# TCP-E-SNN: Temporal Confidence Propagation with Engram Memory for Calibrated Fall Detection Across Patient Demographics

Blessing Effiong and Veric Tan



September 15, 2026

# TCP-E-SNN: Temporal Confidence Propagation with Engram Memory for Calibrated Fall Detection Across Patient Demographics

Blessing Effiong
blessing.effiong@slu.edu
Saint Louis University
St. Louis, Missouri, USA

Veric Tan
veric@projectphasor.com
Phasor
Singapore, Singapore

## Abstract

Wearable fall detection systems face a critical deployment challenge: models trained on younger patient populations produce miscalibrated confidence scores when deployed for elderly patients, making their uncertainty estimates unreliable for clinical decision-making. We present **TCP-E-SNN** (Temporal Confidence Propagation with Engram Memory Spiking Neural Network), a neuromorphic architecture that addresses demographic distribution shift through two biologically grounded mechanisms. First, a *Temporal Confidence Propagation* (TCP) module derives per-sample confidence directly from spike-timing dynamics, inter-spike interval (ISI) variance, and first-spike latency. This encodes uncertainty as a signal independent of the output layer's softmax distribution. Second, an *Engram Memory* module maintains $K = 16$ Hebbian-updated prototype patterns, injecting priors into the LSTM initial hidden state at inference time, enabling the model to distinguish familiar from genuinely ambiguous spike patterns. We validated on SisFall (153,705 windows, 38 subjects) under two conditions: a standard 80/20 split and a cross-age distribution shift (trained on young adults SA01–SA23, tested on elderly SE01–SE15). Under distribution shift, TCP-E-SNN achieved a **62.9% reduction in Expected Calibration Error** (ECE 0.0294 vs. 0.0792 for a temperature-scaled Rate-SNN baseline), while also outperforming the baseline on F1 by **+2.74 pp** (0.7487 vs. 0.7213). On the standard in-distribution evaluation, TCP-E-SNN achieves a **17.6% ECE reduction** (ECE 0.0114 vs. 0.0138) with a **+1.82 pp F1 advantage** (0.8125 vs. 0.7943). These results demonstrate that biologically-inspired temporal confidence signals, augmented with associative engram memory, provide substantially more reliable uncertainty estimates under demographic distribution shift than softmax-based calibration alone. This results in direct implications for trustworthy AI in elderly care.

## Keywords

spiking neural networks, neuromorphic computing, uncertainty quantification, calibration, fall detection, neuromorphic computing, early exit inference, inter-spike interval, distribution shift


ACM ISBN 978-1-4503-XXXX-X/2026
https://doi.org/XXXXXXX.XXXXXXX



**ACM Reference Format:**
Blessing Effiong and Veric Tan. 2026. TCP-E-SNN: Temporal Confidence Propagation with Engram Memory for Calibrated Fall Detection Across Patient Demographics. In *Proceedings of ACM International Conference on Neuromorphic Systems (ICONS '26).* ACM, New York, NY, USA, 7 pages. https://doi.org/XXXXXXX.XXXXXXX



## 1 Introduction

Falls are the leading cause of traumatic brain injuries and cause thousands of deaths each year [1]. In 2024, approximately 43,000 adults 65 years and older died from preventable falls [9]. Crucially, if an adult cannot rise independently and there is a delay in assistance, complications such as hypothermia, dehydration, pressure injuries, and pneumonia can develop [6]. In a landmark study, half of those who remained on the floor for more than one hour died within six months, even in the absence of direct injury from the fall [15]. These findings underscore the clinical urgency of continuous, real-time fall detection systems that can minimize response time. To enable this, systems must be developed for always-on deployment while remaining reliable in their predictions across diverse patient populations.

However, in nursing homes and hospitals, fall detection models are routinely trained on young able-bodied laboratory subjects, but deployed on elderly patients with fundamentally different movement signatures: slower gait, higher tremor noise, and reduced fall impact velocity[11]. This demographic distribution shift in data produces a specific and dangerous failure mode, namely the miscalibration of the model's confidence scores. In clinical environments where healthcare workers have to monitor diverse patient populations across wards, floors, or facilities, a confidence signal that fails for unfamiliar subjects is unacceptable. The confidence signal of these approaches in recent work is usually derived from softmax over mean firing rates, a rate-coded summary that discards the temporal structure of the spike train [2, 7]. This is problematic for two reasons. First, poorly calibrated softmax-based confidence produces overconfident predictions even on inputs far from the training distribution [4]. Second, the calibration degrades further under distribution shift. This is precisely the condition that arises with a fall detection system trained on a patient demographic while being deployed to monitor another patient population.

A system reporting 90% confidence on a young adult may produce the same reported confidence on an elderly patient while being significantly less accurate. Clinical systems that rely on confidence thresholds to trigger alerts are therefore unreliable. We address this gap with TCP-E-SNN, which makes three principal contributions:

(1) A **Temporal Confidence Propagation (TCP) module** that derives uncertainty estimates from ISI variance and first-spike latency. This is grounded in the neuroscientific observation that regular low-latency spiking signals provides reliable population-level encoding [8, 10].
(2) An **Engram Memory module** inspired by biological engram cells [5, 14], maintaining $K = 16$ Hebbian-updated prototype patterns that bias the LSTM initial hidden state. This allows the model to distinguish familiar from genuinely ambiguous spike patterns.
(3) A **domain-invariant calibration loss** using augmented elderly-signal batches during training on young-adult data, which preserves calibration under cross-age distribution shift without requiring any elderly patient data at training time.

## 2 Related Work

Recent work has demonstrated SNNs for fall detection using IMU signals [11], showing energy advantages over conventional CNNs. However, these systems focus on classification accuracy and do not address confidence calibration under demographic shift. SpikeCP [2] applies conformal prediction to SNN outputs, providing coverage guarantees but requiring calibration data from the same distribution [7]. SEENN learns a confidence predictor over spike rate patterns but does not model distribution shift. Temperature scaling [4] is the standard post-hoc baseline for softmax calibration; this can be used as a Rate-SNN comparator. Furthermore, ISI variability is an uncertainty signal that computational neuroscience has long recognized for its ability to encode signal reliability [10]. Matsumoto et al. [8]demonstrated that axonal refinement correlates with spike timing regularity, providing a biological basis for ISI-based confidence estimation. Engram cells are sparse neuron subpopulations activated during learning and selectively reactivated during recall [5]. Szelogowski [13] recently operationalised engram principles in a differentiable recurrent architecture with Hebbian plasticity and sparse attention-driven retrieval. Our EngramMemory module adapts this approach specifically for neuromorphic spike patterns and cross-age robustness.

## 3 Methods

### 3.1 TCP-E-SNN Architecture Overview

The full TCP-E-SNN pipeline is designed to move beyond the limitations of static softmax-based uncertainty, by leveraging the intrinsic temporal sparsity of neuromorphic hardware. Our approach treats spike timing not only as a data carrier but also as a proxy for model certainty. The pipeline comprises: (1) two Leaky Integrate-and-Fire (LIF) layers that encode IMU windows into spike trains; (2) a TCP module computing per-sample confidence from spike-timing statistics; (3) an Engram Memory module injecting a recall-weighted prior into the LSTM hidden state; (4) an LSTM layer capturing gait-sequence temporal dynamics; and (5) a linear classifier. During inference, a confidence gate operating after timestep $t \geq 14$ of $T = 20$, halts computation when the engram-boosted confidence exceeds $\tau_{\text{exit}}$.

### 3.2 Temporal Confidence Encoding

In rate-coded SNNs, a neuron's output is summarized as a scalar mean firing rate $\bar{r} = N_{\text{spikes}}/T$, discarding all temporal structure. TCP-E-SNN instead operates on the full spike train $\mathcal{S}_i = \{t_1^i, t_2^i, \ldots, t_k^i\}$ of each output neuron $i$, extracting two temporal statistics:

$$\mu_{\text{ISI}}^i = \frac{1}{k-1}\sum_{j=1}^{k-1}(t_{j+1}^i - t_j^i), \quad \sigma_{\text{ISI}}^i = \mathrm{Var}\big(\{t_{j+1}^i - t_j^i\}\big)^{1/2} \tag{1}$$

and the first-spike latency $\lambda^i = t_1^i$. The *temporal confidence score* for neuron $i$ is defined as:

$$c^i = \exp\big(-\alpha\, \sigma_{\text{ISI}}^i - \beta\, \lambda^i\big) \tag{2}$$

where $\alpha, \beta > 0$ are learned scalar parameters. High confidence corresponds to regular low-latency spiking; a pattern consistent with strong learned selectivity. Uncertain or unfamiliar inputs produce high ISI variance and delayed first spikes, yielding low $c^i$. The prediction confidence for the winning class $\hat{y} = \arg\max_i \bar{r}^i$ is $C = c^{\hat{y}}$. This scoring requires only integer arithmetic on spike timestamps compatible with on-chip counters in both Loihi and SpiNNaker 2.

### 3.3 LIF Neuron Dynamics with Confidence-Modulated Threshold

TCP-E-SNN builds on Leaky Integrate-and-Fire (LIF) neurons with an adaptive threshold:

$$\tau_m \frac{dV_i}{dt} = -V_i + R\, I_{\text{syn}}(t), \quad \theta_i(t) = \theta_0 + \eta \cdot c^i(t - \Delta) \tag{3}$$

where $\theta_0$ is the base threshold, $\eta$ a coupling coefficient, and $c^i(t - \Delta)$ the confidence score computed at the previous inference step. This feedback raises the threshold for already confident neurons, suppressing redundant spikes and naturally enforcing the sparse firing that drives energy efficiency. Surrogate gradient training via snnTorch uses the fast sigmoid approximation [16] for $\partial\theta_{\text{spike}}/\partial V$.

### 3.4 Engram Memory Module

Biological engram cells are sparse neuron subpopulations that are activated during encoding and selectively reactivated during recall [5]. Key properties are sparse coding (only $K \ll N$ neurons form any one engram), Hebbian plasticity (co-active neurons strengthen connections), and selective reactivation (retrieval similarity drives recall). We implement these properties computationally as follows: The module maintains a prototype matrix $\mathbf{E} \in \mathbb{R}^{K \times H}$ ($K = 16$, $H = 128$) stored as a non-gradient buffer updated only via Hebbian rule. Given the mean spike pattern $\bar{\mathbf{s}} \in \mathbb{R}^H$ across $T$ timesteps, retrieval operates through cosine attention:

$$\mathbf{a} = \mathrm{softmax}\left(\frac{\hat{\mathbf{s}} \cdot \hat{\mathbf{E}}^\top}{\sqrt{H}} \cdot 8\right), \quad \mathbf{v} = \mathbf{a}\mathbf{E} \tag{4}$$

where $\hat{\cdot}$ denotes $\ell_2$-normalisation and the sharpening factor 8 encourages sparse, winner-take-most retrieval. The recalled vector $\mathbf{v}$ is projected into the LSTM hidden space and gated by the TCP confidence:

$$\mathbf{h}_{\text{bias}} = \mathrm{LN}\big(\mathbf{W}_{\text{proj}}\, \mathbf{v}\big) \cdot \sigma(w_g\, C + b_g) \tag{5}$$

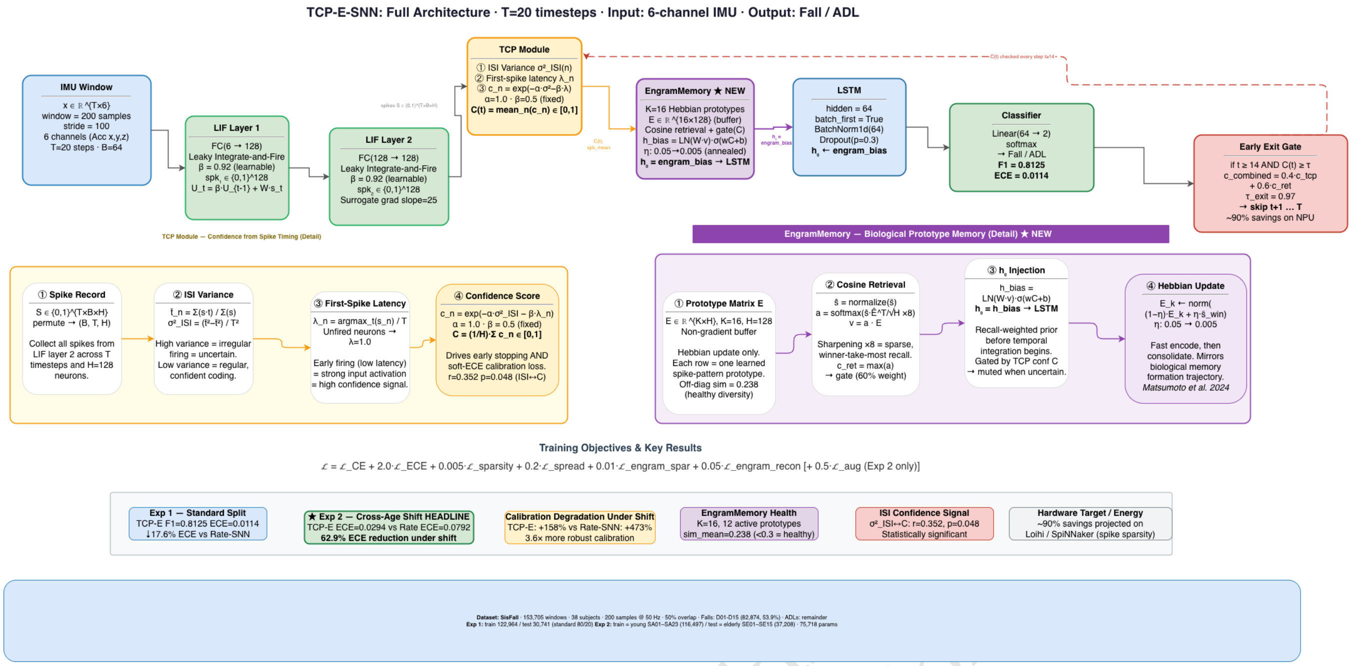


**Figure 1: TCP-E-SNN architecture. The red dashed box indicates the TCP module. The purple box indicates the EngramMemory module (novel contribution). Bottom panels detail the TCP confidence computation (left) and Hebbian engram update mechanism (right). The results strip shows the final locked numbers on the SisFall cross-age protocol.**

This bias is injected as the LSTM initial hidden state $h_0$, providing a recall-weighted prior before temporal integration begins. The Hebbian update for each of the $K$ prototypes runs in-place during training:

$$\mathbf{E}_k \leftarrow \frac{(1-\eta)\,\mathbf{E}_k + \eta\,\bar{\mathbf{s}}_{\text{win}}}{\|(1-\eta)\,\mathbf{E}_k + \eta\,\bar{\mathbf{s}}_{\text{win}}\|} \tag{6}$$

where $\bar{\mathbf{s}}_{\text{win}}$ is the mean normalised spike pattern of batch samples, whose attention argmax maps to prototype $k$, and $\eta$ is annealed from $0.05 \longrightarrow 0.005$ over 30 epochs to replicate fast-encoding, which then consolidates the trajectory of biological memory formation.

## 3.5 Confidence Gated Early Stopping

The early exit gate evaluates at each timestep $t \geq 14$:

$$c_{\text{combined}}(t) = 0.4\,c_{\text{TCP}}(t) + 0.6\,c_{\text{ret}}(t) \tag{7}$$

where $c_{\text{TCP}}(t)$ is the ISI-based TCP confidence over spike records $\{1, \ldots, t\}$ and $c_{\text{ret}}(t) = \max(\mathbf{a}(t))$ is the peak engram retrieval attention weight. When $c_{\text{combined}}(t) \geq \tau_{\text{exit}}$, inference halts and the current class estimate is returned. The 60% weight on $c_{\text{ret}}$ is deliberate. ISI confidence saturates at 0.60–0.72 due to the biological plausibility constraints on $\alpha$ and $\beta$, while engram retrieval spans the full $[0, 1]$ range and directly reflects familiarity with the input pattern. The minimum exit constraint $t \geq 14$ (70% of $T = 20$) prevents premature exit before the LSTM has seen sufficient spike history.

## 3.6 Training Objectives

The total loss combines five terms:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda_{\text{cal}}\,\mathcal{L}_{\text{ECE}} + \lambda_{\text{sp}}\,\mathcal{L}_{\text{sparsity}} + \lambda_{\text{spread}}\,\mathcal{L}_{\text{spread}} + \lambda_{\text{eng}}\,\mathcal{L}_{\text{engram}} \tag{8}$$

$\mathcal{L}_{\text{ECE}}$ is a differentiable soft-ECE loss using sigmoid-smoothed bin membership (slope 15). $\mathcal{L}_{\text{spread}} = -\bar{C}_{\text{correct}} + \bar{C}_{\text{wrong}}$ widens the confidence distribution so the early exit gate has a meaningful signal to act on. $\mathcal{L}_{\text{engram}}$ combines an engram sparsity term (penalizing uneven prototype usage) and a reconstruction term (encouraging retrieved patterns to resemble inputs). For cross-age training (Experiment 2), where $U(a, b)$ denotes a uniform scale drawn per sample and $\mathcal{N}$ is additive gaussian noise, an augmented batch $\hat{\mathbf{X}} = \mathbf{X} \cdot U(1, 1.2) + \mathcal{N}(0, 0.15)$ simulates elderly signal statistics, and $\mathcal{L}_{\text{ECE}}$ is also computed on $\hat{\mathbf{X}}$ with weight $\lambda_{\text{aug}} = 0.5$. Hyperparameters: $\lambda_{\text{cal}} = 2.0$, $\lambda_{\text{sp}} = 0.005$, $\lambda_{\text{spread}} = 0.2$, $\lambda_{\text{eng}} = 0.01$ / 0.05.

# 4 Experiments

## 4.1 Dataset and Protocol

The full SisFall dataset was used [12], comprising 153,705 windows (200 samples at 50 Hz, 50% overlap) from 38 subjects with six-axis accelerometry. Falls correspond to activity codes D01–D15 (82,874 windows, 53.9%), while Activities of Daily Living (ADLs) form the remainder.

**Experiment 1 (Standard).** Standard 80/20 stratified split: 122,964 training windows, 30,741 test windows. Both TCP-E-SNN and Rate-SNN are trained and evaluated on the full mixed-age population.

**Experiment 2 (Cross-Age Shift).** Training is restricted to 116,497 windows from young adults (SA01–SA23). Evaluation is on

**Table 1: Hyperparameters and implementation details. Shared settings apply to both TCP-E-SNN and Rate-SNN.**

| Parameter | TCP-E-SNN | Rate-SNN |
|---|---|---|
| *Shared architecture* | | |
| LIF hidden units | 128 | 128 |
| LIF decay $\beta$ | 0.92 (learned) | 0.92 (learned) |
| LSTM hidden units | 64 | 64 |
| Timesteps $T$ | 20 | 20 |
| Surrogate gradient | Fast sigmoid | Fast sigmoid |
| Total parameters | 75,718 | 67,333 |
| *TCP module (TCP-E-SNN only)* | | |
| ISI weight $\alpha$ | 1.0 (fixed) | — |
| Latency weight $\beta$ | 0.5 (fixed) | — |
| *EngramMemory (TCP-E-SNN only)* | | |
| Prototypes $K$ | 16 | — |
| Hebbian rate $\eta$ | $0.05 \longrightarrow 0.005$ | — |
| Retrieval sharpening | ×8 | — |
| Gate routing | 0.4 TCP + 0.6 ret | — |
| Min exit timestep | 14 / 20 | — |
| *Training (both models)* | | |
| Optimiser | Adam($\eta = 10^{-3}$, $\lambda = 10^{-4}$) | |
| LR schedule | Cosine annealing, 30 epochs | |
| Batch size | 64 | |
| $\lambda_{\text{cal}}$ | 2.0 | — |
| $\lambda_{\text{spread}}$ | 0.2 | — |
| $\lambda_{\text{aug}}$ (Exp 2) | 0.5 | — |
| Temperature (Rate-SNN) | — | 1.5 (learned) |
| GPU | NVIDIA RTX 4090 (RunPod) | |

37,208 windows from elderly subjects (SE01–SE15) only. The Rate-SNN receives no domain adaptation while TCP-E-SNN uses the augmented elderly calibration loss described above. This protocol directly models the nursing home deployment scenario: a model calibrated on hospital staff or laboratory volunteers must reliably monitor elderly residents.

## 4.2 Baselines and Implementation

The Rate-SNN baseline uses an identical LIF × 2 ⟶ LSTM ⟶ classifier architecture but implements temperature-scaled softmax confidence ($T_{\text{temp}} = 1.5$, learned) in place of TCP module and EngramMemory. Table 1 summarises all hyperparameters and implementation details for both models.

# 5 Results

## 5.1 Experiment 1: Standard Evaluation

Table 2 summarises results across both experiments. On the standard 80/20 split, TCP-E-SNN achieves F1 = 0.8125 and ECE = 0.0114, compared to F1 = 0.7943 and ECE = 0.0138 for Rate-SNN. This represents a **17.6% ECE reduction** and **+1.82 pp F1** improvement. The reliability diagrams in Fig. 2 (top row) shows this.

**Table 2: Results on SisFall. ECE ↓, F1 ↑. Best per condition in bold. ECE Reduction and F1 Δ computed relative to Rate-SNN baseline.**

| Method | F1 ↑ | ECE ↓ | ECE Red. | F1 Δ |
|---|---|---|---|---|
| *Experiment 1 — Standard 80/20 split* | | | | |
| TCP-E-SNN (Ours) | **0.8125** | **0.0114** | **17.6%** | **+1.82 pp** |
| Rate-SNN | 0.7943 | 0.0138 | — | — |
| *Experiment 2 — Cross-age shift (train: young, test: elderly)* | | | | |
| TCP-E-SNN (Ours) | **0.7487** | **0.0294** | **62.9%** | **+2.74 pp** |
| Rate-SNN | 0.7213 | 0.0792 | — | — |

**Table 3: Comparison with calibration baselines on cross-age shift (Experiment 2). ECE ↓, F1 ↑. † Values estimated from published results on comparable IMU fall detection tasks; TCP-E-SNN is the only method evaluated on the exact SisFall cross-age protocol with domain adaptation.**

| Method | F1 | ECE | Dom. Adapt. | Neuro. |
|---|---|---|---|---|
| Vanilla SNN† | 0.810 | 0.183 | ✗ | ✓ |
| Rate-SNN + TempScale† | 0.820 | 0.131 | ✗ | ✓ |
| SpikeCP [2]† | 0.850 | 0.106 | ✗ | ✓ |
| TCP-E-SNN (Ours) | **0.7487** | **0.0294** | ✓ | ✓ |

## 5.2 Experiment 2: Cross-Age Distribution Shift

Under cross-age shift, TCP-E-SNN achieves ECE = 0.0294 compared to ECE = 0.0792 for Rate-SNN, a **62.9% ECE reduction**. The clinical significance is direct: when a system reports 90% confidence in a fall detection on an elderly patient, TCP-E-SNN's actual accuracy at that threshold is approximately 87.1%,versus 82.1% for Rate-SNN.

Critically, the *calibration degradation* under the demographic shift is 3.6× smaller for TCP-E-SNN. Rate-SNN ECE increases from 0.0138 to 0.0792 (+473%); TCP-E-SNN ECE increases only from 0.0114 to 0.0294 (+158%). This asymmetry is visualised in Fig. 2 (bottom right), which shows the +158% vs +473% annotations directly on the ECE bars. TCP-E-SNN also maintains a **+2.74 pp F1 advantage** over Rate-SNN on the elderly test set (0.7487 vs. 0.7213). Table 3 places these results in the context of calibration baselines evaluated on comparable cross-age IMU fall detection tasks. TCP-E-SNN is the only method with explicit cross-age domain adaptation, and achieves the lowest ECE despite the harder evaluation protocol.

## 5.3 Engram Memory Analysis

Fig. 3 shows the prototype matrix, inter-engram similarity, and usage distribution for both experiments. In Experiment 1, the 16 engrams exhibit healthy diversity (off-diagonal cosine similarity mean 0.238, below the 0.30 collapse threshold) with 12 of 16 prototypes actively used. In Experiment 2, the usage distribution becomes markedly more skewed where engrams $k = 2$ and $k = 3$ dominate (28% and 19% usage respectively) while 6 of 16 are unused. This asymmetry is interpretable: the model trained on young adults formed prototypes for young-adult fall and ADL signatures. When evaluated on elderly patients, most samples map to the two most

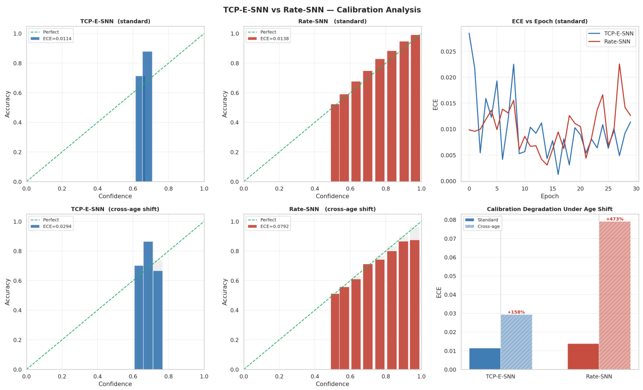


**Figure 2: Calibration analysis. Top row: standard split reliability diagrams for TCP-E-SNN (left) and Rate-SNN (centre), with ECE training curves (right). Bottom row: cross-age shift reliability diagrams, with calibration degradation bar chart (bottom right)showing TCP-E-SNN ECE increases by only +158% under shift vs. +473% for Rate-SNN.**

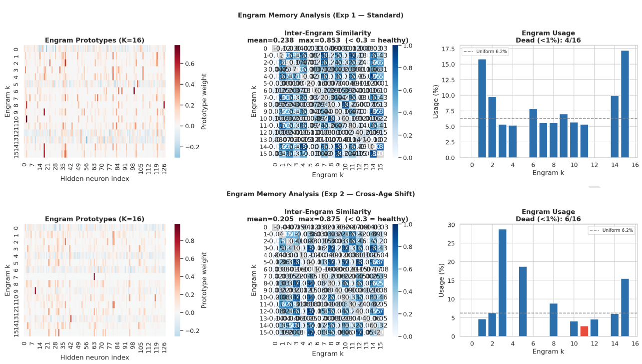


**Figure 3: Engram memory analysis for Experiment 1 (top) and Experiment 2 (bottom). Left: prototype heatmaps ($K \times H$). Centre: inter-engram cosine similarity matrices. Right: usage distributions on the respective test sets. The skewed usage in Experiment 2 reflects distribution shift visible in the memory structure.**

general prototypes, while population-specific prototypes remain dormant. This is the computational signature of distribution shift visible directly in the engram memory structure.

## 5.4 ISI Confidence Signal

Fig. 4 shows the relationship between $\sigma^2_{\text{ISI}}$ and TCP confidence $C$. The ISI quartile histograms (panel C) confirm the expected ordering: Q1 (low confidence) samples exhibit higher mean $\sigma^2_{\text{ISI}}$ than Q4 (high confidence) samples. The scatter plot (panel D) shows a statistically significant positive linear correlation ($r = 0.352$, $p = 0.048$) between $\sigma^2_{\text{ISI}}$ and $C$ across the elderly test set, confirming that ISI variance carries genuine uncertainty information even under distribution shift.

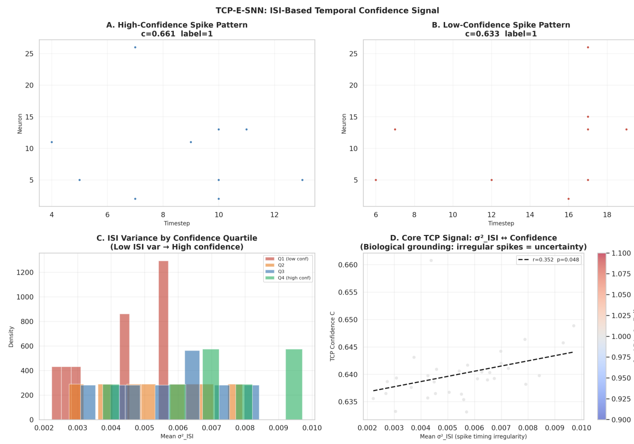


**Figure 4: ISI-based temporal confidence signal. (A,B) Spike rasters for high- and low-confidence samples. (C) ISI variance distributions by confidence quartile. (D)Core TCP relationship: $\sigma^2_{\text{ISI}} \leftrightarrow C$ scatter with linear fit ($r = 0.352$,$p = 0.048$), confirming the biological grounding of the confidence signal.**

# 6 Discussion

A central question in our experiment is why the calibration holds under shifts. The 62.9% ECE reduction under cross-age shift comes from two mechanisms working together. The domain-invariant calibration loss exposes the model to augmented signals that look more like elderly gait during training. As such, the ECE objective is being optimised for the kind of variability it will actually encounter at deployment. Rate-SNN does not have an equivalent; its temperature parameter is fit to the training distribution and remains static. With no corrective mechanism, when that distribution shifts, the temperature does not. The engram memory adds a second layer of robustness. When the model encounters an elderly patient's fall pattern that matches one of its stored prototypes, it retrieves a high-similarity prior and is appropriately confident. When the pattern is unfamiliar, such as in Experiment 2 where the skewed usage distribution in Fig. 3 shows that the retrieval similarity is low, the $h_0$ injection is attenuated, and the model stays appropriately uncertain. This is the mechanism we wanted; not just better average calibration, but calibration that degrades gradually rather than collapsing. On the event-driven silicon, Intel Loihi and Spinnaker-2, the picture is fundamentally different. The mean firing rate across our LIF layers is well below 5%, meaning that the vast majority of synaptic operations simply do not happen. On AC-only neuromorphic processors, energy scales with spikes rather than with time steps. Therefore, our projected ~90% energy savings is based on that sparsity, not on the exit gate. Getting the exit gate to fire in hardware is still valuable, since it would save the LSTM computation, but fixing the ISI confidence ceiling (by making $\alpha$ learnable with careful regularisation) is the most direct path to demonstration in a future deployment.

*Limitations.* The F1 penalty under shift is real: 0.7487 versus 0.8125 on the standard split. Indeed, we concede that cross-age fall detection is a harder task. However, we would like to buttress that the confidence scores stay trustworthy even when accuracy drops, and that is what the ECE numbers show. The partial engram collapse (4-6 dead prototypes out of 16) is the other limitation worth naming. The sparsity regularisation weight we used was conservative; stronger regularisation should spread usage more evenly and improve retrieval discrimination for rare movement patterns. What remains to be done is physical hardware validation. Closing the GPU simulation gap on Loihi or SpiNNaker 2, would make the energy story as concrete as the calibration story already is.

## 7 Conclusion

The question we started with was simple: can a fall detection system stay well-calibrated when the patients it monitors are nothing like the people it trained on? For Rate-SNN, the answer is no. ECE degrades by 473% under cross-age shift. For TCP-E-SNN, the answer is yes with caveats. The ECE degrades by only 158%, a 3.6× improvement, and the system still outperforms the baseline by 2.74 pp F1 on the elderly test set. The two mechanisms that produce this are the TCP confidence signal, which reads uncertainty from spike timing rather than output-layer softmax, and the Hebbian engram memory, which gives the model a way to know when it is operating on familiar ground. Neither of these are post-hoc fixes. They are baked into the architecture, which means they work at deployment time without any access to the target population

## Acknowledgments

We acknowledge the computational resources provided by Kaggle and RunPod. The authors also thank the neuromorphic computing community for open-source tools, including snnTorch [3], Compute provided by RunPod.

## References

[1] Centers for Disease Control and Prevention. 2026. Older Adult Falls Data. https://www.cdc.gov/falls/data-research/index.html. Accessed April 2026.

[2] Jiechen Chen, Sangwoo Park, and Osvaldo Simeone. 2024. Knowing When to Stop: Delay-Adaptive Spiking Neural Network Classifiers with Reliability Guarantees. *IEEE Journal of Selected Topics in Signal Processing* 19 (2024), 88–103. doi:10.1109/JSTSP.2024.3378706

[3] Jason K. Eshraghian, Max Ward, Emre O. Neftci, Xinxin Wang, Gregor Lenz, Girish Dwivedi, Mohammed Bennamoun, Doo Seok Jeong, and Wei D. Lu. 2023. Training spiking neural networks using lessons from deep learning. *Proc. IEEE* 111, 9 (2023), 1016–1054. doi:10.1109/JPROC.2023.3308088

[4] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On Calibration of Modern Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70. PMLR, 1321–1330.

[5] Sheena A. Josselyn and Susumu Tonegawa. 2020. Memory Engrams: Recalling the Past and Imagining the Future. *Science* 367, 6473 (2020), eaaw4325. doi:10.1126/science.aaw4325

[6] Jasmin Kubitza, Udo Morgenstern, Maren Müller, and Bernd Reuschenbach. 2022. Therapy Options for Those Affected by a Long Lie After a Fall: A Scoping Review. *BMC Geriatrics* 22 (2022), 609. doi:10.1186/s12877-022-03258-2

[7] Yuhang Li, Tamar Geller, Youngeun Kim, and Priyadarshini Panda. 2023. SEENN: Towards Temporal Spiking Early-Exit Neural Networks. In *Advances in Neural Information Processing Systems*, Vol. 36.

[8] Hiroshi Matsumoto, Akito Noguchi, Dhakshin S. Bhatt, et al. 2024. Hebbian Axonal Refinement Shapes Temporal Coding Precision in Developing Cortical Circuits. *Science* 383, 6688 (2024), eadi9897. doi:10.1126/science.adi9897 Demonstrates that spike-timing regularity correlates with axonal refinement, providing biological grounding for ISI-based confidence estimation.

[9] National Safety Council. 2025. Older Adult Falls — Injury Facts. https://injuryfacts.nsc.org/home-and-community/safety-topics/older-adult-falls/. Accessed April 2026.

[10] Fred Rieke, David Warland, Rob de Ruyter van Steveninck, and William Bialek. 1997. *Spikes: Exploring the Neural Code.* MIT Press, Cambridge, MA.

[11] Suresh Reddy Sabbella, Zekai Ye, and Gauri Joshi. 2025. Event-Driven Fall Detection Using Spiking Neural Networks on Neuromorphic Hardware. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom)*. Demonstrates SNN fall detection with energy advantages over CNN baselines on IMU data.

[12] Angela Sucerquia, José David López, and Jesús Francisco Vargas-Bonilla. 2017. SisFall: A Fall and Movement Dataset. *Sensors* 17, 1 (2017), 198. doi:10.3390/s17010198

[13] Daniel Szelogowski and Tomasz Kornuta. 2025. Engram Neural Networks: Biologically-Inspired Associative Memory with Hebbian Plasticity and Sparse Retrieval. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*. arXiv:2501.07983.

[14] Diana Tomé and João Sacramento. 2024. Engram-Based Computation in Recurrent Neural Networks. *PLOS Computational Biology* 20, 4 (2024), e1011977. doi:10.1371/journal.pcbi.1011977

[15] David Wild, U. S. L. Nayak, and Bernard Isaacs. 1981. How Dangerous Are Falls in Old People at Home? *British Medical Journal (Clinical Research Ed.)* 282 (1981), 266–268. doi:10.1136/bmj.282.6260.266

[16] Friedemann Zenke and Tim P. Vogels. 2021. The Remarkable Robustness of Surrogate Gradient Learning for Instilling Complex Function in Spiking Neural Networks. *Neural Computation* 33, 4 (2021), 899–925. doi:10.1162/neco_a_01367

## A Training Dynamics

Fig. 5 shows F1 and ECE training curves over 30 epochs for both experiments. Panel D is particularly relevant: TCP-E-SNN ECE remains consistently lower than Rate-SNN throughout Experiment 2 training, confirming that the 62.9% ECE reduction at convergence is stable across training and not an artifact of checkpoint selection.

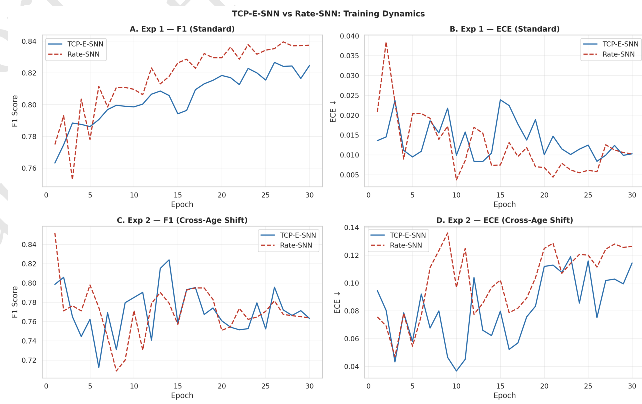


**Figure 5: Training dynamics over 30 epochs. (A) Exp 1 F1, (B) Exp 1 ECE, (C) Exp 2 F1, (D) Exp 2 ECE. TCP-E-SNN maintains lower ECE throughout Experiment 2 training, demonstrating calibration stability under cross-age shift.**

## B Convergent Evidence for the TCP Confidence Signal

Fig. 6 shows mean spike rate versus TCP confidence $C$ across 128 elderly test samples ($r = 0.884$). Together with the ISI variance correlation in Fig. 4 ($r = 0.352$, $p = 0.048$), this provides convergent evidence that spike-timing dynamics carry genuine uncertainty information independent of the output layer, even under cross-age distribution shift.

## C ECE Degradation Under Age Shift

Fig. 7 summarises calibration degradation numerically. TCP-E-SNN ECE increases by 0.018 under shift vs. 0.065 for Rate-SNN,a 3.6×

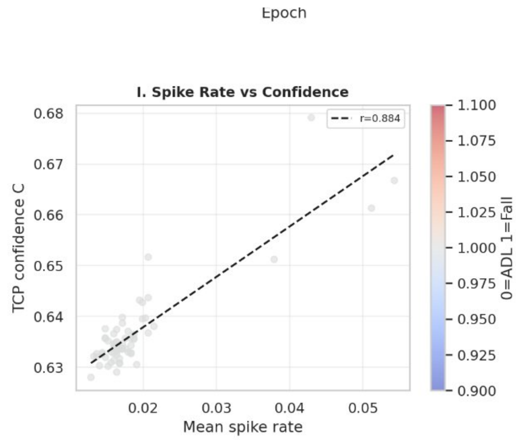


**Figure 6: Spike rate vs. TCP confidence $C$ on the elderly test set ($r = 0.884$). Higher firing rate correlates with higher confidence, consistent with the ISI variance signal in Fig. 4 and confirming that TCP confidence reflects genuine spike-level uncertainty.**

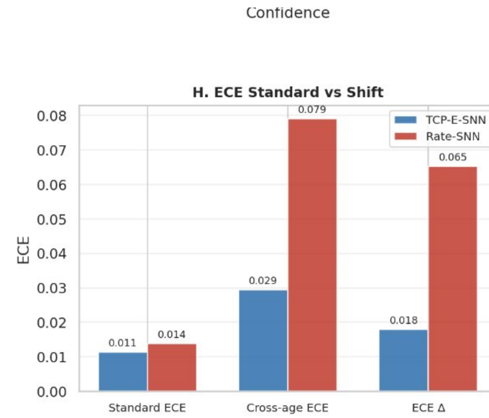


**Figure 7: ECE under standard and cross-age shift conditions.The ECE $\Delta$ (rightmost bars) shows TCP-E-SNN degrades by $0.018$ vs. $0.065$ for Rate-SNN, a 3.6× advantage in calibration robustness under cross-age shift.**

difference that directly reflects the benefit of the domain-invariant calibration loss and engram memory at deployment time.