SARE: Structure-Aware Reliability Evaluation and Processing for Event-Based Neural Systems

Haiyu Li
School of Electrical and Electronic Engineering
The University of Sheffield
Sheffield, United Kingdom
hli108@sheffield.ac.uk

Charith Abhayaratne
School of Electrical and Electronic Engineering
Centre for Machine Intelligence
The University of Sheffield
Sheffield, United Kingdom
c.abhayaratne@sheffield.ac.uk

***Abstract*—Event-based neuromorphic sensors produce sparse and asynchronous streams whose reliability depends not only on global event statistics, but also on the preservation of underlying spatio-temporal structure. In practice, noise or localized signal loss or aggressive pre-processing can disrupt structural continuity while leaving global statistics largely unchanged, yielding events that appear plausible yet are structurally invalid. Such mismatches can lead to abrupt performance collapse or silent failure in neural systems. In this paper, we propose a structure-preserving local–global graph processing framework for *reliability diagnosis* in event-based neural systems. Rather than serving as a denoising or enhancement module, it is designed as a diagnostic intervention operator that applies controlled structure-aware graph transformations to probe model sensitivity under progressive structural degradation. The framework jointly captures local neighborhood coherence and global geometric organization, enabling targeted structural interventions without directly optimizing task accuracy. We further introduce a structure-centric evaluation protocol based on geometry-driven objects, multi-scale motion settings and standardized degradation mechanisms that explicitly decouple structural corruption from global event statistics. Experiments show that structural integrity degrades much earlier than global statistical indicators and reliably predicts abrupt failure in both spiking and non-spiking neural models. Compared with purely local or purely global strategies, the proposed local–global formulation is more robust under severe structural degradation and exposes failure modes overlooked by conventional preprocessing and task-driven evaluation. The dataset and codes associated with this work are publicly available at https://github.com/Haiyu12716/SARE and https://doi.org/10.15131/shef.data.31159726.**

***Index Terms*—Event-based vision, Neuromorphic vision, Structure-aware reliability, Failure diagnosis, Structural degradation, noisy events.**

I. Introduction

Modern neural and spiking systems increasingly operate on structured, high-dimensional inputs whose reliability depends not only on statistical validity, but on the preservation of underlying spatio-temporal structure. Localized corruption or preprocessing may leave global statistics largely unchanged while silently degrading structural organization, leading to abrupt inference failure despite statistically plausible inputs [1]. Event-based neuromorphic sensors offer a revealing testbed for such failure modes by asynchronously emitting sparse and locally organized event streams with high temporal resolution and dynamic range, exposing neural systems to regimes where structural integrity critically governs the inference stability [2], [3].

The effectiveness of event streams as neural inputs depends not only on global statistics, but more critically on the preservation of spatio-temporal structure. Event data are inherently vulnerable to noise corruption, aggressive preprocessing and localized signal loss [4], [5]. Importantly, many forms of *structural degradation* arise locally or progressively without inducing substantial changes in global event statistics such as event count, polarity balance, or temporal histograms [6], [7]. As a result, event streams with fundamentally different structural integrity may remain statistically indistinguishable, causing neural and spiking models to operate on inputs that are statistically plausible yet structurally invalid. Similar discrepancies between statistical validity and representational integrity have been observed more broadly in neural systems, where accuracy- or loss-driven evaluation may fail to reflect impending inference instability.

From a system-level perspective, this discrepancy exposes a critical limitation of existing event-based processing and evaluation pipelines, which typically assess robustness through downstream task performance [8], [9]. Such task-driven evaluation is insufficient to diagnose whether the geometric and temporal organization of events remains intact. Structural degradation may accumulate silently and only become apparent as an abrupt performance collapse once a critical threshold is crossed, leading to *silent failure* modes that are difficult to anticipate or interpret [10].

Addressing this challenge requires explicit mechanisms to *diagnose* structural reliability in neural and spiking learning systems, rather than relying solely on task-level outcomes. Event-based perception offers a concrete and controllable setting in which such diagnostic mechanisms can be systematically studied. Reliable neuromorphic perception therefore calls for processing strategies that jointly capture local event coherence and global geometric organization, as structural violations may arise at either scale [11]. Fig. 1 overviews our proposed *Structure-Aware Reliability Evaluation and Processing* framework, which is designed as a diagnostic pipeline rather than a performance-enhancement method.

In summary, the main contributions of this work are as follows:

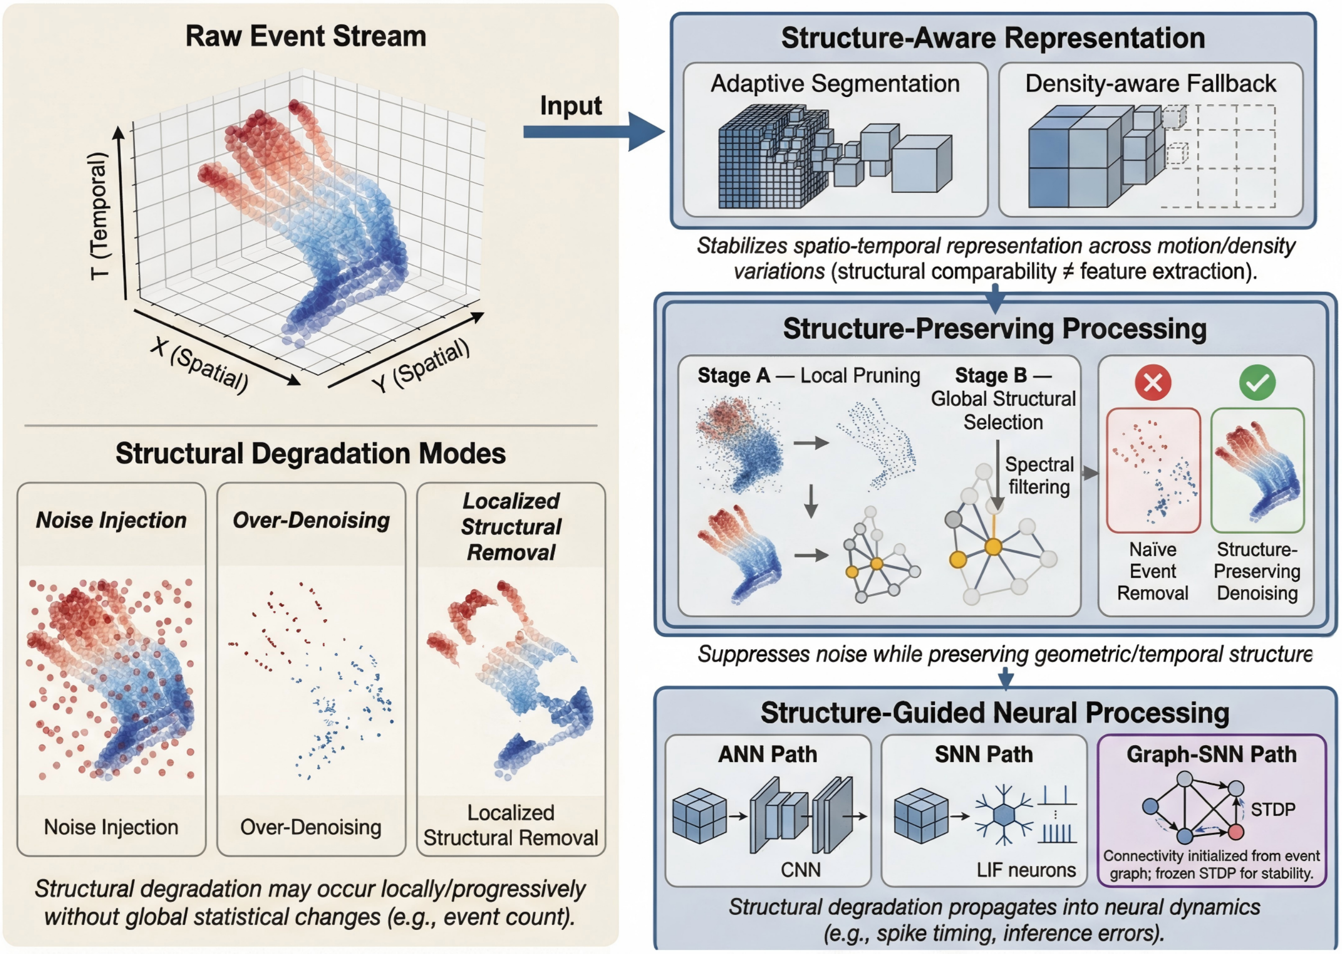

Fig. 1. Overview of the proposed *Structure-Aware Reliability Evaluation and Processing* framework. Controlled degradation (noise, over-denoising, localized removal) can impair spatio–temporal structure without obvious changes in global statistics. The pipeline combines structure-aware representation, structure-preserving processing, and structure-guided neural inference (ANN/SNN/Graph-SNN) to enable reliability-oriented diagnosis of silent failure in event-based neural systems.

- We propose a *structure-aware local–global event graph processing* mechanism for diagnosing representational reliability in neural and spiking systems, preserving spatio–temporal structural fidelity under identical global event statistics.
- We introduce a *neural system analysis tool* that applies controlled, structure-preserving graph transformations prior to inference and traces their effects through both spiking and non-spiking neural models, revealing localized structural failures that remain invisible to accuracy- or loss-based evaluation.
- We establish a *structure-centric evaluation protocol* that decouples local structural degradation from global event statistics and enables principled detection of *silent failure* by linking structural fidelity to neural response dynamics in event-based systems.

II. Related Work

Event-based cameras have attracted increasing attention due to their high temporal resolution, low latency, and sparse asynchronous output, which make them well suited for perception under high-speed motion, in challenging illumination conditions [2] and privacy preserving monitoring [12]. Event-based sensing also shares a close connection to neuromorphic computing and spiking neural networks (SNN) [3].

Beyond raw event streams, numerous representations have been proposed to encode local spatio–temporal structure, including time surfaces, voxel grids, and spike-based accumulations,and graph-based formulations [13], enabling conventional learning models to process asynchronous data [14], [15]. However, these representations typically emphasize statistical aggregation and may obscure fine-grained structural discontinuities.

To support algorithm development, a wide range of event-based datasets have been introduced for task-driven benchmarks such as object recognition, optical flow estimation, and visual reconstruction. Representative datasets including N-CARS [8], CIFAR10-DVS [16], Caltech101-DVS [3], MVSEC [17], and DSEC [18] have been instrumental in advancing event-based learning models through evaluation based primarily on downstream task performance.

Despite their success, most existing benchmarks implicitly assume that event streams which are statistically plausible also preserve the spatio–temporal structure required for reliable neural processing. In practice, degradation mechanisms such as aggressive denoising, sensor noise, or spatially localized signal loss may disrupt local geometric or temporal structure while leaving global statistics (e.g., event count or polarity balance) largely unchanged [4], [10], [19]. Under such conditions, task performance alone cannot reliably distinguish between model limitations and failures induced by compromised input structure [6], [7].

Several works have explored controlled or synthetic event data to study sensing regimes, temporal coding, or neural dynamics under simplified conditions [5], [20]. While these datasets improve interpretability, evaluation remains predominantly task-oriented, with structural degradation reflected only indirectly through performance metrics or qualitative inspection. In our previous work [1], the importance of structural awareness in event camera data denoising prior to machine learning is examined. It further, demonstrates that global-statistics-based metrics fail to capture spatio-temporal structural degradation and motivates the need for structure-sensitive evaluation for neural network-based understanding of event data.

From a reliability perspective, recent studies have begun to examine robustness, noise tolerance, and failure behavior in event-based and spiking systems [9], [21]. In particular, localized or progressive structural degradation may lead to *silent failure* modes that remain invisible to global descriptors and downstream accuracy metrics [11]. These observations highlight the need for diagnostic benchmarks and evaluation protocols that explicitly perturb spatio–temporal structure while maintaining comparable global statistics, motivating structure-preserving processing and reliability-oriented diagnostic frameworks for event-based neural systems.

III. Methodology: Structure-Preserving Local–Global Graph Processing

Rather than proposing a task-specific denoising or reconstruction model, the proposed method is formulated as a *structure-aware intervention operator* on event streams. The objective is to regulate spatio–temporal structural consistency for reliability diagnosis, rather than to optimize downstream task performance.

Event-based cameras generate asynchronous streams of events $e_i = (x_i, y_i, t_i, p_i)$, where $(x_i, y_i) \in \mathbb{R}^2$ denotes spatial location, $t_i \in \mathbb{R}^+$ is the timestamp, and $p_i \in \{-1, +1\}$ represents event polarity. Unlike frame-based sensing, event streams are sparse, temporally precise, and organized around local spatial neighborhoods coupled with motion-induced temporal continuity. As a result, the suitability of event streams as neural inputs depends primarily on the preservation of spatio-temporal structure rather than on global event statistics. We use ϵ_{time} to prevent temporal degeneracy in density estimation and ϵ_{norm} as a numerical stabilizer in normalization; the two constants serve distinct roles and are not interchangeable.

a) Structural reliability and diagnostic objective.: We study *structural reliability*, defined as the degree to which an event stream preserves coherent spatial organization, stable local connectivity, and consistent temporal ordering under noise and preprocessing. Structural degradation arises when these properties are locally disrupted or progressively eroded, even when global statistics (e.g., event density or polarity balance) remain approximately invariant. The goal of this work is to diagnose how such degradation alters event organization and neural response, motivating a processing strategy that is explicitly sensitive to both local and global structural organization.

b) Locality of structural degradation.: In practical sensing scenarios, structural degradation is predominantly local. Noise corruption, aggressive filtering, or region-specific signal loss typically affect confined spatial or temporal neighborhoods, selectively breaking edges, contours, or motion trajectories. While such localized failures can severely impair global structural coherence, they often leave aggregate statistics nearly unchanged. Consequently, event streams with fundamentally different structural integrity may appear indistinguishable under global evaluation metrics, necessitating diagnostic methods that explicitly capture *local–global structural discrepancies*.

c) Structure-stabilized event representation.: To enable meaningful comparison across sensing conditions, we first normalize the spatio–temporal organization of events prior to diagnostic processing. Given an event set containing N events, the normalized spatio–temporal density is defined as

$$\rho_{\text{norm}} = \frac{N}{|\Omega_{xy}|\ \max\big(\Delta t,\, \varepsilon_{\text{time}}\big)}. \tag{1}$$

Here $|\Omega_{xy}| = (x_{\max} - x_{\min})(y_{\max} - y_{\min})$ denotes the spatial support area of the event stream, $\Delta t = t_{\max} - t_{\min}$ is the temporal extent, and $\varepsilon_{\text{time}}$ is a small constant ($\varepsilon_{\text{time}} = 10^{-6}$) introduced to prevent numerical degeneracy when the temporal span collapses.

Based on ρ_{norm}, the number of events per segmentation window is determined as

$$N_{\text{win}} = \max\Big(N_{\min},\ \big\lceil \rho_{\text{norm}}\ \log_{10}(N+1) \big\rceil\Big), \tag{2}$$

where $N_{\min}$ enforces minimal structural support ($N_{\min} = 50$ in all experiments). This adaptive rule stabilizes segmentation granularity across event streams with different densities and motion profiles. When temporal variation collapses ($t_{\max} - t_{\min} < \varepsilon_{\text{time}}$), segmentation is performed purely in the spatial domain to avoid degenerate temporal partitioning. This stage normalizes structural units rather than enhancing task accuracy, and operates linearly with respect to the number of events.

d) Structure-preserving local–global graph processing.: Within each stabilized voxel, events are represented as an undirected weighted graph $G = (V, E)$, where each node $v_i \in V$ corresponds to an event $e_i = (x_i, y_i, t_i, p_i)$. Edges are

Algorithm 1 Structure-Preserving Local–Global Graph Processing

Require: Event stream $E = \{(x_i, y_i, t_i, p_i)\}_{i=1}^{N}$; k; $N_{\min}$; $\varepsilon_{\text{time}}$; ε_{vel}; ε_{dir}; $\varepsilon_{\text{norm}}$; $(\alpha, \beta, \gamma, \delta)$; candidate thresholds $\mathcal{T}$; reduction size m.

Ensure: Processed graph $\hat{G}$.

$(\Delta x, \Delta y, \Delta t) \leftarrow \textsc{Span}(E)$

$\rho_{\text{norm}} \leftarrow \dfrac{N}{\Delta x\, \Delta y\, \max(\Delta t, \varepsilon_{\text{time}})}$

$N_{\text{win}} \leftarrow \max\Big(N_{\min}, \lceil \rho_{\text{norm}} \log_{10}(N+1) \rceil\Big)$

$\mathcal{S} \leftarrow \textsc{Partition}(E, N_{\text{win}})$

$G \leftarrow \textsc{kNNGraph}(\mathcal{S}, k)$

$\tilde{W} \leftarrow \textsc{ProxyWeights}(G; \alpha, \beta, \gamma, \delta, \varepsilon_{\text{vel}}, \varepsilon_{\text{dir}})$

for all $\tau \in \mathcal{T}$ **do**

  $G_\tau \leftarrow \textsc{PruneEdges}(G, \tilde{W}, \tau)$

  $\phi_\tau \leftarrow \textsc{NodeResponse}(G_\tau)$

  $\bar{\phi}_\tau \leftarrow \dfrac{\phi_\tau}{\max(\phi_\tau) + \varepsilon_{\text{norm}}}$

  $\sigma_\tau^2 \leftarrow \mathrm{Var}(\bar{\phi}_\tau)$

end for

$\tau^\star \leftarrow \arg\min_{\tau \in \mathcal{T}} \sigma_\tau^2$

$G_{\text{loc}} \leftarrow G_{\tau^\star}$

$(G_{\text{red}}, \pi) \leftarrow \textsc{ReduceGraph}(G_{\text{loc}}, m)$

$c_{\text{red}} \leftarrow \textsc{Centrality}(G_{\text{red}})$

$c \leftarrow \textsc{Lift}(c_{\text{red}}, \pi)$

$\hat{G} \leftarrow \textsc{GlobalRefine}(G_{\text{loc}}, c)$

return $\hat{G}$

constructed via k-nearest-neighbor search in spatio–temporal space ($k = 8$), yielding a sparse local connectivity pattern.

For each connected node pair (i, j), a proxy edge weight is defined as

$$\tilde{w}_{ij} = \alpha D_{ij} + \beta \Delta v_{ij} + \gamma \theta_{ij} + \delta P_{ij}, \tag{3}$$

where D_{ij} denotes the normalized spatial distance between events. Local motion consistency is captured using the velocity estimate

$$\mathbf{v}_{ij} = \left(\frac{x_j - x_i}{t_j - t_i + \varepsilon_{\text{vel}}}, \frac{y_j - y_i}{t_j - t_i + \varepsilon_{\text{vel}}} \right), \tag{4}$$

with ϵ_t preventing numerical instability. The velocity magnitude deviation is defined as $\Delta v_{ij} = |\|\mathbf{v}_{ij}\| - \bar{v}|$, where $\bar{v}$ denotes the voxel-wise mean velocity magnitude. Directional inconsistency is measured by

$$\theta_{ij} = 1 - \frac{\mathbf{v}_{ij} \cdot \bar{\mathbf{v}}}{\|\mathbf{v}_{ij}\|\, \|\bar{\mathbf{v}}\| + \varepsilon_{\text{dir}}}, \tag{5}$$

where $\bar{\mathbf{v}}$ denotes the mean velocity vector in the local neighborhood, and $\varepsilon_{\text{dir}} > 0$ is a small constant to avoid numerical instability. All components are normalized to $[0, 1]$ prior to aggregation. Unless otherwise stated, $\alpha = \beta = \gamma = \delta = 1$, reflecting a diagnostic rather than task-optimized formulation.

To diagnose structural integrity, a hybrid *local–global* strategy is adopted. Local structural inconsistency is quantified by the variance of the normalized node-degree response

$$\sigma_\tau^2 = \mathrm{Var}\left(\frac{\phi_\tau}{\max(\phi_\tau) + \varepsilon_{\text{norm}}} \right), \tag{6}$$

where Φ_τ denotes the node-degree sequence of G_τ. Low variance indicates homogeneous local connectivity, whereas elevated variance signals fragmentation caused by structural degradation. The complete pipeline is summarized in Algorithm 1.

e) Neural sensitivity probing via structure-guided inference.: To examine how structural degradation propagates into neural dynamics, both non-spiking and spiking neural models are employed as *structure-sensitive probes*. In the spiking setting, a graph-structured SNN is constructed by initializing synaptic connectivity from the processed event graph. Each node is mapped to a neuron governed by leaky integrate-and-fire dynamics:

$$I_i(t) = \sum_{j \in \mathcal{N}(i)} \tilde{A}_{ij} s_j(t-1), \tag{7}$$

$$V_i(t) = \lambda V_i(t-1) + I_i(t), \tag{8}$$

$$s_i(t) = \mathbb{I}[V_i(t) \geq \theta], \tag{9}$$

where $\tilde{A}_{ij}$ denotes the normalized adjacency, $\lambda = 0.95$ is the membrane decay factor, and θ is a fixed firing threshold. We use the symmetrically normalized adjacency $\tilde{A} = D^{-1/2} A D^{-1/2}$ of $\hat{G}$. Optional STDP is applied only to low-confidence edges, while high-confidence ones are frozen, as the neural models are used to expose structure-coupled failure signatures rather than to optimize accuracy.

f) Reliability-oriented methodology.: Overall, the proposed methodology integrates structure-stabilized representation, local–global graph processing, and neural inference into a unified diagnostic framework. By explicitly enforcing and perturbing structural consistency at both local and global scales, the method enables systematic analysis of how structural degradation emerges and how it manifests in neural behavior.

IV. Dataset Construction

We present SAED (*Structure-Aware Event Dataset*)[1], which is a *structure-centric diagnostic dataset* specifically designed for controlled analysis of local–global spatio–temporal structure in event streams. Unlike task-driven benchmarks, it prioritizes interpretable structural variation over semantic diversity or task difficulty. Representative geometry-driven objects and acquisition conditions are shown in Fig. 2.

The dataset is constructed using geometry-driven objects selected for their clear and interpretable structural signatures in the event domain, such as well-defined edges, contours, and connectivity patterns. To enable systematic local–global structural analysis, object geometry is held fixed while controlled variations are introduced through object scale and motion.

[1] Dataset download link: https://doi.org/10.15131/shef.data.31159726

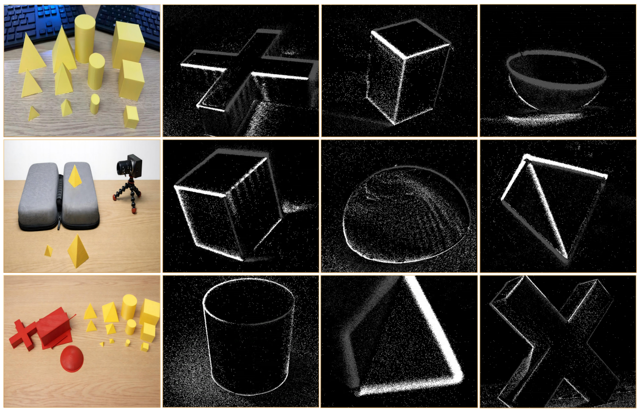

Fig. 2. **Qualitative overview of SAED dataset.** Geometry-driven objects and corresponding event-stream visualizations. Clear edges and contours are preserved in the event domain, supporting structure-centric analysis under noise, over-denoising, and localized damage.

Specifically, the dataset includes ten 3D-printed geometric objects, comprising five multi-scale shapes fabricated at three relative scales (S/M/L = $0.5\times$/$1.0\times$/$1.5\times$) and five single-scale shapes. All objects are recorded under two motion regimes: object motion with a static camera and camera motion with a static object, producing comparable yet distinct spatio–temporal event structures. All recordings are acquired using a DVXplorer event camera with a spatial resolution of 640×480 under fixed illumination conditions. Multiple sequences are captured for each object–scale–motion configuration, resulting in approximately 27,000 event clips in total.

All sequences are collected in a controlled indoor environment to minimize confounding sensing factors. The dataset is organized in a factor-aware manner across object geometry, spatial scale, and motion condition, with standardized formats and associated metadata.

V. Performance Evaluation

A. Evaluation Goals and Evidence Chain

The evaluation is designed to validate the proposed *structure-preserving local–global event graph processing* method for reliability diagnosis in event-based neural systems. Rather than maximizing task performance, the goal is to establish an interpretable evidence chain demonstrating that (i) structural degradation can remain statistically inconspicuous, (ii) preserving local–global structure improves robustness beyond event suppression, and (iii) neural failure behavior is aligned with violations of local–global structural consistency.

Specifically, the experiments are organized to support three claims.

(C1) Silent failure arises when local–global structure collapses while global statistics remain stable. We examine whether progressive and localized structural corruption can severely impair neural inference while global event statistics (e.g., event count, polarity balance, and average event rate) evolve smoothly or remain comparable.

(C2) Robustness gains arise from preserving local–global structure, not merely removing events. We test whether the proposed hybrid local–global strategy yields higher structural fidelity and delayed failure thresholds under matched sparsification levels, thereby distinguishing structure preservation from event-count reduction.

(C3) Spiking systems exhibit structure-coupled failure signatures beyond accuracy. We analyze whether SNN dynamics, including spike rate and active neuron ratio, exhibit characteristic breakpoints aligned with structural fidelity collapse, and compare these behaviors with non-spiking models.

These claims are evaluated through a causal evidence chain comprising: **(i)** standardized perturbations that explicitly decouple structural corruption from global statistics; **(ii)** local–global graph processing variants (Local-only / Global-only / Hybrid) to isolate the contribution of each structural cue; and **(iii)** neural readouts (SNN) used as *structure-sensitive probes* to connect structural fidelity with inference behavior. Ablation studies further verify which components are necessary for the observed diagnostic trends.

a) Evaluation overview: Structural reliability is evaluated using a dual-track protocol that combines synthetic data for causal control with our dataset for real-world validation, under standardized degradation settings designed to decouple structural corruption from global event statistics.

B. Dataset, Factors, and Degradation Protocols

a) Dual-track evaluation: We adopt a dual-track evaluation strategy. Synthetic event streams with explicit geometric primitives and motion models are used to enable precise and causal control of local–global structural degradation, while our dataset provides real-world validation under controlled sensing conditions. For the dataset, each clip is associated with a geometry category $g \in \mathcal{G}$, physical scale $s \in \mathcal{S}$, and motion condition, $m \in \mathcal{M}$ (`OBJECT-MOVE`, `CAMERA-MOVE`, `BOTH-MOVE`). Unless otherwise stated, models predict g, while s and m are treated as nuisance factors.

b) Representation and degradation protocols: Raw event streams are segmented into clips of duration Δt and converted into representation units compatible with grid-based or graph-based models. Adaptivate event points segmentation (AEPS) [22] is optionally applied as a *measurement stabilizer* to improve comparability across motion and scale regimes.

To decouple structural degradation from global statistics, three standardized perturbation protocols are applied to both synthetic and real data: (i) noise injection with ratio r_n, (ii) over-denoising stress tests via threshold sweeps $\tau \in \mathcal{T}$, and (iii) localized structural removal with ratio r_ℓ. Performance is reported as response curves over τ, r_n, and r_ℓ to identify reliability margins and structural breakpoints.

C. Models and Pipelines

a) Local–global graph processing (LGG): Given an event set, an event graph $G = (V, E)$ is constructed, where nodes correspond to events and edges encode pairwise structural relationships. LGG performs hybrid local–global selection: Stage A applies fast local pruning to generate a sparse candidate threshold set, and Stage B evaluates global structural consistency on the reduced graph to select the final threshold. Stage-A-only, Stage-B-only, and full Hybrid variants are evaluated to isolate the roles of local adaptivity, global consistency, and computational efficiency.

b) Neural inference as sensitivity probes: Artificial Neural Network (ANN) and SNN models are employed as diagnostic probes to examine how structural degradation affects inference behavior. These include a voxel-based ANN and an LIF-based SNN trained with surrogate gradients. A graph-structured SNN (GSSNN) is further evaluated, with synaptic connectivity initialized from the processed event graph to strengthen structure–spiking coupling. All neural models are treated as diagnostic readouts rather than task-optimized architectures.

c) Pipelines compared: We compare end-to-end pipelines (preprocessing $\rightarrow$ graph processing $\rightarrow$ neural inference): Raw, Classical filter, Local-only, Global-only, and Hybrid.

D. Evaluation Metrics

a) Global statistics: Event count, polarity balance, and average event rate are reported as global statistical baselines, which are by construction insensitive to localized structural failures.

b) Structural fidelity: Structural integrity is quantified using geometry- and point-set-based metrics, specifically Chamfer Distance (CD) and Hausdorff Distance (HD), between processed outputs and clean references under consistent sampling. CD measures average bidirectional proximity between point sets, reflecting overall geometric alignment, while HD captures the maximum deviation, revealing worst-case structural discontinuities.

c) Neural outcomes: Accuracy curves over τ and r_ℓ are reported to identify failure thresholds and robustness margins. For SNNs, spiking indicators such as spike rate and active neuron ratio are also reported to expose failure signatures beyond accuracy.

E. Main Evaluation and Results

1) Evaluation 1: Silent Failure as Local–Global Structure Collapse:

a) Setup: Progressive structural degradation is evaluated using (i) over-denoising sweeps $\tau \in \mathcal{T}$ and (ii) localized structural removal ratios $r_\ell \in \mathcal{R}_\ell$. For each setting, we report global statistics, structural fidelity (CD/HD), and neural readouts, including ANN/SNN accuracy and spiking indicators. The analysis emphasizes *ordering*, examining whether structure-aware metrics and neural signatures exhibit discernible changes before global statistics show substantial variation.

b) Findings and interpretation: As shown in Fig. 3, global statistics and structural fidelity diverge markedly as degradation increases. While global descriptors evolve smoothly across perturbation levels, CD and HD exhibit clear early breakpoints, indicating that global statistics are insufficient proxies for input validity. Neural behavior mirrors this structural breakdown: ANN accuracy degrades gradually, whereas SNNs undergo abrupt performance collapse once structural fidelity falls below a critical threshold. This regime, in which inputs remain statistically plausible but become structurally invalid, characterizes a *silent failure* phenomenon that is invisible to global statistical evaluation.

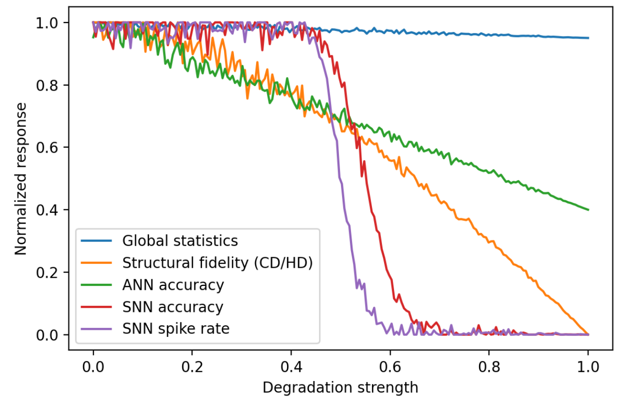

Fig. 3. Evaluation 1: Silent failure as local–global structural collapse. Global statistics, CD/HD, and ANN/SNN readouts are reported over degradation strength. CD/HD are evaluated on spatio-temporal point sets $(x, y, \tilde{t})$ with normalized timestamps to measure structural continuity.

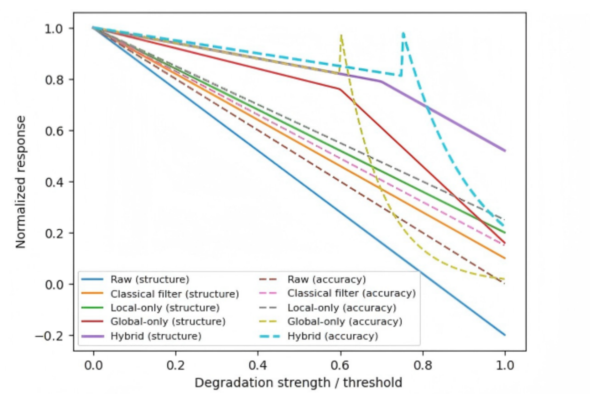

Fig. 4. Evaluation 2: Hybrid local–global graph processing improves structural fidelity and robustness under controlled perturbations.

2) Evaluation 2: Hybrid Local–Global Graph Processing Preserves Structure Beyond Event Suppression:

a) Setup: We compare multiple pipelines (Raw, Classical filter, Local-only, Global-only, and Hybrid) under noise injection and over-denoising settings. Structural fidelity (CD/HD), global statistics, and ANN/SNN outcomes are evaluated. To disentangle *event suppression* from *structure preservation*, all pipelines are additionally compared at matched event-count levels whenever applicable.

b) Findings and interpretation: As illustrated in Fig. 4, the compared pipelines differ fundamentally in whether they preserve structural organization or merely suppress events. Classical filtering and Local-only pruning primarily reduce event density and exhibit pronounced geometric fragmentation

TABLE I

ABLATION SUMMARY ON OUR SAED DATASET UNDER MATCHED DEGRADATION SETTINGS (MEAN ± STD).

Ablation	Variant	CD ↓	HD ↓	Acc. ↑	Runtime ↓
A1	Stage A only	0.142 ± 0.006	0.311 ± 0.012	84.6 ± 0.7	0.18 ± 0.01 s
	Stage B only	0.121 ± 0.005	0.284 ± 0.010	86.9 ± 0.6	1.92 ± 0.06 s
	Hybrid	$\mathbf{0.098 \pm 0.004}$	$\mathbf{0.241 \pm 0.009}$	$\mathbf{89.7 \pm 0.5}$	0.54 ± 0.02 s
A2	Dense sweep	0.097 ± 0.004	0.240 ± 0.009	89.8 ± 0.5	3.86 ± 0.10 s
	Sparse candidates	0.099 ± 0.004	0.242 ± 0.010	89.6 ± 0.5	$\mathbf{0.54 \pm 0.02}$ s
A3	Full (all factors)	$\mathbf{0.098 \pm 0.004}$	$\mathbf{0.241 \pm 0.009}$	$\mathbf{89.7 \pm 0.5}$	0.54 ± 0.02 s
	without D	0.136 ± 0.006	0.305 ± 0.012	85.1 ± 0.8	0.49 ± 0.02 s
	without Δv	0.120 ± 0.005	0.276 ± 0.011	86.8 ± 0.7	0.50 ± 0.02 s
	without θ	0.114 ± 0.005	0.268 ± 0.011	87.4 ± 0.6	0.50 ± 0.02 s
	without polarity	0.104 ± 0.004	0.249 ± 0.010	89.1 ± 0.5	0.52 ± 0.02 s
A4	**AEPS full**	$\mathbf{0.098 \pm 0.004}$	$\mathbf{0.241 \pm 0.009}$	89.7 ± 0.5	0.54 ± 0.02 s
	without fallback	0.106 ± 0.008	0.255 ± 0.018	89.2 ± 1.2	0.52 ± 0.02 s
	fixed voxel count	0.112 ± 0.012	0.270 ± 0.026	88.8 ± 1.8	$\mathbf{0.48 \pm 0.02}$ s
A5	GSSNN: no STDP	0.128 ± 0.006	0.295 ± 0.012	83.9 ± 0.9	0.41 ± 0.02 s
	GSSNN: STDP	0.121 ± 0.010	0.282 ± 0.021	86.1 ± 2.0	0.57 ± 0.03 s
	GSSNN: STDP+freeze	$\mathbf{0.112 \pm 0.007}$	$\mathbf{0.263 \pm 0.014}$	$\mathbf{87.9 \pm 1.0}$	0.59 ± 0.03 s
A6	Global-only eval	0.101 ± 0.004	0.246 ± 0.010	89.3 ± 0.5	0.12 ± 0.01 s
	Local+Global eval	$\mathbf{0.098 \pm 0.004}$	$\mathbf{0.241 \pm 0.009}$	89.7 ± 0.5	0.16 ± 0.01 s

TABLE II

BREAKPOINT SUMMARY UNDER PROGRESSIVE OVER-DENOISING. τ^*_{STRUCT} IS DERIVED FROM CD/HD CURVES, WHILE τ^*_{ANN} AND τ^*_{SNN} MARK ACCURACY COLLAPSE. HIGHER τ^* INDICATES GREATER ROBUSTNESS.

Pipeline	τ^*_{struct} ↑	τ^*_{ANN} ↑	τ^*_{SNN} ↑
Raw	0.53	0.59	0.55
Classical filter	0.56	0.62	0.57
Local-only	0.60	0.66	0.60
Global-only	0.63	0.69	0.62
Hybrid (Stage A + B)	**0.70**	**0.75**	**0.68**

under severe perturbation. Global-only processing improves robustness but incurs substantially higher computational cost. In contrast, the Hybrid strategy consistently achieves higher structural fidelity and delayed failure thresholds, including under matched event-count conditions, demonstrating that its robustness gains arise from preserving local–global structure rather than from aggressive sparsification.

AEPS further reduces confounding effects by adjusting segmentation granularity across sensing conditions. This leads to more consistent structural and neural response trends, supporting the role of AEPS as a *measurement stabilizer* rather than a task-optimizing component.

To summarize robustness under progressive over-denoising, Table II reports the structural breakpoint τ^*_{struct}, together with the corresponding ANN and SNN failure thresholds τ^*_{ANN} and τ^*_{SNN}, across all pipelines. The Hybrid local–global strategy consistently attains the highest breakpoints under both structural and neural criteria, indicating superior robustness prior to accuracy collapse.

3) Evaluation 3: Structure-Coupled Spiking Failure Signatures and Graph-Guided Connectivity:

a) Setup: We evaluate an ANN, a conventional SNN, and a graph-structured SNN under progressive structural degradation. Accuracy curves and spiking indicators, including layer-wise spike rate and active neuron ratio, are reported to analyze how graph-guided connectivity modulates failure sensitivity.

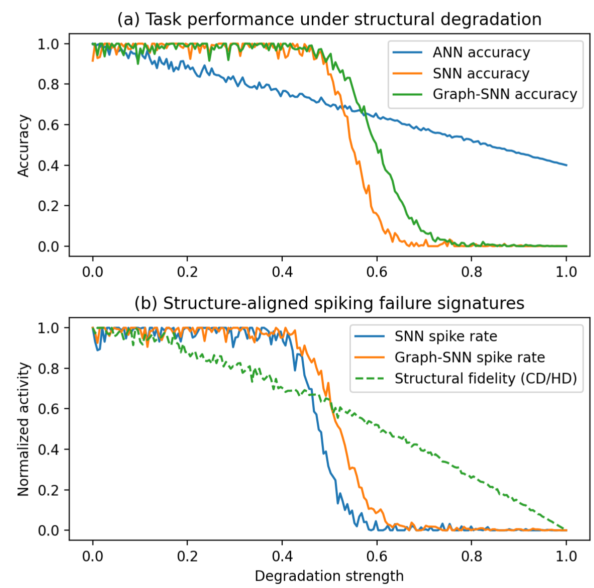

Fig. 5. Evaluation 3: Structure-coupled spiking failure signatures. Accuracy and spiking dynamics are reported over degradation strength.

b) Findings and interpretation.: As shown in Fig. 5, structural degradation manifests not only as accuracy loss but also as distinct spiking failure signatures. Conventional SNNs exhibit abrupt collapse near the CD/HD structural breakpoints, consistent with the loss of coherent spatio–temporal organization. In contrast, the GSSNN demonstrates a clearer and more stable correspondence between structural fidelity and spiking dynamics, owing to its structure-guided connectivity. These results indicate that spiking failure behavior is tightly coupled to local–global structural integrity, and that explicit graph-guided topology enhances the interpretability and sensitivity of spiking neural responses under progressive structural degradation.

F. Ablation Studies and Component Analysis

We perform ablation studies by comparing the Local-only (Stage A), Global-only (Stage B), and the proposed Hybrid (Stage A+B) variants. As summarized in Table I, Local-only processing is computationally efficient but degrades under uneven event distributions, whereas Global-only processing improves robustness at substantially higher computational cost. The Hybrid formulation consistently achieves the best trade-off among structural fidelity, robustness, and efficiency, highlighting the necessity of coupling local adaptivity with global structural consistency.

Restricting Stage B to the sparse candidate threshold set $\mathcal{D}_{\mathrm{cand}}$ generated by Stage A incurs negligible degradation in CD/HD trends while substantially reducing computational overhead. By contrast, dense threshold sweeps yield only marginal robustness gains at significantly higher cost, confirming the efficiency advantage of the proposed hybrid strategy.

Additional ablations show that spatial proximity and motion consistency are the dominant factors in preserving geometric continuity and temporal ordering, while polarity consistency improves robustness under localized structural damage. Removing AEPS fallback or adaptive granularity reintroduces representation instability and increases variance in CD/HD and neural responses, confirming their stabilizing role for reliability-oriented evaluation rather than task optimization.

1) Neural Sensitivity and Diagnostic Necessity: We further ablate plasticity mechanisms in the GSSNN. Moderate spike-timing-dependent plasticity (STDP) improves robustness under progressive structural degradation, while freezing high-confidence connections stabilizes learning dynamics and preserves structure-coupled failure trends.

Finally, we compare global-only reporting with local–global sensitivity analysis under localized structural removal. Global-only statistics fail to reflect severe localized damage, whereas local–global structural measures align closely with neural failure behavior.

VI. Conclusions

In this paper, we have presented a structure-preserving local–global event graph framework for reliability evaluation in event-based neural systems, emphasizing *structural reliability* over task performance. We showed that the structural degradation can emerge locally and progressively despite largely unchanged global event statistics. Controlled perturbations revealed a *silent failure* regime in which event streams remain statistically plausible yet become structurally invalid. In this regime, Chamfer Distance (CD) and Hausdorff Distance (HD) exhibit clear breakpoints aligned with abrupt inference collapse, especially in spiking neural models. Our results show that reliable neuromorphic perception requires explicit monitoring of local–global spatio–temporal structure rather than reliance on task accuracy or global statistics.

References

[1] H. Li and C. Abhayaratne, "Is structural awareness the key to event camera data cleansing for enhancing veracity?" in *Proceedings of the British Machine Vision Conference (BMVC)*, 2025. [Online]. Available: https://bmvc2025.bmva.org/proceedings/1007/

[2] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, D. Scaramuzza, and Y. Aloimonos, "Event-based vision: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 1, pp. 154–180, 2022.

[3] G. Orchard, A. Jayawant, G. Cohen, and N. Thakor, "Converting static image datasets to spiking neuromorphic datasets using saccades," *Frontiers in Neuroscience*, vol. 9, p. 437, 2015.

[4] R. Baldwin, M. Almatrafi, J. Ritchie, and M. R. Azimi-Sadjadi, "Noise characterization of dynamic vision sensors," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018.

[5] E. Mueggler, H. Rebecq, G. Gallego, T. Delbrück, and D. Scaramuzza, "The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and SLAM," *The International Journal of Robotics Research*, vol. 36, no. 2, pp. 142–149, 2017.

[6] D. Gehrig, H. Rebecq, G. Gallego, and D. Scaramuzza, "End-to-end learning of representations for asynchronous event-based data," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.

[7] H. Rebecq, R. Ranftl, V. Koltun, and D. Scaramuzza, "Events-to-video: Bringing modern computer vision to event cameras," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[8] A. Sironi, M. Brambilla, N. Bourdis, X. Lagorce, and R. Benosman, "HATS: Histograms of averaged time surfaces for robust event-based object classification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[9] Y. Hu, S. Liu, and T. Delbrück, "Event-based denoising and background activity suppression," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 5756–5770, 2022.

[10] J. Binas, D. Neil, S.-C. Liu, and T. Delbrück, "DDD20: A neuromorphic dataset for continuous object recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[11] D. Scaramuzza and F. Fraundorfer, "Visual odometry [tutorial]," *IEEE Robotics & Automation Magazine*, vol. 18, no. 4, pp. 80–92, 2011.

[12] S. Al-Obaidi, H. Al-Khafaji, and C. Abhayaratne, "Making sense of neuromorphic event data for human action recognition," *IEEE Access*, vol. 9, pp. 82 686–82 700, 2021.

[13] Y. Bi and S.-C. Zhu, "Graph-based event stream representation for human action recognition," in *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)*, 2018.

[14] X. Lagorce, G. Orchard, F. Galluppi, B. E. Shi, and R. Benosman, "HOTS: A hierarchy of event-based time-surfaces for pattern recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 7, pp. 1346–1359, 2017.

[15] A. Zhu, L. Yuan, and G. Shi, "A time-efficient method for event-based optical flow estimation," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2019.

[16] H. Li, H. Liu, X. Ji, G. Li, and L. Shi, "CIFAR10-DVS: An event-stream dataset for object classification," *Frontiers in Neuroscience*, vol. 11, p. 309, 2017.

[17] A. Zhu, D. Yuan, Y. Chi, and G. Shi, "The multi-vehicle stereo event camera dataset: An event camera dataset for 3d perception," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018.

[18] M. Gehrig, D. Gehrig, J. Hidalgo-Carrió, and D. Scaramuzza, "DSEC: A stereo event camera dataset for driving scenarios," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4947–4954, 2021.

[19] X. Clady, S.-H. Ieng, and R. Benosman, "Asynchronous event-based corner detection and matching," *Neural Networks*, vol. 66, pp. 91–106, 2015.

[20] C. Reinbacher, G. Munda, T. Pock, and T. Pock, "Real-time intensity-image reconstruction for event cameras using manifold regularisation," in *Proceedings of the British Machine Vision Conference (BMVC)*, 2016.

[21] D. Neil, M. Pfeiffer, and S.-C. Liu, "Phased LSTM: Accelerating recurrent network training for long or event-based sequences," in *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, 2016.

[22] H. Li and C. Abhayaratne, "AW-GATCN: Adaptive weighted graph attention convolutional network for event camera data joint denoising and object recognition," in *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)*, 2025, pp. 1–8.