Computation using Latency Dynamics in Living and Artificial Neural Networks

Riley T. Zeller-Townson, G. Kumar Venayagamoorthy, Steve M. Potter
Georgia Institute of Technology, Atlanta, GA

Spiking Neural Networks use the precise timing of action potentials to convey meaning. The conduction delays between neurons are one set of parameters that can be tuned to improve network performance on computational tasks, however no biologically inspired delay learning rules have been adopted by the artificial neural network community. This work shows the computational properties of delay update rules that are based on how delay change in living neural networks, as well as how the actual biological data can be used to improve performance for a prediction task.

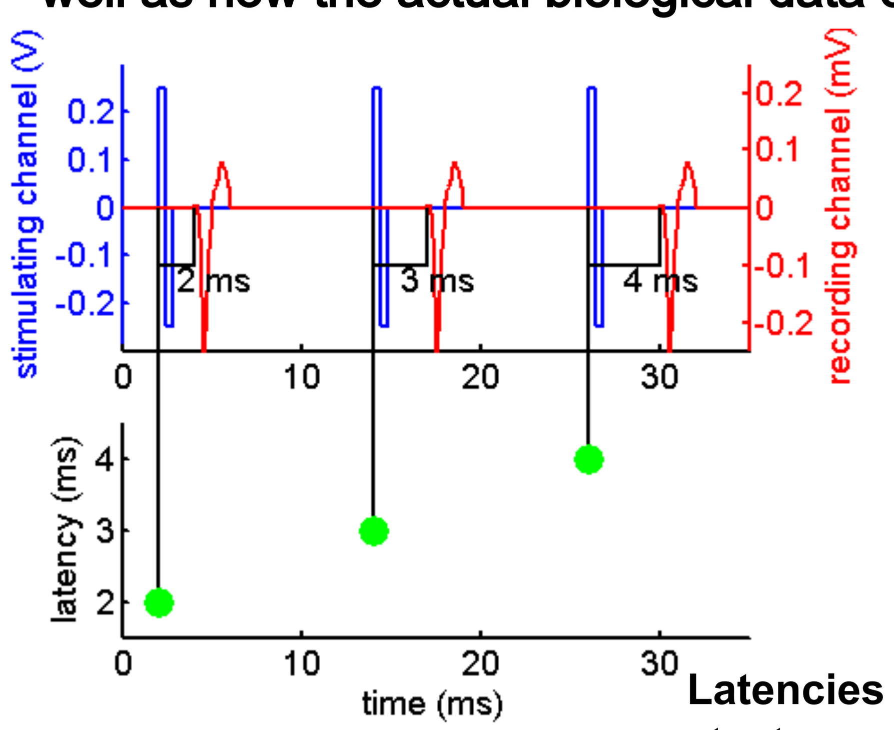

Conduction delays were measured through extracellular stimulation and recording of action potentials. Dissociated cortical cultures were grown on Micro Electrode Arrays (MEAs). Electrical stimulation and extracellular recording were performed with the open-source electrophysiology system NeuroRighter. Stimulation was performed in the presence of synaptic blockers to prevent spontaneous action potentials. Delays were measured by finding the time between pulsatile electrical stimulation (simulating spike initiation) on one electrode in the dish, and detection of a reliably evoked action potential elsewhere in the MEA. This time is referred to as action potential latency (left), and is a direct measure of how long an action potential takes to travel along an section of axon.

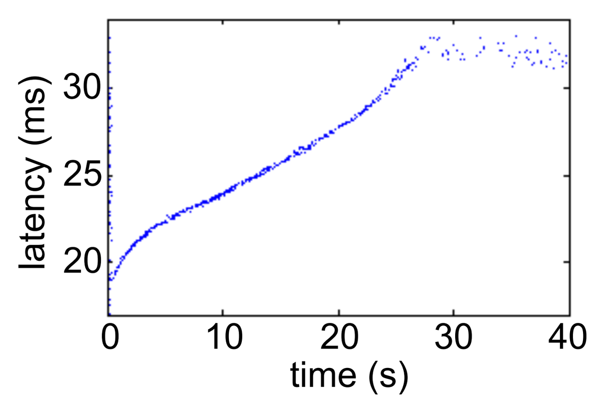

Latencies tend to increase during repetetive stimulation (left). When stimulating a neuron at rates greater than 5 hz, latencies will increase until reaching an 'intermittent regime' of behavior, where latency stops increasing and becomes more variable, and conduction becomes less reliable. While generally a decrease in stimulation frequency is necessary to bring a neuron out of this mode of behavior, we have observed latencies spontaneously decreasing under constant stimulation frequency (right).

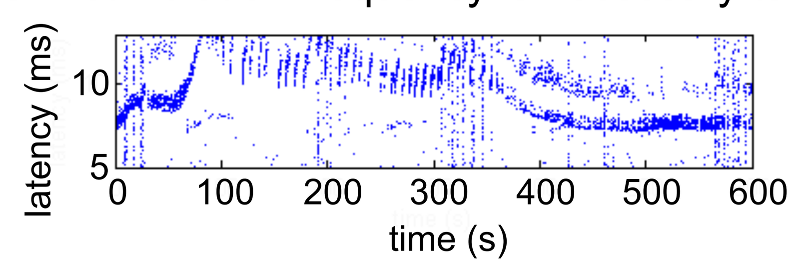

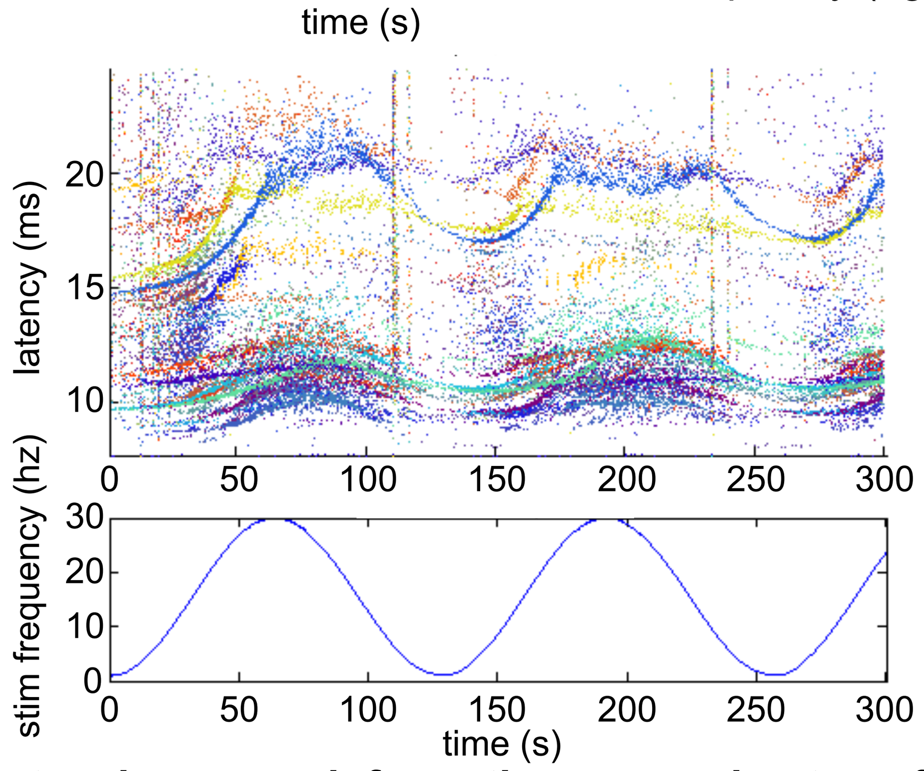

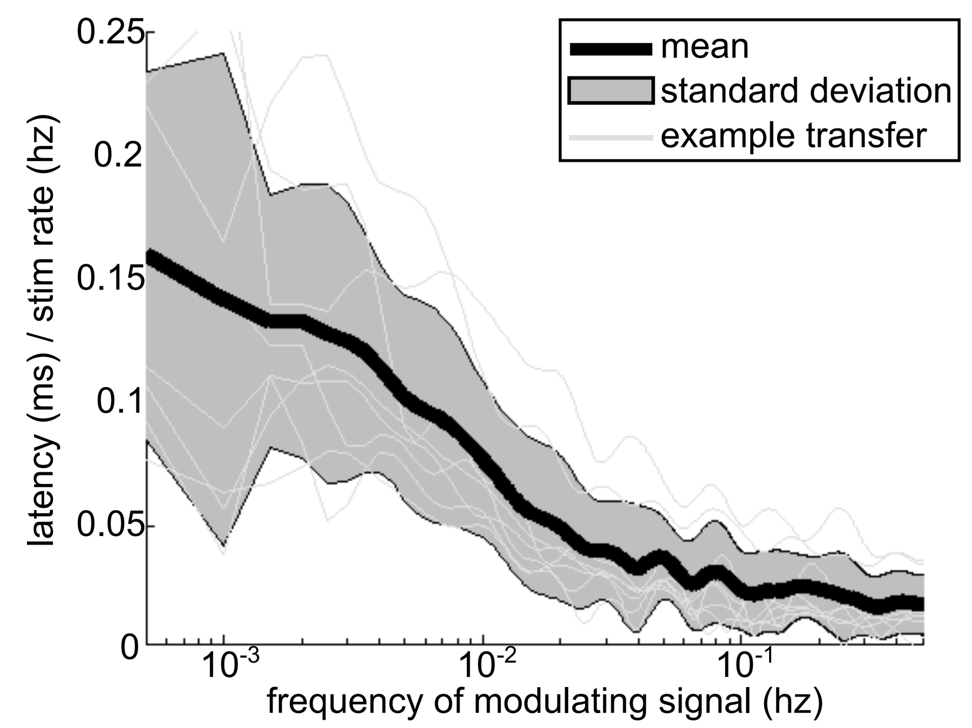

Latencies are an information perserving transform of firing rate. Stimulation frequency can be modulated by altering the duration of the pauses between stimuli. In the figure above, a low frequency sinusoid (1/128 hz) is used to modulate stimulation frequency between 1 and 30 hz. (Above, left) latencies recorded on several recording electrodes (each electrode is given a different color) replicate aspects of the stimulation modulating signal. In this way, latencies can be thought of as a transform of stimulation frequency, or firing rate. Though this transform clearly has non-linear components, if one choses stimulation sequences carefully to prevent latencies from entering the intermittent regime, this transform closely resembles a low-pass filter (above, right).

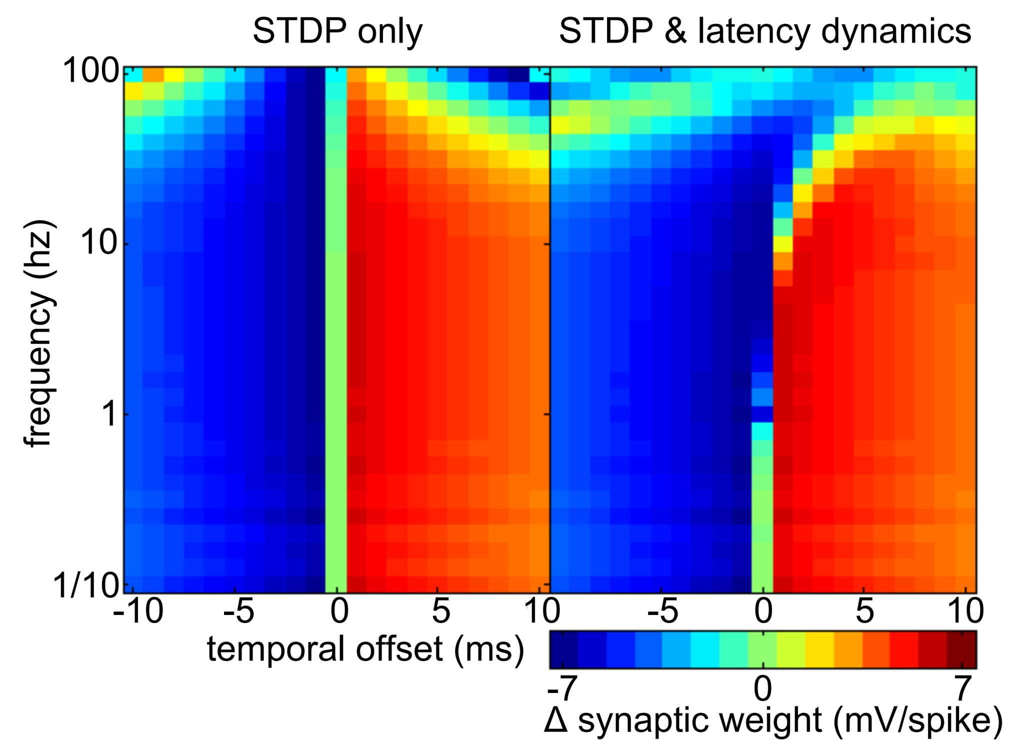

Latency dynamics interact with Spike Timing Dependent Plasticity. By adding a simple, linear model of latency changes to a common model of STDP, firing rate becomes a factor that can be used to manipulate synaptic plasticity, leading to an increase in synaptic depression at high firing rates.

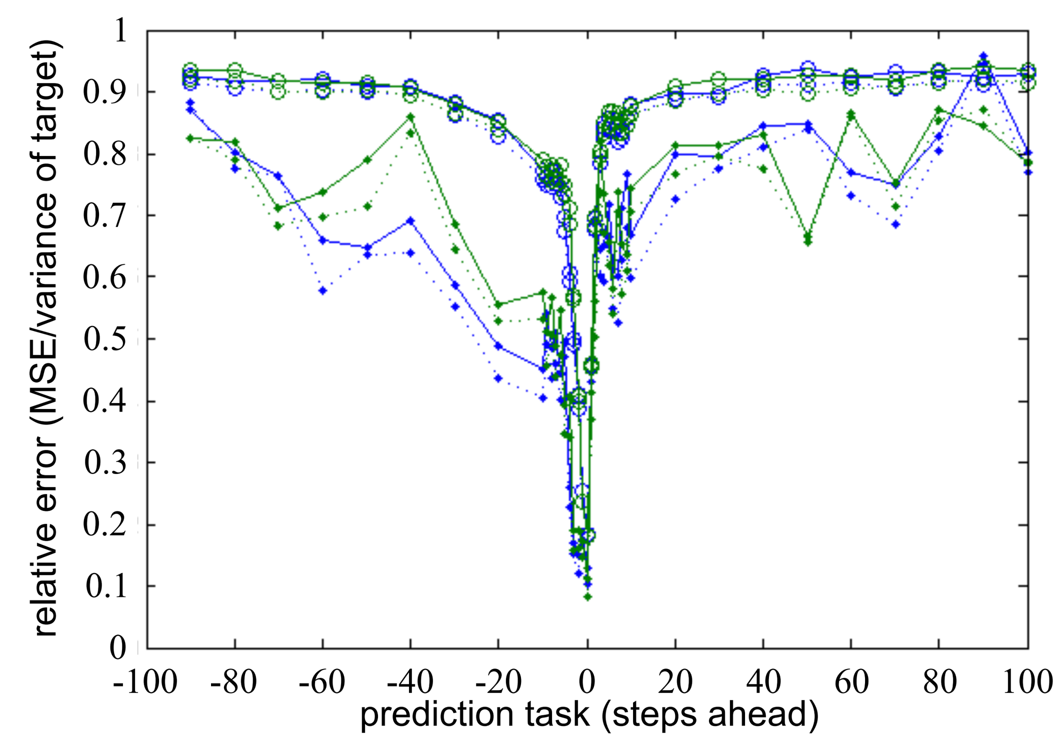

Latency dynamics increase network memory in a Biologically-Inspired Artificial Neural Network. We adopted the method presented in Johnson et al for performing prediction with a biologically inspired artificial neural network. Briefly, 96 spiking neurons were used to encode terminal voltages and turbine speeds for a 4 generator simulated power grid. These 'sensory neurons' projected into a recurrent 'reservoir' of 1000 spiking neurons. Firing rates for all 1096 neurons were collected and mapped onto a prediction using either linear regression (open dots) or a MLP (closed dots). In this experiment, both prediction (regression target is the the terminal voltage of a generator some number of samples ahead) and memory (regression target is the terminal voltage of a generator some number of samples in the past) were tried. The network that used a simple model of dynamic delays (blue) performed equally well to the network that did not use dynamic delays (green) on the prediction task. However, the dynamic delay network performed with 21% less error, on average, on the memory task.

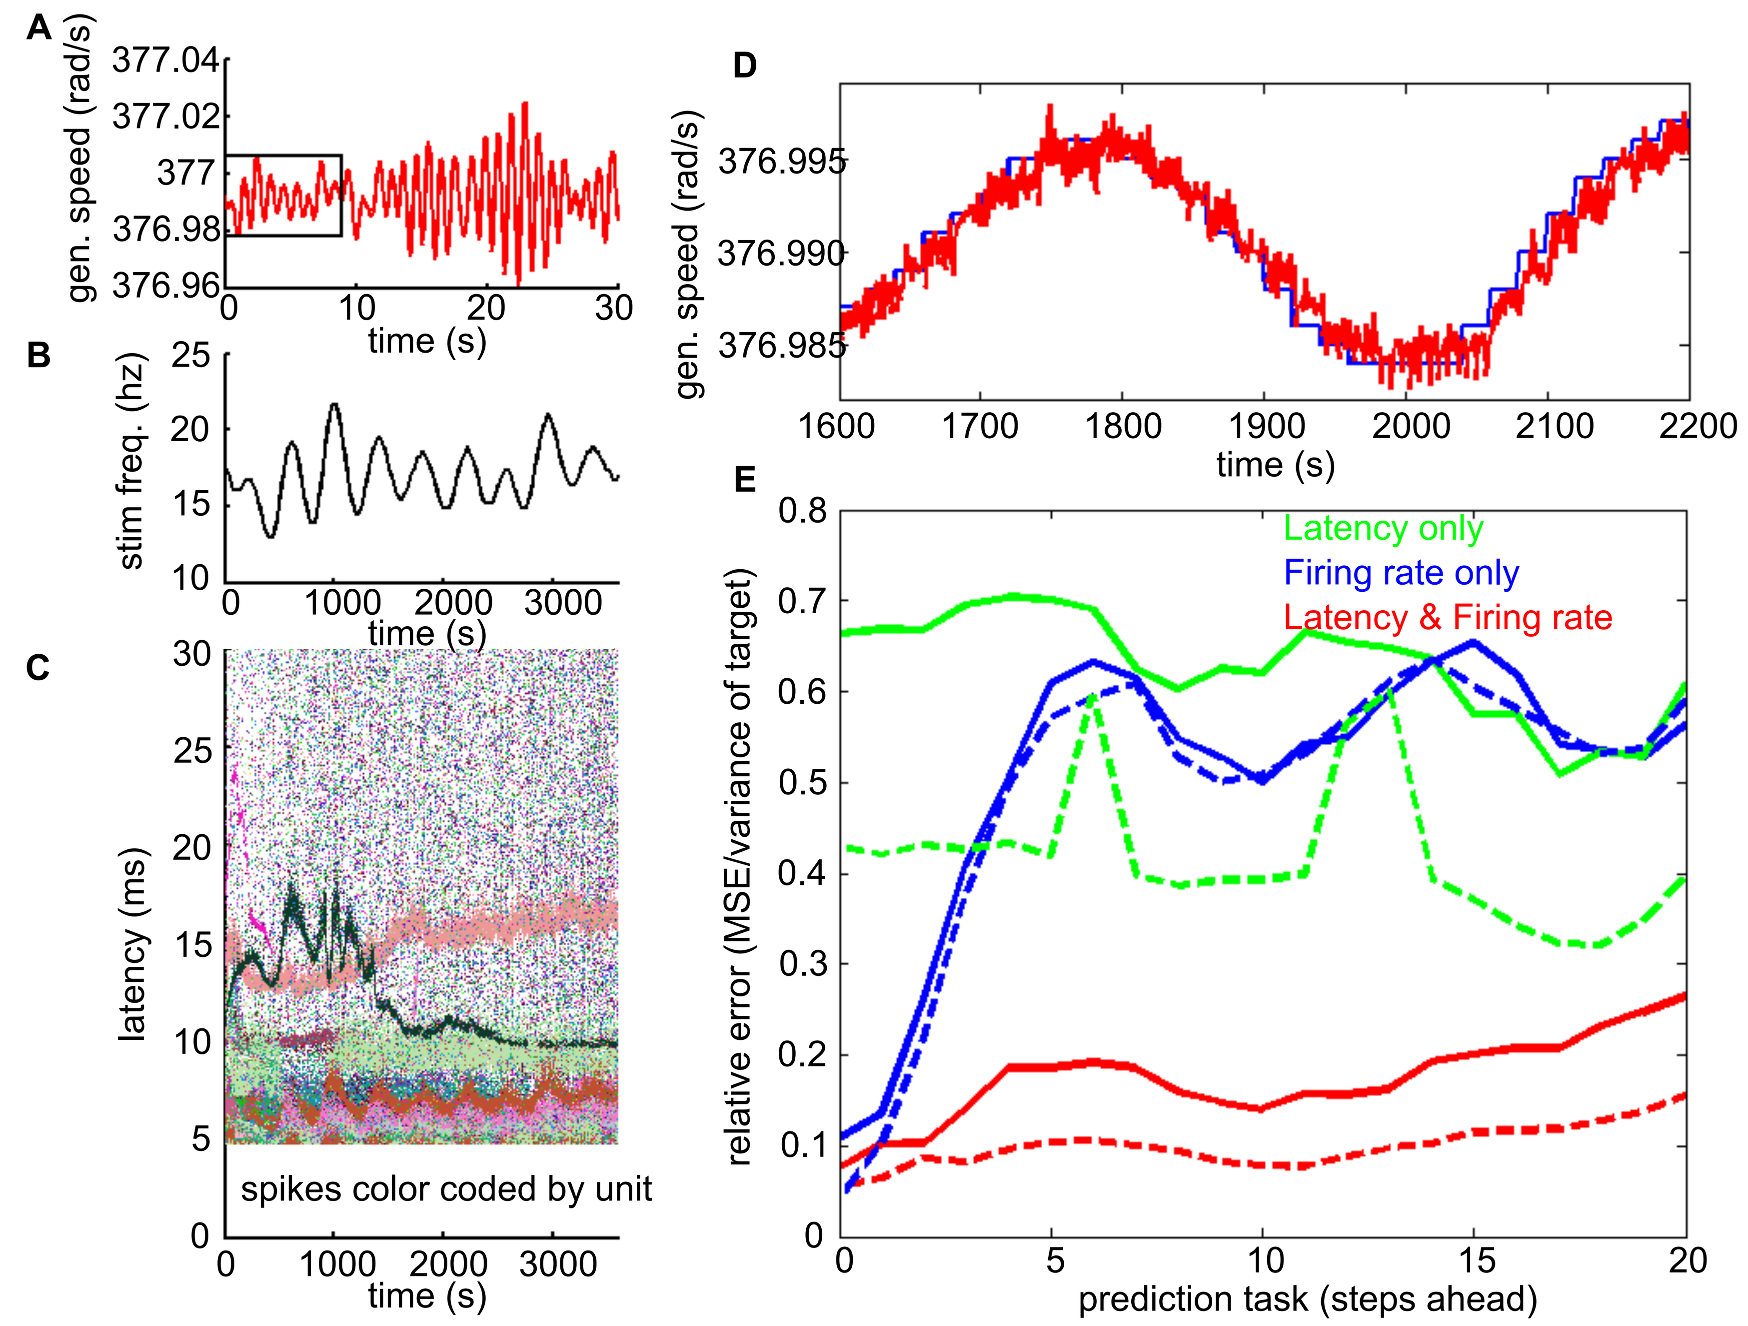

The rate-latency transform can be used to project low dimensional inputs into a high dimensional feature space. Latencies measured from live neurons growing in an MEA were used as part of a prediction task. A) generator speed was encoded by first slowing the signal down by a factor of 400, and then mapping the signal onto a stimulation rate. B) A pulsatile stimulation train was then built using this modulation scheme. C) Latencies of action potentials were recorded in vitro responded to this modulated stimulation. D) Latencies were then filtered, and mapped onto generator speed using linear regression- blue is the target, red is the estimate based on regression. E) This task was performed at different 'difficulties' by predicting multiple time steps ahead (horizontal axis), and using different feature sets (latencies only, firing rate only, or both sets). It was found that while firing rate alone was better at prediction than latency alone, the combination of the two substantially decreased error when predicting several time steps ahead, indicating that latency contains information that is complementary to firing rate for this prediction task.

The reservoir computing approaches shown here are preliminary work that shows what sort of information is passed, and what sorts of transforms are being applied by these dynamics. The more interesting problem is the next one, which is how living neural systems use these dynamics. At the very least, we have presented a suggestion for how delays could be used in artificial systems. The more useful conclusion is that we have shown that there are ways these dynamics can be used, suggesting that the biology could be exploiting an already present resource. We doubt that the methods we have shown here are representative of exactly how biological systems use these dynamics, but they are a proof of concept that these dynamics could be a vital part of neural computation.

CLEMSON UNIVERSITY
LABORATORY FOR NEUROENGINEERING
MISSOURI S&T