The predictive neuron, how active dendrites enable spatiotemporal computation in the neocortex

Subutai Ahmad*, Jeff Hawkins
sahmad@numenta.com, jhawkins@numenta.com

SUMMARY

Sequence learning is ubiquitous in cortex

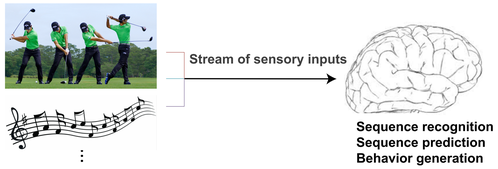

What are the neural mechanisms for sequence learning?

In this poster we show that:

1) Neurons can recognize temporal context and make predictions using active dendrites
2) Network of neurons with active dendrites form a powerful sequence memory
3) The network can learn both external temporal sequences and sensorimotor sequences
4) The network demonstrates excellent performance on real world problems

HTM = Hierarchical Temporal Memory

ACTIVE DENDRITES ENABLE A POWERFUL SEQUENCE MEMORY

The Predictive Neuron

Pyramidal neuron

5K to 30K synapses
10% proximal, 90% distal

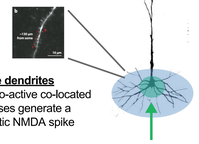

Active dendrites
8-15 co-active co-located synapses generate a dendritic NMDA spike

Sustained depolarization of soma but doesn't lead to action potential.

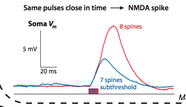

HTM neuron

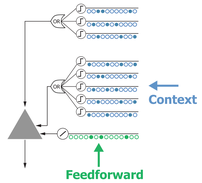

Neuron model

Proximal synapses define the feedforward receptive field

Distal dendritic segments detect context and put the cell into a depolarized or "predictive" state. Depolarized cells fire sooner and inhibit nearby neurons.

Major, Larkum and Schiller 2013

Network Model for Sequence Learning

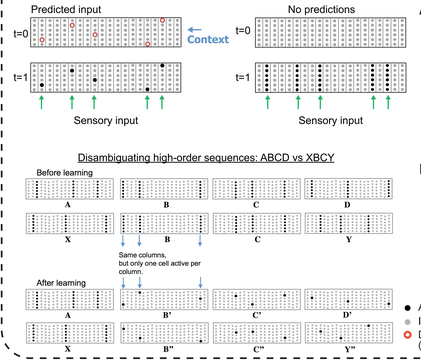

Activation rules
- Cells within a minicolumn learn same receptive field
- Distal dendritic segment learn contextual patterns
- Active segments act as predictions and bias cells

- With sensory input, these cells fire first and inhibit others. This creates a very sparse code representing this input in this particular context.

- Without predictions, all cells in minicolumn become active

Branch specific Hebbian-style plasticity
- If predicted cell becomes active, reinforce active segment
- If predicted cell does not become active, apply decay
- If no cell in active minicolumn was predicted, reinforce most active segment.

FAULT TOLERANCE

Robustness to damage
- Input stream contained high-order sequences mixed with random elements.
- The maximum prediction accuracy for this data stream is 50%
- After the network reached stable performance we inactivated a random subset of neurons.
- Network can handle significant faults
- Network rapidly recovers performance

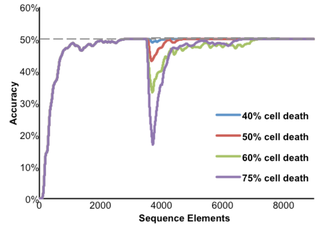

EXPERIMENTALLY TESTABLE HYPOTHESES

Impact of NMDA spikes
- Dendritic NMDA spikes cause cells to fire faster than they would otherwise.
- Fast local inhibitory networks (e.g. minicolumns) inhibit cells that don't fire early.
- Sparser activations during a predictable sensory stream.
- For predictable natural stimuli, dendritic spikes will be more frequent than APs.
  (Vinje & Gallant, 2002; Smith et al, 2013; Wilmes et al, 2016; Moore et al, 2017)

Branch specific plasticity
- Strong LTP in dendritic branch when NMDA spike followed by back action potential (bAP).
- Weak LTP (without NMDA spike) if synapse cluster becomes active followed by a bAP.
- Weak LTD when an NMDA spike is not followed by an action potential/bAP.
  (Holthoff et al, 2004; Losonczy et al, 2008; Yang et al, 2014; Cichon & Gang, 2015)

Correlation structure
- Low pair-wise correlations between cells but significant high-order correlations.
- High order assembly correlated with specific point in a predictable sequence.
- Unanticipated inputs leads to a burst of activity, correlated within minicolumns.
- Activity during predicted inputs will be a subset of activity during unpredicted inputs.
- Neighboring mini-columns will be uncorrelated.
  (Ecker et al, 2010; Smith & Häusser, 2010; Schneidman et al, 2006; Miller et al, 2014; Homann et al, 2017)

RESULTS

Real World Sequence Prediction Tasks

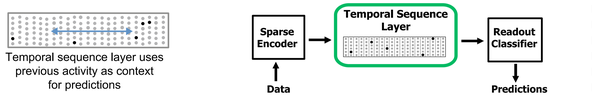

Excellent performance
- HTM sequence memory on par with best deep learning system
- Network responds far better than those systems when statistics change

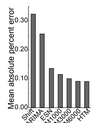

HTM adapts quickly to changes
- HTM adapts more quickly to changes in statistics due to its continuous unsupervised Hebbian learning rule

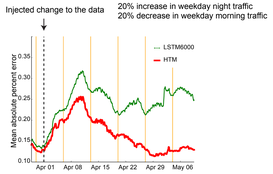

Task: predict taxi passenger count in NYC

(Cui et al., 2016)

Data sources: http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml

Sensorimotor Sequences

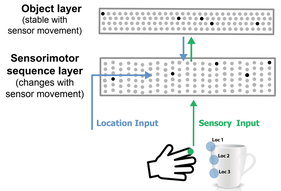

Object layer
- Pools over sparse activity from sensorimotor layer
- Stable while sensing an object

Sensorimotor layer
- Sequence memory layer
- Uses motion derived location signal as context for predictions
- Sparsely active cells represent sensory features at specific locations

YCB Object Benchmark

Yale-CMU-Berkeley (YCB) Object Benchmark
- 80 objects designed for robotics grasping tasks
- Includes high-resolution 3D CAD files
  (Calli et al, 2017)

Sensorimotor simulations
- Created a virtual hand using Unity game engine
- Curvature based sensor on each fingertip
- Used sensorimotor sequence and object layers
- 98.7% recall accuracy (77/78 uniquely classified)
- Convergence time depends on object, sequence of sensations

Untangling Sequences

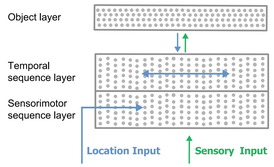

There are two fundamental types of sequences
- Sensory inputs can change due to 1) external factors, or 2) our own behavior

- These sequences are often mixed and must be untangled

- Network containing temporal and location contextual inputs automatically separates them

Inferring combined sensorimotor and temporal sequence streams

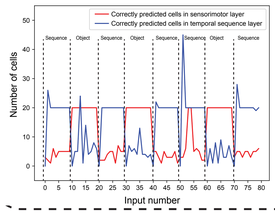

Simulation with mixed random sequences
- Each sequence is either temporal or sensorimotor (50 temporal sequences and 50 sensorimotor)

- Same sensory features for both sequence types

- Location signal is random for temporal sequences

REFERENCES

Hawkins, J., Ahmad, S., & Cui, Y. (2017) A Theory of How Columns in the Neocortex Enable Learning the Structure of the World. Front. Neural Circuits, 11, 81.

Hawkins, J. & Ahmad, S. (2016) Why Neurons Have Thousands of Synapses, a Theory of Sequence Memory in Neocortex. Front. Neural Circuits, 10, 1–13.

Cui, Y., Ahmad, S., & Hawkins, J. (2016) Continuous online sequence learning with an unsupervised neural network model. Neural Comput., 28, 2474–2504.

Ahmad, S. & Hawkins, J. (2017) Untangling Sequences: Behavior vs. External Causes. bioRxiv, 190678.

All code available as open source: https://github.com/numenta