Hierarchically Gated Spiking Neural Networks for Interpretable Arrhythmia Detection

Rahul Pemmaraju, Anshul Voleti, Aaron D. Milstein
Rutgers University
Piscataway, NJ, USA

While traditional deep learning algorithms have achieved impressive performance for automatic arrhythmia detection, such methods typically require significant computational resources, making them infeasible for continuous on-device monitoring. In this work we introduce a neuromorphic architecture, a hierarchically gated recurrent spiking neural network (HG-RSNN). Our network augments a typical recurrent spiking neural network (RSNN) with learnable gates and top-down connectivity. Learnable gates are used to restrict feature sharing between layers and across time, allowing for the extraction of multi-timescale features. Visualization of gate activity highlights important regions of ECG waveforms, providing a method for understanding information flow through the model. We evaluate model performance on a binary arrhythmia detection task using ECG recordings acquired from the PhysioNet MIT-BIH Arrhythmia Database. Our model achieves average area under the receiver-operator curve of 0.902 showing an improvement of +0.120 over a standard recurrent spiking neural network. We further assess model performance when trained using a variant of eligibility propagation (e-prop). e-prop is a biologically plausible and memory efficient approximation of the backpropagation-through-time algorithm (BPTT), the standard algorithm for training recurrent neural networks. e-prop condenses past information into constant-memory eligibility traces, reducing memory requirements. We show that our model can be effectively trained using e-prop. We see marginal loss in performance compared to BPTT, illustrating how our method could be used for direct on-device learning.

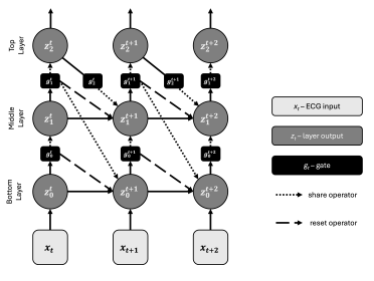

Figure 1. HG-RSNN architecture unrolled over several timesteps. Information is passed through learnable gates.

Model	AUROC	Sensitivity	Specificity
RSNN	0.78 ± 0.07	0.79 ± 0.17	0.65 ± 0.24
HG-RSNN	0.90 ± 0.03	0.91 ± 0.05	0.65 ± 0.29
HG-RSNN (e-prop)	0.89 ± 0.02	0.91 ± 0.05	0.60 ± 0.22

Table 1. Average classification performance of a standard RSNN and HG-RSNN either trained with BPTT (top) or e-prop (bottom) across 5 train-validation splits. Average values ± standard deviation are reported.