A piecewise approximation theory for spiking neural networks

Dominik Dold and Philipp Petersen

Faculty of Mathematics and Research Network DataScience @ Uni Vienna, University of Vienna, Austria

How expressive are spiking neural networks (SNNs)? And what spike encoding is best? We present two recent theoretical results that address these questions: **(i)** A novel metric that provides lower bounds of the approximation rate for SNNs using single-spike coding [1], and **(ii)** a transference principle showing that bounds on approximation rates for single-spike neural networks transfer directly to multi-spike neural networks [2].

Our first result is centred around the concept of *causal pieces*, which is inspired from *affine pieces* used to analyse ReLU artificial neural networks (ANNs). A causal piece is a region of the input and parameter space of an SNN where the same set of neurons and synapses cause the spikes of the output neurons. These pieces partition the input and parameter space into a mosaic of distinct pieces (Figure A, shown only for inputs). Within each piece, the SNN behaves as a continuous function of the input, making the overall SNN a piecewise continuous, non-linear function.

We prove that in the case of leaky integrate-and-fire neurons (LIF) with large membrane time constants and at most one spike per neuron, the number of causal pieces p lower bounds the error SNNs achieve: $\|\psi - f\|_\infty > c \cdot p^{-2}$, where ψ is the SNN, $f \in C^3$ a non-affine function, $\|\cdot\|_\infty$ the supremum-norm, and $c > 0$ a constant. Consequently, we hypothesise that initialising SNNs such that the training data samples fall into many causal pieces is beneficial for training, which we demonstrate in simulation for the Yin-Yang, Fashion MNIST, and EuroSAT RGB datasets (Figure B,C). Preliminary results for standard single-spike LIF neurons on Yin-Yang confirm this trend.

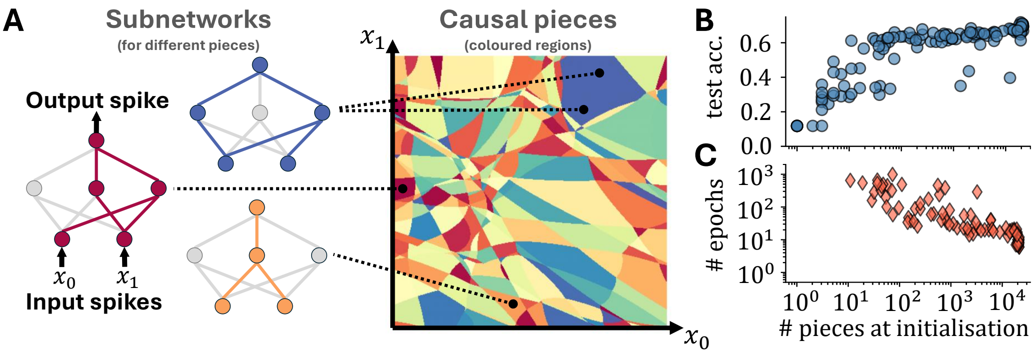

Figure 1: **A.** Every coloured region is a causal piece, fully partitioning the input space (here: 2D inputs). Within each piece, the same subnetwork causes the SNN output (left, coloured like the corresponding piece). **B.** More pieces at initialisation correlates with training success (shown for EuroSAT). **C.** Number of epochs required to reach a test accuracy of 0.5 (shown for EuroSAT).

Both for studying biology and in machine learning, neurons that spike multiple times are usually of interest though. By introducing specific circuit constructions, we prove that single-spike and multi-spike neural networks are equivalent in their approximation capabilities. Specifically, we show that any bound on approximation rates for single-spike neural networks also applies to multi-spike ones, and vice versa – with an at most linear (in the maximum number of spikes multi-spike neurons can emit over a time interval) increase in network size. This result is valid for a large class of neuron models, such as LIF with subtractive reset mechanism. Thus, every approximation bound for single-spike neural networks (e.g., [3]) also applies to multi-spike ones, including the one presented above.

[1] D. Dold and P.C. Petersen, arXiv preprint, arXiv:2504.14015, 2025.

[2] D. Dold and P.C. Petersen, arXiv preprint, arXiv:2603.13478, 2026.

[3] A. M. Neuman et al., JMLR, 26(246), 1-49, 2025.