

Event-Integrated Spline Readouts and Finite-Precision Synaptic Realizations

Daniel Schmitter

Abstract

Analytic dynamics can replace time stepping when accumulating a continuous learning objective, but exact integration need not outperform a structure-aware numerical control. We study cubic spline readouts of an observed exponentially decaying trace and derive cellwise sufficient statistics for ridge regression. In 48 synthetic records grouped into six independent seeds, analytic accumulation meets all prescribed accuracy tolerances, whereas the tested fixed-order quadrature rules fail on long quiet intervals. A subsequently constructed quadrature control removes the constant contribution analytically and matches accuracy at essentially equal runtime. We then separate continuous-time realization equivalence from finite-precision behavior in a grounded diffusion model of synaptic memory. Physical and modal realizations agree to floating-point roundoff before quantization; after equal-budget stochastic rounding, the modal realization has larger output distortion in every tested seed at both six and eight effective bits. These results establish an accurate event-integrated learning primitive and two negative computational findings. They do not establish a trained spiking network, a neuromorphic hardware advantage, or increased memory capacity.

1 Introduction

A dynamical learning system may evolve continuously while receiving observations only at events. Simulating every quiet interval on a fine clock can waste work when the between-event trajectory is known. The relevant computational question is not merely whether that trajectory has a closed form, but whether the learning objective can be accumulated accurately and cheaply along it. A second question arises when the state is compressed into a few quantized coordinates: equivalent differential equations may cease to be equivalent implementations.

We examine these questions in two bounded systems. First, a cubic spline maps a scalar exponential trace to a supervised output. Known local polynomial pieces permit analytic accumulation of the normal equations. Second, a stable diffusion chain represents a synaptic consolidation kernel. Diagonalization converts the chain into an exponential bank, providing an exact real-arithmetic comparator. The contribution is a derivation and controlled assessment of these representation choices, including their failure to establish the hypothesized speed and quantization benefits.

2 Related work

Event-based updating of synaptic learning rules is established. Stapmanns and colleagues give event-based implementations of voltage-dependent plasticity rules [1]. EventProp computes event-based gradients in spiking networks [2]. The present problem is different: supplied event times and a known scalar trajectory define a fixed-feature continuous regression objective. We neither differentiate spike times nor learn a recurrent spiking network.

Benna and Fusi motivate synaptic memory through interacting dynamical variables [3]. Diagonalizing a stable symmetric similarity transform is classical linear systems algebra, not a new memory-capacity mechanism. Likewise, integrating local spline pieces is standard numerical analysis. Our experiments test the consequences of combining those facts under explicit observation, timing, and persistent-state budgets. The comparisons deliberately provide the same known dynamics to the numerical controls.

3 Architecture and information flow

In the event learner, time is eliminated only between known events. An impulse determines a new scalar trace state; its known decay determines when spline cells are crossed. Each crossing interval contributes to the continuous normal equations under a held teacher value. A later solve fits readout coefficients. The experiment neither learns event times nor supplies a general recurrent credit-assignment rule.

Let events update the learning objective



Read left to right above, then right to left below. Arrows show data or computation flow.

Figure 1: Let events update the learning objective. Boxes distinguish supplied information, fitted components, and the quantity evaluated. Arrows show computation or data dependence, not a newly trained deep network.

The long quiet tail isolates a numerical singular-looking term that has a simple decomposition. Removing the constant contribution analytically lets low-order quadrature integrate the remaining polynomial exactly on the zero cell. This control uses the same physical information as the analytic routine, and removes the apparent practical speed advantage. In the separate synaptic study, diagonalization preserves ideal input-output dynamics but changes the orientation of quantization cells; real-arithmetic equivalence does not predict rounded performance.

4 Continuous regression along an event trace

Between supplied positive impulses, the trace is nonnegative and satisfies the first equation below. A teacher value is observed at each event and deliberately held constant until the next event. Let the feature vector contain zero-extended cardinal cubic basis functions on an amplitude grid of spacing h . All four active coefficients are retained at every observed amplitude; there is no periodic wrapping or clipping.

$$x(t+s) = x(t)e^{-s/\tau}, \quad f_c(x) = c^T \phi(x), \quad (1)$$

$$G = \int \phi(x(t))\phi(x(t))^T dt, \quad b = \int \phi(x(t))y(t)dt. \quad (2)$$

For a ridge penalty with coefficient λ , the minimizer is obtained from the normal equations. The target-energy constant is unnecessary for the minimizer but is necessary for the absolute objective value. Retaining G and b alone therefore does not retain every possible query about the original history.

$$\mathcal{J}(c) = c^T G c - 2c^T b + \int y(t)^2 dt + \lambda \|c\|^2, \quad (3)$$

$$c_* = (G + \lambda I)^{-1} b. \quad (4)$$

In a cell with local coordinate u , the four cubic weights are the entries of the following vector. Substitution of amplitude for time yields a weighted polynomial integral. Every interval is partitioned at its actual knot crossings before either analytic or quadrature accumulation.

$$p(u) = \frac{1}{6} \begin{pmatrix} (1-u)^3 \\ 3u^3 - 6u^2 + 4 \\ -3u^3 + 3u^2 + 3u + 1 \\ u^3 \end{pmatrix}, \quad u = x/h - \lfloor x/h \rfloor, \quad dt = -\tau dx/x. \quad (5)$$

Products of cubic weights have degree six. On a nonzero amplitude cell, their integral against the reciprocal amplitude can be evaluated by polynomial moments and a logarithmic term; stable short-interval evaluation may use a convergent series. This is analytic integration implemented in floating point, not symbolic exact arithmetic. Since only four features are active, the continuous Gram matrix has local overlap structure for this scalar readout. The prototype nevertheless stores it densely, so a banded-memory saving is not a measured result.

5 A structure-aware quadrature control

The cell adjacent to zero contains the long quiet tail. Direct fixed-order time quadrature can miss the changing part or misallocate its nearly constant contribution. Write its vector polynomial as $p(u)=a+u d(u)$, where d is quadratic and a is the value at zero. For a segment decreasing from u_{hi} to u_{lo} , its quadratic statistic decomposes as

$$\int p(u)p(u)^T dt = \Delta t aa^T + \tau \int_{u_{lo}}^{u_{hi}} [ad(u)^T + d(u)a^T + ud(u)d(u)^T] du. \quad (6)$$

The remaining integrand has degree at most five and is exactly integrated by three-point Gauss quadrature in real arithmetic. The linear statistic has a similarly polynomial remainder. This removes the difficult tail without a general analytic-moment implementation. Eight-point amplitude quadrature is retained on nonzero cells. This comparator was designed after seeing the initial failures; its results are a disclosed explanatory audit, not independent confirmation.

6 Equivalent synaptic kernels and quantized states

Let C be a positive diagonal capacity matrix and K a grounded symmetric positive-definite conductance matrix. A physical state obeys a stable linear diffusion equation. Its symmetric similarity transform has an orthonormal eigenbasis, giving independent exponential modes.

$$\dot{u} = -C^{-1}Ku + C^{-1}e_1 I(t), \quad S = -C^{-1/2}KC^{-1/2} = Q\Lambda Q^T, \quad (7)$$

$$z = Q^T C^{1/2}u, \quad \dot{z} = \Lambda z + Q^T C^{-1/2}e_1 I(t). \quad (8)$$

With readout from the first physical coordinate, the impulse-response residues are squares of the modal input coefficients and are nonnegative. Applying the exact similarity transform therefore changes coordinates but not the continuous input-output kernel. Quantization inserts a nonlinear operation into the state update. Rounding in modal coordinates and transforming back generally differs from rounding directly in physical coordinates; neither operation commutes with the similarity transform.

7 Experimental methods

The event panel uses six confirmation seeds, 915001 through 915006. Each seed generates 256 intervals at each mean interarrival-time to tau ratio in $\{0.02, 0.2, 2, 20\}$ and each spacing in $\{0.25, 1\}$, for 48 records. Tau equals one, initial trace equals 0.4, and impulses are uniform on $[0.05, 0.5]$. The interval-start teacher is $\tanh(2x)$ minus $0.3x$ plus Gaussian noise of standard deviation 0.02. Holding that noisy value across an interval is an experimental assumption, not a property of general sensors.

All arms use the same crossing partition, features, target, and ridge coefficient of 0.0001 times total duration. Controls use 4-, 8-, or 16-point Gauss quadrature, or event-aware midpoint time stepping with step/tau ratios 0.1, 0.01, or 0.001. Three randomized-order repetitions are summarized by their median. Timing

includes crossing discovery and accumulation but excludes synthetic generation and reference construction. High-order quadrature is checked against independent adaptive quadrature.

Accuracy requires relative G and b error at most 10^{-9} , prediction RMS discrepancy at most 10^{-7} on 257 domain queries, and no material negative Gram eigenvalue. The speed hypothesis requires at least a twofold median advantage over qualified quadrature in three of four rate groups, without a group slowdown beyond 1.25. The six seed groups, not the 48 related conditions, are the uncertainty units.

The realization identity panel has nine cases: dimensions 4, 8, and 12 under three seeds, with 256 signed impulses. The subsequent independent quantization panel fixes dimension eight, unit-time exact transitions, 2048 synapses, 2048 signed unit-impulse updates, and seeds 917101 through 917106. Six- and eight-bit stochastic rounding use scales set before sampling from each coordinate’s stationary variance. The quantizer step is eight stationary standard deviations divided by the maximum positive signed code. Both realizations use the same inputs and ideal unquantized modal reference.

The primary quantization metric is normalized output RMS distortion over the final 512 updates. The prospect criterion requires a modal/physical ratio at most 0.5 in at least five of six seeds at both precisions, no ratio above one, and clipping below 0.1%. Persistent states are int8 even for the six-bit arm. Shared coefficients, random-generator state, and float64 update temporaries are additional costs; no neuromorphic processor is used.

8 Results

Table 1: Event accumulation: all 48 records; timing is the overall median, not an accuracy-matched speed ratio.

Method	Accuracy passes	Time (ms)	Largest relative G error
Analytic	48/48	2.118	9.92×10^{-16}
Gauss 4	0/48	2.062	1.35×10^{-2}
Gauss 8	24/48	2.119	1.76×10^{-3}
Gauss 16	32/48	2.210	1.04×10^{-4}
Midpoint 0.1	0/48	12.194	2.24×10^{-1}
Midpoint 0.01	0/48	16.422	6.46×10^{-3}
Midpoint 0.001	5/48	55.658	7.80×10^{-5}

Analytic accumulation passes every accuracy check, with maximum prediction discrepancy 2.01×10^{-14} . Yet qualified quadrature/analytic time ratios are 0.965, 0.966, and 1.053 in the first three rate groups; only eight of twelve cases have a qualified quadrature control in the third group, and none do in the fourth. The predeclared speed criterion fails. A lack of a qualified comparator is not an infinite speedup.

The post-exposure constant-tail control passes all 48 cases. Its maximum relative Gram error is 6.97×10^{-13} and prediction discrepancy is 1.47×10^{-12} . Group timing ratios relative to analytic evaluation range from 0.980 to 1.001. It therefore explains the original quadrature failure without supporting a practical analytic-kernel speed advantage. Process peak RSS in the original panel is 117.3 MiB; dense statistics occupy 160 to 49,296 bytes depending on the visited domain.

Before quantization, the physical and modal implementations agree with relative output error at most 1.54×10^{-15} and state error at most 2.96×10^{-15} . After quantization, modal/physical distortion exceeds one in all twelve seed-by-precision cases. Median ratios are 1.403 at six bits and 1.750 at eight bits, with descriptive paired-seed bootstrap intervals [1.379,1.447] and [1.727,1.811]. Median normalized physical/modal errors are 0.767/1.084 and 0.257/0.453 respectively.

Clipping is rare: its maximum six-bit rate is 3.70×10^{-6} and the eight-bit rate is zero. The failure is therefore not explained by frequent saturation. Both implementations use 64 physical bits of int8 state per synapse, including the nominally 48-bit six-bit case. Shared numeric arrays occupy 640 bytes in physical coordinates and 192 bytes in modal coordinates, excluding the random generator. Thirty-six campaign tests pass; retained-state replay reproduces all twelve primary scores.

9 Discussion and reproducibility

The positive result is an accurate continuous-objective primitive under known dynamics and supplied supervision. The negative result is that a strong quadrature control achieves essentially the same computation




Figure 2: Archived accumulation accuracy and finite-precision realization results. Failed accuracy controls are retained. Modal-to-physical distortion greater than one favors the physical coordinates; the hypothesized modal advantage is absent.

cost. Similarly, diagonalization exposes an interpretable exponential bank but worsens this fixed quantization experiment. Neither outcome rules out other event-learning or finite-precision designs; neither licenses tuning the revealed panel into a favorable claim.

The supplementary package retains the protocols, full aggregate records, per-seed quantization metrics, and source fingerprints. It separates the independent panels from the post-exposure comparator. There is no learned recurrent task, delayed-reward policy, classification benchmark, energy measurement, or consolidation-capacity scaling result. Multi-trace interactions could also destroy the single-field Gram locality used here. The scope is a reproducible numerical-method study, including negative evidence about two proposed advantages.

10 Application boundary and research implication

The architecture supplies reusable objective accumulation under declared dynamics. Extending it to interacting traces changes both crossing complexity and feature overlap. Extending it to hardware requires measured quantized arithmetic, event traffic, and energy. The reported negative controls specify those gaps rather than obscuring the accurate integration result.

11 Conclusion

Known between-event dynamics allow spline readout objectives to be accumulated without a fine simulation clock. Once comparable structure is supplied to quadrature, however, the measured speed advantage disappears. Exact equivalence of diffusion and exponential coordinates likewise does not survive equal-budget rounding, and the modal coordinates are worse in the tested regime. Representation structure is useful, but its computational value must be established against implementations that exploit the same information.

References

- [1] J. Stapmanns et al. Event-Based Update of Synapses in Voltage-Based Learning Rules. *Frontiers in Neuroinformatics* 15, 609147, 2021. <https://doi.org/10.3389/fninf.2021.609147>
- [2] T. C. Wunderlich and C. Pehle. Event-Based Backpropagation can compute Exact Gradients for Spiking Neural Networks. *Scientific Reports* 11, 12829, 2021. <https://arxiv.org/abs/2009.08378>
- [3] M. K. Benna and S. Fusi. Computational principles of synaptic memory consolidation. *Nature Neuroscience* 19, 1697-1706, 2016. <https://doi.org/10.1038/nn.4401>