

Audio Event Classification with Spiking Neural Networks and Eligibility Propagation

Audio event classification is commonly addressed with Convolutional Neural Networks (CNNs) trained by backpropagation, yet such approaches remain biologically implausible and can be computationally demanding for temporal signals. Eligibility propagation, or e-prop (Bellec *et al.*, 2020), offers an online alternative for training recurrent Spiking Neural Networks (SNNs) by combining local eligibility traces with neuron-specific learning signals, approximating gradient-based learning without backpropagation through time. In this work, we investigate whether this learning principle can be applied to an audio classification problem: multilabel drum onset detection.

We implement a recurrent SNN with leaky and adaptive leaky integrate-and-fire neurons, and train it using e-prop to classify 14 drum-event classes from time-frequency audio representations. The same task is also addressed with a CNN trained with backpropagation, providing a conventional baseline. The models are evaluated using micro- and macro-F1 scores, with particular attention to class imbalance and low-data training regimes.

Preliminary results suggest that SNNs trained with e-prop can approximate the overall performance of CNNs trained with backpropagation while using half the number of training examples, remaining approximately 10 percentage points below the CNN in micro-F1. In low-data experiments, including one-shot and five-shot settings, the e-prop model outperforms the CNN by 5.5 and 7.5 percentage points, respectively, in micro-F1. These findings support the potential of eligibility-based learning as a biologically plausible and data-efficient alternative for temporal audio classification, and motivate further work on online learning in neuromorphic auditory systems.