LEARNING LOCAL MIXTURES OF TEMPORAL CODES FOR SPIKING NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Temporal codes determine how a static image drives a spiking neural network (SNN) over time. A pointwise encoder applies the same coding rule across an image, although equal normalized intensities can occur in different local structures. This raises an allocation question: how should a bank of temporal codes be combined across spatial locations? We introduce FAPE, a framework that learns this allocation from local image context. FAPE constructs binary Burst, time-to-first-spike (TTFS), and Rate candidates in parallel and uses multiscale features to predict normalized mixture weights at each location and simulation step. The resulting continuous currents drive an unchanged SNN backbone, with optional temporal conditioning parameterized separately. On validation-selected CIFAR-100 models, spatial-only FAPE achieves $69.21 \pm 0.19\%$ official-test top-1 accuracy across three training seeds, exceeding parameter-count-matched input-global and fixed-input gates by 1.90 and 3.18 percentage points. Both paired gains are positive for every seed. On the 1.28-million-image ImageNet32 benchmark, spatial-only FAPE exceeds the globally learned mixture by 2.86 points under a shared 90-epoch protocol. Direct min–max current reaches $74.24 \pm 0.41\%$ on CIFAR-100, identifying a stronger alternative input representation. Follow-up retraining with time-averaged, repeated inputs improves selected validation accuracy in all nine matched method–seed comparisons, while reducing the mean allocation contrasts. Together, these results connect local-allocation gains to the candidate representation and its presentation to the backbone.

1 INTRODUCTION

Spiking neural networks (SNNs) process information through discrete spikes and evolving neuronal states. For static-image recognition, input encoding converts images into sequences of firing rates, first-spike times, or spike bursts. These coding choices complement advances in SNN training and architecture (Wu et al., 2019; Deng et al., 2022; Fang et al., 2021). Hybrid approaches also assign different codes to network blocks (Chen et al., 2024).

Conventional pointwise codes leave local structure outside the coding decision: equal normalized intensities in flat regions, edges, and textures receive the same deterministic pattern or sampling distribution. Spatially shared mixture weights retain a common code preference across locations, even when conditioned on the image. Temporal-code mixtures therefore need allocation rules that can express differences in both spatial position and local image content.

FAPE learns local allocation over parallel binary Burst, time-to-first-spike (TTFS), and Rate candidates (Figure 1). Multiscale spatial filters provide context for gates that predict normalized mixture weights at each location and simulation step. The resulting continuous currents drive an unchanged SNN backbone. Separately parameterized temporal factors support generic, image-shared modulation and phase–amplitude coupling (PAC)-inspired, image-conditioned modulation.

Spatial-only FAPE reaches $69.21 \pm 0.19\%$ official-test top-1 on validation-selected CIFAR-100 models, with positive gains over parameter-count-matched input-global and fixed-input gates in all three seeds. It also exceeds the globally learned mixture by 2.86 points on ImageNet32 under a shared 90-epoch budget. Direct min–max current reaches $74.24 \pm 0.41\%$ on CIFAR-100. A validation-only follow-up improves all nine matched method–seed scores through repeated-mean retraining, with

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

smaller mean allocation contrasts. These findings connect local-allocation gains to the candidate representation and its presentation to the backbone.

2 Related Work

2.1 Temporal input codes and hybrid encoding

Rate-based ANN-to-SNN conversion approximates activations through firing rates (Diehl et al., 2015); time-based coding uses spike timing (Han & Roy, 2020), and burst coding uses short spike groups (Naud & Sprekeler, 2018). Hybrid approaches assign codes across network blocks (Chen et al., 2024), while multisynaptic neurons incorporate spatiotemporal coding into neuron design (Fan et al., 2025). FAPE learns local-context-dependent weights over a fixed Burst, TTFS, and Rate bank.

2.2 Adaptive input encoding and conditional mixtures

Gated Attention Coding (GAC) combines temporal and spatial–channel attention to gate a convolutional spiking encoder (Qiu et al., 2024). Mixture-of-experts networks use input-dependent gates over subnetworks (Shazeer et al., 2017), and conditional convolution uses example-dependent kernel coefficients (Yang et al., 2019). FAPE weights fixed temporal-code operators at each location and simulation step, evaluating all three sequences before mixing.

2.3 Frequency-aware processing in SNNs

Analyses of leaky integrate-and-fire (LIF) dynamics motivate high-frequency enhancement (Fang et al., 2025), while FSTA-SNN introduces frequency-based spatial and temporal attention (Yu et al., 2025). FEEL-SNN reconstructs input sequences using time-dependent Fourier masks and learns evolutionary leak factors for noise robustness (Xu et al., 2024). FAPE constructs candidates from image intensities and conditions their mixture weights on learned multiscale neighborhood features.

2.4 Spiking backbones and PAC-inspired conditioning

Residual SNNs enable deep spiking networks (Fang et al., 2021); Spikformer and Spike-Driven Transformer introduce attention-based architectures (Zhou et al., 2023; Yao et al., 2023). We evaluate FAPE's separately parameterized input allocation with residual and VGG-style SNNs, retaining their neuron dynamics.

Neuroscientific PAC relates higher-frequency amplitude to lower-frequency phase (Canolty & Knight, 2010), inspiring FAPE's optional conditioning. Two learned responses of scale-ordered image context produce bounded code-weight factors. Separate allocation and modulation parameters support spatial-only, generic, and PAC-inspired comparisons.

3 Adaptive Code Allocation

3.1 Problem setup

Given channel-standardized images $x \in \mathbb{R}^{N\times C\times H\times W}$ and T steps, FAPE constructs $Y_{0:T-1}$ by mixing binary Burst, TTFS, and Rate candidates. The core spatial-only model predicts location-dependent gates from learned spatial filters. Optional generic and PAC-inspired factors provide image-shared and image-conditioned modulation, respectively; the output interface connects the mixture to the backbone (Figure 1).

For the backbone, let I_t denote the current entering a LIF neuron after preceding layer transformations. With decayed input and zero hard reset, the discrete dynamics are

$$u_t = \beta u_{t-1}(1 - s_{t-1}) + (1 - \beta)I_t, \qquad s_t = \mathbb{1}[u_t \geq V_{\text{th}}], \tag{1}$$

where u_t is the pre-reset potential and s_t is the binary spike. The reported time constant $\tau_{\text{LIF}} = 2$ gives $\beta = 1-1/\tau_{\text{LIF}} = 0.5$. From zero initial state, omitting thresholding and reset yields the linear

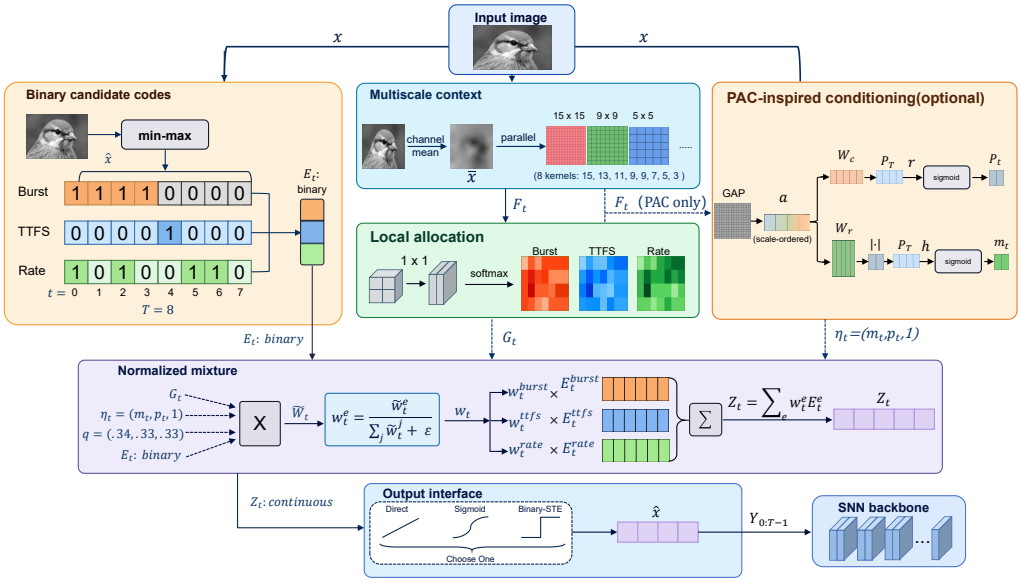


Figure 1: FAPE combines parallel Burst, TTFS, and Rate candidates using location-dependent gates and optional temporal factors. Generic factors are shared across images; PAC-inspired factors use each image's scale-ordered context. The direct mixture serves ImageNet32, VGG-11, and the validation-selected CIFAR-100 cohort. Historical residual CIFAR uses a sigmoid mixture, and Binary-STE applies straight-through thresholding.

charging-path transfer function

$$H(z) = \frac{1-\beta}{1-\beta z^{-1}}, \tag{2}$$

which characterizes temporal filtering along simulation steps. This temporal response motivates examining how candidate codes distribute events across steps.

3.2 Candidate Temporal Codes

The candidate set is $\mathcal{E} = \{\text{burst}, \text{ttfs}, \text{rate}\}$, with tensors stacked in Rate–TTFS–Burst order. Each image is independently rescaled over its channel and spatial entries:

$$\hat{x}_n = \frac{x_n - \min_{c,h,w} x_{nchw}}{\max_{c,h,w} x_{nchw} - \min_{c,h,w} x_{nchw} + 10^{-6}} \in [0,1]^{C\times H\times W}. \tag{3}$$

For a zero-based simulation index $\tau \in \{0, \ldots, T-1\}$, the three branches are

$$\begin{aligned} b(\hat{x}) &= \text{clip}(\text{round}(T\hat{x}), 0, T)\,, & E_\tau^{\text{burst}} &= \mathbb{1}[\tau < b(\hat{x})], \\ \tau^*(\hat{x}) &= \text{clip}(\text{round}((T-1)(1-\hat{x})), 0, T-1)\,, & E_\tau^{\text{ttfs}} &= \mathbb{1}[\tau = \tau^*(\hat{x})], \\ p(\hat{x}) &= \text{clip}(\hat{x} r_{\max} \Delta t, 0, 1), & E_\tau^{\text{rate}} &\sim \text{Bernoulli}(p(\hat{x})). \end{aligned} \tag{4}$$

The reported experiments use $T = 8$, a maximum Burst length of T, and $r_{\max} = \Delta t = 1$. Burst and TTFS use nearest-integer rounding with ties to even. At $\hat{x} = 0$, Burst and Rate produce zero events, and TTFS produces one event at $\tau = T - 1$. Rate draws are independent across steps and tensor entries. Each candidate has shape $N \times C \times H \times W$. Candidates use $\hat{x}$; local context uses the standardized image x.

3.3 Multiscale Local Allocation

The channel mean $\bar{x} = C^{-1} \sum_c x_c$ provides a shared spatial input for the allocation module. At each simulation step, a learned spatial filter K_t produces a feature map and its global average:

$$F_t = K_t * \bar{x}, \qquad a_t = \text{GAP}(F_t). \tag{5}$$

Here GAP averages over spatial positions, giving one descriptor per image and step. For $T = 8$, kernel widths $\{15, 13, 11, 9, 9, 7, 5, 3\}$ are assigned in order to steps $t = 0, \ldots, 7$, linking broad-to-local spatial context to the simulation sequence. All filters act on the same static image, with symmetric zero padding preserving spatial dimensions. A step-specific 1×1 convolution maps F_t to three logits, followed by a softmax over candidate codes:

$$G_t = \mathrm{softmax}_e(C_t(F_t)) \in [0, 1]^{N \times 3 \times H \times W}. \tag{6}$$

The resulting allocation field varies with image, spatial location, and simulation step. Each code's gate weight is shared across the input color channels at a given location.

Fixed gates use $G_t^e = 1/3$; globally learned gates use a softmax-normalized three-way vector for each step. Code-mixture variants share the fixed positive branch prior $q = (0.33, 0.33, 0.34)$ in Rate–TTFS–Burst order. Equal gates with unit temporal factors therefore produce effective weights approximately equal to q, with the stabilizer specified in Equation 9.

3.4 Optional temporal conditioning

For PAC-inspired conditioning, $a = (a_0, \ldots, a_{T-1})$ collects the scale-ordered descriptors of each static image. The sequence index follows the spatial scales assigned to simulation steps. Two unconstrained one-dimensional filters produce a context response and an absolute-valued response:

$$r = \mathcal{P}_T(W_c * a), \qquad h = \mathcal{P}_T(|W_r * a|), \tag{7}$$

followed by the sigmoid factors

$$p_t = \sigma(\gamma_c r_t), \qquad m_t = \sigma(\gamma_h h_t). \tag{8}$$

The operator $\mathcal{P}_T$ applies adaptive average pooling when the response length differs from T. For $T = 8$, W_c uses width 4 and padding 2, producing nine outputs that are pooled to eight. The filter W_r uses width 7 and padding 3 to directly produce eight outputs. The reported configuration applies the sigmoid directly to these responses with fixed gains $\gamma_c = \gamma_h = 1$.

PAC-inspired factors modulate TTFS by p_t and Burst by m_t, with a unit factor for Rate. These image-dependent factors satisfy $p_t \in (0, 1)$ and $m_t \in [0.5, 1)$ under the reported settings. Generic conditioning uses two sigmoid-transformed learned logits per step, shared across images. Spatial-only allocation uses unit factors for all three candidate codes.

3.5 Code mixing and output interfaces

Code weights combine the local gate, the positive branch prior, and the temporal factors. With $\eta_t = (1, p_t, m_t)$ in Rate–TTFS–Burst order, the effective weights are

$$\tilde{w}_t^e = G_t^e q_e \eta_t^e, \qquad w_t^e = \frac{\tilde{w}_t^e}{\sum_{j \in \mathcal{E}} \tilde{w}_t^j + \epsilon}. \tag{9}$$

These weights combine the binary candidates into a generally continuous mixture:

$$Z_t = \sum_{e \in \mathcal{E}} w_t^e E_t^e. \tag{10}$$

Writing $S_t = \sum_e G_t^e q_e \eta_t^e$, the weights sum to $S_t/(S_t + \epsilon)$, with $\epsilon = 10^{-6}$. The mixture satisfies $Z_t \in [0, 1]$, and its weight sum approaches one when $S_t \gg \epsilon$. Generic factors can be absorbed into step-specific gate biases under ideal normalization. They provide a parameterization and optimization control, with the exact relation and finite-ϵ qualification given in Appendix C.

The *mixture interface* passes Z_t directly to the backbone before its first convolution:

$$Y_t^{\mathrm{mix}} = Z_t. \tag{11}$$

ImageNet32, Spiking VGG-11, and the validation-selected CIFAR-100 mixture cohort use this interface. Historical residual CIFAR experiments use a sigmoid transformation of the mixture:

$$Y_t^{\mathrm{sig}} = \sigma(Z_t), \tag{12}$$

giving continuous currents within $[0.5, \sigma(1)]$. When $E_t^e = 0$ for every candidate e, the direct and sigmoid interfaces supply 0 and 0.5, respectively. The Binary-STE configuration provides an event-valued boundary through thresholding:

$$Y_t^{\mathrm{ste}} = \mathbb{1}[Z_t \geq 0.5], \qquad \frac{\partial Y_t^{\mathrm{ste}}}{\partial Z_t} \approx 1, \tag{13}$$

where the straight-through estimator (STE) carries gradients through the threshold during optimization. Allocation ablations use matched mixture interfaces; comparisons with binary single-code baselines additionally vary the input representation.

3.6 Optimization and optional regularization

The fixed candidates are non-differentiable and parameter-free. Allocation and conditioning receive gradients through the mixture and output interface; the backbone uses surrogate-gradient backpropagation through time with detached reset spikes. Binary-STE adds Equation 13 at the boundary. Every forward pass computes all three candidates.

The optional Information Sparse Loss (ISL) augments cross-entropy as

$$\mathcal{L} = \mathcal{L}_{\mathrm{CE}} + \lambda_1 \mathcal{L}_{\mathrm{MI}} + \lambda_2 \mathcal{L}_{\mathrm{sp}}. \tag{14}$$

The InfoNCE-style bidirectional term $\mathcal{L}_{\mathrm{MI}}$ pairs image features with encoder features, pooled backbone features, or logits, optionally using a projection head. The activity term constrains mean encoder activation $\bar{a}$ to $[\rho_{\mathrm{lo}}, \rho_{\mathrm{hi}}]$:

$$\mathcal{L}_{\mathrm{sp}} = \begin{cases} (\rho_{\mathrm{lo}} - \bar{a})^2, & \bar{a} < \rho_{\mathrm{lo}}, \\ (\bar{a} - \rho_{\mathrm{hi}})^2, & \bar{a} > \rho_{\mathrm{hi}}, \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

Here $\bar{a}$ measures mean continuous output or, for Binary-STE, event rate. Factorization, ImageNet32, VGG-11, and validation-selected CIFAR-100 use cross-entropy alone. Auxiliary CIFAR studies evaluate ISL through recognition accuracy (Appendix F); Appendix B gives the forward pass.

4 Experiments

4.1 Experimental design

Data and backbones. We evaluated FAPE on CIFAR-10 and CIFAR-100 (Krizhevsky, 2009) and on ImageNet32 (Chrabaszcz et al., 2017). ImageNet32 retains 1,000 classes, 1,281,167 training images, and a 50,000-image validation split at 32×32 resolution. The primary backbone was a CIFAR-stem Spiking ResNet-18 (Fang et al., 2021); Spiking VGG-11 on CIFAR-100 provided a non-residual convolutional comparison. All models used $T = 8$ simulation steps, a LIF time constant of 2.0, hard reset, and multi-step execution.

Validation-selected CIFAR-100 cohort. The validation-selected cohort comprised ten methods and three training seeds, $\{42, 1337, 2024\}$, with a fixed 45,000/5,000 training/validation split. Each run used 150 epochs of SGD training, with parameter-group settings specified in Appendix D. Mean validation top-1 over five encoding draws selected each checkpoint, with the earliest epoch breaking ties. All 30 selected models were frozen before evaluation on the official 10,000-image test split. Test reporting averaged five encoding draws within each model, then calculated means and sample standard deviations across the three training seeds. Appendix D documents the controls, numerical-recovery amendment, and final-evaluation procedure.

Historical protocols and reporting. Historical residual CIFAR and VGG-11 runs used 150-epoch SGD, with five warm-up epochs and cosine decay for the main parameter group. ImageNet32 used Adam for 90 epochs with constant learning rates. Appendix A provides the parameter-group settings, neuron thresholds, and output-interface map. ImageNet32 and VGG-11 mixtures used Equation 11; residual CIFAR mixtures used the sigmoid interface. Historical CIFAR scores used the standard test split for checkpoint selection and reporting, yielding best-epoch evaluation accuracies.

Table 1: Official-test accuracy of the validation-selected CIFAR-100 cohort with Spiking ResNet-18. Entries are mean ± sample standard deviation across three training seeds after averaging five encoding draws per model. The upper block shares the candidate code bank and continuous-mixture boundary; the lower block evaluates alternative continuous inputs. All methods use eight steps, the same data split, and 150 epochs per run. Summaries were reconciled against all 30 model records (Appendix D).

Input configuration	Top-1 (%) ↑	Top-5 (%) ↑
Fixed mixture	65.78 ± 0.17	88.45 ± 0.15
Global learnable mixture	67.76 ± 0.14	89.80 ± 0.07
Input-global gate	67.31 ± 0.77	89.46 ± 0.44
Fixed-input gate	66.03 ± 1.02	88.56 ± 0.61
FAPE (spatial only)	69.21 ± 0.19	90.34 ± 0.14
FAPE (generic temporal)	68.57 ± 0.74	90.21 ± 0.27
FAPE (PAC-inspired)	69.31 ± 0.67	90.71 ± 0.51
Direct standardized current	$\mathbf{74.69 \pm 0.13}$	92.94 ± 0.18
Direct min–max current	74.24 ± 0.41	$\mathbf{93.00 \pm 0.34}$
Learned continuous current	74.31 ± 0.05	92.83 ± 0.21

ImageNet32 checkpoints were selected on its official validation split, with top-5 taken from the same top-1-selected epoch. Each comparison applied a common selection rule, and Appendix E provides seed-level values. Unless stated otherwise, historical entries report means and sample standard deviations across the same three training seeds.

Comparison families. Within each protocol, allocation comparisons shared the candidate code bank, continuous-mixture interface, backbone, and training recipe. The fixed mixture used effective Rate–TTFS–Burst weights close to $(0.33, 0.33, 0.34)$. Global learnable mixtures and local FAPE variants supplied the allocation alternatives. The validation-selected cohort added input-global and fixed-input gates as parameter-count-matched controls, together with direct and learned continuous-input alternatives. Binary Rate, Burst, and TTFS baselines supplied end-to-end comparisons across encoder representations. Each protocol formed a separate evidence block, preserving its checkpoint-selection and interface conditions.

4.2 Local allocation and continuous-input alternatives on CIFAR-100

Spatial-only FAPE exceeded the input-global and fixed-input gates by 1.90 and 3.18 percentage points in mean official-test top-1 (Table 1). Both prespecified paired differences were positive for every training seed (Appendix D.1). All three configurations had 808 active encoder parameters and shared the candidate branches. The fixed-input control conditioned its gate on a fixed texture while constructing candidate codes from the actual image. These comparisons support image-conditioned, location-dependent allocation within the matched gate construction.

The continuous-input alternatives achieved higher mean accuracy than every candidate-mixture variant. Direct min–max current exceeded spatial-only FAPE by 5.03 points while sharing its per-image normalization. Direct standardized current and learned continuous current also ranked above the mixture family. The learned-current control retained step-dependent gates, allowing its continuous representation to vary across simulation steps. Within the mixture family, PAC-inspired FAPE had a three-seed mean 0.10 points above spatial-only FAPE. Together, these comparisons locate the allocation benefit within the candidate bank and show the stronger recognition performance of the tested continuous-input alternatives.

4.3 Allocation under temporal averaging

A validation-only follow-up examined temporal structure in the realized mixture. We applied reversal, a fixed permutation, and temporal averaging to the nine frozen spatial, input-global, and fixed-input models. The averaging condition repeated the time-averaged mixture at all eight steps. We then trained nine new models with this transformation throughout training and validation, retaining the

Encoding strategy	Top-1 (%)	Top-5 (%)
Rate (binary)	10.08 ± 0.07	25.38 ± 0.06
Burst (binary)	12.77 ± 0.17	31.66 ± 0.62
TTFS (binary)	13.74 ± 0.15	33.26 ± 0.38
Fixed prior $(.33/.33/.34)$	14.53 ± 0.11	34.42 ± 0.22
Global learnable mixture	16.49 ± 0.26	37.96 ± 0.50
FAPE (PAC-inspired)	18.86 ± 0.27	41.42 ± 0.30
FAPE (generic temporal)	19.08 ± 0.39	41.65 ± 0.50
FAPE (spatial only)	$\mathbf{19.35 \pm 0.37}$	$\mathbf{42.39 \pm 0.66}$

Table 2: ImageNet32 validation accuracy with Spiking ResNet-18 under a shared 90-epoch protocol. Entries are mean $\pm$ sample standard deviation across three training seeds. The upper and lower blocks use binary single-code and continuous-mixture interfaces, respectively. Allocation comparisons use the lower block, whose methods share the candidate code bank and output interface.

original split, paired initializations, 808 encoder parameters, and 150-epoch budget (Appendix J). All scores in this follow-up use the 5,000-image checkpoint-selection validation split.

Frozen spatial models achieved $69.17 \pm 0.46\%$ with their original inputs and $46.69 \pm 6.59\%$ after averaging. With matched retraining, spatial allocation reached $71.44 \pm 0.54\%$; all nine retrained method–seed combinations exceeded their original validation scores. Under repeated-mean inputs, the spatial gain over input-global allocation averaged 0.85 points and was positive in every seed. Its gain over fixed-input allocation averaged 0.52 points, with paired differences of -0.24, $+0.99$, and $+0.82$ points. These mean contrasts were smaller than the original validation contrasts of 1.61 and 2.68 points. The frozen-model drops characterize sensitivity to a changed input distribution, while retraining establishes effective recognition with time-constant encoded inputs. Temporal averaging changes both across-step variation and stochastic fluctuations, retaining the time-integrated mixture and the backbone's neuronal dynamics.

4.4 Large-sample controlled comparison on ImageNet32

On ImageNet32, spatial-only FAPE exceeded global learnable and fixed mixtures by 2.86 and 4.82 percentage points in mean validation top-1 (Table 2). Every local variant exceeded the fixed mixture in each training seed. These comparisons evaluate the combined addition of image conditioning, spatial variation, and gate parameters to an input-independent mixture. Spatial-only FAPE also had the highest mean among the local variants on both top-1 and top-5. The three-seed matrix extends the local-mixture comparison to the 1.28-million-image training set under a shared 90-epoch budget.

4.5 Allocation across datasets and backbones

The historical factorization matrix showed setting-dependent effects of additional temporal conditioning (Table 3). Spatial-only allocation had the highest mean on CIFAR-10 and ImageNet32. On residual CIFAR-100, generic conditioning led spatial-only FAPE by 0.18 points; on VGG-11, PAC-inspired conditioning led by 0.10 points. The configuration with the highest mean varied across the evaluated datasets, backbones, and interfaces.

On CIFAR-100 with VGG-11, the fixed mixture achieved $62.10 \pm 0.14\%$, compared with $65.85 \pm 0.85\%$ for spatial-only FAPE and $65.95 \pm 0.24\%$ for PAC-inspired FAPE. PAC-inspired FAPE gained 3.85 points over the fixed mixture, with a positive difference in all three seeds. This comparison extends the local-allocation evidence to a non-residual convolutional backbone.

Appendix F reports the auxiliary residual CIFAR single-code benchmark and component studies under their respective protocols. Each experiment family is interpreted separately, preserving its selection and optimization conditions.

4.6 Associations between learned gates and image statistics

We examined allocation fields on 1,280 CIFAR-10 evaluation images using one seed-42 checkpoint per variant. Table 4 reports correlations with blurred grayscale intensity and the absolute residual

Local variant	C10 / R18	C100 / R18	ImageNet32 / R18	C100 / VGG-11
FAPE (spatial only)	$\mathbf{92.57 \pm 0.07}$	71.97 ± 0.28	$\mathbf{19.35 \pm 0.37}$	65.85 ± 0.85
FAPE (generic temporal)	92.24 ± 0.01	$\mathbf{72.15 \pm 0.14}$	19.08 ± 0.39	–
FAPE (PAC-inspired)	91.60 ± 1.14	71.73 ± 0.02	18.86 ± 0.27	$\mathbf{65.95 \pm 0.24}$

Table 3: Historical factorization of local allocation and temporal conditioning, with mean $\pm$ sample standard deviation across three training seeds. C10/C100 denote CIFAR-10/CIFAR-100; R18 denotes Spiking ResNet-18. CIFAR values are test-selected best-epoch evaluation accuracies; ImageNet32 values are selected validation accuracies. Residual CIFAR uses the sigmoid interface, while ImageNet32 and VGG-11 use the direct mixture. Each column defines a separate comparison with its own dataset, backbone, and interface.

Variant	Code	Absolute blur residual	Blurred intensity
Spatial only	Rate	0.133 ± 0.009	0.848 ± 0.005
	TTFS	-0.179 ± 0.010	-0.893 ± 0.004
	Burst	0.183 ± 0.010	0.893 ± 0.004
Generic temporal	Rate	-0.258 ± 0.011	0.633 ± 0.010
	TTFS	-0.178 ± 0.009	-0.272 ± 0.007
	Burst	0.240 ± 0.010	0.133 ± 0.007
PAC-inspired	Rate	-0.301 ± 0.015	0.580 ± 0.018
	TTFS	-0.152 ± 0.008	-0.468 ± 0.014
	Burst	0.238 ± 0.009	0.366 ± 0.012

Table 4: Effective code weights versus two image statistics on CIFAR-10, using one seed-42 checkpoint per variant. Entries are mean within-image Pearson correlations $\pm$ the half-width of a 95% normal-approximation interval. The half-width is $1.96s/\sqrt{n}$, with $n = 1{,}280$ images for every entry. Intervals summarize across-image uncertainty conditional on the selected checkpoint. Correlations are computed separately for each statistic, without adjustment for the other.

between grayscale input and its Gaussian-blurred version. Burst weights had positive mean residual correlations, while TTFS weights had negative means in all three variants.

For spatial-only FAPE, Burst weights correlated more strongly with blurred intensity (0.893) than with the absolute residual (0.183). Shared normalization couples the branch weights, linking their associations with either statistic. These observations describe dependence on local image statistics, with intensity and residual contrast remaining intertwined in this diagnostic. Appendix H specifies the calculation and the normalization relation.

We probed the three frozen CIFAR-100 spatial encoders with 29 equal-histogram patterns (Appendix J.3). A one-pixel shift of width-one vertical stripes changed spatially averaged gates by a mean L1 distance of 0.0031, versus 1.1735 for same-coordinate gate maps. Spatial averaging thus concealed local response changes. These probes vary both local intensity and its neighborhood; the same-coordinate metric uses unaligned positions.

4.7 Event-valued boundary and computational accounting

Binary-STE used Equation 13 to produce an event-valued boundary and achieved $90.99 \pm 0.26\%$ on CIFAR-10. Its per-seed differences from continuous PAC-inspired FAPE had mixed signs. These results establish binary-interface feasibility in the evaluated CIFAR-10 setting.

In CIFAR-100 validation profiling, spatial, input-global, and fixed-input gates each used 0.803M encoder Conv/Linear multiply–accumulate operations (MACs) per image. Spatial allocation yielded 408.63M estimated internal binary SynOps, versus 381.18M for input-global and 487.49M for direct min–max current. Appendix I reports seed variability, activity statistics, and counting scope. These activity-scaled estimates complement the accuracy comparisons; hardware energy and latency were not measured.

5 DISCUSSION

FAPE makes code allocation testable through a fixed candidate bank and matched gates. Each coefficient weights a defined temporal-code operator at a specific location and step. The CIFAR-100 controls support image-conditioned, location-dependent allocation. Repeated-mean retraining improves all nine selected validation scores while narrowing mean allocation contrasts. Spatial–input-global gains remain positive in every seed; fixed-input comparisons have mixed signs. These results connect allocation gains to the chosen input-sequence construction.

Gate diagnostics relate allocation to intensity, residual contrast, and spatial arrangement. Burst and TTFS weights have opposite residual correlations; spatial-only Burst weights associate more strongly with blurred intensity. Shared normalization couples these responses, and equal-histogram probes reveal local changes hidden by spatial averages.

ImageNet32 and CIFAR-100 VGG-11 extend the local-versus-fixed mixture evidence across sample scale and convolutional architecture. Temporal modulation has setting-dependent effects: spatial-only allocation leads the historical CIFAR-10 and ImageNet32 comparisons, while generic and PAC-inspired variants lead residual CIFAR-100 and VGG-11, respectively. FAPE's modular design allows these contributions to be examined separately.

Direct and learned continuous currents achieve higher accuracy than candidate mixtures in the validation-selected CIFAR-100 cohort. The fixed bank makes local code preferences explicit, while continuous-input controls test the representation itself. Together, these comparisons distinguish gains from allocating a given bank from gains obtained by changing its input forms.

6 CONCLUSION

FAPE learns local mixtures of Burst, TTFS, and Rate codes through multiscale gates while preserving backbone dynamics. Matched CIFAR-100 controls support spatial and image-conditioned allocation; ImageNet32 and VGG-11 extend local-versus-fixed comparisons across sample scale and architecture. Validation-only repeated-mean retraining improves recognition with smaller allocation contrasts. Stronger continuous-current controls establish input representation as a key companion to allocation in evaluating SNN encoders.

REPRODUCIBILITY STATEMENT

Section 3 specifies the method and interfaces; the appendices document optimization, neurons, software, pseudocode, selection rules, and seed-level results. A separate anonymous evidence supplement provides A10/A30 model records, the 30-model cost/activity audit, full-precision tables, historical summaries, and record-level recomputation scripts. Its path-redacted derivatives have a separate integrity manifest. Experiment configurations, launch scripts, and retained training artifacts will be released upon publication.

AI USE STATEMENT

Generative AI tools assisted with conceptual framing, experiment-planning and diagnostic feedback, interpretation and organization of author-provided results, Chinese–English translation, manuscript restructuring and language editing, code drafting and debugging, visual drafting, and consistency checks. The tools did not autonomously conduct experiments or fabricate or alter measurements. The authors reviewed and tested AI-assisted code, checked prose and graphics against implementations, configurations, logs, and result files, manually reviewed citations, and take responsibility for the final text, claims, equations, code, and artifacts.

REFERENCES

Ryan T. Canolty and Robert T. Knight. The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, 14(11):506–515, 2010.

X. Chen, Q. Yang, J. Wu, H. Li, and K. C. Tan. A hybrid neural coding approach for pattern recognition with spiking neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3064–3078, 2024.

Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A downsampled variant of ImageNet as an alternative to the CIFAR datasets. *arXiv preprint arXiv:1707.08819*, 2017.

Shikuang Deng, Yuhang Li, Shanghang Zhang, and Shi Gu. Temporal efficient training of spiking neural network via gradient re-weighting. In *International Conference on Learning Representations*, 2022.

Peter U. Diehl, Daniel Neil, Jonathan Binas, Matthew Cook, Shih-Chii Liu, and Michael Pfeiffer. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *International Joint Conference on Neural Networks*, 2015.

L. Fan, H. Shen, X. Lian, Y. Li, M. Yao, G. Li, and D. Hu. A multisynaptic spiking neuron for simultaneously encoding spatiotemporal dynamics. *Nature Communications*, 16(1):7155, 2025.

Wei Fang, Zhaofei Yu, Yanqi Chen, Tiejun Huang, Timothee Masquelier, and Yonghong Tian. Deep residual learning in spiking neural networks. In *Advances in Neural Information Processing Systems*, 2021.

Wei Fang, Yanqi Chen, Jianhao Ding, Zhaofei Yu, Timothee Masquelier, Ding Chen, Liwei Huang, Huihui Zhou, Guoqi Li, and Yonghong Tian. SpikingJelly: An open-source machine learning infrastructure platform for spike-based intelligence. *Science Advances*, 9(40):eadi1480, 2023.

Yuetong Fang, Deming Zhou, Ziqing Wang, Hongwei Ren, ZeCui Zeng, Lusong Li, Shibo Zhou, and Renjing Xu. Spiking neural networks need high frequency information, 2025.

Bing Han and Kaushik Roy. Deep spiking neural network: Energy efficiency through time based coding. In *European Conference on Computer Vision*, 2020.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

Richard Naud and Henning Sprekeler. Sparse bursts optimize information transmission in a multiplexed neural code. *Proceedings of the National Academy of Sciences*, 115(26):6669–6674, 2018.

Adam Paszke et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.

Xuerui Qiu, Rui-Jie Zhu, Yuhong Chou, Zhaorui Wang, Liang-Jian Deng, and Guoqi Li. Gated attention coding for training high-performance and efficient spiking neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(1):601–610, 2024. doi: 10.1609/aaai.v38i1.27816. URL https://ojs.aaai.org/index.php/AAAI/article/view/27816.

Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc V. Le, Geoffrey E. Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations*, 2017.

Yujie Wu, Lei Deng, Guoqi Li, Jun Zhu, and Luping Shi. Spatio-temporal backpropagation for training high-performance spiking neural networks. *Neural Networks*, 121:331–343, 2019.

Mengting Xu, De Ma, Huajin Tang, Qian Zheng, and Gang Pan. FEEL-SNN: Robust spiking neural networks with frequency encoding and evolutionary leak factor. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-2917. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/a73474c359ed523e6cd3174ed29a4d56-Abstract-Conference.html.

Brandon Yang, Gabriel Bender, Quoc V. Le, and Jiquan Ngiam. CondConv: Conditionally parameterized convolutions for efficient inference. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Man Yao, JiaKui Hu, Zhaokun Zhou, Li Yuan, Yonghong Tian, Bo Xu, and Guoqi Li. Spike-driven transformer. In *Advances in Neural Information Processing Systems*, 2023.

K. Yu, T. Zhang, H. Wang, and Q. Xu. FSTA-SNN: Frequency-based spatial-temporal attention module for spiking neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025.

Zhaokun Zhou, Yuesheng Zhu, Chao He, Yaowei Wang, Shuicheng Yan, Yonghong Tian, and Li Yuan. Spikformer: When spiking neural network meets transformer. In *International Conference on Learning Representations*, 2023.

A Historical implementation and protocol details

The shared data pipeline first applies dataset-level channel standardization. The spatial context and gates use this standardized tensor directly, whereas each candidate encoder independently applies the same per-image min–max rescaling from Equation 3 over all channel and spatial entries. Candidate branches otherwise share the timestep budget, augmentation, and classifier within a comparison. Experiments were executed under Linux on NVIDIA RTX A6000 GPUs (48 GB) with Python 3.10.18, PyTorch 2.6.0+cu124, TorchVision 0.21.0+cu124, and SpikingJelly 0.0.0.0.14 (Paszke et al., 2019; Fang et al., 2023).

The PAC-inspired conditioner uses one-dimensional kernel widths $k_c = 4$ and $k_r = 7$, fixed gains $\gamma_c = \gamma_h = 1$, and no response normalization. Symmetric padding gives the width-4 context response length 9 at $T = 8$; adaptive average pooling returns it to eight entries before modulation. Only the spatial 1×1 `gate_convs` parameters form the dedicated constant-learning-rate optimizer group. Global gate logits, the temporal filters, and generic temporal logits remain in the scheduled main group. The dedicated spatial-gate learning rate is 10^{-4} for residual CIFAR and VGG-11 and 3×10^{-4} for ImageNet32. Its weight decay is 0.1 for residual CIFAR and 5×10^{-4} for ImageNet32 and VGG-11.

Setting	Residual CIFAR	ImageNet32	CIFAR-100 / VGG-11
Timesteps	8	8	8
Epochs	150	90	150
Batch size	128	128	128
Optimizer	SGD	Adam	SGD
Initial learning rate	0.1	3×10^{-4}	0.1
Momentum	0.9	–	0.9
Main-group weight decay	5×10^{-4}	5×10^{-4}	5×10^{-4}
Spatial-gate weight decay	0.1	5×10^{-4}	5×10^{-4}
Spatial-gate learning rate	10^{-4}	3×10^{-4}	10^{-4}
Main-group schedule	warm-up + cosine	none	warm-up + cosine
Spatial-gate schedule	constant	constant	constant
LIF time constant	2.0	2.0	2.0
LIF threshold	1.0	0.75	0.75
Output interface	sigmoid mixture	direct mixture	direct mixture

Table 5: Historical optimization and neuron settings. These are not the optimizer roles or selection protocol of the new CIFAR-100 cohort in Appendix D.

A.1 OUTPUT-MODE MAP

Table 6 records the numerical boundary used by each evidence block. This distinction is necessary because a normalized mixture of binary candidates is continuous even before an optional sigmoid. Within every matched mixture comparison, all methods use the same output mode.

Evidence block	Encoder output	Interpretation
Residual CIFAR benchmark/factorization	$\sigma(Z_t)$	continuous current
ImageNet32 mixture family	Z_t	continuous current
CIFAR-100 / VGG-11 mixture family	Z_t	continuous current
CIFAR-10 Binary-STE	$\mathbb{1}[Z_t \geq 0.5]$	binary events
Individual Rate/Burst/TTFS	candidate E_t^e	binary events

Table 6: Output interfaces used by the reported experiments.

A.2 SOURCE-DATA SCOPE

The A10 files `a10_official_test_summary.csv` and `a10_official_test_paired_differences.csv` agree cell-for-cell with the returned server summaries. Appendix D.1 describes their completed record-level reconciliation. The anonymous evidence supplement supplies per-model records and full-precision summaries separately from the compilation source archive.

Eight historical CSVs cover ImageNet32, VGG-11, residual CIFAR, component, Binary-STE, gate-diagnostic, and capacity-screen results. Means and sample standard deviations are recomputed wherever seed values are available. ImageNet32 top-5 and benchmark/component aggregates retain their packaged summaries. Each family specifies its selection rule; selected top-5 follows the top-1-selected epoch. Appendix I separately reports a new validation-only cost/activity audit of the frozen CIFAR-100 cohort. Uncertainty is reported at each study's stated sampling level.

B FORWARD ALGORITHM

The code stores candidate channels in Rate–TTFS–Burst order. The names r_t and h_t below denote learned responses; neither filter is constrained to implement a predefined frequency response.

Algorithm 1 PAC-inspired FAPE encoder with explicit output mode

Require: Image x, timesteps T, code set $\mathcal{E} = \{\mathrm{rate}, \mathrm{ttfs}, \mathrm{burst}\}$, output mode o

$\hat{x} \leftarrow \mathrm{MinMax}_{c,h,w}(x)$
Generate $E^{\mathrm{rate}}, E^{\mathrm{ttfs}}, E^{\mathrm{burst}}$ from $\hat{x}$
$\bar{x} \leftarrow \mathrm{mean}_c(x)$
for $t = 0$ to $T-1$ **do**
  $F_t \leftarrow K_t * \bar{x}$; $a_t \leftarrow \mathrm{GAP}(F_t)$
end for
$r \leftarrow \mathcal{P}_T(W_c * a)$; $h \leftarrow \mathcal{P}_T(|W_r * a|)$
for $t = 0$ to $T-1$ **do**
  $G_t \leftarrow \mathrm{softmax}_e(C_t(F_t))$
  $p_t \leftarrow \sigma(\gamma_c r_t)$; $m_t \leftarrow \sigma(\gamma_h h_t)$
  $(\eta_t^{\mathrm{rate}}, \eta_t^{\mathrm{ttfs}}, \eta_t^{\mathrm{burst}}) \leftarrow (1, p_t, m_t)$
  $\tilde{w}_t^e \leftarrow G_t^e q_e \eta_t^e$; $w_t^e \leftarrow \tilde{w}_t^e / (\sum_j \tilde{w}_t^j + \epsilon)$
  $Z_t \leftarrow \sum_{e \in \mathcal{E}} w_t^e E_t^e$
  if $o =$ MIXTURE **then**
    $Y_t \leftarrow Z_t$
  else if $o =$ SIGMOID **then**
    $Y_t \leftarrow \sigma(Z_t)$
  else
    $Y_t \leftarrow \mathbb{1}[Z_t \geq 0.5]$ with a straight-through gradient
  end if
end for
return $Y_{0:T-1}$

C GATE REPARAMETERIZATION AND NUMERICAL NORMALIZATION

Let ℓ_t^e denote the logits before the spatial softmax and let $S_t = \sum_e G_t^e q_e \eta_t^e$. For positive priors and factors, the implemented weights satisfy

$$w_t^e = \frac{S_t}{S_t + \epsilon}\ \mathrm{softmax}_e(\ell_t + \log q + \log \eta_t). \tag{16}$$

For the input-independent generic factors, $\log \eta_t$ can be absorbed into the bias of the step-specific 1×1 gate convolution. This identifies a reparameterization of relative mixture weights, not identical optimization trajectories. With finite ϵ, the factor $S_t/(S_t + \epsilon)$ also affects output magnitude and can vary with the image. Exact trained-checkpoint equivalence has not been established. Without temporal modulation, $S_t \geq 0.33$, giving a normalization deficit below 3.04×10^{-6}. The same lower bound need not hold for modulated gates.

D VALIDATION-SELECTED CIFAR-100 PROTOCOL

Status and separation from historical evidence. This cohort contains ten methods and seeds $\{42, 1337, 2024\}$, with 150 epochs per run. All 30 validation-selected models were frozen on 11 September 2026, and official-test evaluation completed on 12 September. Table 1 reports summaries reconciled with the returned per-model records. The official test set had informed historical development; this cohort uses validation-only model selection and post-freeze testing. Historical CIFAR tables retain their separate selection, optimization, and output-interface protocols.

Data and random streams. The 50,000 official training images are stratified into 45,000 training and 5,000 validation images, using split seed 27001. Each class contributes 450 training and 50 validation images. The 10,000 official test images are excluded from training and selection. Training applies a horizontal flip with probability 0.5, zero-padding by four pixels, and a random 32×32

crop, in that order. All splits use channel means $(0.5071, 0.4867, 0.4408)$ and standard deviations $(0.2675, 0.2565, 0.2761)$. Augmentation is keyed by seed, epoch, and original image index; sampling order and worker initialization use separate random streams. The same training seed fixes the initial backbone across methods. These controls do not imply identical subsequent optimization trajectories.

Mixture controls. All seven candidate-mixture methods use the same binary code bank, $T = 8$, branch prior, and continuous output $Y_t = Z_t$. The first two use equal gates (`uniform`) or learned input-independent per-step gate vectors (`global`). The next three share the spatial feature extractor and gate parameterization:

- `input_global` spatially averages the already-softmaxed gate probabilities and broadcasts them: $G_t^{\mathrm{IG}} = \mathrm{Broadcast}(\mathrm{GAP}(G_t))$. It does not average logits before softmax.
- `static_input` drives the gate with one fixed random $1 \times 1 \times 32 \times 32$ texture, generated with seed 314159 and shared across images. Candidate codes still come from the actual image.
- `spatial` uses the image-dependent local gate without an additional temporal conditioner. `generic` and `pac` add the factors defined in Section 3.

The two prespecified primary contrasts are `spatial` minus `input_global` and `spatial` minus `static_input`. They target spatial variation and image dependence within this gate construction; they do not identify a general frequency-specific mechanism.

Continuous-input controls. `direct_standard` repeats the standardized image x over eight steps, whereas `direct_minmax` repeats $\hat{x}$ from Equation 3. The former is not restricted to $[0, 1]$. `learned_current` retains the spatial gate but replaces event candidates with three continuous bases,

$$B_1(z) = z, \quad B_2(z) = z^2, \quad B_3(z) = \frac{\sqrt{z + \epsilon_c} - \sqrt{\epsilon_c}}{\sqrt{1 + \epsilon_c} - \sqrt{\epsilon_c}}, \qquad \epsilon_c = 10^{-6}, \tag{17}$$

where $z = \hat{x}$. Its output is $\sum_{j=1}^{3} G_t^j B_j(z)$, without the code prior or denominator shrinkage. These controls test alternative input interfaces, not only the location of code allocation. Active encoder parameter counts are 0 for each direct input and the fixed mixture, 24 for global, 808 for input-global, static-input, spatial, and learned-current, 824 for generic, and 819 for PAC-inspired. Equal parameter counts do not make their computations identical.

Backbone and optimization. The backbone is CIFAR-stem Spiking ResNet-18 with three input channels, 100 classes, multi-step LIF neurons, $\tau_{\mathrm{LIF}} = 2$, threshold 0.75, and detached zero hard reset. The stem uses a 3×3, stride-one, 64-channel convolution, batch normalization, and identity max-pooling. SpikingJelly 0.0.0.0.14 supplies the sigmoid surrogate with slope parameter 4 and input decay enabled. Neuronal states reset before each forward pass, and class logits are averaged over time. Cross-entropy training uses batch size 128, eight data workers, and SGD with momentum 0.9, Nesterov momentum, and weight decay 5×10^{-4}. Backbone and scale-filter parameters have base learning rate 0.1; allocation and modulation parameters have base rate 10^{-4}. Generic temporal logits belong to allocation. Every active role uses the same five-epoch linear warm-up and cosine multiplier, ending at 10^{-3} of its base rate. Unused encoder parameters are frozen. Historical protocols instead use a constant spatial-gate learning rate.

Numerical amendment and recovery. The encoder executes in FP32, while the backbone uses automatic mixed precision with initial loss scale 128. The original finite-gradient guard stopped three uniform runs after epochs 60, 57, and 57. The amended protocol lets GradScaler skip an overflowing optimizer update and reduce its scale. The next minibatch proceeds without replaying the skipped batch; non-finite parameter or buffer states still cause failure. More than 16 consecutive skipped updates or a skip fraction above 5% in an epoch also causes failure. Thus 150 epochs fixes data passes, not the number of successful optimizer updates. Per-epoch and cumulative attempted, applied, and skipped update counts are retained.

The three uniform runs resumed from their retained finite checkpoints; the other 27 runs started from initialization under the amended protocol. For resumed runs, numerical audits cover only the

post-migration segment; historical state lineage is retained separately. Migration does not relabel historical epochs as runs of the amended code. Production preflight and recovery checks precede the resumed cohort. The amendment changes overflow handling, not methods, seeds, splits, the learning-rate schedule, or the checkpoint-selection rule.

Selection and final reporting. Validation evaluates five encoding draws with seeds $\{101, 211, 307, 401, 503\}$ without perturbing the training random state. The highest mean validation top-1 selects the checkpoint, with the earliest epoch breaking ties. Freezing binds all 30 selected model artifacts before official testing. Final evaluation uses the frozen cohort and five draws per model; the report first averages within each model and then summarizes across the three training seeds. Encoding draws are repeated measurements, not 15 independently trained models. Primary comparisons use same-seed paired differences. Deterministic interfaces may have zero between-draw variability without implying zero training-seed variability. No checkpoint is selected or replaced using official-test scores.

D.1 Prespecified paired official-test comparisons

Table 7 retains all three paired differences for each primary contrast. Each model score is the mean of five encoding draws. The positive directions across three seeds support the within-family allocation comparison in this cohort; they are not a significance test or a guarantee of transfer to other settings.

Table 7: Same-seed official-test top-1 differences in percentage points, transcribed from the author-provided paired summary. The bottom rows summarize three training-seed differences, not fifteen encoding draws.

Training seed	Spatial − input-global	Spatial − fixed-input
42	+2.576	+1.814
1337	+1.412	+3.640
2024	+1.708	+4.084
Mean	+1.899	+3.179
Sample SD	0.605	1.203

Returned-artifact reconciliation. Reconciliation passed 4,065 checks covering 37 transferred files, 30 model records, and the 58-file source inventory. Recorded aggregates, paired differences, selected epochs, and identity bindings agree with the supplied summaries. This record-level audit recomputes saved metrics; checkpoint tensors and dataset bytes were outside the transfer. The anonymous evidence supplement includes sanitized records, full-precision tables, the reconciliation report, and a numerical verification script.

E Historical seed-level results

E.1 ImageNet32

ImageNet32 entries are selected at the epoch with the highest top-1 validation accuracy within each 90-epoch run. The aggregate top-5 values in Table 2 are taken from those same top-1-selected epochs, not independently maximized. The packaged summary contains aggregate, but not seed-level, top-5 values.

Configuration	Seed 42	Seed 1337	Seed 2024
Rate	10.01	10.14	10.10
Burst	12.58	12.88	12.85
TTFS	13.68	13.91	13.63
Fixed prior $(.33/.33/.34)$	14.46	14.47	14.66
Global learnable mixture	16.78	16.43	16.27
FAPE (PAC-inspired)	19.05	18.98	18.56
FAPE (generic temporal)	19.52	18.79	18.92
FAPE (spatial only)	18.98	19.36	19.72

Table 8: Seed-level ImageNet32 top-1 validation accuracies (%).

E.2 RESIDUAL CIFAR FACTORIZATION

Dataset	Configuration	Seed 42	Seed 1337	Seed 2024
CIFAR-10	FAPE (spatial only)	92.64	92.50	92.58
	FAPE (generic temporal)	92.25	92.24	92.24
	FAPE (PAC-inspired)	92.46	92.02	90.31
CIFAR-100	FAPE (spatial only)	72.29	71.76	71.86
	FAPE (generic temporal)	72.02	72.14	72.30
	FAPE (PAC-inspired)	71.75	71.73	71.72

Table 9: Historical residual CIFAR factorization: test-selected best-epoch top-1 evaluation accuracy (%). Comparisons use a common protocol within each dataset.

E.3 SPIKING VGG-11 TRANSFER

Configuration	Seed 42	Seed 1337	Seed 2024
Fixed prior $(.33/.33/.34)$	61.95	62.14	62.22
FAPE (spatial only)	64.94	66.63	65.99
FAPE (PAC-inspired)	65.68	66.09	66.09

Table 10: Historical Spiking VGG-11 results: test-selected best-epoch CIFAR-100 top-1 evaluation accuracy (%).

F AUXILIARY CIFAR EXPERIMENTS

The CIFAR single-code benchmark, factorization matrix, and component study use distinct protocols and are interpreted within their respective groups. The packaged files provide summaries and available seed values. Run-bound logs and checkpoints have partial coverage for the single-code benchmark, component study, and seed-42 capacity screen. These historical blocks provide secondary descriptive evidence. The primary allocation contrasts use the validation-selected A10 cohort; ImageNet32, residual factorization, and VGG-11 provide additional experimental settings.

F.1 SINGLE-CODE BENCHMARK GROUP

Encoder	CIFAR-10	CIFAR-100
Burst	90.56 ± 0.23	67.50 ± 0.15
TTFS	89.40 ± 0.16	65.85 ± 0.30
Rate	84.30 ± 0.24	58.14 ± 0.11
FAPE (PAC-inspired) + CE	92.40 ± 0.18	71.41 ± 0.07
FAPE (PAC-inspired) + ISL	92.28 ± 0.08	71.41 ± 0.07

Table 11: Historical residual CIFAR single-code benchmark. Values are mean $\pm$ standard deviation of best-epoch evaluation accuracy over three runs. FAPE uses the sigmoid-mixture interface, whereas single codes are binary. This is end-to-end context, not a matched allocation ablation.

F.2 SEED-42 MIXTURE SCREEN

At seed 42, the fixed-prior and globally learned mixtures both obtain 90.45% on CIFAR-10, while multiscale spatial allocation obtains 92.64%. This single-seed screen is not parameter matched and does not isolate capacity from input or spatial dependence. The three-seed ImageNet32 mixture family supplies the main historical configuration comparison.

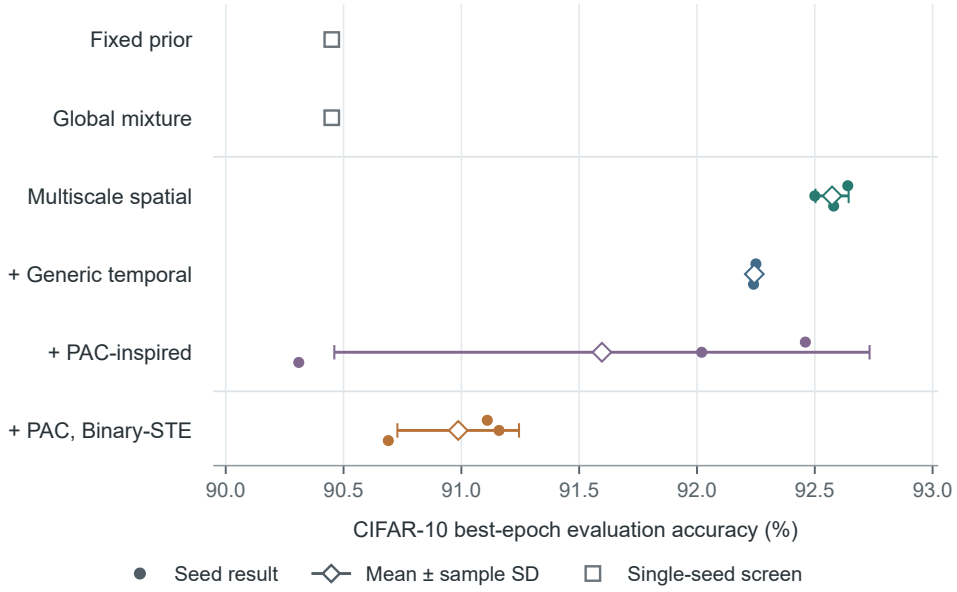


Figure 2: Historical CIFAR-10 allocation and interface comparisons, using best-epoch evaluation accuracy. Hollow squares show the single-seed mixture screen; dots are seed-level results, diamonds are means, and whiskers are sample standard deviations. The Binary-STE row uses an event-valued boundary.

F.3 COMPONENT STUDY GROUP

Variant	Accuracy (%)	Within-group Δ
FAPE (PAC-inspired) + ISL	92.28 ± 0.08	–
FAPE (PAC-inspired) + CE only	92.11 ± 0.31	-0.17
No sigmoid output	92.15 ± 0.39	-0.13
No multiscale sequence	92.05 ± 0.18	-0.23
No response rectification	$91.83 \pm 0.10^{\dagger}$	-0.45
Fixed spatial gate	89.52 ± 0.21	-2.76

Table 12: Historical CIFAR-10 component study using best-epoch evaluation accuracy. Δ is measured only against the PAC-inspired FAPE+ISL row in this run group. †Two completed runs are summarized; a third run diverged. This failure is not evidence of a universal stability effect.

The CE and ISL rows have similar mean accuracy in this study. Without direct measurements of retained information or activity reduction, we make no empirical claim that ISL improves either quantity; it remains an optional objective defined in Section 3.

G BINARY-STE SEED RESULTS

The Binary-STE variant obtains 91.11, 91.16, and 90.69 on CIFAR-10 for seeds 42, 1337, and 2024, respectively ($90.99 \pm 0.26\%$). The continuous PAC-inspired configuration in the same factorization group obtains 92.46, 92.02, and 90.31. The signed pairwise differences are therefore not uniform across seeds, which is why the main text presents Binary-STE as an interface feasibility result rather than a fixed accuracy cost.

H GATE–FREQUENCY DIAGNOSTIC PROTOCOL

The quantitative gate analysis uses the seed-42 CIFAR-10 checkpoints and 20 evaluation mini-batches of 64 images, for 1,280 images per variant. We first undo input normalization and convert RGB values to grayscale as $0.2989R + 0.5870G + 0.1140B$. A 5×5 Gaussian filter with standard deviation 1.0 defines the low-frequency map; the high-frequency statistic is the absolute residual between grayscale input and its blurred version. Effective branch weights from Equation 9 are averaged over the eight simulation steps. For each image and branch, we compute the spatial Pearson correlation between the resulting weight map and the image statistic. Cases with a zero denominator are recorded as undefined, and only finite correlations enter the aggregate. Table 4 reports the across-image mean and normal-approximation half-width, $1.96s/\sqrt{n}$, for both statistics, using the retained sample count n for each entry. All nine branch–variant combinations are retained, and the table is not repeated here. This is an association diagnostic rather than a significance test or causal intervention.

For exactly normalized weights and any statistic H, $\sum_e \mathrm{Cov}(w^e, H) = 0$. Opposite branch associations can therefore arise from the shared normalization constraint; the correlation coefficients themselves need not sum to zero. The finite-ϵ weights additionally include the scale factor described in Appendix C. These relationships do not identify which gate variation causes an accuracy difference.

I CIFAR-100 OPERATION AND ACTIVITY AUDIT

We profiled all 30 frozen CIFAR-100 models from Appendix D on the 5,000-image validation split, using five encoding draws per model. This post-hoc diagnostic used training-derived validation data throughout and left the frozen checkpoints unchanged. Each model processed 25,000 image-draws over 200 batches: 195 of size 128 and five of size eight. Table 13 averages draws within each model and reports means and sample standard deviations across three training seeds.

Operation counts use executed Conv1d, Conv2d, and Linear call shapes, normalized per image across all eight steps. Continuous-input operations contribute dense MACs. For each binary-input call, estimated SynOps equal its dense MAC potential multiplied by the input's global nonzero fraction. This activity-scaled estimate approximates event-driven accumulation without spatial fan-out weighting. Counts exclude bias, normalization, pooling, LIF scalar dynamics, residual additions, activations, random-number generation, and memory movement. Zero encoder Conv/Linear MACs therefore describe only the counted operator class.

Input configuration	Active encoder parameters	Encoder MACs (M/image)	Estimated binary SynOps (M/image)	LIF activity (%)
Fixed mixture	0	0	357.77 ± 9.12	10.53 ± 0.43
Global learnable mixture	24	0	375.64 ± 5.93	11.24 ± 0.28
Input-global gate	808	0.802816	381.18 ± 13.66	11.57 ± 0.47
Fixed-input gate	808	0.802816	382.90 ± 16.32	11.79 ± 0.68
FAPE (spatial only)	808	0.802816	408.63 ± 12.66	12.71 ± 0.39
FAPE (generic temporal)	824	0.802816	415.23 ± 9.98	12.72 ± 0.23
FAPE (PAC-inspired)	819	0.802908	420.38 ± 25.64	13.13 ± 0.83
Direct standardized current	0	0	494.93 ± 3.49	16.10 ± 0.17
Direct min–max current	0	0	487.49 ± 6.99	15.76 ± 0.26
Learned continuous current	808	0.802816	452.05 ± 4.98	14.39 ± 0.18

Table 13: CIFAR-100 validation-only cost and activity audit of the frozen A10 cohort. M denotes 10^6. MACs cover encoder Conv/Linear operations; SynOps estimate internal binary-input backbone operations. LIF activity is the aggregate spike fraction over recorded neuronal outputs. Entries with variability are mean $\pm$ sample standard deviation across three training seeds. Active parameter and MAC counts are identical across seeds. Official-test accuracy is reported separately in Table 1.

All models have 11,220,132 active backbone parameters and 14.565376M continuous backbone MACs per image. The first convolution contributes 14.155776M and the pooled classifier

0.409600M. Their $4{,}443.7504$M backbone dense-MAC potential is a separate all-dense reference; adding it to the mixed-operation counts would double-count computation. Non-direct encoder objects store 859 parameters, with Table 13 counting only parameters enabled by each configuration.

Spatial, input-global, and fixed-input allocation match in active parameters and counted encoder MACs. Spatial allocation has higher estimated internal SynOps than these mixture controls and lower estimates than direct-current inputs, alongside the accuracy ordering in Table 1. The seven mixture variants share a 70.88% nonzero continuous-input fraction under the same images and draws. Positive mixture weights preserve the candidates' union of active positions; allocation changes amplitudes and downstream activity. This occupancy describes continuous inputs, while LIF activity describes binary neuronal outputs. All three candidates execute in every mixture variant.

The anonymous supplement includes per-model layer counts, activity aggregates, full-precision summaries, and a record-level checker. Reconciliation passed 21,960 checks of cohort identities, call-shape arithmetic, denominators, and aggregate statistics; local verification used saved records and source bytes. These are analytical operation counts and measured activity summaries. Hardware latency and energy were not measured.

J Temporal controls and equal-histogram gate probes

This follow-up protocol was fixed on 14 September 2026 after the original CIFAR-100 official-test summaries had been observed. It comprises frozen-model temporal interventions, nine matched retraining runs, and synthetic gate probes. All recognition scores in this appendix use the same 5,000-image validation partition from the official CIFAR-100 training set. Checkpoint selection also uses this partition. The follow-up contains no official-test evaluation. Each model is evaluated with five encoding draws, with draw seeds $\{101, 211, 307, 401, 503\}$; means and sample standard deviations across training seeds use $n = 3$, with seeds $\{42, 1337, 2024\}$.

J.1 Frozen-model temporal interventions

We select the spatial, input-global, and fixed-input models from the original frozen cohort. Before intervention, all nine identity evaluations reproduce their stored selected validation top-1 exactly at the recorded precision; the acceptance tolerance is 10^{-8} percentage points. The original encoder and backbone weights remain frozen. Transformations act on the realized FP32 post-mixture sequence $Z \in \mathbb{R}^{8\times N\times C\times H\times W}$ before the first backbone convolution:

$$Y_t^{\text{original}} = Z_t, \qquad Y_t^{\text{reverse}} = Z_{7-t}, \tag{18}$$

$$Y_t^{\text{perm}} = Z_{\pi_t}, \qquad Y_t^{\text{mean}} = \frac{1}{8}\sum_{s=0}^{7} Z_s, \tag{19}$$

where $\pi = (2, 5, 0, 7, 3, 1, 6, 4)$ is a single fixed permutation and the mean is repeated at every step. All conditions use common encoding draw seeds and validation images. The reverse and permutation conditions change input order; mean repetition removes across-step variation and changes stochastic fluctuations while preserving the integrated input up to floating-point precision. All conditions retain eight-step SNN dynamics.

Allocation	Input	Seed 42	Seed 1337	Seed 2024	Top-1	Top-5
Spatial	Original	69.152	68.720	69.632	69.17 ± 0.46	90.53 ± 0.34
	Reverse	42.284	61.456	59.520	54.42 ± 10.55	80.26 ± 7.58
	Permutation	49.504	58.620	56.848	54.99 ± 4.83	81.43 ± 3.60
	Mean repeat	39.208	51.608	49.260	46.69 ± 6.59	72.85 ± 7.03
Input-global	Original	66.556	68.124	67.996	67.56 ± 0.87	89.34 ± 0.51
	Reverse	56.696	57.400	52.968	55.69 ± 2.38	80.86 ± 2.26
	Permutation	59.512	57.524	54.952	57.33 ± 2.29	82.74 ± 2.05
	Mean repeat	50.404	44.540	49.972	48.31 ± 3.27	74.75 ± 3.00
Fixed-input	Original	67.740	66.016	65.696	66.48 ± 1.10	88.73 ± 0.49
	Reverse	52.624	52.628	54.604	53.29 ± 1.14	79.40 ± 0.98
	Permutation	54.232	61.452	58.120	57.93 ± 3.61	83.15 ± 2.71
	Mean repeat	47.976	49.576	54.528	50.69 ± 3.42	77.49 ± 3.27

Table 14: Complete frozen-model temporal diagnostic. The three seed columns give validation top-1 (%); the last two columns give mean $\pm$ sample standard deviation across three training seeds. Each seed value averages five paired encoding draws. No parameters are updated between conditions.

All three interventions reduce mean validation accuracy for each allocation variant (Table 14). For spatial allocation, the mean drops are 14.75, 14.18, and 22.48 points for reversal, permutation, and mean repetition, respectively. These measurements characterize sensitivity of frozen models to an altered input distribution. The following training-matched experiment evaluates how the models adapt to the repeated-mean presentation.

J.2 Matched repeated-mean retraining

We train spatial, input-global, and fixed-input encoders from fresh paired initializations, with three seeds per arm. The post-mixture transformation Y_t^{mean} is present throughout training and validation. Gradients propagate through the temporal average; the transformation adds no parameters, detachment, or normalization. Each encoder retains 808 active parameters. The 45,000/5,000 split, augmentations, batch size 128, optimizer groups, learning-rate schedules, neuron settings, and 150-epoch budget match the amended protocol in Appendix D. The highest five-draw validation mean selects the checkpoint, with the earliest epoch breaking ties. Initial backbone hashes match across arms and the corresponding original training seed.

All nine runs complete 150 epochs. Each run attempts 52,800 optimizer steps, with 13–16 isolated AMP overflow skips handled by the fixed skip/backoff policy. The largest consecutive-skip count is one, and all applied updates pass the finite-gradient guard. Table 15 reports every selected epoch and score.

Allocation	Seed	Epoch	Top-1	Top-5	Original Top-1	Change (pp)
Spatial	42	146	70.880	91.076	69.152	+1.728
	1337	137	71.952	91.432	68.720	+3.232
	2024	146	71.492	91.648	69.632	+1.860
Input-global	42	139	70.472	90.796	66.556	+3.916
	1337	143	70.440	91.008	68.124	+2.316
	2024	146	70.872	91.208	67.996	+2.876
Fixed-input	42	149	71.116	91.452	67.740	+3.376
	1337	144	70.964	90.872	66.016	+4.948
	2024	150	70.676	91.132	65.696	+4.980

Table 15: Nine repeated-mean training runs. Top-1 and top-5 are five-draw validation accuracies (%). "Original" denotes the corresponding frozen model's identity validation score. Change is retrained minus original, in percentage points. All nine changes are positive.

Allocation	Original Top-1	Retrained Top-1	Retrained Top-5
Spatial	69.17 ± 0.46	71.44 ± 0.54	91.39 ± 0.29
Input-global	67.56 ± 0.87	70.59 ± 0.24	91.00 ± 0.21
Fixed-input	66.48 ± 1.10	70.92 ± 0.22	91.15 ± 0.29

Table 16: Validation-selected results before and after repeated-mean retraining, summarized across the same three training seeds as mean $\pm$ sample standard deviation (%). Original and retrained models use the same validation split and selection rule.

Paired allocation contrasts. Let $a^c_{m,s}$ be selected validation top-1 for method m, seed s, and training condition $c \in \{\text{original}, \text{mean}\}$. For control b, we report $d^c_{b,s} = a^c_{\text{spatial},s} - a^c_{b,s}$ and its paired change $d^{\text{mean}}_{b,s} - d^{\text{original}}_{b,s}$. Table 17 reports both prespecified controls and all seeds.

Spatial minus	Condition	Seed 42	Seed 1337	Seed 2024	Mean $\pm$ SD
Input-global	Original	$+2.596$	$+0.596$	$+1.636$	1.609 ± 1.000
	Retrained	$+0.408$	$+1.512$	$+0.620$	0.847 ± 0.586
	Change in contrast	-2.188	$+0.916$	-1.016	-0.763 ± 1.567
Fixed-input	Original	$+1.412$	$+2.704$	$+3.936$	2.684 ± 1.262
	Retrained	-0.236	$+0.988$	$+0.816$	0.523 ± 0.663
	Change in contrast	-1.648	-1.716	-3.120	-2.161 ± 0.831

Table 17: Paired allocation differences in validation top-1 percentage points. "Change in contrast" is retrained minus original. All uncertainty entries are sample standard deviations of the three paired differences. These validation contrasts are distinct from the original cohort's official-test contrasts.

Repeated-mean training narrows the mean spatial–input-global contrast by 0.76 points and the spatial–fixed-input contrast by 2.16 points. The first change varies in sign across seeds; the second is negative in all three. After retraining, spatial allocation exceeds input-global allocation in each seed and exceeds fixed-input allocation in two of three seeds. The resulting pattern supports a local-allocation contribution under both input presentations, with its magnitude depending on the complete temporal-averaging and retraining intervention. This descriptive follow-up reports selected validation outcomes; the paired values and their variability specify the observed effect across the three training replications.

J.3 Equal-histogram synthetic gate probes

We probe the three frozen original spatial encoders with a fixed bank of 29 synthetic RGB images of size 32×32. Every channel contains 512 pixels at intensity 64 and 512 at intensity 192, followed by the same CIFAR-100 channel standardization. The bank comprises vertical stripes, horizontal stripes, and checkerboards, each at widths $w \in \{1, 2, 4, 8\}$ pixels and offsets $p \in \{0, 1\}$ pixel, plus five fixed histogram-preserving shuffles with seeds $\{101, 211, 307, 401, 503\}$. Offset one means a one-pixel displacement at every width. These artificial inputs control the global intensity histogram while varying arrangement and local neighborhoods.

The measured quantity is the raw three-way softmax distribution $g_{s,j,t,h,\ell}$ before prior multiplication and mixture normalization. Here s indexes training seeds, j indexes probes, and (h, ℓ) are pixel coordinates. The actual spatial mixture weights subsequently follow Equation 9 with unit temporal factors. Probe $j = 0$ is the width-one vertical stripe pattern at offset zero. For each model and step, spatially average the gate to obtain $\bar{g}_{s,j,t}$. We calculate

$$L_1^{\text{avg}}(s, j) = \frac{1}{T} \sum_t \|\bar{g}_{s,j,t} - \bar{g}_{s,0,t}\|_1, \tag{20}$$

$$L_1^{\text{pos}}(s, j) = \frac{1}{THW} \sum_{t,h,\ell} \|g_{s,j,t,h,\ell} - g_{s,0,t,h,\ell}\|_1. \tag{21}$$

The corresponding $\mathrm{JS}^{\mathrm{avg}}$ and $\mathrm{JS}^{\mathrm{pos}}$ replace the L1 distance by Jensen–Shannon divergence, $\mathrm{JS}(u,v) = \frac{1}{2}\mathrm{KL}(u\|m) + \frac{1}{2}\mathrm{KL}(v\|m)$ with $m = (u+v)/2$, using natural logarithms. Same-coordinate metrics compare identical array positions without inverse-shift alignment. Each metric is aggregated over the three trained encoders as mean and sample standard deviation. Pattern identities and time steps are repeated probes within an encoder.

Table 18 reports the full bank in generation order. Raw-gate responses vary across all non-reference patterns. The width-one vertical phase shift has a small spatial-average L1 distance (0.0031) and a large same-coordinate distance (1.1735), illustrating the distinction between aggregate code preference and local gate arrangement. The diagnostic measures encoder responses to synthetic patterns; the classifier is not evaluated, and no recognition score is attached to these probes.

Pattern	(w,p) / seed	L_1^{avg}	L_1^{pos}	$\mathrm{JS}^{\mathrm{avg}}$	$\mathrm{JS}^{\mathrm{pos}}$
Vertical	(1,0)	0.000 ± 0.000	0.000 ± 0.000	0.00000 ± 0.00000	0.00000 ± 0.00000
Vertical	(1,1)	0.003 ± 0.002	1.174 ± 0.031	0.00001 ± 0.00001	0.32326 ± 0.02082
Vertical	(2,0)	0.079 ± 0.043	0.677 ± 0.063	0.00540 ± 0.00525	0.16722 ± 0.01500
Vertical	(2,1)	0.074 ± 0.041	0.670 ± 0.062	0.00473 ± 0.00472	0.16601 ± 0.01458
Vertical	(4,0)	0.137 ± 0.074	0.695 ± 0.067	0.01337 ± 0.00856	0.15403 ± 0.01805
Vertical	(4,1)	0.131 ± 0.066	0.679 ± 0.065	0.01221 ± 0.00782	0.15098 ± 0.01735
Vertical	(8,0)	0.175 ± 0.093	0.725 ± 0.066	0.01844 ± 0.01131	0.15164 ± 0.01957
Vertical	(8,1)	0.169 ± 0.087	0.711 ± 0.064	0.01723 ± 0.01054	0.14908 ± 0.01912
Horizontal	(1,0)	0.044 ± 0.052	0.655 ± 0.081	0.00247 ± 0.00402	0.17173 ± 0.01667
Horizontal	(1,1)	0.045 ± 0.049	0.655 ± 0.080	0.00244 ± 0.00392	0.17142 ± 0.01622
Horizontal	(2,0)	0.083 ± 0.032	0.672 ± 0.056	0.00497 ± 0.00328	0.16204 ± 0.00850
Horizontal	(2,1)	0.078 ± 0.028	0.671 ± 0.057	0.00430 ± 0.00266	0.16303 ± 0.00814
Horizontal	(4,0)	0.140 ± 0.076	0.702 ± 0.063	0.01327 ± 0.00847	0.15427 ± 0.01518
Horizontal	(4,1)	0.131 ± 0.066	0.699 ± 0.062	0.01183 ± 0.00735	0.15502 ± 0.01445
Horizontal	(8,0)	0.177 ± 0.096	0.727 ± 0.065	0.01860 ± 0.01144	0.15113 ± 0.01860
Horizontal	(8,1)	0.168 ± 0.090	0.723 ± 0.062	0.01700 ± 0.01032	0.15173 ± 0.01722
Checker	(1,0)	0.126 ± 0.063	0.649 ± 0.035	0.01180 ± 0.00885	0.13993 ± 0.01383
Checker	(1,1)	0.126 ± 0.063	0.646 ± 0.037	0.01180 ± 0.00885	0.13951 ± 0.01457
Checker	(2,0)	0.064 ± 0.032	0.677 ± 0.057	0.00307 ± 0.00230	0.17185 ± 0.00079
Checker	(2,1)	0.065 ± 0.033	0.680 ± 0.058	0.00314 ± 0.00237	0.17311 ± 0.00104
Checker	(4,0)	0.084 ± 0.028	0.675 ± 0.043	0.00516 ± 0.00265	0.16058 ± 0.00594
Checker	(4,1)	0.080 ± 0.026	0.675 ± 0.039	0.00474 ± 0.00244	0.16171 ± 0.00524
Checker	(8,0)	0.141 ± 0.078	0.702 ± 0.052	0.01310 ± 0.00777	0.15181 ± 0.01429
Checker	(8,1)	0.133 ± 0.073	0.698 ± 0.048	0.01205 ± 0.00708	0.15228 ± 0.01336
Shuffle	101	0.060 ± 0.017	0.664 ± 0.037	0.00279 ± 0.00171	0.16563 ± 0.00878
Shuffle	211	0.061 ± 0.014	0.680 ± 0.038	0.00277 ± 0.00140	0.16932 ± 0.00799
Shuffle	307	0.060 ± 0.017	0.651 ± 0.038	0.00269 ± 0.00169	0.16227 ± 0.00750
Shuffle	401	0.061 ± 0.016	0.677 ± 0.035	0.00270 ± 0.00152	0.16968 ± 0.00876
Shuffle	503	0.062 ± 0.017	0.671 ± 0.040	0.00294 ± 0.00191	0.16729 ± 0.00827

Table 18: All 29 equal-histogram probes on the three frozen original spatial encoders. Values are mean $\pm$ sample standard deviation across training seeds. w and p denote width and one-pixel offset; shuffle rows give the permutation seed. The first row is the self-reference and is identically zero. JS entries use nats, with five-decimal rounding; full-precision values are included in the source CSV. Distances measure raw gates, with no inverse-shift alignment for same-coordinate metrics.

Artifact reconciliation. The returned records comprise all 36 frozen-model conditions, all 87 seed–probe combinations, and 1,350 epoch records from nine new runs. Reconciliation recomputes five-draw aggregates, training-seed summaries, paired differences, and earliest-maximum checkpoint selections, and checks source, protocol, split, and receipt bindings. The results-only transfer contains tensor digests and paths, with checkpoint and raw-gate tensors excluded. The anonymous evidence supplement provides sanitized records, full-precision derived tables, and numerical verification instructions; its audit covers saved metrics and metadata.