EVENTCERT: CERTIFIED ROBUSTNESS OF EVENT-BASED CLASSIFIERS TO STRUCTURED DELETION

Anonymous authors
Paper under double-blind review

ABSTRACT

Event cameras record the changes in a scene as a stream of events, which is turned into a tensor so that standard models can process it efficiently. In practice, packet loss, sensor faults, or occlusion can remove whole blocks of this tensor at once. In a traditional camera, such a loss does limited damage because neighbouring pixels carry similar information. Such redundancy is almost absent in event streams. As a result, a lost block is lost for good and can change a classifier's prediction, posing serious risks in safety-critical applications such as robotics and driving. It is therefore essential to prove how many blocks can be removed while the prediction is guaranteed to stay the same, a question existing certification methods have largely left unexplored. We close this gap with EventCert, a method that proves, with high probability, how many blocks each prediction can lose without changing. It mirrors how events are deleted and certifies how many blocks an adversary may remove. Because only blocks that contain events can change the prediction, it turns this worst-case certificate into the first exact guarantees for random and bursty channel loss, and for an adversary and a lossy channel acting together, all from a single certification run. On five event-camera and neuromorphic-audio benchmarks, we compare EventCert with three existing certification methods: Gaussian smoothing, randomized ablation, and random deletion of single events. EventCert certifies 3.8–5.6$\times$ more deleted blocks than the best of them, while the other two certify almost none. On N-MNIST, 89% of test inputs stay certified even when an adversary removes any 8 blocks. While standard metrics only tell us how often a model is right, EventCert tells us, for each prediction, whether it will stay the same when part of the data is lost. By telling when an event-camera decision can still be trusted and when it cannot, EventCert can make self-driving cars safer, drones more reliable, and autonomous robots more effective.

1 INTRODUCTION

Event cameras do not record full frames. Each pixel reports a small *event* when its brightness changes. This makes the sensor fast and energy-efficient, and it is increasingly used in robotics, driving, and edge devices (Gallego et al., 2022). In practice, however, parts of the event stream are often lost: network packets drop, buffers overflow, sensor regions fail, or objects block the view (Wang et al., 2026). Most classifiers first turn the event stream into a tensor. Each of these failures then removes a whole block of that tensor: a short time window, an area of the sensor, and its polarity channels. We call such a block a *spatio-temporal (SST) group*. A deployed system needs to know how many groups can be lost before its prediction changes.

EventCert introduces a new way to certify event-based classifiers, built in three steps. First, instead of adding pixel noise as in standard randomized smoothing (Cohen et al., 2019), it smooths by deleting whole SST groups, which ensures that the noise has exactly the same structure as real event loss (Levine & Feizi, 2020b). Using standard Monte Carlo estimation and Clopper–Pearson confidence bounds, it certifies a worst-case radius: the number of groups an informed adversary may delete without changing the prediction. Second, and central to our approach, we introduce an active-support bridge. Because deleting an empty group leaves the tensor unchanged, only event-carrying groups count against this radius. This insight lets us convert the worst-case certificate into a physical-loss certificate, a guarantee that existing certificates, limited to an adversary, cannot give. Third, we make this certificate exact for realistic links. We derive a new dynamic program over a

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Gilbert burst channel for bursty loss, a hypergeometric tail for uniform loss, and a worst-case analysis for hybrid threats in which an adversary and a lossy channel act together. All of these certificates are computed from a single certification run, without evaluating the classifier again. Our three main contributions are:

- **Guarantees for Bursty and Mixed Loss:** To our knowledge, EventCert is the first method to certify robustness to random, bursty channel loss, and to an adversary and a lossy channel acting together, all from one certification run.
- **Much Stronger Certificates:** On five benchmarks, EventCert certifies 3.8–5.6$\times$ more deleted groups than the strongest existing method, while two other existing methods certify almost none.
- **Bursts Break Naive Guarantees:** Guarantees that assume independent losses are far too optimistic: on N-MNIST at a 5% loss rate, certified accuracy drops from 85.6% to 14.5% when losses come in bursts.

The rest of the paper is organized into five sections. Section 2 reviews existing certification methods and the gaps they leave. Section 3 defines event tensors, SST groups, and the two kinds of loss we certify against. Section 4 presents EventCert and its guarantees for worst-case, random, bursty, and hybrid loss. Section 5 evaluates EventCert on five benchmarks against existing certification methods. Finally, Section 6 discusses limitations and concludes; proofs and full results are in the appendix.

2 Related Work

Randomized smoothing turns any classifier into one with a provable guarantee by classifying many noisy copies of the input (Cohen et al., 2019), and adversarial training strengthens these certificates (Salman et al., 2019). These methods certify a ball of small additive noise, but deleting a block of events is not a small change: it sets many tensor entries from 1 to 0, so even one block lies far outside the certified ball. A second line of work matches the noise to a deletion threat: randomized ablation and derandomized smoothing certify changed pixels and image patches (Levine & Feizi, 2020b;a), RS-Del certifies deleted tokens (Huang et al., 2023), PointGuard and PointCert certify changed points (Liu et al., 2021; Zhang et al., 2023), and hierarchical smoothing certifies changes to a subset of entities (Scholten et al., 2023). These works establish that the noise must match the threat, but their units are never blocks spanning time, space, and polarity, and they certify only a worst-case adversary, not random loss on a real link.

Table 1: Research gaps and where EventCert closes them (references in the text). ✓: covered; $\times$: not covered.

Method family	Certified unit	Spatio-temporal unit	Adversary	Random loss	Bursty loss	Adversary + loss
Gaussian smoothing	ℓ_2 perturbation	$\times$	✓	$\times$	$\times$	$\times$
Ablation / patch smoothing	pixels, patches	$\times$	✓	$\times$	$\times$	$\times$
Sequence deletion	tokens	$\times$	✓	$\times$	$\times$	$\times$
PointGuard / PointCert	points	$\times$	✓	$\times$	$\times$	$\times$
Hierarchical smoothing	entity subsets	$\times$	✓	$\times$	$\times$	$\times$
Event / SNN robustness	ℓ_p / empirical	$\times$	partial	$\times$	$\times$	$\times$
EventCert	SST groups	✓ Sec. 3	✓ Eq. 5	✓ Eq. 7	✓ Prop. 1	✓ Prop. 2

For event cameras and spiking networks, prior work designs attacks on raw events (Yao et al., 2024), studies empirical robustness (Ding et al., 2025), and certifies spiking networks against ℓ_p perturbations (Mukhoty et al., 2024; Wu et al., 2024), but none certifies the deletion of events or models the packet loss these systems face (Wang et al., 2026). EventCert applies the matching principle to the unit in which event data is actually lost, the spatio-temporal group, and goes beyond prior certificates by also certifying random and bursty loss, and an adversary combined with a lossy channel. Table 1 summarizes these gaps and where EventCert closes them.

3 THREAT MODEL

An event stream is a list of events (t_i, x_i, y_i, s_i): a time, a pixel position, and a polarity s_i (brightness up or down). A fixed preprocessing step turns the stream into a binary tensor $X \in \{0,1\}^{T\times C\times H\times W}$. An entry is 1 if at least one event falls into that time bin t, channel c, and pixel (h, w):

$$X[t, c, h, w] = \mathbf{1}\big[\exists i : \ \mathcal{B}(t_i) = t,\ (x_i, y_i) = (w, h),\ \pi(s_i) = c\big], \tag{1}$$

where $\mathcal{B}$ maps a time to its bin and π maps a polarity to its channel. All guarantees are stated for this tensor. We split the tensor into m non-overlapping SST groups; each group covers a few time bins, a block of pixels (or frequency channels for audio), and all channels. $G_i(X)$ denotes the entries of group i, and $\mathcal{D}_S(X)$ is the tensor with every group in a set $S \subseteq \{1, \dots, m\}$ set to zero. A group is *active* if it contains at least one event. The set of active groups is $\mathcal{N}(X) = \{i : G_i(X) \neq 0\}$. Deleting an inactive group does not change X. We consider two kinds of loss. (1) ***Worst-case loss.*** An adversary who knows X and the partition deletes any set S of at most k groups, i.e. it may produce any tensor in $\mathcal{A}_k(X) = \{\mathcal{D}_S(X) : |S| \leq k\}$. It can always choose active groups, so sparsity does not help against it. (2) ***Physical loss.*** A random process Π, such as a lossy channel, deletes a random set S. This set may contain active and inactive groups. What matters is the number of active groups it removes, $H = |S \cap \mathcal{N}(X)|$. EventCert first certifies against the adversary and then uses H to certify against physical loss.

4 EVENTCERT

EventCert keeps each group independently with probability p_{keep} and deletes it otherwise. For a random mask $M \in \{0,1\}^m$ with $M_i \sim \text{Bernoulli}(p_{\text{keep}})$, let $X^{(M)}$ be X with all unkept groups set to zero. The smoothing noise therefore has exactly the same form as the threat. Given a base classifier f, which we assume to be deterministic, the smoothed classifier g predicts the class that f outputs most often:

$$p_c(X) = \Pr_M\big[f(X^{(M)}) = c\big], \qquad g(X) = \arg\max_c p_c(X). \tag{2}$$

4.1 ESTIMATING CLASS PROBABILITIES

As in Cohen et al. (2019), we sample in two independent rounds. The first n_0 masks pick the likely top class $\hat{A}$. Then n fresh masks give a lower confidence bound $p^L_{\hat{A}}$ for $\hat{A}$ and upper bounds for all other classes; p^U_{comp} is the largest of these. We use Clopper–Pearson bounds (Clopper & Pearson, 1934) with a Bonferroni correction, so that all bounds hold together with probability at least $1 - \alpha$.

4.2 WORST-CASE CERTIFICATE

The key observation is simple. Deleting a set S of k groups and then smoothing gives the same output distribution as smoothing X and keeping only the masks that already delete all of S. Such masks occur with probability $q_k = (1 - p_{\text{keep}})^k$. Splitting all masks into those that delete S and those that do not gives $p_c(X) = q_k\, p_c(\mathcal{D}_S(X)) + (1 - q_k)\beta_c$ for some $\beta_c \in [0, 1]$. Using $\beta_c \geq 0$ for other classes and $\beta_{\hat{A}} \leq 1$ for $\hat{A}$:

$$p_c(\mathcal{D}_S(X)) \leq \frac{p_c(X)}{q_k}, \qquad p_{\hat{A}}(\mathcal{D}_S(X)) \geq \frac{p_{\hat{A}}(X) - (1 - q_k)}{q_k}. \tag{3}$$

Combining these bounds with the confidence bounds, $\hat{A}$ stays the top class after *any* deletion of at most k groups if

$$p^L_{\hat{A}} - p^U_{\text{comp}} > 1 - (1 - p_{\text{keep}})^k. \tag{4}$$

The right-hand side grows with k, so the certified ***worst-case radius*** is the largest k that satisfies the condition:

$$k_{\text{wc}}(X) = \max\big\{k \in \{0, \dots, m\} : \ p^L_{\hat{A}} - p^U_{\text{comp}} > 1 - (1 - p_{\text{keep}})^k\big\}. \tag{5}$$

If even $k = 0$ fails, EventCert abstains. With probability at least $1 - \alpha$, $g(\mathcal{D}_S(X)) = \hat{A}$ for every S with $|S| \le k_{\mathrm{wc}}(X)$ (proof in Appendix F). A small p_{keep} allows larger radii but makes the classifier see heavily deleted inputs; a large p_{keep} does the opposite.

4.3 From worst-case to random loss

Deleting an inactive group changes nothing, so $\mathcal{D}_S(X) = \mathcal{D}_{S\cap\mathcal{N}(X)}(X)$. Hence the prediction is safe whenever the number of deleted *active* groups, $H = |S \cap \mathcal{N}(X)|$, is at most $k_{\mathrm{wc}}(X)$. This does not enlarge the adversarial radius, because an adversary can pick active groups. For random loss, however, it gives a guarantee: if $\Pr_{S\sim\Pi}[H > k_{\mathrm{wc}}(X)] \le \delta$, the prediction survives the loss with probability at least $1 - \delta$. When r groups are lost uniformly at random, H follows a hypergeometric distribution. With $n_X = |\mathcal{N}(X)|$ and the *operational radius* r_δ defined as the largest certified r,

$$\Pr[H > k] = \sum_{j=k+1}^{\min(r,n_X)} \frac{\binom{n_X}{j}\binom{m-n_X}{r-j}}{\binom{m}{r}}, \tag{6}$$

$$r_\delta(X) = \max\{r : \Pr[H > k_{\mathrm{wc}}(X)] \le \delta\}. \tag{7}$$

Because random loss also hits inactive groups, r_δ can exceed k_{wc} for sparse inputs. For example, $k_{\mathrm{wc}} = 8$ gives $r_{0.05} = 20$ on the sparse N-TIDIGITS18 tensors but only 8 on the dense DVS-Gesture tensors (Appendix A).

4.4 Bursty channels

Real links lose packets in bursts. We send the groups in time order over a Gilbert channel (Gilbert, 1960; Elliott, 1963): a two-state Markov chain that loses every packet while it is in its "bad" state and none while it is "good". It has two parameters, the average loss rate ρ and the mean burst length L. Let $u_t = 1$ if the t-th group sent is active, and $B_t = 1$ if the channel is bad at step t. The number of lost active groups is $H = \sum_t u_t B_t$. We compute $\Pr[H > k]$ exactly by stepping through the m groups while tracking the channel state and the number of lost active groups (capped at $k + 1$), in $O(mk)$ time.

Proposition 1 (bursty loss). *If $\Pr_{\rho,L}[H > k_{\mathrm{wc}}(X)] \le \delta$, then g keeps its prediction under the channel's losses with probability at least $1 - \delta$ (at least $1 - \delta - \alpha$ when the sampling error is included).* We report the largest certified loss rate $\rho_\delta(X; L)$. Independent loss is the special case $H \sim \mathrm{Binomial}(n_X, \rho)$.

4.5 Hybrid threats

An adversary first deletes a groups of its choice. The channel then loses r of the remaining $m - a$ groups uniformly at random. If b of the adversary's groups are active, the total number of deleted active groups is $b + H'$, where $H' \sim \mathrm{Hypergeom}(m - a,\, n_X - b,\, r)$. This total can only grow with b, so the worst adversary uses $b = \min(a, n_X)$ active groups.

Proposition 2 (hybrid threat). *If $\Pr[b + H' > k_{\mathrm{wc}}(X)] \le \delta$ for $b = \min(a, n_X)$, then g keeps its prediction for every choice of a adversarial deletions followed by r random losses, with probability at least $1 - \delta$ (at least $1 - \delta - \alpha$ with sampling error).* Setting $a = 0$ recovers r_δ, and $r = 0$ recovers the worst-case certificate. All three certificates are computed after the fact from $k_{\mathrm{wc}}(X)$ and $\mathcal{N}(X)$, without evaluating the classifier again. Proofs are in Appendix F.11.

5 Experiments

5.1 Setup

We use three event-vision datasets, N-MNIST, DVS-Gesture (Orchard et al., 2015; Amir et al., 2017) and N-Caltech101 (Orchard et al., 2015), and two neuromorphic audio datasets, SHD (Cramer et al., 2022) and N-TIDIGITS18 (Anumula et al., 2018). Each stream becomes a binary tensor over a fixed time window. Each group covers two time bins, a block of pixels or frequency channels, and all channels (Appendix D). The classifiers are a small CNN for vision and a temporal CNN

Table 2: EventCert on five event benchmarks (VisionCNN/AudioCNN, group-mask training). Mean±std over seeds; accuracies in %. p_{keep} is the validation-selected retention probability (mode over seeds). ACR is the average certified radius (groups), $\tilde{k}_{\text{wc}}$ the median worst-case radius over correctly classified, certified test inputs, and $\bar{r}_{0.05}/\bar{k}_{\text{wc}}$ the ratio of mean operational to mean worst-case radius under uniform random group loss ($\delta = 0.05$). $|\mathcal{N}|/m$ is the median active-group fraction.

Dataset	m	$\lvert\mathcal{N}\rvert/m$	p_{keep}	Clean	Smoothed	Cert@1	Cert@8	ACR	$\tilde{k}_{\text{wc}}$	$\bar{r}_{0.05}/\bar{k}_{\text{wc}}$
N-MNIST	150	0.73	0.1	99.5±0.2	99.0±0.3	97.9±0.2	89.2±0.7	18.21±0.19	17	1.20
DVS-Gesture	384	1.00	0.1	77.0±3.5	70.2±1.8	61.5±1.4	33.6±2.7	7.59±0.81	8	1.00
N-Caltech101	384	0.75	0.1	53.5±0.7	46.5±0.9	42.5±0.6	34.5±0.5	10.97±0.11	27	1.32
SHD	350	0.61	0.3	85.0±0.8	82.1±2.0	63.8±1.5	6.1±0.9	2.49±0.07	3	1.11
N-TIDIGITS18	160	0.26	0.5	79.9±2.3	77.3±2.4	52.5±3.7	0.0±0.0	1.19±0.12	1	1.88

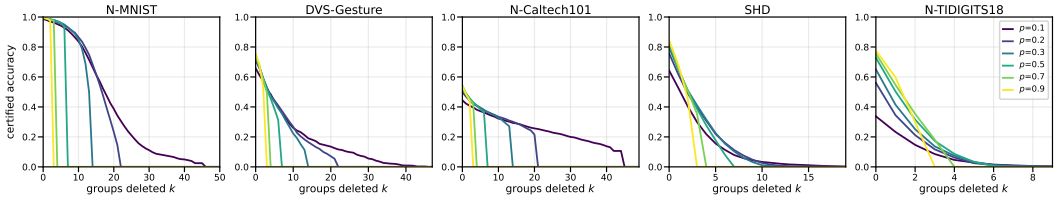

Figure 1: Certified accuracy against deletion of any k groups, for each p_{keep} (mean over seeds 0–2, which are certified at every p_{keep}). Smaller p_{keep} gives larger radii but lower accuracy.

for audio; the appendix adds a spiking network and a study of the group size. As is standard in randomized smoothing (Cohen et al., 2019), each model is trained on inputs with the same random deletions used for smoothing. Certification uses $n_0 = 128$, $n = 2048$, and $\alpha = 0.001$. We choose $p_{\text{keep}} \in \{0.1, 0.2, 0.3, 0.5, 0.7, 0.9\}$ on validation data to maximize the *average certified radius* (ACR: the mean of k_{wc}, counting wrong or abstained inputs as 0). *Cert@k* is the fraction of test inputs that are correct and certified for k groups. We certify a fixed random set of 1000 test inputs per dataset (all 264 for DVS-Gesture) and report mean±std over five seeds. Details are in Appendix C.

5.2 Certified robustness

Table 2 shows that EventCert certifies many inputs against multi-group deletion. On DVS-Gesture, 33.6% of test gestures are certified against the deletion of *any* 8 of the 384 groups, and on N-MNIST the average radius is 18.21 groups. Validation picks a small p_{keep} for vision and a larger one for audio, where each active group carries more of the signal. Fig. 1 shows the trade-off: a smaller p_{keep} certifies larger radii but lowers accuracy.

5.3 Comparison with existing certification methods

No earlier work certifies the deletion of events, so there are no published certified results on these benchmarks. We therefore apply established certification methods to the same data and compare them with EventCert under one protocol: the same models, tuning, and sampling budget, with each method trained under its own noise (Table 3; per-seed variation in Appendix D). Gaussian smoothing (Cohen et al., 2019) certifies an ℓ_2 ball, Bernoulli deletion as in RS-Del (Huang et al., 2023) certifies deleted tensor entries, and randomized ablation (Levine & Feizi, 2020b), applied to our groups, keeps a fixed number of groups (group ablation). EventCert has the highest ACR and Cert@1 on every dataset, with 3.8–5.6× the ACR of group ablation (Fig. 2a). Group ablation deletes the same units, but its bound needs about twice the margin ($2(1 - q_k)$ instead of $1 - q_k$). Gaussian smoothing and entry deletion do work: they certify an ℓ_2 radius of about 0.65 or up to tens of entries. But one active group contains more occupied entries than that (Table 22), so their certificates rarely cover even one whole group. Only EventCert also certifies random, bursty, and mixed physical loss.

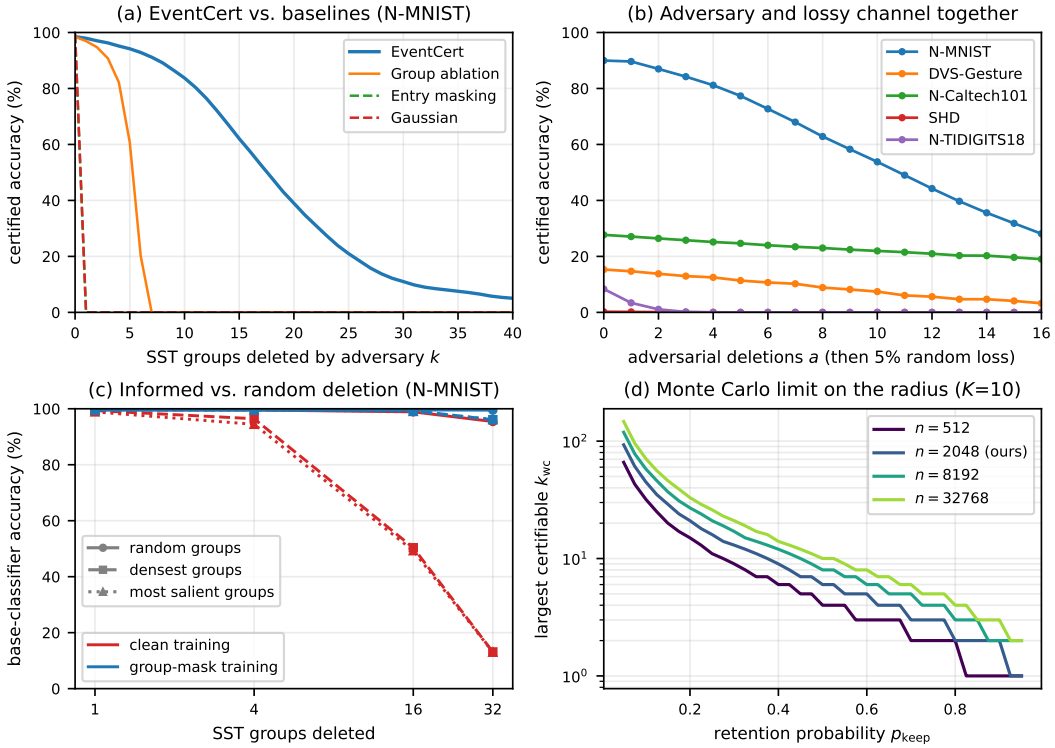

Figure 2: Summary of the main findings (mean over five seeds). **(a)** Certified accuracy against an adversary deleting k groups on N-MNIST: EventCert certifies far larger deletions than group ablation, while entry masking and Gaussian smoothing certify almost none. **(b)** Hybrid threat: certified accuracy when an adversary deletes a groups and 5% of the remaining groups are then lost at random. **(c)** Without certification, deleting the densest or most salient groups breaks a clean-trained model, while training with group deletion keeps it accurate. **(d)** The number of Monte Carlo samples n caps the certifiable radius; smaller p_{keep} and larger n allow larger radii.

Table 3: Comparison with established certification methods on the same data and protocol. Each cell: average certified radius in groups (Cert@1 in %), mean over five seeds; best in bold. Only EventCert also certifies physical loss.

Method	**N-MNIST**	**DVS-Gesture**	**N-Caltech101**	**SHD**	**N-TIDIGITS18**	**Physical loss**
Gaussian smoothing (Cohen et al., 2019)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	×
Bernoulli deletion, RS-Del (Huang et al., 2023)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.03 (2.6)	0.09 (7.5)	×
Randomized ablation (Levine & Feizi, 2020b)	4.45 (96.9)	1.92 (55.1)	1.97 (40.6)	0.66 (45.4)	0.22 (18.7)	×
EventCert (ours)	**18.21** (**97.9**)	**7.59** (**61.5**)	**10.97** (**42.5**)	**2.49** (**63.8**)	**1.19** (**52.5**)	✓

5.4 Clean training is not enough

The certificate only needs class predictions, so it also applies to a model trained on clean tensors (Table 4). Such a model, however, breaks down when most groups are deleted: its Cert@1 is only 8.6–20.8%, against 42.5–97.9% after training with group deletion. As with Gaussian smoothing, useful certificates need training under the smoothing noise.

Table 4: EventCert applied post hoc to classifiers trained on clean tensors versus classifiers trained with SST group masks. The certificate is valid for both. Accuracies in %, mean±std over seeds.

	Clean training (post hoc)				Group-mask training		
Dataset	Clean	Smoothed	Cert@1	ACR	Smoothed	Cert@1	ACR
N-MNIST	99.6±0.1	11.3±0.7	11.0±0.0	4.08±1.31	99.0±0.3	97.9±0.2	18.21±0.19
DVS-Gesture	79.7±1.7	10.2±1.7	9.7±1.2	3.53±1.16	70.2±1.8	61.5±1.4	7.59±0.81
N-Caltech101	58.3±0.9	8.9±0.5	8.6±0.1	2.99±1.03	46.5±0.9	42.5±0.6	10.97±0.11
SHD	86.8±0.7	84.8±0.7	20.8±1.8	0.23±0.02	82.1±2.0	63.8±1.5	2.49±0.07
N-TIDIGITS18	82.9±0.8	68.0±31.4	17.1±5.5	0.21±0.03	77.3±2.4	52.5±3.7	1.19±0.12

Table 5: Operational certified accuracy (%) at $\delta = 0.05$, computed exactly from each input's k_{wc} and active set. Channel: each group is lost with probability 5% (independently, or in Gilbert bursts of mean length L). Uniform/hybrid: exactly $\mathrm{round}(0.05(m-a))$ groups are lost uniformly without replacement after an informed adversary deletes a groups (a=0: random loss only). The two random models differ (random vs. fixed number of losses), so only columns within a group are directly comparable. Mean over seeds.

	Worst case	Channel loss 5%			Uniform/hybrid (a + 5% random)		
Dataset	Cert@1	indep.	L=4	L=16	a=0	a=2	a=4
N-MNIST	97.9	85.6	61.1	14.5	90.0	87.0	81.2
DVS-Gesture	61.5	10.2	1.1	0.0	15.3	13.8	12.5
N-Caltech101	42.5	25.0	20.0	8.2	27.7	26.4	25.1
SHD	63.8	0.1	0.0	0.0	0.2	0.1	0.0
N-TIDIGITS18	52.5	6.1	0.0	0.0	8.4	1.1	0.0

5.5 Random, bursty, and hybrid loss

Table 5 and Fig. 3 apply the certificates of Section 4 to every test input. First, random loss is well covered for N-MNIST: 85.6% of inputs are certified when each group is lost independently with probability 5%. The gain over the worst-case radius is largest for the sparsest data (N-TIDIGITS18, last column of Table 2) and absent for the fully dense DVS-Gesture tensors. Second, *bursts matter*: at the same 5% loss rate, certified accuracy on N-MNIST falls to 61.1% for bursts of mean length 4 and to 14.5% for length 16, because one burst removes many neighbouring groups at once. A certificate that assumes independent loss would overstate robustness several times over. Third, the guarantees combine: 81.2% of N-MNIST inputs stay certified when an adversary deletes 4 groups and 5% of the rest are then lost at random, and certified accuracy decreases smoothly as the adversary deletes more groups (Fig. 2b). For dense data and for small radii (audio), the operational certificates are correspondingly weaker.

5.6 Stress tests and sanity check

Without a certificate, informed deletions are much more harmful than random ones: a clean-trained N-MNIST model keeps 98.9% accuracy when 16 random groups are deleted, but only 50.4% when the 16 fullest groups are deleted (Fig. 2c, Table 21). This is why EventCert certifies against an informed adversary. Training with group deletion closes most of this gap (99.3%). As a sanity check, deleting exactly k_{wc} of the most salient or fullest active groups from 966 certified test inputs changed 0 smoothed predictions (Table 23).

6 Limitations and Conclusion

Event cameras are increasingly used in safety-critical systems, yet packet loss, sensor faults, and occlusion routinely remove whole blocks of their data. Because event data has little redundancy, such losses can silently change a classifier's prediction, and until now there was no way to prove how much loss a prediction can survive.

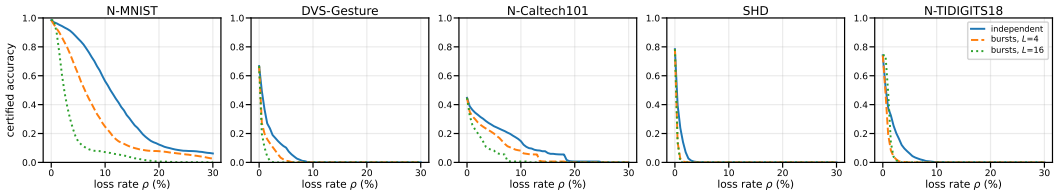

Figure 3: Certified accuracy ($\delta = 0.05$) against the channel loss rate ρ, for independent loss and for bursts of mean length L. Bursts sharply reduce the loss rate that can be certified.

EventCert resolves this problem. On the theoretical side, smoothing with the same group deletions as the threat yields a finite-sample certificate against an informed adversary (Section 4). Since only event-carrying groups count against this certificate, it extends to exact guarantees for random, bursty, and combined adversarial-and-random loss, all proved in Appendix F and computed from a single certification run. The experiments confirm that these guarantees are both strong and needed. EventCert certifies 3.8–5.6$\times$ larger radii than the strongest existing method on five benchmarks, keeps 89% of N-MNIST inputs certified when an adversary deletes any 8 groups, and changed none of 966 certified predictions under informed attacks. It also reveals that guarantees assuming independent loss are unsafe: at a 5% loss rate, certified accuracy drops from 85.6% to 14.5% once losses arrive in bursts.

EventCert also has limitations. It certifies deletions after tensorization only, not added events, time shifts, or attacks on the preprocessing step. Its channels are simulated at the tensor level rather than measured on real links, and deletions that cut through a group are covered only if the whole group is treated as deleted. The number of Monte Carlo samples caps the certified radius, and to keep certification affordable on one GPU we use compact models and a 1000-input test subset, which leaves radii on the audio datasets small. These limitations point to clear next steps: certificates for inserted and time-shifted events, calibration against measured network traces and packet-to-group mappings of real systems, and more efficient sampling that allows larger models and larger certified radii. Together, these would move EventCert from benchmark evidence to deployed systems, helping self-driving cars, drones, and robots know when their event-camera decisions can still be trusted.

Reproducibility Statement

The certificate is fully specified in Section 4 with complete proofs in Appendix F. Datasets, tensorization, SST partitions, architectures, training, smoothing-parameter selection, and certification settings are listed in Appendix C. The anonymous supplementary archive contains all code, dataset preparation scripts, the job queues used for every run, per-seed logs with the Monte Carlo class counts of every certified test input, and the scripts that regenerate every table, figure, and number in this paper from those logs (Appendix B).

Ethics Statement

This work uses only public benchmark datasets and involves no human-subject data collection. It studies certified robustness for event-based perception; the empirical deletion stress tests are evaluation tools and do not constitute a practical attack on deployed systems. Certificates hold only for the stated threat model and should not be read as general safety guarantees.

AI use statement

We used a large language model (LLM) coding and writing assistant during this work. The research problem, the EventCert certification approach (structured spatio-temporal group deletion with a worst-case certificate and its connection to physical loss), the choice of datasets, and the original manuscript were developed by the authors. The LLM assisted with (i) implementing, debugging, and running the experimental pipeline, including the re-run evaluation reported in this version, and writing scripts that generate all tables and figures from the stored run logs; (ii) developing and proving the extensions to bursty-channel and hybrid adversarial-plus-random loss (Propositions 1–2), which are covered by unit tests in the supplement; (iii) finding and checking published results cited for comparison; and (iv)

substantially editing and restructuring the manuscript for clarity. All reported numbers are generated programmatically from logged runs included in the supplement, and all AI-assisted derivations, code, and text were reviewed by the authors. The authors take full responsibility for the content of this paper, including all claims, proofs, and artifacts.

REFERENCES

Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, et al. A low power, fully event-based gesture recognition system. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7243–7252, 2017.

Jithendar Anumula, Daniel Neil, Tobi Delbruck, and Shih-Chii Liu. Feature representations for neuromorphic audio spike streams. *Frontiers in neuroscience*, 12:23, 2018.

Charles J Clopper and Egon S Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934.

Jeremy M. Cohen, Elan Rosenfeld, and J. Zico Kolter. Certified adversarial robustness via randomized smoothing. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1310–1320. PMLR, 2019.

Benjamin Cramer, Yannik Stradmann, Johannes Schemmel, and Friedemann Zenke. The Heidelberg spiking data sets for the systematic evaluation of spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(7):2744–2757, 2022.

Jianhao Ding, Zhaofei Yu, Jian K. Liu, and Tiejun Huang. Neuromorphic computing paradigms enhance robustness through spiking neural networks. *Nature Communications*, 16:10175, 2025. doi: 10.1038/s41467-025-65197-x.

E. O. Elliott. Estimates of error rates for codes on burst-noise channels. *Bell System Technical Journal*, 42(5):1977–1997, 1963.

Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180, 2022.

Edgar N. Gilbert. Capacity of a burst-noise channel. *Bell System Technical Journal*, 39(5):1253–1265, 1960.

Zhuoqun Huang, Neil G Marchant, Keane Lucas, Lujo Bauer, Olya Ohrimenko, and Benjamin I. P. Rubinstein. RS-Del: Edit distance robustness certificates for sequence classifiers via randomized deletion. *Advances in Neural Information Processing Systems*, 36:18676–18711, 2023.

Alexander Levine and Soheil Feizi. (De)Randomized smoothing for certifiable defense against patch attacks. In *Advances in Neural Information Processing Systems*, volume 33, 2020a.

Alexander Levine and Soheil Feizi. Robustness certificates for sparse adversarial attacks by randomized ablation. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, pp. 4585–4593, 2020b.

Hongbin Liu, Jinyuan Jia, and Neil Zhenqiang Gong. PointGuard: Provably robust 3D point cloud classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6186–6195, 2021. URL https://openaccess.thecvf.com/content/CVPR2021/html/Liu_PointGuard_Provably_Robust_3D_Point_Cloud_Classification_CVPR_2021_paper.html.

Bhaskar Mukhoty, Hilal AlQuabeh, Giulia De Masi, Huan Xiong, and Bin Gu. Certified adversarial robustness for rate encoded spiking neural networks. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=5bNYf0CqxY.

Garrick Orchard, Ajinkya Jayawant, Gregory K Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in neuroscience*, 9:437, 2015.

Hadi Salman, Jerry Li, Ilya Razenshteyn, Pengchuan Zhang, Huan Zhang, Sebastien Bubeck, and Greg Yang. Provably robust deep learning via adversarially trained smoothed classifiers. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Yan Scholten, Jan Schuchardt, Aleksandar Bojchevski, and Stephan Günnemann. Hierarchical randomized smoothing. In *Advances in Neural Information Processing Systems*, 2023. URL https://papers.neurips.cc/paper_files/paper/2023/hash/9c0efc0d84c263972af72bf70a2de533-Abstract-Conference.html.

Haoyang Wang, Ruishan Guo, Pengtao Ma, Ciyu Ruan, Xinyu Luo, Wenhua Ding, Tianyang Zhong, Jingao Xu, Yunhao Liu, and Xinlei Chen. Event camera meets mobile embodied perception: abstraction, algorithm, acceleration, application. *ACM Computing Surveys*, 58(8):1–41, 2026.

Keming Wu, Man Yao, Yuhong Chou, Xuerui Qiu, Rui Yang, Bo Xu, and Guoqi Li. RSC-SNN: Exploring the trade-off between adversarial robustness and accuracy in spiking neural networks via randomized smoothing coding. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 2748–2756, 2024. doi: 10.1145/3664647.3680639.

Yanmeng Yao, Xiaohan Zhao, and Bin Gu. Exploring vulnerabilities in spiking neural networks: Direct adversarial attacks on raw event data. In *Computer Vision – ECCV 2024*, volume 15130 of *Lecture Notes in Computer Science*, pp. 412–428. Springer, 2024. doi: 10.1007/978-3-031-73220-1_24.

Jinghuai Zhang, Jinyuan Jia, Hongbin Liu, and Neil Zhenqiang Gong. PointCert: Point cloud classification with deterministic certified robustness guarantees. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9496–9505, 2023. URL https://arxiv.org/abs/2303.01959.

A Additional Method Details

A.1 Certification procedure

Algorithm 1 gives one certification. It draws n_0 masks to select the most frequent class $\hat{A}$, draws n fresh masks to bound the class probabilities with simultaneous Clopper–Pearson intervals, abstains if $\hat{A}$ cannot be certified even for $k = 0$, and otherwise returns the worst-case radius k_{wc} of Eq. 5. The random, bursty, and hybrid certificates (Eq. 7 and Propositions 1–2) are then computed from k_{wc} and the active set $\mathcal{N}(X)$ alone, without further classifier evaluations.

Algorithm 1 EventCert certification

Require: classifier f, input X, m groups, p_{keep}, samples (n_0, n), confidence α, optional loss process Π and level δ
Ensure: predicted class $\hat{A}$, worst-case radius k_{wc}, optional operational radius r_δ
Draw n_0 masks; $\hat{A} \leftarrow$ most frequent prediction
Draw n fresh masks; count predictions $C[c]$ for every class $c \in \{1, \dots, K\}$
$p^L_{\hat{A}} \leftarrow \mathrm{CPLower}(C[\hat{A}], n, \alpha/2)$; $p^U_{\mathrm{comp}} \leftarrow \max_{c \neq \hat{A}} \mathrm{CPUpper}(C[c], n, \alpha/[2(K-1)])$
if $p^L_{\hat{A}} \leq p^U_{\mathrm{comp}}$ **then**
    return abstain
end if
$k_{\mathrm{wc}} \leftarrow \max\{k : p^L_{\hat{A}} - p^U_{\mathrm{comp}} > 1 - (1 - p_{\mathrm{keep}})^k\}$
if Π is given **then**
    $r_\delta \leftarrow \max\{r : \Pr_{S \sim \Pi_r}[|S \cap \mathcal{N}(X)| > k_{\mathrm{wc}}] \leq \delta\}$
end if
return $(\hat{A}, k_{\mathrm{wc}}, r_\delta)$

A.2 How sparsity enlarges the operational radius

Random loss also hits groups without events, so it rarely removes more than k_{wc} event-carrying groups when few groups are active. Fig. 4 evaluates Eq. 6 exactly with the median active fraction

$|\mathcal{N}|/m$ of each dataset. At $k_{\mathrm{wc}} = 8$, it gives $r_{0.05} = 20$ groups on N-TIDIGITS18 ($|\mathcal{N}|/m = 0.26$) and 10 on SHD (0.61), but only 8 on N-MNIST (0.73), N-Caltech101 (0.75), and the fully dense DVS-Gesture (1.00). The gain depends on both the active fraction and k_{wc}, and never enlarges the adversarial radius.

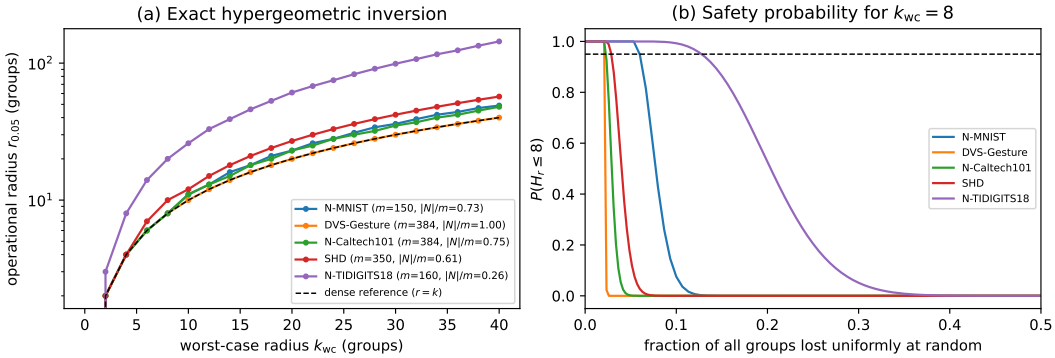

Figure 4: Operational versus worst-case radius under uniform random group loss, computed exactly from the measured active fractions $|\mathcal{N}|/m$. *(a)* $r_{0.05}$ as a function of k_{wc} (log scale). *(b)* $\Pr[H_r \leq 8]$ as the fraction of lost groups grows.

B Reproducibility Details

The certificate is fully specified by Eqs. 4–6 and Appendix F; an independent implementation needs only the class counts, confidence level, retention probability, SST partition, and physical-loss model. Every certified radius reported in this paper can be recomputed from the tuple

$$(C_{\hat{A}}, \{C_c\}_{c \neq \hat{A}}, n, \alpha, p_{\mathrm{keep}}, m, |\mathcal{N}(X)|, \Pi_r),$$

which the artifact stores for every test input.

B.1 Anonymous artifact

The supplementary package contains the certification code (confidence bounds, worst-case and operational radii), the SST partition and deletion operators, all baseline smoothing schemes and their group conversions, model definitions, dataset tensorization scripts with fixed configurations for all five datasets, the training/certification runner, the job queue used for all runs, aggregation scripts that generate every table and figure in this paper directly from the run logs, unit tests, and a frozen environment. For every run it stores a manifest (arguments, data hashes, software versions), the training history, the selected checkpoint, split indices, validation tuning scores, and one log line per test input with the Monte Carlo class counts from both stages.

B.2 Compute environment

All experiments use Python 3.11, PyTorch 2, Tonic, NumPy, SciPy, and matplotlib on a single Apple M5 Pro laptop GPU (PyTorch MPS backend, 24 GB unified memory). Dependency versions are frozen in the artifact. The 125 runs of the main evaluation took about 4.5 hours of wall-clock time on this machine with three to eight concurrent jobs.

C Experimental Protocol and Additional Details

This appendix provides the dataset, model, training, tensorization, certification, and baseline settings used in the empirical evaluation. These details complement the reproduction workflow in Appendix B and define the exact experimental scope of EventCert.

Table 6: Contents of the anonymous EventCert artifact.

Component	Role
Certification implementation	Clopper–Pearson bounds, k_{wc}, exact hypergeometric r_δ.
Tensorization/group utilities	Fixed binary tensorization, SST partition, deletion operator.
Baselines	Entry masking, Gaussian smoothing, fixed-cardinality group ablation, with group conversions.
Stress-test utilities	Informed and geometric deletion processes.
Models and training	Architectures, group-mask (noise-matched) training, checkpoint selection.
Configurations	Tensor shape, time window, and SST partition for all five datasets.
Run logs and checkpoints	Per-seed manifests, tuning scores, per-input class counts, selected weights.
Aggregation scripts	Regenerate every table and figure from the run logs.
Tests	Exhaustive small-classifier deletion checks, hypergeometric enumeration, bound coverage.

C.1 DATASETS

EventCert is evaluated on five public event-based datasets. The datasets cover both event-vision and neuromorphic audio-like streams, allowing the certification behavior to be tested across different sparsity patterns and temporal structures.

Table 7: Datasets used in the EventCert evaluation.

Dataset	Task type	Role in evaluation	Split policy
N-MNIST (Orchard et al., 2015)	Event vision (10 classes)	Certification, baselines, stress tests	Standard train/test
DVS-Gesture (Amir et al., 2017)	Event vision (11 classes)	Certification, baselines, stress tests	Standard train/test
N-Caltech101 (Orchard et al., 2015)	Event vision (101 classes)	Harder, many-class setting	Stratified 80/20 (seed 2026)
SHD (Cramer et al., 2022)	Neuromorphic audio (20 classes)	Certification, baselines, stress tests	Standard train/test
N-TIDIGITS18 (Anumula et al., 2018)	Neuromorphic audio (11 classes)	Sparse temporal setting	Standard train/test

C.2 TENSORIZATION AND GROUP CONSTRUCTION

All event streams are converted into binary tensors using the fixed preprocessing operator in Eq. 1. The same tensorization, padding, and group construction rule is used for EventCert and all baselines within a dataset. Each structured spatio-temporal group spans a local temporal window, a spatial block, and all polarity channels.

Table 8: Tensorization and SST group construction.

Item	Setting
Input representation	Binary event tensor $X \in \{0,1\}^{T\times C\times H\times W}$
Temporal handling	Events are assigned to fixed temporal bins using $\mathcal{B}(\cdot)$
Polarity handling	Event polarity is mapped to channel index using $\pi(\cdot)$
Padding rule	Tensors are padded before grouping when needed so that groups are disjoint and cover the tensor
Group structure	Each SST group spans a local temporal window, a spatial block, and all polarity channels
Deletion operator	$\mathcal{D}_S(X)$ zeros out all groups indexed by S and leaves the remaining groups unchanged
Active groups	$\mathcal{N}(X) = \{i \in \mathcal{I} : G_i(X) \neq 0\}$
Effective deletion count	$h_X(S) = \lvert S \cap \mathcal{N}(X)\rvert$

C.3 MODEL ARCHITECTURES

The certificate is classifier-agnostic. We use one compact convolutional model per modality. Models are deliberately small because certification requires $n_0 + n = 2176$ forward passes per input and retention value.

Table 9: Model families used in the evaluation.

Model	Used for	Definition
VisionCNN	N-MNIST, DVS-Gesture, N-Caltech101	Time bins stacked as channels; five 3×3 conv–BN–ReLU layers (32/32/64/64/128), max pooling, global average pooling, dropout 0.2 ($\approx$0.15M parameters)
AudioCNN	SHD, N-TIDIGITS18	Three temporal conv–BN–ReLU layers (128 channels, kernel 3) over time bins, global average pooling, dropout 0.2

C.4 TRAINING AND MODEL SELECTION

Unless stated otherwise, each classifier is trained under the corruption family of the smoothing scheme it will be certified with (*noise-matched training*, standard in randomized smoothing (Cohen et al., 2019; Levine & Feizi, 2020b)): for every training example, a parameter is drawn uniformly from that scheme's selection grid and the corresponding corruption is applied. For EventCert this means training on inputs with independently deleted SST groups. We also report EventCert on classifiers trained on clean tensors (post-hoc setting, Table 4).

Table 10: Training and model-selection settings.

Item	Setting
Optimizer	AdamW, learning rate 10^{-3}, weight decay 10^{-4}, cosine schedule, batch 128
Epochs (early-stopping patience)	N-MNIST 15 (5); N-Caltech101, SHD, N-TIDIGITS18 40 (10); DVS-Gesture 60 (15)
Random seeds	$\{0, 1, 2, 3, 4\}$
Validation split	Stratified 10% of the training split
Checkpoint selection	Best validation accuracy under the training corruption
Smoothing-parameter selection	Validation average certified radius (ACR); ties by Cert@1
Test split use	Final reporting only
Summary statistic	Mean $\pm$ standard deviation over five seeds unless a caption states otherwise

C.5 CERTIFICATION SETTINGS

Certification uses two independent Monte Carlo stages. The first selects the candidate class and the second estimates simultaneous confidence bounds for the candidate and all competing classes. All smoothing schemes use the same Monte Carlo budget and confidence level. Following common practice for Monte Carlo certification (Cohen et al., 2019), we certify a fixed random subset of at most 1000 test inputs per dataset (identical across seeds and methods; all 264 test inputs for DVS-Gesture); validation tuning uses 300 validation inputs and $n = 512$.

C.6 BASELINE ADAPTATION

All baselines use the same tensors, SST partition, architecture, noise-matched training, validation selection, and Monte Carlo budget as EventCert; only the smoothing distribution and certificate differ. Certificates stated in other units are converted to whole-group deletions conservatively: an adversary deleting k groups is assumed to delete the k groups with the most occupied entries.

Retaining SST groups at fixed positions and replacing deleted groups by zeros induces exactly the EventCert smoothing distribution; such a group-token deletion variant is therefore not a separate baseline.

Table 11: Certification protocol used across EventCert and baselines.

Item	Setting
Class-selection samples	$n_0 = 128$
Confidence-estimation samples	$n = 2048$ (tuning: 512)
Failure probability	$\alpha = 10^{-3}$ per input
Selected-class lower bound	Clopper–Pearson lower bound with $\alpha/2$
Competing-class upper bounds	Bonferroni-corrected Clopper–Pearson upper bounds with $\alpha/[2(K-1)]$
Retention-probability grid	$p_{\text{keep}} \in \{0.1, 0.2, 0.3, 0.5, 0.7, 0.9\}$ (EventCert, group ablation, entry masking)
Gaussian noise grid	$\sigma \in \{0.25, 0.5, 1, 2\}$
Worst-case radius	k_{wc} from Eq. 5
Operational radius	$r_{0.05}$ for uniform random group loss (Eq. 6), $\delta = 0.05$
Bursty-channel certificate	largest certified loss rate $\rho_{0.05}$; independent loss and Gilbert bursts with $L \in \{4, 16\}$, time-major order
Hybrid certificate	$a \in \{0, 2, 4\}$ adversarial deletions followed by $\text{round}(0.05(m-a))$ uniform random losses
Test inputs	fixed subset of ≤ 1000 per dataset (all 264 for DVS-Gesture)

Table 12: Baselines and their conversion to SST group deletion.

Method	Smoothing and certificate
Entry masking	Independent Bernoulli retention of tensor entries (event-level Bernoulli deletion; the binary-tensor analogue of RS-Del (Huang et al., 2023)). Certifies k_e deleted entries; converted to the largest k such that the k fullest groups contain at most k_e occupied entries.
Gaussian smoothing	Isotropic Gaussian noise (Cohen et al., 2019) with the multiclass ℓ_2 radius from the same simultaneous bounds. For binary tensors, deleting a set of groups has ℓ_2 norm $\sqrt{\text{occupied entries}}$; converted as above.
Group ablation	Retains exactly $s = \text{round}(m\, p_{\text{keep}})$ uniformly chosen groups (fixed-cardinality ablation in the style of (Levine & Feizi, 2020b)); certifies k if the margin exceeds $2(1 - \binom{m-k}{s}/\binom{m}{s})$.
EventCert	Independent Bernoulli retention of SST groups; certificate of Eq. 5.

C.7 Empirical deletion protocols

We use eight empirical deletion processes as stress tests of the base classifier (Table 21); they are not certificates. Four vary how informed the deletion is, and four impose a spatio-temporal geometry. The certified operational processes (uniform loss, Gilbert bursts, hybrid) are listed in Table 11.

Table 13: Empirical deletion processes used for stress testing.

Protocol	Definition
Random raw deletion	Samples deleted groups from all groups in $\mathcal{I}$; may remove active or inactive groups.
Random active deletion	Samples deleted groups only from $\mathcal{N}(X)$.
Density-based deletion	Ranks active groups by number of occupied entries and deletes the fullest groups first.
Saliency-guided deletion	Ranks active groups by gradient sensitivity and deletes high-saliency groups first.
Temporal burst deletion	Removes contiguous groups along the temporal dimension.
Spatial block deletion	Removes coherent sensor-region groups.
Spatio-temporal cuboid deletion	Removes a spatial region over a temporal interval.
Packet-burst deletion	Removes short burst-like groups that approximate packet-loss behavior.

C.8 Reported metrics

Table 14: Metrics reported in the main paper and appendix.

Metric	Meaning
Clean accuracy	Accuracy of the base classifier on unmodified test tensors.
Smoothed accuracy	Fraction of test inputs whose selected class $\hat{A}$ is correct.
Cert@k	Fraction of test inputs that are correct and certified with $k_{\mathrm{wc}} \geq k$.
ACR	Average certified radius: mean of k_{wc} over test inputs, counting incorrect or abstained inputs as 0.
$\tilde{k}_{\mathrm{wc}}$	Median k_{wc} over correct, certified test inputs.
$r_{0.05}$	Largest number of uniformly lost groups certified at failure probability $\delta = 0.05$.
Stress accuracy	Base-classifier accuracy under empirical deletion processes (no certificate).
Time/1000	Wall-clock certification time per 1000 inputs.

D Full Empirical Results

All numbers in this appendix are generated from the stored run logs by the aggregation script in the artifact.

Table 15: Tensorization and SST partition per dataset. All groups span every polarity channel.

Dataset	Tensor (T, C, H, W)	Group $(\Delta t, \Delta h, \Delta w)$	Window	m	Test N
N-MNIST	(12,2,34,34)	(2,7,7)	0.31 s	150	1000
DVS-Gesture	(12,2,64,64)	(2,8,8)	1.8 s	384	264
N-Caltech101	(12,2,64,64)	(2,8,8)	0.31 s	384	1000
SHD	(20,1,1,700)	(2,1,20)	1 s	350	1000
N-TIDIGITS18	(20,1,1,64)	(2,1,4)	1 s	160	1000

Table 16: Largest attainable worst-case radius k_{wc} (unanimous votes, $n = 2048$, $\alpha = 10^{-3}$). Regardless of the classifier, the Monte Carlo budget caps the radius at these values; lower retention probabilities permit larger radii but require classifiers that remain accurate under heavier masking.

Classes	p_{keep}=0.1	p_{keep}=0.2	p_{keep}=0.3	p_{keep}=0.5	p_{keep}=0.7	p_{keep}=0.9
10	45	21	13	6	3	2
11	45	21	13	6	3	2
20	44	21	13	6	3	2
101	44	20	13	6	3	2

D.1 Comparison with certification methods: per-seed variation

Table 17 gives the per-seed variation (mean$\pm$std) of the comparison in Table 3.

D.2 Certified accuracy at all budgets

Table 18 lists EventCert's certified accuracy at every budget k, together with the abstention rate and certification time.

D.3 Spiking classifier

EventCert only needs class predictions, so it applies unchanged to spiking networks. Table 19 certifies a leaky integrate-and-fire classifier (128 LIF neurons, leak 0.9, threshold 1, surrogate-gradient training

Table 17: Certified robustness to SST group deletion for smoothing schemes trained and tuned identically (same architecture, noise-matched training, validation-selected parameter, n_0=128, n=2048, α=10^{-3}). Entry and Gaussian certificates are converted to whole-group deletions conservatively (Appendix C). Cert@1 in %, ACR in groups; mean±std over seeds.

	EventCert		Group ablation		Entry masking		Gaussian	
Dataset	Cert@1	ACR	Cert@1	ACR	Cert@1	ACR	Cert@1	ACR
N-MNIST	**97.9±0.2**	**18.21±0.19**	96.9±0.3	4.45±0.03	0.0±0.0	0.00±0.00	0.0±0.0	0.00±0.00
DVS-Gesture	**61.5±1.4**	**7.59±0.81**	55.1±3.4	1.92±0.12	0.0±0.0	0.00±0.00	0.0±0.0	0.00±0.00
N-Caltech101	**42.5±0.6**	**10.97±0.11**	40.6±1.2	1.97±0.03	0.0±0.0	0.00±0.00	0.0±0.0	0.00±0.00
SHD	**63.8±1.5**	**2.49±0.07**	45.4±3.3	0.66±0.03	2.6±1.3	0.03±0.01	0.0±0.0	0.00±0.00
N-TIDIGITS18	**52.5±3.7**	**1.19±0.12**	18.7±4.0	0.22±0.03	7.5±0.7	0.09±0.01	0.0±0.0	0.00±0.00

Table 18: EventCert certified accuracy (%) at every budget k (VisionCNN/AudioCNN, selected p_{keep}), abstention rate, and mean certification time. Mean over seeds.

Dataset	Cert@1	Cert@2	Cert@4	Cert@8	Cert@16	Cert@32	Abstain	s/1000
N-MNIST	97.9	97.0	95.2	89.2	57.5	9.0	0.8	400
DVS-Gesture	61.5	58.1	49.5	33.6	18.3	4.7	8.3	699
N-Caltech101	42.5	41.0	38.3	34.5	28.8	18.8	12.7	606
SHD	63.8	50.6	31.4	6.1	0.1	0.0	9.0	153
N-TIDIGITS18	52.5	33.6	9.4	0.0	0.0	0.0	7.4	62

with SST group masks, membrane state reset for every smoothing sample) under the same protocol as the primary models.

Table 19: EventCert with a leaky integrate-and-fire spiking classifier (membrane state reset for every smoothing sample). Mean±std over seeds; accuracies in %.

Dataset	Clean	Smoothed	Cert@1	ACR
N-MNIST	95.0±0.3	91.6±0.2	84.3±0.3	7.75±0.11
DVS-Gesture	61.2±2.6	62.2±4.6	47.9±4.1	3.88±0.51
N-Caltech101	41.2±1.3	35.5±0.8	33.0±0.8	9.11±0.15
SHD	51.9±0.8	26.5±1.3	13.0±0.3	0.40±0.02
N-TIDIGITS18	38.0±4.2	14.8±3.7	9.9±0.9	0.67±0.15

D.4 Group granularity

The SST partition fixes the unit of the certificate: larger groups make each certified group cover more of the input but reduce the number of groups m. Holding tensorization, architecture, training, and selection fixed, Table 20 varies the group size on DVS-Gesture and SHD; each partition trains its own group-masked classifier. Clean accuracy and Cert@1 change little across partitions. Finer partitions certify more groups, but each group covers less of the tensor; measured in covered tensor volume (ACR times group volume), coarser partitions certify more. The partition should therefore be chosen to match the physical loss unit (for example, the tensor region carried by one packet) rather than to maximize the radius in groups.

D.5 Empirical stress tests

The stress tests evaluate the deterministic base classifier (no certificate) under deletion processes of increasing information. *Random* deletes uniformly chosen groups; *active* deletes uniformly chosen active groups; *density* deletes the active groups with the most occupied entries; *saliency* deletes the active groups with the largest summed input-times-gradient magnitude of the true-class loss. Geometric processes delete a temporally contiguous run of groups (*temporal*), a spatially contiguous block (*spatial*), a compact spatio-temporal box (*cuboid*), or a contiguous run in flattened group order

Table 20: Effect of the SST group size (group-mask training, validation-selected p_{keep}). Mean±std over seeds; accuracies in %.

Dataset	Group $(\Delta t, \Delta h, \Delta w)$	m	Clean	Cert@1	ACR
DVS-Gesture	(4,16,16)	48	75.5±9.3	59.4±6.8	3.39±0.46
DVS-Gesture	(2,16,16)	96	74.7±3.0	57.0±2.0	5.27±0.34
DVS-Gesture	(2,8,8)	384	77.0±3.5	61.5±1.4	7.59±0.81
DVS-Gesture	(1,8,8)	768	72.9±5.8	58.7±4.0	9.54±1.07
SHD	(4,1,50)	70	83.9±0.7	63.0±2.0	1.55±0.03
SHD	(2,1,50)	140	81.6±0.8	59.0±1.3	1.64±0.04
SHD	(2,1,20)	350	85.0±0.8	63.8±1.5	2.49±0.07
SHD	(1,1,20)	700	84.1±0.4	62.4±0.7	2.87±0.05

(*packet*). Randomized processes use four draws per input; actual effective active-group counts are logged.

Table 21: Base-classifier accuracy (%) after deleting $b = 16$ SST groups with each process, for classifiers trained on clean tensors and with SST group masks. Mean over seeds.

Dataset	Training	Clean	Random	Active	Density	Saliency	Temporal	Spatial	Cuboid	Packet
N-MNIST	clean	99.6	98.9	97.7	50.4	49.3	96.8	94.7	92.7	97.6
N-MNIST	group masks	99.5	99.6	99.6	99.3	99.0	99.4	99.5	99.5	99.5
DVS-Gesture	clean	79.7	78.7	78.1	57.6	58.9	76.1	72.1	75.3	75.4
DVS-Gesture	group masks	77.0	76.4	76.6	71.8	73.4	75.9	75.9	75.9	76.3
N-Caltech101	clean	58.3	54.6	52.9	37.9	36.6	56.0	47.2	54.0	50.9
N-Caltech101	group masks	53.5	53.9	53.8	53.4	53.6	53.5	53.7	53.6	53.7
SHD	clean	86.8	77.9	73.4	33.0	27.7	68.1	75.8	73.1	78.1
SHD	group masks	85.0	84.4	84.2	76.9	73.5	83.5	78.6	83.1	80.7
N-TIDIGITS18	clean	82.9	71.3	40.3	18.6	17.4	57.8	73.5	60.9	72.4
N-TIDIGITS18	group masks	79.9	78.5	70.1	57.8	54.3	76.4	71.6	74.1	72.9

D.6 Native-unit baseline certificates

Table 22 compares the native certificates of entry masking and Gaussian smoothing with the size of one SST group.

Table 22: Why entry-level and ℓ_2 certificates rarely cover one SST group. For correctly classified, certified test inputs we report the median native radius of entry masking (entries) and Gaussian smoothing (ℓ_2), next to the median number of occupied entries in the fullest active group (and its ℓ_2 norm when deleted). Seed 0.

Dataset	Entry radius (entries)	Entries in fullest group	Gaussian ℓ_2 radius	ℓ_2 norm of fullest group
N-MNIST	45	83	0.66	9.11
DVS-Gesture	33	242	0.66	15.56
N-Caltech101	41	198	0.65	14.07
SHD	8	40	0.66	6.32
N-TIDIGITS18	2	8	0.64	2.83

D.7 Empirical check of the certificates

Table 23 attacks certified inputs with exactly k_{wc} informed group deletions; no prediction changed.

E Runtime

Every method performs the same $n_0 + n = 2176$ classifier evaluations per input; confidence bounds and radius searches are negligible. Table 24 reports measured wall-clock time on the Apple M5 Pro

Table 23: Empirical check of the certificates. For correctly classified inputs certified with $k_{\mathrm{wc}} \geq 1$ (seed 0, selected p_{keep}), we delete exactly k_{wc} active groups ranked by saliency or by event density and re-estimate the smoothed prediction with fresh Monte Carlo samples. A change of prediction would indicate a certificate violation.

Dataset	Inputs	Median k_{wc}	Saliency flips	Density flips
N-MNIST	200	17	0	0
DVS-Gesture	166	8	0	0
N-Caltech101	200	26.5	0	0
SHD	200	3.5	0	0
N-TIDIGITS18	200	2	0	0

GPU. Runs shared the GPU with between three and eight concurrent jobs, and the EventCert runs were executed first, under the heaviest sharing, so differences between columns reflect load rather than method cost; the times are indicative only.

Table 24: Mean certification time per 1000 test inputs (seconds; $n_0 + n = 2176$ classifier evaluations per input) at the selected smoothing parameter.

Dataset	EventCert	Group ablation	Entry	Gaussian
N-MNIST	400	172	163	164
DVS-Gesture	699	287	272	293
N-Caltech101	606	304	293	300
SHD	153	52	47	45
N-TIDIGITS18	62	11	11	11

F Proof of the EventCert Certificate

This appendix proves the certificate condition used in Eq. 4, the effective active-group guarantee, and the operational physical-loss certificate. The proof is written for a multi-class deterministic classifier and uses simultaneous Clopper–Pearson confidence bounds over all competing classes.

F.1 Assumptions and Notation

The proof follows the setting of Sections 3–4. A fixed preprocessing operator Φ maps an event stream to a binary event tensor $X = \Phi(\mathcal{E})$. The tensor is partitioned into a fixed set of disjoint SST groups indexed by $\mathcal{I} = \{1, \ldots, m\}$. For a deletion set $S \subseteq \mathcal{I}$, the operator $\mathcal{D}_S(X)$ zeros out all groups in S and leaves all remaining groups unchanged.

EventCert samples an independent Bernoulli mask $M \in \{0,1\}^m$, where each group is retained with probability $p_{\mathrm{keep}} \in (0,1)$ and removed with probability $1 - p_{\mathrm{keep}}$. The masked input is

$$X^{(M)} = \mathcal{D}_{\{i:M_i=0\}}(X).$$

For a deterministic base classifier f, the smoothed class probability is

$$p_c(X) = \Pr_M[f(X^{(M)}) = c],$$

and the smoothed classifier is

$$g(X) = \arg\max_c p_c(X),$$

with deterministic tie-breaking if needed.

F.2 Monte Carlo Confidence Bounds

EventCert uses two independent Monte Carlo stages. The first stage selects a candidate class $\hat{A}$. The second stage draws n fresh masks and records class counts $C_1, \ldots, C_K$ for K classes. For the

selected class $\hat{A}$, EventCert computes a lower Clopper–Pearson bound

$$p_{\hat{A}}^{L} = \mathrm{CPLower}(C_{\hat{A}}, n, \alpha/2).$$

For every competing class $c \neq \hat{A}$, it computes an upper Clopper–Pearson bound

$$p_{c}^{U} = \mathrm{CPUpper}\left(C_c, n, \frac{\alpha}{2(K-1)}\right).$$

The simultaneous competitor bound is

$$p_{\mathrm{comp}}^{U} = \max_{c \neq \hat{A}} p_{c}^{U}.$$

Because the second-stage samples are independent of the class-selection stage, the Clopper–Pearson bounds are valid for the selected class and all competitors. By the union bound, with probability at least $1-\alpha$ over the Monte Carlo sampling,

$$p_{\hat{A}}(X) \geq p_{\hat{A}}^{L} \quad \text{and} \quad p_c(X) \leq p_c^{U} \;\; \forall c \neq \hat{A}.$$

All guarantees below are conditioned on this confidence event.

F.3 Conditional Distribution Induced by Deletion

Let S be a deletion set with $|S| = \ell$, and let $E_S = \{M_i = 0 \; \forall i \in S\}$ be the event that the mask already deletes every group in S. Mask entries are independent, so $q_\ell := \Pr(E_S) = (1-p_{\mathrm{keep}})^{\ell} > 0$. Once $\mathcal{D}_S$ has removed the groups in S, the mask values on those groups no longer matter, and all other groups are masked exactly as before. Smoothing $\mathcal{D}_S(X)$ is therefore the same as smoothing X conditioned on E_S: $p_c(\mathcal{D}_S(X)) = \Pr_M[f(X^{(M)}) = c \mid E_S]$ for every class c.

F.4 Class-Probability Bounds After Deletion

By the law of total probability, $p_c(X) = q_\ell \, p_c(\mathcal{D}_S(X)) + (1-q_\ell)\beta_c$ with $\beta_c = \Pr(f(X^{(M)}) = c \mid E_S^c) \in [0,1]$, so $p_c(\mathcal{D}_S(X)) = \big(p_c(X) - (1-q_\ell)\beta_c\big)/q_\ell$. For the selected class, $\beta_{\hat{A}} \leq 1$ and $p_{\hat{A}}(X) \geq p_{\hat{A}}^{L}$ give $p_{\hat{A}}(\mathcal{D}_S(X)) \geq \big(p_{\hat{A}}^{L} - (1-q_\ell)\big)/q_\ell$. For every competitor $c \neq \hat{A}$, $\beta_c \geq 0$ gives $p_c(\mathcal{D}_S(X)) \leq p_c(X)/q_\ell \leq p_c^{U}/q_\ell \leq p_{\mathrm{comp}}^{U}/q_\ell$. The selected class therefore stays above every competitor after deleting S whenever $\big(p_{\hat{A}}^{L} - (1-q_\ell)\big)/q_\ell > p_{\mathrm{comp}}^{U}/q_\ell$, which, since $q_\ell > 0$, is equivalent to $p_{\hat{A}}^{L} - p_{\mathrm{comp}}^{U} > 1 - q_\ell$.

F.5 Worst-Case Group-Deletion Certificate

The main paper certifies all deletion sets of size at most k. For any S with $|S| = \ell \leq k$,

$$q_\ell = (1-p_{\mathrm{keep}})^{\ell} \geq (1-p_{\mathrm{keep}})^{k} = q_k.$$

Therefore,

$$1 - q_\ell \leq 1 - q_k.$$

If the certificate condition in Eq. 4 holds,

$$p_{\hat{A}}^{L} - p_{\mathrm{comp}}^{U} > 1 - (1-p_{\mathrm{keep}})^{k} = 1 - q_k,$$

then it also holds for every $\ell \leq k$:

$$p_{\hat{A}}^{L} - p_{\mathrm{comp}}^{U} > 1 - q_\ell.$$

Combining this with the class-probability bounds above gives

$$p_{\hat{A}}(\mathcal{D}_S(X)) > p_c(\mathcal{D}_S(X)) \qquad \forall c \neq \hat{A},$$

for every deletion set $S \subseteq \mathcal{I}$ with $|S| \leq k$. Hence,

$$g(\mathcal{D}_S(X)) = \hat{A} \qquad \forall S \subseteq \mathcal{I}, \; |S| \leq k.$$

This proves the worst-case certificate used by EventCert. The certified worst-case radius is

$$k_{\mathrm{wc}}(X) = \max\left\{k \in \{0, \ldots, m\} : p_{\hat{A}}^{L} - p_{\mathrm{comp}}^{U} > 1 - (1-p_{\mathrm{keep}})^{k}\right\},$$

when the set is non-empty.

If the set is empty, no certificate is obtained and the procedure abstains; in the artifact this is recorded as $k_{\mathrm{wc}} = -1$ and counted as uncertified in all metrics. When the $k = 0$ condition $p_{\hat{A}}^{L} > p_{\mathrm{comp}}^{U}$ holds, this certifies the smoothed prediction at the original input.

F.6 Effective Active-Group Guarantee

Let

$$\mathcal{N}(X) = \{i \in \mathcal{I} : G_i(X) \neq 0\}$$

be the active group set. If $i \notin \mathcal{N}(X)$, then $G_i(X) = 0$, so deleting group i leaves the tensor unchanged. Therefore, for any deletion set S,

$$\mathcal{D}_S(X) = \mathcal{D}_{S\cap\mathcal{N}(X)}(X).$$

Define the effective active-group deletion count as

$$h_X(S) = |S \cap \mathcal{N}(X)|.$$

If

$$h_X(S) \leq k_{\mathrm{wc}}(X),$$

and $k_{\mathrm{wc}}(X)$ is a valid certified radius satisfying Eq. 4, then the worst-case certificate applies to the active subset $S \cap \mathcal{N}(X)$, and therefore

$$g(\mathcal{D}_S(X)) = g(\mathcal{D}_{S\cap\mathcal{N}(X)}(X)) = \hat{A}.$$

This proves the effective active-group guarantee.

Importantly, this is not a larger raw worst-case adversarial radius. A knowledgeable adversary may choose only active groups, so the raw worst-case radius remains k_{wc}. The active-group identity only states that deletions of inactive groups do not consume the certified budget because they do not change the tensor.

F.7 Operational Certificate for Physical Deletion

Let $S \sim \Pi_r$ be a physical deletion process at raw deletion level r, and define

$$H_r = |S \cap \mathcal{N}(X)|.$$

From the effective active-group guarantee, prediction invariance holds whenever

$$H_r \leq k_{\mathrm{wc}}(X),$$

provided that $k_{\mathrm{wc}}(X)$ is a valid certified radius. Therefore, conditioned on the Monte Carlo confidence event,

$$\Pr_{S\sim\Pi_r}\left[g(\mathcal{D}_S(X)) \neq \hat{A}\right] \leq \Pr_{S\sim\Pi_r}\left[H_r > k_{\mathrm{wc}}(X)\right].$$

Thus, if

$$\Pr_{S\sim\Pi_r}\left[H_r > k_{\mathrm{wc}}(X)\right] \leq \delta,$$

then the physical deletion process is certified at failure probability at most δ, conditioned on the Monte Carlo confidence event. Equivalently, with probability at least $1-\alpha$ over certification, the physical deletion failure probability is at most δ.

The operational radius is therefore

$$r_\delta(X) = \max\left\{r : \Pr_{S\sim\Pi_r}\left[H_r > k_{\mathrm{wc}}(X)\right] \leq \delta\right\}.$$

This is the operational certificate used in Eq. 7. It shows how sparsity can increase the tolerated raw physical deletion level without increasing the raw worst-case adversarial radius.

F.8 Uniform Random Deletion as a Hypergeometric Case

For uniform random deletion without replacement, the physical process selects r groups from the total group set $\mathcal{I}$ of size m. Let

$$n_X = |\mathcal{N}(X)|$$

be the number of active groups. Then H_r counts how many active groups are selected when r groups are sampled without replacement from m total groups containing n_X active groups. Hence,

$$H_r \sim \mathrm{Hypergeometric}(m, n_X, r).$$

The probability that more than k active groups are removed is

$$\Pr[H_r > k] = \sum_{j=k+1}^{\min(r,n_X)} \frac{\binom{n_X}{j}\binom{m-n_X}{r-j}}{\binom{m}{r}}.$$

This is the expression used in Eq. 6 and in the analytical adversarial–operational gap shown in Fig. 4. Other physical deletion processes use the same definition of H_r, with the probability evaluated under the corresponding process, e.g. exactly by dynamic programming for the Gilbert channel of Proposition 1.

F.9 Sample Complexity and Attainable Radius

The finite-sample budget limits the largest radius attainable by this independent-mask conditioning certificate. Under the most favorable second-stage counts, the maximum simultaneous margin is

$$M_{\max}(n, K, \alpha) = \left(\frac{\alpha}{2}\right)^{1/n} + \left(\frac{\alpha}{2(K-1)}\right)^{1/n} - 1.$$

Hence,

$$k_{\max} = \max\left\{k \in \mathbb{N} : 1 - (1 - p_{\text{keep}})^k < M_{\max}(n, K, \alpha)\right\}.$$

For large n,

$$1 - M_{\max} \approx \frac{\log(2/\alpha) + \log(2(K-1)/\alpha)}{n},$$

so attaining radius k requires, up to constants,

$$n = \Omega\left(\frac{\log(K/\alpha)}{(1 - p_{\text{keep}})^k}\right).$$

The exponential dependence on k comes from the conditioning probability $q_k = (1 - p_{\text{keep}})^k$ and is a limitation of this certificate construction, not a claim about all deletion certifiers.

F.10 Boundary Misalignment and Model Misspecification

A physical deletion need not align perfectly with the SST partition. For a deleted physical region B, define the conservative SST cover

$$\Gamma(B) = \{i : G_i \cap B \neq \emptyset\}.$$

The whole-group deletion theorem applies to $\mathcal{D}_{\Gamma(B)}(X)$ when

$$|\Gamma(B) \cap \mathcal{N}(X)| \leq k_{\text{wc}}(X).$$

It does not directly certify the partially deleted tensor produced by B: a deterministic classifier need not respond monotonically to deletion within a group. Applying the whole-group theorem therefore requires an explicit operation that zeros every intersected group before classification, or a separate certificate for within-group corruption. Merely counting intersected groups does not establish invariance for an unsanitized, partially deleted input.

The operational guarantee is conditional on the specified process Π_r. If the true deletion process is Q_r, the worst-case radius k_{wc} is unchanged, but r_δ should be recalibrated. More generally, if

$$d_{\text{TV}}(Q_r, \Pi_r) \leq \varepsilon,$$

then for the failure event $\{H_r > k_{\text{wc}}\}$,

$$\Pr_{Q_r}[H_r > k_{\text{wc}}] \leq \Pr_{\Pi_r}[H_r > k_{\text{wc}}] + \varepsilon.$$

Thus, an operational bound of δ under Π_r becomes at most $\delta + \varepsilon$ under this total-variation mismatch.

F.11 Bursty-Channel and Hybrid Certificates

Let E denote the Monte Carlo event on which the confidence bounds hold, $\Pr[E] \geq 1-\alpha$. On E, the worst-case certificate and the active-support identity give $g(\mathcal{D}_S(X)) = \hat{A}$ for every S with $|S \cap \mathcal{N}(X)| \leq k_{\mathrm{wc}}(X)$.

Proposition 1. Let the channel losses S be random and independent of the Monte Carlo samples. On E, $\{g(\mathcal{D}_S(X)) \neq \hat{A}\} \subseteq \{|S \cap \mathcal{N}(X)| > k_{\mathrm{wc}}(X)\} = \{H > k_{\mathrm{wc}}\}$, so $\Pr[g(\mathcal{D}_S(X)) \neq \hat{A} \mid E] \leq \Pr[H > k_{\mathrm{wc}}] \leq \delta$ and $\Pr[E \wedge g(\mathcal{D}_S(X)) = \hat{A}] \geq 1-\alpha-\delta$. For the Gilbert channel, (B_t) is a Markov chain started from its stationary law $(1-\rho, \rho)$ with transition probabilities $\Pr[B_{t+1}{=}1 \mid B_t{=}0] = \rho/(L(1-\rho))$ and $\Pr[B_{t+1}{=}0 \mid B_t{=}1] = 1/L$. The pair $(B_t, \min(\sum_{s \leq t} u_s B_s, k{+}1))$ is itself a (time-inhomogeneous) Markov chain on $2(k{+}2)$ states, so propagating its law over $t = 1, \ldots, m$ gives $\Pr[H > k]$ exactly. For independent loss, $H \sim \mathrm{Bin}(|\mathcal{N}(X)|, \rho)$. For fixed L and $\rho \leq 1 - 1/L$ both transition kernels are stochastically monotone and the bad-state entry probability increases with ρ, so the chains can be coupled with $B_t(\rho) \leq B_t(\rho')$ for $\rho \leq \rho'$; hence H and its tail are nondecreasing in ρ, and bisection over $[0, 1-1/L]$ yields the largest certified rate. For independent loss, monotonicity of the binomial tail is immediate.

Proposition 2. Let the adversary choose A with $|A| = a$ (possibly depending on X and on the classifier, but not on the random losses), and let S be a uniformly random r-subset of $\mathcal{I} \setminus A$. With $b = |A \cap \mathcal{N}(X)|$, the number of deleted active groups is $b + |S \cap (\mathcal{N}(X) \setminus A)|$, and $|S \cap (\mathcal{N}(X) \setminus A)| \sim \mathrm{Hypergeom}(m-a, |\mathcal{N}(X)| - b, r)$. Applying Proposition 1 to this count and taking the maximum over all feasible b bounds the failure probability for every adversarial choice of A. The maximum is attained at $b = \min(a, |\mathcal{N}(X)|)$: if the adversary deletes one more active group ($b \to b+1$), one active group of the random pool becomes inactive; coupling the two random draws on the same r positions, H' decreases by at most one, so $b + H'$ does not decrease pathwise and its tail can only grow.

F.12 Summary of the Guarantee

With probability at least $1-\alpha$ over the Monte Carlo certification procedure, whenever Eq. 4 holds for a radius k, EventCert guarantees that every deletion set S with $|S| \leq k$ preserves the smoothed prediction $\hat{A}$. It also guarantees prediction invariance for any deletion set whose effective active-group deletion count satisfies $h_X(S) \leq k$. For a physical deletion process Π_r, the operational certificate further guarantees failure probability at most δ whenever

$$\Pr_{S \sim \Pi_r}[H_r > k] \leq \delta.$$

Setting $k = k_{\mathrm{wc}}(X)$ gives the two robustness guarantees reported in the main paper—worst-case SST certification and operational physical-loss certification—with $h_X(S)$ serving as the input-dependent bridge between them.

G Datasets, Licences, and Released Assets

All datasets are public event-based benchmarks, cited to their original releases and used only for academic evaluation: N-MNIST and N-Caltech101 (Orchard et al., 2015), DVS-Gesture (Amir et al., 2017), SHD (Cramer et al., 2022), and N-TIDIGITS18 (Anumula et al., 2018). They are downloaded from their publishers and not redistributed; their original licences and terms apply. The released package contains only code, configuration, trained weights of the small classifiers used here, and run logs; it contains no dataset files, personal data, or human-subject annotations. The software builds on Python, PyTorch, Tonic, NumPy, SciPy, and matplotlib under their standard open-source licences.