HTM spatial pooler

Yuwei Cui
ycui@numenta.com

Numenta

February 2017

Background

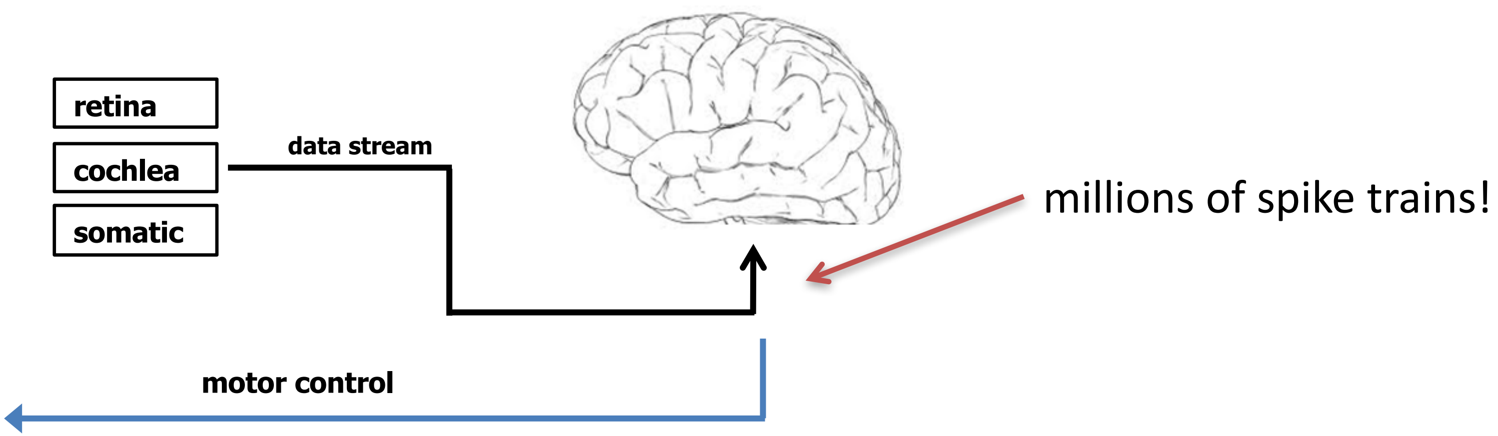

- How do individual neurons learn to respond to specific input patterns?
- How do populations of neurons represents input features?

Online sequence learning with HTM

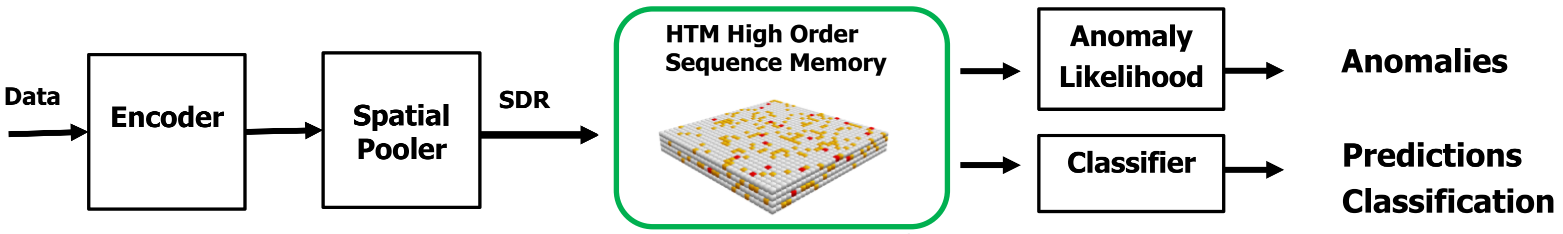

HTM Spatial Pooler

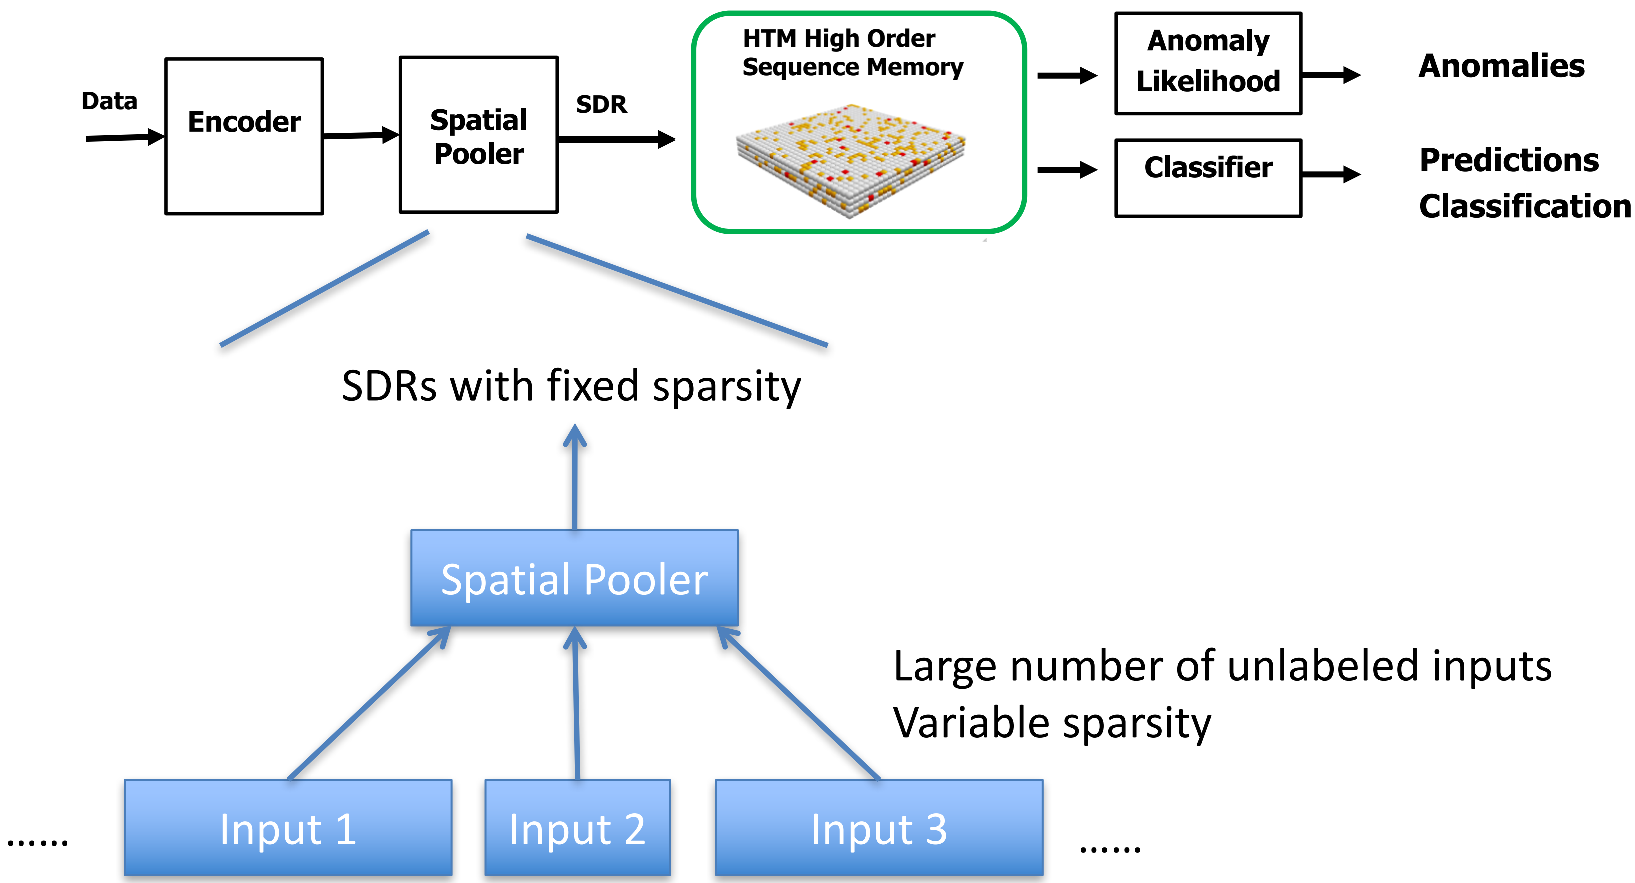

What properties do SP achieve?

Outline

- Background
- HTM spatial pooler
- Properties and metrics
- Simulation results

HTM spatial pooler

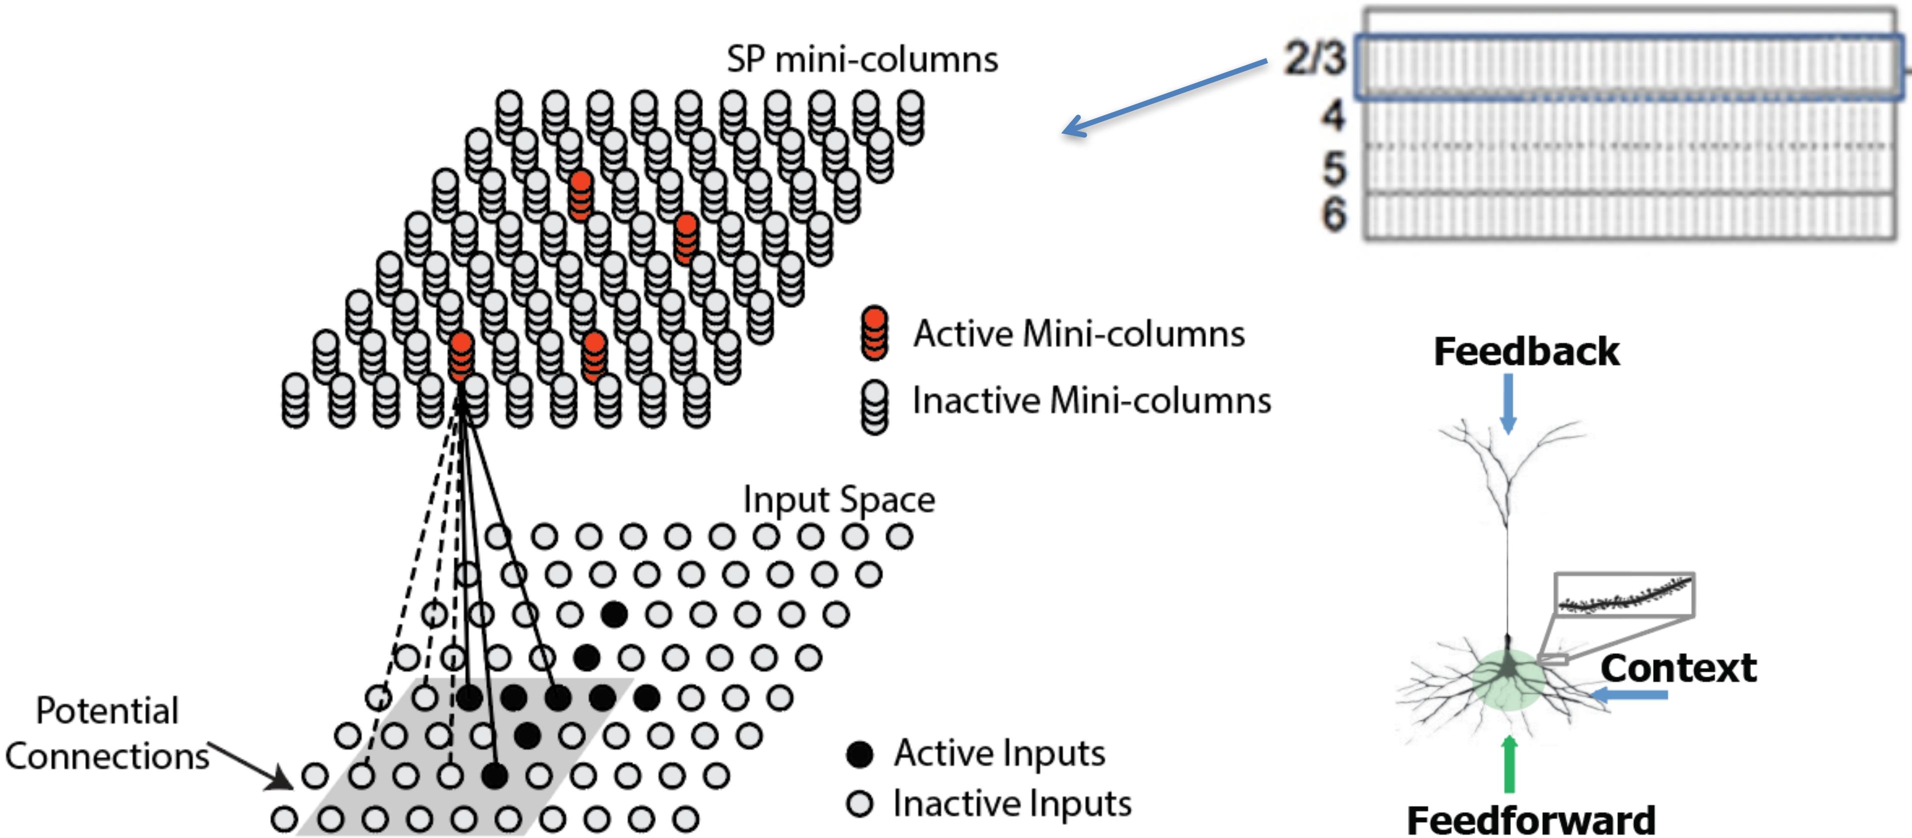

HTM spatial pooler – winner take all

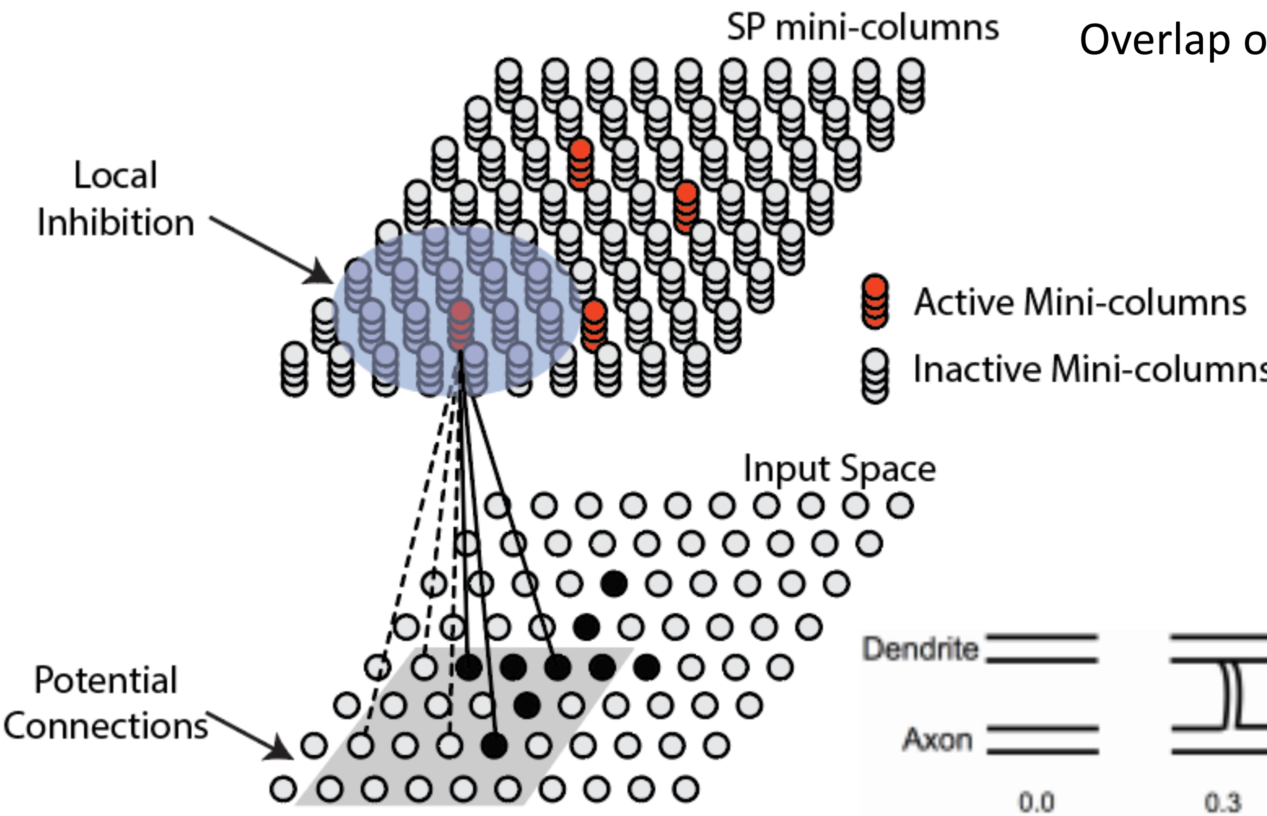

Overlap of the *i*th MC

*j*th input

$$o_i = b_i \sum_j W_{ij} z_j$$

Synaptic connection from *j*th input to *i*th MC

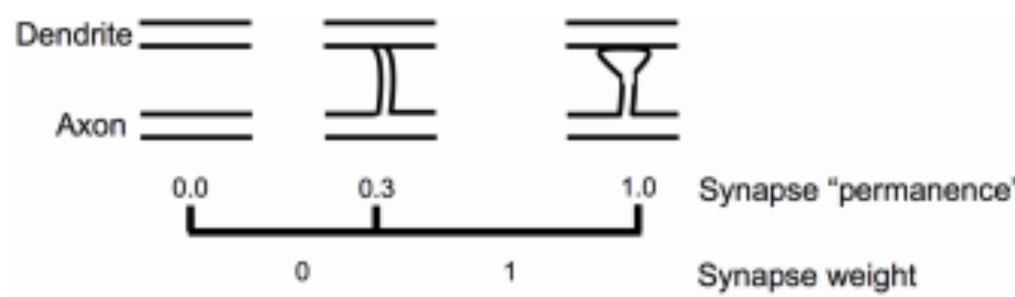

HTM spatial pooler – boosting

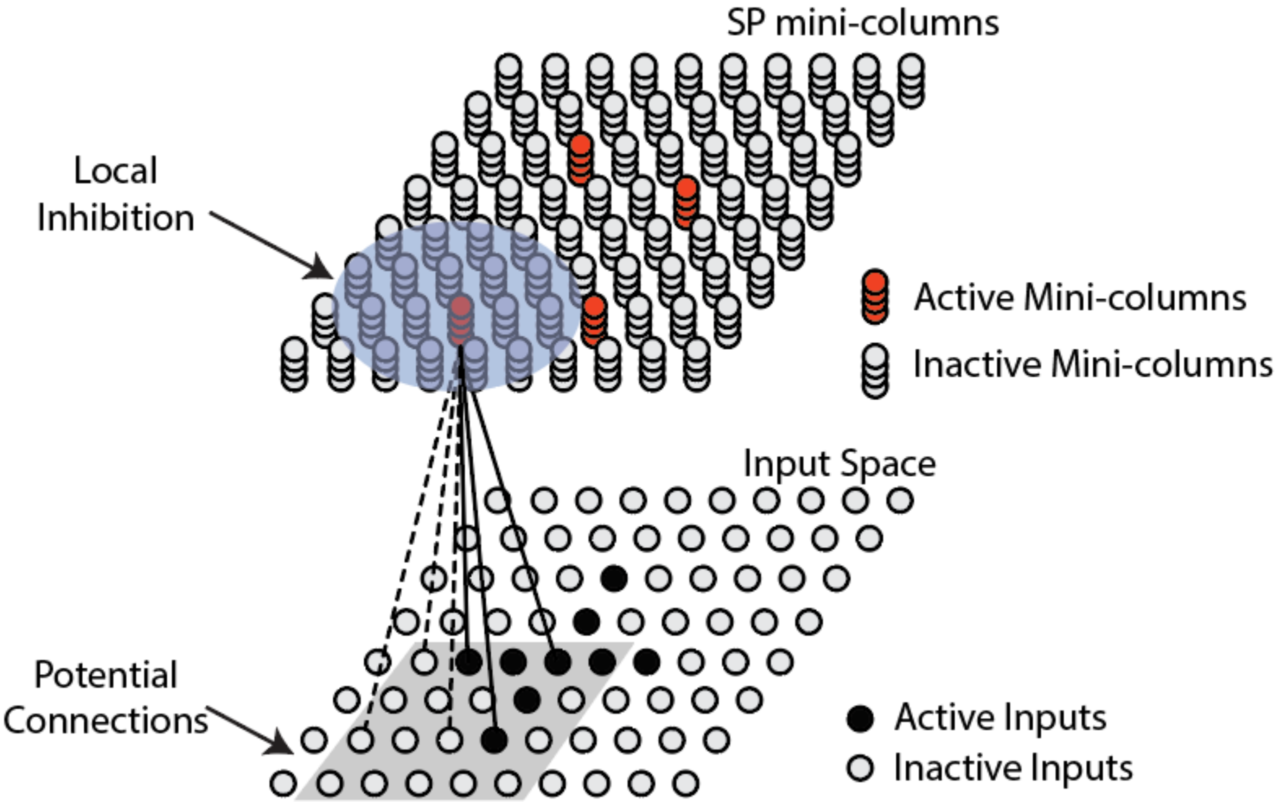

$$o_i = b_i \sum_j W_{ij} z_j$$

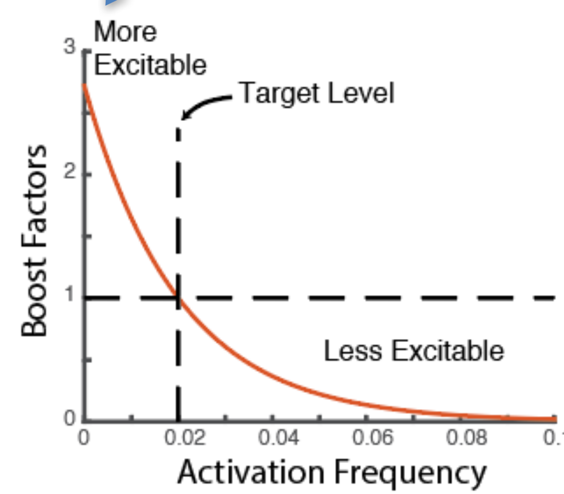

HTM spatial pooler – learning

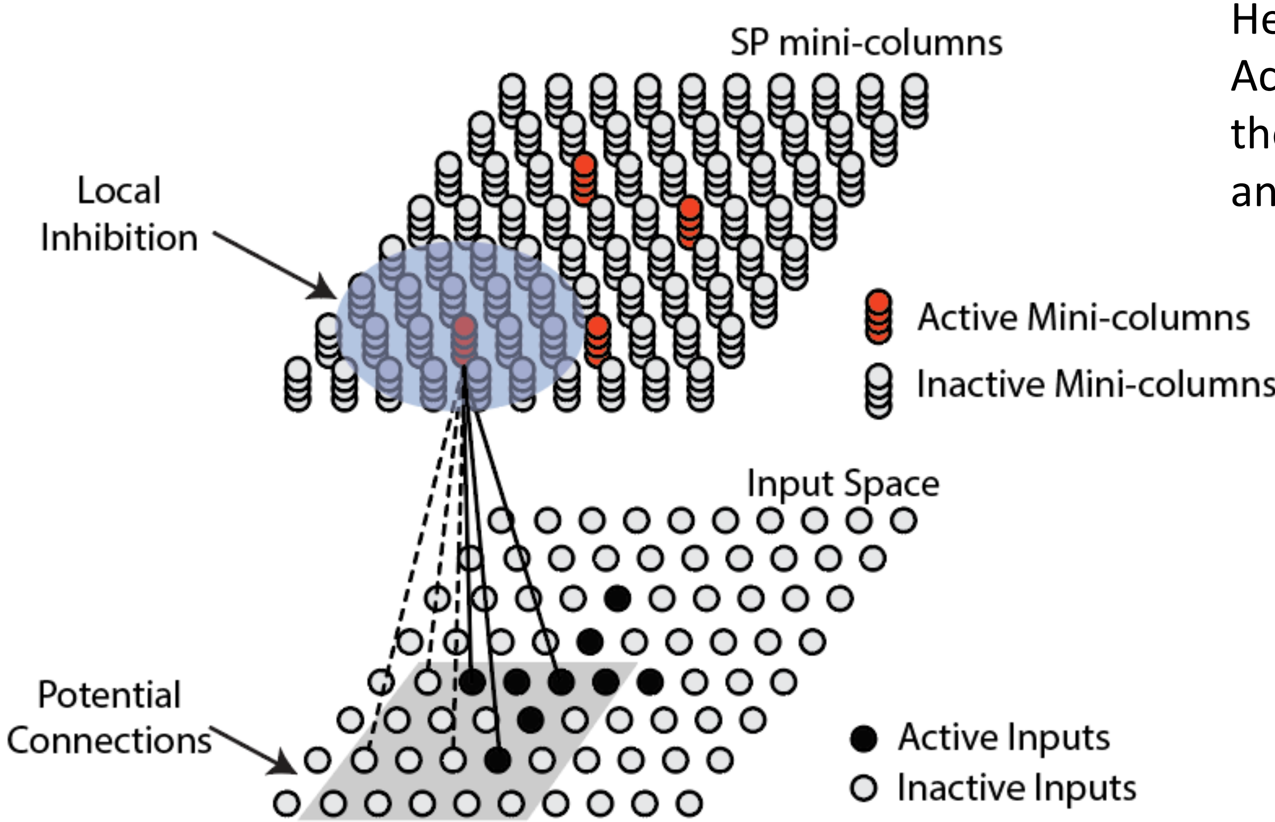

Hebbian Learning Rule:
Active (winner) MCs reinforce their active input connections, and depress inactive inputs

Learning in HTM Spatial Pooler

- Why do we need learning in SP?

If inputs are random, an untrained random SP will do just as good as any trained SP

However, real inputs are structured.
Input SDRs occurs with non-equal probabilities.
The actual inputs should be "better" represented than random inputs after learning

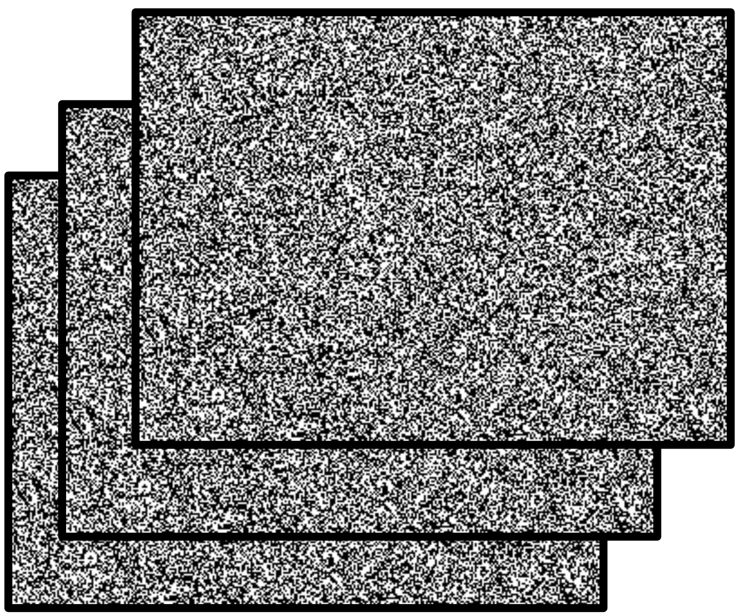

Properties of HTM spatial pooler

- Fixed-sparseness
- Distributed coding
- Preserving semantic similarity
- Noise robustness / Fault tolerance
- Continuous learning
- Stability

Properties of SP– fixed sparseness

Population sparseness:

$$s^t = \frac{1}{N}\sum_{i=1}^{N} a_i^t$$

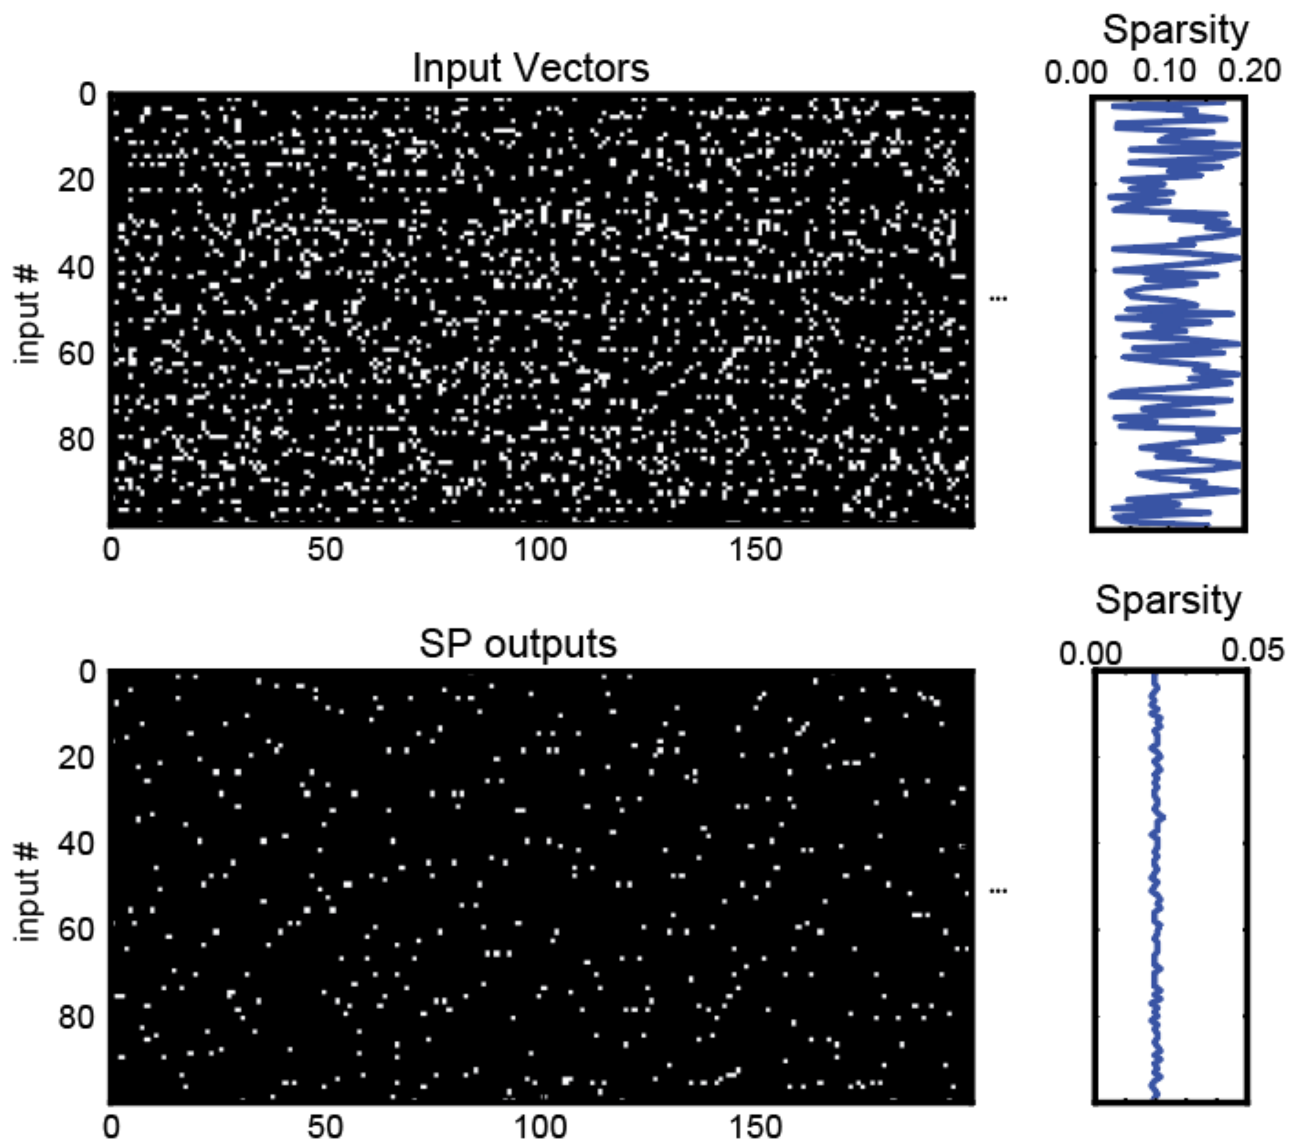

Properties of SP – distributed coding

Activation prob. of the *i*th MC $$P(a_i) = \frac{1}{M}\sum_{t=1}^{M} a_i^t$$

Entropy of the *i*th MC $$S_i = -P(a_i)\log_2 P(a_i) - (1 - P(a_i))\log_2(1 - P(a_i))$$

$$Entropy = \sum_{i=1}^{N} S_i$$

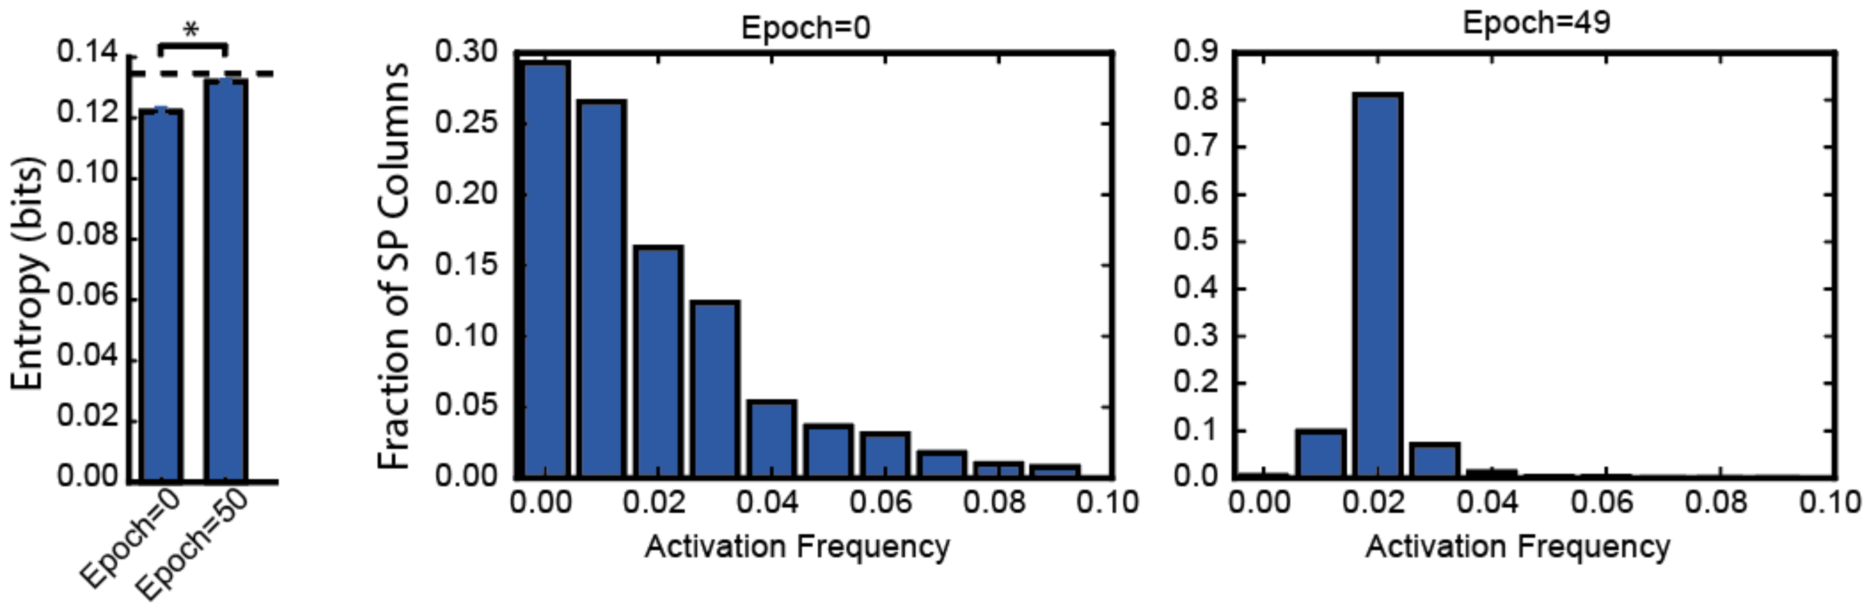

Properties of SP – noise robustness

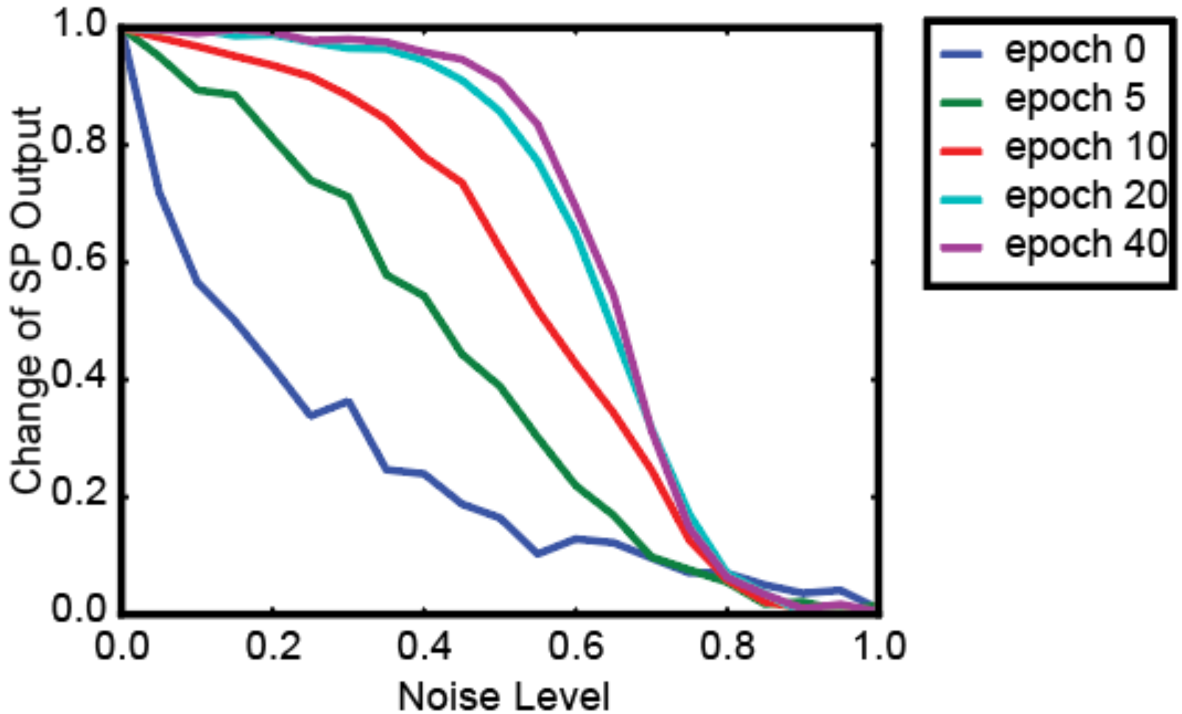

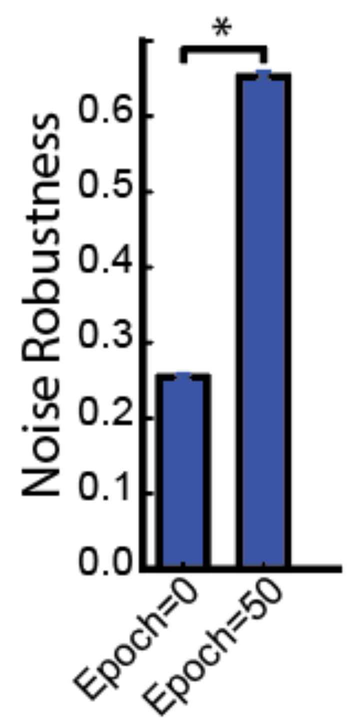

$$NoiseRobustness = \frac{1}{M}\sum_{i=1}^{M}\int_{k=0}^{1}\frac{\|\mathbf{a}_i \circ \mathbf{a}_i'(k)\|_0}{\|\mathbf{a}_i\|_0}\mathrm{d}k$$

Properties of SP – continuous learning

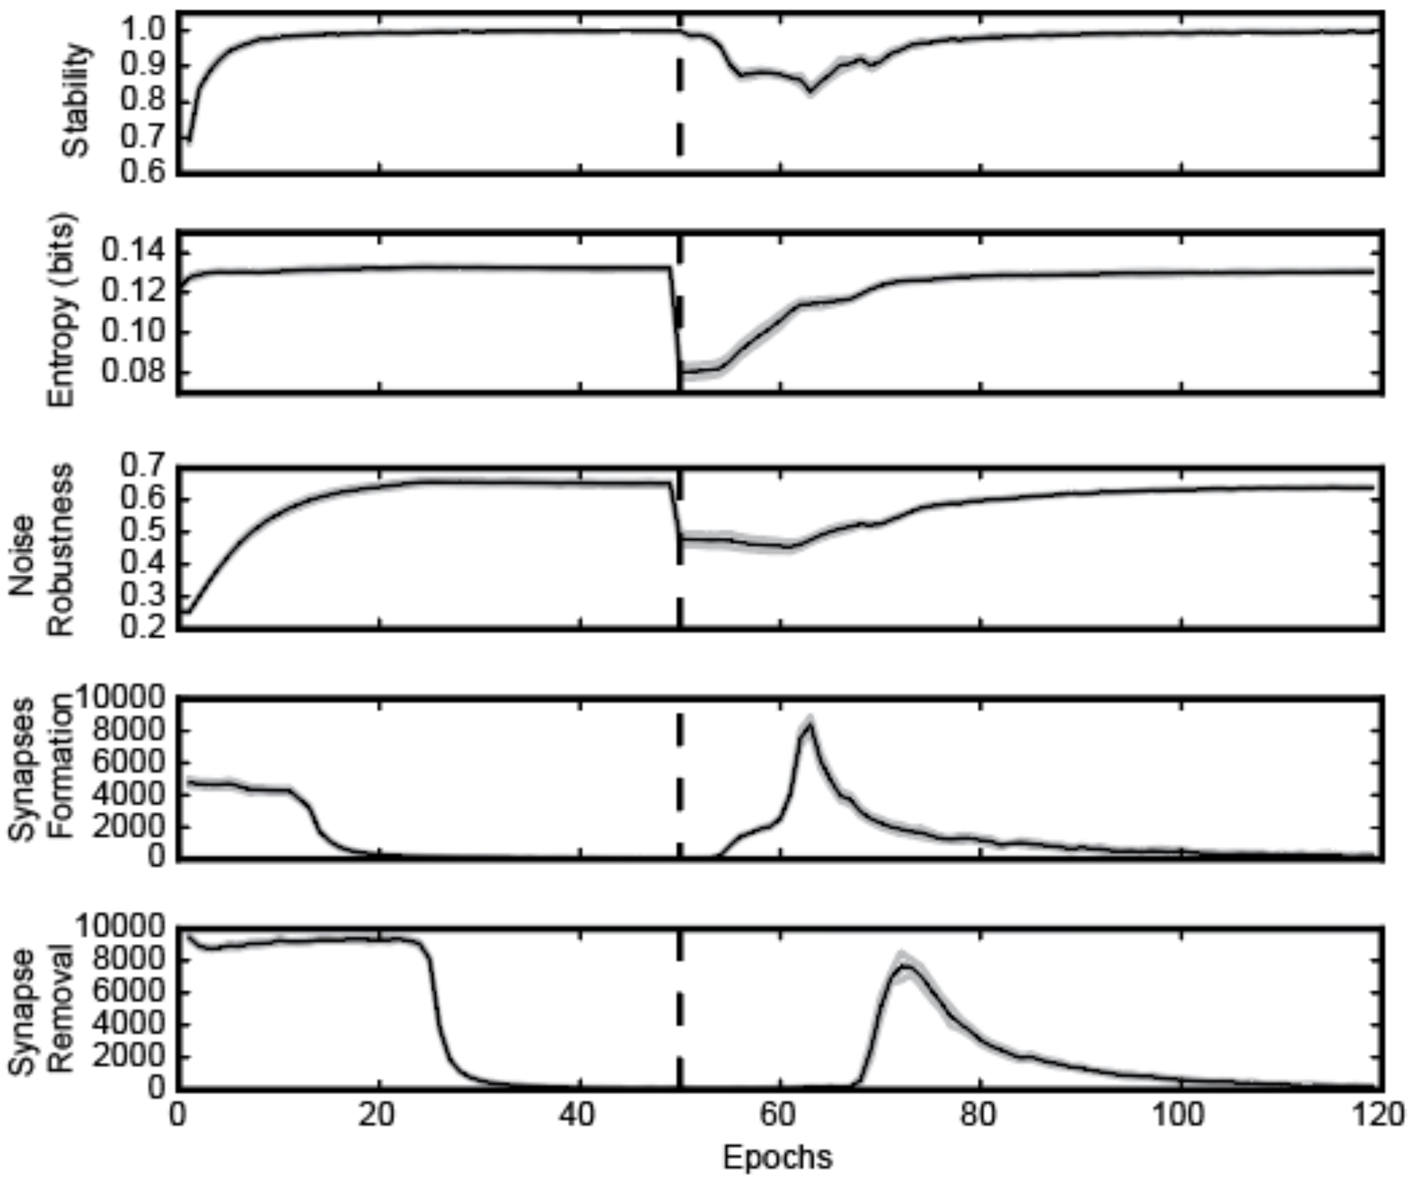

Properties of SP – continuous learning

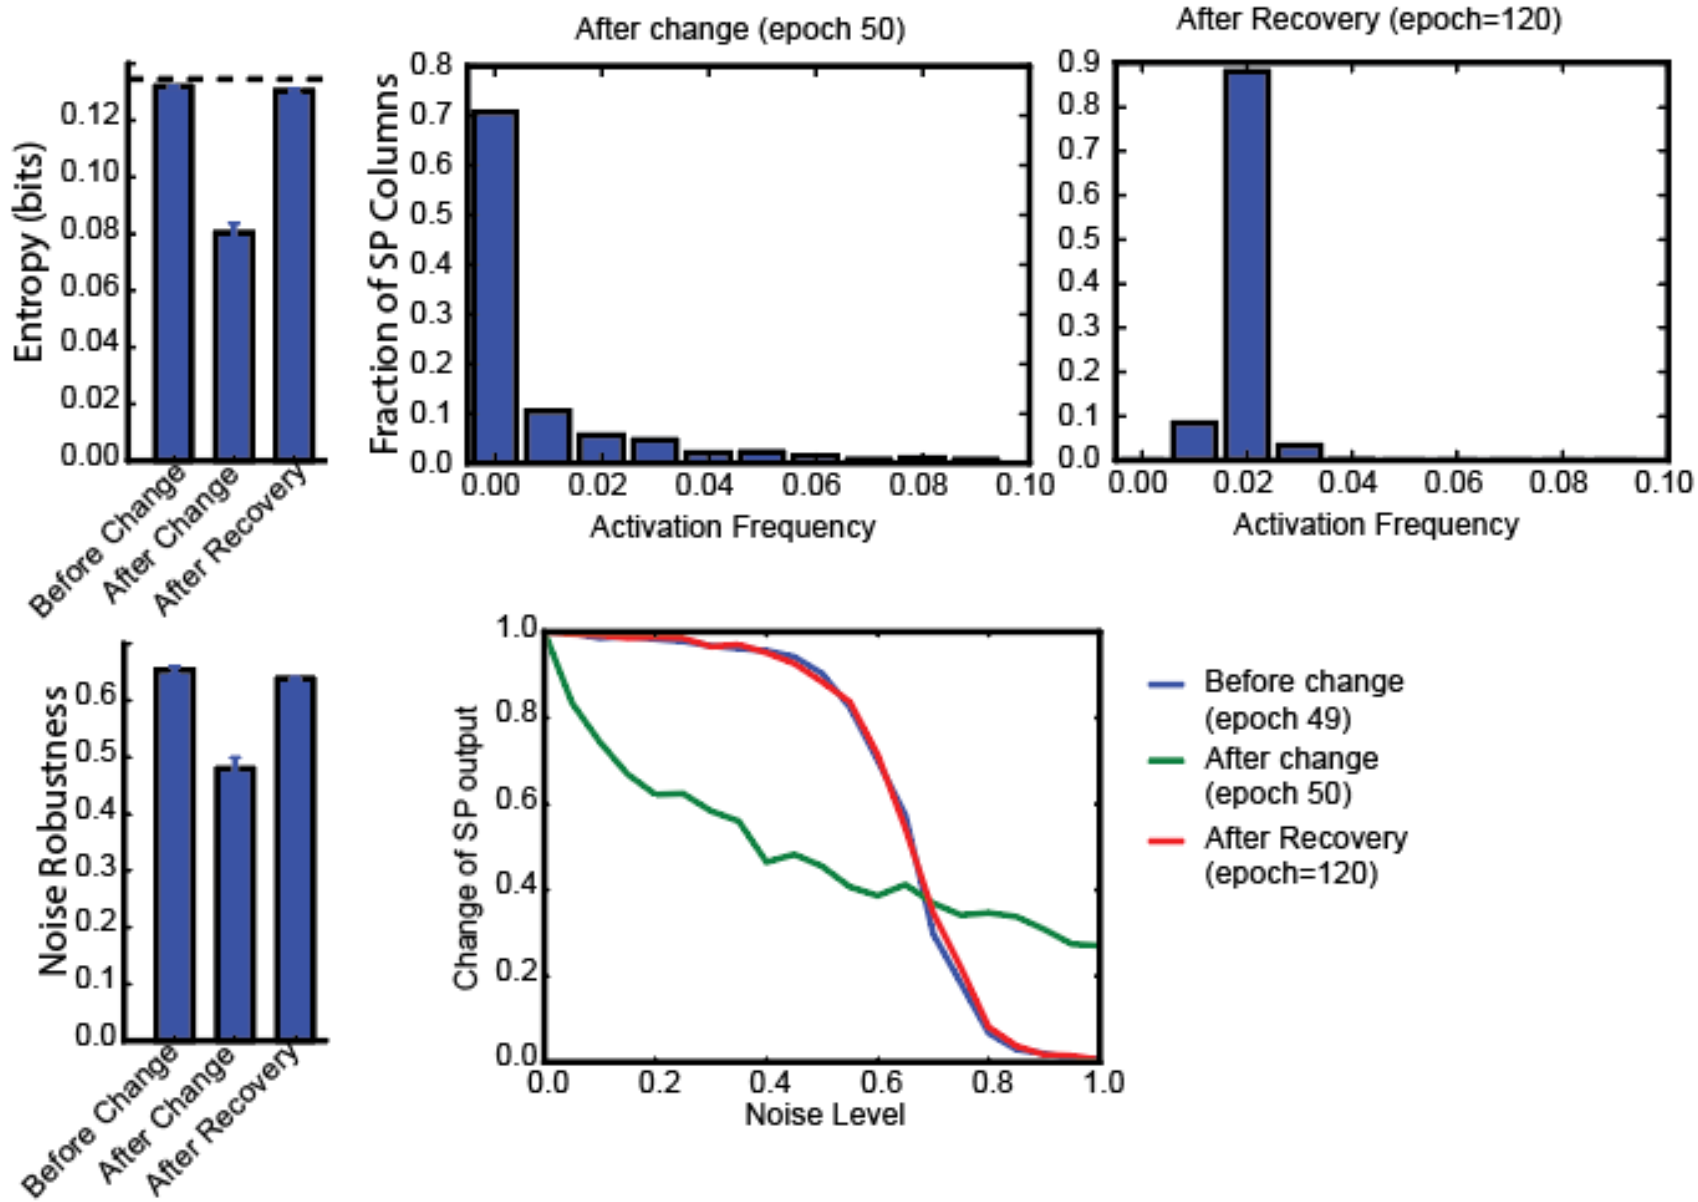

Properties of SP – fault tolerance

Experiment 1: lesion of SP MCs

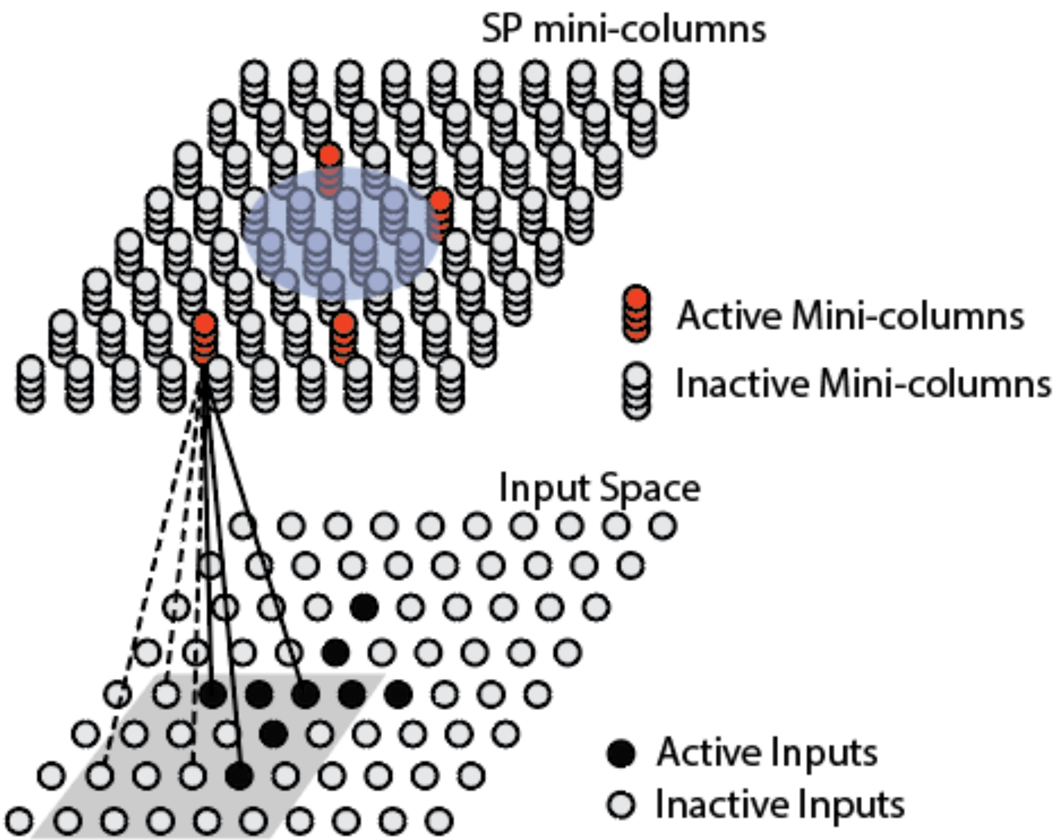

We monitor:

- RF center of SP MCs
- Coverage of input space
- Avg boost factors
- Growth & elimination of synapses

Properties of SP – fault tolerance

Experiment 1: lesion of SP MCs

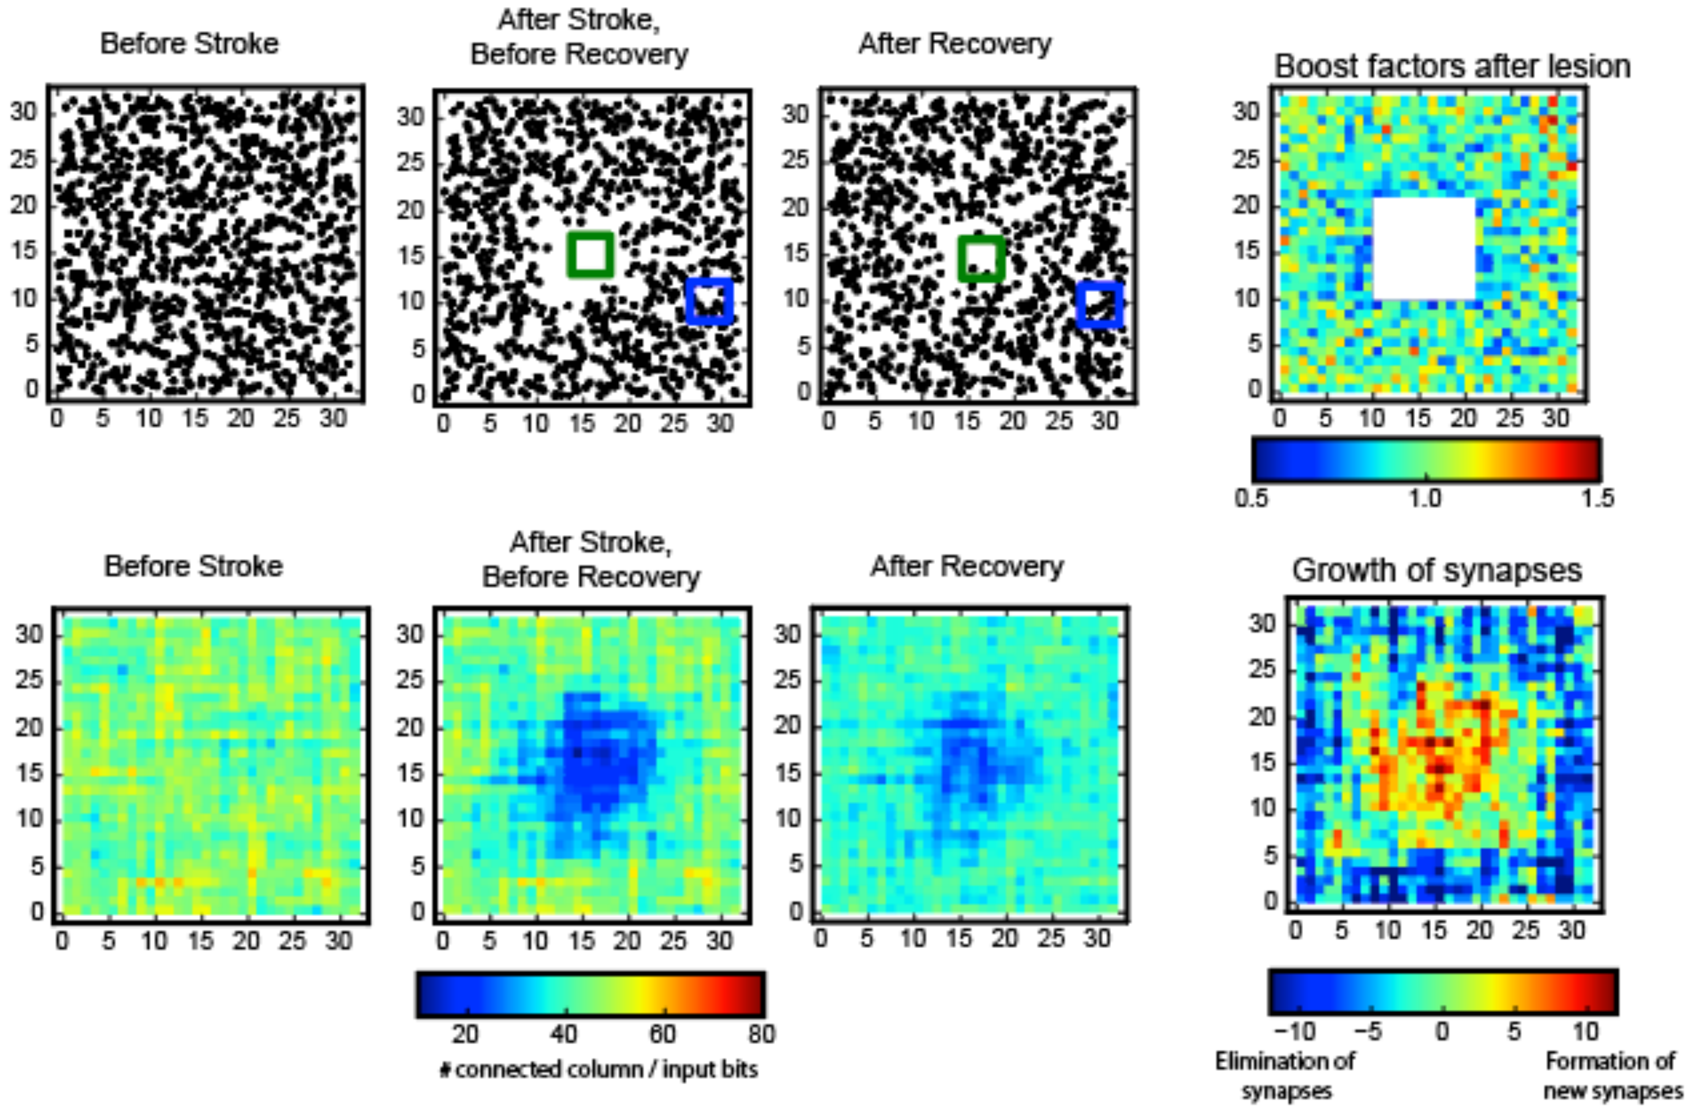

Properties of SP – fault tolerance

Experiment 1: lesion of SP MCs

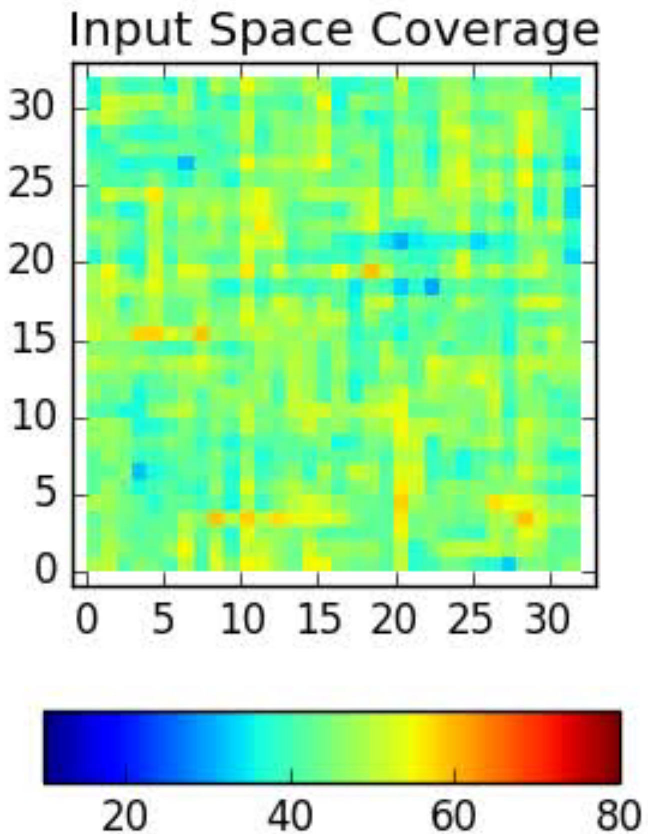

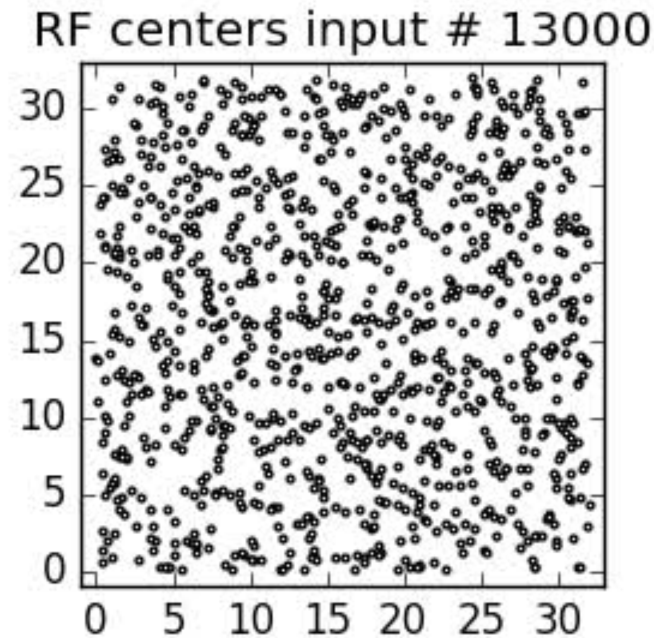

Properties of SP – fault tolerance

Experiment 2: lesion of afferent inputs

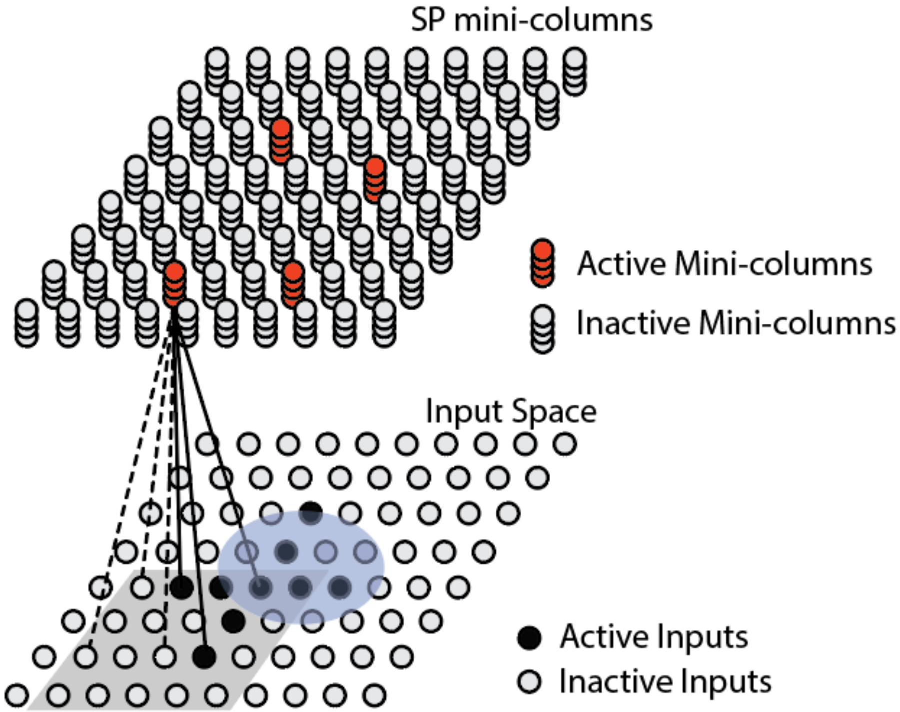

We monitor:

- RF center of SP MCs
- Coverage of input space
- Avg boost factors
- Growth & elimination of synapses

Properties of SP – fault tolerance

Experiment 2: lesion of afferent inputs

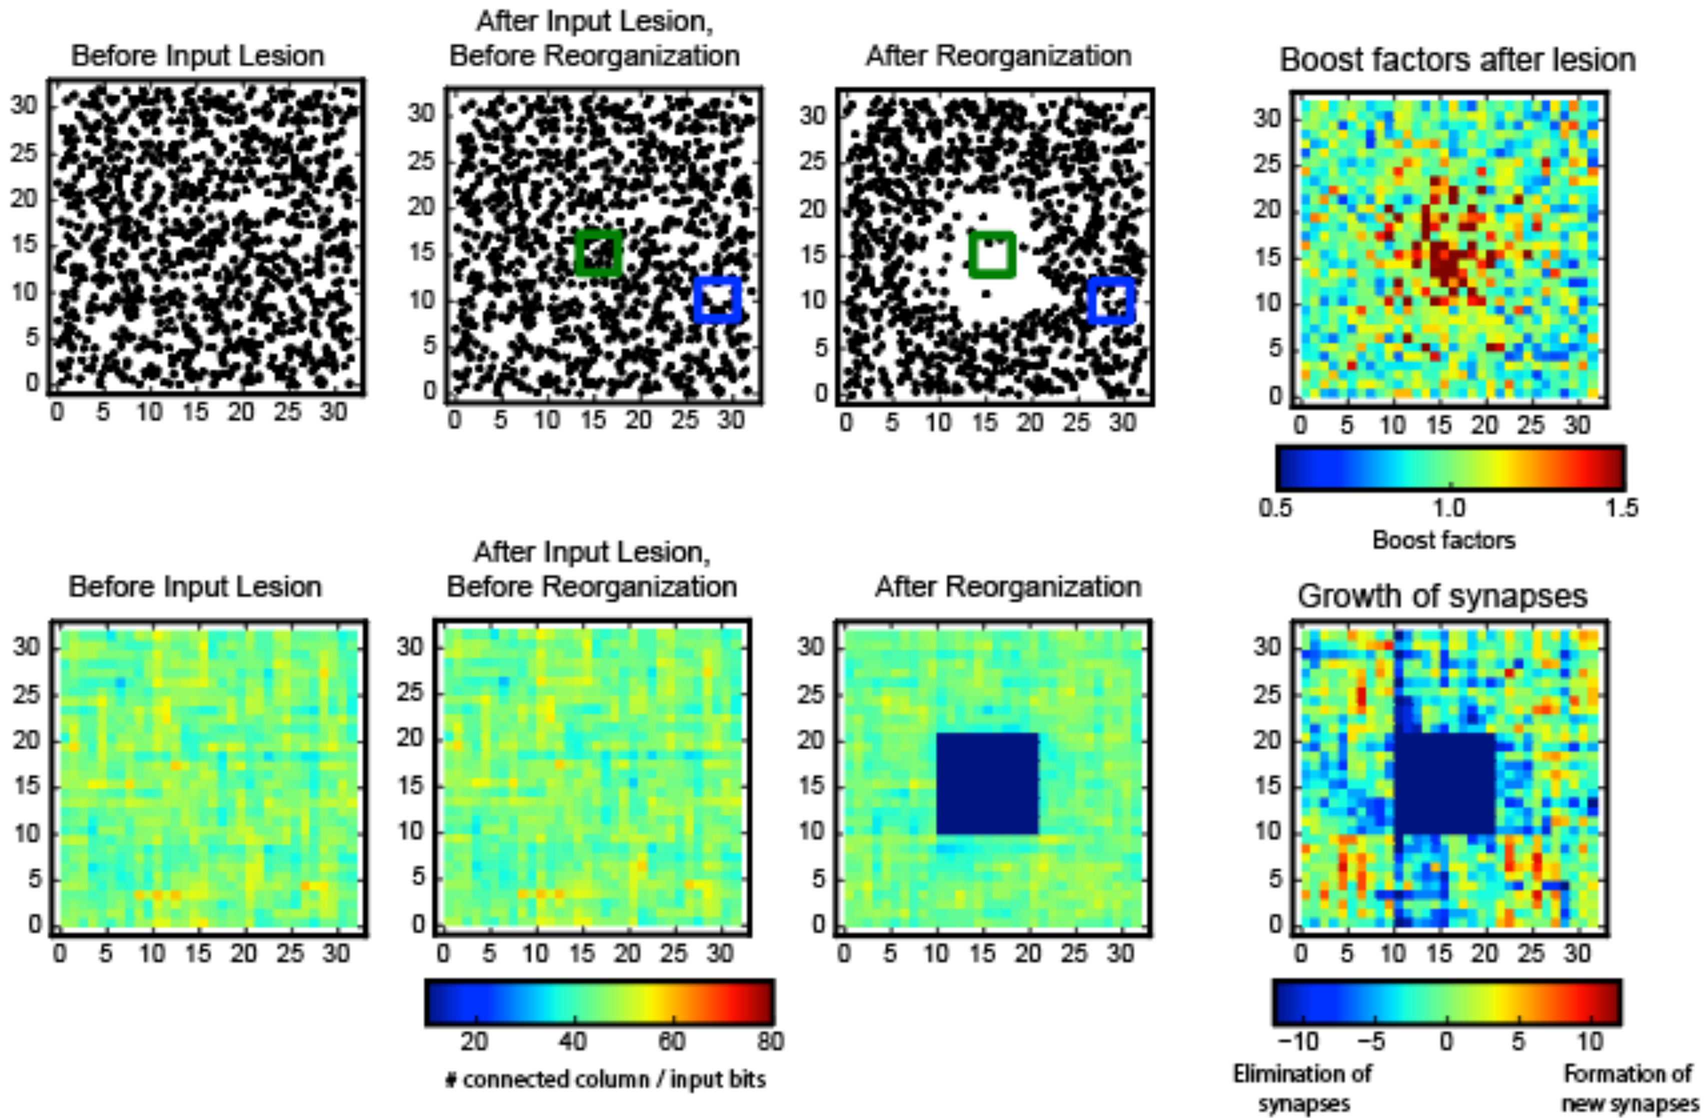

Properties of SP – fault tolerance

Experiment 2: lesion of afferent inputs

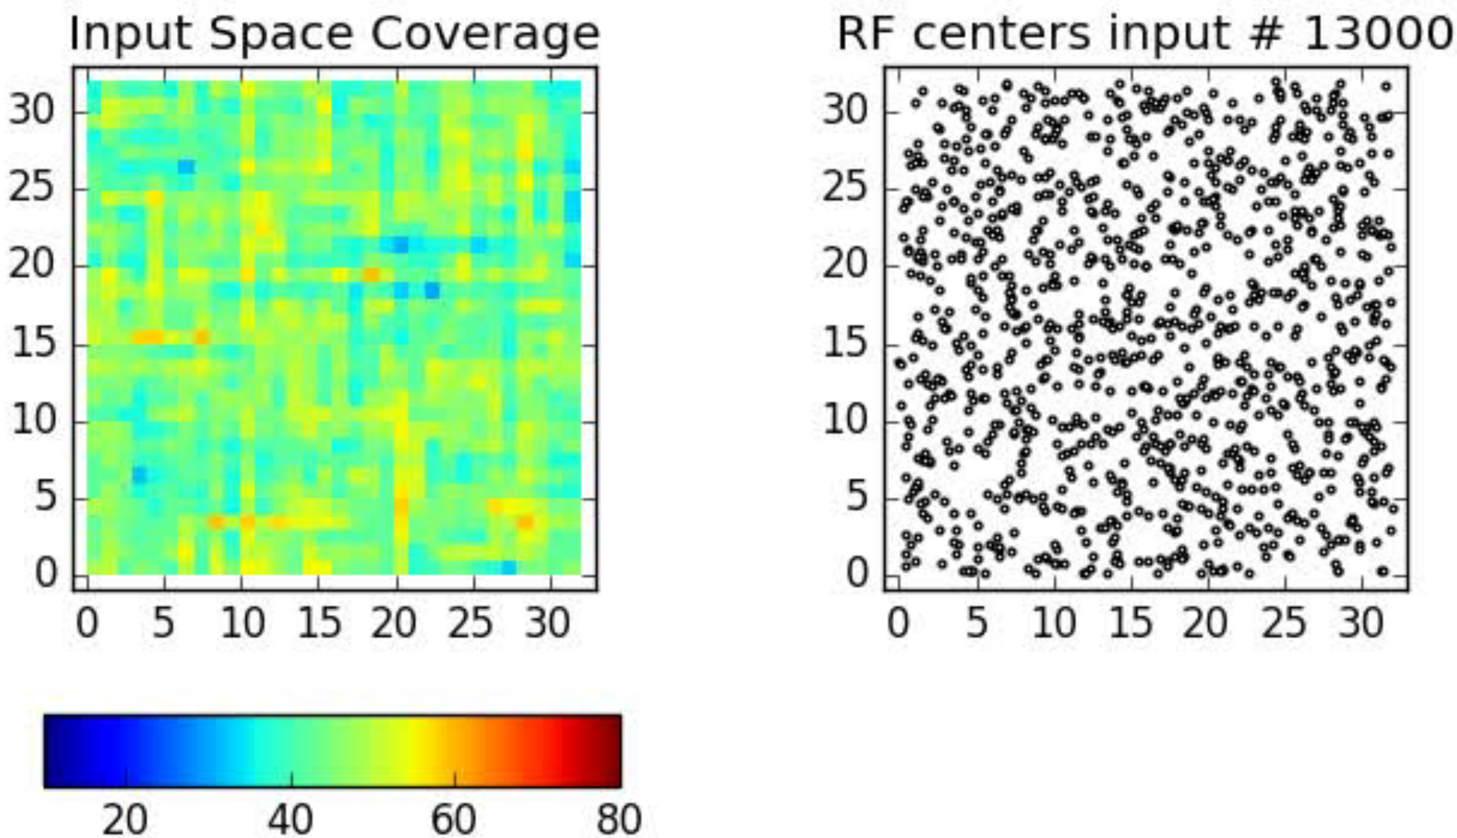

Properties of HTM spatial pooler

- Fixed-sparseness
- Distributed coding
- Preserving semantic similarity
- Noise robustness / Fault tolerance
- Continuous learning
- Stability

Spatial Pooler Resources

- Cui Y, Ahmad S, Hawkins J (under review) The HTM Spatial Pooler: A Neocortical Algorithm for Online Sparse Distributed Coding. bioRxiv. DOI: 10.1101/085035
- HTM School: http://numenta.org/htm-school/
- Spatial Pooler in NUPIC:
  https://github.com/numenta/nupic/blob/master/src/nupic/research/spatial_pooler.py
- Spatial Pooler pseudocode:

Neural mechanisms of SP