BIO-INSPIRED AUDITORY PROCESSING: COUPLING PARALLEL RESONATE-AND-FIRE DYNAMICS WITH TERNARY RECURRENT NETWORKS FOR SPEECH RECOGNITION

Anonymous authors
Paper under double-blind review

ABSTRACT

Spiking Neural Networks (SNNs) promise energy-efficient inference through sparse, event-driven computation, yet remain constrained on continuous, context-dependent tasks such as automatic speech recognition (ASR). A key bottleneck is the Leaky Integrate-and-Fire (LIF) neuron, whose membrane decay erases information across pauses and phonetic boundaries while it must both extract features and retain long-range context. Inspired by cortical separation of nonlinear filtering and somatic integration, we introduce the Spiking Ternary GRU (stGRU): a ternary-weight gated recurrent unit performs temporal filtering, while downstream Resonate-and-Fire (RF) neurons generate the binary spike stream driving the next block's event pathway. This makes binary event transmission a native training constraint rather than a separate deployment-time activation quantization step; together with ternary weights, learned projections avoid dense floating-point matrix MACs. Trained end-to-end on LibriSpeech-960, Small, Medium, and Large stGRUs reach 8.09%, 7.04%, and $6.44 \pm 0.17\%$ WER on test-clean without an external language model. To our knowledge, these are the first matrix-MAC-free spiking networks without a learned ANN sequence decoder to achieve single-digit WER on LibriSpeech. The architecture also achieves 10.88% test CER on Aishell-1 (9.28% on dev) when trained from scratch. Internally, trained networks develop a depth-graded hierarchy: shallow layers act as fast acoustic filters, while deep layers form slower semantic integrators whose latent representations segregate into part-of-speech manifolds and word-level spike events. Deep-layer firing rates follow lognormal distributions with a sparsity gradient, mirroring hierarchical processing observed in human cortex in vivo. Our biologically grounded architecture combines competitive ASR with favorable modeled compute and memory cost while exposing interpretable internal dynamics.

1 INTRODUCTION

Spiking Neural Networks (SNNs) promise highly energy-efficient inference by leveraging sparse, event-driven computations Roy et al. (2019). Yet they remain fundamentally constrained on continuous, context-dependent tasks such as automatic speech recognition (ASR) Bittar & Garner (2022); Song et al. (2025); Wu et al. (2020). The primary bottleneck lies in the standard Leaky Integrate-and-Fire (LIF) neuron Gerstner & Kistler (2002); Maass (1997), whose intrinsic membrane decay erases critical context across pauses and phonetic boundaries Bellec et al. (2018); Zenke & Ganguli (2018). Consequently, while continuous ANNs achieve state-of-the-art ASR performance Gulati et al. (2020); Vaswani et al. (2017), their reliance on dense multiply-accumulate (MAC) operations incurs prohibitive memory and energy costs Horowitz (2014)-leaving a critical gap for energy-constrained applications.

Prior spiking ASR systems largely resorted to hybrid architectures pairing a spiking encoder with a dense, floating-point ANN sequence decoder, forfeiting the energy benefits of a truly edge-ready system Bittar & Garner (2022); Song et al. (2025); Wu et al. (2020). Deep event-driven architectures are further hindered by an $\mathcal{O}(L^2)$ training bottleneck arising from sequential recurrent dependencies

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.

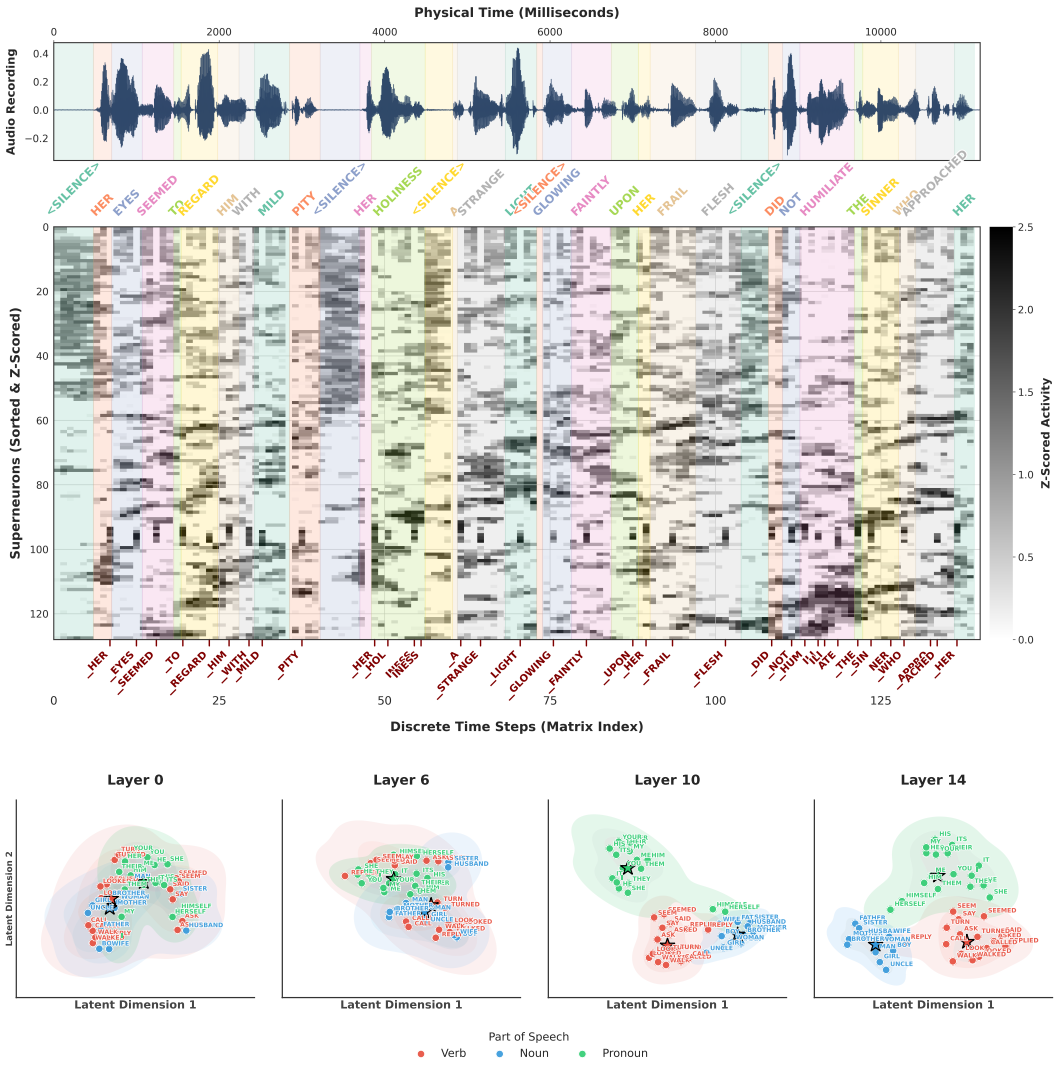

Figure 1: **Acoustic alignment, internal population activity, and semantic manifold emergence in stGRU. (Top)** Input acoustic waveform (LibriSpeech test utterance) aligned with ground-truth word boundaries (color-shaded spans) and internal spiking representations across physical time (top axis) and discrete network timesteps (bottom axis). The central heatmap displays the Z-scored activity of 128 superneurons (co-active neural clusters sorted via Rastermap Stringer et al. (2025)), with darker pixels indicating higher population firing. Red labels beneath the raster mark subword BPE token emissions predicted by the CTC readout layer (leading underscores denote new word onsets; consecutive identical tokens are collapsed into a single emission; CTC blanks omitted for clarity). **(Bottom)** 2D UMAP projections of latent representations across network depth (layers 0, 6, 10, and 14). Internal states progressively disentangle into distinct semantic and syntactic clusters, color-coded by part of speech: verbs (red), nouns (blue), and pronouns (green).

Neftci et al. (2019); Werbos (1990) and discontinuous spike resets Lee et al. (2016); Neftci et al. (2019). Together, these limitations motivate separating rich temporal filtering from discrete event communication while keeping learned projections matrix-MAC-free.

A complementary limitation appears in MatMul-free recurrent models: ternary weights remove dense learned matrix products, but the MLGRU backbone retains continuous activations, so deployment on low-precision neuromorphic hardware requires a separate activation-quantization stage Zhu et al. (2024). This motivates our central question: can the activation constraint itself be learned while retaining gated long-range temporal filtering?

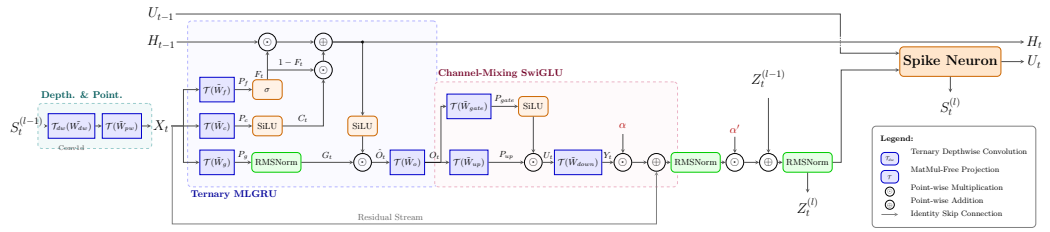

(a) Proposed Architecture

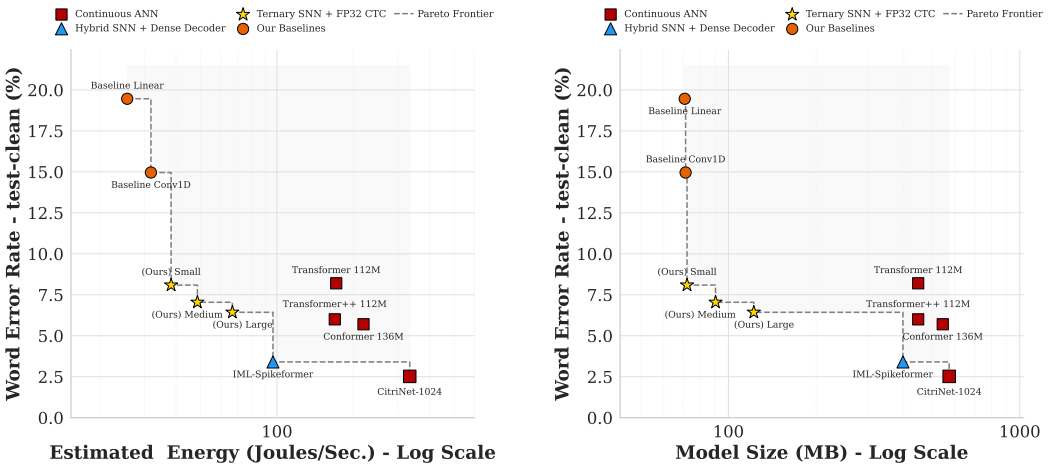

(b) Energy Efficiency Frontier

(c) Memory Efficiency Frontier

Figure 2: **Architecture and efficiency comparisons. (a)** stGRU block architecture. **(b)** LibriSpeech test-clean WER versus estimated arithmetic energy over the 19,440 s evaluation set. **(c)** WER versus theoretical parameter storage in decimal MB. The Large stGRU WER is a six-seed mean; energy estimates correspond to the reported profiling run.

We answer this by treating spiking as a trainable activation interface for the event-driven pathway. In stGRU, a ternary MLGRU performs nonlinear temporal filtering, while a downstream RF neuron emits the binary spike stream that drives the next block's event pathway. In the ternary backbone, weights discretize what the network stores; RF spikes discretize what its event pathway communicates. Crucially, the spike constraint is present during optimization rather than imposed only at deployment. This separation also mirrors cortical neurons, where gated nonlinear processing precedes somatic integration Gidon et al. (2020); Häusser & Mel (2003); London & Häusser (2005); Mel (1991; 1992); Shai et al. (2014; 2015). This storage–communication duality is central to stGRU. Our primary contributions are as follows:

Train-Time Dual Discretization & matrix-MAC-free ASR.

stGRU pairs ternary learned weights with an RF-generated binary event stream, moving activation discretization into training rather than leaving it solely to deployment. To our knowledge, it is the first matrix-MAC-free SNN without a learned ANN sequence decoder to achieve single-digit WER on LibriSpeech Panayotov et al. (2015): Small, Medium, and Large stGRUs reach 8.09%, 7.04%, and $6.44 \pm 0.17\%$ on test-clean without an external LM. Training from scratch on Aishell-1 Bu et al. (2017) yields 10.88% test CER (9.28% dev), close to the VGG-BiLSTM baseline (10.7% test CER) but below the strongest dense and hybrid systems in Table S5, providing cross-lingual evidence that the architecture transfers from English to Mandarin ASR.

Exceptional Performance-to-Model-Size Trade-off. Operating at 1.95 mJ/s (Small) and 3.42 mJ/s (Large), stGRU operates far below continuous Transformers (8.80 mJ/s) and Conformers (11.28 mJ/s) Gulati et al. (2020); Horowitz (2014). For the Large Conv1D–MLGRU–ReLU baseline, ternarizing the weights changes test-clean WER from 5.83% to 6.49%, an increase of 0.66 percentage points, while the modeled arithmetic-energy rate changes from 33.48 to 4.10 mJ/s (seed 3407). Ternary INT2 weights reduce the Large models to 122.22 MB on LibriSpeech and 167.84 MB on Aishell-1, yielding a favorable modeled energy–storage trade-off.

Emergent ACC-like Error-Monitoring. Layer L11 autonomously develops an error-monitoring signature resembling an ACC-related error-monitoring signature Alexander & Brown (2018); Bryden et al. (2019): incorrect words trigger elevated firing, peak sparsity (0.83), and heavy-tailed lognormal rate statistics ($\Delta R^2 = +0.091$ over Gaussian).

Decision Confidence & Autonomous ACC→PFC Handoff. At L14, the coding regime inverts into a prefrontal cortex (PFC)-like commitment state Kiani & Shadlen (2009); Rutishauser et al. (2015): high-confidence correct words elicit peak firing, sparsity drops, and distributions shift to Gaussian. Together, these signatures suggest an autonomous ACC→PFC-like functional handoff Alexander & Brown (2018) emerging within a single artificial sequence model.

2 Related Work

While dense continuous models (Transformers, Conformers, CitriNet) define state-of-the-art ASR Hou et al. (2024); Majumdar et al. (2021); Radford et al. (2022), their reliance on dense multiply-accumulate (MAC) operations incurs prohibitive compute and memory costs Horowitz (2014). MatMul-free MLGRU combines ternary weights with continuous activations; its Loihi 2 mapping therefore separately quantizes activations and retains 16-bit precision after 8-bit quantization substantially degrades accuracy Zhu et al. (2024).

Motivated by biological efficiency, Spiking Neural Networks (SNNs) have been adapted for speech via surrogate gradients Bittar & Garner (2022); Wu et al. (2020), TCNs Bai et al. (2018), and Spike-driven Transformers Yao et al. (2023; 2024). However, existing ternary spiking speech models remain heavily restricted: QT-SNN Wang et al. (2025) targets only isolated keyword spotting, whereas DyTr-SNN Wang et al. (2023) is evaluated solely on single-speaker LJSpeech without reporting WER and relies on a dense continuous 6-layer Transformer decoder. More broadly, leading SNN ASR baselines (GSU Hao et al. (2024), Spike-driven Transformer Yao et al. (2024), IML-Spikeformer Song et al. (2025)) depend on continuous 6-layer ANN decoders to maintain competitive accuracy Song et al. (2025). While their encoders use sparse accumulate (AC) operations, continuous decoders reintroduce a floating-point bottleneck and serialized execution, compromising edge viability.

stGRU bridges these two lines of work: it retains ternary MLGRU filtering but trains an RF-generated binary event interface end-to-end, making event discretization part of the model rather than a post-training hardware conversion. Deferred scaling keeps learned backbone projections matrix-MAC-free, while direct CTC decoding avoids a learned ANN sequence decoder.

3 Methodology

3.1 Ternary Projections with Deferred Weight Scaling

We use ternary BitLinear projections Ma et al. (2024); Zhu et al. (2024). For a full-precision shadow-weight matrix $W \in \mathbb{R}^{m\times d}$, let $\gamma_W = (md)^{-1}\sum_{i,j}|W_{ij}|$ and quantize by

$$\widetilde{W} = \mathrm{RoundClip}\left(\frac{W}{\gamma_W + \epsilon}, -1, 1\right), \qquad \widetilde{W}_{ij} \in \{-1, 0, +1\}. \tag{1}$$

The shadow weights are trained with a straight-through estimator (STE) Bengio et al. (2013) and QAT warmup (Appendix C.1).

Let V denote the input presented to a projection after its shared input normalization. The three ablation modes differ only in the placement or omission of the weight scale:

$$\begin{aligned} Y_{\mathrm{naive}} &= V(\gamma_W \widetilde{W})^\top, \\ Y_{\mathrm{proposed}} &= (V\widetilde{W}^\top)\gamma_W, \\ Y_{\mathrm{none}} &= V\widetilde{W}^\top. \end{aligned} \tag{2}$$

The naive and proposed modes are algebraically equivalent for fixed V, $\widetilde{W}$, and γ_W. Naive scaling folds γ_W into the weights, producing a floating-point weight matrix. Deferred scaling instead accumulates through ternary weights using additions and subtractions, then applies γ_W to the output.

The None mode omits this scale entirely. Thus the proposed mode avoids dense floating-point matrix MACs in the learned projection, but retains an output scaling operation. We write this operator as

$$\mathcal{T}(V, \widetilde{W}, \gamma_W)_i = \gamma_W \sum_{j=1}^{d} \widetilde{W}_{ij} V_j, \qquad i = 1, \ldots, m. \tag{3}$$

The same scale placement applies to a temporal depthwise projection. Let $V \in \mathbb{R}^{T \times C}$ be the input to that projection and $\widetilde{W}_{dw} \in \{-1, 0, +1\}^{K \times C}$ its ternary kernel. With a weight scale γ_{dw}, we define

$$\mathcal{T}_{dw}(V, \widetilde{W}_{dw}, \gamma_{dw})_{t,c} = \gamma_{dw} \sum_{k} \widetilde{W}_{dw,k,c} V_{t+k,c}. \tag{4}$$

The temporal sum therefore uses additions and subtractions through ternary weights; γ_{dw} is applied after that sum. This avoids folding a floating-point scale into every kernel weight. In the ablation (Table S8), deferred scaling reduces the modeled energy rate by about $5\times$ relative to naive scaling, with comparable test-clean WER (Small: 8.09% vs. 8.15%; Large: 6.14% vs. 6.44%). Omitting the scale instead worsens test-clean WER to 9.44% and 7.51%, respectively.

3.2 Parallel Spiking-Ternary Block Architecture

To bypass the prohibitive $\mathcal{O}(L^2)$ training bottleneck of recurrent networks and discontinuous resets Werbos (1990), the Spiking-Ternary GRU Block decouples recurrent filtering from spike emission, enabling global parallel unrolling in $\mathcal{O}(L \log L)$ time while avoiding dense learned matrix MACs at inference.

Acoustic features are initially mapped by a pre-encoder into discrete spike trains $S_t^{(0)} \in \{0, 1\}^C$. At block l, temporal features are extracted from the incoming discrete spike sequence $S^{(l)} \in \{0, 1\}^{T \times C}$ via a depthwise-separable operator pair consisting of a 1D temporal depthwise convolution and a pointwise channel projection:

$$X_{dw} = \mathcal{T}_{dw}(S^{(l)}, \widetilde{W}_{dw}, \gamma_{dw}), \quad X = \mathcal{T}(X_{dw}, \widetilde{W}_{pw}, \gamma_{pw}) \tag{5}$$

where $\widetilde{W}_{pw} \in \{-1, 0, +1\}^{C \times C}$ denotes the ternary pointwise weights and γ_{pw} their scale.

We adapt the MLGRU of Zhu et al. (2024) with deferred-scaling ternary projections, RMS-normalized gating, and a ternary output projection. For brevity, we denote gate-specific ternary projections as $\mathcal{T}_k(V) \triangleq \mathcal{T}(V, \widetilde{W}_k, \gamma_k)$ for $k \in \{c, f, g\}$. For input representation $X_t \in \mathbb{R}^C$, the gates are computed as $C_t = \text{SiLU}(\mathcal{T}_c(X_t))$, $F_t = \sigma(\mathcal{T}_f(X_t))$, and $G_t = \text{RMSNorm}(\mathcal{T}_g(X_t))$, updating state and output via:

$$H_t = F_t \odot H_{t-1} + C_t \odot (1 - F_t) \tag{6}$$

$$O_t = \mathcal{T}(G_t \odot \text{SiLU}(H_t), \widetilde{W}_o, \gamma_o) \tag{7}$$

Because Equation 6 is strictly element-wise (no dense Wh_{t-1} products), intermediate representations are computed globally across sequence length T via an associative parallel scan Gu et al. (2021); Smith et al. (2022).

Channels are mixed via a ternary SwiGLU block stabilized by learnable ReZero vectors $\alpha \in \mathbb{R}^C$ Bachlechner et al. (2021) (Appendix C.2):

$$U_t = \text{SiLU}(\mathcal{T}(O_t, \widetilde{W}_{\text{gate}}, \gamma_{\text{gate}})) \odot \mathcal{T}(O_t, \widetilde{W}_{\text{up}}, \gamma_{\text{up}}) \tag{8}$$

$$Z_t = X_t + \alpha \odot \mathcal{T}(U_t, \widetilde{W}_{\text{down}}, \gamma_{\text{down}}) \tag{9}$$

To prevent vanishing gradients and stabilize variance across depth, the continuous residual stream $Z^{(l)} \in \mathbb{R}^C$ updates via:

$$Z_t^{(l)} = \text{RMSNorm}(Z_t^{(l-1)}) + \alpha' \odot \text{RMSNorm}(Z_t) \tag{10}$$

The stabilized residual $Z_t^{(l)}$ serves as the input current to the Parallel Resonate-and-Fire (PRF) neuron Huang et al. (2024). Unlike standard LIF neurons which enforce sequential BPTT due to hard reset

Table 1: Comprehensive Architectural Comparison on LibriSpeech-960 ASR. All reported WER results evaluate the standalone acoustic model without an external language model. The Large RF WER is mean ± standard deviation over six seeds; other in-house rows and arithmetic-energy rates correspond to seed 3407.

Architecture	Params (M)	Weight Type	Memory (MB)	Energy rate (mJ/s) (mJ/s)	WER (%) test-clean	WER (%) test-other
Continuous Artificial Neural Networks (ANNs)						
Moritz et al. Moritz et al. (2019)	70.0	FP32	280.0†	-	5.4	16.1
Park et al. Park et al. (2019)	360.0	FP32	1440.0†	-	2.8	6.8
Whisper-Small Radford et al. (2022)	244.0	FP16	488.0†	-	3.1	7.4
ResNet-BiGRU Reza et al. (2022)	33.0	FP32	132.8†	-	15.6	-
Quantized RNN-T Zhen et al. (2022)	27.8	FP8	27.8†	-	8.70	22.36
Quantized RNN-T Zhen et al. (2022)	27.8	FP4	13.9†	-	8.78	22.54
Transformer * Hou et al. (2024)	112.0	FP32	448.0†	8.80††	8.2	15.2
Transformer++ * Hou et al. (2024)	112.0	FP32	448.0†	8.69††	6.0	11.8
Conformer * Hou et al. (2024)	136.0	FP32	544.0†	11.28††	5.7	10.8
CitriNet-1024 Majumdar et al. (2021)	143.1	FP32	572.4†	17.18††	2.52	6.22
Hybrid Spiking Neural Networks (SNNs) & Dense ANN Decoder/LM						
Deep FC SNN (LS-360h) Wu et al. (2020)	-	FP32	-	-	9.4	-
1 Spiking RNN + 3 LSTM Bittar & Garner (2022)	-	FP32	-	-	3.51	-
4 Spiking RNN + 0 LSTM Bittar & Garner (2022)	-	FP32	-	-	9.94	-
GSU (+ 6-layer ANN Dec) ** Hao et al. (2024)	185.0	FP32	740.4†	-	12.8	32.1
Spike-driven Trans. (+ 6-layer ANN Dec) ** Yao et al. (2024)	99.0	FP32	397.2†	-	10.5	25.7
IML-Spikeformer (+ 6-layer ANN Dec) Song et al. (2025)	99.0	FP32	397.2†	4.96††	**3.4**	**7.9**
Ours (Structural Baselines, FP32 & W/O Learned ANN Sequence Decoder)						
Large (Linear → RF)	77.01	FP32‡	117.28	1.30	17.34	33.64
Large (Conv1D → RF)	77.58	FP32‡	119.56	2.22	11.64	26.20
Ours (Parameter-Matched Baselines, Continuous Dense ReLU Activations)						
Small (Conv1D → MLGRU → ReLU)	132.66	FP32	530.63	22.12	7.25	19.14
Medium (Conv1D → MLGRU → ReLU)	204.16	FP32	816.63	26.25	6.41	17.29
Large (Conv1D → MLGRU → ReLU)	329.39	FP32	1317.55	33.48	5.83	16.11
Small (Conv1D → MLGRU → ReLU)	132.66	Ternary INT2	72.01	2.60	7.74	20.26
Medium (Conv1D → MLGRU → ReLU)	204.16	Ternary INT2	90.16	3.13	6.68	17.71
Large (Conv1D → MLGRU → ReLU)	329.39	Ternary INT2	121.96	4.10	6.49	17.40
Ours (matrix-MAC-free, Ternary Spiking Networks & W/O Learned ANN Sequence Decoder)						
Ours Large (Streaming) (Conv1D → MLGRU → RF)	329.45	Ternary INT2	122.22	3.51	9.11	22.13
Ours Small (Conv1D → MLGRU → RF)	132.68	Ternary INT2	72.09	1.95	8.09	20.79
Ours Medium (Conv1D → MLGRU → RF)	204.20	Ternary INT2	90.31	2.49	7.04	18.74
Ours Large (Conv1D → MLGRU → RF)	329.45	Ternary INT2	122.22	3.42	6.44 ± 0.17	17.49 ± 0.23

† Theoretical memory size computed based on parameter count and format precision (FP32 = 4 Bytes, FP16 = 2 Bytes, INT2 = 0.25 Bytes).

†† Our Reproduced results by publicly available data.

* Trained on 2.5k hours of various datasets as cited in Hou et al. (2024).

** Reproduced by Song et al. (2025).

‡ Structural baselines use an INT2 ternary downsampler and FP32 downstream weights; unmarked FP32 variants also use an FP32 downsampler.

subtractions ($V_{th}S_{t-1}$), PRF folds resets into a complex decay factor $\tilde{\gamma} = -\frac{1}{\tau} + i\theta$. Under Euler discretization over timestep Δ, membrane integration unfolds linearly:

$$U_t = \exp\left(-\frac{\Delta}{\tau}\right)(\cos(\Delta\theta) + i\sin(\Delta\theta))\,U_{t-1} + \Delta Z_t^{(l)} \tag{11}$$

Because temporal integration is independent of output spikes, membrane potentials $U_{1:T}$ evaluate globally in the frequency domain via Fast Fourier Transform ($\mathcal{F}$) Cooley & Tukey (1965); Gu et al. (2021) using decay kernel $K_{1:T} = (\Delta A^0, \ldots, \Delta A^{T-1})$ with $A = \exp(\Delta\tilde{\gamma})$:

$$U_{1:T} = \mathcal{F}^{-1}(\mathcal{F}(Z_{1:T}^{(l)}) \odot \mathcal{F}(K_{1:T})) \tag{12}$$

Finally, the continuous potentials are folded back into the binary domain for the $(l+1)$-th block by applying a Heaviside step function strictly to the real component:

$$S_t^{(l+1)} = \mathcal{H}(\Re\{U_t\} - V_{th}) \tag{13}$$

Surrogate gradients (Multi-Gaussian Yin et al. (2021)) train this binary event interface end-to-end, so the next block is optimized on the same discrete spike stream used at inference. The MLGRU recurrence admits an associative scan and PRF integration an $\mathcal{O}(L \log L)$ FFT implementation; empirically, training is 3.1×–3.4× faster than the parameter-matched LIF variants (Tables S7, S9).

4 Experiments

We evaluate the Spiking Ternary GRU (stGRU) on LibriSpeech-960 and Aishell-1 Panayotov et al. (2015); Bu et al. (2017) implemented in SpeechBrain Ravanelli et al. (2022) across 8 AMD Instinct

MI250X GCDs on the LUMI supercomputer. Acoustic inputs (80-dimensional log-Mel filterbanks at 50 Hz) undergo $4\times$ temporal subsampling via a ternary 2D CNN frontend (512 and 1024 channels), followed by an MLGRU layer coupled with a Parallel Resonate-and-Fire (PRF) neuron acting as a pre-encoder that emits discrete spikes at a reduced rate of 12.5 Hz. The subsequent stGRU backbone applies 1D depthwise convolutions with kernel sizes up to $K = 63$ across 1024 channels (Appendix A), trained end-to-end via CTC loss Graves et al. (2006) with Semi-Resonant Cosine Annealing (Appendix C.3). All primary metrics evaluate the standalone acoustic model; an optional OpenSLR 4-gram language model[1] is decoded (Table S6) via non-neural CPU hash lookups and is therefore excluded from neural energy profiling.

Performance and Efficiency Benchmarks. We benchmark against continuous ANNs (Transformer, Conformer, CitriNet Vaswani et al. (2017); Hou et al. (2024); Gulati et al. (2020); Majumdar et al. (2021)) and the hybrid SNN IML-Spikeformer Song et al. (2025), calculating theoretical energy via 45 nm CMOS constants (4.6 pJ/FP-MAC, 0.9 pJ/sparse-AC, 17.5 pJ/PRF-update Horowitz (2014); Huang et al. (2024)). As reported in Table 1, Small, Medium, and Large stGRUs achieve standalone WERs of 8.09%, 7.04%, and $6.44\pm0.17\%$ on LibriSpeech test-clean without an external LM (optional 4-gram LM yields $5.95 \pm 0.16\%$; Table S6), setting the first single-digit WERs for matrix-MAC-free SNNs without a learned ANN sequence decoder. A bounded-latency streaming variant with 320 ms look-ahead attains 9.11% WER at 3.51 mJ/s (Table 1). A separately trained stGRU achieves 10.88% test CER on Aishell-1 and 9.28% on dev (Table S5).

Parameter-matched ablations isolate individual design choices (Table 1, Table S7): quantizing from FP32 to ternary weights cuts memory by $7.4\times$–$10.8\times$ with negligible loss ($5.83\% \rightarrow 6.49\%$). At the Large scale, discrete PRF spikes surpass continuous ternary ReLUs in both WER and energy (6.14% at 3.42 mJ/s vs. 6.49% at 4.10 mJ/s; widening to 33.48 mJ/s for FP32 ReLUs; Table S7). Under the same analytical 45 nm arithmetic-cost model, our estimate for IML-Spikeformer allocates 2.73 mJ/s to its FP32 frontend and 1.68 mJ/s to its ANN decoder, whereas the stGRU estimate includes 0.05 mJ/s for active-weight accumulation in its FP32 CTC readout (Table S3). These are operation-based estimates, not measured hardware energy. Upgrading from Linear $\rightarrow$ RF to the full ternary stack adds < 5 MB and $+2.12$ mJ/s while reducing test-clean WER by 11.20 percentage points for seed 3407 ($17.34\% \rightarrow 6.14\%$).

5 Neuromorphic Interpretation

Having established favorable modeled efficiency on English and Mandarin ASR, we now examine whether the network's internal representations exhibit functional analogies to cortical decision signals. Echoing the parallel, distributed processing of human auditory cortex Hamilton et al. (2021); Bhaya-Grossman & Chang (2022), deep layers (L10–L14) autonomously partition into functionally distinct subpopulations whose activity statistics mimic those of neurons in human cortex during perceptual decisions. We substantiate this below.

5.1 Error-Monitoring Signature in L11 Mimics the Human Medial Frontal Cortex

At L11, incorrect words result in significantly higher population firing than either high- or low-confidence correct words (Fig. 3, row 4) analogous to error-related activity associated with anterior cingulate cortex (ACC) Alexander & Brown (2018); Bryden et al. (2019). The error code is carried by a subpopulation of neurons: incorrect-preferring neurons outnumber high-confidence-correct neurons ($n = 194$ vs. 64; Fig. 3, top). L11 also exhibits the sparsest spiking activity from all layers (Fig. 3, row 3), placing it in the recruitment-based regime characteristic of cortical error monitoring. The response is also temporally appropriate for an error signal: incorrect-preferring neurons exhibit a sharp post-onset transient that decays before the CTC emission window, on the timescale of error detection rather than sustained stimulus encoding. Per-neuron firing rates at L11 follow a lognormal distribution (Fig. S2) as observed in cortical recall populations. It follows that directionality, recruitment, timing, and distribution align with an ACC-like role of L11.

[1] https://www.openslr.org/11/

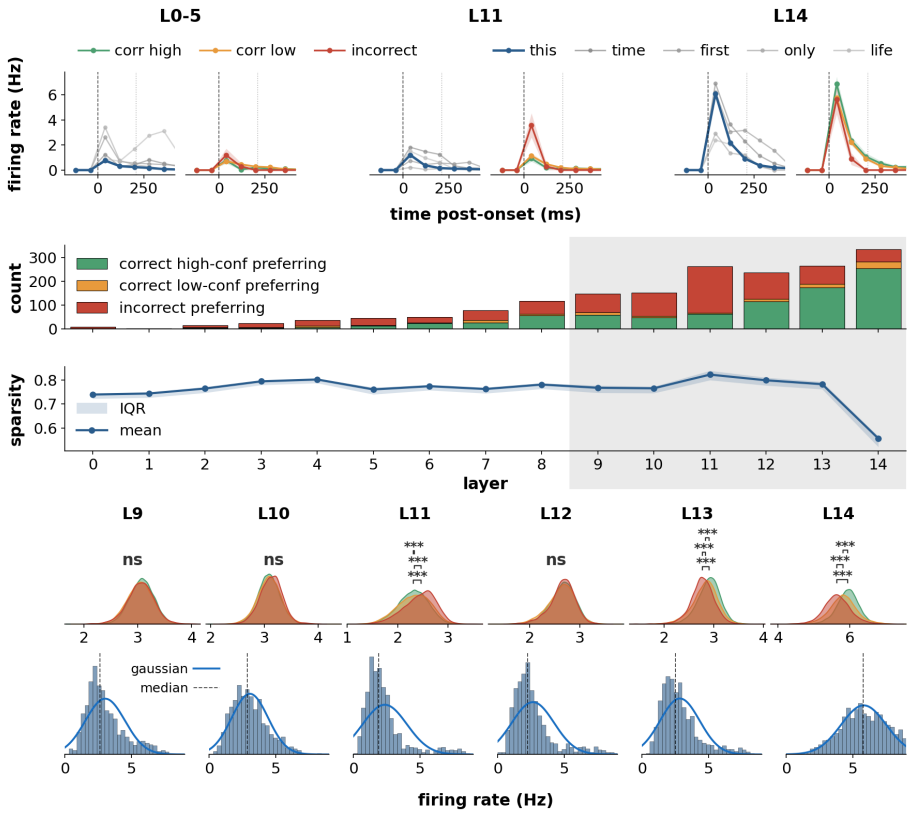

Figure 3: **Depth-graded coding emerges in trained Spiking-Ternary networks.** Row 1: PSTH of word-selective neurons at shallow (L0–5), uncertainty (L11), and commitment (L14) depths. Each panel pair shows responses to a target word ("this," blue) against acoustically similar competitors (gray), with the right sub-panel split by outcome (green: high-conf correct; orange: low-conf correct; red: incorrect). Row 2: Per-layer counts of outcome-selective neurons (matched-AUC, FDR $q < 0.001$); correct-preferring neurons are sub-classified by high- vs. low-confidence via median CTC log-probability. Selectivity grows from L0–L5 to L11–L14 (< 50 vs. > 250 neurons), and the bipartite partition reverses across depth: incorrect-preferring neurons peak at L11 ($n = 194$), high-confidence-correct-preferring at L14 ($n = 254$). Row 3: Population sparsity (per-word silent fraction) across encoder depth, with shaded IQR. Row 4: Population firing rate distributions for L9–L14 split by outcome (brackets: pairwise Mann–Whitney U-test with Bonferroni correction; `ns`: panels with at least one non-significant comparison). Row 5: Per-neuron firing rate histograms for L9–L14 with Gaussian fit overlaid.

5.2 Decision-Confidence Signature in L14 Mimics the Prefrontal Cortex

L14, in turn, exhibits different behavior to L11. In L14 the highest spike rate corresponds to high-confidence correct words while low-confidence correct words are between the high-confidence and incorrect distributions. This graded confidence axis is absent at L11 and mimics the behavior observed in primate lateral intraparietal (LIP) and human medial temporal lobe (MTL) Kiani & Shadlen (2009); Rutishauser et al. (2015). At the level of decision hierarchy, prefrontal cortex (PFC) is the area where confidence signals from upstream regions are converted into committed action. The composition of the selective population inverts in parallel: high-confidence-correct neurons dominate ($n = 254$ vs. 58 incorrect-preferring; Fig. 3, top). PSTH responses broaden in turn, replacing the narrow error-only transient at L11 with large, sustained post-onset peaks across all three outcome conditions—the response shape expected of a layer broadcasting a committed token rather than gating a selective monitoring signal. Moreover, population sparsity decreases significantly from L11 to L14 and the per-neuron firing-rate changes from lognormal to near-Gaussian ($\Delta R^2 = -0.011$, the only negative value in the recall hierarchy; Fig. S2). The latter indicates broad recruitment rather than selective monitoring with L14 emulating human PFC behavior.

5.3 Autonomous ACC-to-PFC Handoff Within a Single Network

The transition between L11 and L14 proceeds gradually through L12–L13 in a functional handoff. L12 acts as the inflection layer where responses to all categories become similar (Fig. 3, bottom, `ns`), the population code's directionality reverses, and the spike rate lognormality gradually decays from L11 ($+0.091$) through L12 ($+0.064$) and L13 ($+0.041$) before vanishing at L14 (-0.011). Selectivity composition, sparsity, firing-rate directionality, and distribution shape move together along this trajectory, consistent with a continuous functional gradient. This structure mimics the ACC with its error and surprise signals Alexander & Brown (2018) while PFC implements the downstream control that converts those signals into committed decisions. Critically, this handoff is implemented *within* a single recurrent network rather than *across* brain regions: a network trained end-to-end on CTC, with no architectural prior favoring functional segregation, partitions its deep layers into error-monitoring and confidence-readout subpopulations whose responses, sparsity, and firing-rate statistics jointly track signatures observed in human cortex *in vivo*. We note that this partition is not incidental to depth or spiking activity alone: an ablated variant with MLGRU gating removed shows the same deep-layer outcome selectivity collapse by 95.5% once empty-emission trials are controlled for, versus 8.1% in the full model (Supp. Fig. S3). This ablation supports a role for MLGRU gating in the observed outcome-selective activity. This mirrors cortical pyramidal neurons, where nonlinear dendritic filtering precedes somatic integration London & Häusser (2005); Mel (1992); Poirazi et al. (2003); Shai et al. (2015) and subthreshold resonance bridges temporal context across pauses Bellec et al. (2018); Giraud & Poeppel (2012): upstream recurrent gating provides the temporal substrate necessary for downstream somatic integrators to develop decision-level outcome selectivity.

Limitations. Our energy estimates use analytical operation counts and 45 nm CMOS constants. QAT also retains full-precision shadow weights and their training overhead: the forward weights become ternary after warmup, while the straight-through estimator updates the continuous weights during backpropagation. Furthermore, while our design eliminates all dense $\mathcal{O}(C^2)$ floating-point matrix multiplications, the model is not entirely multiplication-free: continuous element-wise multiplications persist in recurrent gating, residual scaling, RMSNorm, and neuron dynamics. These operations are element-wise rather than dense $O(C^2)$ matrix multiplications, but deploying them on pure addition-only or integer-only edge silicon still requires specialized arithmetic units or approximations. Performance remains weaker on `test-other` than on `test-clean`. Finally, mitigating the capacity limits of 2-bit weights required wide channels and temporal kernels of up to $K = 63$, despite the low precision of their weights.

6 Conclusion

We introduced stGRU, an ASR architecture combining deferred-scaling ternary projections, gated recurrent filtering, and resonant spiking integration. At inference, its learned projections avoid dense matrix MACs, while the Large model reaches $6.44 \pm 0.17\%$ LibriSpeech test-clean WER without an external language model and a separately trained Aishell-1 model reaches 10.88% test CER. The modeled arithmetic-energy and parameter-storage results show a favorable efficiency trade-off, subject to the accounting limits above.

Conceptually, stGRU couples complementary low-precision constraints in the event-driven backbone: ternary learned projections constrain synaptic weights, while RF neurons make the inter-block event stream binary during training rather than relying on a separate deployment-time activation-quantization step. Spiking thus serves not only as neuronal dynamics but as an end-to-end optimized binary communication bottleneck. Despite this constraint, deep layers still develop error-to-confidence selectivity and structured sparsity, suggesting that efficient event-driven communication can coexist with rich internal organization.

Reproducibility Statement

Complete implementation details, network specifications across all model scales, and exact random seed configurations are documented in Appendices A, B, and Table S6. Training runs were conducted using SpeechBrain on standardized public benchmarks (LibriSpeech-960 and Aishell-1). All analytical energy profiling formulations and CMOS constants are detailed in Appendix D. Anonymized

code, environment specifications, and configuration scripts to reproduce all reported experiments are provided in the supplementary material.

AI USE DISCLOSURE

Generative AI tools were used to assist with the implementation and debugging of training and visualization code, and with drafting, editing, and refining the presentation and interpretation of portions of the manuscript. Visualization code assisted by AI was executed by the authors to produce figures from experimental data using Python; no figures or experimental data were generated directly by generative AI. All AI-assisted code was reviewed, tested, and executed by the authors, and all reported experiments and results were independently verified. The authors reviewed all AI-assisted text and take full responsibility for the methods, results, claims, and final content of this work.

REFERENCES

William H. Alexander and Joshua W. Brown. Frontal cortex function derives from hierarchical predictive coding. *Scientific Reports*, 2018.

Thomas C. Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W. Cottrell, and Julian McAuley. Rezero is all you need: Fast convergence at large depth. *Conference on Uncertainty in Artificial Intelligence*, 2021.

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv: Learning*, 2018.

G. Bellec, Darjan Salaj, Anand Subramoney, R. Legenstein, and W. Maass. Long short-term memory and learning-to-learn in networks of spiking neurons. *Neural Information Processing Systems*, 2018.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv: Learning*, 2013.

Ilina Bhaya-Grossman and Edward F. Chang. Speech computations of the human superior temporal gyrus. *Annual Review of Psychology*, 2022. doi: 10.1146/annurev-psych-022321-035256.

Alexandre Bittar and Philip N. Garner. Surrogate gradient spiking neural networks as encoders for large vocabulary continuous speech recognition. *ArXiv*, 2022. doi: 10.48550/arxiv.2212.01187.

Daniel W. Bryden, Adam T. Brockett, Elizabeth Blume, Kerry Heatley, Andrew Zhao, and Matthew R. Roesch. Single neurons in anterior cingulate cortex signal the need to change action during performance of a stop-change task that induces response competition. *Cerebral Cortex*, 2019.

Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao Zheng. Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline. *Oriental COCOSDA International Conference on Speech Database and Assessments*, 2017. doi: 10.1109/icsda.2017.8384449.

J. Cooley and J. Tukey. An algorithm for the machine calculation of complex fourier series. 1965. doi: 10.1090/s0025-5718-1965-0178586-1.

W. Gerstner and W. M. Kistler. Spiking neuron models: Single neurons, populations, plasticity. 2002. doi: 10.1017/cbo9780511815706.

Albert Gidon, Timothy Adam Zolnik, Pawel Fidzinski, Felix Bolduan, Athanasia Papoutsi, Panayiota Poirazi, Martin Holtkamp, Imre Vida, and Matthew Evan Larkum. Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science*, 367(6473):83–87, 2020. doi: 10.1126/science.aax6239.

Anne-Lise Giraud and David Poeppel. Cortical oscillations and speech processing: emerging computational principles and operations. *Nature Neuroscience*, 2012. doi: 10.1038/nn.3063.

A. Graves, Santiago Fernández, Faustino J. Gomez, and J. Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. *ICML*, 2006. doi: 10.1145/1143844.1143891.

Albert Gu, Karan Goel, and Christopher R'e. Efficiently modeling long sequences with structured state spaces. *International Conference on Learning Representations*, 2021.

Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. Conformer: Convolution-augmented transformer for speech recognition. *Interspeech*, 2020. doi: 10.21437/interspeech.2020-3015.

Liberty S. Hamilton, Yulia Oganian, Jeffery A. Hall, and Edward F. Chang. Parallel and distributed encoding of speech across human auditory cortex. *Cell*, 2021. doi: 10.1016/j.cell.2021.07.019.

Xiang Hao, Chenxiang Ma, Qu Yang, Jibin Wu, and Kay Chen Tan. Towards ultra-low-power neuromorphic speech enhancement with spiking-fullsubnet. *arXiv (Cornell University)*, 2024. doi: 10.48550/arxiv.2410.04785.

M. Horowitz. 1.1 computing's energy problem (and what we can do about it). *IEEE International Solid-State Circuits Conference*, 2014. doi: 10.1109/isscc.2014.6757323.

Zejiang Hou, Goeric Huybrechts, Anshu Bhatia, Daniel Garcia-Romero, Kyu J. Han, and Katrin Kirchhoff. Revisiting convolution-free transformer for speech recognition. *Interspeech*, 2024. doi: 10.21437/interspeech.2024-588.

Yulong Huang, Zunchang Liu, Changchun Feng, Xiaopeng Lin, Hongwei Ren, Haotian Fu, Yue Zhou, Hong Xing, and Bojun Cheng. Prf: Parallel resonate and fire neuron for long sequence learning in spiking neural networks. *arXiv.org*, 2024. doi: 10.48550/arxiv.2410.03530.

Michael Häusser and Bartlett W. Mel. Dendrites: bug or feature? *Current Opinion in Neurobiology*, 2003. doi: 10.1016/s0959-4388(03)00075-8.

Roozbeh Kiani and Michael N. Shadlen. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 2009.

Jun Haeng Lee, Tobi Delbruck, and Michael Pfeiffer. Training deep spiking neural networks using backpropagation. *Frontiers in Neuroscience*, 2016. doi: 10.3389/fnins.2016.00508.

Michael London and Michael Häusser. Dendritic computation. *Annual Review of Neuroscience*, 2005. doi: 10.1146/annurev.neuro.28.061604.135703.

Shuming Ma, Hongyu Wang, Lingxiao Ma, Lei Wang, Wenhui Wang, Shaohan Huang, Lifeng Dong, Ruiping Wang, Jilong Xue, and Furu Wei. The era of 1-bit llms: All large language models are in 1.58 bits. *arXiv.org*, 2024. doi: 10.48550/arxiv.2402.17764.

Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Transactions of The Society for Computer Simulation International*, 1997. doi: 10.5555/281543.281637.

Somshubra Majumdar, Jagadeesh Balam, Oleksii Hrinchuk, Vitaly Lavrukhin, Vahid Noroozi, and Boris Ginsburg. Citrinet: Closing the gap between non-autoregressive and autoregressive end-to-end models for automatic speech recognition. *arXiv: Audio and Speech Processing*, 2021.

Bartlett W. Mel. The clusteron: Toward a simple abstraction for a complex neuron. In J. Moody, S. Hanson, and R. P. Lippmann (eds.), *Advances in Neural Information Processing Systems*, volume 4, pp. 35–42. Morgan Kaufmann, 1991.

Bartlett W. Mel. NMDA-based pattern discrimination in a modeled cortical neuron. *Neural Computation*, 4(4):502–517, 1992. doi: 10.1162/neco.1992.4.4.502.

Niko Moritz, Takaaki Hori, and Jonathan Le Roux. Triggered attention for end-to-end speech recognition. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2019. doi: 10.1109/icassp.2019.8683510.

Emre Ozgur Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *arXiv: Neural and Evolutionary Computing*, 2019. doi: 10.1109/msp.2019.2931595.

Vassil Panayotov, Guoguo Chen, Daniel Povey, and S. Khudanpur. Librispeech: An asr corpus based on public domain audio books. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2015. doi: 10.1109/icassp.2015.7178964.

Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin Dogus Cubuk, and Quoc V. Le. Specaugment: A simple data augmentation method for automatic speech recognition. *conference of the international speech communication association*, 2019. doi: 10.21437/interspeech.2019-2680.

Panayiota Poirazi, Terrence Brannon, and Bartlett W. Mel. Pyramidal neuron as two-layer neural network. *Neuron*, 37(6):989–999, 2003. ISSN 0896-6273. doi: https://doi.org/10.1016/S0896-6273(03)00149-1. URL https://www.sciencedirect.com/science/article/pii/S0896627303001491.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, C. McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. *International Conference on Machine Learning*, 2022. doi: 10.48550/arxiv.2212.04356.

Mirco Ravanelli, Titouan Parcollet, Peter Plantinga, Aku Rouhe, Samuele Cornell, Loren Lugosch, Cem Subakan, Nauman Dawalatabad, Abdelwahab Heba, Jianyuan Zhong, Ju-chieh Chou, Sung-Lin Yeh, Szu-Wei Fu, Chien-Feng Liao, Elena Rastorgueva, Francois Grondin, William Aris, Hwidong Na, Yan Gao, Renato De Mori, and Yoshua Bengio. Speechbrain: A general-purpose speech toolkit. *arXiv: Audio and Speech Processing*, 2022.

Selim Reza, Marta Campos Ferreira, J.J.M. Machado, João Manuel R. S. Tavares, Selim Reza, Marta Campos Ferreira, J.J.M. Machado, and João Manuel R. S. Tavares. A customized residual neural network and bi-directional gated recurrent unit-based automatic speech recognition model. *Expert systems with applications*, 2022. doi: 10.1016/j.eswa.2022.119293.

Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 2019. doi: 10.1038/s41586-019-1677-2.

Ueli Rutishauser, Shengxuan Ye, Matthieu Koroma, Oana Tudusciuc, Ian B. Ross, Jeffrey M. Chung, and Adam N. Mamelak. Representation of retrieval confidence by single neurons in the human medial temporal lobe. *Nature Neuroscience*, 2015.

Adam S. Shai, Christof Koch, and Costas A. Anastassiou. Spike-timing control by dendritic plateau potentials in the presence of synaptic barrages. *Frontiers in Computational Neuroscience*, 8:89, 2014. doi: 10.3389/fncom.2014.00089.

Adam S. Shai, Costas A. Anastassiou, Matthew E. Larkum, and Christof Koch. Physiology of layer 5 pyramidal neurons in mouse primary visual cortex: Coincidence detection through bursting. *PLOS Computational Biology*, 11(3):e1004090, 2015. doi: 10.1371/journal.pcbi.1004090.

Jimmy Smith, Andrew Warrington, and Scott W. Linderman. Simplified state space layers for sequence modeling. *arXiv.org*, 2022. doi: 10.48550/arxiv.2208.04933.

Zeyang Song, Shimin Zhang, Yuhong Chou, Jibin Wu, and Haizhou Li. Iml-spikeformer: Input-aware multi-level spiking transformer for speech processing. *arXiv.org*, 2025. doi: 10.48550/arxiv.2507.07396.

Matei Ioan Stan and Oliver Rhodes. Learning long sequences in spiking neural networks. *arXiv.org*, 2023. doi: 10.48550/arxiv.2401.00955.

Carsen Stringer, Lin Zhong, Atika Syeda, Fengtong Du, Maria Kesa, and Marius Pachitariu. Rastermap: a discovery method for neural population recordings. *Nature Neuroscience*, 28 (1):201–212, 2025. doi: 10.1038/s41593-024-01783-4.

Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Neural Information Processing Systems*, 2017.

Qingyu Wang, Tielin Zhang, Minglun Han, Yi Wang, Duzhen Zhang, and Bo Xu. Complex dynamic neurons improved spiking transformer network for efficient automatic speech recognition. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'23/IAAI'23/EAAI'23. AAAI Press, 2023. ISBN 978-1-57735-880-0. doi: 10.1609/aaai.v37i1.25081. URL https://doi.org/10.1609/aaai.v37i1.25081.

Shuai Wang, Dehao Zhang, Ammar Belatreche, Yichen Xiao, Hongyu Qing, Wenjie Wei, Malu Zhang, and Yang Yang. Ternary spike-based neuromorphic signal processing system. *Neural Networks*, 187:107333, 2025. ISSN 0893-6080. doi: https://doi.org/10.1016/j.neunet.2025.107333. URL https://www.sciencedirect.com/science/article/pii/S0893608025002126.

P. Werbos. Backpropagation through time: What it does and how to do it. *Proceedings of the IEEE*, 1990. doi: 10.1109/5.58337.

Jibin Wu, Emre Yilmaz, Malu Zhang, Haizhou Li, and Kay Chen Tan. Deep spiking neural networks for large vocabulary automatic speech recognition. *Frontiers in Neuroscience*, 2020. doi: 10.3389/fnins.2020.00199.

Man Yao, Jiakui Hu, Zhaokun Zhou, Liuliang Yuan, Yonghong Tian, Boxing Xu, and Guoqi Li. Spike-driven transformer. *arXiv.org*, 2023. doi: 10.48550/arxiv.2307.01694.

Man Yao, Jiakui Hu, Tianxiang Hu, Yifan Xu, Zhaokun Zhou, Yonghong Tian, Boxing Xu, and Guoqi Li. Spike-driven transformer v2: Meta spiking neural network architecture inspiring the design of next-generation neuromorphic chips. *arXiv.org*, 2024. doi: 10.48550/arxiv.2404.03663.

Bojian Yin, Federico Corradi, and Sander M. Bohte. Accurate and efficient time-domain classification with adaptive spiking recurrent neural networks. *Nature Machine Intelligence*, 2021. doi: 10.1038/s42256-021-00397-w.

Friedemann Zenke and Surya Ganguli. Superspike: Supervised learning in multilayer spiking neural networks. *Neural Computation*, 2018. doi: 10.1162/neco_a_01086.

Kai Zhen, Martin H. Radfar, H. Nguyen, Grant P. Strimel, Nathan Susanj, and A. Mouchtaris. Sub-8-bit quantization for on-device speech recognition: A regularization-free approach. *Spoken Language Technology Workshop*, 2022. doi: 10.1109/slt54892.2023.10022821.

Rui-Jie Zhu, Yu Zhang, Ethan Sifferman, Tyler Sheaves, Yiqiao Wang, Dustin Richmond, Peng Zhou, and J. Eshraghian. Scalable matmul-free language modeling. *arXiv.org*, 2024. doi: 10.48550/arxiv.2406.02528.

A Architectural Configurations and Streaming Setup

Table S1 details the layer-by-layer configurations for the Small, Medium, and Large variants of our proposed stGRU architecture. The network hierarchically scales its temporal receptive field, mapping sequences from basic phonetic structures up to extended sentence and paragraph-level contexts. Across all scales, the temporal convolutions utilize a fixed channel dimension of 1024, a stride of 1, and a dilation of 1. Capacity is scaled by uniformly increasing the depth within each linguistic stage, culminating in a 15-layer Large variant that introduces an exclusive fifth stage with a receptive field of $K = 63$.

A.1 Streaming Implementation via Asymmetric Padding

To enable the 320 ms look-ahead streaming variant evaluated in Section 4, we bound future context. Following the $4\times$ spatio-temporal subsampling of the 50 Hz acoustic features, our pre-encoder operates at an internal frame rate of 12.5 Hz, where each discrete temporal step represents exactly 80 ms of audio. Standard temporal convolutions utilize symmetric padding, inherently violating streaming constraints by continually pulling from future acoustic frames. To resolve this, we adopt an

Table S1: **Detailed architectural configurations for the Small, Medium, and Large stGRU variants.** The receptive field expands progressively across stages to capture hierarchical linguistic features.

Linguistic Stage	Channels	Kernel Size	Number of Layers		
			Small	Medium	Large
Stage 1 (Phonemes)	1024	11	1	2	3
Stage 2 (Syllables)	1024	21	1	2	3
Stage 3 (Words)	1024	39	1	2	3
Stage 4 (Sentences)	1024	51	1	2	3
Stage 5 (Extended Context)	1024	63	–	–	3
Total stGRU Layers			**4**	**8**	**15**

asymmetric padding strategy. By eliminating right-sided padding across the majority of the network and deliberately restricting it to accumulate exactly 4 future frames at specific intermediate layers, we bound the network's total future receptive field to exactly 320 ms (4×80 ms). This bounds algorithmic look-ahead compliance without altering the core matrix-MAC-free MLGRU dynamics.

B Reproducibility, Training Infrastructure, and Hyperparameters

All model training and empirical evaluations were executed on the LUMI supercomputer cluster using a single LUMI-G node equipped with 4 AMD Instinct MI250X modules, which expose 8 Graphics Compute Dies (GCDs) addressed as independent devices by the runtime. The execution pipeline was containerized using Singularity. Distributed Data Parallel (DDP) training was conducted across the 8 GCDs using `torchrun` with the `c10d` backend.

Table S2: **Key hyperparameters for PRF neuronal dynamics, optimization, tokenization, and CTC decoding**.

Component	Hyperparameter	Value / Initialization
PRF Dynamics	Membrane Time Constant (τ)	Target 1.0 (with $\mathcal{N}(0, 0.05)$ jitter on τ^{-1})
	Resonant Frequency (θ)	Linear spread $\in [0, \pi/2]$ + $\mathcal{N}(0, 0.02)$ jitter
	Integration Step (dt)	Target 1.0 + $\mathcal{N}(0, 0.05)$ jitter
	Spike Threshold (V_{thr})	Target 1.0 + $\mathcal{N}(0, 0.05)$ jitter
Optimization	SRCA Peak Learning Rate ($\eta_{\max}$)	10^{-3}
	SRCA Valley Learning Rate ($\eta_{\min}$)	10^{-5}
	QAT Warmup Steps	10,990 steps (10% of total)
Tokenization	LibriSpeech-960	SentencePiece BPE (10,000 Vocabulary)
	Aishell-1	BERT-base-Chinese tokenizer (21,128-token vocabulary)
CTC Decoding	Beam Size	512
	Language Model Weight (α)	0.8
	Word Insertion Penalty (β)	1.2

Note: All four PRF parameters ($\tau, \theta, dt, V_{thr}$) are defined as trainable variables jointly optimized during backpropagation.

C Algorithmic Derivations and Pseudocode

C.1 Quantization-Aware Training (QAT) Warmup Strategy

To mitigate the optimization instability and "quantization shock" associated with aggressively discretizing a randomly initialized network, we employ a progressive QAT warmup strategy. Let W be the continuous full-precision weight matrix. We compute the scaled weights normalized by the

absolute mean γ_W:

$$W_{\text{scaled}} = \frac{W}{\gamma_W + \epsilon}, \quad W_{\text{ternary}} = \max\left(-1, \min\left(1, \text{round}(W_{\text{scaled}})\right)\right) \tag{S1}$$

During forward passes, we construct an effective weight matrix $W_{\text{effective}}$ using a scheduling parameter $\alpha_{\text{qat}} \in [0, 1]$:

$$W_{\text{effective}} = (1 - \alpha_{\text{qat}})W_{\text{scaled}} + \alpha_{\text{qat}} W_{\text{ternary}}, \quad \alpha_{\text{qat}} = \min\left(1, \frac{\text{step}}{N_{\text{warmup}}}\right) \tag{S2}$$

At $\alpha_{\text{qat}} = 0$, the network operates in a purely continuous domain. As $\alpha_{\text{qat}} \to 1$, the network is progressively constrained to ternary representations. Backpropagation passes gradients directly through $W_{\text{effective}}$ via the Straight-Through Estimator (STE) to update the continuous latent weights W.

C.2 Deep Residual Stabilization via Channel-Wise Scaling

To ensure stable signal propagation across deep layers without adding matrix multiplications, we introduce a layer-wise learnable parameter vector $\alpha \in \mathbb{R}^C$. Let $\mathcal{F}(X_t)$ represent the discrete transformation of the l-th Spiking-Ternary block. The residual update is modulated via element-wise multiplication:

$$Z_t = X_t + \alpha \odot \mathcal{F}(X_t) \tag{S3}$$

The vector α is initialized to zero (or 10^{-4}), preserving identity gradient flow during early epochs. Because α is a 1D vector applied element-wise, it requires only C sparse floating-point multiplications per timestep, completely bypassing dense matrix operations.

C.3 Semi-Resonant Cosine Annealing (SRCA)

Let $\eta_{\max}$ and $\eta_{\min}$ denote the peak and anchored valley learning rates. Defining S_{hold} as the holding period, $S_{\text{decay}} = S_{\text{total}} - S_{\text{hold}}$, and $c = 2.75$ as the number of cosine cycles, the cycle progress for $t > S_{\text{hold}}$ is $C_t = c(t - S_{\text{hold}})/S_{\text{decay}}$. The learning rate η_t is governed by:

$$\eta_t = \begin{cases} \eta_{\max} & \text{if } t \leq S_{\text{hold}} \\ \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min}) \cdot \gamma^{C_t} \cdot \left(1 + \cos(2\pi(C_t + \phi_{\text{start}}))\right) & \text{if } t > S_{\text{hold}} \end{cases} \tag{S4}$$

where $\gamma = 0.75$ is the exponential decay factor and $\phi_{\text{start}} = 0.75$ initializes the scheduler at the valley phase for a warmup sweep (Figure S1).

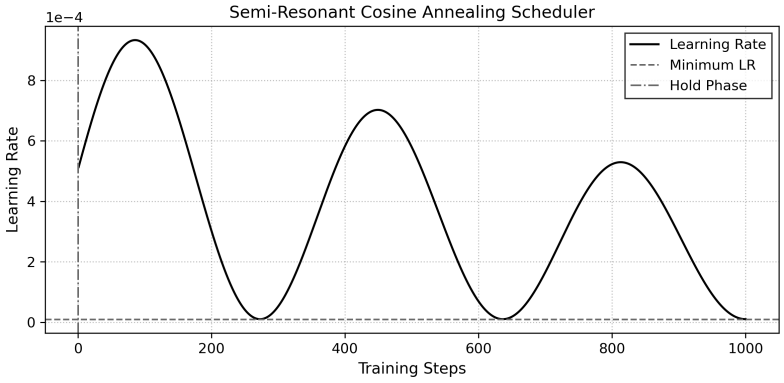

Figure S1: **Semi-Resonant Cosine Annealing (SRCA) learning rate schedule** with initial holding period ($S_{\text{hold}} = 150$), followed by 2.75 cosine cycles with exponential peak decay ($\gamma = 0.75$) and initial phase offset ($\phi_{\text{start}} = 0.75$).

Algorithm S1 End-to-End Spiking-Ternary Block Forward Pass (with Normalized Residuals)

Require: $S^{(l-1)} \in \{0,1\}^{T\times C}$, Continuous residual $Z^{(l-1)} \in \mathbb{R}^{T\times C}$
Require: Ternary weights $\widetilde{W}$ and structural scales γ for all projections
Require: $\alpha, \alpha' \in \mathbb{R}^C$, Initial MLGRU hidden state H_0, PRF decay kernel $K_{1:T}$

// 1. Temporal Feature Extraction & MLGRU Update
for $t = 1$ **to** T **do**
  $X_{dw,t} \leftarrow \mathcal{T}_{dw}(S^{(l-1)}, \widetilde{W}_{dw}, \gamma_{dw})_t$ ▷ Depthwise matrix-MAC-free operator
  $X_t \leftarrow \mathcal{T}(X_{dw,t}, \widetilde{W}_{pw}, \gamma_{pw})$ ▷ Pointwise projection
  $C_t \leftarrow \text{SiLU}(\mathcal{T}(X_t, \widetilde{W}_c, \gamma_c))$
  $F_t \leftarrow \sigma(\mathcal{T}(X_t, \widetilde{W}_f, \gamma_f))$
  $G_t \leftarrow \text{RMSNorm}(\mathcal{T}(X_t, \widetilde{W}_g, \gamma_g))$
  $H_t \leftarrow F_t \odot H_{t-1} + C_t \odot (1 - F_t)$ ▷ Element-wise parallel scan recurrence
  $O_t \leftarrow \mathcal{T}(G_t \odot \text{SiLU}(H_t), \widetilde{W}_o, \gamma_o)$
end for
// 2. SwiGLU Channel Mixing & Residual Highway
for $t = 1$ **to** T **do**
  $U'_t \leftarrow \text{SiLU}(\mathcal{T}(O_t, \widetilde{W}_{\text{gate}}, \gamma_{\text{gate}})) \odot \mathcal{T}(O_t, \widetilde{W}_{\text{up}}, \gamma_{\text{up}})$
  $Z_t \leftarrow \text{RMSNorm}(X_t + \alpha \odot \mathcal{T}(U'_t, \widetilde{W}_{\text{down}}, \gamma_{\text{down}}))$
  $Z_t^{(l)} \leftarrow \text{RMSNorm}(Z_t^{(l-1)}) + \alpha' \odot Z_t$
end for
// 3. Parallel Resonate-and-Fire (PRF) Somatic Integration
$U_{1:T} \leftarrow \mathcal{F}^{-1}\big(\mathcal{F}(Z^{(l)}_{1:T}) \odot \mathcal{F}(K_{1:T})\big)$ ▷ Frequency-domain global parallel integration
$S^{(l)}_{1:T} \leftarrow \mathcal{H}(\Re\{U_{1:T}\} - V_{\text{th}})$ ▷ Binary spike emission for next block
return $S^{(l)}_{1:T}, Z^{(l)}_{1:T}$

C.4 Deferred Scaling in Temporal Depthwise Projections

For an input $V \in \mathbb{R}^{T\times C}$, ternary kernel $\widetilde{W}_{dw} \in \{-1,0,+1\}^{K\times C}$, and scalar γ_{dw}, the naive and deferred forms are

$$\begin{aligned} X^{\text{naive}}_{t,c} &= \sum_k (\gamma_{dw}\widetilde{W}_{dw,k,c}) V_{t+k,c}, \\ X^{\text{deferred}}_{t,c} &= \gamma_{dw} \sum_k \widetilde{W}_{dw,k,c} V_{t+k,c}. \end{aligned} \tag{S5}$$

They are algebraically equal. In the deferred form, a kernel coefficient of $+1$, -1, or 0 selects addition, subtraction, or omission of the corresponding input value. The accumulated output is then scaled by γ_{dw}. Input normalization, when used, precedes both forms and does not distinguish the two modes. Deferred scaling removes floating-point multiplication by learned kernel weights inside the temporal sum; it does not remove the post-accumulation scaling operation.

D Hardware-Calibrated Energy and Memory Profiling

D.1 Dynamic Energy Formulation and Readout Derivation

Synaptic projection energy scales proportionally with input firing rate $FR \in [0,1]$ and non-zero weight density $NZ \in [0,1]$. In standard 45 nm CMOS technology, floating-point Multiply-Accumulate (MAC) operations consume $E_{\text{MAC}} = 4.6\,\text{pJ}$, while floating-point Accumulate (AC) additions consume $E_{\text{AC}} = 0.9\,\text{pJ}$ Horowitz (2014).

MLGRU Gate Energy (over T timesteps):

1. Spike-Ternary Depthwise Conv ($\mathcal{T}_{dw}$): $\mathcal{O}(K \cdot C)$ operation requiring $T \cdot K \cdot C \cdot FR \cdot NZ$ operations at E_{AC}.
2. Pointwise Conv ($\mathcal{T}$): $\mathcal{O}(C^2)$ operation requiring $T \cdot C^2 \cdot NZ$ operations at E_{AC}.

3. Gate Projections ($\mathcal{T}$): Candidate and gate representations require $3T \cdot C^2 \cdot NZ$ operations at E_{AC}.
4. Output Projection ($\mathcal{T}$): Gated recurrent activations require $T \cdot C^2 \cdot NZ$ operations at E_{AC}.

SwiGLU Channel Mixing Energy: With hidden expansion $E_{\text{swi}} = 4.0$, Up and Gate projections require $8T \cdot C^2 \cdot NZ$ operations, and Down projection requires $4T \cdot C^2 \cdot NZ$ operations, totaling $12T \cdot C^2 \cdot NZ \cdot E_{\text{AC}}$.

Neuronal Update Energy: Discrete membrane updates per PRF neuron require 4 multiplications and 3 additions per timestep Huang et al. (2024). Under 45 nm CMOS (3.7 pJ/mult, 0.9 pJ/add), $E_{\text{PRF}} \approx 17.5$ pJ. Across C channels and T timesteps, this contributes a deterministic $T \cdot C \cdot E_{\text{PRF}}$ cost.

Connectionist Temporal Classification (CTC) Readout Energy: The CTC linear readout projects the final layer's hidden representations to the vocabulary dimension ($V = 10{,}000$ for LibriSpeech; $V = 21{,}128$ for Aishell-1). Because this projection operates on unquantized full-precision weights $W_{\text{ctc}} \in \mathbb{R}^{V \times C}$ (accounting for the 10.24 M non-ternary parameter remainder, or 40.96 MB at FP32), its computational costing requires explicit disambiguation:

- **Input Representation:** The input entering the readout is the binary spike train emitted by the final PRF somatic layer $S^{(L)} \in \{0,1\}^{T \times C}$.
- **Multiplier-Free Accumulation:** For any output vocabulary token $v \in [1, V]$ at timestep t, the logit evaluates as:

$$y_{t,v} = \sum_{c=1}^{C} W_{\text{ctc},v,c} \cdot S_{t,c}^{(L)} = \sum_{c:S_{t,c}^{(L)}=1} W_{\text{ctc},v,c} \tag{S6}$$

  Because $S_{t,c}^{(L)} \in \{0,1\}$, scalar weight multiplication by zero is bypassed, and multiplication by one reduces to an identity fetch. Consequently, the ideal sparse implementation of the linear readout fetches active FP32 weights and accumulates them without floating-point multiplication. This statement concerns readout arithmetic; weight-fetch energy is not included in our estimate.
- **Arithmetic Verification:** LibriSpeech `test-clean` comprises 5.40 hours (19,440 s). Downsampled at 12.5 Hz, the evaluation sequence spans $N = 19{,}440 \times 12.5 = 243{,}000$ discrete frames. The dense dimension product is $N \cdot C \cdot V = 243{,}000 \times 1024 \times 10{,}000 \approx 2.488 \times 10^{12}$ theoretical operations. In continuous FP32 baselines, dense activations trigger continuous MACs at 4.6 pJ:

$$E_{\text{baseline}} = 2.488 \times 10^{12} \times 4.6\,\text{pJ} \approx 11.45\,\text{J} \quad (0.59\,\text{mJ/s}) \tag{S7}$$

  which directly matches the 11.51 J reported for our parameter-matched FP32 baselines in Table S4. In contrast, for our spiking network, the active operation count scales by the empirical final-layer firing rate $FR \approx 42.85\%$ and executes at $E_{\text{AC}} = 0.9$ pJ:

$$E_{\text{stGRU}} = \big(2.488 \times 10^{12} \times 0.4285\big) \times 0.9\,\text{pJ} \approx 1.066 \times 10^{12} \times 0.9\,\text{pJ} \approx \mathbf{0.96}\,\text{J} \quad (\mathbf{0.05}\,\text{mJ/s}) \tag{S8}$$

  This confirms the exact derivation of the 0.96 J figure reported in Tables S3 and S4.

D.2 Theoretical Energy Estimation Methodology for Baselines

To benchmark dynamic energy consumption against continuous ANNs and hybrid SNNs, we dynamically calculate Multiply-Accumulate (MAC) and Accumulate (AC) operations per audio sample following canonical analytical profiling conventions Horowitz (2014); Hou et al. (2024); Song et al. (2025). All continuous baselines replace autoregressive generation with an FP32 Connectionist Temporal Classification (CTC) linear readout costing $N \cdot d \cdot V$ MACs Graves et al. (2006); Hou et al. (2024). The baseline architectures are decomposed as follows:

- **Vanilla Transformer Vaswani et al. (2017); Hou et al. (2024):** The core encoder operations consist of Multi-Head Attention (MHA) and a single Feed-Forward Network (FFN) per

block. The MHA mechanism scales with sequence length N and hidden dimension d, requiring $4d^2N + 2dN^2$ MACs ($3d^2N$ for Q, K, V projections, d^2N for output projection, and $2dN^2$ for attention context matrix products). The standard FFN (expansion factor 4) requires $8d^2N$ MACs. Front-end acoustic subsampling uses dense strided 2D convolutions.

- **Conformer Gulati et al. (2020); Hou et al. (2024):** To capture local acoustic context, the Conformer block introduces a depthwise convolution module adding $K \cdot d \cdot N$ MACs per block for kernel size K. Furthermore, it adopts the Macaron layout, sandwiching attention between two half-step FFN modules, doubling FFN compute to $16d^2N$ MACs per block.
- **Transformer++ Hou et al. (2024):** Discards 2D convolutional down-sampling in favor of linear frame-stacking ($r \cdot d_{\text{input}} \times d$) and removes internal encoder convolutions, while retaining the Macaron dual-FFN structure ($16d^2N$ MACs).
- **Hybrid Spiking Architectures (IML-Spikeformer) Song et al. (2025):** Features are downsampled via a continuous FP32 2D frontend incurring 53.13 J (2.73 mJ/s) of dense MACs. Its spiking encoder uses HD-RepSSA (ND^2 ops), ChannelMLP ($8ND^2$), and attention matrix ops ($2N^2D$) scaled across simulation window $T = 6$ and average firing rate $R \approx 18\%$ (10.58 J / 0.54 mJ/s). However, total inference remains heavily dominated by its dense 6-layer FP32 autoregressive Transformer decoder (32.72 J / 1.68 mJ/s). Whereas IML repeats dynamics over $T = 6$ simulation timesteps per frame, stGRU operates at $T = 1$ without multi-step latency.

Algorithm S2 Profiled Arithmetic-Energy Estimate

Require: Model $\mathcal{M}$, Dataset $\mathcal{D}$, Process constants $E_{\text{AC}} = 0.9\,\text{pJ}$, $E_{\text{PRF}} = 17.5\,\text{pJ}$
Ensure: Total Dynamic Energy E_{total}

$E_{\text{total}} \leftarrow 0$
for each batch $X \in \mathcal{D}$ **do**
  Record forward-pass intermediate activations across all layers $\mathcal{L} \in \mathcal{M}$
  for each layer $\mathcal{L} \in \mathcal{M}$ **do**
    if $\mathcal{L}$ is a Synaptic Ternary Layer **then**
      $NZ \leftarrow \text{count}(\mathcal{L}.weight \neq 0)/\text{total}(\mathcal{L}.weight)$
      $FR \leftarrow (\text{input} \in \{0, 1\})\ ?\ [\text{count}(\mathcal{L}.input = 1)/\text{total}(\mathcal{L}.input)] : 1.0$
      $\text{ActiveOps} \leftarrow \text{TheoreticalOps}(\mathcal{L}) \cdot FR \cdot NZ$
      $E_{\text{total}} \leftarrow E_{\text{total}} + (\text{ActiveOps} \cdot E_{\text{AC}})$
    else if $\mathcal{L}$ is a Neuronal PRF Layer **then**
      $E_{\text{total}} \leftarrow E_{\text{total}} + (\text{total}(\mathcal{L}.output) \cdot E_{\text{PRF}})$
    end if
  end for
end for
return E_{total}

Table S3: **Estimated arithmetic-energy decomposition on LibriSpeech test-clean (5.40 hours; 19,440 s).** For the continuous ANN comparisons, the downstream column is a modeled FP32 CTC readout; for IML-Spikeformer it is a six-layer ANN decoder; for stGRU it is active-weight accumulation in an FP32 CTC readout driven by binary spikes; see Appendix D.1). Estimates omit normalization, activation functions, residual scaling, memory traffic, and decoding/search.

Architecture	Estimated Arithmetic Energy (Joules over 19,440 s)				Estimated Arithmetic-Energy Rate (mJ/s)				Dense Matrix MACs
	Sub-Sampler	Encoder	Decoder	Total	Sub-Sampler	Encoder	Decoder	Total	
Continuous Artificial Neural Networks (ANNs)									
Transformer 112M Hou et al. (2024)	53.13	117.01	1.16	171.31	2.73	6.01	0.06	8.80	Yes
Transformer++ 112M Hou et al. (2024)	0.18	167.73	1.16	169.08	0.009	8.623	0.06	8.69	Yes
Conformer 136M Hou et al. (2024)	53.13	165.19	1.16	219.49	2.73	8.49	0.06	11.28	Yes
Hybrid Spiking Neural Networks (SNNs) & Dense ANN Decoder									
IML-Spikeformer Song et al. (2025)	53.13	10.58	32.72	96.45	2.73	0.54	1.68	4.96	Yes (Hybrid)
Ours (Ternary SNNs)									
Ours Small (Conv1D → MLGRU → RF)	27.46	9.86	0.92	38.24	1.41	0.50	0.04	**1.95**	**No**
Ours Medium (Conv1D → MLGRU → RF)	27.03	20.58	0.94	48.55	1.39	1.05	0.05	2.49	**No**
Ours Large (Conv1D → MLGRU → RF)	26.73	39.01	0.96	66.70	1.374	2.006	0.05	**3.42**	**No**

Table S4: **Stage-by-Stage Dynamic Energy Profiling across Model Scales and Architectures for LibriSpeech and Aishell-1**.

Scale	Architecture Variant	Prec.	Memory (MB)	Estimated Arithmetic Energy (Joules)				Estimated Arithmetic-Energy Rate (mJ/s)			
				Sub	Enc	Dec	Total	Sub	Enc	Dec	Total
Dataset: LibriSpeech test-clean (5.40 Hours — 19,440 s)											
Small	Linear → RF	FP32‡	70.79	24.62	0.22	0.82	25.66	1.26	0.01	0.04	1.31
	Conv1D → RF	FP32‡	71.28	26.25	4.72	0.86	31.83	1.35	0.24	0.04	1.63
	MLGRU → ReLU	FP32	513.32	338.47	75.48	11.51	425.46	17.40	3.88	0.59	21.87
	MLGRU → RF	INT2	70.98	24.02	9.20	0.87	34.09	1.23	0.47	0.04	1.74
	Conv1D → MLGRU → ReLU	FP32	530.63	338.47	80.34	11.51	430.32	17.40	4.13	0.59	22.12
	Conv1D → MLGRU → RF	**INT2**	72.09	27.46	9.86	0.92	38.24	1.41	0.50	0.04	**1.95**
Medium	Linear → RF	FP32‡	87.69	23.22	0.36	0.84	24.42	1.19	0.02	0.04	1.25
	Conv1D → RF	FP32‡	88.69	26.41	9.44	0.88	36.73	1.35	0.48	0.04	1.87
	MLGRU → ReLU	FP32	782.01	338.47	150.97	11.51	500.95	17.40	7.76	0.59	25.75
	MLGRU → RF	INT2	88.09	25.32	19.36	0.89	45.57	1.30	0.99	0.04	2.33
	Conv1D → MLGRU → ReLU	FP32	816.63	338.47	160.68	11.51	510.66	17.40	8.26	0.59	26.25
	Conv1D → MLGRU → RF	**INT2**	90.31	27.03	20.58	0.94	48.55	1.39	1.05	0.05	**2.49**
Large	Linear → RF	FP32‡	117.28	24.00	0.57	0.88	25.45	1.23	0.03	0.04	1.30
	Conv1D → RF	FP32‡	119.56	24.77	17.70	0.93	43.40	1.27	0.91	0.04	2.22
	MLGRU → ReLU	FP32	1252.23	338.47	283.07	11.51	633.05	17.40	14.55	0.59	32.54
	MLGRU → RF	INT2	118.02	27.95	36.86	0.86	65.67	1.43	1.89	0.04	3.36
	Conv1D → MLGRU → ReLU	FP32	1317.55	338.47	301.40	11.51	651.38	17.40	15.49	0.59	33.48
	Conv1D → MLGRU → RF	**INT2**	122.22	26.73	39.01	0.96	66.70	1.374	2.006	0.05	**3.42**
Dataset: Aishell-1 test (10.03 Hours — 36,108 s)											
Small	Linear → RF	FP32‡	116.41	57.55	0.40	2.90	60.85	1.59	0.011	0.08	1.68
	Conv1D → RF	FP32‡	116.91	60.20	8.78	2.70	71.68	1.66	0.24	0.07	1.97
	MLGRU → ReLU	FP32	558.98	629.77	140.45	45.27	815.49	17.44	3.89	1.24	22.57
	MLGRU → RF	INT2	116.61	59.71	18.33	2.98	81.02	1.65	0.50	0.08	2.23
	Conv1D → MLGRU → ReLU	FP32	576.25	629.77	149.49	45.27	824.53	17.44	4.14	1.25	22.83
	Conv1D → MLGRU → RF	**INT2**	117.72	63.58	19.81	3.00	86.39	1.76	0.55	0.08	**2.39**
Medium	Linear → RF	FP32‡	133.32	59.64	0.51	2.55	62.70	1.65	0.01	0.07	1.73
	Conv1D → RF	FP32‡	134.32	60.01	17.56	2.84	80.41	1.66	0.48	0.08	2.22
	MLGRU → ReLU	FP32	827.64	629.77	280.90	45.27	955.94	17.44	7.77	1.25	26.46
	MLGRU → RF	INT2	133.71	59.83	36.15	3.00	98.98	1.65	1.00	0.08	2.73
	Conv1D → MLGRU → ReLU	FP32	862.26	629.77	298.98	45.27	974.02	17.44	8.28	1.25	26.97
	Conv1D → MLGRU → RF	**INT2**	135.94	62.26	39.78	3.01	105.05	1.72	1.10	0.08	**2.90**
Large	Linear → RF	FP32‡	162.91	57.76	0.78	2.59	61.13	1.60	0.02	0.07	1.69
	Conv1D → RF	FP32‡	165.18	59.19	32.93	2.99	95.11	1.63	0.91	0.08	2.62
	MLGRU → ReLU	FP32	1297.86	629.77	526.69	45.27	1201.73	17.44	14.58	1.25	33.28
	MLGRU → RF	INT2	163.65	62.23	68.72	3.06	134.01	1.72	1.90	0.08	3.70
	Conv1D → MLGRU → ReLU	FP32	1363.17	629.77	560.79	45.27	1235.83	17.44	15.53	1.25	34.23
	Conv1D → MLGRU → RF	**INT2**	167.84	61.22	76.29	2.95	140.46	1.69	2.11	0.08	**3.88**

‡ Structural baselines use an INT2 ternary downsampler and FP32 downstream weights; unmarked FP32 variants also use an FP32 downsampler.

E Neuromorphic Interpretation and Analysis

Per-neuron firing rates were evaluated across the LibriSpeech evaluation set. Neurons firing below 0.1 Hz were excluded as recording-duration artifacts ($\leq$ 2 of 1024 neurons per layer; $n = 1022$–1024 per layer). For each layer, we fit lognormal and Gaussian distributions to empirical firing rates and evaluated R^2 as the squared Pearson correlation between empirical and fitted CDFs. We report $\Delta R^2 = R^2_{\text{ln}} - R^2_{\mathcal{N}}$ as the per-layer measure of lognormal superiority. Lognormality peaks at L11 ($+0.091$), persists through L12–L13 ($+0.064$, $+0.041$), and slightly favors Gaussian at L14 (-0.011), capturing the transition from an error-monitoring regime to a broad recruitment readout (Figure S2).

Table S5: **Comprehensive Architectural Comparison on Aishell-1 Mandarin ASR.** Evaluated on Character Error Rate (CER %). Memory sizes are theoretical footprints based on parameter counts and precision formats.

Architecture	Params (M)	Weight Type	Memory (MB)	Energy rate (mJ/s)	CER (%) dev	CER (%) test
Continuous Artificial Neural Networks (ANNs)						
VGG-BiLSTM Ravanelli et al. (2022)	93.30	FP32	373.0†	–	9.7	10.7
Binary S4D* Stan & Rhodes (2023)	21.00	FP32	84.0†	–	10.3	11.9
Transformer* Stan & Rhodes (2023)	61.90	FP32	247.6†	–	5.1	5.6
Hybrid Spiking Neural Networks (SNNs) & Dense ANN Decoder/LM						
LIF* Song et al. (2025)	22.96	FP32	91.84†	–	13.8	15.7
Spiking TCN (+ 6-layer ANN Dec)* Bai et al. (2018)	23.35	FP32	93.40†	–	11.4	13.7
GSU (+ 6-layer ANN Dec)* Hao et al. (2024)	60.74	FP32	242.96†	–	10.9	12.6
Spike-driven Trans. (+ 6-layer ANN Dec)* Yao et al. (2023)	30.35	FP32	121.40†	–	10.2	11.9
IML-Spikeformer (+ 6-layer ANN Dec) Song et al. (2025)	61.89	FP32	247.56†	–	**5.2**	**5.7**
Ours (Structural Baselines, FP32 & W/O Learned ANN Sequence Decoder)						
Large (Linear → RF)	88.42	FP32‡	162.91	1.69	19.95	22.92
Large (Conv1D → RF)	88.98	FP32‡	165.18	2.62	10.81	12.52
Ours (Parameter-Matched Baselines, Continuous Dense ReLU Activations)						
Small (Conv1D → MLGRU → ReLU)	144.06	FP32	576.25	22.83	10.37	11.84
Medium (Conv1D → MLGRU → ReLU)	215.56	FP32	862.26	26.97	11.00	12.60
Large (Conv1D → MLGRU → ReLU)	340.79	FP32	1363.17	34.23	10.62	12.24
Small (Conv1D → MLGRU → ReLU)	144.06	Ternary INT2	117.64	3.56	10.19	11.68
Medium (Conv1D → MLGRU → ReLU)	215.56	Ternary INT2	135.79	4.11	10.60	12.35
Large (Conv1D → MLGRU → ReLU)	340.79	Ternary INT2	167.58	5.18	9.98	11.42
Ours (matrix-MAC-free, Ternary Spiking Networks & W/O Learned ANN Sequence Decoder)						
Ours Small (Conv1D → MLGRU → RF)	144.08	Ternary INT2	**117.72**	**2.39**	10.49	12.36
Ours Medium (Conv1D → MLGRU → RF)	215.60	Ternary INT2	135.94	2.90	9.66	11.44
Ours Large (Conv1D → MLGRU → RF)	340.86	Ternary INT2	167.84	3.88	**9.28**	**10.88**

† Theoretical memory computed from parameter count and format precision.

* Reproduced results based on publicly available data. Reported numbers correspond to seed 3407.

‡ Structural baselines use an INT2 ternary downsampler and FP32 downstream weights; unmarked FP32 variants also use an FP32 downsampler.

Table S6: **Large Model WER (%) on LibriSpeech across 6 Random Seeds**.

Test Set	4-gram LM	Seed 1	Seed 10	Seed 100	Seed 1000	Seed 10000	Seed 3407	Mean ± Std
test-clean	Without	6.53	6.60	6.33	6.46	6.58	6.14	$\mathbf{6.44 \pm 0.17}$
	With	6.01	6.12	5.86	5.98	6.09	5.66	5.95 ± 0.16
test-other	Without	17.49	17.74	17.18	17.74	17.55	17.28	$\mathbf{17.49 \pm 0.23}$
	With	16.54	16.77	16.28	16.75	16.58	16.30	16.53 ± 0.21

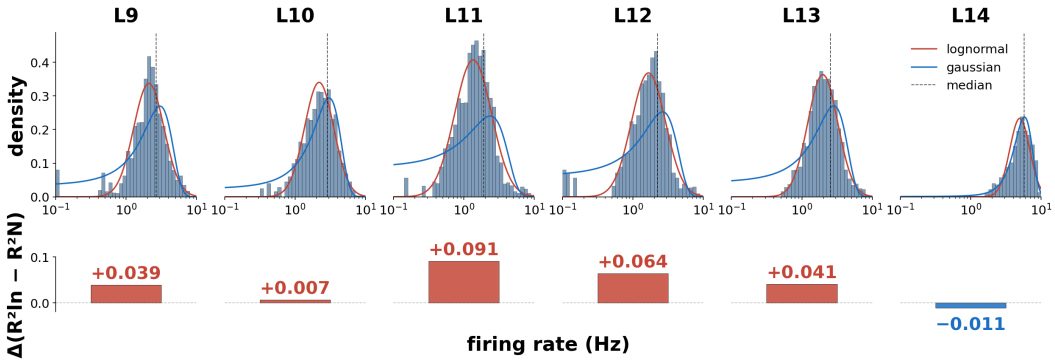

Figure S2: **Depth-dependent change in firing-rate fits (L9–L14).** Top: log-axis firing-rate histograms with lognormal (red) and Gaussian (blue) fits. Bottom: $\Delta R^2 = R^2_{\ln} - R^2_{\mathcal{N}}$; the lognormal fit has its largest advantage at L11 ($+0.091$), while the Gaussian fit is slightly better at L14 (-0.011).

Table S7: **Component-Wise Ablation: Model Scale, Architecture, and Weight Precision across LibriSpeech and Aishell-1**. Isolates individual contributions of low-precision weights, sequence gating, and discrete spiking activations. **Reported numbers correspond to seed 3407.**

Scale	Architecture Variant	Precision	Params (M)	Memory (MB)	Error Rate (%)		Energy rate (mJ/s)
					test-clean / Dev	test-other / Test	
Part I: LibriSpeech-960 Evaluation (WER %)							
Small	Linear → RF	FP32‡	65.39	70.79	19.46	35.40	1.31
	Conv1D → RF	FP32‡	65.51	71.28	14.96	30.39	1.63
	MLGRU → ReLU	FP32	128.33	513.32	14.46	28.75	21.87
	MLGRU → RF	Ternary	128.35	70.98	15.95	30.87	1.74
	Conv1D → MLGRU → ReLU	FP32	132.66	530.63	7.25	19.14	22.12
	Conv1D → MLGRU → ReLU	Ternary	132.66	72.01	7.74	20.26	2.60
	Conv1D → MLGRU → LIF	Ternary	132.67	72.07	8.18	20.65	1.87
	Conv1D → MLGRU → RF	Ternary	132.68	72.09	8.09	20.79	1.95
Medium	Linear → RF	FP32‡	69.61	87.69	18.38	34.72	1.25
	Conv1D → RF	FP32‡	69.86	88.69	13.20	28.10	1.87
	MLGRU → ReLU	FP32	195.50	782.01	12.87	26.87	25.75
	MLGRU → RF	Ternary	195.54	88.09	14.50	29.18	2.33
	Conv1D → MLGRU → ReLU	FP32	204.16	816.63	6.41	17.29	26.25
	Conv1D → MLGRU → ReLU	Ternary	204.16	90.16	6.68	17.71	3.13
	Conv1D → MLGRU → LIF	Ternary	204.18	90.27	7.30	18.72	3.21
	Conv1D → MLGRU → RF	Ternary	204.20	90.31	7.04	18.74	2.49
Large	Linear → RF	FP32‡	77.01	117.28	17.34	33.64	1.30
	Conv1D → RF	FP32‡	77.58	119.56	11.64	26.20	2.22
	MLGRU → ReLU	FP32	313.06	1252.23	12.13	25.24	32.54
	MLGRU → RF	Ternary	313.12	118.02	13.53	28.09	3.36
	Conv1D → MLGRU → ReLU	FP32	329.39	1317.55	5.83	16.11	33.48
	Conv1D → MLGRU → ReLU	Ternary	329.39	121.96	6.49	17.40	4.10
	Conv1D → MLGRU → LIF	Ternary	329.43	122.15	7.05	18.36	3.21
	Conv1D → MLGRU → RF	Ternary	329.45	122.22	6.14	17.28	3.42
	Large (320ms Look-Ahead)	Ternary	329.45	122.22	9.11	22.13	3.51
Part II: Aishell-1 Evaluation (CER %)							
Small	Linear → RF	FP32‡	76.79	116.41	21.20	24.51	1.68
	Conv1D → RF	FP32‡	76.92	116.91	12.68	14.77	1.97
	MLGRU → ReLU	FP32	139.74	558.98	18.32	20.67	22.57
	MLGRU → RF	Ternary	139.76	116.61	17.66	19.87	2.23
	Conv1D → MLGRU → ReLU	FP32	144.06	576.25	10.37	11.84	22.83
	Conv1D → MLGRU → ReLU	Ternary	144.06	117.64	10.19	11.68	3.56
	Conv1D → MLGRU → LIF	Ternary	144.08	117.69	9.15	10.59	2.18
	Conv1D → MLGRU → RF	Ternary	144.08	117.72	10.49	12.36	2.39
Medium	Linear → RF	FP32‡	81.02	133.32	21.20	24.51	1.73
	Conv1D → RF	FP32‡	81.27	134.32	11.63	13.68	2.22
	MLGRU → ReLU	FP32	206.91	827.64	17.52	19.77	26.46
	MLGRU → RF	Ternary	206.95	133.71	18.03	20.51	2.73
	Conv1D → MLGRU → ReLU	FP32	215.56	862.26	11.00	12.60	26.97
	Conv1D → MLGRU → ReLU	Ternary	215.56	135.79	10.60	12.35	4.11
	Conv1D → MLGRU → LIF	Ternary	215.59	135.90	9.19	10.76	2.72
	Conv1D → MLGRU → RF	Ternary	215.60	135.94	9.66	11.44	2.90
Large	Linear → RF	FP32‡	88.42	162.91	19.95	22.92	1.69
	Conv1D → RF	FP32‡	88.98	165.18	10.81	12.52	2.62
	MLGRU → ReLU	FP32	324.46	1297.86	16.40	18.58	33.28
	MLGRU → RF	Ternary	324.53	163.65	16.69	18.98	3.70
	Conv1D → MLGRU → ReLU	FP32	340.79	1363.17	10.62	12.24	34.23
	Conv1D → MLGRU → ReLU	Ternary	340.79	167.58	9.98	11.42	5.18
	Conv1D → MLGRU → LIF	Ternary	340.84	167.77	8.45	10.05	3.71
	Conv1D → MLGRU → RF	Ternary	340.86	167.84	9.28	10.88	3.88

‡ Structural baselines use an INT2 ternary downsampler and FP32 downstream weights; unmarked FP32 variants also use an FP32 downsampler.

Table S8: **Placement of the weight scale in Small and Large stGRUs.** All modes share input normalization. Naive folds γ into the weights; Proposed applies it after ternary accumulation; None omits it. Proposed has substantially lower modeled arithmetic-energy rate than Naive, while omitting γ worsens test-clean WER at both scales. Energy rates use the stated 45 nm operation costs and exclude an external language model.

Model	Scaling Mode	WER (%)		Sparsity (%)		Energy rate
		test-clean	test-other	Activation	Weight	(mJ/s)
Small	Naive $(X(\widetilde{W}\gamma)^\top)$	8.15	20.92	12.93	41.53	9.78
	Proposed $((X\widetilde{W}^\top)\gamma)$	8.09	20.79	12.84	41.76	**1.95**
	None $(X\widetilde{W}^\top)$	9.44	22.61	14.44	39.81	2.21
Large	Naive $(X(\widetilde{W}\gamma)^\top)$	6.44	17.43	17.36	38.57	17.74
	Proposed $((X\widetilde{W}^\top)\gamma)$	**6.14**	**17.28**	17.08	39.15	**3.42**
	None $(X\widetilde{W}^\top)$	7.51	19.20	19.64	38.16	3.75

Table S9: **Empirical Training Throughput for 10 Epochs on LibriSpeech-960**: PRF vs. LIF on 8 AMD Instinct MI250X GCDs.

Model Size	Neuron Type	Training Time (10 epochs)
Small	LIF	8h 53m
Small	PRF	**2h 40m** (3.3× speedup)
Medium	LIF	14h 35m
Medium	PRF	**4h 44m** (3.1× speedup)
Large	LIF	25h 00m
Large	PRF	**7h 21m** (3.4× speedup)

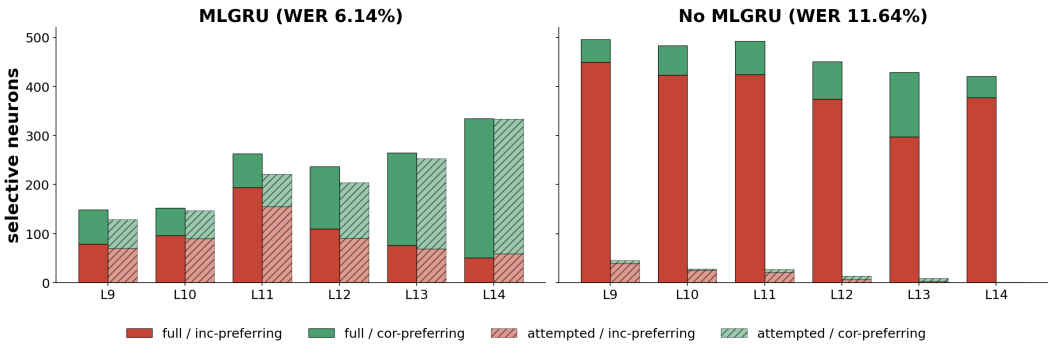

Figure S3: **Outcome-selective neuron counts (L9–L14) for full vs. attempted-emission-only trial pools, MLGRU vs. no-MLGRU (Large variants)**. Selective neuron counts per layer, split by outcome preference (incorrect-preferring, red; correct-preferring, green). Left: Full MLGRU (Conv1D → MLGRU → RF). Empty-emission rate on incorrect words is 13.82%; restricting to attempted emissions reduces aggregate selective count by only 8.1%. Right: Ablated variant without MLGRU (Conv1D → RF). Empty-emission rate is 57.76%; restricting to attempted emissions collapses aggregate selective count by 95.5%, supporting a role for MLGRU gating for the emergent error-monitoring property.