Encoding and Decoding Temporal Signals with Spiking Bandpass Wavelets

Jens E. Pedersen, Tony Lindeberg, and Peter Gerstoft

arXiv:2605.09770v2 [cs.NE] 2 Oct 2026

***Abstract*—Spike-based encodings are sparse and energy-efficient. However, previous theory for such spiking representations is largely disconnected from most signal processing literature and does not extend to multi-scale frames. We recast spike-encoders as time-causal overcomplete wavelet frames with closed-form bandwidth tiling and reconstruction error bounds. We show that scale covariance, frame bounds, and reconstruction guarantees hold for any finite-energy input in both encoding and decoding steps. The proposed spiking bandpass wavelets preserve the sparsity and locality of spiking representations: A threshold parameter sets the event rate, and the reconstruction error grows linearly with it, up to spike quantization and time discretization. On ECG and audio datasets, we reach reconstruction accuracy comparable to continuous wavelet transforms, multi-channel integrate-and-fire time encoding machines, and Sigma-Delta modulation. The experiments demonstrate robustness of the proposed spiking bandpass wavelets to noise and spike-time jitter, and that the spike activity shifts across channels as predicted by the covariance property under temporal rescaling. Every operation reduces to a leaky integrate-and-fire difference with thresholding, and both the encoder and the decoder map directly to existing neuromorphic hardware.**

***Index Terms*—Spiking neural networks, wavelet frames, time-causal scale space, event-based encoding, sparse coding, neuromorphic signal processing.**

I. Introduction

Analog-to-digital conversion (ADC) underpins every digital signal processing pipeline. Uniform Nyquist-Shannon sampling [1, Thm. 1] is a well-understood approach, but it is agnostic to signal structure: it indiscriminately samples both silence and activity.

Biological sensory systems sample differently. Rather than measuring at a fixed rate, spikes are emitted only when the input changes, producing sparse, event-driven, and highly energy-efficient representations [2]. However, their energy advantage over dense activations is not automatic, but depends on the event rate and hardware [3]. A leaky integrate-and-fire (LIF) neuron implements a time-causal convolution (19) followed by a threshold comparator, encoding change as spikes rather than amplitude as samples. Neuromorphic hardware strives to exploit this efficiency, but a principled signal-processing framework for spike-based encoding is missing.

The encoded representations should be (1) localized in frequency, (2) scale covariant, to enable structurally similar processing over different temporal scales, and (3) equipped with formal bounds that guarantee stable reconstruction regardless of the input. Wavelet filter banks satisfy all three properties by construction [4] and are the natural mathematical home for a principled spike-based ADC. A spiking analogue is missing: a wavelet frame whose coefficients are emitted as spike events rather than continuous amplitudes. Such a frame would inherit the covariant, multi-resolution structure of classical wavelet theory, while producing sparse, event-driven output.

Neuroscience-driven models are typically presented with information-theoretic [5] or probabilistic [6] foundations rather than signal processing guarantees. Recent work explores the integrate-and-fire model from a signal processing perspective [7], but does not extend to multi-scale frames. Sigma-delta ($\Sigma\Delta$) modulators [8] offer recovery guarantees but are single-channel or restricted to shift-invariant subspaces. Time encoding machines [9] operate on fixed shift-invariant subspaces rather than scale-parameterized families. Quantization in overcomplete frame expansions is well studied in its own right [10], but for fixed frames rather than for a covariant, scale-parameterized family generated by a spiking encoder. Prior work [11], [12] established covariance guarantees for the LIF response and sketched a spiking wavelet construction. We expand this with two time-causal wavelet families realized as spiking analysis and synthesis operators: difference-of-truncated-exponentials (DoE), built from a single leaky integrator per scale, and difference-of-time-causal-limit-kernels (DoT), using cascades of leaky integrators (Fig. 1). The spiking wavelets rely on bandpass wavelets [13], [14]. But where spiking autoencoders learn their encodings from data [15]–[17], both families here exist in closed form, so their bandwidth and reconstruction bounds follow analytically, and we can derive reconstruction guarantees, error bounds, and quantitative evaluation.

Our contributions are threefold:

(i) Spiking wavelets: two spiking time-causal scale-covariant wavelet families, based on difference-of-truncated-exponentials (DoE) (29) and difference-of-time-causal-limit-kernels (DoT) (28).

(ii) Guarantees: a closed-form reconstruction error bound (47) for the causal spiking wavelets: unconditionally for the DoT, and under mild assumptions on the signal spectrum for the DoE. This extends Time Encoding Machines [9], [18], [19] to overcomplete multi-scale frames.

(iii) Experiments: reconstruction on ECG signals (MIT-BIH [20]) and audio (LibriSpeech [21]) reaching reconstruction RMSE comparable to classical and continuous wavelets, and to the MC-IFTEM and $\Sigma\Delta$ event encoders at a matched spike rate. Additionally, the DoE wavelets export to existing neuromorphic

J. E. Pedersen and P. Gerstoft are with the Department of Electrical and Photonics Engineering, Technical University of Denmark, Denmark (e-mail: jegpe@dtu.dk; pegers@dtu.dk).

T. Lindeberg is with the Department of Computational Science and Technology, KTH Royal Institute of Technology, Sweden (e-mail: tony@kth.se).

Support from the Novo Nordisk Foundation (NNF24OC0089302) and the Swedish Research Council (2022-02969) is gratefully acknowledged.

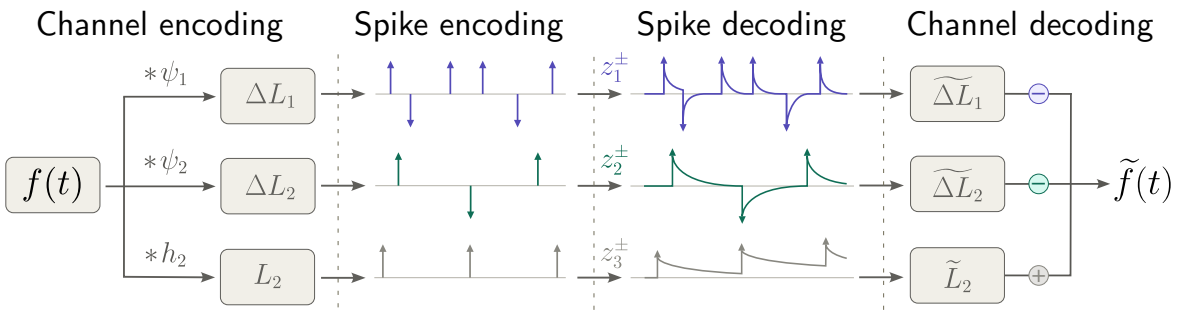

Fig. 1: Proposed Algorithm 1, demonstrated with three channels. A signal $f(t)$ is encoded into three channels: two bandpass representations ΔL by wavelets ψ (21) and a lowpass representation L with smoothing kernel h (1). The channels are then encoded as sparse polarity spikes $z_k^{\pm}(t)$ (17). To recover an approximation of the original signal $f(t)$, the three channels are recreated from the spikes (36) and combined to form $\widetilde{f}(t)$ (41).

backends via the Neuromorphic Intermediate Representation [22], enabling hardware-backed execution in the future.

II. Background

This section covers scale-space theory, including time-causal scale-space representations, wavelets, wavelet frames, and spiking neuron models.

A. Scale spaces

Scale-space theory parameterizes a signal $f\colon \mathbb{R} \to \mathbb{R}$ over a scale parameter $\sigma \in \mathbb{R}_+$ into the scale-space representation $L\colon \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$ [23]. Computing the scale-space representation L corresponds to convolving the signal $f(t)$ with smoothing kernels $h(t;\ \sigma)$ [24, (1.2)]

$$L(t;\ \sigma) = h(t;\ \sigma) * f(t) = \int_{\xi=-\infty}^{\infty} h(\xi;\ \sigma)\, f(t-\xi)\, \mathrm{d}\xi, \tag{1}$$

where "$*$" denotes convolution, with the zero-scale condition

$$L(t;\ \sigma=0) = f(t) \qquad \text{thus} \quad h(t;\ \sigma=0) = \delta(t). \tag{2}$$

Scale-space kernels h are required to diminish variations over scales and to form a continuous semi-group under convolution, leaving the Gaussian with mean 0 and standard deviation σ as the unique admissible kernel over spatial domains

$$h_{\text{Gauss}}(t;\ \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{t^2}{2\sigma^2}}. \tag{3}$$

The Gaussian kernel is physically interpreted as the Green's function of the heat equation, over the variance σ^2 [23, p. 370]. Scale-space representations are scale covariant [25], or scale equivariant. Consider the scaling transformation $t' = s\,t$ and the corresponding scale parameter transformation $\sigma' = s\,\sigma$ for any scaling factor $s > 0$. A scale-covariant representation satisfies

$$L'(t';\ \sigma') = L(t;\ \sigma), \tag{4}$$

meaning that the scale-space representation L' over the transformed domain captures the same structural information as L, but observed at a different scale.

Representing a signal at multiple scales requires a set of scale levels σ_k, $k \in \mathbb{N}$, between some minimum scale σ_0 and maximum scale σ_K. A logarithmic spread of scales retain self-similarity [25] and describe temporal relations in memory [26]. Denote the ratio between adjacent scale levels as c, with the default choices $c \in \{\sqrt{2}, 2\}$:

$$\sigma_k = c\,\sigma_{k-1}. \tag{5}$$

We write $L(t;\ \sigma_k, c)$ for the scale-space representation (1) at the k-th level of this grid, c denoting the distance to neighbouring scales.

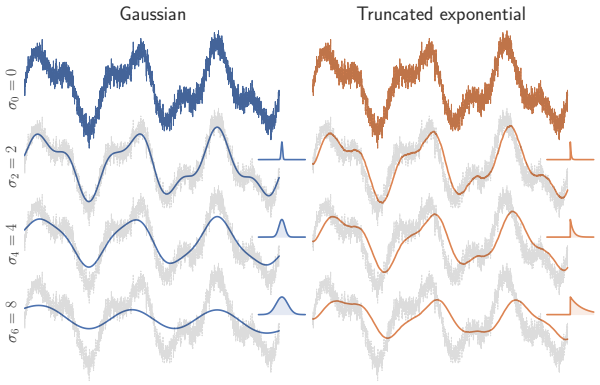

Fig. 2: Scaled representations of a noisy signal smoothed by Gaussian (3) and truncated exponential (6) kernels, shown in the small insets. The original signal (top row) at scale $\sigma = 0$ (2) smoothens as σ or μ increases.

1) Time-causal scale spaces: For time-causal operations, a continuous semi-group structure is not possible: the causality constraint requires one-sided kernels which cannot generate a continuous semi-group of smoothing kernels while retaining variation-diminishing properties [14], [25]. The scale parameter must be discretized with the unique variation-diminishing kernel with one-sided support, the *truncated exponential* [27]

$$h_{\exp}(t;\ \mu) = \begin{cases} \mu^{-1} \exp\left(-t/\mu\right) & t > 0 \\ 0 & t \leqslant 0, \end{cases} \tag{6}$$

where μ is the time constant and also the delay. Time-causal scale spaces are built by cascading K such kernels with geometrically spaced time constants as (5).

Fig. 2 illustrates the scale representation at varying temporal scales using the Gaussian (3) and truncated exponential (6) kernels. The causal representation thus has temporal delays that increase with the temporal scale, see Fig. 2 right.

B. Wavelets and frames

Wavelets decompose a signal $f(t)$ into scale and shift components using the wavelet generating function

$$\psi(t; a, b) = |a|^{-1/2}\, \psi\left((t-b)/a\right) \tag{7}$$

with scale a and temporal shift b. The wavelet transform is defined as

$$\begin{aligned}\big(T(f)\big)(t; a, b) &= \langle f(t), \psi(t; a, b)\rangle \\ &= \int_{-\infty}^{\infty} f(t)\, |a|^{-1/2}\, \overline{\psi\left((t-b)/a\right)}\, \mathrm{d}t,\end{aligned} \tag{8}$$

where $\langle \cdot, \cdot \rangle$ denotes the inner L^2 product, and $\overline{\psi}$ denotes the complex conjugate of ψ [4, p. 24].

For ψ to serve as a mother wavelet, it must satisfy the admissibility criterion [4, (2.4.1)]

$$C_\psi = \int_0^\infty \frac{|\hat{\psi}(\omega)|^2}{\omega}\, \mathrm{d}\omega < \infty, \tag{9}$$

where $\hat{\psi}$ is the Fourier transform of ψ, which implies the zero-mean condition [4, (1.2.3)]

$$\int_{-\infty}^{\infty} \psi(t)\, \mathrm{d}t = 0. \tag{10}$$

1) Frame operator: Frames are possibly overcomplete, continuous-time expansions. A family of wavelets ψ_k, indexed by $k \in K$ scales, is a frame if there exist bounds A and B, where $0 < A \leqslant B < \infty$, such that for all f in a Hilbert space [4, (3.1.2)]

$$A\, \|f\|^2 \leqslant \sum_{k \in K} |\langle f, \psi_k\rangle|^2 \leqslant B\, \|f\|^2. \tag{11}$$

When $A = B$, the frame is tight, providing uniform representations across scales. When $A < B$, the frame is overcomplete, offering robustness, but at the cost of redundant parameters. For a filterbank of shift-invariant channels, the bounds are given by the energy capture function [28, Thm. 5.11]

$$S(\omega) = \sum_{k \in K} |\hat{\psi}_k(\omega)|^2, \tag{12}$$

with $A = \inf_\omega S(\omega)$ and $B = \sup_\omega S(\omega)$. A frame operator Λ transforms a continuous-time signal from $L_2(\mathbb{R}) \to \ell_2(K)$ as $\Lambda(f(t))_{k \in K} = \langle f(t), \psi_k\rangle$. Since Λ is bounded, the Hermitian adjoint is well defined $\Lambda^* : \ell_2(K) \to L_2(\mathbb{R})$

$$\Lambda^*\big(\Lambda(f(t))\big) = \sum_k \langle f(t), \psi_k\rangle\, \psi_k. \tag{13}$$

This we use to define a dual family of vectors $\tilde{\psi}_k \in L_2(\mathbb{R})$

$$\tilde{\psi}_k = (\Lambda^*\Lambda)^{-1}\, \psi_k. \tag{14}$$

$\tilde{\psi}_k$ is itself a frame with bounds B^{-1} and A^{-1} (11) [4, (3.2.6)], permitting signal reconstruction [4, (3.2.8)]

$$f(t) = \sum_{k \in K} \langle f, \psi_k\rangle\, \tilde{\psi}_k. \tag{15}$$

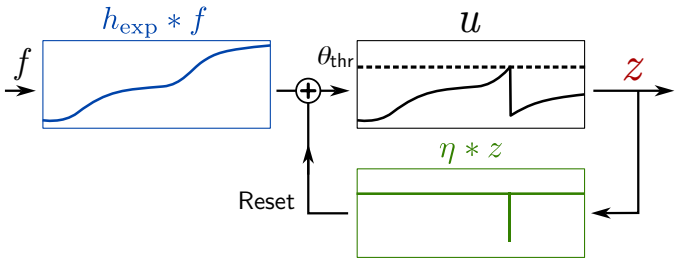

Fig. 3: The leaky integrate-and-fire (LIF) model (19) as a composition of linear systems. An integrator $h_{\exp}$ convolves a signal f, captured by a voltage membrane u that outputs spikes (z) when $u \geqslant \theta_{\text{thr}}$.

C. Spiking neuron models

Physical neurons implement $h_{\exp}$ (6) directly [6, (1.6)]. The membrane voltage u of a leaky integrator dissipates according to the RC circuit derived from Kirchhoff's law, $C\dot{u} = -(u - u_{\text{rest}})/R + I(t)$, where C is the membrane capacitance, R the membrane resistance, u_{rest} the resting potential, and $I(t)$ the input current in amperes. Measuring u relative to u_{rest}, scaling the input current to units of voltage as $f(t) = RI(t)$, and identifying the membrane time constant $\mu = RC$ with the scale parameter of (6) gives

$$\mu\, \dot{u} = -u + f(t). \tag{16}$$

The impulse response of (16) is exactly $h_{\exp}(t; \mu)$, so for arbitrary input $f(t)$ the membrane voltage implements the time-causal scale-space convolution [11, (22)] $u(t; \mu) = (h_{\exp}(\cdot; \mu) * f)(t)$. In biological neurons, the continuous membrane voltage u is communicated to downstream neurons via discrete spike events. A spike is emitted when u hits a threshold θ_{thr}, after which u resets. The spike train $z(t)$ is a sum of Dirac impulses at firing time $t_{\text{spike}} \in \mathcal{S}$, where $\mathcal{S}$ is the set of all firing times [6, (1.14)]

$$z(t) = \sum_{t_{\text{spike}}} \delta(t - t_{\text{spike}}), \tag{17}$$

after which the membrane resets by a reset kernel η

$$u(t) = (\eta * z)(t). \tag{18}$$

Combining the leaky integrator (16), the spike function (17) (left), and the reset term (18), we arrive at the LIF model (Fig. 3) [6, (6.28)]

$$\begin{aligned}u(t) &= (\eta * z)(t) + (h_{\exp} * f)(t) \\ &= \underbrace{-\theta_{\text{thr}}\, z(t)}_{\text{Reset integrand}} + \underbrace{\int_{\xi=0}^{\infty} f(t-\xi)\, \frac{1}{\mu} e^{-\xi/\mu}\, \mathrm{d}\xi}_{\text{Input integrand}}.\end{aligned} \tag{19}$$

Defining the reset kernel $\eta(t) = -\theta_{\text{thr}}\, \delta(t)$ feeds each spike back into the membrane, closing the loop in Fig. 3. Setting the smoothing kernel from (1) $h = h_{\exp}$, the membrane voltage is the time-causal scale-space representation, $u(t; \mu) = L(t; \mu)$, so the construct is scale covariant [11, (24)]. Under temporal scaling operations $t' = s\, t$ and $\mu' = s\, \mu$ with $f'(t') = f(t)$, the

substitution $\xi' = s\,\xi$ in (19) leaves the reset and input terms unchanged, so

$$L'(t';\ \mu') = L(t;\ \mu). \tag{20}$$

Table I summarizes the symbols used in this paper along with their meaning and their initial appearance.

Symbol	Meaning	Appearance
$\hat{}$	Fourier transform	(9)
$\bar{}$	Channel normalization ($\bar{x}_k$, $\bar{\psi}$); average ($\bar{r}$, $\bar{G}$)	(24) (54)
$\widetilde{}$	Reconstruction ($\widetilde{f}$, $\widetilde{x}_k$); dual ($\tilde{\psi}_k$)	(41) (14)
A, B	Lower and upper frame bounds	(11)
c	Ratio between adjacent scale levels	(5)
$f(t)$	Real-valued signal at time t	(1)
$\widetilde{f}(t)$	Reconstructed signal at time t	(41)
h_{Gauss}	Gaussian	(3)
h_{exp}	Truncated exponential	(6)
h_Ψ	Time-causal limit kernel	(27)
K	Number of bandpass channels. Scale levels run σ_0 (finest) to σ_K (coarsest)	(5)
N	Number of cascaded truncated exponentials in h_Ψ	(27)
$\lvert f\rvert$	Number of samples in $f(t)$	(54)
p_i	Polarity of spike i	(38)
R_k	Spike-synthesis kernel of channel k	(36)
$S(\omega)$	Energy capture function	(12)
θ_{thr}	Firing threshold of the LIF quantizer	(19)
Δt	Simulation time step	(49)
Ω_k	Essential bandwidth of channel k. The channel is not band-limited, but all of its spectral mass beyond $\lvert\omega\rvert > \Omega_k$ integrates to at most Δ_Ω	(42)
$L(t; \sigma_k, c)$	Scale-space representation at time t, at the k-th level of the c-spaced grid	(1)
μ	Time constant, temporal scale resolution	(6) (27)
Λ	Frame operator	(13)
σ	Standard deviation, scale resolution	(1) (5)
$\mathcal{S}$	The set of firing times in a spike train $z(t)$	(17)
ψ_k	Bandpass wavelet of the channel at the k-th scale	(21)
$(\psi_k)_{k\in K}$	Wavelet dictionary with K scales	(8) (22)
$z(t)$	A sum of Dirac impulses at firing times $\mathcal{S}$	(17)
MC-IFTEM	Multi-channel integrate-and-fire time-encoding machine (baseline)	(51)
$\Sigma\Delta$	Sigma-delta modulator (baseline)	(53)

TABLE I: Symbols and their meanings used throughout the paper with references to their first appearances.

III. Scale-covariant bandpass representations

Sensory and physical signals carry features over multiple temporal scales. To account for all scales in a spiking, or quantized, encoder, the representation must decompose into discrete scale channels, be covariant under temporal rescaling, respecting the self-similarity of multi-scale signals, and remove constant components so the channels admit wavelet frames; the latter making the frame bounds [14] possible. Following the bandpass representation based on a Laplacian pyramid [29], the notion of scale-space representation [23], the multi-resolution decomposition of wavelet frames [30], and specifically the time-causal bandpass representation in [13], we obtain all three by taking adjacent differences of scale-space representations: L (1) is covariant by construction, and the difference channels ensure that the responses vanish at $\omega = 0$. Taking adjacent differences at the geometric scale spacing (5) yields a covariant representation,

$$\begin{aligned}\Delta L(t;\ \sigma_k, c) &= L(t;\ \sigma_k, c) - L(t;\ \sigma_{k-1}, c)\\ &= \psi(t;\ \sigma_k, c) * f(t),\end{aligned} \tag{21}$$

where ψ is a bandpass kernel. In the Gaussian case (3), differences of adjacent scale levels approximate the Laplacian-of-Gaussian, the prototype bandpass filter [31].

An exact reconstruction of a signal at scale σ_k, is achieved by summing the representations along with a lowpass residual at the coarsest scale σ_K, seen by inserting (21) into (1), [13, (102)]

$$\begin{aligned}L(t;\ \sigma_k, c) &= L(t;\ \sigma_K, c) - \sum_{j=k+1}^{K} \Delta L(t;\ \sigma_j, c)\\ &= h(t;\ \sigma_K) * f(t) - \sum_{j=k+1}^{K} \psi(t;\ \sigma_j, c) * f(t).\end{aligned} \tag{22}$$

Setting $\sigma_0 = 0$ gives $L(t;0) = f(t)$ (2), and we can reconstruct a wavelet frame (15) based on (22) by summing over subsequent scales: [13, (103)]:

$$\tilde{f}(t) = L(t;0,c) = h(t;\ \sigma_K) * f(t) - \sum_{k=1}^{K} \psi(t;\ \sigma_k, c) * f(t). \tag{23}$$

K is the number of bandpass channels and the full filterbank has $K+1$ channels (K bandpass and one lowpass residual at scale σ_K). We construct an overcomplete frame, where the lowpass residual corresponds to the scaling function in a finite-level multi-resolution analysis [30, (9)].

A. Channel normalization

Write x_k for the response of channel k (the bandpass response $\Delta L(t;\ \sigma_k, c)$ (21) for $k = 1, \ldots, K$, and x_{lowpass} for the lowpass residual at scale σ_K). The channel norm $\|\psi(t;\sigma_k, c)\|_2$ varies with the scale σ_k and with the spacing c (5), so we normalize each channel by its own norm

$$\bar{x}_k(t) = x_k(t) / \|\psi(t;\ \sigma_k, c)\|_2, \tag{24}$$

as in the L^2 convention of the continuous wavelet transform, and likewise $\bar{x}_{\text{lowpass}} = x_{\text{lowpass}} / \|h(t;\sigma_K)\|_2$ for the lowpass residual. The norm is a scalar, so this equals normalizing the kernel, $\bar{\psi} = \psi / \|\psi(t;\ \sigma_k, c)\|_2$ with $\bar{x}_k = \bar{\psi} * f$. By Plancherel's theorem, the same scalar normalizes the channel in both the time and frequency domain

$$\|\bar{\psi}(t;\ \sigma_k, c)\|_2^2 = \frac{1}{2\pi} \int |\hat{\bar{\psi}}(\omega;\ \sigma_k, c)|^2 \,\mathrm{d}\omega = 1, \tag{25}$$

equalizing the energy of the channels, but not their peak gain.

B. Bandpass wavelets

We rely on three scale- and shift-parameterized bandpass kernels from [13], [14] as wavelets.

1) Difference-of-Gaussians (DoG): The DoG wavelet

$$\psi_{\mathrm{DoG}}(t;\ \sigma_k, c) = h_{\mathrm{Gauss}}(t;\ \sigma_k) - h_{\mathrm{Gauss}}(t;\ \sigma_{k-1}), \quad (26)$$

has zero mean and is admissible according to (9).

2) Difference-of-time-causal-limit-kernels (DoT): By convolving $h_{\exp}$ N times, we construct a time-causal limit kernel with mean and standard deviation [25, (8, 10)]

$$h_\Psi(t;\ \sigma, c) = *_{n=1}^{N}\ h_{\exp}(t;\ \mu_n),$$
$$\mathrm{mean}[h_\Psi] = \sum_{n=1}^{N} \mu_n, \quad \mathrm{std}[h_\Psi] = \sigma = \sqrt{\sum_{n=1}^{N} \mu_n^2}. \quad (27)$$

When $N \to \infty$, we define bandpass filters at a specific scale k

$$\psi_{\mathrm{DoT}}(t;\ \sigma_k, c) = h_\Psi(t;\ \sigma_k, c) - h_\Psi(t;\ \sigma_{k-1}, c). \quad (28)$$

3) Difference-of-truncated-exponentials (DoE): For $N = 1$,

$$\psi_{\mathrm{DoE}}(t;\ \mu_k, c) = h_{\exp}(t;\ \mu_k) - h_{\exp}(t;\ \mu_{k-1}). \quad (29)$$

Signal reconstruction for both the DoE and DoT wavelets follows from (22). Both the DoT and DoE are wavelets, because they satisfy the admissibility criterion (9). The DoG is non-causal and can "react" backwards in time, while the DoT and DoE will necessarily lag behind. See Fig. 4 for an overview of time, frequency, and spectral decay for the three wavelet families.

C. Spectral properties and frame geometry

Consider the frequency-to-bandwidth ratio (or quality factor) of a bandpass filter

$$Q = \omega_{\mathrm{peak}}/\mathrm{BW} = \omega_{\mathrm{peak}}/(\omega_+ - \omega_-), \quad (30)$$

where $\omega_\pm$ are the upper and lower -3 dB band edges. The geometric spacing in (5) implies that each channel is a temporal dilation of the previous by c, $\psi_{k+1}(t) = c^{-1}\psi_k(t/c)$. In turn, the tiling is constant-Q and spans a frequency range with ratio c^{K-1} [14, Sec. 4]. The spectral decay is [14, Tab. 1]

$$\mathrm{DoG} : e^{-\sigma^2\omega^2/2} \quad \mathrm{DoT} : O(\omega^{-N}) \quad \mathrm{DoE} : O(\omega^{-1}). \quad (31)$$

The Gram matrix $\boldsymbol{G} \in \mathbb{R}^{K\times K}$ has the entries

$$G_{jk} = \langle \psi_j, \psi_k \rangle = \frac{1}{2\pi} \int_{-\infty}^{\infty} \widehat{\psi}_j(\omega)\, \overline{\widehat{\psi}_k(\omega)}\, \mathrm{d}\omega, \quad (32)$$

and the normalized Gram matrix is

$$G_{\mathrm{norm},jk} = G_{jk}/\sqrt{G_{jj}\, G_{kk}}, \quad (33)$$

with $G_{jj} = \langle \psi_j, \psi_j \rangle = \|\psi_j\|^2$ [14, Sec. 5]. All three families form overcomplete, non-tight frames [14]. For the $K+1$ channels here, the energy capture function (12) instantiates as $S(\omega) = |\widehat{h}(\omega;\ \sigma_K)|^2 + \sum_{k=1}^{K} |\widehat{\psi}_k(\omega;\ \sigma_k, c)|^2$, where the lowpass term keeps A from vanishing at $\omega = 0$, and the bounds are taken over the tiled band $\omega \in [0,\ 1/\sigma_0)$. The frame bound ratio B/A governs the sensitivity to noise and quantization, with the values [14, Tab. 1]

$$\mathrm{DoG} : 8^* \quad \mathrm{DoT} : 2.38 \quad \mathrm{DoE} : 11.63, \quad (34)$$

at $c = \sqrt{2}$ and where * is a numerically evaluated quantity since the DoG bounds have no closed form.

IV. Spiking scale-covariant wavelets

We define spiking wavelets by quantizing bandpass wavelet responses as streams of signed spikes, whose dual can be approximated to recover the original signal up to a bounded error that we find to be linear in the spiking threshold.

A. Spiking polarity channels

We split each channel into two polarity channels to capture sign information. Let μ_m denote the membrane time constant, distinct from the channel scale μ_k since the scale selectivity resides in ψ_k (we choose $\mu_m = \mathrm{d}t$ in our simulations). The integrative terms based on the leaky integrator (16) become

$$\mu_m \dot{u}_k^+ = -u_k^+(t) + \bar{x}_k(t), \quad \mu_m \dot{u}_k^- = -u_k^-(t) - \bar{x}_k(t), \quad (35)$$

and likewise for $\bar{x}_{\mathrm{lowpass}}$ on the coarsest channel. When $u_k^\pm$ crosses θ_{thr}, the neuron emits a spike and resets (19), acting as a quantizer with two spike trains (17) $z_k^+(t)$ and $z_k^-(t)$ for the positive and negative polarity channels, respectively.

B. Reconstructing from polarity spikes

To reconstruct the signal, we first reconstruct each channel independently, then sum the channels to obtain the full reconstruction via the bandpass decomposition (22). Each bandpass channel with scale σ_k is encoded as the signed spike difference $z_k(t) = z_k^+(t) - z_k^-(t)$, which is passed through a leaky integrator $h_{\exp}(t;\mu_r)$ (6) with reconstruction time constant $\mu_r = \mu_k$. The per-channel reconstruction kernel becomes

$$R_k(t;\ \sigma_k, c) = h_{\exp}(t;\ \mu_r) * \psi(t;\ \sigma_k, c), \quad (36)$$

where ψ is either ψ_{DoG} (26), ψ_{DoT} (28) or ψ_{DoE} (29). This yields the following reconstruction, where h is either the Gaussian (3), h_Ψ (27), or $h_{\exp}$ (6), and z_{lowpass} is the spike train of the low-pass filter (1) at scale K:

$$\tilde{f}(t) = (h_{\exp}(t;\ \mu_r) * h(t;\ \sigma_K, c)) * z_{\mathrm{lowpass}}(t) - \sum_{k=1}^{K} R_k(t;\ \sigma_k, c) * z_k(t). \quad (37)$$

In Equation (37), each spike contributes a uniform amplitude from the threshold θ_{thr} (17).

To restore the amplitude and absorb any discretization corrections, we solve for a per-spike weight by least squares. For each channel, we build a matrix whose columns are the channel's own bandpass reconstruction kernel $R_k = h_{\exp}(\mu_r) * \psi_k$ from (36), shifted to each of channel k's spike times $\mathcal{S}_k$,

$$A_{t,i}^{(k)} = p_i\, R_k\big(t - t_{\mathrm{spike}}^{(i)};\ \sigma_k, c\big), \qquad i \in \mathcal{S}_k, \quad (38)$$

where $p_i \in \{+1, -1\}$ is the polarity of spike i, and solve

$$\boldsymbol{w}^{(k)*} = \arg\min_{\boldsymbol{w}} \big\| x_k - \boldsymbol{A}^{(k)}\boldsymbol{w} \big\|_2^2 \ = \ \boldsymbol{A}^{(k)\dagger} x_k,$$
$$\widetilde{x}_k = \boldsymbol{A}^{(k)}\boldsymbol{w}^{(k)*}. \quad (39)$$

The solution is regularized. Defining $\boldsymbol{G}^{(k)} = \boldsymbol{A}^{(k)\top}\boldsymbol{A}^{(k)}$ and $\bar{G}^{(k)}$ as the mean of its diagonal, we solve the normal equations

$$\boldsymbol{w}^{(k)*} = \big(\boldsymbol{G}^{(k)} + \boldsymbol{\Gamma}^{(k)}\big)^{-1} \boldsymbol{A}^{(k)\top} x_k$$
$$\Gamma_{ii}^{(k)} = \rho\, \max\big(G_{ii}^{(k)}, \gamma\, \bar{G}^{(k)}\big), \quad (40)$$

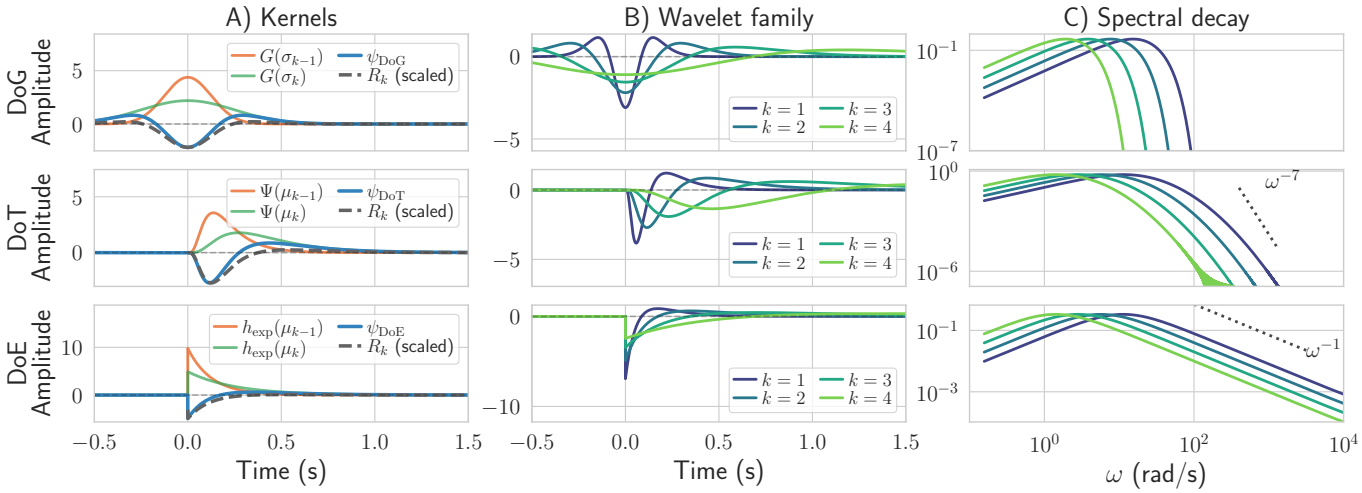

Fig. 4: The three wavelet families at $c = 2$ over $K = 4$ bandpass channels. **A)** Each ψ (26), (28) ($N = 7$), (29) as the difference of two adjacent smoothing kernels, normalized by the single scalar $\|\psi\|_2$ (24), with the spike-synthesis kernel R_k (36) for the same channel, matched to the peak of ψ, which it follows up to the delay. Broadening is a result of the membrane $h_{\exp}(\mu_r)$. **B)** The normalized channels $\bar{\psi}_k$ (24) of each family. **C)** The same channels in the frequency domain with the asymptotes of (31) as dotted guides.

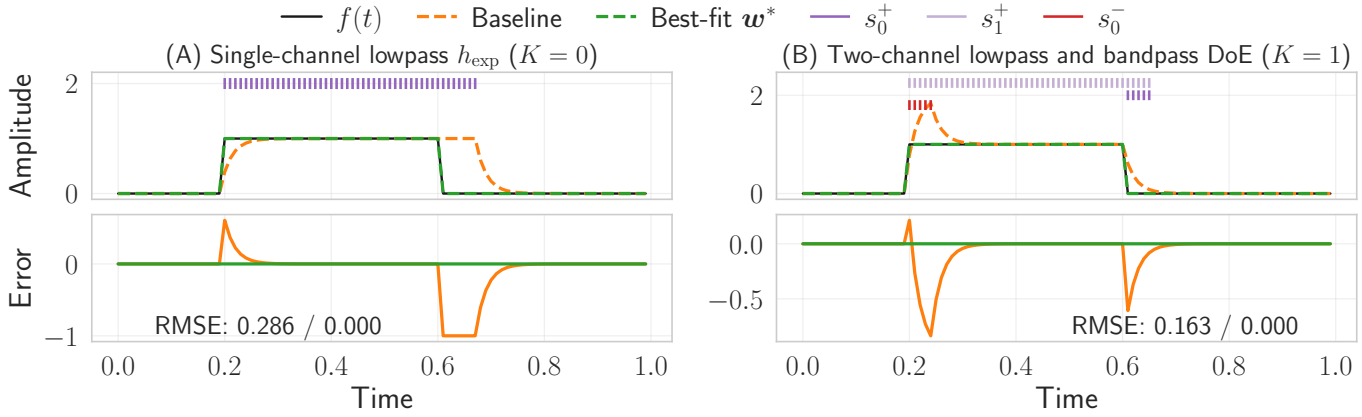

Fig. 5: Reconstruction using identity weights (61) (baseline) or best-fit (39) of a boxcar signal using a lowpass filter ($K = 0$, a) and a lowpass combined with a single bandpass channel DoE wavelet ($K = 1$, b). RMSE is quoted per panel as baseline / best-fit. We use $\mathrm{d}t = 0.01$, $\mu = 2\mathrm{d}t = 0.02$ and $\theta_{\mathrm{thr}} = 0.01$.

where $\rho = 10^{-4}$ and $\gamma = 10^{-1}$ are constants. The full reconstruction is the sum from (22),

$$\widetilde{f}(t) = \widetilde{x}_{\mathrm{lowpass}}(t) - \sum_{k=1}^{K} \widetilde{x}_k(t). \tag{41}$$

Fig. 5 compares the identity weights (61) with the best-fit (39) on a boxcar signal at coarse time resolution ($\mu = 2\mathrm{d}t$): a single lowpass channel ($K = 0$) cannot track the boxcar edges under either weighting, while adding a bandpass channel ($K = 1$) lets the spikes track the transitions and reduces both errors substantially.

This approximate reconstruction offers three advantages:
(1) The encoding is entirely time-recursive and no external history or state is required to propagate the filter states, so encoding runs in real time with bounded memory.
(2) The implementation is time-causal and implementable as differences of leaky integrators, which maps directly to neuromorphic hardware.
(3) The overall scheme preserves scale covariance, since both the encoding wavelets and the reconstruction filters scale covariantly with temporal scaling (20).

Algorithm 1 lists the exact instructions to carry out the encoding and decoding proposed in this paper. The encoding is split into a scale-space encoding and spike generation part, where the channels (K bandpass channels and 1 lowpass), formed by either DoG (26), DoT (28), or DoE (29) representations, are quantized into polarized spike trains (35) using LIF (19). Rather than producing a bandpass representation directly from the difference kernels ψ, we build the bandpass representations from differences of lowpass representations (21), which we reuse for the different bandpass scales for efficiency. Note that the separate spike train for the lowpass channel corresponds to

the signal representation $L(t;\ \sigma_K, c)$ (1) at the coarsest scale K. The decoder first recovers the channel representations from the spike trains, from which f is built by summation (41). Again, the lowpass channel is processed separately from the bandpass channels to carry the coarsest signal structure.

Algorithm 1 Spiking wavelet encoding and reconstruction with lowpass and bandpass channels.

Input: timeseries $f(t)$, wavelet kernel $h \in \{h_{\mathrm{Gauss}}, h_{\Psi}, h_{\mathrm{exp}}\}$, K, spike threshold θ_{thr}
Output: spike times $\{\mathcal{S}_1^{\pm}, \cdots, \mathcal{S}_K^{\pm}, \mathcal{S}_{\mathrm{lowpass}}^{\pm}\}$ (encoder), timeseries $\widetilde{f}(t)$ (decoder)

Scale-space lowpass encoding (ENCODER)
for $k = 0, 1, \ldots, K$ **do**
  $L_k(t;\sigma_k, c) \leftarrow h(t;\sigma_k) * f(t)$ (1)
end for

Bandpass and spike generation (ENCODER)
for $k = 1, \ldots, K$ **do**
  $\Delta L_k(t;\sigma_k, c) \leftarrow L_k(t;\sigma_k, c) - L_k(t;\sigma_{k-1}, c)$ (21)
  $x_k \leftarrow \Delta L_k / \|\psi(t;\sigma_k, c)\|_2$ (24)
  $\mathcal{S}_k^{\pm} \leftarrow \{t\colon u_k^{\pm}(t) = \theta_{\mathrm{thr}}\}$ (35)
end for
$x_{\mathrm{lowpass}} \leftarrow L_K / \|h(t;\sigma_K)\|_2$ (24)
$\mathcal{S}_{K,\mathrm{lowpass}}^{\pm} \leftarrow \{t\colon u_{\mathrm{lowpass}}^{\pm}(t) = \theta_{\mathrm{thr}}\}$ (35)
  ▷ LIF on lowpass channel $L(t;\sigma_K, c)$

Per-channel decoding (DECODER)
for $k = 1, \ldots, K$ **do**
  $A_{t,i}^{(k)} \leftarrow p_i R_k\left(t - t_{\mathrm{spike}}^{(i)}\right),\ i \in \mathcal{S}_k^{\pm}$ (38)
  $\boldsymbol{w}^{(k)*} \leftarrow \arg\min_{\boldsymbol{w}} \left\|\Delta L_k(t;\sigma_k, c) - \boldsymbol{A}^{(k)}\boldsymbol{w}\right\|_2^2$ (39)
  $\widetilde{\Delta L}_k \leftarrow \boldsymbol{A}^{(k)}\boldsymbol{w}^{(k)*}$
end for
$A_{t,i}^{(\mathrm{lp})} \leftarrow p_i R_{K,\mathrm{lowpass}}\left(t - t_{\mathrm{spike}}^{(i)}\right),\ i \in \mathcal{S}_{K,\mathrm{lowpass}}^{\pm}$
$\boldsymbol{w}^{(\mathrm{lp})*} \leftarrow \arg\min_{\boldsymbol{w}} \left\|L_{K,\mathrm{lowpass}}(t;\sigma_K, c) - \boldsymbol{A}^{(\mathrm{lp})}\boldsymbol{w}\right\|_2^2$
$\widetilde{L}_{K,\mathrm{lowpass}}(t;\sigma_K, c) \leftarrow \boldsymbol{A}^{(\mathrm{lp})}\boldsymbol{w}^{(\mathrm{lp})*}$
  ▷ Low-pass spike channel decoding

Signal synthesis (DECODER)
$\widetilde{f}(t) \leftarrow \widetilde{L}_{K,\mathrm{lowpass}}(t;\sigma_K, c) - \sum_{k=1}^{K} \widetilde{\Delta L}_k(t;\sigma_k, c)$ (41)

C. Reconstruction and quantization errors

The quality of the reconstruction depends on the spectral decay of the bandpass channels and the quantization error introduced by the spiking encoding. The faster spectral decay for the DoT (31) means that each DoT channel has a well-defined effective bandwidth Ω_k, which guarantees that the spectral mass outside Ω_k is bounded for any desired tolerance. For the DoE, the slower decay requires additional assumptions on the signal.

There are two sources of error in the spiking case: a mismatch in amplitude introduced by the discrete spikes and a mismatch in time, introduced by the simulated temporal discretization. We begin by studying the continuous case, where the spectral decay of time-scaled responses provides bandwidth limits for each channel which, in turn, bounds the amplitude error for the individual channels. We then proceed to study the temporal discretization, which we combine with the amplitude error above to bound the overall spiking reconstruction process.

1) Spike amplitude error: Let Ω be the *bandwidth* of signal f, and $\neg\Omega = \{\omega\colon |\omega| > \Omega\}$ be the region outside it, not restricted to the -3 dB margin in (30). The spectral mass in $\neg\Omega$ is bounded from above by some tolerance Δ_Ω [32, Sec. 2.4]. For each bandpass channel, the filtered signal $\widehat{\Delta L}_k(\omega;\ \sigma_k, c)$ satisfies

$$\begin{aligned}\int_{\omega\in\neg\Omega_k} |\widehat{\Delta L}_k(\omega;\ \sigma_k, c)|\,\mathrm{d}\omega &= \int_{\omega\in\neg\Omega_k} |\widehat{\psi}(\omega;\ \sigma_k, c)|\,|\widehat{f}(\omega)|\,\mathrm{d}\omega \\ &\leqslant C\int_{\omega\in\neg\Omega_k} |\omega|^{-N}\,|\widehat{f}(\omega)|\,\mathrm{d}\omega, \qquad (42)\end{aligned}$$

for some $C \in \mathbb{R}^+$, where the last step holds for $|\omega|$ sufficiently large by (31). A finite Ω_k meeting the tolerance Δ_Ω exists whenever the right-hand side is finite and decreasing in Ω_k. That depends on N and on the decay of $\widehat{f}$ jointly, which separates the three families.

For $N = 1$ (DoE), the $O(|\omega|^{-1})$ decay of the bandpass kernel is not sufficient to guarantee convergence of (42) from the kernel alone. Additional assumptions on $\widehat{f}$, or the use of reconstruction error feedback mechanisms, are required to ensure that the approximate bandwidth condition is met. For the non-causal DoG, the Gaussian decay (31) dominates, so the tail integral (42) converges to an essential bandwidth.

Assume the bandpass signal $\Delta L_k(t;\ \sigma_k, c)$ has essential bandwidth Ω_k at tolerance $\Delta_\Omega = \theta_{\mathrm{thr}}$, where θ_{thr} is the firing threshold in (35). A spectral tail carrying less than one threshold of charge is invisible to the encoder. Any two signals producing the same spike train have running integrals (e.g. the leaky integrator (16)) agreeing to within θ_{thr} and differentiation of a signal with bandwidth Ω_k increases its supremum by at most a factor Ω_k. This bounds the per-channel amplitude encoding error [32, Th. 2.2]

$$\|\epsilon_{\mathrm{amp},k}\|_\infty = \|\Delta L_k(t;\ \sigma_k, c) - \widetilde{\Delta L}_k(t;\ \sigma_k, c)\|_\infty \leqslant C\,\theta_{\mathrm{thr}}\,\Omega_k, \qquad (43)$$

where $C > 0$ is a constant absorbing leakage and the tail outside the band. This bound is model-agnostic and holds for any (spiking) reconstruction method.

The error ϵ between the original signal f and the reconstruction $\tilde{f}$ from (22) is then

$$\begin{aligned} f(t) - \widetilde{f}(t) &= L_K - \widetilde{L}_K - \sum_{k=1}^{K}\left(\Delta L_k - \widetilde{\Delta L}_k\right) \\ &= \epsilon_{\mathrm{amp},K} - \sum_{k=1}^{K}\epsilon_{\mathrm{amp},k}. \end{aligned} \qquad (44)$$

Inserting (43), where the lowpass residual contributes the essential bandwidth $Q\,\Omega_K$ (30) with the constant-Q factor absorbed into C, we get the total reconstruction error

$$\|f(t) - \widetilde{f}(t)\|_\infty \leqslant C\ \theta_{\mathrm{thr}}\ \left(\Omega_K + \sum_{k=1}^{K}\Omega_k\right). \qquad (45)$$

Given that the bandpass filters are self-similar over scales, a reasonable measure of the bandwidth Ω_k of the bandpass filters must obey the following variability over scales

$$\Omega_k = \Omega_1 / c^{k-1}, \tag{46}$$

which when inserted into (45) by the sum of a geometric series gives the closed-form expression

$$\|f(t) - \widetilde{f}(t)\|_\infty \leqslant C\ \theta_{\mathrm{thr}}\ \left(\frac{\Omega_1}{c^{K-1}} + \Omega_1 \left(\frac{1 - \frac{1}{c^K}}{1 - \frac{1}{c}} \right) \right). \tag{47}$$

For $K \to \infty$ this gives the error bound

$$\|f(t) - \widetilde{f}(t)\|_\infty \leqslant C\ \theta_{\mathrm{thr}} \left(\frac{\Omega_1\, c}{c-1} \right). \tag{48}$$

The bound decreases with c because it counts quantization error but not rate, so it cannot select c or K. That trade-off is measured in Section V-C.

2) Discretization error: In simulations, we are limited to a time-granularity of Δt. Additionally, the temporal discretization introduces another error from the temporal uncertainty of the spike times: in a causal simulation, a spike firing at $t_0 + \epsilon$, for some number $\epsilon > 0$, is only registered at $t_0 + \Delta t$, giving a worst-case delay of Δt per spike. For m_k spikes in a channel k, the per-channel error, combining amplitude quantization (43) and temporal discretization is

$$\|\epsilon_{\mathrm{disc},k}\|_\infty = \|\epsilon_{\mathrm{amp},k}\|_\infty + m_k\, \Delta t\, \Omega_k \leqslant C \theta_{\mathrm{thr}}\, \Omega_k + C\, m_k\, \Delta t\, \Omega_k. \tag{49}$$

The total reconstruction from a time-discretized reconstruction $\widetilde{f}_{\mathrm{disc}}$ is then

$$\begin{aligned} \|f(t) - \widetilde{f}_{\mathrm{disc}}(t)\|_\infty &\leqslant \|\epsilon_{\mathrm{amp},K}\|_\infty + \sum_{k=1}^{K} \|\epsilon_{\mathrm{disc},k}\|_\infty \\ &\leqslant C\ \theta_{\mathrm{thr}}\ \left(\Omega_K + \sum_{k=1}^{K} \Omega_k \right) + C\, \Delta t \sum_{k=1}^{K} m_k\, \Omega_k. \end{aligned} \tag{50}$$

V. Experiments

We test the bandpass and spiking wavelets against three properties that an encoder-decoder pair should satisfy: frequency localization, scale covariance, and formal bounds for stable reconstruction. Frequency localization follows analytically from the channel tiling and channel overlap (see Section III-C), so the experiments below address scale covariance and the reconstruction bounds. The reconstruction is measured against two baseline continuous-time wavelets and two event encoders. For the continuous-time wavelets, we use the non-causal Morlet wavelet that scales and shifts Gaussians [4, (3.3.26)], and the causal Szu wavelet that are restricted to observing past inputs [33, (10)].

For the event encoders we use $\Sigma\Delta$ (53) and a multi-channel integrate-and-fire time-encoding machine (MC-IFTEM) [9]. MC-IFTEM fires the mth of its M integrate-and-fire channels at trigger times t_k^m set by the t-transform that maps amplitudes into the time domain [9, (2)]

$$\int_{t_k^m}^{t_{k+1}^m} \big(f(t) + \beta\big)\, \mathrm{d}t = 2\kappa_{\mathrm{tem}}\, \theta_{\mathrm{tem}}, \qquad m = 1, \ldots, M, \tag{51}$$

where β is a bias, κ_{tem} the integrator constant, and θ_{tem} the threshold (δ in [9]). MC-IFTEM recovers $f(t)$ from the M trains by solving a dense frame system for the coefficients $\boldsymbol{a}$ of the interval-integral sinc frame against the measurements $\boldsymbol{q}$,

$$G\, \boldsymbol{a} = \boldsymbol{q}, \qquad G \in \mathbb{R}^{P \times P}, \quad P = |\mathcal{S}| - M \tag{52}$$

one measurement per interspike interval, at $\mathcal{O}(|\mathcal{S}|^3)$ cost in the pooled event-trigger count $|\mathcal{S}|$.

A $\Sigma\Delta$ modulator [32] instead integrates the difference between $f[t]$ and its one-bit feedback $b[t]$ through two states v_1, v_2, binarized by sign (second-order form) [34, (11)]

$$\begin{aligned} v_1[t] &= v_1[t-1] + f[t] - b[t-1], \\ v_2[t] &= v_2[t-1] + v_1[t] - b[t-1], \\ b[t] &= \operatorname{sign}\big(v_2[t]\big). \end{aligned} \tag{53}$$

The event encoders are compared against the wavelets at a matched spike rate. For an encoder E and a signal $f(t)$, let $\mathcal{S}_E(f)$ be the set of events over channels ($2(K+1)$ for the polarized spiking wavelet (35), the M trigger trains in the MC-IFTEM (51), or the one-bit output of $\Sigma\Delta$ (53)). The spike rate is the event count per input sample, averaged over the evaluation set $\mathcal{D}$,

$$\bar{r}(E) = \frac{1}{|\mathcal{D}|} \sum_{f \in \mathcal{D}} \frac{\big|\mathcal{S}_E(f(t))\big|}{|f|}, \tag{54}$$

where $|f|$ is the number of samples in $f(t)$, i.e. one second at the dataset's sampling rate ($|f| = 360$ for MIT-BIH, $|f| = 16\,000$ for LibriSpeech). The Nyquist-rate sampling of a signal band-limited to $f_{\max}$ acquires $2f_{\max}$ samples per second [1, Thm. 1], i.e., $2f_{\max}\, \Delta t$ events per timestep in the units of (54). In our experiments, $f_{\max} = 1/(2\Delta t)$ is the highest frequency representable at each dataset's sampling rate (180 Hz for MIT-BIH, 8 kHz for LibriSpeech), so the Nyquist rate is exactly one event per timestep. Each baseline is a one-parameter family E_λ, with λ the oversampling ratio for $\Sigma\Delta$ and the threshold θ_{tem} (51) for MC-IFTEM, and $\bar{r}(E_\lambda)$ is monotone in λ. Spike-rate matching selects $\lambda^\star$ such that

$$\big|\bar{r}(E_{\lambda^\star}) - \bar{r}(E_{\mathrm{ref}})\big| \leqslant \varepsilon\, \bar{r}(E_{\mathrm{ref}}), \qquad \varepsilon = 0.05, \tag{55}$$

by bisection on λ, with E_{ref} the spiking DoT at $\theta_{\mathrm{thr}} = 0.1$. This keeps the number of events per input sample on-par with the DoT reference wavelet, comparing fidelity at a controlled event rate.

The experiments are conducted on the MIT-BIH heart arrhythmia dataset [20] and the LibriSpeech audiobook recordings [21]. The datasets are chosen for their different characteristics (see Fig. 6 for signal samples): MIT-BIH ECG features bursts of activity sampled at 360 Hz, while LibriSpeech contains dense speech signals sampled at 16 kHz. We split the data into one-second samples.

A. Experimental details

All simulations are clocked to the native sampling rate of each dataset. The DoG (26), DoE (29), and DoT (28) wavelets represent the signal exactly. The continuous wavelet transforms (CWT) and spiking kernels are continuous in time,

but evaluated on a discrete grid, introducing discretization error that vanishes at $\Delta t \to 0$. We choose Δt as the inverse of the sampling rate, $\Delta t = \{1/360, 1/16000\}$ for MIT-BIH and LibriSpeech, respectively. In all experiments, the DoT wavelet uses a finite cascade of $N = 7$ truncated-exponential stages (27) to approximate the time-causal limit kernel. For fixed bandwidth minimum and maximum $[\Omega_{\min}, \Omega_{\max}]$, we set the scale ratio c according to [14, (29)].

As error metric, we rely on the root mean squared error (RMSE) standardized to zero mean and unit variance. Standardization permits us to compare reconstruction error across LibriSpeech and MIT-BIH ECG datasets. For non-spiking wavelets, $\widetilde{f}$ comes from frame synthesis (15). For spiking systems, $\widetilde{f}$ is the per-spike least-squares reconstruction (38).

The probe of the scale-covariance experiment (56) uses $n_{\mathrm{cyc}} = 3$ and is centered at one tenth of the window. We place the carrier at $f_0 = 72$ Hz for MIT-BIH and $f_0 = 3$ kHz for LibriSpeech, high in each band so that dilation has room to move the activity toward the coarser channels.

We chose the MIT-BIH ECG [20] and LibriSpeech [21] datasets because they express two distinct characteristics. The MIT-BIH dataset consists of heart arrhythmia recordings that mostly oscillate around 0, with short, transient bursts. The LibriSpeech contains recordings of audiobooks by English speakers. Fig. 6 visualizes the different nature of the signals. The LibriSpeech has much more structure and requires a denser sampling to recover the 16 kHz signal.

The experimental details and code are openly available at github.com/jegp/swavelet. The spiking DoE and DoT encoders additionally export to the Neuromorphic Intermediate Representation [22] as a graph of leaky-integrator, linear, and LIF primitives for deployment on neuromorphic backends.

B. Reconstruction experiment

Using Algorithm 1, we measure the ability of the wavelet systems to reconstruct signals from their respective encodings. We group the wavelets into four classes: bandpass wavelets (DoG, DoE, and DoT), continuous wavelets (CWT), spiking wavelets, and event encoders (MC-IFTEM and $\Sigma\Delta$). The spiking DoG serves as a non-causal control: it shares the full encoding and decoding pipeline with the causal spiking DoE and DoT, so their comparison isolates the cost of causality within the same construction. For the non-spiking DoG (26), DoE (29), and DoT (28) wavelets, the reconstruction is exact by construction. For the Morlet wavelets (non-causal) and the Szu wavelets (causal), we compute dense coefficients $c_k(t) = \langle x, \psi_k(\cdot - t)\rangle$ for each channel k and reconstruct via a least-squares pseudoinverse.

All event encoders are decoded by the same oracle regularized least squares (40) fitted against the clean signal. The spiking wavelet fits its per-spike weights onto the bandpass channel outputs x_k, $\Sigma\Delta$ its moving-average atoms, and the MC-IFTEM its sinc frame. The reconstructions therefore upper-bound what a decoder using these atoms can recover from the spike trains. All event encoders are compared across the same event crossing threshold. The spiking wavelets share the firing threshold $\theta_{\mathrm{thr}} = 0.1$, and the MC-IFTEM threshold θ_{tem} and the $\Sigma\Delta$ oversampling ratio are set once by rate-matching (55) to the spiking DoT. The event rates are therefore an outcome of the encoding, reported per entry in Table II.

TABLE II: Reconstruction RMSE across 100 samples from the MIT-BIH ECG [20] and LibriSpeech [21] datasets. Event rates (54) are reported in parentheses.

	MIT-BIH			
Wavelet	c=2, K=8		c=$\sqrt{2}$, K=15	
DoG/DoE/DoT	exact		exact	
Morlet	0.079 ± 0.038		0.060 ± 0.035	
Szu (causal)	0.291 ± 0.096		0.190 ± 0.095	
Spiking DoG	0.068 ± 0.043	(6.2)	0.068 ± 0.044	(11.3)
Spiking DoE	0.085 ± 0.042	(5.9)	0.118 ± 0.051	(8.9)
Spiking DoT	$\mathbf{0.059 \pm 0.035}$	(7.0)	0.070 ± 0.036	(11.6)
MC-IFTEM	0.096 ± 0.168	(6.3)	$\mathbf{0.047 \pm 0.109}$	(11.3)
$\Sigma\Delta$	0.181 ± 0.118	(7.0)	0.106 ± 0.065	(12.0)
	LibriSpeech			
	c=2, K=6		c=$\sqrt{2}$, K=12	
DoG/DoE/DoT	exact		exact	
Morlet	0.079 ± 0.081		$\mathbf{0.064 \pm 0.074}$	
Szu (causal)	0.529 ± 0.123		0.355 ± 0.122	
Spiking DoG	0.072 ± 0.040	(3.5)	0.070 ± 0.040	(6.3)
Spiking DoE	0.092 ± 0.047	(3.5)	0.139 ± 0.059	(4.8)
Spiking DoT	$\mathbf{0.067 \pm 0.041}$	(4.6)	0.103 ± 0.115	(7.4)
$\Sigma\Delta$	0.277 ± 0.048	(4.0)	0.138 ± 0.026	(8.0)

Table II shows the results for the exact bandpass wavelets (DoG, DoE, and DoT), as well as the continuous, but discretized, wavelets (Morlet, Szu, and spiking versions of DoG, DoE, and DoT) and the event-encoder baselines (MC-IFTEM (51) and $\Sigma\Delta$ (53)). Despite the spike quantization, the spiking wavelets stay close to the non-causal Morlet baseline, with the DoT matching or exceeding it at $c = 2$. This is notable because the spiking DoT carries both quantization and discretization error, whereas the non-causal Morlet's error is purely due to discretization from the digital sampling.

Moving from $c = 2$ to the denser $c = \sqrt{2}$ grid leaves the spiking DoG essentially unchanged, while the error of the spiking causal DoE and DoT wavelets grows. This is attributed to the increasing overlap between channels: as $c \to 1$ the bandpass responses overlap more strongly, the off-diagonal entries of the normalized Gram matrix (33) grow, and the frame bound ratio B/A rises (see Section III-C).

The causal Szu wavelet is subject to the same sampling issues as the Morlet wavelet, but the causality requirement incurs a 3–7× RMSE penalty. This strengthens the results of the spiking DoE and DoT wavelets, that are both causal and competitive in the reconstruction metric.

Of the event encoders in Table II, MC-IFTEM matches the spiking wavelets but $\Sigma\Delta$ does not, trailing the best spiking wavelet by between 1.5 and 4× at the same event rate. For the MIT-BIH ECG, the matched rate yields $\sim 10^3$ events, small enough that the $\mathcal{O}(|\mathcal{S}|^3)$ joint decode (52) is tractable. For LibriSpeech it reaches $|\mathcal{S}| \sim 10^5$ events, at which the decode is intractable; we therefore omit MC-IFTEM there.

Fig. 7 sweeps the firing threshold θ_{thr}: the measured RMSE

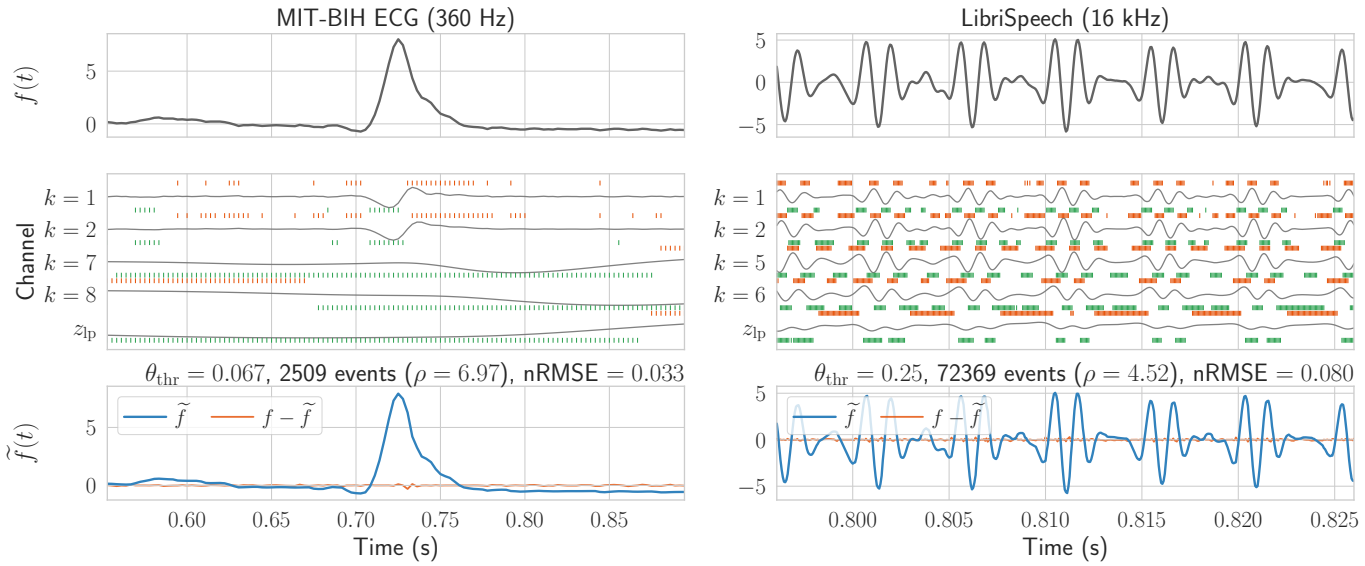

Fig. 6: Spiking DoT encoding and reconstruction of real signals, at the rates Table II reports for $c = 2$. **Top**: the signal $f(t)$. **Middle**: polarity spikes $z_k^{\pm}$ (17) for the two finest and two coarsest bandpass channels and the lowpass residual, with the normalized channel response $\bar{x}_k$ (24) drawn behind each row, scaled individually. The threshold crossing itself is shown in Fig. 5. **Bottom**: the reconstruction $\widetilde{f}$ (41) and the residual $f - \widetilde{f}$.

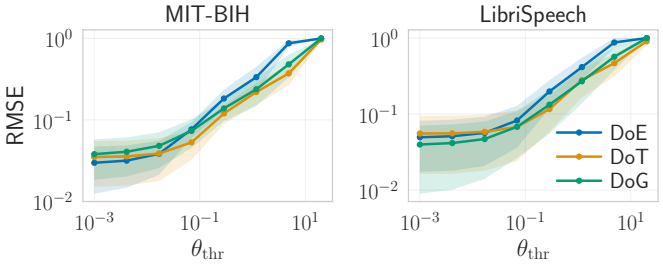

Fig. 7: RMSE versus θ_{thr}, bounded linearly by θ_{thr} following (48). The graded area shows one standard deviation of the RMSE across 10 samples.

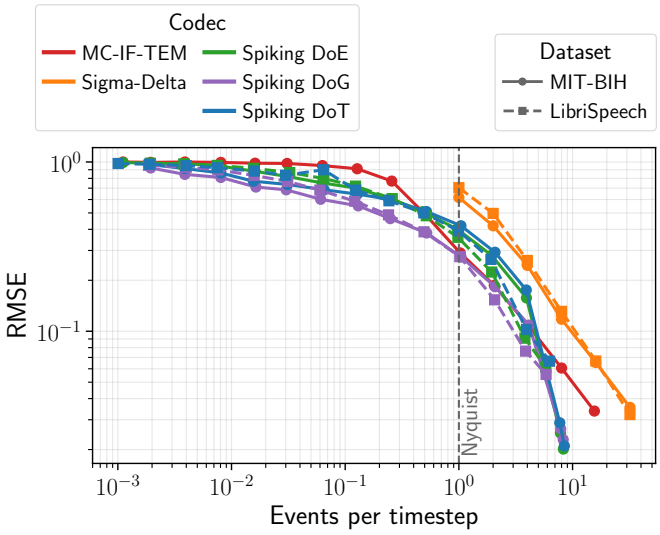

Fig. 8: Operational rate-distortion: reconstruction RMSE versus events per timestep for the spiking wavelets (DoT, DoE, DoG), a second-order $\Sigma\Delta$ modulator (53), and the MC-IFTEM (51), on MIT-BIH (solid) and LibriSpeech (dashed). Every codec is matched onto the same grid of target rates (55). The dashed vertical marks the Nyquist rate $2f_{\max}\Delta t = 1$ (54). MC-IFTEM is shown on MIT-BIH only, as its $\mathcal{O}(|\mathcal{S}|^3)$ joint decode (52) is intractable at the 16kHz LibriSpeech rate.

stays within the linear envelope of the reconstruction bound (48), demonstrating the predicted linear dependence. Below $\theta_{\mathrm{thr}} \approx 0.02$ the curves flatten to a level determined by the temporal discretization (49).

C. Spike rate

Table II fixes a single firing threshold, but any encoder can spend more spikes to lower its error, so one operating point is only a snapshot. We therefore report the operational rate-distortion: reconstruction RMSE as a function of the spike rate, the code cost in events per timestep, which asks which encoder reaches a given fidelity at the lowest rate. Fig. 8 plots this trade-off for the spiking wavelets, a second-order $\Sigma\Delta$ modulator (53) [32], and MC-IFTEM (51), all at $c = 2$. The sweep therefore does not cover the $c = \sqrt{2}$ setting at which MC-IFTEM leads Table II, where it encodes with $M = 16$ shifted integrators rather than the $M = 9$ used here. The $\Sigma\Delta$ curve stops at one event per timestep: the modulator (53) emits one bit per input sample at its lowest rate setting, and every emitted bit counts as an event.

The spiking wavelets improve steadily with rate throughout. They reach a given error at one and a half to four times fewer events than $\Sigma\Delta$, which cannot encode below the Nyquist rate at all, while MC-IFTEM matches them below $\sim$4 events per

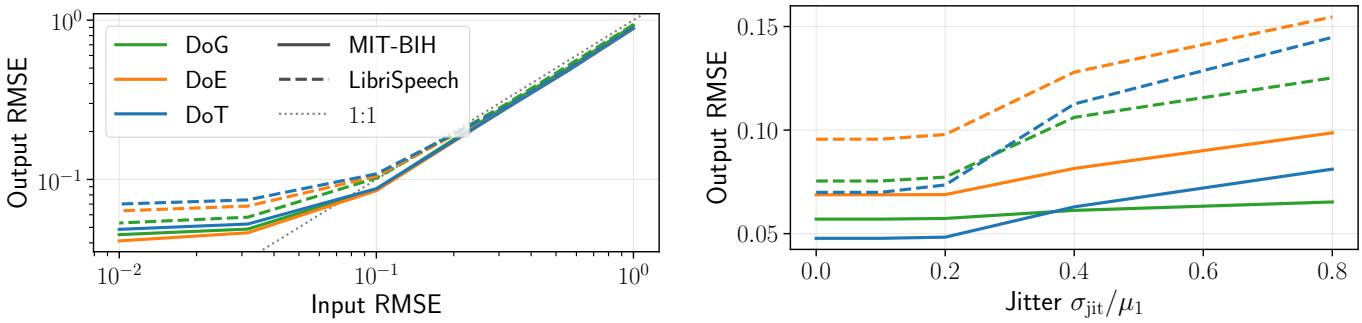

Fig. 9: Robustness of the spiking wavelets, scored against the clean reference. **Left**: output RMSE versus input RMSE under additive white Gaussian noise. **Right**: output RMSE versus Gaussian spike-timing jitter of standard deviation σ_{jit}, as a fraction of the finest channel scale μ_1.

timestep on MIT-BIH and falls behind above it.

D. Robustness

The reconstruction bound (47) holds for any finite-energy input, so the encoding should degrade gracefully under noise. To test this, we add white Gaussian noise to the input across a range of noise levels, shown as input RMSE, and score the reconstruction $\widetilde{f}$ in RMSE against the clean reference (Fig. 9, left). Both axes are amplitude ratios in RMSE, so an identity codec would sit on the 1:1 line. The output RMSE tracks the input RMSE roughly one-to-one down to ≈ 0.1, then saturates at the codec's own reconstruction floor, 0.049 on ECG and 0.070 on speech. The encoding therefore neither amplifies the input noise nor saturates above that floor, which is a precondition for use on real, noisy signals.

Especially on hardware, spike times are subject to jitter, which can be modeled as a random perturbation of the spike times. We perturb the encoded spike-times with a zero-mean Gaussian jitter of standard deviation σ_{jit}, expressed as a fraction of the finest channel scale μ_1, shown in Fig. 9, right. The output RMSE stays constant while σ_{jit} stays under ~ 0.1 of the finest channel scale μ_1, then grows sub-linearly in σ_{jit}, staying within about twice the jitter-free error throughout. A spike displaced by σ_{jit} contributes a per-channel error of order $\sigma_{\text{jit}}/\mu_k$, so the coarse channels are intrinsically robust and the finest channel dominates.

E. Scale-covariant encoding

A scale-covariant encoder should shift the spike distribution across channels predictably when the input is dilated in time, where a non-covariant encoder would not. To demonstrate this, we encode a sine wave with a Gaussian centered at t_0,

$$g(t; f_0) = \frac{1}{\|g\|_2} e^{\frac{-(t-t_0)^2}{2\sigma_g^2}} \sin(2\pi f_0 t), \quad \sigma_g = n_{\text{cyc}}/f_0, \quad (56)$$

dilate it by factors $a \in [1, 8]$, and measure the displacement of the per-channel spike-activity centroid. Here σ_g is the standard deviation of the Gaussian envelope and n_{cyc} the number of oscillations per σ_g, so every frequency carries the same number of oscillations and stretching the packet in time returns the same shape at f_0/a. The packet is normalized to unit energy so that spike counts stay comparable across a; the centroid is the spike-count-weighted average across channels. See Section V-A for n_{cyc} and the frequencies f_0.

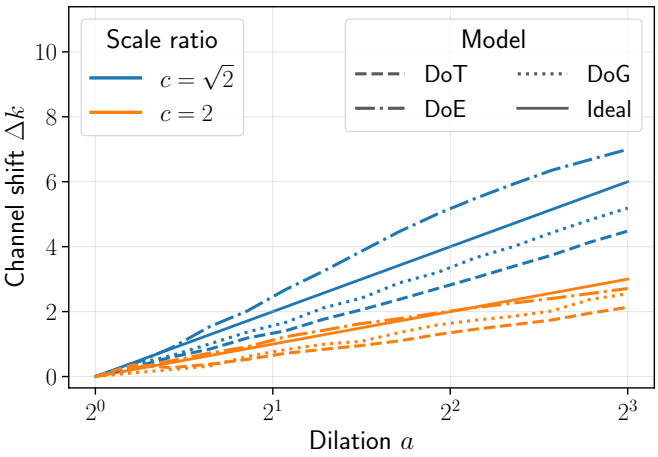

Fig. 10: Scale covariance under temporal rescaling. By dilating the envelope (56) by a factor a and encoded with the spiking wavelets, the spike-activity centroid migrates by Δk, as predicted by the covariance law $\log_c a$ (solid).

Fig. 10 shows that the measured shift follows the covariance law $\Delta k = \log_c a$: the two scale ratios separate closely along the expected slope, with finer spacing $c = \sqrt{2}$ shifting the activities by twice as many channels as $c = 2$. All three spiking wavelet families are covariant to within one channel at $c = 2$ and 1.5 at $c = \sqrt{2}$, and the residual grows at the largest a as the activity centroid approaches the coarsest channel and can migrate no further.

VI. Discussion

We contributed two time-causal and scale-covariant families of spiking wavelets as sparse continuous-time encoders. We derived closed-form reconstruction error bounds for the spiking DoE and DoT wavelets and demonstrated matching or exceeding performance compared to non-causal Morlet wavelet at $c = 2$ and the causal Szu wavelet, while remaining causal and sparse per channel.

A strength of this construction is its generality: the frame bounds and reconstruction guarantees hold for any finite-energy input, with no assumptions on signal statistics or stationarity. The spiking reconstruction bound (47) additionally requires each channel to have an effective bandwidth, which the DoT cascade guarantees through its $O(|\omega|^{-N})$ decay. The slower $O(|\omega|^{-1})$ decay of the DoE requires further assumptions on the signal spectrum. The guarantees presented here are deterministic and worst-case, complementing an information-theoretic, average-case account of spike coding [5], [35].

Continuous-time spiking dynamics are simulated on a digital substrate at fixed time intervals. The method's natural home is event-driven neuromorphic hardware, where continuous dynamics are native and the discretization artifacts disappear. Our study establishes the feasibility of spiking wavelets for signal encoding and reconstruction, with a least-squares decoder that sets the upper bound any causal decoder must approach. The baseline comparisons to $\Sigma\Delta$ and MC-IFTEM rely on the same decoder for fairness, but no deployable spike-only decoder is evaluated. Closing the gap to online, streaming decoders is now a well-posed follow-up, as is quantifying the energy, latency, and memory benefits on analog or mixed-signal platforms.

Because every operation in the spiking wavelets reduces to a leaky-integrator difference with spike thresholding, the encoders map directly onto primitives common to neuromorphic platforms, positioning spiking wavelets as a principled building block for neuromorphic analog-to-digital conversion. The wavelets encode time-varying signals into spike trains that are sparse within each of the $2(K+1)$ polarity channels, with provable reconstruction guarantees. A single threshold θ_{thr} sets the event rate, and the reconstruction error grows linearly with it (48). Based on this theory and experiments, we propose the presented spiking DoE and DoT wavelets as fundamental spiking representations for neuromorphic processing.

REFERENCES

[1] C. Shannon, "Communication in the Presence of Noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.

[2] K. Roy, A. Jaiswal, and P. Panda, "Towards spike-based machine intelligence with neuromorphic computing," *Nature*, vol. 575, no. 7784, pp. 607–617, Nov. 2019.

[3] M. Dampfhoffer, T. Mesquida, A. Valentian, and L. Anghel, "Are SNNs Really More Energy-Efficient Than ANNs? an In-Depth Hardware-Aware Study," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 7, no. 3, pp. 731–741, Jun. 2023.

[4] I. Daubechies, *Ten Lectures on Wavelets*. SIAM, 1992.

[5] D. Warland, F. Rieke, R. d. R. van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code*, ser. Computational Neuroscience Series. MIT Press, Dec. 1996.

[6] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, 2014.

[7] B. A. Moser, A. Werzi, and M. Lunglmayr, "Integrate-and-fire from a mathematical and signal processing perspective," *2024 58th Asilomar Conference on Signals, Systems, and Computers*, p. 567–571, Oct. 2024.

[8] Y. C. Yoon, "LIF and simplified SRM neurons encode signals into spikes via a form of asynchronous pulse sigma-delta modulation," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 5, p. 1192–1205, May 2017.

[9] K. Adam, A. Scholefield, and M. Vetterli, "Sampling and reconstruction of bandlimited signals with multi-channel time encoding," *IEEE Transactions on Signal Processing*, vol. 68, pp. 1105–1119, 2020.

[10] V. K. Goyal, J. Kovačević, and J. A. Kelner, "Quantized frame expansions with erasures," *Applied and Computational Harmonic Analysis*, vol. 10, no. 3, p. 203–233, May 2001.

[11] J. E. Pedersen, J. Conradt, and T. Lindeberg, "Covariant spatio-temporal receptive fields for spiking neural networks," *Nature Communications*, vol. 16, no. 1, p. 8231, 2025.

[12] J. E. Pedersen, T. Lindeberg, and P. Gerstoft, "Scale-covariant spiking wavelets," *IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP 2026)*, pp. 20 347–20 351, 2026.

[13] T. Lindeberg, "Time-causal and time-recursive wavelets," *Frontiers in Signal Processing*, 2026, to appear, preprint at arXiv:2510.05834.

[14] J. E. Pedersen, T. Lindeberg, and P. Gerstoft, "Frame bounds and bandwidth tiling for time-causal bandpass wavelets," *arXiv preprint arXiv:2609.17201*, 2026.

[15] I.-M. Comşa, L. Versari, T. Fischbacher, and J. Alakuijala, "Spiking autoencoders with temporal coding," *Frontiers in Neuroscience*, vol. 15, p. 712667, 2021.

[16] H. Kamata, Y. Mukuta, and T. Harada, "Fully spiking variational autoencoder," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 6, p. 7059–7067, June 2022.

[17] Z. Fang, Y. Liu, and G. Ye, "Event-driven signal reconstruction through neuromorphic compressive sensing," Sep. 2026.

[18] A. A. Lazar and L. T. Tóth, "Perfect recovery and sensitivity analysis of time encoded bandlimited signals," *IEEE Transactions on Circuits and Systems*, vol. 51, no. 10, pp. 2060–2073, 2004.

[19] D. Gontier and M. Vetterli, "Sampling based on timing: Time encoding machines on shift-invariant subspaces," *Applied and Computational Harmonic Analysis*, vol. 36, no. 1, pp. 63–78, 2014.

[20] G. Moody and R. Mark, "The impact of the MIT-BIH Arrhythmia Database," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 3, pp. 45–50, 2001.

[21] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.

[22] J. E. Pedersen, S. Abreu, M. Jobst, G. Lenz, V. Fra, F. C. Bauer, D. R. Muir, P. Zhou, B. Vogginger, K. Heckel *et al.*, "Neuromorphic intermediate representation: A unified instruction set for interoperable brain-inspired computing," *Nature Communications*, vol. 15, no. 1, p. 8122, 2024.

[23] J. J. Koenderink, "The structure of images," *Biological Cybernetics*, vol. 50, no. 5, pp. 363–370, 1984.

[24] T. Lindeberg, *Scale-Space Theory in Computer Vision*. Boston, MA: Springer US, 1994, vol. 256.

[25] ——, "A time-causal and time-recursive scale-covariant scale-space representation of temporal signals and past time," *Biological Cybernetics*, vol. 117, no. 1, pp. 21–59, 2023.

[26] M. W. Howard, Z. G. Esfahani, B. Le, and P. B. Sederberg, "Learning temporal relationships between symbols with Laplace neural manifolds," *Computational Brain & Behavior*, vol. 8, no. 2, pp. 211—232, June 2025.

[27] I. J. Schoenberg, "On variation-diminishing integral operators of the convolution type," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 34, no. 4, pp. 164–169, Apr. 1948.

[28] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed. Burlington, MA: Academic Press, 2009.

[29] P. J. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Transactions on Communications*, vol. COM-31, no. 4, pp. 532–540, 1983.

[30] S. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, Jul. 1989.

[31] D. Marr and E. Hildreth, "Theory of edge detection," *Proceedings of the Royal Society of London. B. Biological Sciences*, vol. 207, no. 1167, pp. 187–217, Feb. 1980.

[32] D. Carbajal and J. L. Romero, "Model agnostic signal encoding by leaky integrate-and-fire, performance and uncertainty," *Applied and Computational Harmonic Analysis*, vol. 83, p. 101856, Mar. 2026.

[33] H. H. Szu, B. A. Telfer, and A. W. Lohmann, "Causal analytical wavelet transform," *Optical Engineering*, vol. 31, no. 9, p. 1825–1829, September 1992.

[34] P. Aziz, H. Sorensen, and J. van der Spiegel, "An overview of Sigma-Delta converters," *IEEE Signal Processing Magazine*, vol. 13, no. 1, pp. 61–84, Jan. 1996.

[35] E. Schneidman, M. J. Berry, R. Segev, and W. Bialek, "Weak pairwise correlations imply strongly correlated network states in a neural population," *Nature*, vol. 440, no. 7087, pp. 1007–1012, Apr. 2006.

APPENDIX A
NEURONAL IMPLEMENTATIONS OF SCALE-SPACE WAVELETS

A. Reconstruction constants

To properly reconstruct the original signal, the contribution of each spiking channel needs to compensate for the scaling terms in both the difference kernels ψ_{DoE} and ψ_{DoT} and the LIF model. Below, we first reverse the scaling from the coupled integration in the wavelet and LIF model. We then derive the initialization scheme for the channel-specific weights based on naive assumptions on the input.

1) Convolution of two truncated exponentials: The spikes generated by the LIF are derived from the combination of two filters with two different temporal scale parameters: μ_k for the ψ_{DoE} with $N = 1$ (29) or σ_m (to first approximation) for the ψ_{DoT} (27) (with a scale factor of exactly $\prod_{n=k}^{\infty} 1/\mu_n$ for the kth time-causal limit kernel) on the one hand, and μ_m for the LIF model on the other. The composition of two truncated exponentials $h_1(t;\ \mu_1) * h_2(t;\ \mu_2)$ does not reduce to a single truncated exponential, but takes the form

$$\begin{aligned}
&\int_{-\infty}^{\infty} h_1(t-\xi;\ \mu_1)\, h_2(\xi;\ \mu_2)\, \mathrm{d}\xi = \frac{1}{\mu_1\,\mu_2} \int_0^t e^{-\frac{t-\xi}{\mu_1}} e^{-\frac{\xi}{\mu_2}}\, \mathrm{d}\xi \\
&= \frac{1}{\mu_1\,\mu_2} e^{-\frac{t}{\mu_1}} \int_0^t e^{\xi(\frac{1}{\mu_2}-\frac{1}{\mu_1})}\, \mathrm{d}\xi \\
&= \frac{-\mu_1\,\mu_2}{\mu_1-\mu_2} \frac{1}{\mu_1\,\mu_2} e^{-\frac{t}{\mu_1}} \left[e^{-\xi\frac{\mu_1-\mu_2}{\mu_1\mu_2}} \right]_0^t \\
&= \frac{1}{\mu_1-\mu_2} \left[e^{-\frac{t}{\mu_1}} - e^{-\frac{t}{\mu_2}} \right]. \qquad (57)
\end{aligned}$$

The factor $1/(\mu_1 - \mu_2)$ is used to recover the scaling that occurs by the composition of the wavelet kernel and the leaky integrator term of the LIF neuron that produces the spikes.

2) Weight initializations: To recover the original signal amplitude, we must recover the original integration value that led to a spike. Recall that the bandpass formulation is based on the difference of exponentials as the generating function. We use C_κ as a normalizing factor

$$\kappa(t;\ \mu_1, \mu_2) = (\mu_1^{-1} e^{-t/\mu_1} - \mu_2^{-1} e^{-t/\mu_2})/C_\kappa. \qquad (58)$$

Since $\int_0^\infty 1/\mu\, e^{-t/\mu} = 1$, we have that

$$C_\kappa = 1/\mu_1 - 1/\mu_2 = \frac{\mu_2 - \mu_1}{\mu_1\,\mu_2}. \qquad (59)$$

We cannot infer the shape of the input signal that caused the spike, since there are infinitely many combinations of the leaky integrator and the input that integrates to θ_{thr}. We can only know that the integration of the difference of the truncated exponentials, beginning at time t_0, reaches θ_{thr} at time t_1,

$$\begin{aligned}
\theta_{\mathrm{thr}} &= \int_{t_0}^{t_1} \kappa(t_1 - t;\ \mu_1, \mu_2) f(t)\ \mathrm{d}t \qquad (60) \\
&= \frac{\mu_2 - \mu_1}{\mu_1\,\mu_2} \int_{t_0}^{t_1} \left(\mu_1^{-1} e^{-(t_1-t)/\mu_1} - \mu_2^{-1} e^{-(t_1-t)/\mu_2}\right)\ f(t)\ \mathrm{d}t.
\end{aligned}$$

Assuming that $f(t) = f_1$ is constant over the interval and setting $w = (t_1 - t)$, which is 0 when $t = t_1$ so that $\mathrm{d}t = -\mathrm{d}w$, we can deduce f_1 from the interval $t_1 - t_0$, μ_1, and μ_2

$$\begin{aligned}
\theta_{\mathrm{thr}} &= f_1 C_\kappa^{-1} \int_{t_0}^{t_1} \mu_1^{-1} e^{-(t_1-t)/\mu_1} - \mu_2^{-1} e^{-(t_1-t)/\mu_2}\ \mathrm{d}t \\
&= f_1 C_\kappa^{-1} \int_{t_1-t_0}^{0} \mu_1^{-1} e^{-w/\mu_1} - \mu_2^{-1} e^{-w/\mu_2}\ (-\mathrm{d}w) \\
&= f_1 C_\kappa^{-1} \left[-e^{-w/\mu_1} - (-e^{-w/\mu_2}) \right]_0^{t_1-t_0} \\
&= f_1 C_\kappa^{-1} \left(e^{-(t_1-t_0)/\mu_2} - e^{-(t_1-t_0)/\mu_1} \right) \iff \\
f_1 &= \theta_{\mathrm{thr}}\ \frac{\mu_2 - \mu_1}{\mu_1\,\mu_2} \frac{1}{e^{-(t_1-t_0)/\mu_2} - e^{-(t_1-t_0)/\mu_1}}. \qquad (61)
\end{aligned}$$