Robust dendritic computations with sparse distributed representations

Subutai Ahmad*, Max Schwarzer, Jeff Hawkins
sahmad@numenta.com, maxa.schwarzer@gmail.com, jhawkins@numenta.com

SUMMARY

Cortex encodes information with sparse distributed patterns. How accurately can neurons recognize sparse patterns?

We show:

1) Scaling laws for computing error probabilities.

2) High dimensional sparse patterns can be classified extremely reliably, even with large amounts of noise.

3) Active dendritic segments can reliably classify patterns using a tiny number of synapses.

4) The equations explain experimentally observed NMDA spike thresholds in active dendrites.

5) Behavior of Poirazi-Mel and HTM neuron models closely match theoretical predictions. Understanding the behavior can lead to dramatically improved accuracies.

MODELING DENDRITES

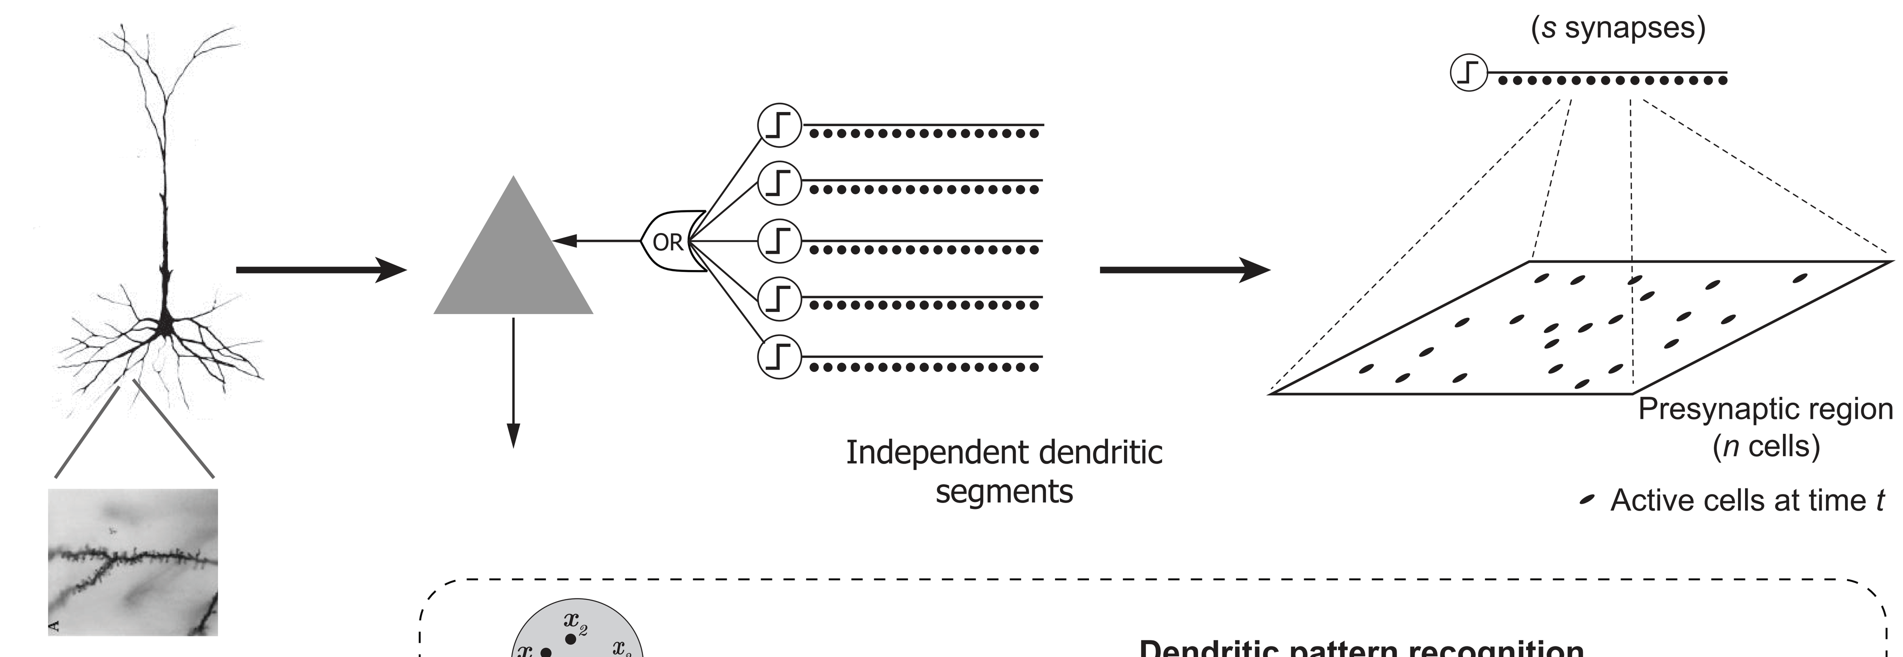

5K to 30K excitatory synapses
90% on active distal segments

Each segment independently recognizes patterns

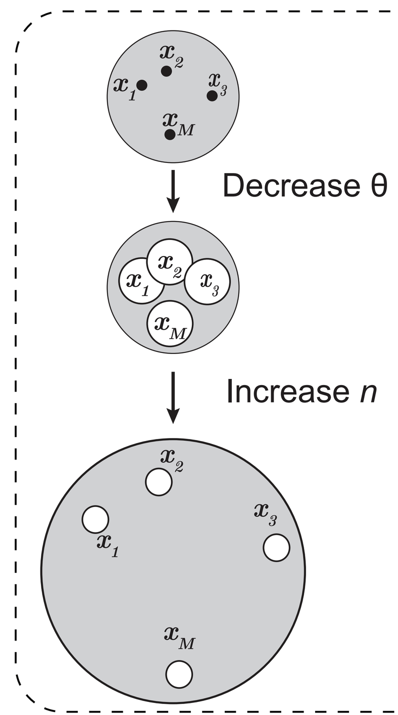

Dendritic pattern recognition

Activity in presynaptic region at time t represented by a binary vector $\boldsymbol{A}_t$ with n components and a_t active cells:

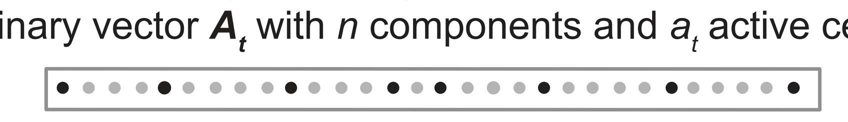

The synapses on a dendrite subsample this activity. Each segment has s synapses and represented by a binary vector $\boldsymbol{D}$ with n components and s "1" bits:

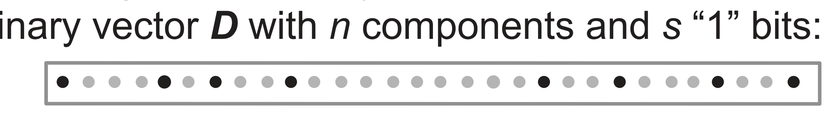

$\boldsymbol{D}$ recognizes the pattern in $\boldsymbol{A}_t$ if at least θ presynaptic inputs are on: $\mathbf{A}_t \cdot \mathbf{D} \geq \theta$

PROBABILITY OF FALSE POSITIVES

Probability of a random input matching a dendrite:

$$P(\mathbf{A}_t \cdot \mathbf{D} \geq \theta) = \frac{\sum_{b=\theta}^{s} |\,\Omega_{\mathbf{D}}(n, a_t, b)\,|}{\binom{n}{a_t}}$$

$|\Omega_{\mathbf{D}}(n,a,b)|$ = number of input vectors that exactly match b synapses:

$$|\Omega_{\mathbf{D}}(n,a,b)| = \binom{s}{b} \times \binom{n-s}{a-b}$$

Number of ways to select exactly b out of s synapses

Number of vectors that have a-b bits on and no overlap with dendrite.

The probability of error decreases dramatically with high dimensionality and input sparsity:

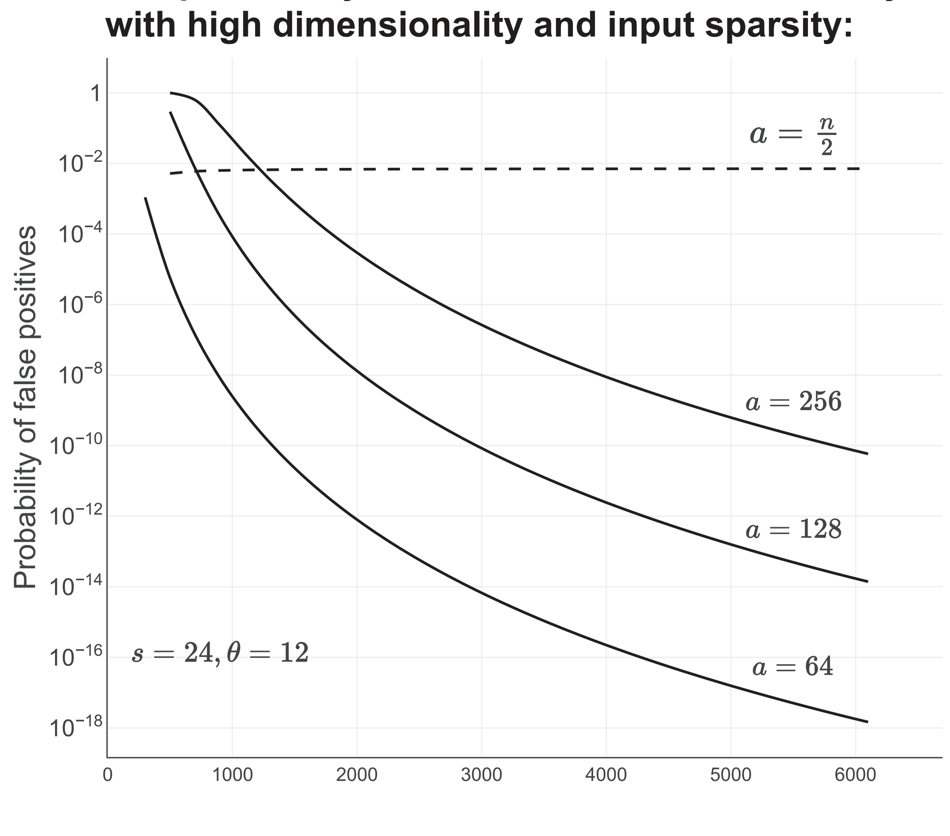

A tiny number of synapses, subsampling from a much larger pattern, is sufficient for robust recognition:

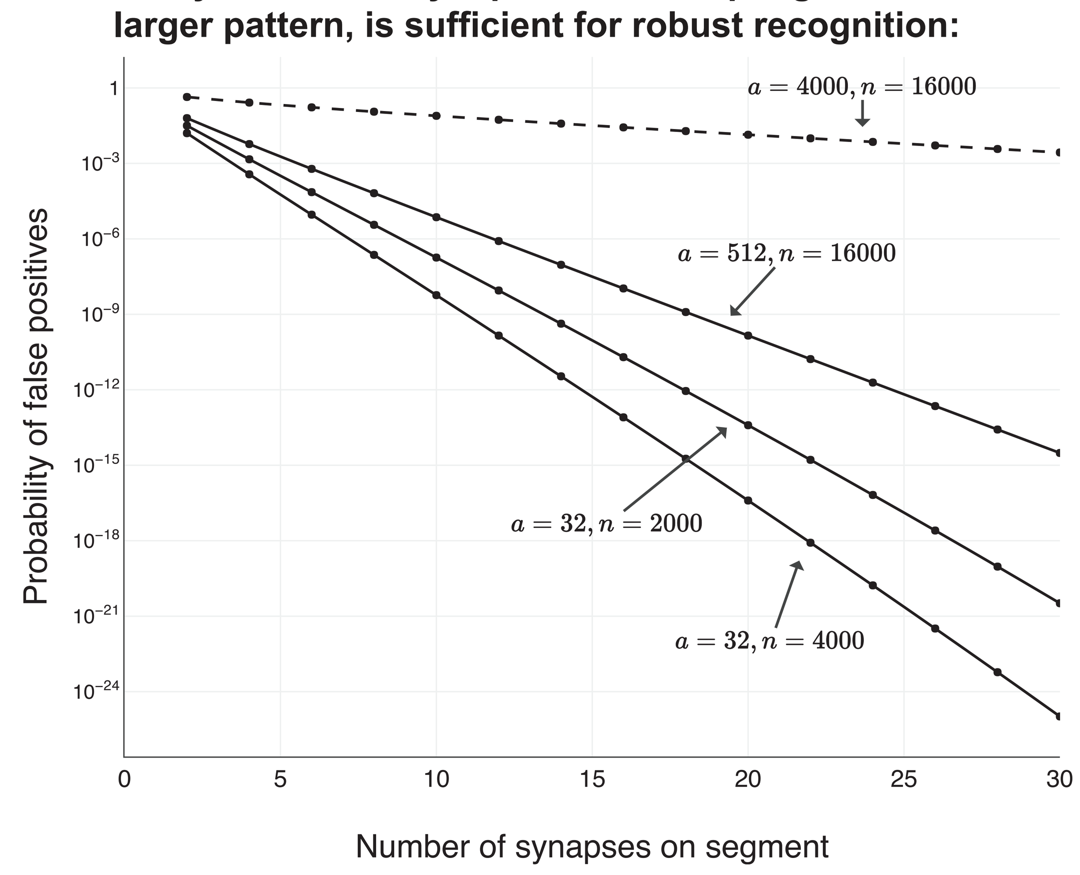

FALSE NEGATIVES

Probability of a corrupted pattern *not* matching a dendrite:

$$P(\mathbf{A}_t^* \cdot \mathbf{D} < \theta) = \frac{\sum_{b=s-\theta+1}^{s} |\Omega_D(a_t^*, v, b)|}{\binom{a_t^*}{v}}$$

A* is a corrupted pattern represented in D, with v bits missing.

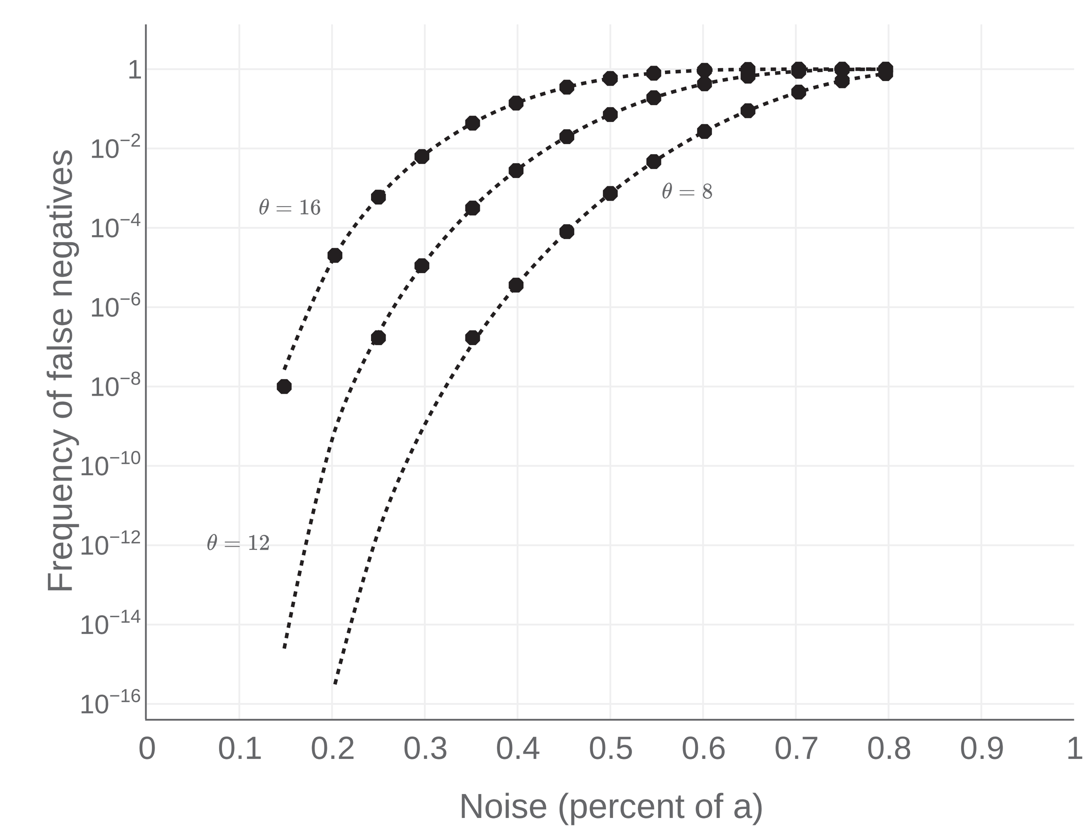

PREDICTS NMDA SPIKE THRESHOLDS

Active dendrites independently recognize patterns: as few as 8 co-active synapses can generate NMDA spike. Neural activity is unreliable and noisy, so how can such a tiny number of synapses reliably detect patterns?

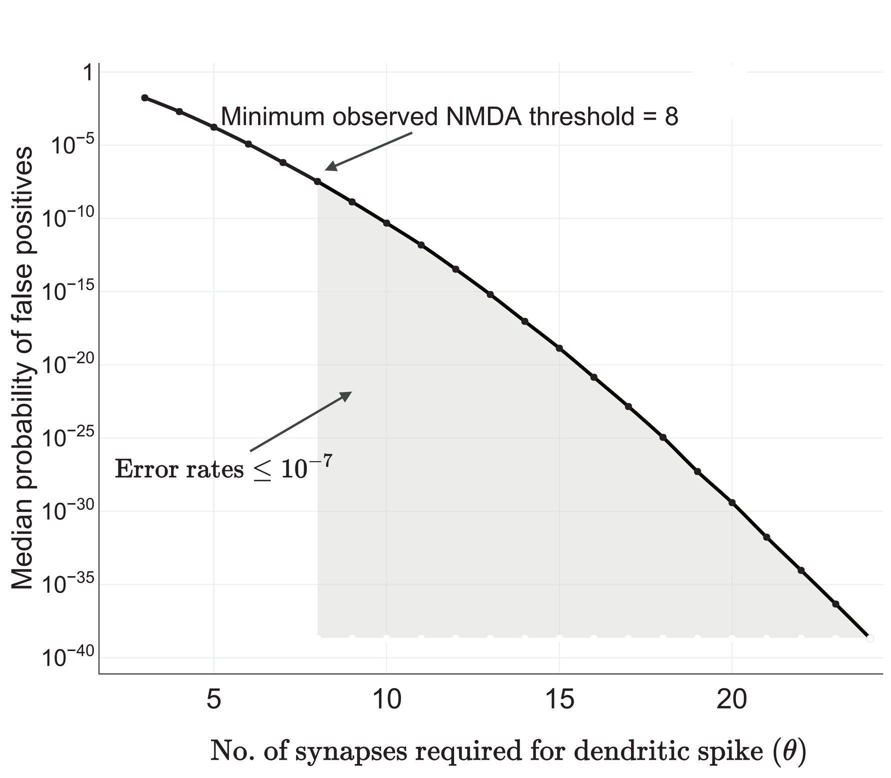

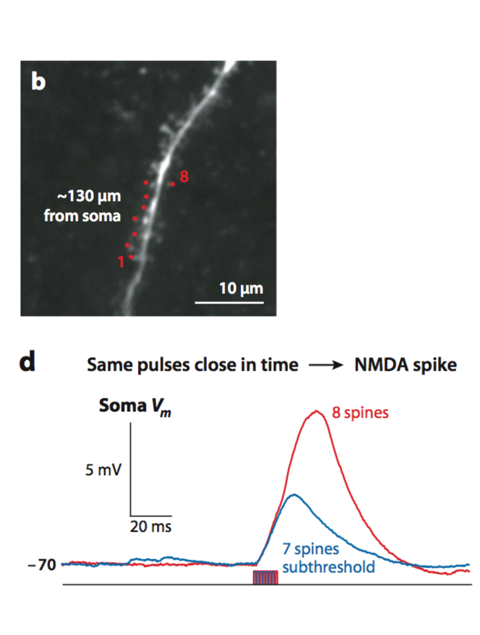

Major, Larkum and Schiller 2013

Our equations show that the small NMDA thresholds observed in biology can be explained by high dimensional sparse representations. NMDA threshold of 8 leads to median error of less than 1 in 10 million with 50% noise.

n:1,000 to 20,000, sparsity:0.5% to 3%, s:20 to 50

VASTLY IMPROVED CLASSIFICATION

Poirazi-Mel neuron with m independent dendritic segments.

Original paper tested populations of these neurons on binary classification tasks, using input dimensionality n=400, of which 40 components were on (10% sparsity).

We tested this neuron model by varying the dimensionality and sparsity of the inputs.

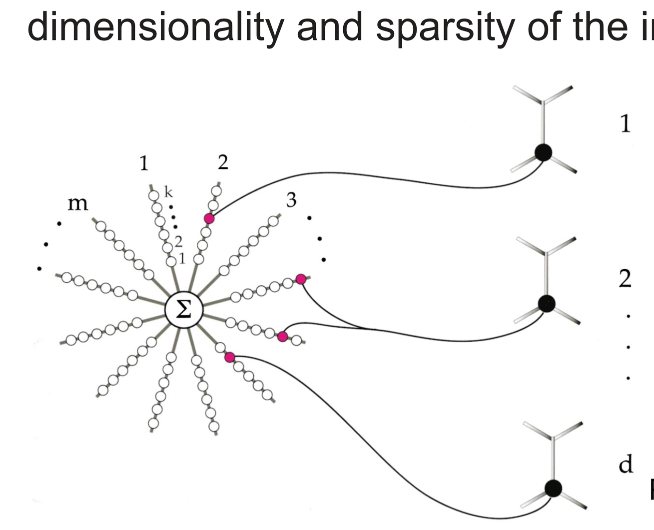

Poirazi and Mel, 2001

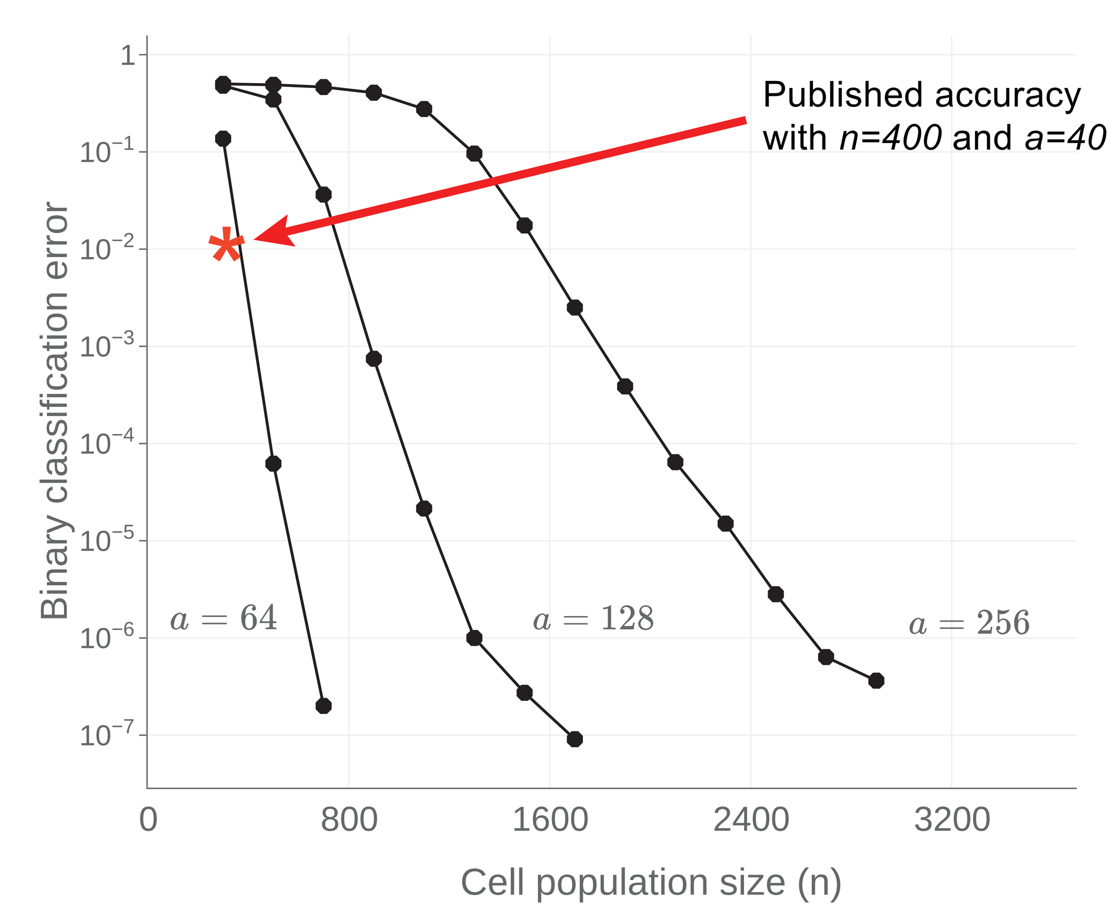

Large improvement with high dimensional sparse representations, even with constant number of synapses, as predicted by our theory.