A theory of sequence memory in the neocortex

Jeff Hawkins, Subutai Ahmad, Yuwei Cui and Chetan Surpur
jhawkins@numenta.com, sahmad@numenta.com, ycui@numenta.com

Numenta

1. Sequence learning is ubiquitous in cortex

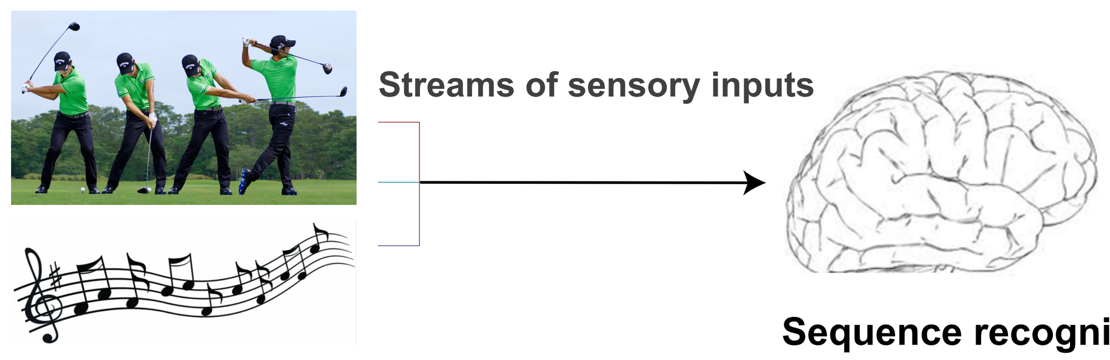

Sequence recognition
Sequence prediction
Behavior generation

What is neural mechanism for sequence learning?

HTM Sequence Memory:

1. Neurons learn to recognize hundreds of patterns using active dendrites.
2. Recognition of patterns act as predictions by depolarizing the cell without generating an immediate action potential.
3. A network of neurons with active dendrites forms a powerful sequence memory

2. HTM neuron models active dendrites

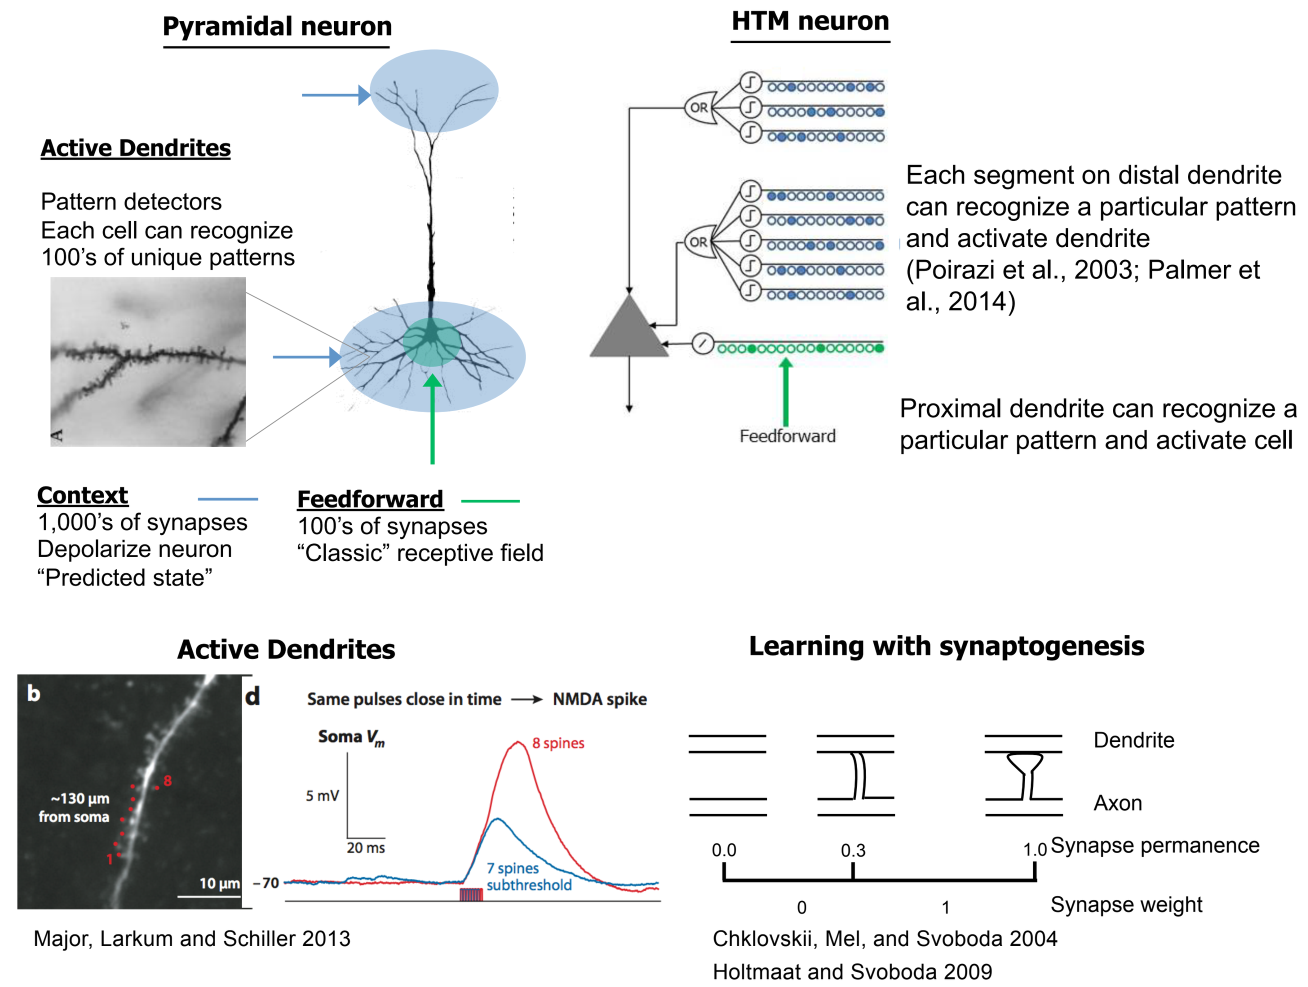

Small sets of synapses in close proximity act as independent pattern detectors.
Detection of a pattern causes an NMDA spike and depolarization at the soma.
Depolarization acts as prediction, causing cell to fire earlier.
Learning occurs by growing new synapses via Hebbian learning rule.

3. HTM network model for sequence learning

Two separate sparse representations

A sparse set of columns becomes active due to intercolumn inhibition

Only one cell active per column, due to intracolumn inhibition

Learning complex high-order sequences
ABCD vs XBCY

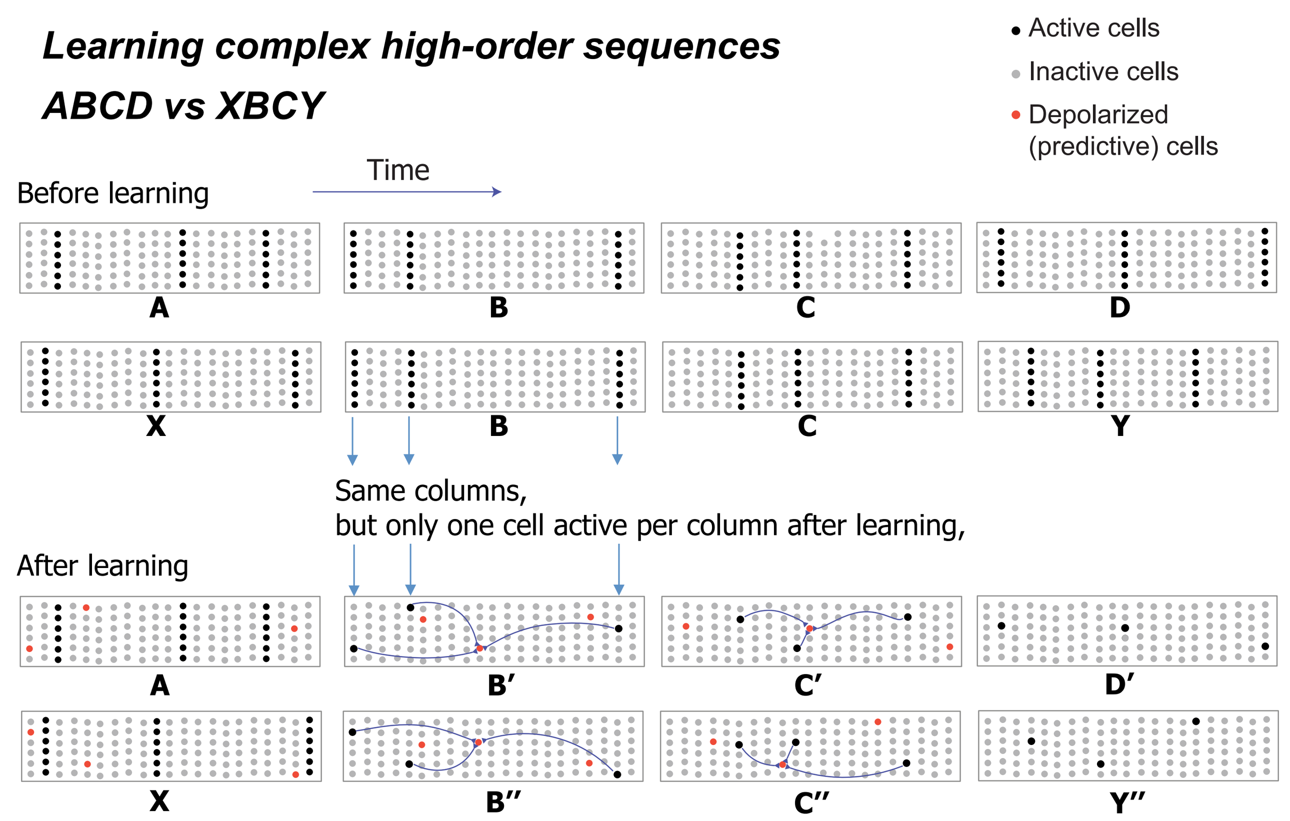

4. Works well on real-world problems

Task: predict taxi passenger count in NYC

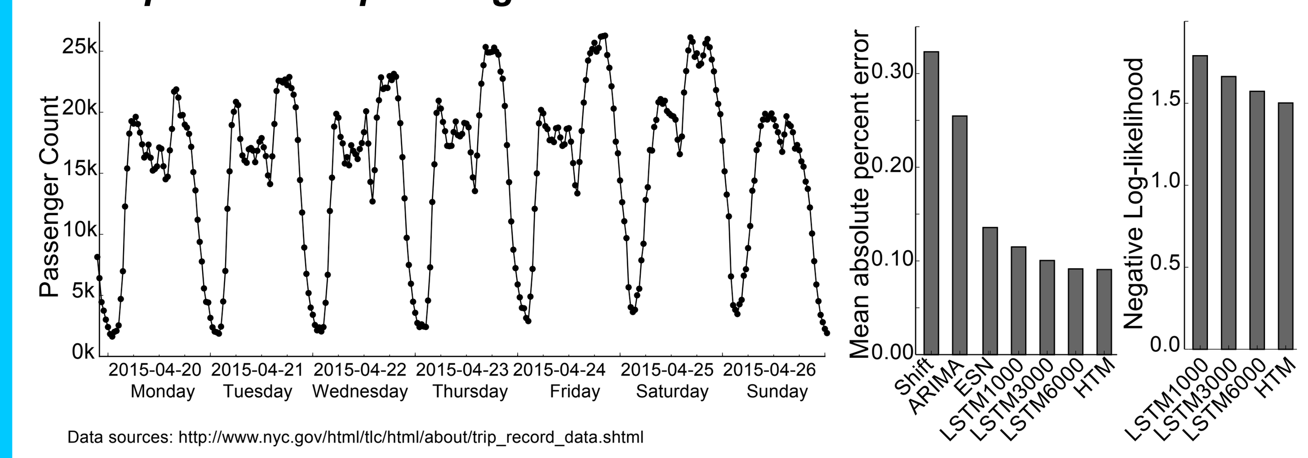

HTM has comparable performance to state-of-the-art algorithms

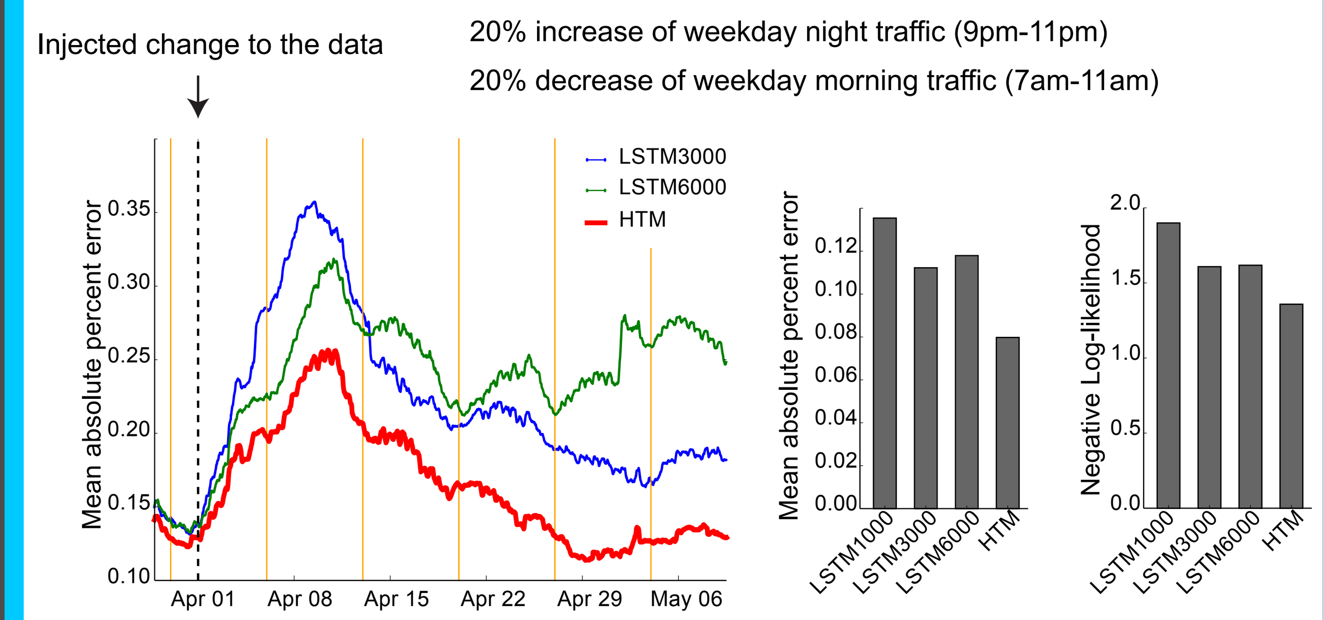

HTM quickly adapts to changes due to its ability to learn continuously

Apical inputs predict entire sequences

It has been speculated that feedback connections implement expectation or bias (Lamme et al., 1998). Our neuron model suggests a mechanism for top-down expectation in the cortex.

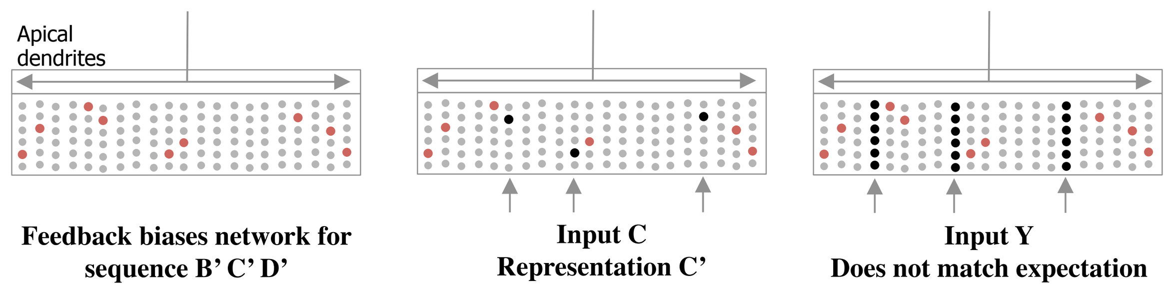

Feedback to the apical dendrites can predict multiple elements simultaneously.
New feedforward input will be intepreted as part of the predicted sequence.

Learning and activation rules

Activation rules:

Select the top 2% of columns with strongest inputs on proximal dendrite as active columns
If any cell in an active column is predicted, only the predicted cells fire
If no cell in an active column is predicted, all cells in the column fire

Unsupervised Hebbian-like learning rules:

If a depolarized cell becomes active subsequently, its active dendritic segment will be reinforced
If a depolarized cell does not become active, we apply a small decay to active segments of that cell
If no cell in an active column is predicted, the cell with the most activated segment gets reinforced

Further implementation detail can be found at https://github.com/numenta/nupic

Continuous learning from sequence streams

Task: sequence prediction with streams of high-order sequences

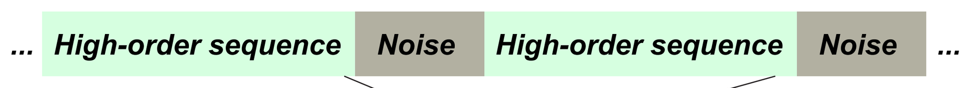

Check prediction at the end of each sequence

Example sequences:	*Modified sequences:*
G, I, H, E, C, D, A	G, I, H, E, C, D, F
B, I, H, E, C, D, F	B, I, H, E, C, D, A
A, J, H, I, F, D, E, B	A, J, H, I, F, D, E, G
C, J, H, I, F, D, E, G	C, J, H, I, F, D, E, B
......	......

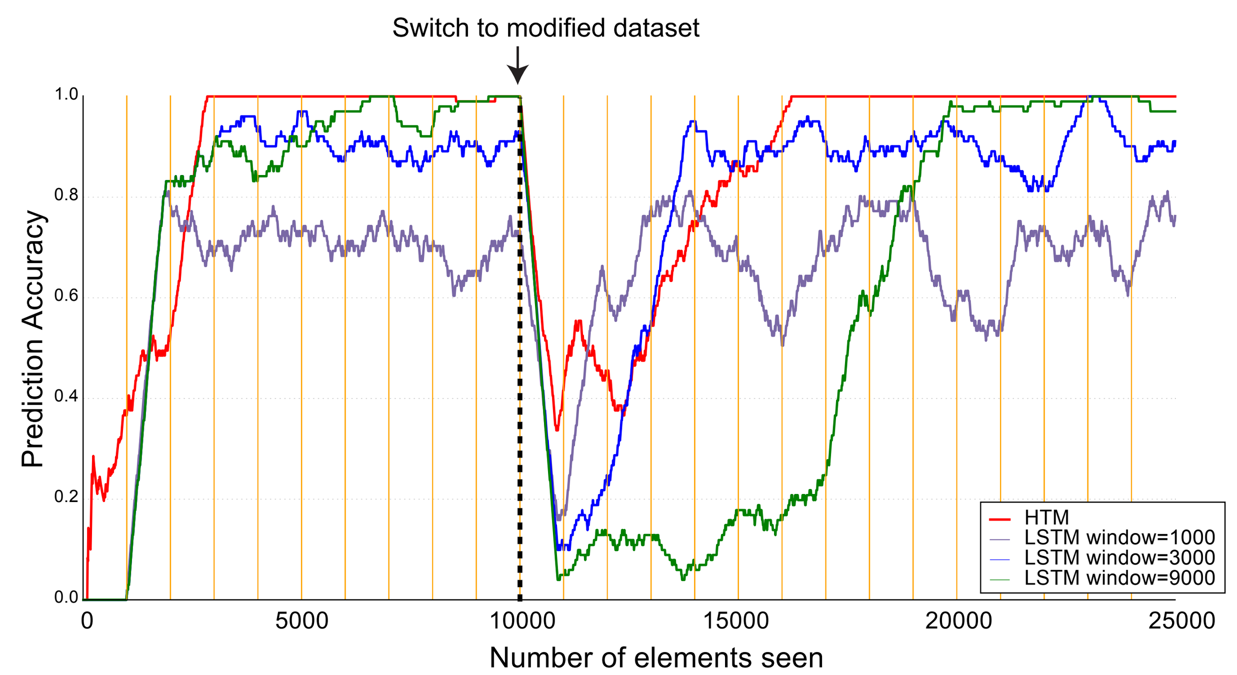

HTM learns continuously, no batch training required.
HTM is more robust and recovers more quickly.

High fault tolerance to neuron death

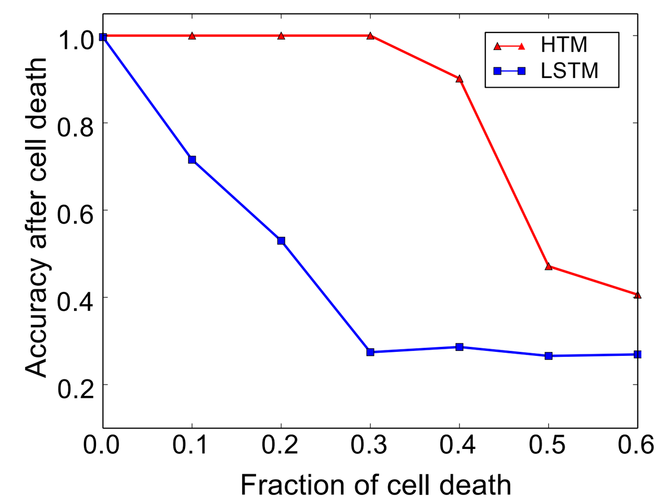

HTM is fault tolerant due to properties of sparse distributed representations (Kanerva 1988) and nonlinear dendritic properties of HTM neurons (Hawkins & Ahmad 2015).

In contrast, LSTM and most other artificial neural networks are sensitive to loss of neurons or synapses (Piuri 2001)

References

Poirazi P, Brannon T, Mel BW (2003) Neuron 37:989–999.
Chklovskii, D. B., Mel, B. W., and Svoboda, K. (2004). Nature 431, 782–8.
Hawkins, J., and Ahmad, S. (2015). arXiv:1511.00083 [q-bio.NC]
Holtmaat A, Svoboda K (2009) Nat Rev Neurosci 10:647–658.
Kanerva, P, Sparse Distributed Memory. The MIT Press, 1988
Lamme, V. A., Supèr, H., and Spekreijse, H. (1998). Curr. Opin. Neurobiol. 8, 529–35.
Losonczy, A., Makara, J. K., and Magee, J. C. (2008). Nature 452, 436–41.
Major, G., Larkum, M. E., and Schiller, J. (2013). Annu. Rev. Neurosci. 36, 1–24.
Piuri, V., J Parallel Distr Com., vol. 61, pp. 18-48.
Smith SL, Smith IT, Branco T, Häusser M (2013) Nature 503:115–120.
Vinje, W. E., and Gallant, J. L. (2002). J. Neurosci. 22, 2904–2915.

Acknowledgements

We thank Scott Purdy, Alex Lavin, Matthew Larkum, Ahmed El Hady for helpful discussions.

Our code is open source

We believe in open research and full transparency. Numenta's research and algorithm code is part of the open-source project Numenta Platform for Intelligent Computing (NuPIC). As a fast growing project, NuPIC currently has more than 4,000 followers and more than 1000 forks on Github.

5. Summary

HTM exhibits many desirable features for sequence learning:

- **Unsupervised learning**
- **Quickly adapts to changes in data**
- **Learns high-order structure in sequences**
- **Robust and fault tolerant**
- **Makes multiple simultaneous predictions**
- **Works well on real-world problems**
- **Accurate biological model**

6. Testable predictions

1) Sparser activations during a predictable sensory stream. (Vinje & Gallant 2002)
2) Unanticipated inputs lead to a burst of activity correlated vertically within mini-columns.
3) Neighboring mini-columns will not be correlated.
4) Predicted cells need fast inhibition to inhibit nearby cells within mini-column.
5) For predictable stimuli, dendritic NMDA spikes will be much more frequent than somatic action potentials. (Smith et al., 2013)
6) Strong LTP in distal dendrites requires bAP and NMDA spike (Losonczy et al., 2008)
7) Weak LTP (in the absence of NMDA spikes) in dendritic segments if a cluster of synapses become active followed by a bAP.
8) Localized weak LTD when an NMDA spike is not followed by a bAP.

February 2016