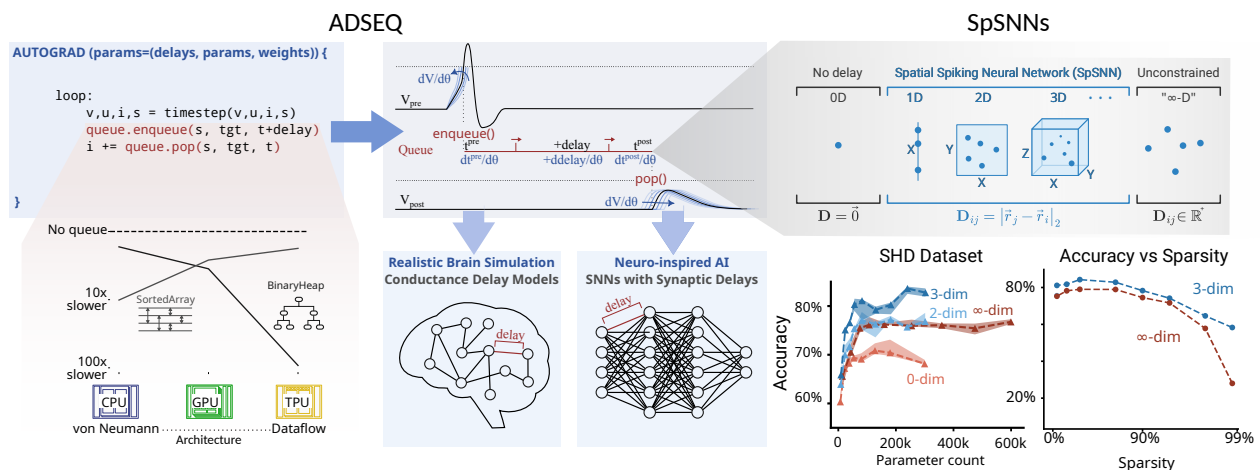


# Autograd-enabled Exact Event-gradients for Spatial Spiking Neural Networks

L. P. L. Landsmeer<sup>(1,2)</sup>, A. Movahedin<sup>(1,2)</sup>, M. Negrello<sup>(1)</sup>, S. Hamdioui<sup>(2)</sup>, C. Strydis<sup>(1,2)</sup>

<sup>(1)</sup> Quantum & Computer Engineering, TU Delft, The Netherlands, <sup>(2)</sup> NeurocomputingLab, Dept. of Neuroscience, Erasmus Medical Center, Rotterdam, The Netherlands

Neurons have been observed to adapt, maintain or learn, not just synaptic weights, but also axonal conduction delays. Recent developments in SNNs for ML have focussed on bringing such delay tuning to neural models, however, here, the focus has been on achieving good ML-task performance given the tuning of arbitrary delay values, for simplified neuron models. In contrast, in biological networks, delays arise from the spatial embedding of neurons, and the physical distance action potentials have to travel. Furthermore, representative neuron models are also much more complex, which are, as of yet, do not support delay learning. To overcome this, we developed a method and library for general Autodifferentiable Spike Event Queues (ADSEQ) for arbitrary neuron models and, built on top of this, introduce Spatial Spiking Neural Networks (SpSNNs).



ADSEQ builds on a generalization of the exact gradients of EventProp by introducing delay tuning and rewriting it to work in an autograd environment, such that the modeller does not have to manually implement gradients throughout the model. On top of that, it includes various queue implementations with different performance tradeoffs according to modelling needs and constraints, and the underlying hardware (CPU, GPU, TPU). In general, capacity-limited FIFOs execute optimally on CPU and memory-constrained GPUs, delay-sized buffers on GPUs and self-sorting queues on TPUs. To show its generality, a differentiable synapse implementation for the Jaxley brain simulator is provided. In SpSNNs, neurons are embedded in a low-dimensional Euclidean space, and synaptic delays are determined by inter-neuronal distance, while synaptic weights and neuron positions are optimized jointly during learning, using the methods introduced in ADSEQ. We evaluated SpSNNs on two neuromorphic temporal processing tasks. Networks with spatially derived delays achieved equal or higher classification accuracy compared to networks with freely learned delays, despite a substantial reduction in the number of trainable parameters. Performance was maximal for networks embedded in the vicinity of two to three spatial dimensions, whereas higher-dimensional embeddings did not yield further improvements, which suggests a regularization effect due to spatial embedding. Together, it is shown that ADSEQ could be key building block for general differentiable brain models with tunable delays, and that on top of that, SpSNNs provide a framework for reducing the resulting increase in parameter counts, while retaining equal or better performance.