# Learning efficient, task-dependent representations with synaptic plasticity

Colin Bredenberg[1], Eero Simoncelli[1*], Cristina Savin[1*]

[*]Equal Contribution
[1]Center for Neural Science, NYU

email:
cjb617@nyu.edu

## Motivation

- Neural circuits cannot perfectly represent the sensory world: their capacity is limited by the number of neurons, response variability (noise) and various biophysical constraints.
- These limited resources should be allocated to represent the stimuli that matter, either because they are frequent or critical for obtaining reward.
- Can a recurrent neural network achieve this goal with simple, biologically plausible synaptic plasticity rules?

## Network architecture

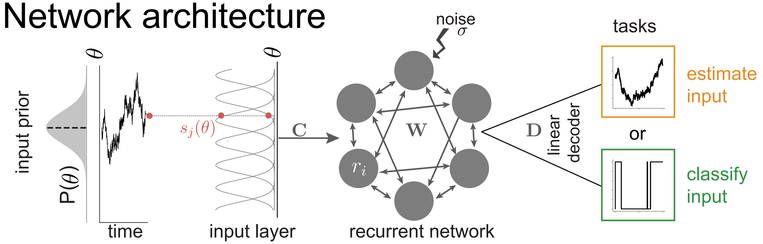


Stochastic nonlinear network dynamics (rate-based):

$$dr_i = \left[-\frac{f^{-1}(r_i)}{R} + \sum_{j=1}^{N} w_{ij} r_j + \sum_{k=1}^{d_s} c_{ij} s_j(\theta) + b_i\right] dt + \sigma dB_i$$

Energy:

$$E = -\frac{1}{2}\sum_{i,j} w_{ij} r_i r_j + \sum_i \frac{1}{R}\int_0^{r_i} f^{-1}(r_i) dr_i - \sum_{ij} c_{ij} s_j r_i - \sum_i b_i r_i \qquad p(\mathbf{r}|\mathbf{s}(\theta),\mathbf{W}) = \frac{1}{Z} e^{-\frac{E(\mathbf{r},\mathbf{s},\mathbf{W})}{\sigma^2}}$$

## Deriving a task-dependent synaptic plasticity rule

Optimize a task-specific objective function:

$$\mathcal{O}(\mathbf{D},\mathbf{W}) = \int \alpha(\mathbf{D}\mathbf{r},\mathbf{s})\, p(\mathbf{r}|\mathbf{s}(\theta),\mathbf{W})\, p(\theta)\, d\mathbf{r}\, d\theta - \lambda \|\mathbf{W}\|_2$$

accuracy — resource constraints

by stochastic gradient ascent:

$$\Delta w_{ij} \propto \frac{\partial}{\partial w_{ij}} \mathcal{O}(\mathbf{D},\mathbf{W})$$

giving a synaptic update:

$$\Delta w_{ij} \propto \alpha(\mathbf{D}\mathbf{r},\mathbf{s}) \left(r_i r_j - \langle r_i r_j \rangle_{p(\mathbf{r}|\mathbf{s})}\right) - \lambda w_{ij}$$

*Similar local update rules for $\mathbf{C}, \mathbf{b}, \mathbf{D}$

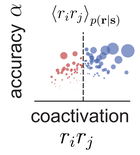


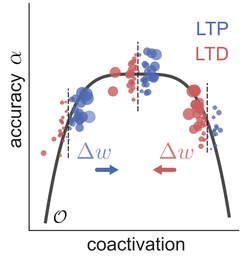


## Learning task-specific representations

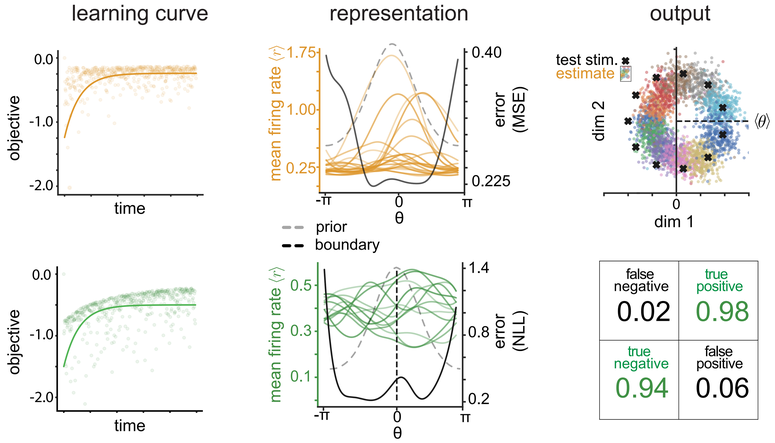


- Learned tuning functions are concentrated in the task-relevant region.
- Better performance for frequent stimuli.

## The effects of internal noise

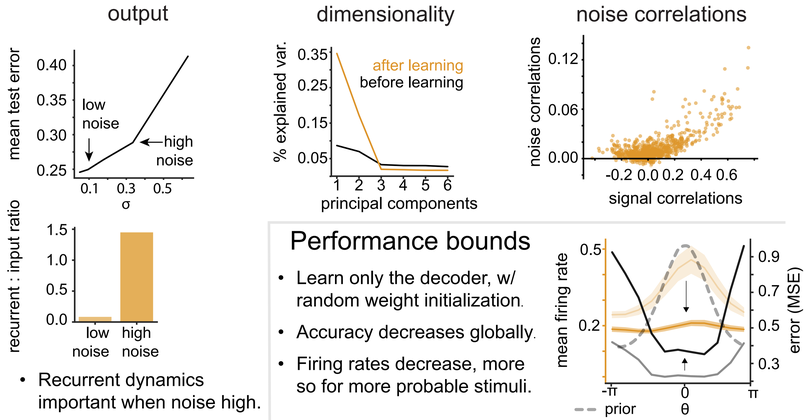


- Recurrent dynamics important when noise high.

### Performance bounds

- Learn only the decoder, w/ random weight initialization.
- Accuracy decreases globally.
- Firing rates decrease, more so for more probable stimuli.




## The effects of prior statistics

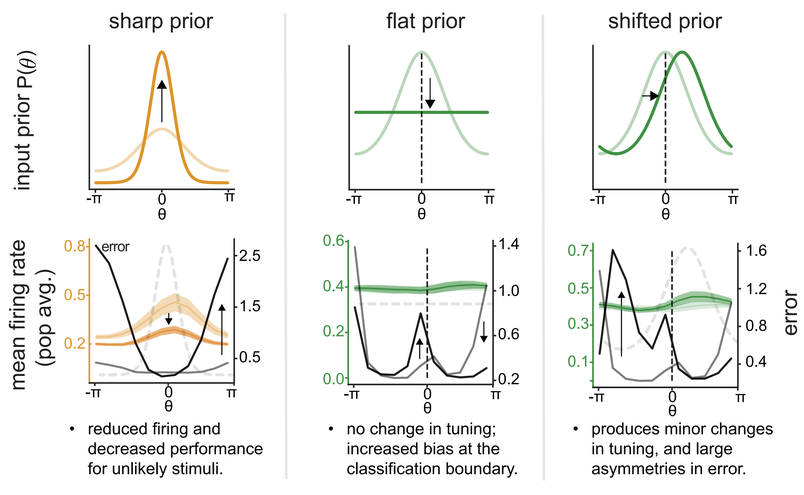


- reduced firing and decreased performance for unlikely stimuli.
- no change in tuning; increased bias at the classification boundary.
- produces minor changes in tuning, and large asymmetries in error.

## Conclusions

- We derive local plasticity rules that optimize task-specific cost functions, under resource constraints (internal noise and metabolic limitations).
- The model takes advantage of intrinsic network noise to perform stochastic gradient ascent on task objective functions.
- The circuit learns to exploit natural input statistics by concentrating neural resources around frequent stimuli, with a corresponding improvement in performance.
- Future work: time-dependence, unsupervised learning, simultaneously learning several tasks ...

## References

[1] D. Ganguli and E.Simoncelli. Efficient sensory encoding and Bayesian inference with heterogeneous neural populations. Neural Comp, 2014.
[2] N. Frémaux and W. Gerstner. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. Frontiers in neural circuits, 9:85, 2016.
[3] C. Savin and J. Triesch. Emergence of task-dependent representations in working memory circuits. Frontiers in Comp. Neurosci., 8:57, 2014.
[4] D. Ackley, G. Hinton, and T. Sejnowski. A learning algorithm for Boltzmann machines. In Readings in Computer Vision, p522–533. Elsevier, 1987.