Alex Lee
al284
5-16-2008

Attempts to Further Understand the STDP network

Izhikevich's proposition that a massively connected neural network implementing a synaptic weight changing algorithm dependent on the timing of the signal was highly interesting in terms of what such a network could be possible of doing. He proposes that Spike Timing Dependent Plasticity (STDP) could create synfire chains which can be used for storage of information.

I attempted to further investigate the details of the network, and how it could be used to learn. Through modification of Izhikevich's neural network, several findings surfaced that may aid future research in this area.

Setup

First of all, the setup involves a network of 800 excitatory and 200 inhibitory neurons, massively interconnected, with STDP as the only rule modifying the synaptic weights. Within the 800 excitatory neurons, random sets of 10 neurons were chosen as “words”, five sets of 10 neurons were presented to the network at predetermined but random intervals within 20 msec, the longest single axon length in the network. This formed the “sentence”. The sentence was presented to the network with noise (manifested as random inputs every millisecond) through direct stimulation. The stimulation caused corresponding neurons to fire ~3 msec later (See *Timing Isssues*).

The network was thus trained for a simulated 10 minutes to get meaningful results without running full 24-hour simulations. After the simulated 10 minutes of noisy presentation, the network was allowed to settle by removing all inputs, then presented with the training sentence again, this time without noise. This was so the precise network engaged by the sentence could be seen. This network was compared with captured statistics during training regarding neuron-firing to come up with a set of “Predictor Neurons”, which are used later to determine if the network acknowledges the presence of the sentence (Figure 1).

Afterward, the network was checked to see if the predictor neurons acknowledged presentation of the sentence or not. First it was presented with just the first word of the sentence (Figure 2). Then it was presented with the full sentence (Figure 3). Each time, the activations of the predictor neurons were captured.

Results

13 Predictor neurons were selected for this documented instance. On presentation of the single word, 11 predictor neurons fired, starting from 18 msec since onset of presentation to 37 msec. On presentation of the entire sentence, 5 predictor neurons fired. On first glance, this result seems contradictory. It is.

However, closer inspection of the predictor neuron behavior showed that although the single word presentation scored more predictor neuron hits, all the hits were within short temporal range (15-34 msec) of the firing of the first neurons. The most important predictor neuron in the set of 13, the last one, temporally located after presentation of the entirety of the sentence, did not fire. However, an inspection of the predictor neurons in the full sentence presentation revealed that 12 of the 13 predictor neurons fired, but many were off the timing by 1-4 msec. Importantly, the last predictor neuron *did* fire in the full sentence condition, albeit 4 msec earlier than expected. But because of the methods used to determine whether or not the neuron fired depended on millisecond accuracy, many of the neurons were missed.

Conclusions

The difference in behavior of Predictor Neurons in the first word presentation and full sentence presentation points to the ability for an STDP network to differentiate between having heard only one word to having heard the entire sentence. It also points to the possibility for the network to extract high-level information such as existence of a sequence through examination of a single neuron. It implies that the last predictor neuron was one in a line of synfire chains that was propagated by both internal and stimulus neurons.

Timing Issues

Regarding the timing of the last Predictor Neuron, there seems to be an issue with the time sensitivity of the network. Because of the nature of the propagation as implemented in the code, the neurons do not necessarily fire right after sufficient input. Because the stabilization steps in the code were operated in half-millisecond steps, the time between input and output could vary from 2-5 msec. Furthermore, this problem worsens as a synfire chain propagates, because the delay in the first leg of the chain delays the rest of the chain and millisecond rounding moves the timing of neurons firing around, as seen with the last predictor neuron. The solution to this problem is most likely going to involve a cleaner reimplementation of the STDP model. Izhikevich's solution is fine for showing such phenomena as synfire chains existing under STDP, but for more precision as demanded by this task, a more precise implementation of neuron firing is needed.

Also because of these timing issues, additional Predictor Neurons which may have been later elements of the synfire chain could have been missed as their activations were spread amongst various timings and failed to break statistical threshold used to determine significance (Figure 5).

Other Findings

During the simulations, I tried many things to see how the network would respond. Some interesting things cropped up in otherwise failed attempts. First of all, the in-between timing for words seem to require them to be within the maximum axonal length. Separation beyond that leads to such a decrease in likelihood that the next word will hit the continuation of a previous synfire chain that it is almost impossible to find normally, despite having 10 neurons in each word. The current implementation guarantees a flat distribution of connections within 20 msec.

One extremely fascinating mistake was, on one trial after the network matured for 10 minutes with noise (no training), and presented with words from the sentence for the first time, several randomly hit on what seemed to be a self-propagating synfire chain (Figure 4). This behavior isn't often seen, but it just shows that such behavior is possible under STDP. This may possibly become an implementation of a clock or a timing mechanism.

Discussion

Because of the limited time and resources available to this investigation, several things that I wanted to try could not be tried. For example, will there be Predictor Neurons capable of differentiating between two sentences if trained together? Will changing the onset timing affect Predictor Neurons' ability to tell if the sentence has been presented? Can use of self-propagating synfire chains lead to time-independence in an otherwise highly timing dependent network? It is my hope that some of the findings presented here will aid further research in STDP networks.

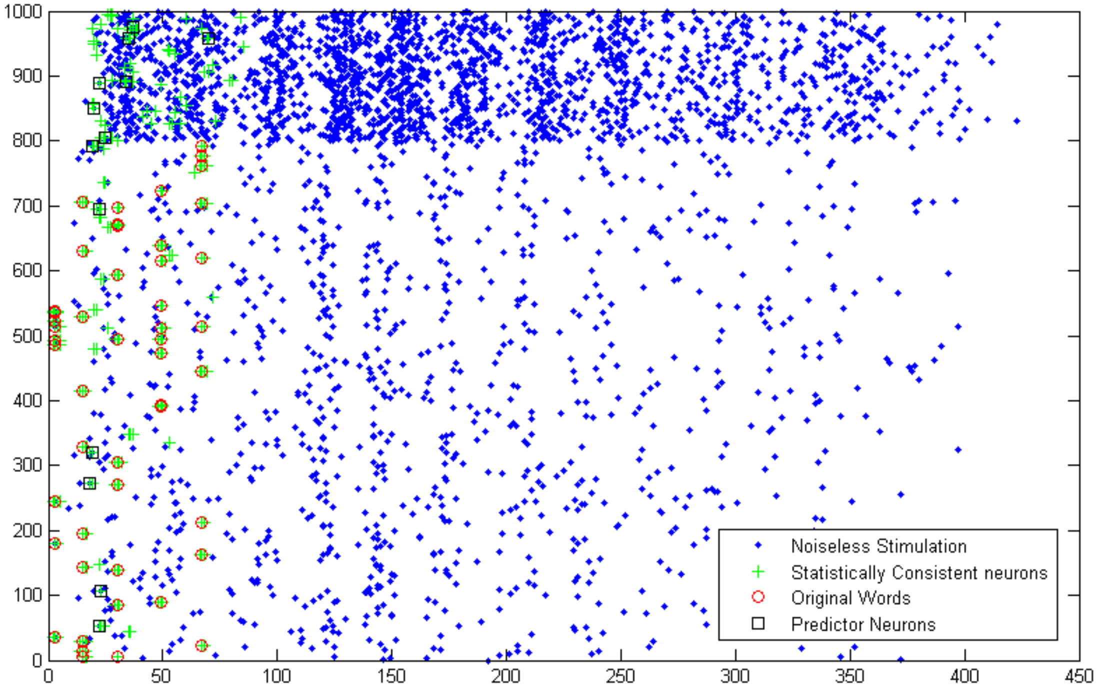

Figure 1. The first 450ms of the onset of presentation of the sentence without noise. The neurons firing are in blue; the stimulated word neurons are in red circles; the statistically high-firing neurons during training are marked in green crosses; the Predictor Neurons are black boxes. Although most Predictor neurons are in the first 40ms since onset, one important one occurs after all the word inputs.

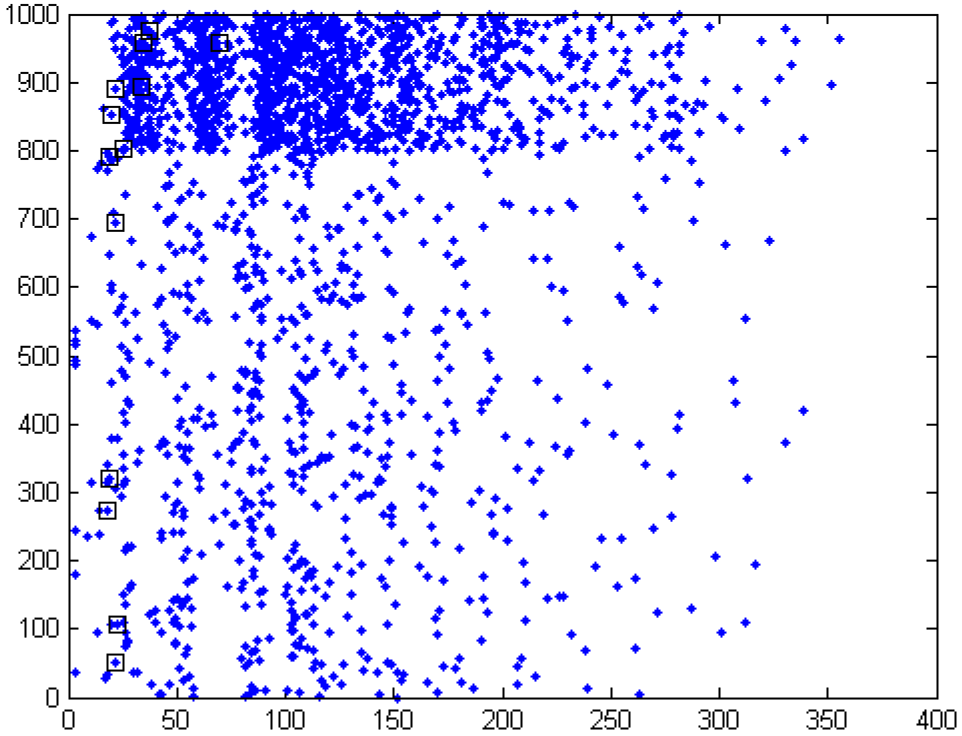

Figure 2. Presentation of the first word of the sentence. The last Predictor Neuron is not active within recent temporal range.

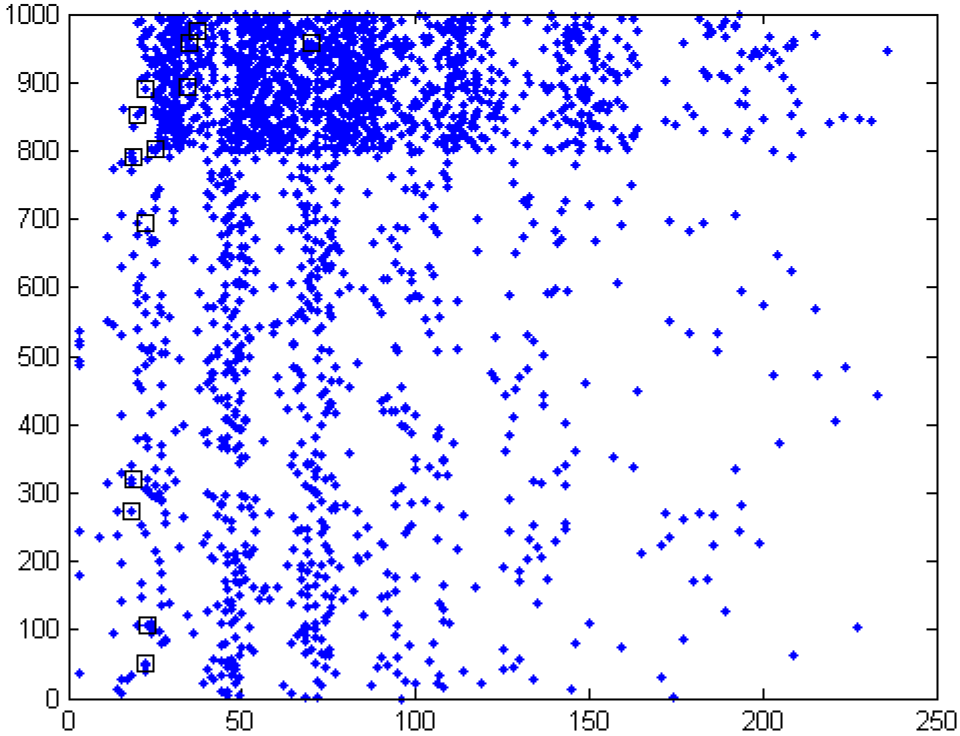

Figure 3. Presentation of the entire sentence. The last Predictor Neuron is active 4ms before the predicted activation.

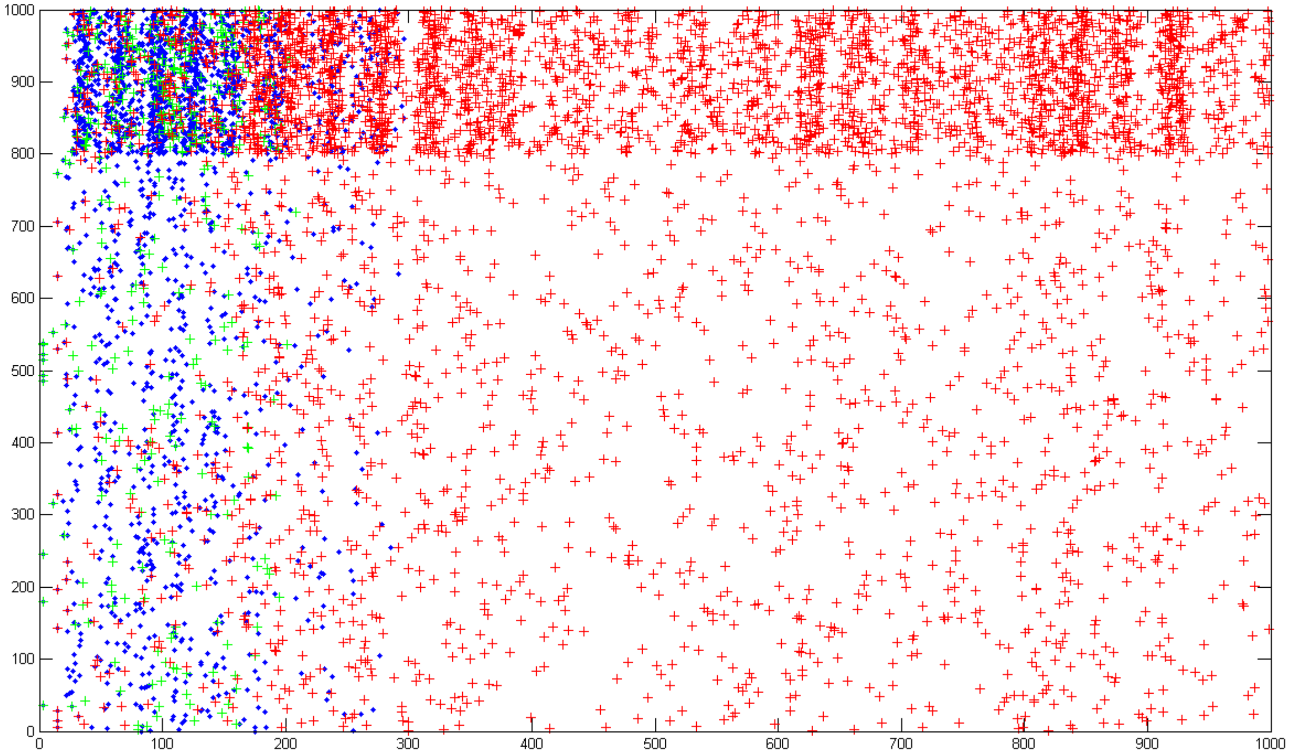

Figure 4. Existence of infinite-length sequences. Blue is full sentence; green is word 1; red is word 3.

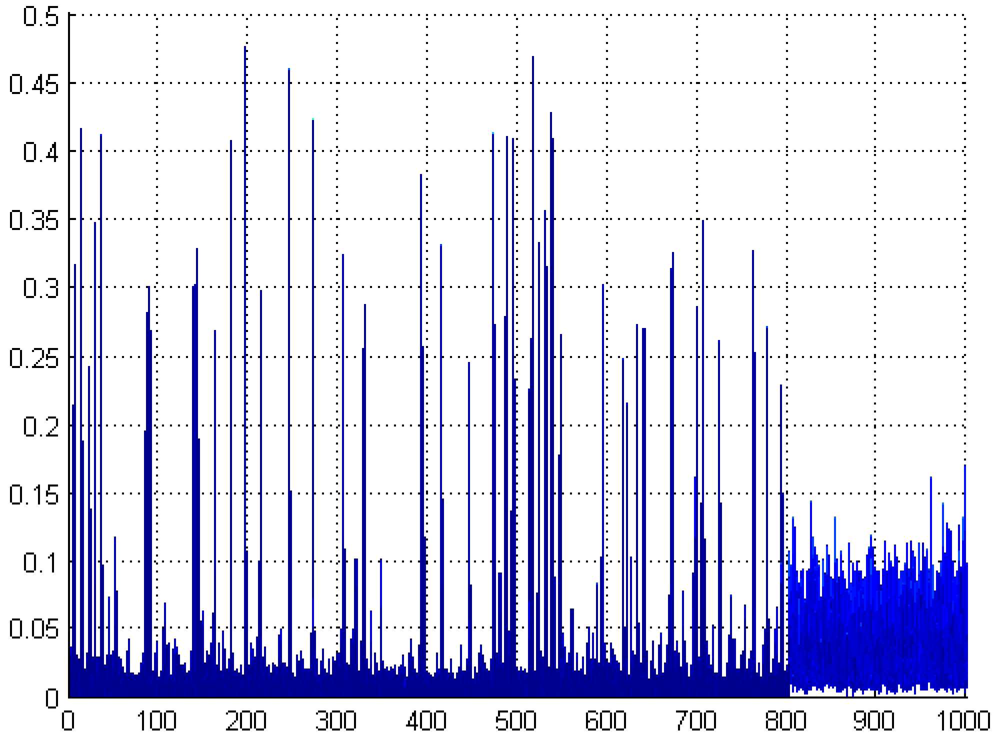

Figure 5. Statistical accumulation of neuron firing rates. The extremely high peaks in the excitatory neurons are the presentations of words. The cutoff for excitatory neurons was 0.06. The cutoff for inhibitory neurons was 0.1.