EXTREMAL GEOMETRY: SIGNAL-INDUCED SUPPORTS FOR TIME SERIES REPRESENTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Local patterns in time series can occur over different durations, yet many representations use fixed windows, scales, or resolutions. We study *extremal geometry*, which represents each local maximum or minimum by a signed magnitude and a two-sided temporal support determined by the surrounding signal. Spikelet is a concrete realization developed for variable-duration time-series analysis. We characterize the prior one-pass Spikelet construction and give a linear-time completion that recovers exact signal-induced supports. The completed supports are noncrossing separately for maxima and minima, inducing an inclusion hierarchy across temporal extents. Magnitude pruning preserves the observed value at every retained apex and creates no new strict extrema. For two-sided interior events, the absolute magnitude also equals the corresponding zero-dimensional persistence lifetime, providing a scalar connection to persistent homology. We further compare extremal geometry with fixed-scale Gaussian smoothing and fixed-resolution Fourier reconstruction: smoothing changes global extremal values, whereas reconstruction can create false extrema even with lower L_2 error. Controlled experiments audit these structural claims, and a real-data audit shows that magnitude-selected events retain widely varying signal-induced temporal supports.

1 INTRODUCTION

Recurring local patterns in time series can vary in duration or be temporally stretched or compressed, as illustrated by DTW and semantic motif studies (Alaee et al., 2021; Imani & Keogh, 2019), yet many representations prescribe temporal span through windows, scales, or resolutions. Fixed-length motif search illustrates the difficulty clearly: a single window can truncate a slow occurrence or include unrelated context around a fast one. More generally, when the duration of a local structure is not known in advance, choosing its temporal span determines what the representation treats as one event. A related issue appears in learned time-series models, where temporal processing is organized through choices such as patch size and available context (Nie et al., 2023; Das et al., 2024; Liu et al., 2025). We ask a more specific representation question: can each local structure's temporal support be inferred from the signal and retained explicitly?

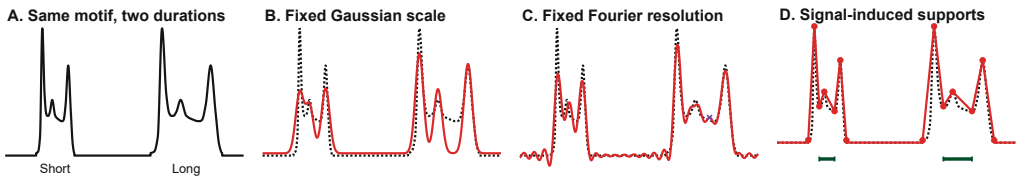

Figure 1: **Variable duration as a representation problem.** A: the same local pattern at two durations. B–C: representative instances using one fixed Gaussian smoothing scale and one fixed Fourier resolution across local structures of different durations. D: Spikelet instead retains observed extrema and assigns each a signal-induced temporal support. The comparison concerns what each representation is designed to preserve, not universal approximation quality.

Spikelet provides a concrete realization of this idea. It retains each observed extremum and assigns it a signed magnitude together with a two-sided temporal support determined by the surrounding

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.

signal (Imamura & Nakamura, 2021). Prior Spikelet work used this decomposition within downstream systems for variable-duration motif discovery and anomaly analysis (Imamura & Nakamura, 2023; 2024; Imamura, 2025); see Appendix A for application background. Here we isolate the representation itself and ask what the prior one-pass construction determines exactly, what support information remains provisional, and how completed supports are organized. We call this support-bearing representation its *extremal geometry*.

We characterize the prior one-pass algorithm, complete provisional supports in linear time, and prove polarity-specific noncrossing, inducing an inclusion hierarchy. For two-sided interior events, Proposition 1 shows that the absolute Spikelet magnitude equals the zero-dimensional persistence lifetime, but only as a scalar: the death level depends only on the higher of the two support-defining minima, whereas Spikelet retains the left and right support endpoints separately. Together, these results characterize a support-bearing representation of local significance and temporal extent.

Our main contributions are:

(1) **Signal-induced support geometry.** For the discrete signal class formalized in Section 3, we characterize precisely what the prior one-pass Spikelet construction determines: every internal extremum receives its exact signed magnitude and exact left support endpoint, while only the right endpoint may remain provisional. We give a linear-time completion and prove polarity-specific noncrossing for this argmin-defined support family, inducing an inclusion hierarchy across temporal extents. (Proposition 2; Theorem 2.)

(2) **Preservation under magnitude pruning.** For the prior magnitude-only Spikelet pruning operation, we prove that every retained apex preserves its observed value and that no new strict extrema are introduced. (Theorem 1.)

(3) **Preservation-oriented comparison with prescribed-scale/resolution representations.** We establish mechanisms by which positive smoothing changes extremal values and fixed-resolution reconstruction can create extrema absent from the input, clarifying the distinction between waveform fidelity and extremal-structure fidelity. (Section 6; Propositions 3–4.)

Controlled experiments audit these theoretical properties, and a real-data audit shows that magnitude-selected events retain supports spanning widely varying temporal extents.

2 Related work and positioning

Fixed-length motif discovery, including the Matrix Profile framework (Yeh et al., 2016; Keogh, 2025), treats subsequence length as an external search choice; prior Spikelet work instead made event support signal-dependent for variable-duration motifs (Imamura & Nakamura, 2021; 2023). The present paper concerns the representation itself, not a new motif-search benchmark. We organize related work by the representation object retained and by how its temporal extent is determined; learned models are discussed separately.

Adaptive-scale and extrema-driven representations. Adaptive scale, extrema-driven decomposition, and event-like coding are established in scale-space/wavelet methods, EMD/ITD, and sparse shift-invariant coding (Witkin, 1983; Lindeberg, 1998; Mallat & Hwang, 1992; Huang et al., 1998; Restrepo et al., 2014; Lewicki & Sejnowski, 1998). Spikelet differs not by adaptivity itself, but by retaining each observed extremum with its own signal-induced two-sided interval in the original signal domain.

Persistence and topological representations. Zero-dimensional persistence represents extrema through birth–death structure (Edelsbrunner et al., 2002); recent 1D work also provides signal-domain windows and nested structure (Biswas et al., 2024), while dynamic persistence structures and merge trees retain richer hierarchical information (Cultrera di Montesano et al., 2024; Pegoraro & Secchi, 2026). The earlier Spikelet scan and linear-time 1D persistence algorithms also share local reductions of alternating extrema but produce different records (Imamura & Nakamura, 2021; Glisse, 2023). Proposition 1 shows that, for two-sided interior extrema, the absolute Spikelet magnitude equals the corresponding persistence lifetime. Spikelet instead studies a different support convention: the oriented argmin-defined interval attached to each original extremum, together with its completion, pruning, and polarity-specific noncrossing laws.

Time-warping and deformation-based representations. Elastic/shape methods align or quotient time warpings; deformation-based methods use invariant dissimilarities or reference-relative deformations, while other approaches impose specified invariances or learn temporal embeddings (Srivastava et al., 2011; Wu et al., 2024; Cantelobre et al., 2022; Germain et al., 2024; Sanborn et al., 2023; Germain et al., 2025; Yue et al., 2022; Fraikin et al., 2024). Spikelet instead retains one support-bearing record per original extremum, including explicit left and right support endpoints in the signal time axis.

Temporal context in learned time-series models. PatchTST uses fixed-length subseries tokens (Nie et al., 2023), while TimesFM and Timer-XL combine patch tokenization with flexible or long context (Das et al., 2024; Liu et al., 2025). Spikelet addresses a different granularity question: whether an event-specific interval can be inferred from the signal and retained explicitly before downstream learning.

Appendix B compares these three representation families by the objects they retain and their relation to extremal geometry.

3 EXTREMAL GEOMETRY

Spikelet provides a concrete realization of the signal-induced-support viewpoint above: each original extremum is retained together with its signed magnitude and a two-sided support determined by the surrounding signal. We call this support-bearing event family its *extremal geometry*. Thus, magnitude describes the vertical contrast of an extremum, while support localizes the temporal structure associated with it; temporal order and support inclusion among maxima and among minima organize these events into a structured representation. Figure 2 illustrates these components for positive events; the negative counterpart is shown in Figure 6 in Appendix A.2.

Let $\mathrm{Ext}(x) = \{p_1, \dots, p_{V(x)}\}$ be the strict internal extrema of a sampled time series $x = (x(1), \dots, x(N))$ in temporal order. For clarity, we state the discrete analysis on the class $\mathcal{X}_{\mathrm{d}}$ of nonconstant time series that are strictly monotone—strictly increasing or decreasing—between consecutive endpoints or strict extrema. Flat runs and plateau extrema are handled by the representative-point convention in Appendix C. Distinct nonadjacent extrema may nevertheless have the same value.

Definition 1 (Spikelet event and excursion support)**.** For a positive apex p, define

$$R(p) = \min\{t > p : x(t) \geq x(p)\}, \qquad L(p) = \max\{t < p : x(t) > x(p)\}.$$

If the right-hand set is empty, set $R(p) = N$; if the left-hand set is empty, set $L(p) = 1$. Let

$$r(p) = \operatorname*{arg\,min}_{[p,R(p)]} x, \qquad \ell(p) = \operatorname*{arg\,min}_{[L(p),p]} x,$$

where, if the minimum is attained at multiple indices, the latest index is chosen. Its two-sided support is $I(p) = [\ell(p), r(p)]$, and its magnitude is

$$m(p) = \min\{x(p) - x(\ell(p)),\, x(p) - x(r(p))\} > 0. \tag{1}$$

The corresponding Spikelet event is $s(p) = (p, \ell(p), r(p), m(p))$. Negative events are defined by applying the same construction to $-x$ and reversing the sign of the resulting magnitude. We refer to positive and negative events as the two polarities.

We first record a useful correspondence for the scalar component of this representation. For two-sided interior events, the absolute Spikelet magnitude coincides exactly with the corresponding zero-dimensional persistence lifetime. Its role here is to identify the scalar information shared with 0D persistence; the signal-domain support itself is treated separately below. Let $d_{\mathrm{PH}}(p)$ denote the 0D persistence death level of the component born at a positive extremum p. Definition 1 determines the Spikelet magnitude from the apex value and the two support-defining minima:

$$m(p) = x(p) - \max\{x(\ell(p)), x(r(p))\}.$$

In zero-dimensional superlevel persistence, the same local maximum p is represented by its birth level $x(p)$ and a lower death level $d_{\mathrm{PH}}(p)$. The difference $x(p) - d_{\mathrm{PH}}(p)$ is its persistence lifetime. Equal-height cases use the κ-refined elder convention specified in Appendix C.

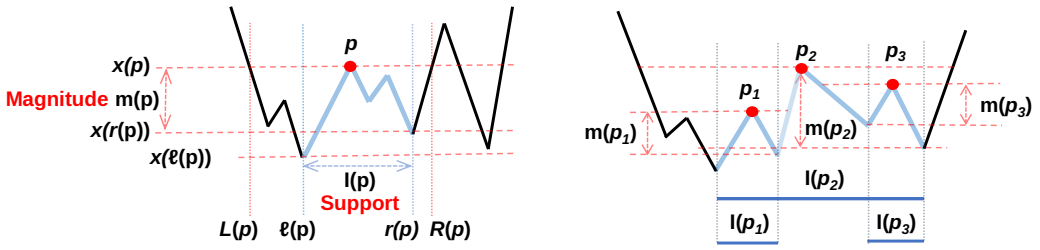

Figure 2: **From a positive extremal event to extremal geometry.** Left: a positive event with apex p. The first-passage limits $L(p)$ and $R(p)$ delimit the two excursion legs; their minima $\ell(p)$ and $r(p)$ define the temporal support $I(p) = [\ell(p), r(p)]$. Right: an illustrative family of positive events. Each apex p_i has its own signal-induced support $I(p_i)$ and magnitude $m(p_i)$; temporal order and support inclusion organize these events into the extremal geometry. The negative-polarity counterpart, obtained by applying the same construction to $-x$, is shown in Figure 6 in Appendix A.2.

Proposition 1 (Persistence correspondence for two-sided extrema)**.** *Let p be the apex of a positive Spikelet event, and suppose that both first-passage sets in Definition 1 are nonempty:*

$$\{t < p : x(t) > x(p)\} \neq \varnothing, \qquad \{t > p : x(t) \geq x(p)\} \neq \varnothing.$$

Let $d_{\mathrm{PH}}(p)$ denote the 0D persistence death level of the component born at p. Then, under the κ-refined elder convention of Appendix C,

$$d_{\mathrm{PH}}(p) = \max\{x(\ell(p)), x(r(p))\},$$

and consequently

$$x(p) - d_{\mathrm{PH}}(p) = \min\{x(p) - x(\ell(p)),\ x(p) - x(r(p))\} = m(p).$$

For a negative event, applying the same construction to $-x$ gives a persistence lifetime equal to $|m(p)| = -m(p)$.

Proof idea. On a one-dimensional path, the two support-defining minima determine the levels at which the component born at p first connects to a κ-elder component on its left and right. The earlier connection occurs at the higher of these two levels, namely $\max\{x(\ell(p)), x(r(p))\}$, which is the persistence death level. Appendix C gives the full proof, including the treatment of equal-height extrema.

Scalar equality versus preserve-first support. Persistence can retain other signal-domain windows and hierarchy (Biswas et al., 2024); J below denotes only the death-level cut. Geometrically, the persistence lifetime corresponds to cutting both sides at the common level $x(p) - m(p) = d_{\mathrm{PH}}(p)$. If one side descends to a deeper support-defining minimum than the other, *this common-level cut stops before reaching that deeper valley, whereas Spikelet retains the two sides separately*.

Formally, for a positive two-sided interior event, let $J(p)$ be the closure of the apex-containing component of $\{t : \widetilde{x}(t) > d_{\mathrm{PH}}(p)\}$, where $\widetilde{x}$ is the piecewise-linear interpolation of the sampled signal. Then $J(p) \subseteq I(p)$, strictly on the deeper side when the support minima differ (Appendix C). Spikelet therefore preserves asymmetric local extent that remains available for later task-dependent processing; a common-level cut can still be imposed afterward if desired.

Proposition 1 applies only when the event is enclosed on both sides before reaching the signal boundary. If one side never reaches the required apex-level first passage, Definition 1 closes the support at the domain boundary instead; such boundary-fallback cases need not have a finite persistence death. The global maximum representing the essential 0D class, i.e., the component with no finite death, likewise falls outside Proposition 1; see Appendix C. Negative events follow by sign reversal.

Definition 1 describes one event at a time. Extremal geometry is the collection of all such events together with their temporal order and support-inclusion relations within each polarity:

Definition 2 (Extremal geometry)**.** Let

$$S(x) = \{\, s(p) = (p, \ell(p), r(p), m(p)) : p \in \mathrm{Ext}(x) \,\}$$

be the family of Spikelet events from Definition 1. Define

$$\mathcal{G}_{\text{ext}}(x) = (S(x), <_t, \preceq_+, \preceq_-),$$

where $<_t$ orders events by their apexes and $\preceq_\pm$ is interval inclusion of their supports within each polarity.

4 Computing and Pruning Extremal Geometry

Section 3 defined the full extremal geometry. We now connect this definition to the prior Spikelet decomposition and then study magnitude pruning as a selection-based coarsening operation. The resulting preservation guarantee provides the reference point for the comparisons in Section 6.

4.1 From the Prior Spikelet Decomposition Algorithm to Exact Supports

First-pass procedure. The prior Spikelet decomposition operates on an ordered anchor sequence consisting of the two endpoints and all strict internal extrema: (i) scan these anchors from left to right; (ii) after each new anchor is pushed, repeatedly test Middle before Left until neither applies; and (iii) after the scan, apply terminal Right reductions to the residual stack. Each reduction records the event(s) determined by the current stack configuration and updates the active stack. Figure 3 illustrates these operations; exact trigger inequalities and implementation-order details are given in Appendix D.1.

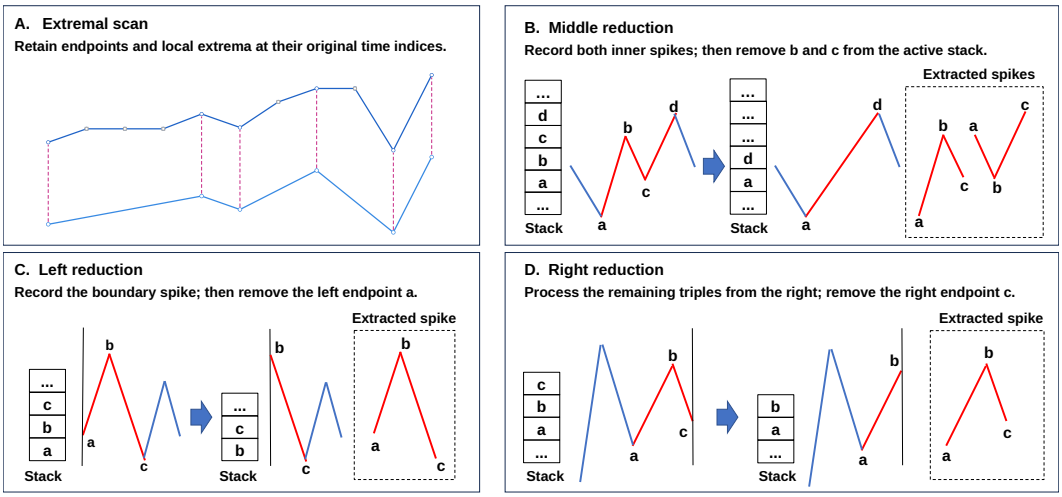

Figure 3: **First-pass Spikelet decomposition of Imamura & Nakamura (2021), expressed in the notation of this paper.** A: extremal scan. B: Middle reduction records the two inner events and removes their active anchors. C–D: Left and terminal Right reductions resolve the boundary portions of the active stack. The exact reduction conditions are given in Appendix D.1.

The two-sided support object itself originates in prior Spikelet work. What was not characterized there is which parts of Definition 1 are already fixed by the implementation-order first pass. Proposition 2 provides this characterization.

Proposition 2 (What one pass fixes)**.** *For every $x \in \mathcal{X}_{\text{d}}$, the prior first-pass, implementation-order reduction assigns every internal extremum exactly the signed magnitude $m(p)$ of Definition 1. If $I^a(p) = [\ell^a(p), r^a(p)]$ is its recorded anchor support, then*

$$\ell^a(p) = \ell(p), \qquad I^a(p) \subseteq I(p).$$

Proof idea. At each recording state, the left neighbor is the exact support minimizer and the right neighbor lies inside the eventual excursion; if the right leg determines the magnitude, that endpoint is already exact. Hence a provisional right endpoint implies that the exact left support endpoint $\ell(p)$ already determines the magnitude. Completion recovers the remaining support geometry in $O(N)$ time by a forward monotone-stack sweep (Appendix D.3).

4.2 PRESERVATION UNDER MAGNITUDE PRUNING

With the complete event family in hand, we next study the magnitude-pruning operation illustrated in Figure 7 and ask what is preserved when events are removed by magnitude. For threshold τ, define

$$S_\tau(x) = \{s(p) \in S(x) : |m(p)| \geq \tau\}, \qquad A_\tau(x) = \{p \in \mathrm{Ext}(x) : s(p) \in S_\tau(x)\},$$

where $A_\tau(x)$ is the set of retained apexes. Reconstruct $\hat{x}_\tau$ by linearly interpolating between consecutive apexes in $A_\tau(x)$, together with the two endpoints. Let $V(z)$ denote the number of strict internal extrema of a signal z.

Theorem 1 (Preservation under magnitude pruning)**.** *For every $x \in \mathcal{X}_\mathrm{d}$ and threshold τ, the reconstruction $\hat{x}_\tau$ satisfies*

$$\hat{x}_\tau(p) = x(p) \quad \textit{for every } p \in A_\tau(x), \qquad \mathrm{Ext}(\hat{x}_\tau) \subseteq A_\tau(x) \subseteq \mathrm{Ext}(x).$$

Consequently,

$$V(\hat{x}_\tau) \leq |A_\tau(x)| \leq V(x).$$

Proof idea. By construction, $\hat{x}_\tau$ passes through every retained apex, so its observed value is unchanged. Between consecutive retained apexes or domain endpoints, $\hat{x}_\tau$ is affine and therefore has no strict internal extremum; hence $\mathrm{Ext}(\hat{x}_\tau) \subseteq A_\tau(x)$. The same argument applies to any interpolation that is monotone between consecutive retained apexes; unconstrained higher-order interpolation need not satisfy this property. A complementary continuous time-reparameterization result is given in Appendix C.2; it provides an idealized continuous view related to the DTW-based motif matching used in prior Spikelet work (Imamura & Nakamura, 2023; 2024).

5 SUPPORT GEOMETRY

Section 4 established how individual events are computed and what is preserved under magnitude pruning. We now ask how the supports of different events are organized collectively. The key result is that same-polarity supports cannot cross, which induces a hierarchy across temporal extents.

Two intervals $[a, b]$ and $[c, d]$ cross if $a < c < b < d$ or $c < a < d < b$; endpoint contact is allowed. A family with no crossing pair is called noncrossing.

Lemma 1 (Excursion-range noncrossing)**.** *For positive apexes p, q, the excursion ranges $[L(p), R(p)]$ and $[L(q), R(q)]$ do not cross.*

Lemma 2 (Endpoint exclusion)**.** *For positive apexes $p \neq q$ of $x \in \mathcal{X}_\mathrm{d}$, if $[L(p), R(p)] \subseteq [L(q), R(q)]$, then, for $p < q$,*

$$R(p) \leq q, \qquad r(q) \geq r(p), \qquad \ell(q) \notin (\ell(p), r(p)),$$

and, for $p > q$,

$$L(p) \geq q, \qquad \ell(q) \leq \ell(p), \qquad r(q) \notin (\ell(p), r(p)).$$

The next result is specific to the argmin-defined signal-domain supports of Definition 1; related nesting phenomena in persistence are discussed in Section 2.

Theorem 2 (Polarity-specific noncrossing)**.** *For every $x \in \mathcal{X}_\mathrm{d}$, the positive excursion supports $\{I(p)\}$ form a noncrossing family, and so do the negative supports. The property is preserved by magnitude pruning; no claim is made for the union of both polarities.*

Proof idea. On a one-dimensional path, the threshold components surrounding same-polarity apexes are nested or disjoint, so their first-passage excursion ranges cannot cross. When one range contains another, the left/right first-passage conditions of Definition 1 and the support-minimum convention exclude interleaving support endpoints. The full proof for the completed supports is in Appendix C; Appendix E separately proves noncrossing for the one-pass anchor supports.

Corollary 1 (Hierarchical organization of same-polarity events)**.** *Within either polarity, the distinct supports $I(p)$, ordered by strict inclusion, form a forest: each support has at most one minimal enclosing parent. Indeed, any two supports that both strictly contain the same support overlap, so noncrossingness forces one to contain the other. After magnitude pruning, the retained supports still form a forest.*

Multiscale localization alone does not imply noncrossing support organization: Appendix F.6, Proposition 8 gives a compactly supported orthonormal wavelet example with crossing supports across adjacent scales.

Thus Spikelet retains more than the magnitude of each extremum: nested supports encode coarse-to-fine temporal organization, while disjoint supports form separate branches. For time-series data mining, this provides a natural way to organize variable-duration local patterns before task-specific motif or anomaly analysis.

6 WHAT IS PRESERVED UNDER TIME-SERIES COARSENING?

Sections 3–5 established a preservation-oriented coarsening: under magnitude pruning, every retained apex keeps its observed value and no new strict extremum is introduced (Theorem 1). To clarify what this guarantee means, we compare it with two standard coarsening mechanisms based on prescribed scale or resolution: Gaussian smoothing and fixed-resolution Fourier reconstruction. The contrast concerns extremal-structure fidelity rather than waveform fidelity alone: positive smoothing can change extremal values, whereas fixed-resolution reconstruction can create extrema absent from the input. Propositions 3–4 formalize these two effects.

Proposition 3 (Positive smoothing contracts extremal values)**.** *Let x be nonconstant and continuous on the circle $\mathbb{T}$, and let $K \in L^1(\mathbb{T})$ satisfy $K > 0$ a.e. and*

$$\int_{\mathbb{T}} K(u)\, du = 1.$$

Then

$$\min x < \min(K * x) \le \max(K * x) < \max x.$$

A positive smoothing kernel averages each value with its surroundings; Proposition 3 shows that, for every nonconstant continuous signal, the global maximum strictly decreases and the global minimum strictly increases. This is the structural effect illustrated by the fixed Gaussian scale in Figure 1B.

Fixed-resolution reconstruction can fail in a different way. A signal may remain monotone while changing very rapidly over a short interval, for example across a narrow monotone transition (ramp). Such a transition contains no strict local extremum in the original signal, but a finite-resolution reconstruction may oscillate around it. The Gibbs phenomenon—overshoot and ringing near sharp transitions in truncated Fourier approximation—is a familiar example (Gottlieb & Shu, 1997). If these oscillations reverse the reconstructed derivative, local extrema can appear inside an originally monotone interval; we call these *false extrema*. Proposition 4 gives a sufficient local condition for this effect.

Proposition 4 (Local oscillation creates false extrema)**.** *Let $P_k x$ be absolutely continuous on an interval where x is monotone, and write its derivative almost everywhere as*

$$(P_k x)' = g' + r'.$$

Suppose $\beta > 0$ and there are three ordered nonempty open subintervals on which, almost everywhere,

$$g' \ge \beta, \qquad g' \le -\beta, \qquad g' \ge \beta,$$

respectively, while $|r'| < \beta$ almost everywhere on all three. Then $P_k x$ contains a false local maximum and a false local minimum on the interval. The sign-reversed pattern $-, +, -$ similarly yields a false local minimum followed by a false local maximum.

Proof idea. Because $|r'| < \beta$, the three witness intervals retain the sign pattern $+, -, +$; the reconstruction therefore rises, falls, and rises on an interval where x itself is monotone, forcing a false local maximum and minimum. Appendix F verifies this local oscillation and controls interference from other transitions for Fourier truncation and periodic spline projection with transition centers on the spline knot grid.

Together, Propositions 3–4 show that coarsening should be evaluated not only by waveform fidelity but also by extremal-structure fidelity—whether extremal values, ordering, and support relations survive. Section 7 audits these effects empirically.

7 EMPIRICAL VERIFICATION

Prior downstream Spikelet results combine the representation with task-specific processing. We therefore use E1–E4 to isolate the structural claims of Sections 3–6 and a real-data audit to test whether nontrivial support geometry remains after magnitude-only selection. Table 1 summarizes these checks.

Experiment	What is checked	Connection to Sections 3–6
E1	One-pass support characterization, preservation under magnitude pruning, and support geometry	Proposition 2, Theorems 1–2, and Proposition 5: exact left endpoints and anchor containment, retained-apex value exactness, $\mathrm{Ext}(\hat{x}_\tau) \subseteq A_\tau(x)$, and same-polarity noncrossing for completed and one-pass anchor supports.
E2	Extremal-value change under positive smoothing	Section 6, Proposition 3: positive smoothing lowers the global maximum and raises the global minimum.
E3	Structural change under fixed-resolution reconstruction	Section 6, Proposition 4, and Corollaries 2–3: the false-extrema mechanism and the method-specific sufficient conditions for Fourier and periodic-spline reconstruction.
E4	Robustness beyond the analytically certified regime	Section 6: robustness of the E3 effect after relaxing alignment and event-shape assumptions.
Real data	Support extent and hierarchy after magnitude-only selection	Sections 3–5: whether magnitude-selected real-data events retain widely varying support lengths and nontrivial same-polarity organization.

Table 1: Theory–experiment map for the controlled empirical checks and the real-data representation audit.

E1–E3 correspond respectively to the Spikelet, smoothing, and fixed-resolution reconstruction cases introduced in Figure 1. For reconstruction experiments, we report waveform error together with extremal-structure diagnostics; full protocols and additional results are in Appendices G–H.

Construction, preservation, and support geometry (E1). E1 audits the support characterization in Proposition 2 and Theorems 1–2 on synthetic signals composed of separated signed Gaussian events with randomized locations and signed amplitudes. At matched nominal budget k, Spikelet applies magnitude pruning, while Fourier retains the $2k+1$ lowest-frequency coefficients. For each Spikelet budget, we count violations of the set inclusion in Theorem 1 as

$$N_{\mathrm{viol}} = \left|\mathrm{Ext}(\hat{x}_\tau) \setminus A_\tau(x)\right|.$$

For the representative run in Figure 4, $N_{\mathrm{viol}} = 0$ across all tested budgets. On 50 generated signals satisfying $\mathcal{X}_{\mathrm{d}}$, the support-characterization checks for Proposition 2, the preservation and completed-support noncrossing checks for Theorems 1–2, and the anchor-support noncrossing check for Proposition 5 all pass; Appendices G–H give the full protocol and results. Figure 4 separately compares waveform approximation by RMSE with the raw strict-extremum change $\Delta_{\mathrm{strict}} = V_{\mathrm{strict}}(y) - V_{\mathrm{strict}}(x)$, where V_{strict} counts raw strict internal extrema without amplitude thresholding. Fourier achieves lower RMSE but introduces additional strict extrema at moderate-to-high budgets, whereas Spikelet maintains $\Delta_{\mathrm{strict}} \leq 0$ throughout. Thus, better waveform approximation does not necessarily imply better preservation of extremal structure.

In the representative run, one-pass supports have exact left endpoints and are contained in their completed supports; both support families show zero same-polarity crossings, while completion changes some right endpoints.

Positive smoothing (E2). Gaussian smoothing decreases the global maximum and increases the global minimum at every tested nonzero scale, as predicted by Proposition 3 (Table 4).

False extrema under fixed-resolution reconstruction (E3). Proposition 4 gives the false-extrema mechanism; Corollaries 2–3 give method-specific sufficient conditions for strict false extrema in Fourier and periodic-spline reconstructions of originally monotone transitions. The high-resolution audit gives finite-resolution cases satisfying these conditions and exhibiting the predicted effect (Appendix H).

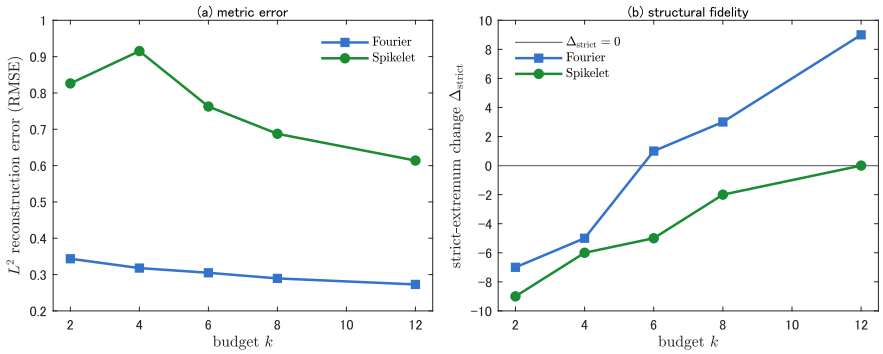

Figure 4: **E1 matched nominal-budget tradeoff.** Fourier achieves lower RMSE (a) but introduces additional strict extrema at moderate-to-high budgets (b), whereas Spikelet maintains $\Delta_{\text{strict}} \leq 0$ at every tested budget.

Robustness beyond the analytically certified regime (E4). The reconstruction effect observed in E3 persists when its analytic assumptions are relaxed: spline projection changes apex values and introduces excess extrema in all 12 tested cases, whereas Spikelet preserves both quantities (Table 7).

Real-data support geometry. We audit three recordings from prior DTW- and semantic-motif studies (Alaee et al., 2021; Imani & Keogh, 2019). With the magnitude-only $K = 3$ heuristic of Imamura (2025), only 1.4%–7.4% of events remain, yet they span broad support ranges; maximum same-polarity hierarchy depth changes from 20–40 to 7–12. Exact support completion changes the one-pass right endpoint for 20.8%–25.6% of selected events. At a representation level, these results help explain how prior Spikelet motif discovery accommodated variable-duration patterns without prescribing one common subsequence length in advance (Imamura & Nakamura, 2023; 2024): after magnitude-based selection, the retained events still have signal-induced supports spanning widely varying temporal extents rather than sharing a single prescribed window length. Even after restricting to two-sided interior events, events in the same persistence-lifetime bin have widely different support lengths under both support conventions. Appendices G–H give the full protocol and results.

8 Discussion

Extremal geometry gives a preservation-oriented view of time-series representation. For two-sided interior events, Spikelet magnitude equals the 0D persistence lifetime, but its completed support retains asymmetric two-sided extent beyond the common death-level cut. Same-polarity supports form a noncrossing hierarchy, and magnitude pruning preserves retained-apex values without introducing new strict extrema; in contrast, smoothing can alter extremal values and fixed-resolution reconstruction can create false extrema.

Broader implication. More broadly, this work connects time-series data mining with topological/geometric representation and classical approximation theory. Although developed for different purposes, these areas can be organized around a common representation question: which properties of the observed signal are retained, discarded, or left available for later processing?

Scope, limitations, and future work. Our results characterize representation-level structure rather than downstream task accuracy, and the real-data audit illustrates retained geometry rather than universal downstream superiority. Prior Spikelet systems selected task-dependent post-processing empirically; determining which operations are important for which combinations of signal class and task remains open (Appendix A). A further direction is to select or learn such processing jointly with deep time-series models, using preserve-first support as an explicit interface between signal structure and learned representations.

Reproducibility Statement

Complete proofs of the theoretical results and detailed descriptions of the Spikelet construction and support-completion algorithms are provided in the appendices. The experimental protocol, including signal generators, real-data selection, parameter settings, evaluation measures, and numerical verification criteria, is specified in Appendix G, with additional results reported in Appendix H. Anonymized source code for reproducing the controlled experiments and rerunning the real-data audits is provided in the supplementary material. The exact processed real-data recordings are not redistributed because redistribution rights for those copies have not been independently verified; instead, the supplementary material includes derived reference outputs for the reported real-data analyses. When authorized local copies of the processed recordings are available, the released code and documented configurations can be used to rerun the real-data audits. All random experiments use fixed seeds; for the E1 theorem-domain robustness audit, candidate seeds are additionally filtered by the deterministic input-only eligibility rule specified in Appendix G before the audited set is frozen.

AI use statement

In this work, we used generative AI tools, primarily OpenAI ChatGPT, with Anthropic Claude additionally used for independent cross-checking. We used these tools to assist in developing and critically examining the conceptual framework; formulate, stress–test, and refine mathematical claims; support proof development and counterexample searches; assist with implementing synthetic data generators based on author-specified experimental designs; provide feedback on experimental methodology; assist with the implementation and debugging of experimental code; and support the analysis and interpretation of reported experimental results. We also used these tools for translation and English–language editing, drafting and restructuring parts of the manuscript, literature search and summarization, and reference formatting. We did not use generative AI tools to collect or clean empirical datasets or to conduct qualitative or thematic data analysis.

The authors reviewed and independently verified all AI–assisted technical work. Mathematical claims and proofs were checked against the stated assumptions and, where relevant, subjected to explicit counterexample searches. AI–assisted code was executed and tested by the authors, and all reported numerical results were obtained from these author-verified programs. Literature claims and reference metadata were manually checked against the cited sources. The authors made all final scientific and editorial decisions and take responsibility for the complete content of this work, including text, claims, code, and artifacts produced with the aid of generative AI.

References

Sara Alaee, Ryan Mercer, Kaveh Kamgar, and Eamonn J. Keogh. Time series motifs discovery under DTW allows more robust discovery of conserved structure. *Data Mining and Knowledge Discovery*, 35(3):863–910, 2021. doi: 10.1007/s10618-021-00740-0. URL https://doi.org/10.1007/s10618-021-00740-0.

Guy Battle. A block spin construction of ondelettes. part I: Lemarié functions. *Communications in Mathematical Physics*, 110(4):601–615, 1987. doi: 10.1007/BF01205550.

Ranita Biswas, Sebastiano Cultrera di Montesano, Herbert Edelsbrunner, and Morteza Saghafian. Geometric characterization of the persistence of 1D maps. *Journal of Applied and Computational Topology*, 8:1101–1119, 2024. doi: 10.1007/s41468-023-00126-9. URL https://doi.org/10.1007/s41468-023-00126-9.

Théophile Cantelobre, Carlo Ciliberto, Benjamin Guedj, and Alessandro Rudi. Measuring dissimilarity with diffeomorphism invariance. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 2572–2596. PMLR, 2022. URL https://proceedings.mlr.press/v162/cantelobre22a.html.

Sebastiano Cultrera di Montesano, Herbert Edelsbrunner, Monika Henzinger, and Lara Ost. Dynamically maintaining the persistent homology of time series. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 243–295. SIAM, 2024. doi:

10.1137/1.9781611977912.11. URL https://doi.org/10.1137/1.9781611977912.11.

Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 10148–10167. PMLR, 2024. URL https://proceedings.mlr.press/v235/das24c.html.

Hoang Anh Dau et al. The UCR time series classification archive, 2018. URL https://www.cs.ucr.edu/~eamonn/time_series_data_2018.

Stephen Demko, William F. Moss, and Philip W. Smith. Decay rates for inverses of band matrices. *Mathematics of Computation*, 43(168):491–499, 1984. doi: 10.1090/S0025-5718-1984-0758197-9.

Herbert Edelsbrunner and Ernst P. Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Transactions on Graphics*, 9(1):66–104, 1990. doi: 10.1145/77635.77639. URL https://doi.org/10.1145/77635.77639.

Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002.

Fuming Fang and Takahiro Shinozaki. Electrooculography-based continuous eye-writing recognition system for efficient assistive communication systems. *PLOS ONE*, 13(2):e0192684, 2018. doi: 10.1371/journal.pone.0192684. URL https://doi.org/10.1371/journal.pone.0192684.

Archibald Felix Fraikin, Adrien Bennetot, and Stéphanie Allassonnière. T-Rep: Representation learning for time series using time-embeddings. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=3y2TfP966N.

Thibaut Germain, Samuel Gruffaz, Charles Truong, Laurent Oudre, and Alain Durmus. Shape analysis for time series. In *Advances in Neural Information Processing Systems 37*, 2024. doi: 10.52202/079017-3029. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/ad86418f7bdfa685cd089e028efd75cd-Abstract-Conference.html.

Thibaut Germain, Chrysoula Kosma, and Laurent Oudre. Time series representations with hard-coded invariances. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 19172–19195. PMLR, 2025. URL https://proceedings.mlr.press/v267/germain25a.html.

Marc Glisse. Fast persistent homology computation for functions on $\mathbb{R}$. *arXiv preprint arXiv:2301.04745*, 2023. doi: 10.48550/arXiv.2301.04745. URL https://arxiv.org/abs/2301.04745.

David Gottlieb and Chi-Wang Shu. On the Gibbs phenomenon and its resolution. *SIAM Review*, 39 (4):644–668, 1997. doi: 10.1137/S0036144596301390.

Norden E. Huang, Zheng Shen, Steven R. Long, Manli C. Wu, Hsing H. Shih, Quanan Zheng, Nai-Chyuan Yen, Chi Chao Tung, and Henry H. Liu. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1971):903–995, 1998. doi: 10.1098/rspa.1998.0193.

Makoto Imamura. Generalized discords for time series anomaly detection with flexible subsequence lengths. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 1013–1024. Association for Computing Machinery, 2025. doi: 10.1145/3711896.3736977. URL https://doi.org/10.1145/3711896.3736977.

Makoto Imamura and Takaaki Nakamura. Spikelet: An adaptive symbolic approximation for finding higher-level structure in time series. In *Proceedings of the 2021 IEEE International Conference on Data Mining (ICDM)*, pp. 1120–1125, 2021. doi: 10.1109/ICDM51629.2021.00132. URL https://doi.org/10.1109/ICDM51629.2021.00132.

Makoto Imamura and Takaaki Nakamura. Parameter-free Spikelet: Discovering different length and warped time series motifs using an adaptive time series representation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 857–866. Association for Computing Machinery, 2023. doi: 10.1145/3580305.3599310. URL https://doi.org/10.1145/3580305.3599310.

Makoto Imamura and Takaaki Nakamura. Efficient discovery of time series motifs under both length differences and warping. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 1188–1198. Association for Computing Machinery, 2024. doi: 10.1145/3637528.3671726. URL https://doi.org/10.1145/3637528.3671726.

Shima Imani and Eamonn J. Keogh. Matrix Profile XIX: Time series semantic motifs: A new primitive for finding higher-level structure in time series. In *2019 IEEE International Conference on Data Mining (ICDM)*, pp. 329–338. IEEE, 2019. doi: 10.1109/ICDM.2019.00043. URL https://doi.org/10.1109/ICDM.2019.00043.

Eamonn Keogh. Finding repeated structures in time series: Algorithms and applications: A unifying view of time series motifs/shapelets/chains and similar primitives. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6054–6062. Association for Computing Machinery, 2025. doi: 10.1145/3711896.3736551. URL https://doi.org/10.1145/3711896.3736551.

Monique A. Ladds, Adam P. Thompson, David J. Slip, David P. Hocking, and Robert G. Harcourt. Seeing it all: Evaluating supervised machine learning methods for the classification of diverse otariid behaviours. *PLOS ONE*, 11(12):e0166898, 2016. doi: 10.1371/journal.pone.0166898. URL https://doi.org/10.1371/journal.pone.0166898.

Michael S. Lewicki and Terrence J. Sejnowski. Coding time-varying signals using sparse, shift-invariant representations. In *Advances in Neural Information Processing Systems 11*, 1998. URL https://proceedings.neurips.cc/paper_files/paper/1998/hash/83f2550373f2f19492aa30fbd5b57512-Abstract.html.

Tony Lindeberg. Feature detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):79–116, 1998.

Yong Liu, Guo Qin, Xiangdong Huang, Jianmin Wang, and Mingsheng Long. Timer-XL: Long-context transformers for unified time series forecasting. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=KMCJXjlDDr.

Stéphane Mallat and Wen Liang Hwang. Singularity detection and processing with wavelets. *IEEE Transactions on Information Theory*, 38(2):617–643, 1992.

Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=Jbdc0vTOcol.

Matteo Pegoraro and Piercesare Secchi. Functional data representation with merge trees. *Journal of Computational and Graphical Statistics*, 2026. doi: 10.1080/10618600.2026.2654769. URL https://doi.org/10.1080/10618600.2026.2654769. Published online 22 June 2026.

Juan M. Restrepo, Shankar Venkataramani, Darin Comeau, and Hermann Flaschka. Defining a trend for time series using the intrinsic time-scale decomposition. *New Journal of Physics*, 16:085004, 2014. doi: 10.1088/1367-2630/16/8/085004. URL https://doi.org/10.1088/1367-2630/16/8/085004.

F. B. Richards. A Gibbs phenomenon for spline functions. *Journal of Approximation Theory*, 66(3): 334–351, 1991. doi: 10.1016/0021-9045(91)90034-8. URL https://doi.org/10.1016/0021-9045(91)90034-8.

Sophia Sanborn, Christian A. Shewmake, Bruno Olshausen, and Christopher J. Hillar. Bispectral neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=xnsg4pfKb7.

Anuj Srivastava, Wei Wu, Sebastian Kurtek, Eric Klassen, and J. S. Marron. Registration of functional data using Fisher–Rao metric. *arXiv:1103.3817*, 2011.

Andrew P. Witkin. Scale-space filtering. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pp. 1019–1022, 1983.

Yuexuan Wu, Chao Huang, and Anuj Srivastava. Shape-based functional data analysis. *TEST*, 33:1–47, 2024. doi: 10.1007/s11749-023-00876-9. URL https://doi.org/10.1007/s11749-023-00876-9.

Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, Hoang Anh Dau, Diego Furtado Silva, Abdullah Mueen, and Eamonn Keogh. Matrix Profile I: All pairs similarity joins for time series: A unifying view that includes motifs, discords and shapelets. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pp. 1317–1322. IEEE, 2016. doi: 10.1109/ICDM.2016.0179. URL https://doi.org/10.1109/ICDM.2016.0179.

Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. TS2Vec: Towards universal representation of time series. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8):8980–8987, 2022. doi: 10.1609/aaai.v36i8.20881.

A Real-data motivation and prior Spikelet representation

This appendix provides application background for readers unfamiliar with motif discovery or earlier Spikelet work. Section A.1 shows why variable duration makes adaptive local support useful; Section A.2 visualizes the prior two-sided Spikelet definition in the notation used in this paper; and Section A.3 illustrates the magnitude-based coarsening operation whose preservation law is established in Theorem 1.

A.1 Variable-duration motifs in real data

A time-series motif is an approximately repeated subsequence embedded in a longer time series. In the standard fixed-length formulation, subsequences are extracted using a sliding window of user-specified length and compared with their nearest neighbors; the Matrix Profile provides a canonical modern framework for this computation (Yeh et al., 2016). For a recent tutorial overview of motif discovery and related repeated-structure primitives, see Keogh (2025).[1] This formulation is natural when the characteristic duration is known in advance.

Two lines of prior motif research motivate the setting considered here. DTW-motif discovery showed that local temporal deformation can make conserved structure difficult to recover with Euclidean matching (Alaee et al., 2021), while semantic motifs explicitly represent repeated structures whose conserved components can be separated by a variable-duration don't-care region (Imani & Keogh, 2019). Thus, variable temporal extent is not introduced here as a purely mathematical construction: in motif discovery, the boundaries and internal timing of a recurring event can themselves be part of the problem. To provide concrete motivation without treating these examples as a comprehensive benchmark, Figure 5 collects representative cases from prior motif studies and prior Spikelet work (Alaee et al., 2021; Imani & Keogh, 2019; Imamura & Nakamura, 2023).

The four domains make the same conceptual point. A recurring behavior need not occupy a single globally prescribed duration: semantically corresponding local patterns can span different time intervals in different occurrences. This motivates assigning signal-induced local temporal support

[1] A visual version of the KDD 2025 tutorial, with many real-world motif examples, is available as a tutorial PDF.

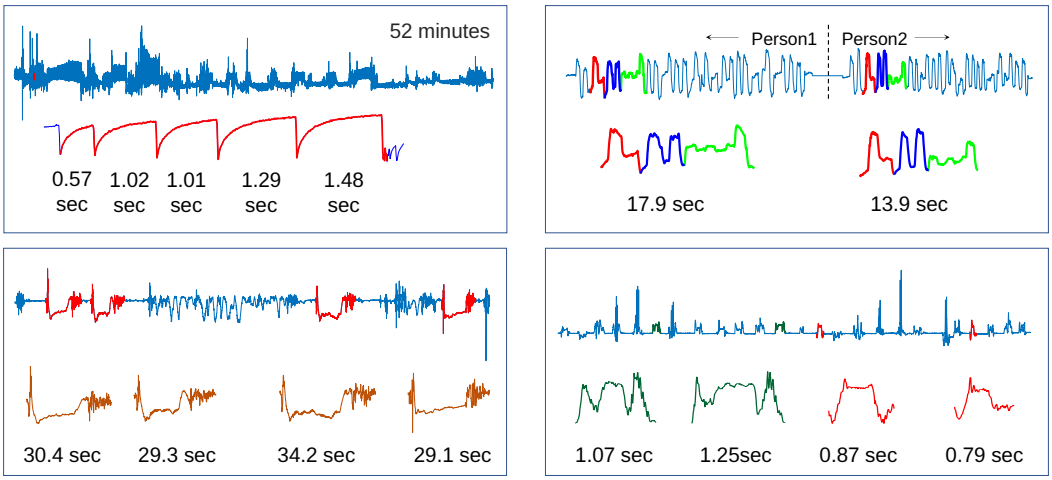

Figure 5: **Real-world variable-duration motifs.** Top left: electrical-penetration-graph telemetry from an Asian citrus psyllid; a sequence of related feeding patterns illustrates recurring structure over 0.57–1.48 s. Top right: eye-movement telemetry from two individuals; two corresponding recurrent patterns extracted from the recordings last 17.9 s and 13.9 s. Bottom left: accelerometer telemetry from an Australian fur seal; four high-ranking occurrences associated with prey-capture acceleration bursts last 30.4, 29.3, 34.2, and 29.1 s. Bottom right: cricket-umpire gesture telemetry; the same communicated gestures may be performed at different speeds, and the illustrated top motif pairs range from 0.79 to 1.25 s. These examples follow the real-data cases reported in Imamura & Nakamura (2023); they motivate variable-duration motif analysis and are not used as evidence for the structural theorems of the present paper.

before nearest-neighbor matching. The main-text Figure 1 abstracts this practical difficulty into the representation question studied here: can event-specific temporal support be inferred from the observed signal and retained explicitly?

Initial decomposition versus application processing. The original ICDM pipeline makes a distinction that is important for the scope of this paper (Imamura & Nakamura, 2021). Initial spike generation produces an adaptive event decomposition; application representations then apply, as needed, global spike reduction, constant-segment extraction, local context-dependent spike reduction, and symbolic mapping before motif or domain analysis. Its real-data examples therefore evaluate processed Spikelet approximations rather than the unreduced initial event family in isolation. We therefore analyze the initial support as a preserve-first representation rather than a universally optimal final window. The present paper studies the common primitive that precedes those task-dependent stages—the initial support-bearing event family and magnitude-only pruning—so controlled audits are more direct tests of the claimed structural laws than a downstream benchmark on the raw event set.

A.2 Geometric reading of the prior Spikelet definition

Figure 6 gives the negative-polarity counterpart of the positive construction shown in Figure 2. A negative event is obtained by applying Definition 1 to $-x$ and reversing the sign of the resulting magnitude. Thus the same first-passage and two-sided-support construction applies to minima, with the orientation reversed. Together, Figures 2 and 6 visualize the two polarities of the extremal geometry. The figures are included as reading aids for the prior Spikelet definition; the new results of this paper characterize the implementation-order construction and the structural laws of the resulting support family, while the persistence correspondence provides a scalar connection.

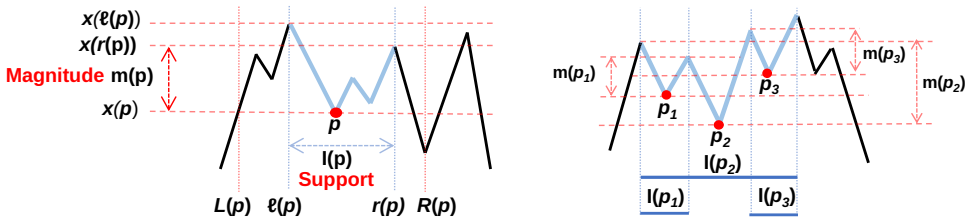

Figure 6: **Negative-polarity counterpart of the extremal geometry in Figure 2.** Left: a negative event centered on a local minimum p, obtained by applying the construction of Definition 1 to $-x$. Its signed magnitude is negative, while $I(p) = [\ell(p), r(p)]$ gives its signal-induced two-sided temporal support. Right: an illustrative family of negative events, showing the same event-specific support organization as in the positive case. Figure 2 shows the corresponding positive-polarity construction.

A.3 Magnitude pruning in prior Spikelet representations

Magnitude-based coarsening predates the present theory, and it is one of a family of operations that a Spikelet pipeline applies after the initial decomposition. Imamura & Nakamura (2021) define three: *global spike reduction*, which discards events below a magnitude threshold; *constant segment extraction*, which identifies near-flat stretches and removes the events they absorb; and *local spike reduction*, which discards an event whose support is contained in that of another event of at least twice its magnitude. Later motif work uses a magnitude threshold (MaT) to select salient events (Imamura & Nakamura, 2023), and the anomaly work selects representative events (Imamura, 2025).

Of these, only global spike reduction and its descendants act by magnitude alone, and Theorem 1 is stated for exactly that class: a threshold on $|m_i|$ and the piecewise-linear reconstruction through the surviving apexes. Constant segment extraction and local spike reduction consult the signal or the support inclusion relation as well, and lie outside the theorem as stated. We regard this as the correct division of labour rather than an omission: the operations are where a pipeline is fitted to a task, and their utility is an empirical matter settled in the works above, whereas what a given operation preserves is a structural question about the object it acts on. Figure 7 shows the same operational idea on a Penguin motion signal used in subsequent Spikelet work (Imamura & Nakamura, 2024): increasing the threshold removes lower-magnitude positive and negative events and exposes progressively coarser extremal structure.

The pruning operation itself is therefore not a new contribution of this paper. What is new here is the structural statement in Theorem 1: for the canonical implementation-order construction and the specified piecewise-linear reconstruction, the value at every retained apex is preserved exactly and no new strict internal extremum is introduced. Thus Figure 5 motivates signal-induced temporal support, whereas Figure 7 illustrates the prior operation whose preservation properties are formalized here. Neither figure is used as empirical evidence for the present theorems.

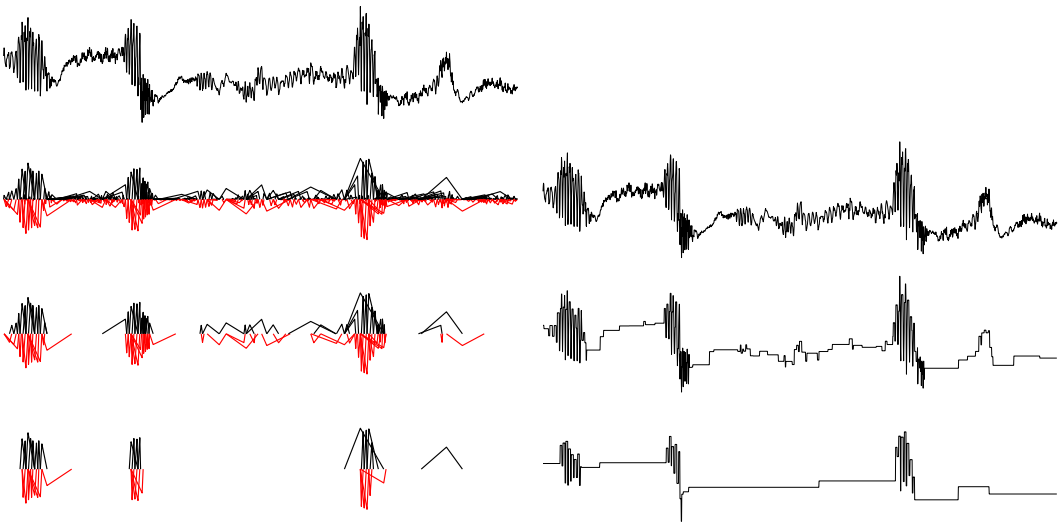

Figure 7: **Spikelet magnitude pruning on a real motion signal.** A snippet of the Penguin motion signal is shown together with pruning at $\tau = 0$, 0.6, and 1.75. Left: retained events at each threshold, with positive events in black and negative events in red. Right: the corresponding time-series representations using both polarities. Increasing τ removes lower-magnitude fluctuations and exposes progressively coarser extremal structure while retaining the observed values at the selected original apexes.

B Representation-object comparison with related work

Table 2 compares the three representation families in Section 2 by the objects they retain and how temporal extent enters the representation. Patch-based and foundation models are discussed separately because the table focuses on event-support geometry, not because such models lack representation objects. The table is intended for positioning, not performance ranking.

	Adaptive-scale and extrema-driven representations	**Persistence and topological representations**	**Time-warping and deformation-based representations**
Representative approaches	**(a) Scale-space / wavelet:** Witkin; Lindeberg; Mallat–Hwang. **(b) Decomposition / atomic coding:** EMD; ITD; sparse shift-invariant coding.	0D persistence; merge trees; 1D persistence geometry; dynamic persistence structures.	Fisher–Rao / elastic representations; deformation-based methods; invariant or temporal embeddings.
Primary representation object	**(a)** transformed extrema or cross-scale features. **(b)** decomposed modes or fitted atoms.	birth–death pairs, persistence windows, merge relations, or persistence-oriented tree structures.	aligned trajectories, quotient representations, deformation-aware features, invariant embeddings, or learned temporal embeddings.
Treatment of local temporal structure	The original extremum is represented through transformed features, decomposed modes, or fitted atoms.	Extrema are characterized through birth–death, window, or merge structure; persistence lifetime provides one vertical scalar summary.	Temporal variation is handled through alignment, quotienting, deformation models, invariant encoding, or learned temporal features.
Relation to Spikelet / extremal geometry	Spikelet retains each **original extremum** and associates it with its own **signal-induced two-sided temporal support**; the resulting support-bearing representation is formalized here as extremal geometry.	For two-sided interior events, the **absolute Spikelet magnitude** equals the corresponding 0D persistence lifetime. Persistence can also retain windows and hierarchy; Spikelet instead analyzes the **original-extremum-centered argmin support** and its completion, pruning, and polarity-specific noncrossing.	**Extremal geometry** retains one support-bearing record per original extremum, including explicit left and right support endpoints in the signal time axis.

Table 2: **Representation-object comparison for the three families in Section 2.** The three representation columns, from left to right, correspond to these families. The first family separates scale-space/wavelet from decomposition/atomic-coding approaches. The comparison concerns what each representation retains, not a performance ranking.

C COMPRESSED-MAIN PROOF DETAILS

This appendix provides proof details for results stated or sketched in the main text.

Roadmap. Section C.1 proves the scalar correspondence with 0D persistence, Section C.2 gives the continuous time-reparameterization result used in Section 4, and Section C.3 proves noncrossing of the completed supports directly from Definition 1.

C.1 PERSISTENCE CORRESPONDENCE

Scope outside the two-sided interior case. The correspondence in Proposition 1 is specific to extrema for which both first-passage sets in Definition 1 are nonempty. If one side reaches the domain boundary before the required first passage occurs, the boundary fallback in Definition 1 does not in general coincide with a finite persistence death. In particular, the global maximum representing the essential 0D class, which has no finite persistence death, falls outside this correspondence. This is why the proposition is restricted to two-sided interior events.

Flat runs and plateau representatives. For a maximal flat run $[a, b]$ with $x(a) = \cdots = x(b)$, we use a deterministic representative point; here we choose the latest index b. Compressing each flat run in this way preserves its amplitude and the temporal order of the retained extrema. Because nesting,

disjointness, and crossing depend only on the relative order of support endpoints, these relations are preserved under the corresponding order-preserving lift back to the original time axis. The structural results are therefore stated for this induced strict-extremum representation.

This plateau convention is distinct from the equal-height apex ordering and the later-index argmin rule below.

Why an event-level tie convention is needed. When birth levels are distinct, the persistence pairing of a local maximum is unambiguous at the event level. Equal-height extrema require more care. Persistent homology itself remains well defined with repeated filtration values, but amplitude alone does not determine which equal-height birth should be identified with a particular Spikelet event. This distinction matters here because Spikelet assigns an explicit signal-domain support to every extremum.

Definition 1 already resolves such cases through its oriented support convention: the left first passage is strict whereas the right first passage is non-strict. To express the same orientation in the persistence pairing, we use the lexicographic order

$$\kappa(i) = (x(i), i),$$

so that among equal-height extrema the later index is elder. We write $d_{\mathrm{PH}}^{\kappa}(p)$ and $\mathrm{pers}_0^{\kappa}(p)$ for the resulting event-level death and lifetime. In the main text we suppress the superscript κ; when birth levels are distinct, the refinement is inactive and these quantities reduce to ordinary zero-dimensional persistence.

Proof intuition. On each side of p, the best bottleneck to a κ-elder component is attained at the corresponding first-passage boundary. Thus the two candidate death levels are $x(\ell(p))$ and $x(r(p))$, and the earlier merge occurs at their maximum.

Proof of Proposition 1. For any sample index $q \neq p$, define the path bottleneck

$$b(p, q) = \min_{j \in [\min\{p,q\}, \max\{p,q\}]} x(j).$$

Under the κ-refined elder rule the death level is

$$d_{\mathrm{PH}}^{\kappa}(p) = \max_{q:\kappa(q)>\kappa(p)} b(p, q).$$

If $q < p$, then $\kappa(q) > \kappa(p)$ iff $x(q) > x(p)$, so every such q lies at or left of $L(p)$ and the best left bottleneck is attained at $L(p)$, with value $x(\ell(p))$. If $q > p$, then $\kappa(q) > \kappa(p)$ iff $x(q) \geq x(p)$, so the best right bottleneck is attained at $R(p)$, with value $x(r(p))$. Hence

$$d_{\mathrm{PH}}^{\kappa}(p) = \max\{x(\ell(p)), x(r(p))\},$$

and therefore

$$\mathrm{pers}_0^{\kappa}(p) = x(p) - d_{\mathrm{PH}}^{\kappa}(p) = m(p).$$

When birth levels are distinct, the κ refinement is inactive, so this reduces to ordinary zero-dimensional persistence. The negative case follows by applying the same argument to $-x$, giving lifetime $|m(p)|$. □

Remark 1 (Why the orientation matters). Lexicographic tie breaking is a standard way to resolve degeneracies, closely related to symbolic perturbation and Simulation of Simplicity (Edelsbrunner & Mücke, 1990). Here it is not introduced to make persistent homology well defined for repeated values; rather, it makes the event-level pairing consistent with Spikelet's oriented support convention. The direction is therefore load-bearing. For example, consider the anchor-value sequence

$$(4, 1, 3, 2, 3, 1, 4)$$

and the positive apex $p = 3$, using one-based indexing. Definition 1 gives

$$L(p) = 1, \qquad R(p) = 5, \qquad x(\ell(p)) = 1, \qquad x(r(p)) = 2,$$

and hence $m(p) = 1$. Under $\kappa(i) = (x(i), i)$, the equal-height apex at index 5 is elder, so

$$d_{\mathrm{PH}}^{\kappa}(p) = 2, \qquad \mathrm{pers}_0^{\kappa}(p) = 1 = m(p).$$

Under the opposite order $\kappa'(i) = (x(i), -i)$, that apex is no longer elder; the best elder connection is instead at level 1, giving

$$d_{\mathrm{PH}}^{\kappa'}(p) = 1, \qquad \mathrm{pers}_0^{\kappa'}(p) = 2 \neq m(p).$$

Thus the opposite orientation is incompatible with Spikelet's event-level pairing. The lifetime remains measured in the original amplitude coordinate.

Death-level cut versus Spikelet support. Work on the strict-extremum representation obtained after the flat-run representative convention above. Let t_i denote the retained original sample indices, let $\sigma \in \{+1, -1\}$ be the event polarity, and let $\widetilde{y}$ be the piecewise-linear interpolation, on the original-index coordinate, of $y_i = \sigma x(t_i)$. For a two-sided interior event with apex p and Spikelet support $I(p) = [\ell(p), r(p)]$, define

$$d_p = \max\{\widetilde{y}(\ell(p)), \widetilde{y}(r(p))\}.$$

Let C_p be the connected component containing p of the strict superlevel set

$$\{t : \widetilde{y}(t) > d_p\},$$

and put $J(p) = \overline{C_p}$.

Equivalently, the endpoints of $J(p)$ are the nearest left and right crossings of level d_p, with linear interpolation between retained sample coordinates. Since both support endpoints have value at most d_p while the apex has value strictly above d_p, $J(p) \subseteq I(p)$. If one support endpoint has value strictly below d_p, the inclusion is strict on that side; hence unequal support-minimum values imply strict inclusion on the deeper side. Equal support-minimum values do not necessarily imply equality because an equal death-level value can occur closer to the apex.

This reduction is not invertible from the reduced record alone. Given the signal on $I(p)$, $J(p)$ is obtained by the death-level cut above. Conversely, the birth–death record together with $J(p)$ does not in general determine $I(p)$ without re-accessing the signal outside $J(p)$. For example,

$$(4, 1, 1.5, 2, 3, 2, 4), \qquad (4, 2.5, 1, 2, 3, 2, 4)$$

have the same positive apex $p = 5$, birth value 3, death level 2, and death-level interval $J = [4, 6]$, but their Spikelet supports are $[2, 6]$ and $[3, 6]$, respectively. The polarity-normalized formulation covers negative events as well.

Support-location ties. A separate tie issue occurs when an interval contains several equal-valued minima. The later-index argmin convention in Definition 1 makes the support location deterministic; it does not change the persistence lifetime, which depends only on the minimum value. Thus the oriented ordering of equal-height apexes and the location rule for equal minima serve different purposes.

C.2 Completeness of ordered extrema

Let $\mathcal{X}_{\mathrm{c}}$ denote the class of nonconstant continuous functions on a compact interval with finitely many strict local extrema and strict monotonicity between consecutive anchors, where the anchors are the two endpoints and the strict local extrema. For $x \in \mathcal{X}_{\mathrm{c}}$, let $\mathcal{E}(x)$ denote the ordered anchors together with their values and, for internal extrema, their polarities. We write $\mathcal{E}(x) \cong \mathcal{E}(y)$ when corresponding anchors have identical values and corresponding internal extrema have the same polarity.

Background completeness fact. For $x, y \in \mathcal{X}_{\mathrm{c}}$,

$$\mathcal{E}(x) \cong \mathcal{E}(y) \quad \Longleftrightarrow \quad y = x \circ \varphi$$

for some orientation-preserving homeomorphism φ. This is included only as background for the continuous time-reparameterization statement used in Section 4; it is not a novelty claim. Its relation to prior shape and deformation-based representations is summarized in Appendix B.

Proof intuition. Between consecutive anchors the signal is strictly monotone, so amplitude itself can serve as the correspondence parameter. Matching equal amplitudes segment by segment therefore constructs the required orientation-preserving time reparameterization.

Proof. If $y = x \circ \varphi$ with orientation-preserving φ, then corresponding anchors preserve order and value, and corresponding internal extrema preserve polarity; strict monotonicity between anchors prevents additional extrema. Conversely, let corresponding anchors be p_i, q_i with $y(q_i) = x(p_i)$. On each inter-anchor interval define

$$\varphi_i(t) = \big(x|_{[p_i, p_{i+1}]}\big)^{-1}\big(y(t)\big).$$

The restrictions are strictly monotone in the same direction, so each φ_i is increasing and maps $[q_i, q_{i+1}]$ homeomorphically onto $[p_i, p_{i+1}]$. Adjacent definitions agree at anchors and glue to an orientation-preserving homeomorphism φ satisfying $y = x \circ \varphi$. □

This gives a complementary continuous characterization of duration variation. After discrete sampling, however, the required correspondence values need not occur at sampled points, so exact value matching need not survive discretization.

C.3 Excursion-support noncrossing

Proof roadmap. The argument has two steps. First, threshold components show that the excursion ranges $[L(p), R(p)]$ are nested or disjoint. Second, when one range contains another, the first-passage conditions and later-index argmin rule exclude the endpoint interleavings that would make the supports cross.

For positive events, let $\kappa(i) = (x(i), i)$ lexicographically, so that $\kappa(i) \le \kappa(p)$ iff $x(i) < x(p)$ or $x(i) = x(p)$ and $i \le p$. Let $A_p = \{i : \kappa(i) \le \kappa(p)\}$, and let C_p be the maximal consecutive block of A_p containing p. Because κ is a total order, the sets A_p are nested by threshold, and the corresponding blocks C_p are nested or disjoint. Write $C_p = \{a_p, \dots, b_p\}$. By Definition 1,

$$[L(p), R(p)] = [\max\{1, a_p - 1\},\ \min\{N, b_p + 1\}].$$

Thus nested blocks give nested excursion ranges. Disjoint blocks cannot be adjacent: if $A_p \subseteq A_q$, then $C_p \subseteq A_q$, so adjacency of C_p and C_q would put them in the same consecutive block of A_q, contradicting the maximality of C_q. Hence disjoint blocks are separated by at least one index, and expanding each block by its adjacent first-passage indices produces ranges that are disjoint or meet only at an endpoint. Therefore the excursion ranges do not cross, proving Lemma 1.

For Lemma 2, suppose $[L(p), R(p)] \subseteq [L(q), R(q)]$. If $p < q$, then $L(q) \le L(p) < p < q$, so Definition 1 gives $x(p) \le x(q)$. The non-strict right-passage condition therefore yields $R(p) \le q$, and hence $r(q) \ge r(p)$.

Suppose, for contradiction, that $\ell(q) \in (\ell(p), r(p))$. Since $R(p) \le q$, the whole interval $[L(p), R(p)]$ lies in $[L(q), q]$. If $\ell(q) \le p$, minimality on $[L(p), p]$ and $[L(q), q]$ gives

$$x(\ell(q)) = x(\ell(p)),$$

but $\ell(q) > \ell(p)$, contradicting the later-index choice of $\ell(p)$. If $\ell(q) > p$, minimality on $[p, R(p)]$ and $[L(q), q]$ gives

$$x(\ell(q)) = x(r(p)),$$

but $r(p) > \ell(q)$, contradicting the later-index choice of $\ell(q)$. Thus $\ell(q) \notin (\ell(p), r(p))$.

If $p > q$, then $p < R(p) \le R(q)$, so p lies before the right first passage of q, which gives $x(p) < x(q)$. Hence q satisfies the strict left-passage condition for p, so $L(p) \ge q$ and $\ell(q) \le \ell(p)$. The symmetric argument, with left and right interchanged, gives $r(q) \notin (\ell(p), r(p))$.

By Lemma 1, any two positive excursion ranges are either disjoint (possibly touching at an endpoint) or nested. In the disjoint case, their supports cannot cross because each support is contained in its own excursion range. In the nested case, after exchanging p and q if necessary, Lemma 2 rules out the endpoint interleavings that would produce a crossing. Hence the positive excursion supports are noncrossing. The negative case follows by applying the same argument to $-x$, and magnitude pruning preserves the property because it selects a subfamily.

D PROOF DETAILS FOR CANONICAL CONSTRUCTION AND PRESERVATION

Figure 3 in the main text summarizes the prior first-pass Spikelet decomposition. This appendix has two main goals: to prove Proposition 2, which characterizes exactly what the prior one-pass reduction determines, and to prove Theorem 1, which gives the preservation guarantee for magnitude pruning. Section D.1 states the exact implementation-order reduction. Section D.2 proves Proposition 2 in two steps: Lemma 3 identifies the information preserved by the active stack, and Lemma 4 then uses this invariant to show that, when an event is recorded, its left stack neighbor is the exact left support endpoint and its right neighbor lies within the eventual right support. Section D.3 completes the only remaining provisional quantity, the right support endpoint, in linear time. Section D.4 proves Theorem 1 from piecewise-linear interpolation through the retained apexes.

D.1 PRIOR SPIKELET DECOMPOSITION ALGORITHM (FIRST PASS)

The initial Spikelet decomposition is inherited from prior Spikelet work (Imamura & Nakamura, 2021). In the notation of this paper, its output is the one-pass event record

$$\mathcal{S}^a(x) = \{s^a(p) = (p, \ell^a(p), r^a(p), m^a(p)) : \ p \in \mathrm{Ext}(x)\},$$

where the superscript a denotes the anchor support recorded by the prior reduction. Proposition 2 establishes the relation between this record and the exact excursion event of Definition 1.

The term *first pass* includes the terminal Right reductions but does *not* include the right-endpoint completion of Section D.3. For the proofs we analyze the default unwindowed implementation, in which the extremal scan retains the two endpoints and every internal extremum. Let

$$a_0 < a_1 < \cdots < a_{s-1}$$

be the current stack. We call each adjacent anchor pair (a_j, a_{j+1}) a *leg*, with magnitude

$$d_j = |x(a_{j+1}) - x(a_j)|.$$

Definition 3 (Implementation–order Spikelet reduction)**.** After each anchor is pushed—including both endpoints and every internal extremum in scan order—the following rules are applied in the stated order.

(i) **Middle reduction.** For the top four anchors $a < b < c < d$, put $d_0 = |x(b) - x(a)|$, $d_1 = |x(c) - x(b)|$, and $d_2 = |x(d) - x(c)|$. Because the anchors alternate in polarity, test the equivalent direct-value predicate

$$\begin{cases} x(a) < x(c) \text{ and } x(d) \geq x(b), & b \text{ is a maximum,} \\ x(a) > x(c) \text{ and } x(d) \leq x(b), & b \text{ is a minimum.} \end{cases} \tag{M}$$

If (M) holds, record

$$s_b = (b, a, c, x(b) - x(c)), \qquad s_c = (c, b, d, x(c) - x(b)),$$

and delete the two inner anchors b, c from the stack.

(ii) **Left reduction.** For the bottom three anchors $a < b < c$, test

$$\begin{cases} x(c) \leq x(a), & b \text{ is a maximum,} \\ x(c) \geq x(a), & b \text{ is a minimum.} \end{cases} \tag{L}$$

If (L) holds, record $s_b = (b, a, c, x(b) - x(a))$ and delete the left endpoint a from the stack. After either (M) or (L), test (M) and then (L) again; continue until neither rule applies.

(iii) **Right reduction.** After the scan has ended, repeatedly take the top three anchors $a < b < c$, record $s_b = (b, a, c, x(b) - x(c))$, and delete the right endpoint c, until only the two boundary anchors remain.

The Middle reduction is the interior reduction described in the original Spikelet algorithm (Imamura & Nakamura, 2021), while the Left and terminal Right reductions make explicit the boundary handling in the public implementation accompanying that work. For alternating extrema, (M) is

algebraically equivalent in exact arithmetic to $d_0 > d_1$ and $d_2 \geq d_1$, while (L) is equivalent to $|x(c) - x(b)| \geq |x(b) - x(a)|$. We use the direct-value forms above in numerical code because they avoid cancellation from subtracting nearly equal leg magnitudes; no numerical tolerance is introduced. The strict/non-strict comparisons and the Middle-before-Left execution order remain part of the implementation convention.

For the stack argument, the two endpoints are assigned the boundary types that continue the maximum/minimum alternation. Middle deletes two consecutive opposite-type anchors, while Left and Right delete an endpoint, so every active stack continues to alternate in type. The execution is deterministic and terminates on every finite extremal sequence: Middle removes two active anchors, Left or Right removes one, and the next test and its location are fixed by Definition 3. The public implementation also contains a `forced_left_reduction` branch for finite-window operation. Because the present analysis concerns the default unwindowed construction, in which all extrema remain eligible during the scan, that window-specific branch is outside our scope.

D.2 What the first pass fixes

We now prove Proposition 2. The proof separates the algorithmic and geometric parts. First, we identify the two properties carried by the active stack: an ordering of surviving same-polarity anchors and an envelope bound between adjacent anchors. We then translate those properties into the signal-domain statement needed by Definition 1: at each recording state the left neighbour is already the exact support minimiser, while the right neighbour lies inside the eventual excursion. The reduction inequalities then determine the recorded magnitude.

Throughout, $x \in \mathcal{X}_{\mathrm{d}}$ and the scan retains both endpoints and every internal extremum. Every internal extremum is recorded exactly once. A Middle reduction records and removes its two inner anchors. A Left reduction records its middle anchor and removes the current left boundary; an internal anchor later removed as a boundary has already been recorded. The terminal Right phase is symmetric. Thus no internal extremum is omitted or recorded twice.

D.2.1 Active-stack invariant

Call a stack (or stack prefix) *settled* if its surviving maxima strictly decrease and its surviving minima strictly increase from left to right.

Intuition. Reductions remove anchors, but they do not discard the information needed to identify support endpoints. The active stack preserves two facts: same-polarity extrema remain ordered in height, which controls where a higher same-polarity extremum can occur, and each adjacent stack pair bounds the values of all samples eliminated between its endpoints. Lemma 4 uses these two facts to determine what the stack neighbors imply about the support of an event when it is recorded.

Lemma 3 (Active-stack invariant)**.** *During the forward scan, immediately before each reduction test—after a push or after an individual Middle or Left reduction—the stack prefix obtained by deleting the current top anchor is settled. When the reduction loop terminates, the whole stack is settled, and terminal Right reductions preserve this property. In addition, at every such state, for consecutive active anchors $q_j < q_{j+1}$ and every $q_j < t < q_{j+1}$,*

$$\begin{cases} x(q_j) < x(t) \leq x(q_{j+1}), & x(q_j) < x(q_{j+1}), \\ x(q_{j+1}) \leq x(t) < x(q_j), & x(q_j) > x(q_{j+1}), \end{cases}$$

and, if the stack contains at least two anchors, every $t < q_0$ lies in the closed value interval spanned by $x(q_0)$ and $x(q_1)$.

The final boundary condition is needed because, after a Left reduction, the bottom active anchor need not be the original signal endpoint. It guarantees that samples already removed to its left remain bounded by the current bottom stack pair, which is the boundary case used in Lemma 4 to identify the exact left support endpoint.

Proof. We first prove the envelope statements. The most recently scanned anchor is never deleted by a forward reduction, so a newly pushed anchor is adjacent to the preceding endpoint or strict extremum in scan order. Hence the displayed adjacent-envelope bounds hold for the new pair, while

older adjacencies are unchanged. A Middle reduction on $a < b < c < d$ replaces the three adjacent pairs by (a, d). If b is a maximum, (M) gives $x(a) < x(c)$ and $x(d) \geq x(b)$; the three old envelope bounds therefore imply $x(a) < x(t) \leq x(d)$ for every $a < t < d$. The minimum case follows by sign reversal. Left and Right reductions create no new interior adjacency.

The boundary part is preserved as well. A push does not change the bottom pair. For a Left reduction $a < b < c$ with b a maximum, (L) gives $x(c) \leq x(a)$; the old boundary bound together with the adjacent-envelope bounds therefore gives $x(c) \leq x(t) \leq x(b)$ for every $t < b$. The minimum case is dual. A Middle reduction changes the bottom pair only when it acts on the entire four-anchor stack. If its second anchor is a maximum, (M) gives $x(d) \geq x(b)$, so the old interval spanned by $x(a), x(b)$ is contained in the new interval spanned by $x(a), x(d)$; the minimum case is dual. Terminal Right reductions only truncate the stack.

It remains to prove the order property. Before a new anchor is pushed, the preceding reduction loop has terminated, so the whole stack is settled. A push creates only one new same-polarity comparison, namely the comparison involving the new top anchor. Suppose the top four anchors are $a < b < c < d$ and b is a maximum. The settled prefix already gives $x(a) < x(c)$, while a violation of the new maximum order is exactly $x(d) \geq x(b)$; together these are precisely (M). The minimum case is dual. If Middle fires, removing b, c leaves the prefix below the new top unchanged and therefore settled, so the same argument can be repeated.

If four or more anchors remain, Left cannot fire: its bottom same-polarity comparison lies entirely in the settled prefix, whereas (L) is exactly a violation of that ordering. If the cascade reaches three anchors, the only same-polarity comparison is the bottom-to-top one, and (L) fires exactly when that comparison violates the required order. Hence the loop terminates with the whole stack settled. Terminal Right reductions only truncate a settled stack. □

The preceding lemma is purely algorithmic. The next lemma gives its geometric meaning at the instant an event is recorded. Intuitively, the settled stack has already fixed the left support minimizer, whereas the right neighbour is guaranteed only to lie inside the eventual right support; it is exact whenever the right adjacent drop is strictly smaller.

D.2.2 Geometry at a recording state

Lemma 4 (Recording-state geometry)**.** *Suppose Definition 3 records an internal extremum q with left and right stack neighbours $a < q < c$. Then*

$$a = \ell(q), \qquad q < c \leq r(q).$$

Moreover, if

$$|x(q) - x(c)| < |x(q) - x(a)|,$$

then $c = r(q)$. The statements for minima follow by applying the same argument to $-x$.

Proof. It is enough to consider a maximum q.

First consider an active maximum q_j in a settled stack or settled prefix. The adjacent-envelope bounds show that no point of $[q_{j-1}, q_j)$ strictly exceeds $x(q_j)$. If $j \geq 2$, the preceding maximum q_{j-2} is strictly higher than q_j, so $L(q_j) \in [q_{j-2}, q_{j-1})$. Applying the envelope bounds on the two adjacent stack intervals shows that every sample from $L(q_j)$ to q_j is at least $x(q_{j-1})$; any equal value before q_{j-1} loses under the later-index tie convention. Hence $q_{j-1} = \ell(q_j)$. For $j = 1$, the boundary part of Lemma 3 shows that no earlier sample strictly exceeds $x(q_1)$ and that q_0 is the later-index minimum on $[1, q_1]$.

This argument covers every recording state. For Middle, the maximum among the two recorded inner anchors lies in the settled prefix below the current top. Left cannot fire with four or more anchors, so a Left recording is exactly the $j = 1$ boundary case. Terminal Right reductions start from a settled residual stack and preserve settledness. Therefore the recorded left neighbour always satisfies $a = \ell(q)$.

For the right neighbour, the adjacent-envelope bounds give $x(c) \leq x(t) < x(q)$ for $q < t < c$. Thus $c \leq R(q)$, and c is the later-index minimum on $[q, c]$. Extending this interval to $[q, R(q)]$ can move the later-index minimiser only to c or to its right, so $c \leq r(q)$.

It remains to prove the final assertion. A strict inequality in favor of the right adjacent drop cannot occur for Left, because (L) makes the left drop no larger. In Middle it can occur for a maximum only when $q = b$ in $a < b < c < d$. Then (M) gives $x(d) \geq x(q)$, hence $R(q) \leq d$. The envelope bounds make c the later-index minimum on $[q, d]$, and since $c \leq R(q) \leq d$, we obtain $r(q) = c$.

For a terminal Right event, view q and c in the settled residual stack at the start of the terminal phase. Same-polarity order places every later maximum strictly below $x(q)$ and every later minimum strictly above $x(c)$, and the envelope bounds extend these inequalities to the samples between successive anchors. Hence no later sample reaches $x(q)$ and no later sample falls below $x(c)$; the later-index rule also favors c over any equal value to its left. Thus $R(q) = N$ and $r(q) = c$. □

Proof of Proposition 2. Let q be a recorded maximum with neighbours $a < q < c$, and set

$$\delta^- = x(q) - x(a), \qquad \delta^+ = x(q) - x(c), \qquad D^- = x(q) - x(\ell(q)), \qquad D^+ = x(q) - x(r(q)).$$

Lemma 4 gives $\delta^- = D^-$ and $\delta^+ \leq D^+$, as well as $\ell^a(q) = \ell(q)$ and $I^a(q) \subseteq I(q)$.

The reduction rules store the smaller adjacent drop. In Middle, (M) is equivalent to $d_0 > d_1$ and $d_2 \geq d_1$, so the common inner leg is the smaller adjacent leg of each recorded inner event. In Left, (L) makes the left leg no larger than the right. In terminal Right, settled same-polarity order makes the right leg smaller. Therefore

$$|m^a(q)| = \min\{\delta^-, \delta^+\}.$$

If $\delta^- \leq \delta^+$, then $D^- = \delta^- \leq \delta^+ \leq D^+$, so the exact magnitude is determined by the already exact left endpoint. If $\delta^+ < \delta^-$, Lemma 4 gives $c = r(q)$ and hence $\delta^+ = D^+$. Thus in either case

$$|m^a(q)| = \min\{D^-, D^+\} = |m(q)|.$$

The sign agrees with the event polarity, and the minimum case follows by sign reversal. Hence the first pass fixes every signed magnitude and every left support endpoint, while the recorded right endpoint may remain provisional. □

D.3 Linear-time support completion

Proposition 2 shows that only the right support endpoint can remain provisional after the first pass. It can be completed for all positive apexes by a forward all-sample monotone-stack sweep.

Process $i = 1, \ldots, N$. Before pushing i, pop top entries q while $x(i) \geq x(q)$. Each active entry stores the later-index minimiser of the contiguous block it represents, and active blocks are adjacent in stack order. During a pop cascade, merge the popped blocks into an accumulator from right to left. When q is popped, its stored block is merged with the accumulator; the resulting block is exactly $[q, i-1]$.

Intuition. For a positive apex q, a pop occurs at the first later sample that reaches height $x(q)$, which identifies $R(q)$. The accumulated block simultaneously stores the later-index minimizer needed for $r(q)$; if no such sample occurs, the final flush gives the boundary case $R(q) = N$.

Lemma 5 (Forward-completion correctness and cost)**.** *For every positive apex q, the forward sweep returns the exact $R(q)$ and $r(q)$. Applying the same construction to $-x$ completes negative events. The total cost is $O(N)$ time and $O(N)$ memory.*

Proof. If a positive apex q is popped at index i, then i is the first later index satisfying $x(i) \geq x(q)$; otherwise q would have been popped at an earlier such index. Hence $i = R(q)$. By the block invariant above, immediately after q is merged into the pop accumulator, that accumulator represents exactly $[q, i-1]$ and stores its later-index minimiser. Since q is a strict maximum and $x(i) \geq x(q)$, adjoining i cannot change that minimum, so the stored index is exactly $r(q)$.

If q survives to the boundary, the final right-to-left flush applies the same block-merging invariant through index N, giving $R(q) = N$ and the later-index minimiser on $[q, N]$. Each index is pushed once and popped at most once, so the sweep uses $O(N)$ stack operations and $O(N)$ memory. Sign reversal handles negative events. □

A symmetric backward sweep, using the strict pop condition $x(i) > x(q)$, can reconstruct $L(q), \ell(q)$ directly from Definition 1 if desired. It is not needed to complete the prior first-pass record because Proposition 2 has already established the exact left endpoint.

D.4 PRESERVATION UNDER MAGNITUDE PRUNING

Proof of Theorem 1. Let $A_\tau(x)$ be the retained-apex set defined in Section 4. By construction, $\hat{x}_\tau(p) = x(p)$ for every $p \in A_\tau(x)$. On each open interval between consecutive interpolation points—the two endpoints and the retained apexes—$\hat{x}_\tau$ is affine and therefore has no strict internal extremum. Hence every strict internal extremum of $\hat{x}_\tau$ must occur at a retained apex, so

$$\mathrm{Ext}(\hat{x}_\tau) \subseteq A_\tau(x) \subseteq \mathrm{Ext}(x).$$

The count inequality follows immediately. The same argument applies to any interpolation that is monotone between consecutive retained apexes; unconstrained higher-order interpolation may introduce additional extrema. The separate continuous time-warp observation used in Section 4 follows instead from the ordered-extrema completeness result in Appendix C.2. □

E PROOF OF POLARITY–SPECIFIC NONCROSSING SUPPORT FAMILIES

Theorem 2 concerns the completed excursion supports of Definition 1; its proof is given in Appendix C. There are two different reasons for noncrossing. For the completed supports $I(p)$, noncrossing is a geometric consequence of Definition 1: Appendix C shows that same-polarity first-passage excursion ranges are nested or disjoint and that the argmin-defined support endpoints cannot interleave. The completion sweep of Appendix D.3 only recovers these already-defined endpoints.

The one-pass anchor supports $I^{\mathrm{a}}(p)$ are different: they are created by the reduction sequence itself. Their noncrossing is therefore an algorithmic invariant. Containment $I^{\mathrm{a}} \subseteq I$ alone is not enough, since shrinking individual intervals does not in general preserve noncrossing. We instead track the local territories carried by the active anchors through the three moves of Definition 3.

Proposition 5 (One-pass anchor-support noncrossing)**.** *For every $x \in \mathcal{X}_{\mathrm{d}}$, the positive anchor supports $\{I^{\mathrm{a}}(p)\}$ recorded by the implementation-order reduction of Definition 3 form a noncrossing family, and so do the negative anchor supports. Magnitude pruning preserves this property.*

The active anchors alternate in maximum/minimum type (Appendix D). For an active anchor q, let $\mathcal{F}(q)$ be the same-polarity supports already recorded and currently carried by q. Initially every such family is empty. Middle transfers the family of each deleted inner anchor to the surviving anchor of the same polarity; Left and Right may instead delete a boundary carrier, in which case its family is *sealed*: it is thereafter permanently separated from all supports recorded later and no longer participates in the induction.

An active internal anchor is open until its own event is recorded. A recorded anchor that remains active can only be a boundary anchor left by a Left or Right reduction; call such a boundary *closed*. Domain endpoints are never closed. For the active stack

$$q_0 < q_1 < \cdots < q_m,$$

define the territory of q_i by

$$T(q_i) = \begin{cases} \{q_0\}, & m = 0, \\ [q_0, q_1], & i = 0,\ m \geq 1,\ q_0 \text{ open}, \\ [1, q_1], & i = 0,\ m \geq 1,\ q_0 \text{ closed}, \\ [q_{i-1}, q_{i+1}], & 0 < i < m, \\ [q_{m-1}, q_m], & i = m,\ m \geq 1,\ q_m \text{ open}, \\ [q_{m-1}, N], & i = m,\ m \geq 1,\ q_m \text{ closed}. \end{cases}$$

Because active types alternate, territories of distinct active anchors of the same polarity have disjoint interiors.

Intuition. Each active anchor owns a local territory for the same-polarity supports it carries. Middle reductions merge such histories into enlarged territories, whereas Left and Right reductions seal a boundary family. Thus noncrossing follows if every carried family remains noncrossing inside its current territory.

Lemma 6 (Territory invariant)**.** *After every push or reduction, for each active anchor q:*

(i) *the supports in $\mathcal{F}(q)$ are noncrossing; and*

(ii) *every support in $\mathcal{F}(q)$ is contained in $T(q)$.*

Moreover, once a family is sealed by a Left or Right reduction, it is permanently separated from every support recorded later.

Proof. Initially all carried families are empty, so the claim is immediate.

Push. A newly pushed anchor carries an empty family. The territory of the previous scan-front anchor can only enlarge, and all other carried families and territories are unchanged.

Middle reduction. Let $a < b < c < d$ be the top four active anchors. The two inner anchors are open, and their current territories are

$$T(b) = [a, c], \qquad T(c) = [b, d].$$

The reduction records exactly these two supports. Adding $[a, c]$ to $\mathcal{F}(b)$ and $[b, d]$ to $\mathcal{F}(c)$ therefore adds an outer interval around each carried family and cannot create a same-polarity crossing.

The enlarged b-family is then transferred to the same-polarity survivor d, and the enlarged c-family to the same-polarity survivor a, before b, c are deleted. Before the move, the territory of d lies to the right of c and meets $[a, c]$ only at c; similarly, the territory of a lies to the left of b and meets $[b, d]$ only at b. Hence each transferred family is noncrossing with the family already carried by its survivor. After the deletion the two survivor territories enlarge to contain the corresponding unions, so both parts of the invariant are preserved. The newly recorded supports $[a, c]$ and $[b, d]$ may cross, but their apexes have opposite polarity.

Left reduction. By Lemma 3, Left cannot fire while four or more anchors remain. Thus at a Left recording the active stack is $a < b < c$, and the recorded support $[a, c]$ is exactly $T(b)$. Adding it preserves noncrossing of $\mathcal{F}(b)$. After deleting a, the recorded apex b becomes the closed left boundary, whose territory expands to $[1, c]$, so containment is preserved.

The family carried by the deleted endpoint a lies in $[1, b]$ and is now sealed. Every support recorded later has both endpoints at or to the right of b, hence lies in $[b, N]$. The sealed family can therefore never cross a future same-polarity support. In particular it is also separated from every family sealed at a later stage.

Right reduction. This is the mirror image of Left. For the top three anchors $a < b < c$, the new support $[a, c]$ is exactly $T(b)$. After deleting c, the recorded apex b becomes the closed right boundary and its territory expands to $[a, N]$. The family carried by c lies in $[b, N]$ and is sealed, while every support recorded later in the terminal Right phase lies in $[1, b]$. Hence the sealed family is permanently separated from all later same-polarity supports.

The strict and nonstrict inequalities in (M) and (L) determine which local move is executed but do not alter these endpoint-order arguments, so equality cases are covered by the fixed implementation convention. □

Proof of Proposition 5. Apply Lemma 6 through the complete deterministic reduction. Supports carried by one active anchor are noncrossing. Supports carried by distinct active anchors of the same polarity lie in territories with disjoint interiors, and every sealed family is permanently separated from all supports recorded later. By chronological induction this also separates any two sealed families. Hence the complete positive anchor-support family is noncrossing, and the same holds for the negative family. Magnitude pruning selects a subfamily, so it preserves noncrossing. Opposite-polarity supports may cross already in a Middle reduction, because $[a, c]$ and $[b, d]$ have interleaving endpoints; no statement is made for their union. □

F PROOFS FOR THE STRUCTURAL COMPARISON RESULTS

Roadmap. Proposition 3 is independent of the false-extrema chain. Proposition 4 gives the abstract criterion: on an interval where the original signal is monotone, a local oscillatory response

that dominates residual interference forces false extrema. Together, the common-response setup and Definition 4 organize the argument around three ingredients: the response identity (R), local oscillation (O), and tail control $(T)_\alpha$. Lemma 7 converts (O)+$(T)_\alpha$ together with transition separation into the dominance condition of Proposition 4. Lemmas 8 and 9 verify these ingredients for Fourier truncation and periodic spline projection, yielding Corollaries 2 and 3. The approximation-theoretic localization ingredients are classical; where useful we cite that theory and retain direct calculations only for the finite-width sign certificates and explicit constants needed by the quantitative E3 audit. Section F.6 collects results independent of this chain.

F.1 Strict extremal contraction by positive smoothing

Reader guide. Proposition 3 is an averaging argument. Because a strictly positive normalized kernel mixes each point with values from the entire period, a nonconstant signal cannot retain its global maximum or minimum exactly: both are pulled strictly inward. The proof below makes this intuition exact; no oscillation argument is involved.

Proof of Proposition 3. Let

$$M = \max_{\mathbb{T}} x.$$

For every $t \in \mathbb{T}$,

$$M - (K * x)(t) = \int_{\mathbb{T}} K(u)\big(M - x(t-u)\big)\, du.$$

The integrand is nonnegative. Since x is continuous and nonconstant,

$$A = \{s \in \mathbb{T} : x(s) < M\}$$

is a nonempty open subset of $\mathbb{T}$ and therefore has positive measure. For each t, its translate

$$A_t = \{u \in \mathbb{T} : t - u \in A\}$$

also has positive measure. On A_t,

$$M - x(t-u) > 0,$$

and $K(u) > 0$ for almost every u. Hence

$$M - (K * x)(t) > 0 \qquad \text{for every } t \in \mathbb{T}.$$

Thus

$$(K * x)(t) < M \qquad \text{for every } t \in \mathbb{T}.$$

Because convolution of an L^1 kernel with a continuous periodic function is continuous, $K * x$ attains its maximum on the compact domain $\mathbb{T}$, and hence

$$\max_{\mathbb{T}}(K * x) < \max_{\mathbb{T}} x.$$

Applying the same argument to $-x$ gives

$$\min_{\mathbb{T}}(K * x) > \min_{\mathbb{T}} x.$$

Combining the two strict inequalities gives

$$\min_{\mathbb{T}} x < \min_{\mathbb{T}}(K * x) \le \max_{\mathbb{T}}(K * x) < \max_{\mathbb{T}} x.$$

□

Remark 2 (Scope of Proposition 3). Strict positivity of K is what upgrades the usual averaging inequalities to strict inequalities for every nonconstant continuous signal. The proposition does not claim that every local-extremum location moves, nor that every nonnegative compactly supported kernel strictly contracts extrema for every signal. It also does not apply to the identity operator. Its direct Gaussian consequence follows because the periodic Gaussian is strictly positive at every nonzero smoothing scale.

F.2 LOCAL OSCILLATION CRITERION AND COMMON-RESPONSE SETUP

Reader guide. The false-extrema comparison has one common mechanism. Monotonicity does not imply slow variation: a transition of amplitude a over width w has slope scale a/w, so even a small-amplitude change can be sharp when concentrated over a short interval. Figure 8 illustrates this mechanism for the transition family used below. The original ramp is monotone and contains no strict local extremum, whereas a fixed-resolution Fourier reconstruction can oscillate around it and create false extrema. Proposition 4 abstracts this mechanism: the local oscillatory response must dominate residual interference from the other transitions.

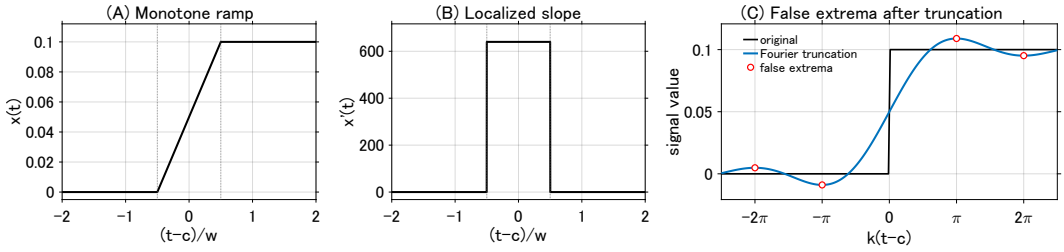

Figure 8: **Reader guide for the false-extrema mechanism.** (A) A narrow ramp is monotone and contains no strict interior local extremum. (B) Its derivative has scale a/w, so even a small amplitude a can correspond to a sharp local transition when the width w is small. (C) A fixed-resolution Fourier reconstruction can oscillate around the transition and create false extrema (red circles) although the original signal remains monotone. Panels A–B use the transition-width coordinate $(t-c)/w$, while Panel C uses the resolution coordinate $k(t-c)$.

Proof of Proposition 4. Because $|r'| < \beta$, the residual cannot reverse the signs of g' on the three witness intervals. Hence $(P_k x)'$ has the sign pattern $+,-,+$ there (or $-,+,-$ in the sign-reversed case). In the $+,-,+$ case, choose $a_1 < a_2$ in the first witness interval, $b_1 < b_2 < c_1 < c_2$ in the second, and $d_1 < d_2$ in the third. Absolute continuity gives

$$(P_k x)(a_2) > (P_k x)(a_1), \qquad (P_k x)(b_1) > (P_k x)(b_2),$$

so the maximum of $P_k x$ on $[a_1, b_2]$ is attained at an interior point, giving a local maximum. Likewise,

$$(P_k x)(c_1) > (P_k x)(c_2), \qquad (P_k x)(d_2) > (P_k x)(d_1),$$

so the minimum on $[c_1, d_2]$ is attained at an interior point, giving a local minimum. Since $b_2 < c_1$, the two extrema are distinct. The $-,+,-$ case is symmetric. Since x is monotone on the ambient interval, both are false relative to x. □

Proposition 4 guarantees the existence of the false local extrema but does not require them to be strict. Strictness in Corollaries 2 and 3 follows from method-specific regularity: real analyticity for Fourier truncation and nonzero alternating cell slopes for the periodic spline projection.

We now specialize this mechanism to an absolutely continuous periodic transition train. Throughout Sections F.2–F.5, let $\mathbb{T} = \mathbb{R}/(2\pi\mathbb{Z})$ and $|u|_{\mathbb{T}} = \min_{n\in\mathbb{Z}} |u + 2\pi n|$. The train is defined, up to an additive constant, by the following derivative identity almost everywhere, with transition intervals interpreted on $\mathbb{T}$:

$$x'(t) = \sum_{j=1}^{m} \frac{h_j}{w} \mathbf{1}_{[c_j - w/2, c_j + w/2]}(t), \qquad m \geq 2, \quad w > 0, \quad h_j \neq 0, \quad \sum_{j=1}^{m} h_j = 0.$$

The centres are circularly ordered, with

$$\delta := \min_{i \neq j} |c_i - c_j|_{\mathbb{T}}, \qquad h_{\min} := \min_j |h_j|, \qquad h_{\max} := \max_j |h_j|.$$

As in the main text, k is the Fourier harmonic cutoff or, for the periodic spline projector, the number of knot intervals.

For a fixed linear reconstruction P_k, let $\Gamma_k \subseteq \mathbb{T}$ be a set of admissible transition centres. We say that P_k has a *common transition response* $\overline{K}_{k,w}$ on Γ_k if, for every such train with $c_j \in \Gamma_k$, $P_k x$ is continuous and absolutely continuous between finitely many breakpoints and

$$(P_k x)'(t) = \sum_{j=1}^{m} g_j'(t), \qquad g_j'(t) := h_j \overline{K}_{k,w}(t - c_j) \quad \text{for a.e. } t. \tag{R}$$

For transition i, write

$$r_i'(t) := \sum_{j \neq i} g_j'(t).$$

Thus g_i' is the local term and r_i' is the interference term in Proposition 4.

Definition 4 (Tail-controlled oscillatory response)**.** A family of common responses $\overline{K}_{k,w}$, indexed by admissible pairs (k, w), is *tail-controlled oscillatory with exponent* $\alpha \geq 1$ if there are constants $r, c_1, C_K > 0$, uniform over all admissible pairs, such that $\rho_k = r/k$ and:

(O) *(local oscillation)* there are three ordered nonempty open intervals $J_{k,w,1} < J_{k,w,2} < J_{k,w,3} \subset [0, \rho_k]$ on which $\overline{K}_{k,w} \geq c_1 k$, $\overline{K}_{k,w} \leq -c_1 k$, and $\overline{K}_{k,w} \geq c_1 k$, respectively, for a positive unit transition;

(T)$_\alpha$ *(tail control)*

$$|\overline{K}_{k,w}(u)| \leq \frac{C_K}{k^{\alpha-1}|u|_{\mathbb{T}}^{\alpha}} \qquad \text{for a.e. } |u|_{\mathbb{T}} \geq \rho_k.$$

Condition (O) gives the local-response magnitude

$$\beta_i := c_1 |h_i| k.$$

Throughout Appendix F we use

$$c_0 := \frac{1}{10}.$$

F.3 A SINGLE DOMINANCE LEMMA FOR SEPARATED TRANSITIONS

The next lemma is the bridge from the method-specific response bounds to Proposition 4. Condition (O) supplies the target local oscillation, while (T)$_\alpha$ and transition separation ensure that interference from all other transitions is too small to reverse it. It is the only place where responses from different transitions are combined.

Lemma 7 (Separated transitions imply local dominance)**.** *Suppose the common response satisfies (O) and (T)$_\alpha$ with constants r, c_1, C_K, and let $m \geq 2$. Define*

$$S_\alpha(m) := \sum_{n=1}^{m-1} n^{-\alpha}, \qquad C_K' := 2^{\alpha+1} C_K, \qquad \kappa_\alpha := \max\left\{ (2r)^\alpha, \frac{C_K'}{c_1} \right\}.$$

If $0 < c_0 \leq r$ and

$$kw \leq c_0, \qquad (k\delta)^\alpha > \kappa_\alpha \frac{h_{\max}}{h_{\min}} S_\alpha(m), \tag{F.1}$$

then the cores $I_i = [c_i - \rho_k, c_i + \rho_k]$ are disjoint, the original transition train is monotone on each I_i, and

$$\operatorname*{ess\,sup}_{t \in I_i} |r_i'(t)| \leq \frac{C_K' h_{\max}}{k^{\alpha-1}\delta^\alpha} S_\alpha(m) < c_1 |h_i| k = \beta_i. \tag{F.2}$$

Hence the hypotheses of Proposition 4 hold on the three translated witness intervals in every core. Moreover,

$$S_1(m) \leq 1 + \ln m, \qquad S_\alpha(m) \leq \zeta(\alpha) \quad (\alpha > 1).$$

Proof. Because $S_\alpha(m) \geq 1$ and $h_{\max}/h_{\min} \geq 1$, the second condition in (F.1), together with $\kappa_\alpha \geq (2r)^\alpha$, gives $\delta > 2\rho_k$. Thus the cores are disjoint; since $kw \leq c_0 \leq r$, no other transition interval meets I_i, so x is monotone there.

For $t \in I_i$ and $j \neq i$,

$$|t - c_j|_{\mathbb{T}} \geq |c_i - c_j|_{\mathbb{T}} - \rho_k \geq \tfrac{1}{2}|c_i - c_j|_{\mathbb{T}}.$$

Since $|c_i - c_j|_{\mathbb{T}} \geq \delta > 2\rho_k$, this also gives $|t - c_j|_{\mathbb{T}} > \rho_k$, so $(\mathrm{T})_\alpha$ applies. Partition the remaining centres according to whether a shortest arc from c_i reaches them clockwise or counterclockwise, breaking an antipodal tie arbitrarily. Within either group, order the centres by increasing shortest-arc distance from c_i. Consecutive centres in that ordering are separated by at least δ, so the nth centre in either group is at circular distance at least $n\delta$. Hence

$$\sum_{j \neq i} |c_i - c_j|_{\mathbb{T}}^{-\alpha} \leq 2 \sum_{n=1}^{m-1} (n\delta)^{-\alpha} = 2\delta^{-\alpha} S_\alpha(m).$$

Using $(\mathrm{T})_\alpha$ and $|h_j| \leq h_{\max}$ therefore gives

$$|r_i'(t)| \leq \frac{2^\alpha C_K h_{\max}}{k^{\alpha-1}} \sum_{j \neq i} |c_i - c_j|_{\mathbb{T}}^{-\alpha} \leq \frac{2^{\alpha+1} C_K h_{\max}}{k^{\alpha-1}\delta^\alpha} S_\alpha(m).$$

The separation condition and $\kappa_\alpha \geq C_K'/c_1$ make this bound strictly smaller than $c_1 h_{\min} k \leq c_1 |h_i| k = \beta_i$, proving (F.2). The stated estimates for S_α are the standard harmonic- and p-series bounds. □

Remark 3 (Fast-tail consequence). If $\alpha > 1$, then $S_\alpha(m) \leq \zeta(\alpha)$. Under the same width restriction $kw \leq c_0$, the separation inequality in (F.1) is implied by the m–independent bound

$$k\delta > \big(\kappa_\alpha \zeta(\alpha)\big)^{1/\alpha} (h_{\max}/h_{\min})^{1/\alpha}.$$

Thus fast localization suppresses cross-transition interference without removing the local oscillation (O).

F.4 Fourier instantiation

The oscillatory behavior of Fourier partial sums near jumps is the classical Gibbs phenomenon (Gottlieb & Shu, 1997). Classical Gibbs theory describes the local behavior near jumps of a fixed function. Here the derivative of a width–w ramp has two opposite jumps separated by $w = O(k^{-1})$, so their responses interact on the same scale as the first Dirichlet side lobes. We therefore retain a short finite-k estimate for the window–averaged Dirichlet response. The standard tail bound is recorded with its explicit constant because that constant enters the E3 certificate.

Proof strategy. Convolution gives (R), the first three averaged Dirichlet side lobes give (O), and Jordan's inequality gives $(\mathrm{T})_1$; Lemma 7 then yields the sufficient condition in Corollary 2.

Lemma 8 (Finite-width Dirichlet response)**.** *If $k \geq 3$ and $kw \leq c_0$, Fourier truncation has (R) for $\Gamma_k = \mathbb{T}$, and its window–averaged Dirichlet response satisfies (O) and $(\mathrm{T})_1$ with*

$$r = \frac{5\pi}{2}, \qquad c_1 = \frac{1}{5\pi^2}, \qquad C_K = \frac{2}{\pi}.$$

Proof. Fourier truncation is convolution: $P_k x = K_k * x$ with $K_k = \frac{1}{\pi} D_k$, where

$$D_k(t) := \frac{\sin((k + \frac{1}{2})t)}{2\sin(t/2)}, \qquad D_k(0) := k + \tfrac{1}{2}$$

by continuous extension. Hence (R) holds for arbitrary transition centres and

$$\overline{K}_{k,w}(u) = \frac{1}{\pi} \overline{D}_{k,w}(u), \qquad \overline{D}_{k,w}(u) := \frac{1}{w} \int_{-w/2}^{w/2} D_k(u - s)\, ds.$$

Finite-width local oscillation. Put $\nu = k + \frac{1}{2}$ and $t_q = (2q+1)\pi/(2\nu)$ for $q = 0, 1, 2$, with $\sigma_q = (-1)^q$. If $kw \leq 0.1$ and $|s| \leq w/2$, then

$$\nu|s| \leq 0.075, \qquad \sigma_q \sin(\nu(t_q - s)) = \cos(\nu s) \geq \cos(0.075) > 0.$$

The denominator is positive on these windows, and $2\sin((t_q - s)/2) \leq t_q - s \leq t_q + w/2$. Therefore

$$\sigma_q \overline{D}_{k,w}(t_q) \geq \frac{2\nu \cos(0.075)}{5\pi + \nu w} \geq \frac{2k \cos(0.075)}{5\pi + 0.15} > \frac{k}{5\pi}.$$

After multiplication by $1/\pi$, the strict inequality and continuity give three ordered open neighbourhoods with signs $+, -, +$ and magnitude at least $c_1 k = k/(5\pi^2)$. Thus (O) holds with $r = 5\pi/2$.

Tail control. Let $d = |u|_{\mathbb{T}} \geq \rho_k = 5\pi/(2k)$. Then $|u - s|_{\mathbb{T}} \geq d - w/2$, and Jordan's inequality gives $|K_k(u - s)| \leq [2(d - w/2)]^{-1}$. Since

$$\frac{w}{2d} \leq \frac{c_0}{5\pi} < 1 - \frac{\pi}{4},$$

one has $[2(d - w/2)]^{-1} \leq (2/\pi)d^{-1}$. Averaging over s proves $(\mathrm{T})_1$ with $C_K = 2/\pi$. □

Corollary 2 (Fourier truncation)**.** *Fix $k \geq 3$, and let P_k retain harmonics $|n| \leq k$. For transition i, define*

$$I_i^{\mathrm{F}} := \left[c_i - \frac{5\pi}{2k},\, c_i + \frac{5\pi}{2k}\right] \quad \text{on } \mathbb{T}.$$

If

$$w \leq \frac{1}{10k}, \qquad \delta > \frac{40\pi}{k}\,\frac{h_{\max}}{h_{\min}}(1 + \ln m),$$

then x is monotone on every I_i^{F} and every such core contains at least one strict false local maximum and at least one strict false local minimum.

Proof. The restriction $k \geq 3$ ensures $\rho_k = 5\pi/(2k) < \pi$, so each Fourier core is represented by a proper local arc of $\mathbb{T}$. Lemma 8 gives $\alpha = 1$, $r = 5\pi/2$, $c_1 = 1/(5\pi^2)$, and $C_K = 2/\pi$, hence

$$C_K' = \frac{8}{\pi}, \qquad \kappa_1 = \max\left\{5\pi, \frac{8/\pi}{1/(5\pi^2)}\right\} = 40\pi.$$

Since $S_1(m) \leq 1 + \ln m$, the stated width and separation conditions imply (F.1). Lemma 7 therefore supplies local dominance, and Proposition 4 gives the false maximum–minimum pair. Since $P_k x$ is a nonconstant trigonometric polynomial, it is real analytic. Any non-strict local extremum would force $P_k x - P_k x(\xi)$ to have zeros accumulating at the extremum ξ, hence to vanish identically. Therefore these extrema are strict. □

F.5 Periodic order–2 spline instantiation

A transition center is called *knot-aligned* when it coincides with a knot of the periodic spline grid.

The spline case follows the same abstract template as Fourier, but the fixed-knot projector is equivariant only under knot shifts, so establishing (R) requires a cyclic calculation. We use the dyadic periodic order–2 spline space associated with the Battle–Lemarié construction (Battle, 1987), although only the projection space enters the argument. Classical spline Gibbs behavior (Richards, 1991) provides the oscillatory background, and inverse-matrix decay (Demko et al., 1984) provides the localization background. The finite-width sign certificate and explicit constants used by E3 are derived directly below.

Corollary 3 (Periodic spline projection)**.** *Fix $k = 2^\ell \geq 8$, and let P_k be the orthogonal projector onto the dyadic order–2 periodic spline space with k knot intervals. The proof below uses only that k is even and $k \geq 8$; the dyadic restriction is retained to match the approximation spaces and experiments considered here. For knot–aligned transition centres, it admits a common single-transition response $\overline{K}_{k,w}$. For transition i, define*

$$I_i^{\mathrm{spl}} := \left[c_i - \frac{6\pi}{k},\, c_i + \frac{6\pi}{k}\right] \quad \text{on } \mathbb{T}.$$

There exists a method constant $C_{\mathrm{spl}} > 0$, independent of k, m, w, δ, and the transition heights, such that if

$$w \leq \frac{1}{10k}, \qquad \delta > \frac{C_{\mathrm{spl}}}{k}\left(\frac{h_{\max}}{h_{\min}}\right)^{1/2},$$

then x is monotone on every I_i^{spl} and every such core contains at least one strict false local maximum and at least one strict false local minimum. An explicit admissible C_{spl} is given below.

Intuition. The inverse spline Gram matrix yields an alternating, geometrically decaying response: after the finite-width correction, the first cells give (O), the same decay gives $(T)_2$, and periodic shift structure gives (R).

Lemma 9 (Periodic order–2 spline response)**.** *Let $k \geq 8$ be even. For knot-aligned transition centres, the periodic spline projector has the common-response form (R). If $kw \leq 0.1$, its finite-width response satisfies (O) with*

$$r = 6\pi, \qquad c_1 = \frac{33 - 19\sqrt{3}}{8\pi}.$$

Moreover, it satisfies $(T)_2$ with

$$C_K^{\mathrm{per}}(2) := \frac{(1+q)^2}{4\pi(1-q^8)} \left(\frac{16\pi}{e\gamma}\right)^2, \qquad q := 2 - \sqrt{3}, \qquad \gamma := -\log q.$$

Proof. The fixed-knot periodic spline space is invariant under integer knot shifts, and its cyclic Gram matrix commutes with cyclic shifts and the cyclic difference operator. The orthogonal projector also fixes constants. We first compute the single-transition slope response on the unit-knot line and its periodization; after that calculation we return to the cyclic normal equations to establish (R) for the balanced transition train. No arbitrary translation covariance is used.

Exact local response. Work first at unit knot spacing. Let B_n be the usual hat functions. Their Gram matrix is

$$G = \tfrac{1}{6}\,\mathrm{tridiag}(1, 4, 1).$$

The general exponential decay of G^{-1} is consistent with the classical decay result for inverses of band matrices (Demko et al., 1984); here the inverse is explicit. With

$$q := 2 - \sqrt{3}, \qquad \rho := -q, \qquad A := \frac{1+q}{2},$$

one has

$$(G^{-1})_{nn'} = \sqrt{3}\,\rho^{|n-n'|}.$$

For the unit step H jumping at the knot 0, let $Gc = \beta$ with $\beta_n := \langle H, B_n \rangle$. The absolutely summable rows of G^{-1} give the unique bounded coefficient sequence

$$c_n := \sum_{n' \in \mathbb{Z}} (G^{-1})_{nn'} \beta_{n'},$$

whose consecutive differences are

$$s_n := c_{n+1} - c_n = A(-1)^n q^n \quad (n \geq 0), \qquad s_{-n-1} = s_n.$$

In particular,

$$s_2 = Aq^2 = \frac{33 - 19\sqrt{3}}{2}.$$

For even period k, periodization means

$$s_n^{(k)} := \sum_{r \in \mathbb{Z}} s_{n+rk}, \qquad 0 \leq n \leq k-1,$$

and summing the two geometric tails gives

$$s_n^{(k)} = A(-1)^n \frac{q^n - q^{k-n-1}}{1 - q^k}, \qquad 0 \leq n \leq k-1, \tag{F.3}$$

and for even $k \geq 8$, $n = 0, 1, 2$, the signs are $+, -, +$ and

$$\frac{q^n - q^{k-n-1}}{1 - q^k} \geq q^n(1 - q^3). \tag{F.4}$$

For the actual ramp, let $h = 2\pi/k$ and $\eta = w/h = kw/(2\pi)$. Under $kw \le 0.1$, $0 < \eta < 1$. Define the unit-height linear ramp centred at 0 by

$$R_\eta(t) = \begin{cases} 0, & t \le -\eta/2, \\ t/\eta + 1/2, & |t| < \eta/2, \\ 1, & t \ge \eta/2, \end{cases}$$

and put $d_\eta := R_\eta - H$. Direct integration against the only three affected hat functions gives

$$\langle d_\eta, B_{-1}\rangle = \frac{\eta^2}{48}, \qquad \langle d_\eta, B_0\rangle = 0, \qquad \langle d_\eta, B_1\rangle = -\frac{\eta^2}{48}. \tag{F.5}$$

The induced coefficient correction is

$$\Delta c_n = \frac{\sqrt{3}\,\eta^2}{48}\left(\rho^{|n+1|} - \rho^{|n-1|}\right).$$

Thus, with $\Delta s_n := \Delta c_{n+1} - \Delta c_n$,

$$\Delta s_n = -\frac{\eta^2}{4} s_n \quad (n \ge 1), \qquad \Delta s_0 = -\frac{\eta^2}{4} s_0 + \frac{\eta^2}{8}, \tag{F.6}$$

with reflected identities on the negative cells. Let $\widetilde{s}_n := s_n + \Delta s_n$ and define its periodization by

$$\widetilde{s}_n^{(k)} := \sum_{r\in\mathbb{Z}} \widetilde{s}_{n+rk}, \qquad 0 \le n \le k-1.$$

Then, for $0 \le n \le k-1$,

$$\widetilde{s}_n^{(k)} = \left(1 - \frac{\eta^2}{4}\right) s_n^{(k)} + \frac{\eta^2}{8}\,\mathbf{1}_{\{0,k-1\}}(n). \tag{F.7}$$

Common-response conclusion. Here G_k denotes the cyclic Gram matrix in the unit-knot coordinate used above. Equivalently, the Gram matrix in the physical coordinate t with knot spacing $h = 2\pi/k$ is hG_k; the same factor h appears in the load vector and therefore cancels from the normal equations. Let ∇ be the cyclic coefficient-difference operator, $(\nabla c)_n := c_{n+1} - c_n$. If the knot-aligned centre satisfies $c_j = n_j h$ modulo 2π, define the cyclic shift by

$$(T_{n_j} v)_n := v_{(n-n_j) \bmod k}.$$

Since G_k is circulant, $G_k\nabla = \nabla G_k$. Hence the cyclic normal equations $G_k c = \beta$ give

$$\nabla c = G_k^{-1}\nabla\beta.$$

For one knot-aligned width–w transition, the infinite-line calculation above gives the coefficient-difference response $\widetilde{s}$. Since $G\nabla = \nabla G$, it satisfies

$$G\widetilde{s} = \nabla\beta^{\mathrm{ramp}}.$$

The sequence $\widetilde{s}$ is absolutely summable and G has finite bandwidth, so periodization gives

$$G_k\widetilde{s}^{(k)} = \mathrm{per}_k(\nabla\beta^{\mathrm{ramp}}),$$

where per_k denotes periodization modulo k.

To identify the load difference for the full train, let Ψ_n be the compactly supported antiderivative of $B_n - B_{n+1}$. Since the two hat functions have equal integral, such a Ψ_n exists, and integration by parts gives

$$(\nabla\beta)_n = \langle x, B_{n+1} - B_n\rangle = \langle x', \Psi_n\rangle.$$

The same identity holds after periodization. Since the derivative of the balanced transition train is the sum of the corresponding cyclic shifts of the ramp derivatives,

$$\nabla\beta^{\mathrm{train}} = \sum_{j=1}^{m} h_j T_{n_j}\,\mathrm{per}_k(\nabla\beta^{\mathrm{ramp}}).$$

Because G_k^{-1} is circulant,

$$\nabla c = \sum_{j=1}^{m} h_j T_{n_j} \widetilde{s}^{(k)}.$$

Constants contribute no derivative term (and the projector fixes constants). Defining the single-transition derivative response by

$$\overline{K}_{k,w}(u) := h^{-1} \widetilde{s}_n^{(k)}$$

on the nth knot cell relative to the transition (with periodic indexing) therefore gives

$$(P_k x)'(t) = \sum_{j=1}^{m} h_j \overline{K}_{k,w}(t - c_j) \quad \text{for a.e. } t,$$

which is (R). For comparison, define $K_{\mathrm{step}}^{(k)}(u) := h^{-1} s_n^{(k)}$ on the corresponding knot cell.

Thus the first three cells retain the signs $+, -, +$. Since $q < 1/3$ and $\eta < 1$,

$$\left(1 - \frac{\eta^2}{4}\right)(1 - q^3) > \frac{3}{4}\frac{26}{27} > \frac{1}{2},$$

so (F.4) yields $|\widetilde{s}_n^{(k)}| > s_2/2$ for $n = 0, 1, 2$. Rescaling derivatives by $1/h$ gives magnitude at least

$$\frac{1}{2}\frac{s_2}{h} = \frac{33 - 19\sqrt{3}}{8\pi}\, k$$

on the first three knot cells. Taking their interiors as witness intervals proves (O) with the stated r and c_1.

Tail from the same geometric factor. For a.e. u with $d := |u|_{\mathbb{T}} \geq 3h$, the corresponding shortest-arc cell is not one of the two exceptional central cells, so (F.7) gives $|\overline{K}_{k,w}(u)| \leq |K_{\mathrm{step}}^{(k)}(u)|$. If $n \geq 2$ is its cell index along a shortest arc, then (F.3), $n \leq k/2$, and $k \geq 8$ give

$$|s_n^{(k)}| \leq \frac{A(1+q)}{1-q^8} q^{n-1}.$$

Since $n - 1 \geq d/h - 2 \geq d/(4h)$ for $d/h \geq 3$, rescaling by $1/h$ yields

$$|\overline{K}_{k,w}(u)| \leq \frac{(1+q)^2 k}{4\pi(1-q^8)} \exp\left(-\frac{\gamma k d}{8\pi}\right).$$

Using

$$e^{-x} \leq \frac{4}{e^2} x^{-2} \qquad (x > 0)$$

gives

$$|\overline{K}_{k,w}(u)| \leq C_K^{\mathrm{per}}(2) k^{-1} d^{-2},$$

which is $(\mathrm{T})_2$. □

Proof of Corollary 3. Lemma 9 gives the common response, (O), and $(\mathrm{T})_2$. Since $S_2(m) \leq \zeta(2)$, one admissible constant is

$$C_{\mathrm{spl}} := \left(\kappa_2 \zeta(2)\right)^{1/2}, \qquad \kappa_2 := \max\left\{(12\pi)^2, \frac{8 C_K^{\mathrm{per}}(2)}{c_1}\right\}, \qquad c_1 = \frac{33 - 19\sqrt{3}}{8\pi}.$$

The stated width and separation conditions are sufficient for (F.1). Lemma 7 therefore gives $|r_i'| < \beta_i$ on the three witness cells, and Proposition 4 gives a false maximum–minimum pair in each knot-aligned core. The total projected spline has nonzero slopes of alternating sign on those cells, so the two intervening knots are strict local extrema. □

Remark 4 (Positive smoothing versus oscillatory reconstruction). Condition (O) is the structural source of the false extrema; $(\mathrm{T})_\alpha$ and the separation condition only ensure that other transitions do not erase this local oscillation. A nonnegative convolution kernel applied to a positive ramp derivative produces a nonnegative response and therefore cannot satisfy the alternating sign condition (O). Such smoothing can nevertheless change extremal values: Proposition 3 shows that a strictly positive normalized averaging kernel contracts the global extremal values of every nonconstant continuous periodic signal.

F.6 SUPPLEMENTARY COMPARISON PRINCIPLES

The remaining results are independent of the false-extrema chain above. They record three complementary limitations of fixed representations: a finite-rank obstruction to universal apex preservation, non-invariance of raw fixed-basis coordinates under time reparameterization, and possible crossing of compact MRA supports.

The following result is intentionally separated from the main experimental claim. Its role is only to identify the minimum rank required for a *fixed* linear reconstruction to preserve $q+1$ independently variable apex values universally.

Proposition 6 (Finite-rank obstruction to universal apex preservation)**.** *Let $\phi_1, \dots, \phi_{q+1} \in C([0,1])$ be continuous unit–height unimodal witness functions, each with a unique apex p_i in the interior of its support, and with pairwise disjoint support interiors. Hence*

$$\phi_i(p_j) = \delta_{ij}.$$

Define

$$V := \operatorname{span}\{\phi_1, \dots, \phi_{q+1}\}, \qquad f_a := \sum_{i=1}^{q+1} a_i \phi_i, \quad a = (a_1, \dots, a_{q+1}) \in \mathbb{R}^{q+1}.$$

Then $f_a(p_j) = a_j$ for every j. For every *fixed linear reconstruction*

$$L : V \to C([0,1]), \qquad \operatorname{rank}(L) \le q,$$

there exist *$a \in \mathbb{R}^{q+1}$ with $\|a\|_\infty = 1$ and an index $j^\star$ such that*

$$f_a(p_{j^\star}) = 1, \qquad (Lf_a)(p_{j^\star}) = 0.$$

Hence no such L can preserve all $q+1$ apex values for every signal in V; for the witness above, the reconstruction error at the selected apex is exactly

$$\left| f_a(p_{j^\star}) - (Lf_a)(p_{j^\star}) \right| = 1.$$

Proof. Evaluating a linear combination at p_j gives

$$f_a(p_j) = \sum_{i=1}^{q+1} a_i \phi_i(p_j) = a_j.$$

Thus the ϕ_i are linearly independent, $\dim V = q+1$, and the coefficient vector a directly specifies the $q+1$ apex-coordinate values.

Let $E : C([0,1]) \to \mathbb{R}^{q+1}$ be the evaluation map

$$E(g) = \big(g(p_1), \dots, g(p_{q+1})\big),$$

and define the reconstructed apex-readout map

$$R(a) := E(Lf_a).$$

Because R factors through L,

$$\operatorname{rank}(R) \le \operatorname{rank}(L) \le q.$$

If L preserved all $q+1$ apex values for every $f_a \in V$, then $R(a) = a$ for every a, so R would be the identity on $\mathbb{R}^{q+1}$ and would have rank $q+1$, a contradiction. This proves the universal impossibility statement.

For the stronger normalized witness, rank–nullity gives a nonzero $a^\star \in \ker R$. Rescale it so that $\|a^\star\|_\infty = 1$, and negate it if necessary so that for some $j^\star$, $a^\star_{j^\star} = 1$. Then

$$f_{a^\star}(p_{j^\star}) = 1, \qquad (Lf_{a^\star})(p_{j^\star}) = R(a^\star)_{j^\star} = 0.$$

Since $p_{j^\star}$ is the unique interior apex of the corresponding disjoint witness event, it is an extremum of $f_{a^\star}$, and the displayed pointwise apex error equals one. □

Remark 5 (Scope of the finite-rank obstruction). The proposition is purely dimensional. Raising the rank to at least $q + 1$ removes this particular impossibility, although it does not by itself guarantee apex preservation. The result therefore does not imply that arbitrary full-rank smoothers must attenuate or destroy an apex. It also does not cover Spikelet pruning in Theorem 1: the retained extrema are selected from the input and then used as interpolation nodes, so the resulting reconstruction is data-adaptive rather than one fixed rank-q linear operator on V. The proposition is consequently a supplementary comparison principle, not a direct target of the controlled experiments.

Proposition 7 (Fixed-basis coordinates are not warp invariants)**.** *Let $x \in C([0,1])$, let $(e_n)_{n\geq 1}$ be any fixed complete orthonormal basis of $L^2([0,1])$, and define $C(x) = (\langle x, e_n\rangle)_{n\geq 1}$. For any orientation–preserving homeomorphism φ such that $x \circ \varphi \neq x$, one has $C(x \circ \varphi) \neq C(x)$. Thus the full raw coefficient map is not invariant under time reparameterizations that change the signal.*

Proof. Completeness of the orthonormal basis makes the coefficient map injective in L^2. Therefore equality of the full coefficient vectors would imply $x \circ \varphi = x$ in L^2. Since both functions are continuous, equality almost everywhere implies equality everywhere, contradicting the hypothesis. □

This does not assert that every finite coefficient subset changes under every warp, nor does it exclude deformation-stable or invariant features constructed on top of a fixed–basis representation.

Proposition 8 (Crossing supports in a compactly supported orthonormal wavelet basis)**.** *Let ψ generate a compactly supported orthonormal MRA wavelet basis, and let $[A, B]$ be the minimal interval containing its support. If $B - A > 1$ in integer translation units, then two wavelets at adjacent levels have crossing support intervals. Consequently, some signal and coefficient threshold activate a crossing pair.*

Proof. Put $L = B - A > 1$ and choose an integer $\ell \in (B, B + L)$. At any level j, the support intervals of $\psi_{j,0}$ and $\psi_{j+1,\ell}$ are $2^{-j}[A, B]$ and $2^{-(j+1)}[\ell + A, \ell + B]$. The choice of ℓ gives $2A < \ell + A < 2B < \ell + B$, so their endpoints interleave strictly. For $x = \psi_{j,0} + \psi_{j+1,\ell}$, orthonormality makes both coefficients equal to one; any threshold in $(0, 1)$ activates the crossing pair. □

The proposition is not a criticism of wavelet localization. It only shows that compact multiscale support alone does not imply the polarity-specific noncrossing organization induced by the Spikelet construction.

G Experimental protocol

This appendix specifies the experimental protocols behind §7, in the same order as Table 1 and Appendix H: E1–E4 followed by the real-data audits.

Common protocol. Spikelet is computed by the implementation-order construction of Definition 3, which records the one-pass anchor supports of Proposition 2. When completed excursion supports are required, they are obtained by the linear-time completion of Section D.3. Each event is represented as $s = (p, \ell, r, m)$ under the support convention stated for the experiment. Audits of Theorem 1 use only magnitude thresholding and piecewise-linear reconstruction through the retained-apex set $A_\tau(x)$ defined in Section 4. No other task-specific Spikelet operation is used in E1–E4; the real-data audit uses only the magnitude-based selection specified below.

All randomness is seeded, and every controlled module reports PASS/FAIL against its stated condition. The numerical Spikelet implementation uses the direct-value forms of (M) and (L) in Definition 3. These are algebraically equivalent in exact arithmetic to the corresponding adjacent-leg comparisons defined there, but avoid floating-point cancellation and require no numerical tolerance.

Where an ε–coarsened extremum count is reported, adjacent extremum pairs whose amplitude difference is below $\varepsilon = 0.02 \times \mathrm{range}$ are cancelled smallest–first. We write

$$\Delta_\varepsilon = V_\varepsilon(\hat{x}) - V_\varepsilon(x).$$

This is a supplementary robustness diagnostic and is not used to test the exact set-inclusion statement of Theorem 1; in the other reconstruction experiments it serves as a persistence-style discrete

surrogate for the genericity assumption of §3. Unless stated otherwise, reconstruction experiments use the corresponding global coarsened count.

E1: construction, preservation, and support geometry. The E1 generator uses $N = 2048$ samples, a zero baseline, and $m = 8$ signed Gaussian components. Component centers are sampled on the integer grid between $0.05N$ and $0.95N$, with redraws until adjacent sorted centers are separated by at least $0.04N$ samples. All components have width $\sigma = 0.006N$; absolute amplitudes are sampled uniformly from $[0.5, 2.0]$ and assigned random signs. The sampled centers and signed amplitudes are stored as reference metadata. The representative run in the main text uses seed 1.

For the theorem-domain audit, candidate seeds are scanned in ascending order, and eligibility is determined from the generated input before any Spikelet decomposition or proposition/theorem diagnostic. Eligibility requires the sampled signal to satisfy the strict-extremum assumptions of $\mathcal{X}_{\mathrm{d}}$: it is nonconstant and strictly monotone between consecutive endpoints or strict extrema. For the finite sampled signals used here, an exact-zero adjacent step therefore makes a signal ineligible. As noted in Section 3, flat runs and plateau extrema are handled by the representative-point convention in Appendix C; $\mathcal{X}_{\mathrm{d}}$ is the canonical form used for the discrete analysis.

The first 50 eligible inputs are frozen before evaluation, so a failed diagnostic would not be replaced by a later seed. Four candidates are excluded beforehand because floating-point underflow in the Gaussian tails produces exact-zero adjacent steps. The plateau representative convention of Appendix C is not used to admit these signals into the theorem-domain audit.

All 50 frozen runs complete. Proposition 2 has exact left endpoints and anchor containment in $50/50$ runs; Theorem 1 has

$$N_{\mathrm{viol}} = |\mathrm{Ext}(\hat{x}_\tau) \setminus A_\tau(x)| = 0$$

and zero retained-apex value violations in $50/50$ runs; and both completed and one-pass anchor supports have zero same-polarity crossings in $50/50$ runs, auditing Theorem 2 and Proposition 5, respectively. The input-only theorem-domain selection log records rejected candidates and their input-domain reasons before any theorem diagnostic.

For the matched nominal-budget comparison, Spikelet uses the magnitude threshold associated with nominal budget k, while Fourier retains the $2k{+}1$ lowest-frequency coefficients. Because tied magnitudes are retained together, the actual number of retained Spikelet events can differ from k. Thus k is a common nominal operating-point index, not a storage-equivalent budget. The direct Theorem 1 test detects strict internal extrema without amplitude thresholding and requires $N_{\mathrm{viol}} = 0$. The representative comparison additionally reports RMSE,

$$\Delta_{\mathrm{strict}} = V_{\mathrm{strict}}(\hat{x}_\tau) - V_{\mathrm{strict}}(x),$$

where V_{strict} counts raw strict internal extrema without amplitude thresholding, together with the supplementary Δ_ε.

E2: positive smoothing. A smooth nonconstant periodic signal is sampled at $N = 4096$ uniform points over one period,

$$x(t) = 0.75\sin(2\pi t) + 0.28\sin(6\pi t + 0.45) + 0.18\cos(8\pi t - 0.30),$$

and mean-centered. Circular convolution uses normalized wrapped-Gaussian kernels with $\sigma \in \{0.02, 0.03, 0.05, 0.08, 0.12\}$ times the period. For each scale we record the smoothed global maximum and minimum, their inward changes from the original values, the remaining signal-range ratio, and the minimum sampled kernel value. PASS requires both strict inequalities of Proposition 3 to exceed

$$100\,\mathrm{eps}(\max\{1, \max|x|\}),$$

where $\mathrm{eps}(a)$ denotes the double-precision floating-point spacing at scale a. This tolerance only excludes differences at numerical-roundoff scale. No Spikelet decomposition is used in E2.

E3: fixed-resolution reconstruction. E3 uses two deterministic branches with $m \in \{2, 4, 6, 8\}$ monotone transitions of alternating unit height.

The finite-resolution illustration uses a nonperiodic spline grid at level $\ell = 6$ ($k = 64$), with $N = 2048$ samples and interior knot spacing $h = N/k$ in sample units. Transitions are knot aligned with

$w/h = 0.08$, and the common radius–$3h$ illustration cores satisfy $\delta/h > 6$. The number of spline knot intervals and the Fourier harmonic cutoff are both 64; this is a common nominal resolution, not an equal number of stored real values. Spikelet uses a nominal event budget of 64 only as the separate global variation-diminishing control. Spline and Fourier counts are reconstructed minus original extrema inside each disjoint common transition core, summed over cores, and the minimum per-core count is also checked. Spikelet is evaluated by the global count $V(\hat{x}) - V(x) \le 0$, because the step train contains plateaus and a core-local comparison can depend on the chosen plateau representative.

The high-resolution theorem audit is periodic on $[0, 2\pi)$ and uses $k = 512$, $N = 131072$ samples, 256 samples per knot interval, and exactly $kw = 0.08$. The numerical Fourier reconstruction uses the discrete Fourier transform, and the spline reconstruction uses discrete least squares on the sampled periodic hat basis. These are numerical discretizations of the continuous projectors analyzed in Appendix F. Transitions are knot aligned and use the method-specific cores

$$I_i^{\mathrm{F}} = \left[c_i - \frac{5\pi}{2k},\, c_i + \frac{5\pi}{2k}\right], \qquad I_i^{\mathrm{spl}} = \left[c_i - \frac{6\pi}{k},\, c_i + \frac{6\pi}{k}\right],$$

matching Corollaries 2 and 3. For each m, the implementation records $k\delta$, $h_{\max}/h_{\min} = 1$, both sides of the corresponding sufficient condition and their ratio, and the minimum number of strict false extrema observed in any theorem core. The original signal has no strict internal extremum in any such core; strict extrema are counted separately in the Fourier and spline theorem cores. For each sampled core segment z, a strict extremum is counted only when its two adjacent differences have opposite signs and both exceed $10^{-12}\max\{1, \|z\|_\infty\}$ in absolute value. This roundoff filter is distinct from ε–coarsening. Raw and ε–coarsened counts are retained as diagnostics. The numerical audit uses $c_0 = 0.1$ and $k = 512$, which satisfy the resolution requirements of Appendix F for both Fourier and spline reconstruction. This branch audits the sufficient-condition inequalities and their method-specific cores; it is not a matched storage-budget comparison.

E4: robustness beyond the analytically certified regime. E4 progressively relaxes the assumptions underlying the spline analysis. The knot-aligned triangular stage places triangular events on the spline knot grid, matching the setting underlying Corollary 3, but does not audit its quantitative sufficient condition; that is the role of E3. The off-grid triangular stage relaxes knot alignment while retaining the triangular event shape, and the off-grid smooth-Gaussian stage additionally replaces the piecewise-linear event shape by a smooth one. The off-grid stages lie outside the knot-alignment hypothesis of Corollary 3 and are handled by phase averaging as robustness checks.

We report

$$\Delta = V_\varepsilon(\hat{x}) - V_\varepsilon(x), \qquad E_{\mathrm{apex}} = \max_{p \in \mathrm{Ext}(x)} |\hat{x}(p) - x(p)|.$$

The spline resolution and Spikelet pruning level follow the experiment-specific settings used to test robustness beyond the analytically certified regime. E4 is therefore not a matched storage-budget comparison; its purpose is to test whether apex values and extremum counts remain preserved after relaxing the analytically certified assumptions.

Real-data support geometry and support-scale robustness. We use the eye-movement, seal, and cricket recordings corresponding to the real-data examples in Appendix A.1. The eye-movement data originate from Fang & Shinozaki (2018) and were used by Alaee et al. (2021); the seal data originate from Ladds et al. (2016) and were used by Imani & Keogh (2019); and the cricket recording is distributed through UCR time-series resources (Dau et al., 2018). We use the same signals as in the prior Spikelet motif study (Imamura & Nakamura, 2023).

The exact processed copies are not redistributed because redistribution rights for those copies have not been independently verified. The released code accepts authorized local CSV copies, while the supplementary material records the expected dimensions, analyzed columns, and checksums and provides derived reference outputs for the reported real-data analyses.

The complete extremal event family is first computed without task-specific reduction. Following the magnitude-only heuristic used in prior Spikelet anomaly analysis (Imamura, 2025), signed magnitudes are clustered directly by $K = 3$ k-means with 20 replicates, at most 1000 iterations, and a fixed random seed of 2027. The cluster whose centroid has the smallest absolute value is removed,

and the other two clusters are retained as selected negative- and positive-magnitude events. No motif matching, anomaly scoring, constant-segment extraction, local context-dependent reduction, or symbolic mapping is applied.

For the full and selected event families we record the 10th percentile, median, and 90th percentile of support length and the maximum same-polarity hierarchy depth; for the selected family we additionally record the retained-event fraction and the fraction whose one-pass right endpoint changes under exact support completion. Support length is measured in native sample indices and is interpreted within each dataset rather than as a cross-dataset physical-duration measure.

For the support-scale robustness audit, we retain only selected two-sided interior events, for which both first-passage sets in Definition 1 are nonempty. We divide $|m|$ into four equal-count groups and summarize support length by its median and 10th–90th percentiles. The analysis is repeated using the same-height excursion interval derived from the persistence death level; this is an alternative support convention, not a canonical persistent-homology output.

These real-data analyses use only three recordings from prior Spikelet studies and a previously used magnitude-only selection heuristic. They characterize representation geometry rather than downstream task accuracy or cross-dataset generalization; the exact theorems make no claim of universal real-data accuracy.

H Extended experimental results

The E1 matched nominal-budget reconstruction is already shown in the main text as Figure 4 and is not repeated here. Table 3 summarizes the frozen 50-signal theorem-domain preservation and support-geometry audit.

E1 theorem-domain audit	Result
Proposition 2: support characterization	$50/50$ PASS
Theorem 1: preservation	$50/50$ PASS
Theorem 2: completed-support noncrossing	$50/50$ PASS
Proposition 5: anchor-support noncrossing	$50/50$ PASS

Table 3: **E1 theorem-domain numerical audit.** All four structural checks pass on the same frozen set of 50 generated inputs satisfying $\mathcal{X}_{\mathrm{d}}$. Appendix G gives the generator, input-only eligibility rule, rejected candidates, and selection and reproducibility details.

σ/period	max drop	min rise	range ratio	min sampled K_σ
0.02	0.01439	0.04222	0.96735	1.87×10^{-138}
0.03	0.03053	0.08939	0.93085	3.12×10^{-63}
0.05	0.07239	0.20606	0.83943	7.51×10^{-25}
0.08	0.14071	0.35003	0.71702	8.02×10^{-12}
0.12	0.21954	0.42282	0.62960	2.76×10^{-7}

Table 4: E2 Gaussian positive-smoothing check for Proposition 3. The original periodic signal has global maximum 0.76737 and global minimum -0.96684. At every tested nonzero scale the maximum decreases and the minimum increases; all five cases pass the strict numerical inequalities. The final column confirms positivity of the sampled wrapped-Gaussian kernel used in each run.

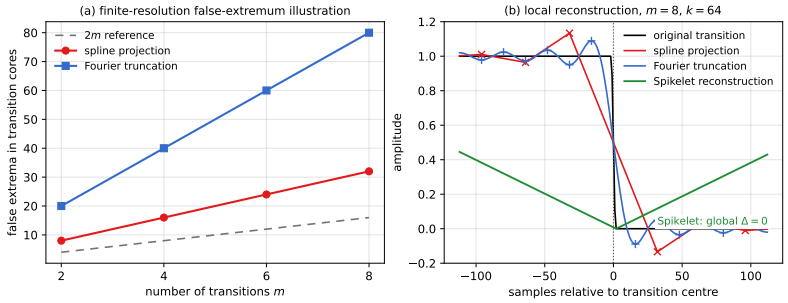

Figure 9: E3 finite–resolution false–extremum illustration and a representative local reconstruction. The $2m$ line is a reference for one local maximum and one local minimum per transition, not a claim that this $k = 64$ branch satisfies the quantitative sufficient conditions of Corollaries 2 and 3. The figure illustrates the E3 failure mode: fixed-resolution reconstruction can introduce local extrema that are absent from the original monotone transition.

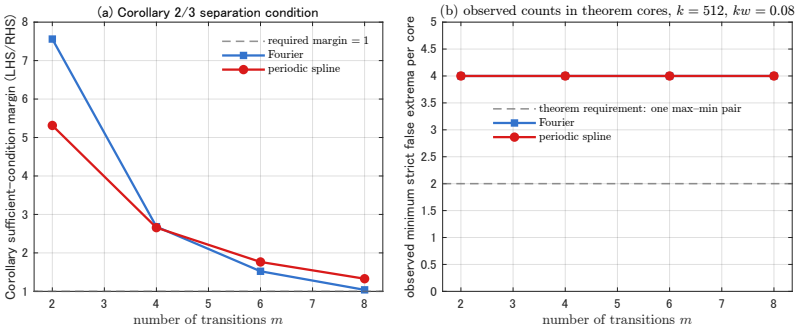

Figure 10: High–resolution E3 sufficient–condition audit. Left: computed condition margins for Fourier truncation and periodic spline projection; the dashed line marks the required margin one. Right: we observe four strict false extrema in every method–specific theorem core, while the theorem guarantees at least one maximum–minimum pair. Thus every tested case both satisfies the method-specific sufficient condition and exhibits the predicted false-extremum effect.

m	$2m$	spline core	Fourier core	Spikelet global Δ
2	4	8	20	0
4	8	16	40	0
6	12	24	60	0
8	16	32	80	0

Table 5: E3 finite–resolution transition–level illustration at $k = 64$. Spline and Fourier counts are core local; Spikelet is the separate global variation–diminishing control. At this illustrative resolution, both fixed-resolution reconstructions create excess local extrema, whereas the Spikelet control does not increase the global extremum count.

m	Fourier margin	spline margin	Fourier observed min/core	spline observed min/core
2	7.560	5.313	4	4
4	2.682	2.657	4	4
6	1.522	1.764	4	4
8	1.039	1.328	4	4

Table 6: High–resolution E3 audit at $k = 512$, $N = 131072$, and $kw = 0.08$. The margin is the left–hand side divided by the right–hand side of the corresponding quantitative sufficient condition in Corollary 2 or Corollary 3; values greater than one pass the numerical inequality. The final columns give the minimum number of strict false extrema observed in any method–specific theorem core. The theorem requires at least two strict false extrema—one maximum–minimum pair—per core; four are observed in every tested core.

Stage	spline projection Δ	spline projection Eapex	Spikelet Δ	Spikelet Eapex
knot-aligned triangular	$6, 14, 30, 32$	0.512–0.515	0	0
off–grid triangular	$6, 14, 20, 25$	0.781–0.785	0	0
off–grid smooth Gaussian	$6, 15, 30, 25$	0.528–0.538	0	0

Table 7: **E4 robustness beyond the analytically certified regime.** Results over $m \in \{2, 4, 8, 12\}$. The three stages progressively relax knot alignment and the piecewise-linear event shape. Spline projection has $\Delta > 0$ and $\mathrm{Eapex} > 0$ in all 12 cases, whereas Spikelet has $\Delta = 0$ and $\mathrm{Eapex} = 0$. Thus the structural contrast observed in E3 persists after relaxing these analytic assumptions.

Real-data support geometry. Figure 11 shows the signed-magnitude distributions and the support-length summaries before and after the magnitude-only $K = 3$ selection. The rows follow the corresponding real-data examples in Appendix A.1: eye movement, seal, and cricket. Table 8 summarizes the retained fractions, support statistics, hierarchy depths, and support-completion rates.

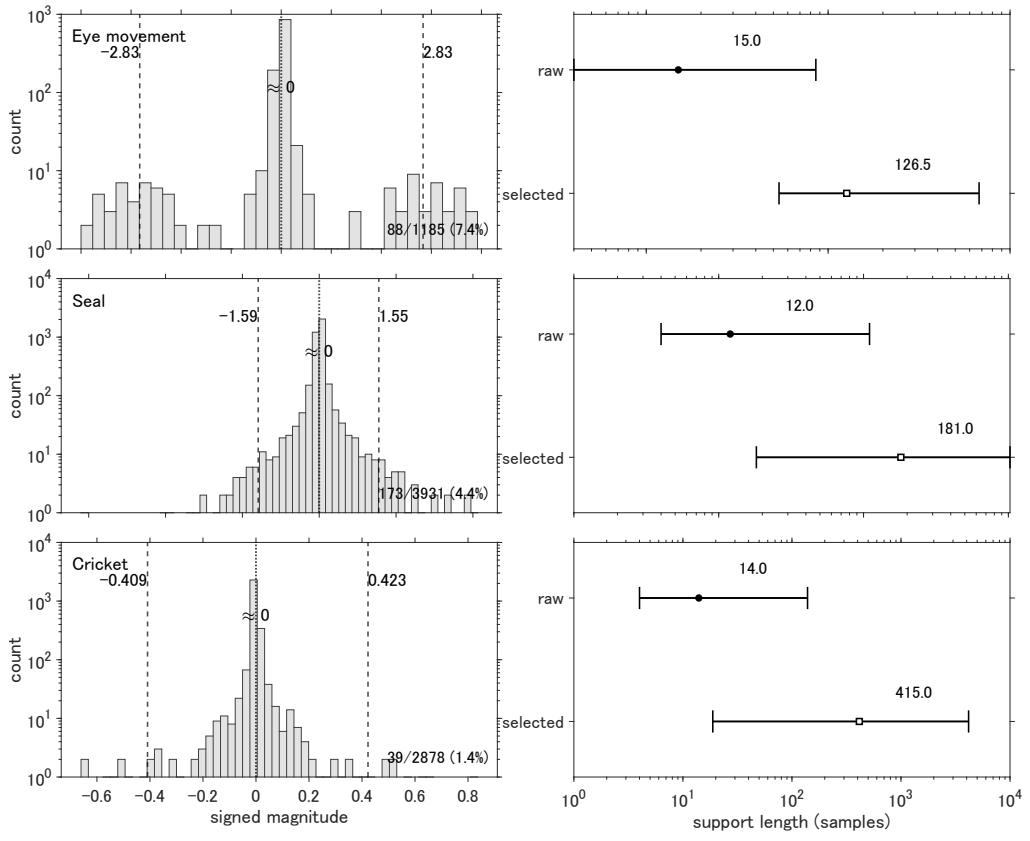

Figure 11: **Real-data support geometry after magnitude-only selection.** Rows show eye-movement, seal, and cricket recordings corresponding to the real-data examples in Appendix A.1. Left: signed event-magnitude distributions. The three vertical lines denote the $K = 3$ k-means centroids; the cluster whose centroid is closest to zero is removed, leaving the selected negative- and positive-magnitude events. The annotation reports the retained-event count and fraction. Right: support lengths of the full and selected event families, summarized by the 10th–90th percentile interval and median on a logarithmic axis. Despite the large reduction in event count, the selected events continue to span widely different temporal extents.

Dataset	retained	full med.	selected med.	max depth	completion
Eye movement	$88/1185$ (7.4%)	15.0	126.5	$20 \to 7$	21.6%
Seal	$173/3931$ (4.4%)	12.0	181.0	$40 \to 12$	20.8%
Cricket	$39/2878$ (1.4%)	14.0	415.0	$39 \to 7$	25.6%

Table 8: **Real-data support-geometry audit.** "Retained" is the number and fraction of events remaining after removing the $K = 3$ cluster whose signed magnitude centroid is closest to zero. Support medians are in native samples. "Max depth" gives the maximum same-polarity support-hierarchy depth before and after selection. "Completion" is the fraction of selected events whose one-pass right support endpoint changes under exact support completion. Thus magnitude-only selection substantially reduces the event count while retaining broad support scales and nontrivial same-polarity hierarchy; the nonzero completion rates also confirm that some selected one-pass right endpoints remain provisional.

Support scale within persistence-lifetime bins. Figure 12 restricts the selected family to two-sided interior events and asks whether the scalar quantity shared with 0D persistence determines temporal extent. Table 9 repeats the comparison under the same-height excursion interval defined at the persistence death level.

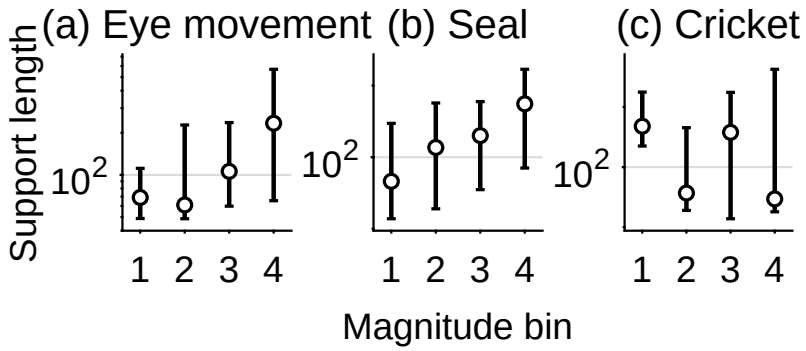

Figure 12: **Temporal support within persistence-lifetime bins.** Selected two-sided interior events are grouped into four equal-count $|m|$ bins, where $|m|$ equals the corresponding 0D persistence lifetime. Vertical lines show the 10th–90th percentile Spikelet support length, and circles show the median. Broad within-bin ranges show that events with comparable persistence lifetimes can have very different support lengths.

Dataset	Spikelet spread	level-cut spread	median I/J	Q90 I/J
Eye movement	2.28–8.69	2.15–9.08	1.10	1.93
Seal	16.97–30.04	21.88–40.81	1.06	3.84
Cricket	7.89–235.22	3.74–144.50	1.04	5.06

Table 9: **Robustness to support convention.** For each support convention, "spread" gives the range across the four equal-count magnitude bins of the within-bin support-length ratio Q90/Q10. I denotes the Spikelet support span and J the same-height excursion span at the persistence death level. Temporal-scale heterogeneity remains under both support conventions. The two supports are typically similar, but can differ substantially for asymmetric events.