toolbox: SLAM for the dynamic world,” *Journal of Open Source Software*, vol. 6, no. 61, p. 2783, 2021.
- [40] Y. Hu, S.-C. Liu, and T. Delbruck, “v2e: From video frames to realistic DVS events,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1312–1321.
- [41] M. Milford, H. Kim, S. Leutenegger, and A. Davison, “Towards Visual SLAM with Event-based Cameras,” in *Robotics: Science and Systems Workshops*, 2015.
- [42] M. J. Milford and G. F. Wyeth, “SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights,” in *IEEE International Conference on Robotics and Automation*, 2012, pp. 1643–1649.
- [43] S. Hussaini, M. Milford, and T. Fischer, “Ensembles of compact, region-specific & regularized spiking neural networks for scalable place recognition,” in *IEEE International Conference on Robotics and Automation*, 2023, pp. 4200–4207.
- [44] A. D. Hines, P. G. Stratton, M. Milford, and T. Fischer, “VPRTempo: A fast temporally encoded spiking neural network for visual place recognition,” in *IEEE International Conference on Robotics and Automation*, 2024.