

Power-of-Two Leaky Integrate-and-Fire Neurons for Event-Driven Vision–Language–Action Inference

Rui Liu He Lv Ziyao Zhang Rui Ma
Rongxiang Nie*
Sivy Limited, Hangzhou, China
www.sivy.ltd

Abstract

Vision–language–action (VLA) models have demonstrated strong capabilities in robotic manipulation, yet their reliance on dense, continuous-valued computation presents a fundamental challenge for energy-constrained deployment. Inspired by the sparse, event-driven signaling and temporal integration mechanisms of biological neurons, a promising alternative is to introduce spiking computation into VLA models. However, enabling such computation without restructuring pretrained policies or compromising closed-loop control remains an open challenge. To address this challenge, we introduce the Power-of-Two Leaky Integrate-and-Fire (PoT-LIF) neuron, which emits signed spikes with power-of-two amplitudes. By combining magnitude-coded spike events with residual-preserving membrane dynamics, PoT-LIF enables continuous information to be progressively represented through discrete spikes. The Heartbeat formulation further localizes this process within weighted projections, bridging biologically inspired spiking computation and pretrained VLA architectures without modifying their external inference interfaces. We integrate PoT-LIF into the vision projection, vision–language projection, and action projection of $\pi_{0.5}$ and evaluate both operating modes over 2,000 paired closed-loop LIBERO episodes per configuration. Standard-mode PoT-LIF VLA achieves success rates of 96.90% and 96.35% at $T = 1$ and $T = 2$, respectively, while Heartbeat mode achieves 97.25% and 96.90% at $T = 2$ and $T = 4$, compared with 96.70% for the original continuous-activation $\pi_{0.5}$ baseline. At $T = 1$ with a 60.01% firing rate, PoT-LIF uses approximately 1/7.67 of the dense-reference arithmetic energy, corresponding to a modeled reduction of 86.95%.

Keywords: spiking neural networks, vision-language-action, robot learning, neuromorphic computing

1 Introduction

Vision–language–action (VLA) models couple pretrained vision–language backbones with action generation modules to enable closed-loop robotic manipulation in open environments [1]. Systems such as $\pi_{0.5}$ [2], Robotics Transformer 2 (RT-2) [3], and OpenVLA [4] have demonstrated strong cross-task generalization. However, VLA inference repeatedly executes dense weighted projections across vision, language, and action modules, resulting in high energy consumption on edge platforms. This motivates an alternative computational interface that preserves representational capacity while reducing arithmetic cost.

Spiking neural networks (SNNs) offer an event-driven computational paradigm [5]. Leaky integrate-and-fire (LIF) neurons are widely adopted in SNNs. They convert continuous membrane potentials into sparse binary spikes, allowing synaptic computation to occur only when

*Corresponding author: nierongxiang@sivy.ltd.

spikes are emitted. [6, 7]. However, a single binary spike does not explicitly encode the magnitude of the underlying membrane potential. To represent fine-grained continuous values, binary SNNs often accumulate information over multiple time steps, increasing computation and scheduling overhead. Recent studies, such as Spike-driven VLA and DS-VLA, have demonstrated the potential of SNNs for embodied control [8, 9]. Their integration strategies include directly trained spiking architectures and dendritic spiking action modules. However, these strategies can spread temporal processing across multiple modules, complicating integration into pretrained VLA architectures. These limitations motivate two key design goals for integrating spiking computation into pretrained VLA models: (1) representing continuous values with a limited number of events, and (2) confining temporal processing to selected projections without altering token layout, residual connections, or key-value (KV) cache organization.

To address the representation requirement, we draw inspiration from biological observations of exponential computation and graded signaling. Gabbiani et al. [10] reported a nearly exponential mapping from membrane potential to firing rate in a locust visual neuron, consistent with multiplication through logarithmic input encoding and exponentiation by active membrane conductances. In cortical circuits, presynaptic membrane potential can also modulate axonal action potentials and the synaptic responses they evoke [11]. Motivated by these mechanisms, we use an exponent-based event representation: the amplitude takes the form b^e where $b > 1$ and e is an integer. Successive exponent values specify multiplicative amplitude changes rather than equally spaced integer levels. We choose $b = 2$ to align this representation with binary digital arithmetic. Residual retention then allows successive events to refine the represented value under a bounded emission budget.

We therefore propose a power-of-two leaky integrate-and-fire (PoT-LIF) neuron. A signed event has the form $s_t = \text{sgn}(u_t)2^{e_t}$, where $\text{sgn}(u_t) \in \{-1, +1\}$ and e_t is an integer. For a floating-point weight w , an event contribution is evaluated through scaling-accumulation (SAC), formally defined in Sec. 3.1. For normal values whose scaled results remain normal and representable, scaling adjusts the weight exponent and combines the signs without a general significand multiplication. A residual membrane state retains the portion not represented by the current event, making the event budget T an explicit control over precision and computational work.

The same PoT-LIF recurrence supports two operating modes: standard mode and Heartbeat. In standard mode, the neuron receives successive inputs and retains its membrane state across input calls. Heartbeat instead applies the recurrence locally within each VLA projection: each projection invocation injects its input once and accumulates the weighted contributions of a bounded event sequence. The residual is retained across internal extraction steps, and the local state is reset at the end of the invocation. This confines the event-time dimension to individual projections while preserving the original token count, residual connections, and KV caches. Standard mode remains available for processing successive inputs and may support future integration with event-based sensors.

We evaluate PoT-LIF in the vision projection, vision-language projection, and action projection of $\pi_{0.5}$ using a fixed paired closed-loop protocol on LIBERO [12]. Across both the main and supplementary evaluations, all tested PoT-LIF configurations of $\pi_{0.5}$ differ from the dense $\pi_{0.5}$ baseline by no more than 0.80 percentage points in success rate. We also estimate the arithmetic energy of PoT-LIF projection computation under a matched FP32 accumulation model, while excluding system-level overheads.

This paper makes three contributions:

- **A unified PoT-LIF formulation with reconstruction guarantees.** We define a unified PoT-LIF recurrence with signed power-of-two emissions and residual retention, describe its standard mode, and formulate Heartbeat as a specific input schedule. We establish residual reconstruction guarantees under leak-free dynamics and exact arithmetic, with the geometric error bound requiring inactive exponent clipping (Sec. 3).
- **PoT-LIF integration into $\pi_{0.5}$ with preserved closed-loop performance.** We integrate PoT-LIF into the vision projection, vision–language projection, and action projection of $\pi_{0.5}$ and evaluate standard mode followed by Heartbeat mode. Across both the main and supplementary evaluations, all tested PoT-LIF configurations differ from the 96.70% dense $\pi_{0.5}$ baseline by no more than 0.80 percentage points in success rate (Sec. 5; paired statistical analyses are provided in Appendix G).
- **Operation-level energy model showing a 7.67-fold arithmetic-energy advantage.** We formulate operation-level energy accounting for multiply–accumulate (MAC) and scaling–accumulation (SAC), the computation adapted to PoT-LIF and formally defined in Sec. 3.1. With matched FP32 accumulation and a 60.01% firing rate at $T = 1$, the modeled PoT-LIF arithmetic energy is approximately reach to 86.95% reduction, excluding system-level overheads.

2 Background and Related Work

In this section, we first review the biological and computational foundations of LIF neurons, including membrane dynamics and surrogate-gradient training. We then survey related work on event representations, efficient arithmetic, and spiking VLA models.

2.1 Spiking neural networks and LIF dynamics

SNNs compute with discrete events rather than dense continuous-valued activations. A leaky integrate-and-fire (LIF) neuron integrates input into its membrane potential, exhibits passive leakage in the absence of input, emits a spike at threshold, and resets after emission ([5], [6], [7]). The LIF dynamics can be written as

$$\tau_m \frac{du}{dt} = -u(t) + R_m I_{\text{syn}}(t). \quad (1)$$

Where $u(t)$ is the membrane potential, τ_m is the membrane time constant, R_m is the membrane resistance, and $I_{\text{syn}}(t)$ is the synaptic current.

At discrete step t , let $\theta > 0$ be the firing threshold. The binary spike occurs when $u_t > \theta$, we denote the spike by adapting the Heaviside function

$$z(t) = H(u_t - \theta) = \begin{cases} 1, & u_t \geq \theta, \\ 0, & u_t < \theta. \end{cases} \quad (2)$$

Assume that r_t denotes the membrane state retained after emission and reset. With zero resting potential, the hard and soft reset rules are given, respectively, by

$$r_t = (1 - z_t)u_t \quad (\text{hard reset}), \quad r_t = u_t - \theta z_t \quad (\text{soft reset}). \quad (3)$$

An exponentially decaying synaptic current $I_{\text{syn}}(t)$ and the passive decay of the membrane potential in the absence of input are described by the following expressions, respectively [13, 14]:

$$I_{\text{syn}}(t) = A \exp(-t/\tau_s), \quad t \geq 0, \quad u(t + \Delta t) = \exp(-\Delta t/\tau_m)u(t), \quad (4)$$

here, A is the synaptic current amplitude and τ_s and τ_m are the synaptic and membrane time constants. Discretizing this decay gives $\lambda = \exp(-\Delta t/\tau_m)$, where Δt is the discrete time step.

A binary spike indicates whether a neuron fires; it does not explicitly encode the membrane-potential magnitude. Consequently, SNNs represent continuous quantities through rate coding, in which firing frequency carries the value, or temporal coding, in which spike timing carries the value [5, 15, 16]. Such representations generally require multiple time steps.

2.2 Surrogate-gradient training

Binary spike generation is modeled by the Heaviside step function, whose derivative is zero almost everywhere and undefined at the threshold. Consequently, direct backpropagation provides no useful gradients for learning. Surrogate-gradient methods enable gradient-based training through discontinuous spike-generation functions [17]. The forward pass retains the binary spike, while the backward pass uses the derivative of a smooth substitute in place of the step-function derivative.

In this paper, PoT-LIF changes the event amplitude and residual retention; its discontinuous event-generation rule can still be trained using the same surrogate-gradient principle. Specifically, the forward pass retains the actual signed power-of-two events, while the backward pass approximates the derivatives of the discontinuous operations. The training configuration used in this work is described in Sec. 5.

2.3 Event representations and efficient arithmetic

Biological neurons exhibit a nearly exponential relationship between membrane potential and firing rate. Presynaptic membrane potential also modulates axonal action potentials and postsynaptic response amplitudes. These observations motivate introducing exponentially spaced firing levels and amplitude-bearing events into LIF dynamics. Theoretical analyses of neural coding under metabolic constraints [18] and experiments on energy-efficient synaptic transmission [19] further motivate considering the information conveyed by each event alongside its computational cost.

Logarithmic representations and DeepShift show how powers of two replace general multiplication with operations suited to binary hardware [20, 21]; AdderNet provides a complementary example of reducing reliance on multiplication through addition-based distance operations [22]. Combining this arithmetic motivation with the biological observations leads to PoT-LIF: exponentially spaced levels follow b^e , where $b > 1$ and e is an integer, and choosing $b = 2$ aligns the event amplitudes with binary digital arithmetic. The resulting events carry a sign and an integer exponent, while the membrane retains the unrepresented residual.

Neuromorphic hardware also supports carrying amplitude information in events. Intel’s Loihi 2 supports graded spikes with integer-valued payloads of up to 32 bits; its technology brief reports that the commonly used 8-bit payload entails negligible extra energy cost relative to binary spike processing [23]. This provides engineering evidence that amplitude-bearing communication can preserve the energy advantages of sparse events, although the cost of a particular encoding depends on its implementation. Together, biological exponential responses, power-of-two arithmetic, and hardware support for graded events motivate the PoT-LIF representation.

2.4 Spiking VLA models

Wang et al. [8] introduce Spike-driven VLA, which directly trains spiking perception, fusion, and action modules. Its 0.15B-parameter model achieves only 92.4% mean LIBERO success, leaving task success as a key limitation.

Lyu et al. [9] introduce DS-VLA, which combines dendritic spiking dynamics with inhibitory gating for robust action control. Although it improves resistance to action perturbations, its nominal LIBERO success rate remains only 91.6%.

We identify limited amplitude information in binary spikes as a key representational bottleneck, which limits success rates in spiking VLA control. This motivates PoT-LIF: signed power-of-two events convey magnitude information at low arithmetic cost, while residual retention refines the representation over successive events. We aim to improve end-to-end task success without sacrificing the arithmetic efficiency of event-based computation.

3 Power-of-Two Leaky Integrate-and-Fire Neurons

In this section, we define the PoT-LIF dynamics and SAC arithmetic. We then introduce the Heartbeat operating mode. Finally, we prove that under leak-free dynamics, Heartbeat guarantees geometric error reduction, exact recovery at zero residual, and reconstruction error bounded by upon threshold stopping.

3.1 PoT-LIF event encoding and SAC

PoT-LIF extends the discrete LIF recurrence in Sec. 2.1 by replacing the binary spike with a signed power-of-two event and retaining the residual.

At each internal time step t , the pre-emission membrane potential u_t is obtained from the input contribution I_t and the retained residual r_{t-1} :

$$u_t = I_t + \lambda r_{t-1}, \quad (5)$$

where λ is the membrane decay factor defined in Sec. 2.1. In our implementation, the threshold is chosen on the event ladder:

$$\theta = 2^h, \quad h \in \mathbb{Z}, \quad e_{\min} \leq h \leq e_{\max}. \quad (6)$$

The full ladder contains $e_{\max} - e_{\min} + 1 = 33$ positive levels, or 66 signed nonzero levels.

Each nonzero event in PoTLIF is encoded as

$$s_t = \begin{cases} \text{sgn}(u_t) 2^{\lfloor \log_2 |u_t| \rfloor}, & |u_t| \geq \theta \text{ and } u_t \neq 0, \\ 0, & |u_t| < \theta \text{ or } u_t = 0, \end{cases} \quad (7)$$

On emitting steps, the ideal exponent is denoted by $e_t^{\text{ideal}} = \lfloor \log_2 |u_t| \rfloor$. In the finite-range implementation, the emitted exponent is clipped to

$$e_t = \begin{cases} e_{\min}, & e_t^{\text{ideal}} < e_{\min}, \\ e_t^{\text{ideal}}, & e_{\min} \leq e_t^{\text{ideal}} \leq e_{\max}, \\ e_{\max}, & e_t^{\text{ideal}} > e_{\max}, \end{cases} \quad (8)$$

in this way, $e_t \in \mathbb{Z}$, and the exponent need not be evaluated when $s_t = 0$. Thus, exponents within the prescribed window remain unchanged, while those outside it are saturated to the

corresponding boundary. The exponent is evaluated only when an event is emitted. Each nonzero event therefore has a magnitude in the discrete set

$$|s_t| \in \{2^e : e \in \mathbb{Z}, e_{\min} \leq e \leq e_{\max}\}. \quad (9)$$

The exponent is carried in a signed 6-bit field, restricted to $[-16, 16]$, while the event sign is determined by $\text{sgn}(u_t)$ and stored in a separate sign bit. This window specifies the event format rather than the weight format.

After emission, the remaining membrane potential is retained as

$$r_t = u_t - s_t. \quad (10)$$

Hence, in exact arithmetic,

$$s_t + r_t = u_t = I_t + \lambda r_{t-1}, \quad (11)$$

the residual r_t is carried forward to the next internal time step. In this way, the pre-emission membrane potential is decomposed into the emitted event and the retained residual. The setting $\lambda = 1$ is the leak-free special case used by our heartbeat configuration.

Let a denote an accumulator, w a floating-point weight, and s_t an event emitted by PoT-LIF. In exact arithmetic, given an event s_t , the scaling-accumulation map (SAC) is defined as

$$\text{SAC}(a, w, s_t) := a + ws_t = \begin{cases} a + \text{sgn}(u_t) \text{Scale}_2(w, e_t), & s_t \neq 0, \\ a, & s_t = 0, \end{cases} \quad (12)$$

where $\text{Scale}_2(w, e_t) := 2^{e_t}w$. The corresponding accumulator update is

$$a \leftarrow \text{SAC}(a, w, s_t). \quad (13)$$

Thus, SAC accumulates weighted event contributions, while the event amplitudes are determined by the PoT-LIF emission rule.

For a receiving neuron j , the accumulator a_j is initialized to zero at each internal time step t . Applying SAC to the presynaptic events $s_{t,i}$ with connection weights w_{ji} produces the input contribution

$$I_{t,j} = a_j = \sum_i \text{SAC}(a, w_{ji}, s_{t,i}). \quad (14)$$

The accumulated input is not restricted to power-of-two values. The receiving neuron then follows the same recurrence:

$$u_{t,j} = I_{t,j} + \lambda r_{t-1,j}, \quad r_{t,j} = u_{t,j} - s_{t,j}. \quad (15)$$

3.2 Heartbeat input schedule

Heartbeat is an input schedule that applies the same recurrence to a single input within a weighted projection. At the start of each projection invocation, the local state is initialized to $r_0 = 0$, and the input is injected once:

$$r_0 = 0, \quad I_t = \begin{cases} x, & t = 1, \\ 0, & t > 1. \end{cases} \quad (16)$$

The neuron then executes at most T event-extraction steps. The resulting approximation is

$$q_T(x) = \sum_{t=1}^T s_t, \quad (17)$$

At each tick, the pre-emission membrane potential, event, and residual are computed using equations (7) and (10). A local accumulator, initialized to zero, is updated by $q \leftarrow q + s_t$ whenever $s_t \neq 0$. Execution stops when the residual becomes zero, when the emission gate produces no event, or when the tick budget T is exhausted.

For a vector input, the recurrence is applied elementwise. With projection weight matrix W , the event contributions are accumulated before the next nonlinearity:

$$Wq_T = \sum_{t=1}^T Ws_t. \quad (18)$$

The projection bias, if present, is added once. The event-time dimension therefore remains internal to the projection. In our configuration, the decay rate is set to $\lambda = 1$, so a residual below the threshold cannot produce a later event in the absence of further input and extraction can stop immediately.

Figure 1 illustrates the two operating modes. In Fig. 1a, presynaptic events carrying signed power-of-two amplitudes arrive through weighted synapses. The neuron integrates these contributions, emits events according to Eq. (7), and retains the residual for subsequent inputs. In Fig. 1b, Heartbeat injects a scalar input once and extracts its event representation over at most T internal steps. With $\lambda = 1$, the residual is retained without leakage between extraction steps and is reset at the end of the projection invocation.




Figure 1: PoT-LIF dynamics and power-of-two event encoding. (a) Standard mode: presynaptic events arrive through weighted synapses and are integrated into the membrane state. Each emitted event carries a sign and an integer exponent specifying its amplitude. The residual remains in the membrane state across successive input calls. (b) Heartbeat mode: a scalar input is injected once and decomposed into a bounded event sequence using the same recurrence with $\lambda = 1$. The residual is retained between internal extraction steps, and the local state is reset at the end of the projection invocation.

3.3 Reconstruction analysis

In the reconstruction analysis, r_T denotes the residual before the local state is reset. If extraction terminates at step $K < T$ under $\lambda = 1$, we set $s_t = 0$ and $r_t = r_K$ for the remaining analytical steps $t = K + 1, \dots, T$.

Proposition 1 (Residual reconstruction for $\lambda = 1$). *Let $x \in \mathbb{R}$, the initialization and input schedule is given in Eq. (16). In exact arithmetic, assume that $\lambda = 1$, the exponent window $[e_{\min}, e_{\max}]$, and the threshold $\theta = 2^h$, where $h \in \mathbb{Z}$ and $e_{\min} \leq h \leq e_{\max}$, the unclipped emission rule in (7). Let $K \leq T$ denote the number of executed ticks, with r_K evaluated before the state reset. Then the input x can be reconstructed:*

$$x = q_T(x) + r_K. \quad (19)$$

In particular:

(i) If an event is emitted at each of the T steps and the residual remains nonzero, then

$$|x - q_T(x)| \leq 2^{-T}|x|. \quad (20)$$

(ii) If the residual reaches zero at some step $t_0 \leq T$, then

$$q_T(x) = x. \quad (21)$$

(iii) If extraction terminates since the remaining residual is below the threshold, then

$$|x - q_T(x)| < \theta. \quad (22)$$

In case (i), the input range condition and $\theta \geq 2^{e_{\min}}$ prevents upper and lower exponent saturation on emitting ticks, respectively. Each nonzero residual is therefore reduced by at least a factor of two. Upper exponent saturation can invalidate this geometric contraction, but it does not invalidate Eq. (19) or cases (ii) and (iii). The proof is given in Appendix A.

For $\lambda < 1$, any nonzero retained state is attenuated before subsequent emissions. If $K \leq T$ ticks are executed, summing the recurrence gives

$$x - q_T(x) = r_K + (1 - \lambda) \sum_{t=1}^{K-1} r_t. \quad (23)$$

The reconstruction error therefore includes both the terminal residual and accumulated leakage. Consequently, $q_T(x)$ generally does not reconstruct the full input, and the reconstruction bound above does not apply in this form.

Our Heartbeat configuration fixes $\lambda = 1$ and set the minimum supported event magnitude as $\theta = 2^{e_{\min}} = 2^{-16}$. This threshold prevents emission when the residual is smaller than the encoding can represent. If a lower threshold were activated, clipping the exponent to $e_{\min}$ could produce an event larger than the residual and reverse its sign.

When upper clipping is inactive, the decomposition successively extracts the nonzero power-of-two components of the input down to the minimum supported amplitude (Sec. 3.1). Exact reconstruction is obtained when all nonzero components lie within the supported exponent range and T is sufficient to extract them. Otherwise, extraction either reaches the budget or stops with a residual smaller than θ . For example, $x = 1.4375$ produces the event sequence (1, 0.25, 0.125, 0.0625). Two events reconstruct 1.25, whereas four reconstruct the input exactly.

The worst-case bound permits a relative error of up to 25% at $T = 2$. For positive inputs, this limit is approached as x approaches 2^{e+1} from below, provided clipping and the threshold do not interfere with extraction. For example, $x = 1.99$ leaves a residual of 0.49, corresponding to a relative error of approximately 24.6%. The rule truncates rather than rounds: the residual contains the binary components below the second extracted nonzero bit. It vanishes when the input has at most two nonzero binary components and both lie within the supported range.

For the significands which distribute over $[1, 2)$ uniformly, the mean relative error is approximately 9.9% at $T = 2$, with a 95th percentile of 23.1%. At $T = 4$, these values decrease to 1.1% and 4.0%, respectively. These statistics describe the specified significand distribution rather than the activation distribution of the VLA model.

4 PoT-LIF Integration into Pretrained Vision–Language–Action Models

We apply the recurrence defined in Sec. 3 to the vision encoder, vision–language backbone, and flow-matching action expert of pretrained VLA models. In this work, we choose $\pi_{0.5}$ [1, 2] as the foundation model. We first evaluate standard mode, in which the event-time dimension extends across layers within a bounded network-level sequence, and then Heartbeat, which confines event generation and weighted accumulation to individual projections. Each selected projection retains its original input and output shapes, preserving token organization, residual connections, KV-cache structure, and the action-generation interface.

4.1 Projection-level event dataflow

Figure 2 shows the projection sites in the three principal components of $\pi_{0.5}$. Figure 2(a) follows the visual inputs through the Sigmoid Loss for Language–Image Pre-training (SigLIP) vision encoder [24] and PaliGemma vision–language model (VLM) [25] to the action expert. Figure 2(b) illustrates how a PoT-LIF encoder supplies events to each selected weighted projection.

We incorporate PoT-LIF into 26 specific projection sites across six functional classes: patch embedding, query/key/value (Q/K/V), attention output, gated feed-forward layers, expert projections, and expert interfaces.

The standard-mode configurations use $T = 1$ or $T = 2$ internal time steps across the network for each image-and-prompt input. Membrane states are reset when that input is supplied and retained between internal steps as events propagate across layers. The T output states are averaged at the velocity-field readout. State does not carry into the next image-and-prompt input.

In Heartbeat mode, Q/K/V, attention-output, and feed-forward projections accumulate the weighted contributions of a bounded event sequence. Attention-score computation, softmax, normalization, and residual addition retain their original operations. Event contributions are combined before the next nonlinearity, so the internal extraction steps do not expand the token sequence or replicate the KV cache. The projection output has the same shape as in the continuous-activation policy.

For consistency with Sec. 3, we write a projection using column-vector notation, $y = Wx$, with bias omitted. Heartbeat produces the approximation $\hat{y} = Wq_T(x)$ by accumulating at most T events per input element:

$$\hat{y} = \sum_{t=1}^T W s_t, \quad W s_t = \sum_{i:s_{t,i} \neq 0} \sigma_{t,i} \text{Scale}_2(W_{:,i}, e_{t,i}). \quad (24)$$

Here, $W_{:,i}$ is the weight column associated with input element i , and Scale_2 applies power-of-two scaling to its floating-point entries. The inner sum includes only nonzero events. For normal, representable results, scaling adjusts the weight exponent and combines the event and weight signs without a general significand multiplication.

Under exact arithmetic, the projection error is $\hat{y} - y = W(q_T(x) - x)$. Thus, the residual bounds in Sec. 3 characterize the input reconstruction error; its effect on the projection also depends on W . Finite-precision scaling and accumulation can introduce additional numerical error. SAC requires event-exponent decoding, exponent addition, and sign handling before accumulation.

At $T = 1$, standard mode and Heartbeat are equivalent under the evaluated design: each projection performs a single emission step per input element from a zero initial state, with no subsequent temporal update. We report this shared configuration as standard mode at $T = 1$. At $T = 2$,

standard mode retains a network-wide temporal sequence, whereas Heartbeat performs both extraction steps locally within each projection.

This scheduling distinction also affects weight movement. On conventional hardware that reloads projection weights for each network pass, standard mode at $T = 2$ requires two rounds of weight movement, in addition to repeated network execution and membrane-state access. Replacing multiplication with SAC therefore does not by itself establish an overall energy advantage for this implementation. This limitation motivates Heartbeat: processing consecutive events within a single projection allows a suitable datapath to reuse resident weight tiles across extraction steps before advancing to the next layer.




Figure 2: (a) PoT-LIF integration into $\pi_{0.5}$. Heartbeat encoders supply signed power-of-two events to selected weighted projections in the vision encoder, vision–language backbone, and action expert. Event contributions are evaluated through scaling–accumulation (SAC) within each projection. Attention computation, the blockwise mask, KV-cache structure, residual paths, and action interfaces retain their original organization. Dashed boxes indicate the three module groups; (b) shows the event–projection pairs within a block. Abbreviations: RGB, red–green–blue; ViT, vision transformer; LM, language model; FFN, feed-forward network; RMSNorm, root mean square normalization; adaRMSNorm, adaptive RMSNorm; GELU, Gaussian error linear unit; bbox, bounding box; proprio., proprioceptive state; MatMul, matrix multiplication.

5 Robot Experiments

We initialize the PoT-LIF policies from pretrained $\pi_{0.5}$ and fine-tune them for 5,000 optimizer updates with an effective batch size of 256. The main configurations use the symmetric emission threshold $\theta = 2^{-16}$ and $\lambda = 1$. The forward pass uses signed power-of-two events, while the backward pass uses a straight-through estimator (STE), which substitutes a prescribed gradient for the discontinuous event-generation rule [17]. The matched-budget continuous control receives the same fine-tuning budget. Detailed training settings are provided in Appendix E.

5.1 Protocol

We evaluate on the Spatial, Object, Goal, and LIBERO-10 (Long-Horizon) suites [12], with ten tasks and 500 closed-loop episodes per suite, giving 2,000 episodes per configuration. Conditions share initial states, paired inference noise, and the action interface (evaluation seed 7). Two RGB views, from the frontal base and left wrist cameras, are resized to 224×224 ; a third wrist slot is inactive. The policy predicts 10-step action chunks and replans every five environment steps. Actions between replanning points are executed open loop. PoT-LIF internal steps, flow-solver steps, and environment steps are distinct time axes.

5.2 Main results

Standard mode achieves 96.90% overall success at $T = 1$ and 96.35% at $T = 2$, compared with 96.70% for the continuous-activation control. Heartbeat achieves 97.25% at $T = 2$ and 96.90% at $T = 4$ (Table 1). Success rates remain essentially unchanged under the evaluated protocol; paired significance tests and detailed results are reported in Appendix G.

Table 1: Main LIBERO results (success rate per suite over 500 episodes; 2,000 per configuration). Δ is the overall percentage-point difference from the continuous-activation control.

Configuration	Spatial	Object	Goal	Long-H.	Overall	Δ
Continuous activation	98.8	98.6	97.0	92.4	96.70	–
Continuous, matched budget	98.0	97.6	97.8	91.6	96.25	–0.45
Standard mode $T = 1$	99.2	98.4	97.6	92.4	96.90	+0.20
Standard mode $T = 2$	98.2	97.8	97.4	92.0	96.35	–0.35
Heartbeat $T = 2$	99.2	97.4	98.2	94.2	97.25	+0.55
Heartbeat $T = 4$	99.2	98.0	97.8	92.6	96.90	+0.20

Performance varies across suites, with no single configuration achieving the highest success rate on all four. The four event-based configurations achieve 97.4–98.4% success on Object, compared with 98.6% for the continuous-activation control. On Long-Horizon, Heartbeat at $T = 2$ achieves 94.2%, compared with 92.4% for the control. Additional ablations of the architecture and training procedure are reported in Appendix C.

5.3 Event statistics and the choice of T

Let n_l be the number of input elements processed at projection site l , a_l the number of nonzero events emitted there, and d_l its output width. The product-weighted event ratio is $\mu = (\sum_l d_l a_l) / (\sum_l d_l n_l)$. Weighting by d_l counts the event–weight pairs across all outputs, including fused QKV projections. For the sampled workload, $\mu_2 = 1.900$ and $\mu_4 = 3.480$. The projection-activity accounting covers 26 instrumented sites under the eight-observation sampling protocol described below. Its dense reference count is $N = \sum_l d_l n_l = 2.271 \times 10^{12}$ MACs, counting one weight–input product and its accumulation for each input–output connection. This count covers the instrumented weighted projections, not all operations in a complete end-to-end policy inference. Heartbeat with $T = 2$ emits 1.240×10^9 events, corresponding to 4.314×10^{12} event–weight pairs; $T = 4$ emits 2.270×10^9 events, corresponding to 7.902×10^{12} pairs. Each nonzero event–weight pair contributes one SAC operation.

We use $T = 2$ as the default Heartbeat setting because it combines 97.25% success with fewer event–weight operations than $T = 4$. The latter requires approximately 1.83 times as many SAC

operations, while scoring 96.90%. Operation counts use eight fixed observations (one environment observation and seven demonstration frames) over the 26 projection sites.

6 Operation-Level Energy Accounting

A reference product is one dense weight–input product; its event-coded replacement may require multiple event–weight products. We compare the arithmetic energy of dense FP32 MAC computation with PoT-LIF SAC computation using matched FP32 accumulation. The comparison uses an illustrative 45-nm arithmetic model and includes only multiplication or exponent scaling and accumulation. Data movement, memory access, control logic, event encoding, and other system-level costs are excluded.

The dense reference uses 4.6 pJ per FP32 MAC [16, 26]. For SAC, we use 0.9 pJ for FP32 addition and a 0.1 pJ integer-addition proxy for exponent addition, giving an illustrative arithmetic cost of 1.0 pJ per event–weight pair. The exponent-addition term uses a 32-bit integer-addition reference as a proxy. The floating-point addition cost represents the accumulation arithmetic, including its associated alignment, normalization, and rounding.

The 2.271 trillion reference MACs come from the 26-site projection-activity accounting in Sec. 5.3, based on eight fixed observations (one environment observation and seven demonstration frames). They count the covered weight–input products, excluding attention-score products, softmax, normalization, and other operations outside those projections. Multiplying this projection workload by 4.6 pJ per MAC gives an estimated dense arithmetic energy of 10.45 J. For the new $T = 1$ setting, using the reported 60.01% firing rate as the effective event–weight activity fraction gives $\mu_1 = 0.6001$. Applied to the same reference workload, this implies approximately 1.363 trillion SAC operations and 1.36 J, corresponding to an 86.95% reduction. The sampled $T = 2$ count of 4.314 trillion SAC operations gives 4.31 J, corresponding to a 58.7% reduction. For $T = 4$, 7.902 trillion SAC operations give 7.90 J, corresponding to a 24.4% reduction.

The arithmetic-energy advantage depends on both the SAC unit cost and the number of events. With $\mu_1 = 0.6001$ at $T = 1$ and an assumed SAC arithmetic cost of 1.0 pJ, the reference-to-SAC energy ratio is $4.6 / (0.6001 \times 1.0) \approx 7.67$, corresponding to an 86.95% energy saving. Increasing the assumed SAC cost to 1.1 or 1.2 pJ gives 0.66011 or 0.72012 pJ per reference product, reducing the estimated energy saving to 85.65% or 84.35%, respectively.

For f reference products per second, modeled arithmetic power is $4.6f$ pW for FP32 MAC and $0.6001f$ pW for PoT-LIF ($T = 1$, firing rate 60.01%).




Figure 3: Average arithmetic energy per reference product with matched FP32 accumulation. One FP32 MAC costs 4.6 pJ. At $T = 1$ and a 60.01% firing rate, the PoT-LIF SAC path averages 0.6001 pJ: 0.06001 pJ for exponent scaling and 0.54009 pJ for accumulation. The reference-to-SAC energy ratio is 7.67, corresponding to a modeled reduction of 86.95%. Other power terms are excluded.

We estimate the energy of PoT-LIF relative to a dense FP8 E4M3 baseline using matched FP8

weights and FP32 accumulation, with a 45 nm reference cost model augmented by engineering proxies for exponent and control operations. For $T = 1$ and $T = 2$ ticks with exponent-window half-width $k = 4$, sampled fanout-weighted activity yields 0.6001 and 1.0565 SAC contributions per original dense projection product, respectively. The modeled projection-arithmetic cost consequently decreases from 1.4646 pJ to 0.5926 pJ and 1.0432 pJ, corresponding to reductions of 59.54% and 28.77%. Including neuron encoding and the other modeled computation gives compute-only savings of 56.25% and 27.16%. Under an identical fixed data-movement budget of 2.518 J per policy inference (ten denoising steps), the modeled dynamic-energy subtotal decreases from 6.038 J to 4.058 J and 5.082 J, yielding savings of 32.79% and 15.83%.




Figure 4: Modeled projection arithmetic energy relative to an FP8 E4M3 baseline. Both methods use FP8 weights and FP32 accumulation. Energy is normalized per original dense projection product using sampled fanout-weighted activity: 0.6001 and 1.0565 SAC contributions for $T = 1$ and $T = 2$, respectively, with $k = 4$. Green denotes multiplication or exponent scaling and control; gray denotes FP32 accumulation. PoT-LIF reduces modeled projection arithmetic energy by 59.54% and 28.77%. These estimates conditionally transfer BF16-derived activity to FP8 weights and exclude neuron encoding, shared computation, and data movement.




Figure 5: Modeled dynamic energy per policy inference with a matched fixed data-movement budget. Each inference includes ten denoising steps; all configurations use FP8 weights and FP32 accumulation. Green denotes projection arithmetic, gold denotes neuron encoding and other modeled computation, and light gray denotes the identical 2.518 J data-movement budget. With $k = 4$, the modeled subtotal decreases from 6.038 J to 4.058 J for $T = 1$ and 5.082 J for $T = 2$, corresponding to savings of 32.79% and 15.83%. Results are conditional on BF16-derived activity and represent partial energy estimates, not hardware measurements or complete system energy.

7 Discussion

As demonstrated, PoT-LIF combines exponent-based event amplitudes with residual membrane dynamics. Standard mode and Heartbeat use the same recurrence but organize its execution differently: standard mode retains temporal processing across layers, while Heartbeat confines event generation and accumulation to individual projections. This local schedule preserves the pretrained policy’s external token and cache organization and allows the event budget to be set within each projection.

The two schedules offer different paths toward deployment. For the current frame-based VLA workload, the repeated weight movement of standard mode at $T = 2$ adds to its energy cost on conventional hardware. Reducing data-movement costs would allow a larger share of the arithmetic savings from power-of-two event coding to translate into end-to-end energy savings. Near-memory and compute-in-memory implementations that support resident weights, local membrane state, and efficient event accumulation offer a route to reducing the repeated weight-access cost of standard mode. Future work will co-design standard-mode event scheduling, membrane-state placement, and weight reuse for these architectures. Standard mode also retains the potential to process asynchronous event data because membrane state evolves with successive inputs.

The reconstruction guarantees in Sec. 3.3 characterize the fidelity of the event representation, whereas closed-loop success also depends on how projection errors propagate through the policy and affect action generation. A larger event budget can therefore improve scalar reconstruction under the stated conditions but need not yield higher task success. In our experiments, increasing the Heartbeat budget from $T = 2$ to $T = 4$ increases the number of event-weight operations without improving the observed success rate. This supports selecting the event budget jointly from task performance and computational work, rather than from reconstruction accuracy alone.

Finally, regarding arithmetic efficiency, under the per-product accounting used here, the $T = 1$ SAC path at a 60.01% firing rate averages 0.6001 pJ per reference product. Its modeled arithmetic energy is approximately reach to an 86.95% reduction.

8 Conclusion

This work introduces PoT-LIF and successfully integrated its signed power-of-two events and residual-preserving dynamics into pretrained $\pi_{0.5}$. The proposed Heartbeat mode achieves 97.25% LIBERO success compared with 96.70% for the original policy, preserving task performance while replacing dense projection multiplication with event-based scaling and accumulation. Under the matched operation-level energy model, the $T = 1$ path at a 60.01% firing rate uses approximately 1/7.67 of the dense-reference arithmetic energy, corresponding to an 86.95% reduction.

PoT-LIF is broadly applicable beyond $\pi_{0.5}$ to VLA models with compatible weighted-projection architectures. By localizing event processing within projections, Heartbeat preserves token layout, residual connections, KV-cache organization, and the action interface. This offers a route to easing the energy demands of edge VLA deployment while maintaining task success, as demonstrated here on $\pi_{0.5}$. In conclusion, PoT-LIF enables energy-efficient VLA inference by preserving closed-loop task performance while reducing arithmetic costs, paving the way for more efficient and sustainable embodied AI.

Acknowledgments

We thank the open-source robot-learning and neuromorphic-computing communities for the software, benchmarks, and published research that enabled this study.

Appendix

A Proof of Proposition 1

Proof. Under the assumptions of Proposition 1, $r_0 = 0$, $I_1 = x$, $\lambda = 1$, and $I_t = 0$ for $t > 1$. Thus $u_1 = x$, while $u_t = r_{t-1}$ for $t > 1$. Summing $r_t = I_t + r_{t-1} - s_t$ gives

$$r_T = x - q_T(x).$$

When an event fires under the unclipped rule,

$$|s_t| = 2^{\lfloor \log_2 |u_t| \rfloor}, \quad |s_t| \leq |u_t| < 2|s_t|.$$

Because the event has the same sign as the potential,

$$|r_t| = |u_t| - |s_t| < \frac{1}{2}|u_t|.$$

(i) If all T steps emit and the residual remains nonzero, repeated application gives $|x - q_T(x)| = |r_T| \leq 2^{-T}|x|$. (ii) If $r_{t_0} = 0$ for some $t_0 \leq T$, all subsequent events are zero and $q_T(x) = x$. (iii) If extraction stops because $|u_t| < \theta$, the residual is held fixed for the remaining analytical steps, so $|x - q_T(x)| = |r_T| < \theta$. For $\lambda < 1$, summing over the K executed ticks gives Eq. (23); the accumulated leakage prevents the same residual-only reconstruction bound. The halving argument requires inactive exponent clipping, as specified in Sec. 3.3. $\square$

B Configurable Response: Definition and Fits

Let $\{\theta_k\}$ be ordered thresholds and c_k the increment associated with threshold k , represented by a bounded sum of signed power-of-two terms. The output response is

$$f_{\Theta, \mathcal{C}}(x) = c_0 + \sum_k c_k \mathbf{1}[x \geq \theta_k], \quad c_k = \sum_{j=1}^{B_k} b_{kj} 2^{e_{kj}}. \quad (25)$$

Here, c_0 is the response for $x < \theta_1$, $b_{kj} \in \{-1, +1\}$, e_{kj} is an integer exponent, and B_k is the local pulse budget. Threshold positions set the segment boundaries, and the pulse codes separately set the increments; c_k is not constrained to equal θ_k . These parameters can be fixed to construct a response or optimized using a straight-through estimator (STE), which replaces a discontinuous forward derivative with a prescribed backward value, [17]. The VLA experiments use the greedy extraction rule of Sec. 3.1 and retain the original network nonlinearities; they do not use this threshold table. The budget B_k per threshold increment is distinct from the event budget T per projection input. Figure 6 in Sec. ?? reports fits to three reference nonlinearities using this construction.

C Supplementary Surrogate and Training Comparisons

Table 2: STE training ablation over 2,000 paired LIBERO episodes per configuration. Both runs start from the same pretrained checkpoint and are fine-tuned separately with the indicated surrogate.

Configuration	Surrogate setting	Success (%)
Heartbeat $T = 4$, identity STE	identity surrogate	96.90
Heartbeat $T = 4$, magnitude STE	magnitude surrogate	96.90

The identity-STE and magnitude-STE Heartbeat models at $T = 4$ are fine-tuned separately from the same initial $\pi_{0.5}$ checkpoint, yielding a distinct final checkpoint for each surrogate. Both achieve 96.90% success under the paired evaluation protocol. The matched-budget continuous control (96.25%) continues fine-tuning $\pi_{0.5}$ with continuous activations for the same number of updates as the event-based variants, providing a reference that accounts for the additional training budget. Standard mode at $T = 2$ achieves 96.35% (Table 1). Relative to Heartbeat, it changes both temporal execution across layers and the velocity-field readout, so this comparison assesses their combined effect.

D Event-Representation Comparison

Table 3 compares our result with the reported nominal LIBERO success rates of other spiking VLA models.

Table 3: Reported nominal LIBERO success of spiking VLA models. Results use each study’s architecture and training protocol; — indicates an unlisted parameter count. The PoT-LIF row repeats the Heartbeat $T = 2$ result in Table 1.

Method	Parameters	Implementation	LIBERO mean
Spike-driven VLA [8]	0.15B	fully spiking, direct training	92.4%
DS-VLA [9]	—	dendritic spiking action control	91.6%
PoT-LIF on $\pi_{0.5}$	3.3B	Heartbeat $T = 2$ projections	97.25%

PoT-LIF achieves the highest reported nominal success among the models in Table 3: 97.25%, exceeding Spike-driven VLA by 4.85 percentage points and DS-VLA by 5.65 points. These results highlight the task-success advantage of combining amplitude-bearing events with a strong pretrained VLA.

Table 4 compares the event representations and integration strategies of the spiking VLA systems discussed in Sec. 2. The two PoT-LIF modes share signed power-of-two events and residual membrane updates. Standard mode propagates events across layers, whereas Heartbeat accumulates their weighted contributions within each projection. The distinction concerns the execution schedule of the recurrence, as described in Secs. 3–4.

Table 4: Event representations and integration strategies. Entries summarize the cited models; the PoT-LIF rows refer to the $\pi_{0.5}$ implementation evaluated here.

Model	Event representation	Integration strategy
Spike-driven VLA [8]	Binary spikes	Spiking perception, fusion, and action modules

Continued on next page

Table 4 (continued)

Model	Event representation	Integration strategy
DS-VLA [9]	Spiking dynamics	Dendritic-inspired architecture for action control
PoT-LIF standard mode	Signed power-of-two events; residual retention	Temporal sequence across layers
PoT-LIF Heartbeat	Signed power-of-two events; residual retention	Event generation and accumulation within each projection

E Run-Mode Parameters and Numerical Formats

The main configurations use bfloat16 (BF16) weights; the main Heartbeat configurations use FP32 membrane state. The supplementary FP8/A24 configuration uses FP8 E4M3FN weights [27], decoded through a BF16 container, and a 24-bit signed block-floating-point persistent accumulator. E4M3FN denotes a finite-value FP8 format with four exponent bits and three fraction bits.

Table 5: Operating-mode parameters. Both modes use the integration and residual balance in Eq. (10) and the emission rule in Eq. (7); the evaluated policies use ten external flow-matching solver steps, distinct from the internal event budget.

Parameter	Standard mode	Heartbeat
Input flow	one current per tick	single injection $I_1 = x$
Decay λ	configurable (main runs: 1)	1
Emission threshold θ	Two-sided; main runs: 2^{-16}	two-sided, code floor $2^{e_{\min}} = 2^{-16}$
State initialization	$r_0 = 0$ for each image-and-prompt input	$r_0 = 0$ per projection invocation; $u_1 = x$
Residual update and reset	Subtractive update each internal step; reset at next image-and-prompt input	Subtractive update; reset after T steps or early termination
Event application	events propagate downstream	applied inside the projection
Cross-layer time axis	present	none

Table 6: Training parameters and numerical formats. FP8 weights and A24 accumulation are supplementary configurations; the main table uses BF16 weights.

Parameter	Setting
Decay λ	1 (no leakage during local decomposition)
Emission threshold θ	two-sided code floor $2^{e_{\min}} = 2^{-16}$; extracts nonzero powers of the input significand
Primary Heartbeat and $T = 1$ surrogate	magnitude STE on the reconstructed sum, $g_T(x) = q_T(x)/x$ for $x \neq 0$
Surrogate at $x = 0$	magnitude: 0 at full strength, $1 - \alpha$ during insertion; identity: 1
Standard-mode surrogate at $T = 2$	constant 0.5 including at zero membrane potential, with cross-tick backpropagation
FP8 weights	E4M3FN, static power-of-two per-tensor scale
A24 accumulator	24-bit signed block-floating-point persistent state
Optimizer	Adam with decoupled weight decay (AdamW) [28], $\beta_1 = 0.9$, $\beta_2 = 0.95$, $\epsilon = 10^{-8}$, weight decay 10^{-10} , grad-norm clip 1.0
Learning-rate schedule	linear warmup 200 updates to 5×10^{-6} , cosine decay to 10^{-6} ; exponential moving average (EMA) decay 0.999

Continued on next page

Table 6 (continued)

Parameter	Setting
PoT insertion	linear ramp over the first 1,000 optimizer updates ($\alpha : 0 \rightarrow 1$)
Membrane precision	FP32 in Heartbeat main runs (BF16 membrane: separate no-fine-tuning variant)
Seeds	training 42, evaluation 7
Fine-tuning	5,000 optimizer updates, effective batch 256 across 4 devices
Forward execution	SAC per nonzero event–weight pair; floating-point power-of-two scaling through exponent addition and sign handling, followed by accumulation inside the projection

During gradual insertion, Heartbeat uses the forward map $x + \alpha(q_T(x) - x)$, where $\alpha = \text{clip}(s/(1000a), 0, 1)$, s is the microstep index, and a is the gradient-accumulation depth. The effective surrogate derivative is $(1 - \alpha) + \alpha g_T(x)$; at zero input it is $1 - \alpha$ for the magnitude surrogate. The shared $T = 1$ configuration uses the magnitude STE (Appendix H). Standard mode at $T = 2$ differentiates through its membrane recurrence with a local gain of 0.5. For the local recurrence with $\lambda = 1$, this gives $\partial s_t / \partial I_k = 0.5 \cdot 0.5^{t-k}$ for $k \leq t$, with internal inputs treated as independent. Membrane states persist across the network’s internal steps and reset for each new image-and-prompt input. For a nonzero weight tensor W , FP8 quantization uses a static power-of-two scale with exponent $e = \lceil \log_2(\max |W|/448) \rceil$ and round-to-nearest-even conversion. Biases, normalization scales, and embeddings retain their original precision. Query–key score products and attention-weighted value products remain floating-point.

F $T = 1$ SAC Arithmetic Accounting

This appendix reports only the operation-level arithmetic energy for one dense reference product. It excludes data transport, static random-access memory (SRAM) and register access, event decoding, sign-control logic, membrane encoding, leakage, clocking, communication, and other system-level costs.

The comparison uses one conventional FP32 multiply–accumulate (MAC) as the reference. The 4.6 pJ MAC cost is decomposed into a 3.7 pJ FP32 multiplication and a 0.9 pJ FP32 accumulation.

For the PoT-LIF $T = 1$ path, each input produces at most one signed power-of-two event. Using the reported 60.01% firing rate as the effective event–weight activity fraction gives an average of 0.6001 SAC operations per reference product, with zero-event contributions skipped. Each event requires one exponent-scaling operation and one FP32 accumulation. Exponent scaling is assigned an illustrative cost of 0.1 pJ per operation, using an integer-addition cost as a proxy.

Table 7: Average arithmetic energy per dense reference product with matched FP32 accumulation; the $T = 1$ PoT-LIF row uses a 60.01% firing rate. All other power terms are excluded.

Path	Multiplication / scaling (pJ)	FP32 accumulation (pJ)	Total (pJ)
FP32 MAC	$1 \times 3.7 = 3.7$	$1 \times 0.9 = 0.9$	4.6
PoT-LIF $T = 1$	$0.6001 \times 0.1 = 0.06001$	$0.6001 \times 0.9 = 0.54009$	0.6001

Thus, under this matched arithmetic model, PoT-LIF $T = 1$ averages 0.6001 pJ per reference product and the FP32 MAC reference uses 4.6 pJ. The energy ratio is $4.6/0.6001 \approx 7.67$, and the arithmetic-energy reduction is $1 - 0.6001/4.6 \approx 86.95\%$.

These values are arithmetic-energy estimates rather than hardware measurements. At an operation rate of f reference products per second, the corresponding modeled arithmetic powers are $4.6f$ pW for the FP32 MAC path and $0.6001f$ pW for the PoT-LIF $T = 1$ path. Full-system power is outside the scope of this accounting.

G McNemar Pairing Statistics

All 25 comparisons use 2,000 paired episodes. For models A and B, b counts A-success/B-failure pairs and c counts A-failure/B-success pairs. The exact two-sided McNemar value [29] is $p = \min(1, 2 \sum_{k=0}^{\min(b,c)} \binom{b+c}{k} 2^{-(b+c)})$. Table 8 groups the comparisons as follows: F1 contains the six dense-control rows, F2 the three matched-budget continuation rows, and E the 16 exploratory rows. Holm correction [30] is applied separately within F1 and F2. F1 includes the four main event configurations and two supplementary configurations: Heartbeat trained with the identity STE and the FP8/A24 variant. The two STE configurations are fine-tuned separately from the same initial checkpoint and evaluated using their respective final checkpoints. F2 compares standard mode at $T = 1$ and Heartbeat at $T = 2$ and $T = 4$ against the matched-budget continuous control to account for the additional fine-tuning updates. Archived FP8/FP32 entries are exploratory; their cross-process reproducibility remains unverified. None of the four main event configurations differs significantly from the original continuous-activation control after Holm correction at the 0.05 level; this does not establish statistical equivalence.

H Expressive range of the event code

We separately examine the representational capacity of power-of-two increments using trainable thresholds whose positions are parameterized separately from the emitted increments. Each ordered threshold sets where an increment is triggered, while the increment is specified by its own signed power-of-two code; its magnitude is not constrained to equal the threshold. This construction is distinct from the single-threshold recurrence used in the VLA projections; its formal definition is given in Appendix B. With 32 thresholds and at most four signed power-of-two terms per increment, it approximates a clipped rectified linear unit (ReLU) scaled by $1/4$, sigmoid, and hyperbolic tangent (tanh) with root mean square error (RMSE) values of 0.020, 0.033, and 0.045, respectively (Fig. 6). Threshold placement controls resolution along the input axis, while the available power-of-two terms constrain the output increments. These function-fitting results illustrate the expressive range of the event code. The ability of trainable thresholds to shape nonlinear responses motivates further investigation of this direction.




Figure 6: Function approximation with trainable threshold positions and separately parameterized power-of-two increments (Appendix B). Orange curves approximate the target functions shown in dark lines using 32 ordered thresholds and at most four signed power-of-two terms per increment. The respective RMSE values are 0.020, 0.033, and 0.045.

Table 8: Paired statistical comparisons. A and B are success rates (%); $\Delta = B - A$ is in percentage points. For F1 and the first nine E rows, A is the 96.70% continuous control; for F2, A is the 96.25% matched-budget control. Other E rows name their reference configuration. p is the unadjusted exact value; p_{Holm} is reported only for F1 and F2. The two STE rows correspond to separately fine-tuned final checkpoints. “BF16 membrane” distinguishes the precision sweeps from the main configurations with FP32 membrane state.

Fam.	Configuration B / reference A	A (%)	B (%)	Δ	b	c	p	p_{Holm}
F1	Heartbeat $T = 4$, identity STE	96.70	96.90	+0.20	38	42	0.738	1.000
F1	Heartbeat $T = 4$, magnitude STE	96.70	96.90	+0.20	38	42	0.738	1.000
E	Heartbeat $T = 4$ FP8/FP32 (archived)	96.70	97.00	+0.30	40	46	0.590	—
F1	Heartbeat $T = 4$ FP8/A24	96.70	97.50	+0.80	38	54	0.117	0.704
F1	Standard mode $T = 2$	96.70	96.35	−0.35	62	55	0.579	1.000
F1	Heartbeat $T = 2$ main	96.70	97.25	+0.55	40	51	0.294	1.000
E	Heartbeat $T = 2$ FP8/FP32 (archived)	96.70	96.90	+0.20	43	47	0.752	—
F1	Standard mode $T = 1$ BF16	96.70	96.90	+0.20	44	48	0.755	1.000
E	Heartbeat $T = 2$ BF16 membrane, $\theta = 2^{-16}$	96.70	97.15	+0.45	38	47	0.386	—
E	Heartbeat $T = 4$ BF16 membrane, $\theta = 2^{-16}$	96.70	97.20	+0.50	33	43	0.302	—
E	Heartbeat $T = 2$ BF16 membrane, $\theta = 2^{-8}$	96.70	97.05	+0.35	43	50	0.534	—
E	Heartbeat $T = 4$ BF16 membrane, $\theta = 2^{-8}$	96.70	97.20	+0.50	35	45	0.314	—
E	$T = 1$ FP8 E4M3FN, $\theta = 2^{-16}$	96.70	97.25	+0.55	40	51	0.294	—
E	$T = 1$ FP8 E4M3FN, $\theta = 2^{-8}$	96.70	96.35	−0.35	57	50	0.562	—
E	$T = 1$ FP8 E4M3FN, $\theta = 2^{-4}$	96.70	97.05	+0.35	47	54	0.551	—
F2	Standard mode $T = 1$ BF16	96.25	96.90	+0.65	44	57	0.232	0.284
F2	Heartbeat $T = 2$ main	96.25	97.25	+1.00	36	56	0.047	0.141
F2	Heartbeat $T = 4$ magnitude STE	96.25	96.90	+0.65	27	40	0.142	0.284
E	$T = 1$ FP8 / BF16 weights, $\theta = 2^{-16}$	96.90	97.25	+0.35	40	47	0.520	—
E	$T = 1$ FP8 / BF16 weights, $\theta = 2^{-8}$	96.90	96.35	−0.55	52	41	0.300	—
E	$T = 1$ FP8 / BF16 weights, $\theta = 2^{-4}$	96.90	97.05	+0.15	40	43	0.826	—
E	$T = 2$ BF16 membrane / main, $\theta = 2^{-16}$	97.25	97.15	−0.10	40	38	0.910	—
E	$T = 4$ BF16 membrane / main, $\theta = 2^{-16}$	96.90	97.20	+0.30	18	24	0.441	—
E	$T = 2$ BF16 membrane / main, $\theta = 2^{-8}$	97.25	97.05	−0.20	39	35	0.728	—
E	$T = 4$ BF16 membrane / main, $\theta = 2^{-8}$	96.90	97.20	+0.30	22	28	0.480	—

I Threshold and Precision Sweeps

The two sweeps below reuse fixed checkpoints without further fine-tuning. Each configuration is evaluated on 500 episodes per suite. Event statistics are computed for each precision path separately.

Standard mode at $T = 1$ with BF16 weights is the shared single-step configuration reported in Table 1. Starting from the continuous-activation checkpoint, it uses the magnitude STE and 5,000 optimizer updates at effective batch size 256. Success rates are 99.2% on Spatial, 98.4% on Object, 97.6% on Goal, and 92.4% on Long-Horizon, giving 96.90% overall (1,938/2,000).

Table 9: Heartbeat with BF16 membrane state and explicit thresholds, using fixed checkpoints and BF16 weights without further fine-tuning. The implementation’s four-bit threshold-selection index maps $c = 15$ to $\theta = 2^{-16}$ and $c = 7$ to $\theta = 2^{-8}$; it is not the event-exponent field. Δ is relative to the main Heartbeat configuration at the same T , which uses FP32 membrane state and $\theta = 2^{-16}$. Rates are percentages; differences are percentage points.

T	θ	Spatial	Object	Goal	Long-H.	Overall	Δ
2	2^{-16}	99.4	98.8	97.6	92.8	97.15	−0.10
2	2^{-8}	98.2	98.6	98.4	93.0	97.05	−0.20
4	2^{-16}	99.4	97.6	98.8	93.0	97.20	+0.30
4	2^{-8}	98.8	98.2	99.0	92.8	97.20	+0.30

Table 10: Standard mode at $T = 1$ with FP8 E4M3FN weights and three explicit thresholds, using the final $T = 1$ checkpoint without further fine-tuning and a static power-of-two scale per weight tensor. Δ is relative to the 96.90% BF16-weight $T = 1$ result. Rates are percentages; differences are percentage points.

θ	Spatial	Object	Goal	Long-H.	Overall	Δ
2^{-16}	98.8	98.6	98.0	93.6	97.25	+0.35
2^{-8}	98.2	98.0	97.6	91.6	96.35	−0.55
2^{-4}	98.6	97.6	97.8	94.2	97.05	+0.15

The BF16-membrane sweep changes overall success by −0.20 to +0.30 percentage points relative to the corresponding main Heartbeat configuration. The FP8-weight sweep changes success by −0.55 to +0.35 points relative to the BF16-weight single-step configuration. These fixed-checkpoint results describe sensitivity to the tested thresholds and numerical formats. The sweep comparisons are exploratory and fall outside the two corrected families in Appendix G.

References

- [1] Kevin Black, Noah Brown, Danny Driess, et al. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [2] Physical Intelligence, Kevin Black, Noah Brown, et al. $\pi_{0.5}$: A vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.
- [3] Anthony Brohan, Noah Brown, Justice Carbajal, et al. RT-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- [4] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, et al. OpenVLA: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [5] Wulfram Gerstner, Werner M. Kistler, Richard Naud, and Liam Paninski. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, 2014. doi: 10.1017/CBO9781107447615.
- [6] Louis Lapicque. Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation. *Journal de Physiologie et de Pathologie Générale*, 9:620–635, 1907. English translation by N. Brunel and M. C. W. van Rossum, *Biological Cybernetics*, 97:341–349, 2007, doi:10.1007/s00422-007-0189-6.
- [7] Nicolas Brunel and Mark C. W. van Rossum. Lapicque’s 1907 paper: From frogs to integrate-and-fire. *Biological Cybernetics*, 97(5–6):337–339, 2007. doi: 10.1007/s00422-007-0190-0.
- [8] Shuai Wang, Malu Zhang, Mingquan Liu, Weihui Dai, Dehao Zhang, Jieyuan Zhang, Yimeng Shan, Zijian Zhou, and Yang Yang. Spike-driven Vision-Language-Action Model. *arXiv preprint arXiv:2609.39514*, 2026.
- [9] Yaxing Lyu, Jingyi Li, Mingkun Xu, and Yujie Wu. DS-VLA: A Dendritic-inspired Vision-Language-Action Model for Robust Action Control. *arXiv preprint arXiv:2609.32253*, 2026.
- [10] Fabrizio Gabbiani, Holger G. Krapp, Christof Koch, and Gilles Laurent. Multiplicative computation in a visual neuron sensitive to looming. *Nature*, 420:320–324, 2002. doi: 10.1038/nature01190.
- [11] Yousheng Shu, Andrea Hasenstaub, Alvaro Duque, Yuguo Yu, and David A. McCormick. Modulation of intracortical synaptic potentials by presynaptic somatic membrane potential. *Nature*, 441:761–765, 2006. doi: 10.1038/nature04720.
- [12] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. In *Advances in Neural Information Processing Systems*, volume 36, pages 44776–44791, 2023.
- [13] Alain Destexhe, Zachary F. Mainen, and Terrence J. Sejnowski. An efficient method for computing synaptic conductances based on a kinetic model of receptor binding. *Neural Computation*, 6(1):14–18, 1994. doi: 10.1162/neco.1994.6.1.14.
- [14] M. V. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94(2):719–723, 1997. doi: 10.1073/pnas.94.2.719.
- [15] Yongqiang Cao, Yang Chen, and Deepak Khosla. Spiking Deep Convolutional Neural Networks for Energy-Efficient Object Recognition. *International Journal of Computer Vision*, 113(1):54–66, 2015. doi: 10.1007/s11263-014-0788-3.
- [16] Abhronil Sengupta, Yuting Ye, Ramakrishnan Wang, Chiao Liu, and Kaushik Roy. Going Deeper in Spiking Neural Networks: VGG and Residual Architectures. *Frontiers in Neuroscience*, 13:95, 2019. doi: 10.3389/fnins.2019.00095.

- [17] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63, 2019. doi: 10.1109/MSP.2019.2931595.
- [18] William B. Levy and Robert A. Baxter. Energy efficient neural codes. *Neural Computation*, 8(3):531–543, 1996. doi: 10.1162/neco.1996.8.3.531.
- [19] Julia J. Harris, Renaud Jolivet, Elisabeth Engl, and David Attwell. Energy-efficient information transfer by visual pathway synapses. *Current Biology*, 25(24):3151–3160, 2015. doi: 10.1016/j.cub.2015.10.063.
- [20] Daisuke Miyashita, Edward H. Lee, and Boris Murmann. Convolutional neural networks using logarithmic data representation. *arXiv preprint arXiv:1603.01025*, 2016.
- [21] Mostafa Elhoushi, Zihao Chen, Farhan Shafiq, Ye Henry Tian, and Joey Yiwei Li. DeepShift: Towards multiplication-less neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 2359–2368, 2021.
- [22] Hanting Chen, Yunhe Wang, Chunjing Xu, Boxin Shi, Chao Xu, Qi Tian, and Chang Xu. AdderNet: Do we really need multiplications in deep learning? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1468–1477, 2020. doi: 10.1109/CVPR42600.2020.00154.
- [23] Intel Labs. Taking Neuromorphic Computing to the Next Level with Loihi 2. Technology brief, Intel Corporation, 2021. URL <https://download.intel.com/newsroom/2021/new-technologies/neuromorphic-computing-loihi-2-brief.pdf>.
- [24] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. *arXiv preprint arXiv:2303.15343*, 2023.
- [25] Lucas Beyer, Andreas Steiner, André Susano Pinto, et al. PaliGemma: A versatile 3B VLM for transfer. *arXiv preprint arXiv:2407.07726*, 2024.
- [26] Mark Horowitz. Computing’s energy problem (and what we can do about it). In *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, pages 10–14, 2014. doi: 10.1109/ISSCC.2014.6757323.
- [27] Paulius Micikevicius et al. FP8 formats for deep learning. *arXiv preprint arXiv:2209.05433*, 2022.
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. arXiv:1711.05101.
- [29] Quinn McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12:153–157, 1947. doi: 10.1007/BF02295996.
- [30] Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979. URL <https://www.jstor.org/stable/4615733>.