This CVPR Workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

Probabilistic Online Event Downsampling

Andreu Girbau-Xalabarder Denso IT Laboratory andreu@mail.d-itlab.co.jp Jun Nagata Denso IT Laboratory Shinichi Sumiyoshi Denso IT Laboratory

Abstract

Event cameras capture scene changes asynchronously on a per-pixel basis, enabling extremely high temporal resolution. However, this advantage comes at the cost of high bandwidth, memory, and computational demands. To address this, prior work has explored event downsampling, but most approaches rely on fixed heuristics or threshold-based strategies, limiting their adaptability. Instead, we propose a probabilistic framework, POLED, that models event importance through an event-importance probability density function (ePDF), which can be arbitrarily defined and adapted to different applications. Our approach operates in a purely online setting, estimating event importance on-thefly from raw event streams, enabling scene-specific adaptation. Additionally, we introduce zero-shot event downsampling, where downsampled events must remain usable for models trained on the original event stream, without taskspecific adaptation. We design a contour-preserving ePDF that prioritizes structurally important events and evaluate our method across four datasets and tasks-object classification, image interpolation, surface normal estimation, and object detection-demonstrating that intelligent sampling is crucial for maintaining performance under event-budget constraints.

1. Introduction

Event cameras, also known as neuromorphic cameras, differ fundamentally from conventional RGB sensors. Instead of capturing images at fixed frame rates, they asynchronously detect changes in intensity, producing a stream of events. Each event encodes its pixel location, precise timestamp (microsecond resolution), and polarity, indicating whether the intensity increased (ON) or decreased (OFF). Key advantages of event cameras include a high dynamic range (up to 140 dB), low power consumption, no motion blur, high temporal resolution, and low latency. These properties make them well-suited for tasks such as video deblurring, high-speed object detection and tracking, and super-slowmotion video synthesis.



Figure 1. *Cup* scene from N-Caltech101. By selectively retaining important events, POLED improves efficiency while maintaining task-relevant details, unlike naive uniform sampling. Top: Original event stream. Middle: Uniformly downsampled events. Bottom: Downsampled events with POLED, which adapts to the scene to preserve structured information, making contours more explicit.

However, the sheer volume of events they generate poses significant challenges. Processing, storing, and transmitting event streams efficiently remains a major bottleneck, particularly in embedded systems or bandwidth-constrained applications like ADAS (Advanced Driver Assistance Systems), or scenarios with limited data transmission capabilities. To address this, event downsampling has emerged as a key strategy, aiming to reduce event rates while preserving the most relevant information in the scene.

To address this, most existing methods reduce the number of events by downscaling spatial dimensions or aggregating events temporally, typically using predefined thresholds. While effective in reducing data, these methods lack adaptability and do not consider the importance of individual events.

We take a different approach, formulating event down-

sampling as a stochastic process. Given an estimated eventimportance probability density function (ePDF), incoming events are sampled probabilistically based on their likelihood of belonging to the distribution. This framework is adaptive, allowing the downsampling strategy to adjust dynamically to scene statistics or to the task at hand by modifying the ePDF on the fly. To mimic realistic scenarioswhere a camera is recording and bandwidth and computational resources are limited-, we adopt an online downsampling perspective, where only past and present events are available when making decisions. In contrast, offline methods can access future events, at the cost of increased memory usage and latency.

Additionally, most prior event downsampling methods have been evaluated on simplistic classification tasks, making it unclear how well they generalize to real-world applications. We argue that a truly effective downsampling strategy should be tested across multiple datasets and tasks, and evaluate our method on classification (NCaltech101 [15]), surface normals estimation (ESFP [14]), frame interpolation (Timelens [17]), and object detection in automotive scenario for Gen1 Automotive Detection dataset [6].

Furthermore, we emphasize zero-shot event downsampling, where the downsampled event stream is directly fed into a pre-trained model without task-specific adaptation. This ensures broad applicability in scenarios where retraining is impractical or infeasible due to limited accessibility, computational resources, time, or data constraints. To complement this, we also investigate the effects of retraining the models on downsampled events, providing insights into how adaptation impacts performance when retraining is an option.

In summary, our contributions are as follows:

- We present a probabilistic online event downsampling framework (POLED), where events are sampled based on an arbitrarily defined event-importance probability density function (ePDF), allowing for adaptable and taskagnostic event reduction.
- We introduce an ePDF that prioritizes the preservation of edges, as well as two baselines.
- We evaluate our approach in a zero-shot setting, where downsampled events are directly fed into pre-trained models without task-specific adaptation, and validate our method on four diverse tasks—classification (N-Caltech101), object detection (Gen1 Automotive Detection), frame interpolation (Timelens), and surface normals estimation (ESFP)—demonstrating that intelligent event sampling is key to preserving performance in eventbudget scenarios.

2. Related work

Modern event cameras produce an overwhelming number of events, posing significant challenges in bandwidth, computation, and memory. Efficient event stream management is therefore crucial, with downsampling emerging as a key strategy to reduce computational and bandwidth demands.

Early works explored event downsampling in spatial and temporal dimensions by scaling event coordinates and timestamps, adapting the sampling strategy to the dataset [5]. Later approaches refined this by integrating events over space and time using a counting strategy with refractory periods [9].

Other methods take inspiration from biological neurons, reducing events based on the activation of multiple sensory unit layers [2]. Spiking Neural Networks (SNNs) have also been explored for downsampling, leveraging neuromorphic processing to optimize event retention [10, 12, 13, 16]. Adaptive compression strategies, such as Huffman encoding that dynamically adjusts to bandwidth constraints, have also been proposed [4], along with pre-processing techniques that use non-uniform spatial sampling via 3D grids [3].

Beyond computational efficiency, research has also examined how downsampling affects human perception of event streams. Studies have compared basic temporal and spatial filtering with more advanced SNN-based approaches to assess how sparsification impacts interpretability [11, 12].

Despite these advances, most existing methods rely on fixed heuristics or task-specific optimizations, limiting their adaptability across different applications.

In contrast to existing approaches, we formulate event downsampling as an online stochastic process, where events are sampled based on their likelihood of belonging to an estimated event distribution. This enables adaptive selection based on event importance and scene statistics, rather than relying on fixed heuristics or thresholds.

The closest work to ours is [1], which uniformly samples events each epoch to train a CNN, studying the effects of event reduction and its interaction with CNN training parameters. However, their approach does not consider event importance or zero-shot applicability. Instead, we focus on the downsampling technique itself, evaluating its performance across independent tasks and models trained on the original event stream. Additionally, we investigate the effects of retraining with downsampled events, reaching similar conclusions to [1] while emphasizing the role of intelligent sampling.

We propose importance-based downsampling using an event-importance probability density function (ePDF), which can be arbitrarily defined and adapted to different settings. To make this framework broadly applicable, we introduce a generic formulation and processing pipeline, namely POLED, capable of handling any valid ePDF. In this work, we present a Poisson-based ePDF designed to prioritize contour preservation, under the premise that contour-related events are more relevant for solving diverse tasks. Furthermore, we approach event downsampling from a purely online perspective, making decisions based only on past and present information, simulating a real-time scenario where future events are unknown.

Finally, prior work has largely focused on simple classification datasets or introduced metrics favoring classification-based evaluation. To provide a more comprehensive assessment, we evaluate our method on four challenging datasets, covering classification, frame interpolation for super-slow-motion video generation, surface normal estimation, and object detection in an automotive setting.

3. Method

In this section, we formalize event downsampling as a stochastic process, present two baseline approaches, one deterministic and one based on uniform sampling, and introduce an adaptive Poisson-based event-importance probability density function (ePDF) designed to prioritize events near edges, which we hypothesize to be generally important across different tasks.

Overview of the proposed downsampling pipeline. The original event stream \mathcal{E} is processed using one of three downsampling techniques: Deterministic, Uniform, or POLED (our probabilistic approach). The resulting downsampled event stream \mathcal{E}_{ω} is then fed into a downstream task model. The downsampling process is designed to retain the most relevant events while reducing data volume for computational efficiency.

3.1. Problem formulation

Event cameras operate asynchronously, responding independently to changes in light intensity at each pixel. A comparator monitors intensity variations relative to a reference level: if the increase surpasses a predefined threshold, an ON event is generated; if the decrease exceeds the threshold, an OFF event is triggered. No event is produced when intensity remains unchanged.

Each event is represented as:

$$e_k = (\mathbf{x}_k, t_k, p_k), \ k \in \mathbb{N}^+ \tag{1}$$

where $\mathbf{x}_k = [x_k, y_k]$ is the pixel location, t_k denotes the timestamp at which the k-th event occurs, and $p_k \in \{\text{ON}, \text{OFF}\}$ specifies the polarity of the change.

A sequence of events, or an event stream, is defined as:

$$\mathcal{E} = \{e_k\}_{k=1}^N \tag{2}$$

We define event downsampling as the problem of finding a mapping Φ_{α} such that:

$$\Phi_{\alpha}: \mathcal{E} \mapsto \mathcal{E}^{\alpha}, \quad \mathcal{E}^{\alpha} \subseteq \mathcal{E}$$
(3)

$$\Phi_{\alpha}(e_k) = \begin{cases} e_k, & \text{if condition holds} \\ \emptyset, & \text{otherwise} \end{cases}$$
(4)

where \mathcal{E}^{α} is the downsampled event stream, and $\alpha \in (0, 1)$ denotes the downsampling ratio with respect to the original event set \mathcal{E} . Note that we define \mathcal{E}^{α} as a subset of the original events \mathcal{E} , meaning that no new events are generated -nor are existing ones modified— only a selection of events from the original set is removed. This is in contrast to previous works, where event information may be altered, aggregated, or even synthesized.

3.2. Online Event Downsampling

To determine whether an incoming event e_k is accepted or rejected, we explicitly adopt an *online* downsampling approach, meaning that only past and present events in $\mathcal{E}_{1\rightarrow k}$ are available at any given time. In contrast, an *offline* setting would allow access to future events, i.e., decisions regarding e_k could be influenced by information in $\mathcal{E}_{1\rightarrow l}$, where l > k. A practical example of offline downsampling occurs in methods that aggregate events within a temporal window extending beyond the current event e_k .

Enforcing event rate constraints. To ensure that downsampling adheres to a predefined event budget—modeled as the downsampling ratio α —we introduce an event capping mechanism. This mechanism enforces the constraint that the number of retained events must not exceed the allowed proportion α of the original stream, simulating real-world constrains on bandwidth or computational resources.

Formally, given the incoming event e_k we define the downsampled event relation as:

$$r = \frac{|\mathcal{E}^{\alpha}_{0 \to k-1}|}{|\mathcal{E}_{0 \to k-1}|} \tag{5}$$

where $|\cdot|$ is the dimension of the set. If at any point $r > \alpha$, the event e_k is dropped, even if it would have otherwise been accepted by the sampling method.

3.3. Baselines

To assess the effectiveness of our probabilistic event downsampling approach, we introduce two baseline strategies: deterministic and uniform downsampling.

3.3.1. Deterministic Downsampling

The first approach we explore is a deterministic downsampling strategy, where small bursts of events are retained based on their timestamps within predefined temporal acceptance windows.

To achieve this, we define a small sliding temporal window T_w and partition it into two sub-intervals: an acceptance window T_a and a rejection window T_r , where $T_a =$



Figure 2. Overview of the proposed downsampling pipeline. To simplify visualization, the downsampling process is illustrated for a single pixel over time rather than across the full spatial-temporal event grid. The original event stream \mathcal{E} is processed using one of three introduced downsampling techniques: Deterministic, Uniform, or POLED (our probabilistic approach). Deterministic applies fixed temporal acceptance windows, Uniform assigns a constant acceptance probability to each event, and POLED estimates an event probability density function (ePDF) to sample events based on importance. The resulting downsampled event \mathcal{E}^{α} stream is then used for downstream tasks, preserving the most relevant events while reducing data volume.

 $T_w \cdot \alpha$ and $T_r = T_w - T_a$. In practice, we use $T_w = 0.1$ milliseconds. The downsampling process is anchored at an initial reference time T_0 , which shifts as $T_0 = nT_w$, defining the start of the *n*-th temporal window. An event $e_k = (\mathbf{x}_k, t_k, p_k)$ is retained if its timestamp falls within the acceptance window, defined as :

$$\Phi_{\alpha}(e_k) = \begin{cases} e_k, & \text{if } T_0 < t_k \le T_0 + T_a \\ \emptyset, & \text{if } T_0 + T_a < t_k \le T_0 + T_w \end{cases}$$
(6)

3.3.2. Uniform Downsampling

The second baseline follows a uniform random sampling strategy, where events are accepted independently with a fixed probability α . Formally:

$$\Phi_{\alpha}(e_k) = \begin{cases} e_k, & \text{with probability } \alpha \\ \emptyset, & \text{with probability } 1 - \alpha \end{cases}$$
(7)

This baseline represents the simplest stochastic approach to event downsampling and serves as a reference point for evaluating our more structured probabilistic methods.

3.4. Probabilistic Online Event Downsampling

We now introduce Probabilistic Online Event Downsampling (POLED), a flexible framework for event reduction based on probabilistic sampling. Instead of applying fixed rules or uniform selection, this approach models event importance using an estimated event Probability Density Function (ePDF). Given any valid ePDF $f(\cdot)$, each incoming event e_k is accepted or rejected based on its likelihood of belonging to the estimated event distribution.

To achieve this, we define the mapping Φ_{α} as:

$$\Phi_{\alpha}(e_k) = \begin{cases} e_k, & \text{with probability } \mathsf{P}_{\mathcal{E}}(\mathbf{x}_k) \\ \emptyset, & \text{with probability } 1 - \mathsf{P}_{\mathcal{E}}(\mathbf{x}_k) \end{cases}$$
(8)

where $P_{\mathcal{E}}(\mathbf{x}_k)$ represents the probability of accepting an event $e_k = (\mathbf{x}_k, t_k, p_k)$ derived from the ePDF.

Pipeline. The downsampling pipeline is as follows. First, we estimate the event probability density function (ePDF), denoted as $f(\mathcal{E}_{1\rightarrow k})$. The way to estimate this PDF depends on how the event importance is defined. In this work, we prioritize events clustered around edges (Sec. 3.4.1), as they are generally more informative across different contexts.

To focus on recent changes in the scene, we estimate the ePDF within a sliding temporal window of size T, using events from the interval $f(\mathcal{E}_{(n-1)T \to nT})$. Following this estimation, we evaluate the likelihood of an incoming event e_k for $t_k > nT$, belonging to the estimated distribution $f_{\mathcal{E}}$. For clarity, we drop the explicit temporal notation and refer to the estimated event distribution as $f_{\mathcal{E}}$.

To convert the likelihood of an event e_k into an acceptance probability, we define a mapping function S that transforms the estimated likelihood $f_{\mathcal{E}}$ into a probability score:

$$\mathbf{P}_{\mathcal{E}}(\mathbf{x}_k) = \mathcal{S}(f_{\mathcal{E}}(\mathbf{x}_k)) \tag{9}$$

The resulting probability $P_{\mathcal{E}}(\mathbf{x}_k)$ is then used as the acceptance probability in a Bernoulli sampling process.

To ensure that $P_{\mathcal{E}}(\mathbf{x}_k)$ is in the valid range [0, 1], the mapping function \mathcal{S} is defined as:

$$\mathcal{S}(\mathbf{x}) = \sigma \left(f_{\mathcal{E}}^*(\mathbf{x}) + \left(\alpha - \overline{f_{\mathcal{E}}^*}(\mathbf{x}) \right); \theta \right)$$
(10)

where $f_{\mathcal{E}}^*$ is the normalized form of $f_{\mathcal{E}}$ by the min-max normalization, its mean value displaced to match the downsampling ratio α , and its values rescaled by the sigmoid function σ to scale it to a probability range [0, 1].

We define the θ parameters to scale and shift the sigmoid function, ensuring that no event has a probability of exactly

0 (allowing all events a chance of being selected) or exactly 1 (allowing even the most likely events to be removed occasionally).

$$\sigma(x;\theta) = \frac{1}{1 + e^{-\theta_1(x-\theta_2)}} \tag{11}$$

In this work, we set $\theta = [\theta_1, \theta_2] = [5, 0.5]$, ensuring that mid-range values follow a near-linear probability scaling, while extreme values are adjusted to prevent hard acceptance or rejection.

On the importance of events. Defining a generic notion of event importance is inherently ill-posed, as relevance depends on the specific objective at hand. Ideally, our goal is to define a $f_{\mathcal{E}}$ that is task-agnostic and generalizable across diverse applications. To this end, in the next section, we introduce an approach that prioritizes edge events, as they tend to be informative across a wide range of vision tasks. However, edge-based downsampling alone may overlook task-specific event distributions. For instance, in object detection, events corresponding to objects of interest—such as cars or pedestrians—are more valuable than background events.

To bridge this gap, we propose incorporating taskspecific prior knowledge into the event downsampling process. This is achieved by integrating a prior distribution that modulates the estimated ePDF, introducing a bias that ensures the selection process aligns with task-relevant event patterns. Whenever a task-specific prior is available, we modify Eq. (10) (omitting x for clarity) as:

$$\mathcal{S}(\mathbf{x}) = \sigma \left((f_{\mathcal{E}} \cdot f_p)^* + (\alpha - \overline{(f_{\mathcal{E}} \cdot f_p)^*}) ; \theta \right)$$
(12)

where f_p is the task-specific prior distribution, and $(f_{\mathcal{E}} \cdot f_p)^*$ is the min-max normalized form of the product of the two distributions. The prior f_p can be derived from dataset statistics, heuristically defined, or learned from data, depending on the task requirements.

3.4.1. Poisson-based Event PDF

While event significance varies across applications, we argue that events that correspond to edges are generally informative across diverse tasks. Building on this intuition, we introduce a Poisson-based method for estimating the ePDF, designed to prioritize edge-preserving downsampling.

Poisson distributions are widely used to model the occurrence of discrete events in space and time, making them well-suited for capturing the sparsity and stochastic nature of event data. Given a rate parameter λ representing the expected number of events in a given region, the probability mass function (PMF) of a Poisson-distributed variable k is:

$$P(k;\lambda) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{13}$$

In our case, we estimate an event probability density function (ePDF) using a spatial Poisson process, where λ is dynamically computed based on pixel-wise event densities within a temporal window T. This allows us to prioritize edge-preserving downsampling by assigning higher probabilities to regions with structured event activity.

To model event importance, we define the ePDF as the probability of at least one event occurring at a given spatial location x within the temporal window [(n-1)T, nT]:

$$f_{\mathcal{E}}(\mathbf{x}) = 1 - P(0; \lambda(\mathbf{x})) = 1 - e^{-\lambda(\mathbf{x})}$$
(14)

where $\lambda(\mathbf{x})$ represents the event density at \mathbf{x} over the temporal window T.

Finally, we convert $f_{\mathcal{E}}$ into an acceptance probability using Eq. 9, ensuring that the sampling process prioritizes edge events while dynamically adapting to scene variations.

4. Experiments

We evaluate our proposed probabilistic event downsampling method (POLED) across four diverse datasets, each representing a distinct vision task. Our goal is to assess both the generalization capabilities of our approach—particularly the Poisson-based ePDF defined in Sec. 3.4.1, named as POLED-Poi in the results—and its impact on downstream performance across various applications.

Our experiments cover the following datasets and tasks:

- N-Caltech101 [15] for object classification, using the model from [7].
- **Gen1** [6] for object detection in driving scenarios, leveraging the approach from [8].
- **Timelens** [17] for frame interpolation, evaluated with its original model.
- **ESFP** [14] for surface normal estimation, with the method introduced in the same work.

These datasets present a diverse set of challenges, allowing us to evaluate the robustness of our downsampling strategy beyond standard classification benchmarks. We compare our method against our deterministic and uniform downsampling baselines, introduced in Sec. 3.3, as well as a recently proposed event downsampling approach EDS [9], analyzing performance in both zero-shot scenarios and retraining settings when applicable.

EDS [9]. To provide a broader comparison, we include EDS(Event DownSampling), a recent method that simulates a lower-resolution event camera by reducing both spatial and temporal dimensions.

While originally designed for a different purpose, EDS offers an interesting point of comparison: it represents

	Events	N-Caltech101 [15]	Gen1 [6]	Timelens [17]		ESFP-Real [14]
		Accuracy ↑	AP50 ↑	Close	Far	AE↓
Original	100%	84.4	68.8	32.47±4.6	32.04±1.8	28.50
EDS [9]	5%	5.2	0.4	29.58±5.0	24.10±1.6	30.05
Deterministic	5%	13.4	7.4	29.75 ± 5.0	24.48±1.7	29.92
Uniform	5%	12.9	8.2	30.23±4.6	26.03±2.1	29.90
POLED-Poi	5%	27.4	10.3	30.39±4.6	$26.87{\pm}2.4$	29.84
EDS	10%	23.7	2.6	29.75±5.0	24.47±1.6	29.76
Deterministic	10%	41.5	18.7	29.96 ± 5.0	24.66±1.7	29.51
Uniform	10%	40.3	19.9	30.89 ± 4.5	27.45 ± 2.0	29.52
POLED-Poi	10%	55.2	24.3	30.99±4.5	$\textbf{28.23}{\scriptstyle\pm2.3}$	29.47
EDS	50%	77.9	61.6	30.97±4.6	27.37±2.3	28.57
Deterministic	50%	82.7	60.5	31.31±4.6	27.90 ± 2.2	28.46
Uniform	50%	82.9	61.1	32.05±4.5	30.80 ± 2.0	28.45
POLED-Poi	50%	83.3	61.8	$32.07{\pm}4.5$	$30.82{\pm}2.1$	28.46

Table 1. Results on a zero-shot setting of the four approaches for different levels of downsampling. For Uniform and POLED-Poi average of several runs. Best in **bold**.

an alternative downsampling strategy—one that naively reduces resolution instead of selecting events based on importance. By evaluating its performance in event-budget scenarios, we can empirically assess whether simple spatiotemporal downsampling is sufficient, or if intelligent event selection is necessary to retain task-critical information.

4.1. Implementation details

We evaluate our approach using existing methods tailored to each dataset and task. For N-Caltech101 and ESFP-Real, we train the classifiers from [7] and [14] respectively on the original event streams, as well as on the downsampled event streams (Sec. 4.4). For Timelens and Gen1, we use the pre-trained models and weights provided in [17] and [8], respectively. Both models are trained on the original event streams and zero-shot tested on the downsampled event streams.

To enforce an online processing constraint, we sequentially load datasets and process events one-by-one. All downsampling methods operate with a temporal window of 6 ms, approximately corresponding to 160 FPS. Downsampling is performed entirely on the CPU, while all methodrelated experiments were conducted on a single NVIDIA Quadro RTX 8000 GPU.

4.2. Introductory results

Table 1 presents the zero-shot performance of the different downsampling techniques accross all datasets and tasks. We report the original performance using 100% of the event stream, along with the zero-shot performance at 5%, 10%, and 50% from the original events.

At low-event budget scenarios (5% and 10%), we observe that POLED-Poi, the Poisson-based ePDF method, outperforms the deterministic and uniform downsampling

baselines in all datasets and tasks, by a considerable margin in some cases such as in N-Caltech101 and Gen1. While events naturally cluster around edges, our results suggest that further emphasizing edge information is crucial for maintaining performance, particularly in low-event regimes. At higher event budgets (50%), performance differences among methods diminish, as essential taskrelevant events are naturally retained due to their inherent concentration around contours.

Among all techinques, EDS [9] consistently performs the worst across all datasets and tasks. We attribute this to three key factors. First, the method applies spatial downsampling, reducing resolution before passing the events to the downstream model. Since our evaluation setup assumes a method-agnostic downsampling process, events must be rescaled back to the original resolution, potentially introducing artifacts and loss of fine-grained details. Second, EDS integrates events within fixed temporal intervals, which may not capture the temporal dynamics of the events accurately, particularly in low-event regimes. Finally, the combined effects of spatial and temporal aggregation may alter the statistical distribution of the event stream, introducing a distribution shift that impacts downstream performance.

4.3. Event importance

As previously discussed, not all events contribute equally to solving a task. While event importance is inherently task-dependent, our results suggest that events concentrated around object contours generally play a crucial role across a wide range of tasks. However, not all event concentrations are equally important. While in Timelens most contours contribute equally to solve the task, in the Gen1 dataset for



Figure 3. Comparison of the distribution of a) bounding boxes in the training set and b) raw event distribution in the Gen1 dataset.

	Events	N-Caltech1	01 (Accuracy ↑)	Gen1 (AP50 ↑)	
		w/o prior	w/ prior	w/o prior	w/ prior
Original	100%	84.4	-	68.8	-
POLED-Poi	5%	25.1	27.4 (+2.3)	9.3	10.3 (+1.0)
POLED-Poi	10%	52.0	55.2 (+3.2)	22.1	24.3 (+2.2)
POLED-Poi	50%	83.1	83.3 (+0.2)	60.9	61.8 (+0.9)

Table 2. Zero-shot performance of POLED (Poisson-based) downsampling on N-Caltech101 and Gen1, comparing results with and without a task-specific prior.

car and pedestrian detection, events concentrated around vehicle edges are far more critical than those outlining sidewalks.

We illustrate this in Figure 3. Comparing the distribution of training set bounding boxes (Fig.3a) with the raw event distribution (Fig.3b), we observe a significant mismatch between event locations and actual detections. This suggests that POLED-Poi, which prioritizes high-density regions, may oversample events in areas like trees on sidewalks rather than focusing on the most relevant contours.

To address this misalignment, we incorporate the training set bounding box distribution as a prior, using Eq. (12) to modulate the ePDF and guide sampling toward the taskrelevant regions. Also, in N-Caltech101, we use a 2D Gaussian distribution as the objects of interest tend to concentrate at the center of the image. Table 2 presents the performance of POLED-Poi with and without a task-specific prior, highlighting the performance gains achieved by incorporating prior knowledge.

4.4. Task specific re-training

While zero-shot event downsampling is ideal for real-world deployment—where re-training may not always be feasible—the performance gap between downsampled and fullevent models remains significant. To better understand the upper bound of downsampling effectiveness, we retrain models on the downsampled event streams for N-Caltech101 and ESFP, allowing them to fully adapt to the reduced event distribution.



Figure 4. Accuracy (\uparrow) vs. sampling rate on N-Caltech101 [15] for zero-shot (dashed) and re-trained (solid) models at different downsampling levels. The purple star indicates the original performance. The x-axis is in logarithmic scale.



Figure 5. Mean Average Error (\downarrow) vs. sampling rate on ESFP-Real [14] for zero-shot (dashed) and re-trained (solid) models at different downsampling levels. The purple star indicates the original performance. The x-axis is in logarithmic scale.

Figures 4 and 5 compare zero-shot (dashed) vs. retrained (solid) performance across different downsampling rates. Notably, in both cases, the amount of events required to solve the task (i.e., achieving full-event performance) can be reduced dramatically: on N-Caltech101, accuracy recovers almost fully with as little as 5% of the original events, and on ESFP, nearly identical performance is achieved with

Method	Total (ms/Kev)	PDF (ms/Kev)	Eval (ms/Kev)
Deterministic	4.90	-	0.45
Uniform	4.56	-	0.14
POLED-Poi	6.55	1.72	0.14

Table 3. Computational time per thousand events (ms/Kev) for different downsampling approaches. *Total* represents the overall downsampling time, *PDF* accounts for computing the event probability density function, and *Eval* measures the cost of the acceptance or rejection decision.

just 0.1% of the original event stream. This aligns with findings from [1], which also demonstrate that models can adapt to highly reduced event streams with appropriate training.

It is also important to note that the assumption that edgerelated events are crucial holds in N-Caltech101 (Fig. 4), where the Poisson-based ePDF outperforms the baselines when both training and testing on downsampled events. Additionally, at higher event retention rates (90%), we observe that downsampled models exceed the original full-event performance, suggesting that our approach introduces a denoising effect, even though it was not explicitly designed for that purpose.

4.5. Run time

Computational efficiency is critical for event-based processing, particularly in real-time applications. Table 3 reports the runtime of our proposed POLED-Poi compared to the deterministic and uniform baselines. As expected, estimating the ePDF introduces additional overhead, making POLED-Poi approximately 1.4 times slower than the baselines. However, our current implementation is written in pure Python with a single blocking thread, representing a worst-case scenario for all techniques. Optimizing the implementation in lower-level languages such as C++ or CUDA would significantly improve runtime across all downsampling methods. Additionally, POLED can benefit from parallelization, as the ePDF computation can be executed in a separate thread, reducing the impact of its overhead.

4.6. Qualitative results

We provide visual insights on how downsampling affects the tasks by showing qualitative results in Figure 6. We highlight the contour preservation of the Poisson-based ePDF (POLED-Poi) estimation in contrast to the Deterministic or Uniform approaches.

5. Limitations and future work

While our proposed event PDF estimation and further downsampling method is generic and can be successfully applied in very diverse tasks, we find that generically deter-



Figure 6. Detailed output of the frame interpolation method [7] by using the event stream downsampled to maintain 10% of the original events. Even at low-event regimes, POLED-Poi is able to keep the contour information (windows).

mine what makes an event *relevant* is an ill-posed problem, as each event has its own relevance depending on the downstream task. An example we found in our research is based on the Gen1 dataset, where the relevant events to solve that concrete task (car and pedestrian detection) would not be the same as the events needed to solve another task (e.g. road lines segmentation).

Another limitation is on the computation time, as the two baselines, Deterministic and Uniform, are much simpler, thus faster, than the proposed PDF estimators. This limitation can be addressed with parallelization and implementation on more appropriate coding languages, and will be considered as future work.

6. Conclusion

We introduced POLED, a probabilistic framework for event downsampling that models the process as an online stochastic sampling problem. By estimating an event importance probability density function (ePDF) on the fly, POLED enables adaptive downsampling that prioritizes the most relevant events. As a specific instance, we proposed a Poissonbased ePDF designed to preserve edge information and compared it against deterministic and uniform baselines.

To evaluate its effectiveness, we tested our method on four diverse tasks: classification, frame interpolation, surface normal estimation, and object detection. Our experiments, conducted in both zero-shot and fine-tuned settings, demonstrate that structured downsampling strategies significantly outperform naive approaches, particularly in lowevent regimes.

Future work includes exploring alternative ePDF formulations, integrating learned event-importance models, and optimizing implementations for real-time deployment on neuromorphic hardware.

References

- Hesam Araghi, Jan van Gemert, and Nergis Tomen. Pushing the boundaries of event subsampling in event-based video classification using cnns. *arXiv preprint arXiv:2409.08953*, 2024. 2, 8
- [2] Juan Barrios-Avilés, Alfredo Rosado-Muñoz, Leandro D Medus, Manuel Bataller-Mompeán, and Juan F Guerrero-Martínez. Less data same information for event-based sensors: A bioinspired filtering and data reduction algorithm. *Sensors*, 18(12):4122, 2018. 2
- [3] Yin Bi, Aaron Chadha, Alhabib Abbas, Eirina Bourtsoulatze, and Yiannis Andreopoulos. Graph-based object classification for neuromorphic vision sensing. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 491–501, 2019. 2
- [4] Anthony Bisulco, Fernando Cladera Ojeda, Volkan Isler, and Daniel Dongyuel Lee. Near-chip dynamic vision filtering for low-bandwidth pedestrian detection. In 2020 IEEE Computer Society Annual Symposium on VLSI (ISVLSI), pages 234–239. IEEE, 2020. 2
- [5] Gregory Cohen, Saeed Afshar, Garrick Orchard, Jonathan Tapson, Ryad Benosman, and Andre van Schaik. Spatial and temporal downsampling in event-based visual classification. *IEEE Transactions on Neural Networks and Learning Systems*, 29(10):5030–5044, 2018. 2
- [6] Pierre De Tournemire, Davide Nitti, Etienne Perot, Davide Migliore, and Amos Sironi. A large scale eventbased detection dataset for automotive. arXiv preprint arXiv:2001.08499, 2020. 2, 5, 6
- [7] Daniel Gehrig, Antonio Loquercio, Konstantinos G. Derpanis, and Davide Scaramuzza. End-to-end learning of representations for asynchronous event-based data. In *Int. Conf. Comput. Vis. (ICCV)*, 2019. 5, 6, 8
- [8] Mathias Gehrig and Davide Scaramuzza. Recurrent vision transformers for object detection with event cameras. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 13884–13893, 2023. 5, 6
- [9] Anindya Ghosh, Thomas Nowotny, and James Knight. Evdownsampling: a robust method for downsampling event camera data. 2, 5, 6
- [10] Anindya Ghosh, Thomas Nowotny, and James C Knight. Insect-inspired spatio-temporal downsampling of eventbased input. In *Proceedings of the 2023 International Conference on Neuromorphic Systems*, pages 1–5, 2023. 2
- [11] Amélie Gruel, Lucía Trillo Carreras, Marina Bueno García, Ewa Kupczyk, and Jean Martinet. Frugal event data: how small is too small? a human performance assessment with shrinking data. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 4092– 4099, 2023. 2
- [12] Amélie Gruel, Jean Martinet, Bernabé Linares-Barranco, and Teresa Serrano-Gotarredona. Performance comparison of dvs data spatial downscaling methods using spiking neural networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 6494–6502, 2023. 2

- [13] Shriya TP Gupta and Basabdatta Sen Bhattacharya. Implementing a foveal-pit inspired filter in a spiking convolutional neural network: a preliminary study. In 2020 International Joint Conference on Neural Networks (IJCNN), pages 1–8. IEEE, 2020. 2
- [14] Manasi Muglikar, Leonard Bauersfeld, Diederik Moeys, and Davide Scaramuzza. Event-based shape from polarization. In *IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*, 2023. 2, 5, 6, 7
- [15] Garrick Orchard, Ajinkya Jayawant, Gregory K Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in neuro-science*, 9:437, 2015. 2, 5, 6, 7
- [16] Charles P Rizzo, Catherine D Schuman, and James S Plank. Neuromorphic downsampling of event-based camera output. In Proceedings of the 2023 Annual Neuro-Inspired Computational Elements Conference, pages 26–34, 2023. 2
- [17] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. Time lens: Event-based video frame interpolation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 16155–16164, 2021. 2, 5, 6