SNNSys – Workshop on Spiking Neural Networks

Bernhard A. Moser*
SCCH, Hagenberg-Linz
bernhard.moser@scch.at

Michael Lunglmayr
JKU Linz, Inst. of Signal Processing
michael.lunglmayr@jku.at

Robert Legenstein
Institute of Theoretical Computer Science, Graz University of Technology
robert.legenstein@tugraz.at

Overview

This volume presents the proceedings of the SNNSys Workshop at the AIRoV Symposium, reflecting a pivotal moment in neuromorphic engineering. The contributions collected here highlight a dual-track progression: the deepening scientific understanding of temporal processing and the pragmatic hurdles of industrial technology adoption.

The Science of Spiking Systems: Sparsity, Recurrence, and Temporal Dynamics

The research papers in this edition push the boundaries of how we model and train spike-based intelligence, moving beyond simple pattern recognition toward a fundamental science of dynamic computation:

- **Mathematical Foundations of Sparsity:** A core scientific pillar of SNNs is the exploitation of event-driven sparsity. Windhager et al. (Paper 43) elevate this from a heuristic to a formal optimization problem. By applying Linearized Bregman Iterations, they demonstrate that sparsity is not just a byproduct of spiking, but a property that can be mathematically promoted via proximal soft-thresholding. Their introduction of AdaBreg bridges the gap between classical sparse optimization and modern adaptive moment estimation (Adam), providing a rigorous framework for training networks that are efficient by design.
- **The Functional Necessity of Recurrence:** A critical scientific debate in the field concerns whether SNNs can be simplified into parallelizable, feed-forward structures. Mayr et al. (Paper 66) provide a compelling counter-argument, proving that nonlinear recurrence and membrane dynamics are functionally necessary for true stateful computation. Their comparison across ECG and state-tracking tasks shows that without these recurrent loops, SNNs lose the ability to generalize in complex, time-dependent scenarios—aligning with the biological principle that state-tracking is an emergent property of recurrent processing.
- **Temporal Memory and Heterogeneous Delays:** Expanding on the science of time, Perrinet (Paper 82) investigates how synaptic delays contribute to working memory. By moving away from uniform delays, the research shows that heterogeneity allows a network to autonomously replay sparse patterns. This suggests that the "memory" of an SNN is stored not just in synaptic weights, but in the specific temporal architecture of its connections.
- **Stochasticity and Gradient Descent:** The science of learning in SNNs often struggles with the non-differentiable nature of spikes. Higuchi et al. (Papers 79 & 80) tackle this

*Double Affiliation: JKU Linz, Inst. of Signal Processing

The Third Austrian Symposium on AI, Robotics, and Vision (AIROV26).

through two lenses: probabilistic modelling and algorithmic simplification. Their work on Probabilistic LIF Neurons suggests that noise and stochasticity improve learning landscapes in recurrent settings, while their rehabilitation of one-step truncated BPTT proves that high-order temporal averaging can be achieved through momentum rather than computational complexity.

Breakout Session about the Reality of Technology Adoption

Despite these scientific strides, our "all-hands" breakout session revealed significant friction in moving SNNs from the lab to the market. While the "Science of SNNs" promises a $100\times$ leap in energy efficiency, the reality of "Technology Adoption" faces a $61B edge AI market dominated by highly optimized conventional pipelines. Key challenges identified include:

- **The Data Ingestion Bottleneck:** Perhaps the most significant "hidden" hurdle is the front-end interface. Today's market is almost 100% dominated by conventional, Shannon-based Analog-to-Digital Converters (ADCs). These converters produce frame-based, dense data streams that negate the inherent sparsity of SNNs from the start. For SNNs to achieve their full efficiency potential, we require a commercial shift toward event-based sensing or neuromorphic ADCs, yet the infrastructure for these remains immature compared to the entrenched Shannon-sampling paradigm.
- **The Hardware-Robustness Trade-off:** The challenge of integrating analogue computation (aiming for 10%–90% hybrid chips) without succumbing to the noise that compromises scientific robustness.
- **The Ecosystem Gap:** The lack of standardized software and the high R&D costs required to compete with mature GPU/TPU ecosystems.
- **Targeting the Breakthrough:** Identifying "killer apps" for 2026—specifically hearing aids and micro-robotics—where $<$1ms latency and $<$100mW power constraints make SNN science a commercial necessity rather than an alternative.

Conclusion

These proceedings demonstrate that while the science of Spiking Neural Systems is maturing into a rigorous discipline of temporal AI, successful adoption requires translating these scientific insights into solutions that fit the strict constraints of the industrial edge.