Correspondence concerning this article should be addressed to Rachel Lee, Department of Mathematical & Computational Science, Stanford University, CA 94309

Email: rachlee@stanford.edu

Website: stanford.edu/~rachlee

Modeling Perceptual Learning with Deep Networks

R. Lee, A. Saxe, J. McClelland

Department of Mathematical & Computational Science, Electrical Engineering, and Psychology, Stanford University

Introduction

- Behavioral performance on orientation discrimination tasks improves with practice
- Debate exists over which visual areas (V1, V4, IT) change to realize these improvements
- Saxe (2014) used a tractable linear multilayer network driven by gradient descent to derive predicted changes across several layers of a deep network. His simplified linear network matches neurobiology data of perceptual learning from Schoups et al. 2001 and Raiguel et al. 2006.

Simple Linear Network (Saxe, 2014)

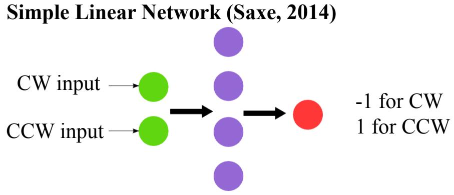

- We investigate if a more biologically realistic, non-linearly activated gradient descent-driven deep network exhibits similar learning dynamics as Saxe's linear network and can also model perceptual learning.

Method

Non-linear Network

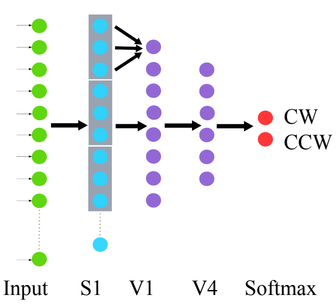

1. Network takes in gabor functions (GF) of different orientations and randomized phase pixel by pixel; goal is to determine if GF is clockwise or counterclockwise to trained orientation of 0°.
2. First layer (simple cell) weights are initialized with rows of GFs from -90 to 90 degrees; second layer 'complex cells' use max-pooling to achieve phase invariance; third layer 'V4' weights are initialized to Gaussian curves; final 'Decision layer' softmax weights are initialized to zero.
3. Orientation Tuning Curves of each layer are determined pre- and post- training

Results

Result 1: Higher layers change more than lower layers.

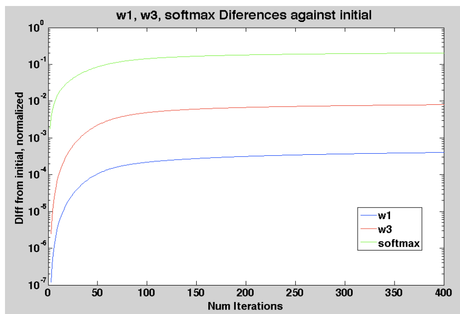

Result 2: Higher layers always change faster than lower layers

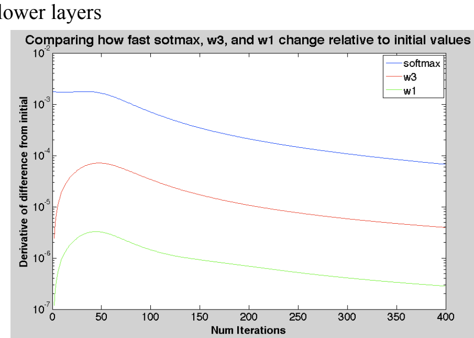

Result 3: Easier discrimination leads to only higher layers changing; Harder discrimination leads to all layers changing.

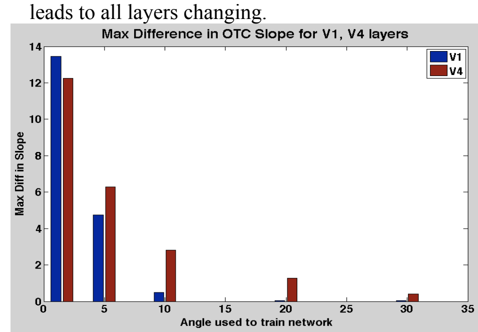

Result 4: The most informative neuron, with prefered orientation equal to std. dev. σ of the Orientation Tuning Curve (OTC) changes the most

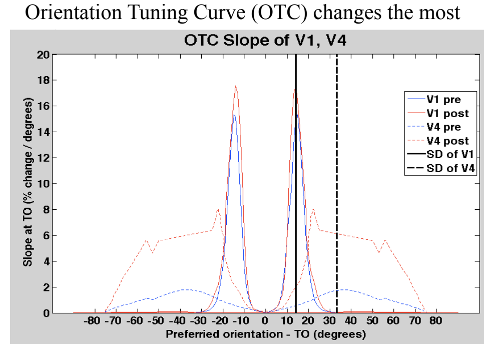

Schoups et al. 2001 on OTC Slope changes for V1

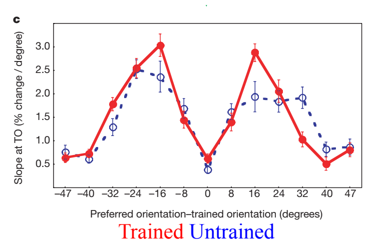

Raiguel et al. 2006 OTC Slope Changes for V4

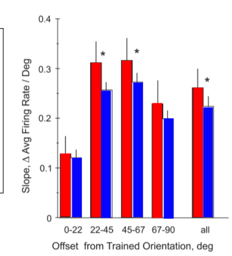

Discussion

- The non-linear network expands Saxe's simple linear network by taking in GFs pixel by pixel, building phase invariant complex cells from phase-sensitive simple cells, and using non-linear activation functions.
- Despite added complexity, the four primary learning dynamic properties found in Saxe's simple linear network can also be found in the non-linear network, suggesting that the linear network has potential for predicting properties of the more complex non-linear network.
- The non-linear network's results qualitatively match neurobiology results for the particular OTC curve changes found in V1 and V4 layers. We expect further parameter tuning to improve the quantitative match.
- We can expand our understanding of the mechanism and properties of perceptual learning by testing intuitions we can more easily deduce from the linear network on the more biologically realistic non-linear network.

Acknowledgements

Many thanks to PI Andrew Saxe, my mentor Jay McClelland, and many more colleagues in the PDP labs. Thank you to Andrew Ng's Deep Learning Tutorial for source code. Thank you to Stanford UROP for funding this summer. Thank you to Sym Sys and Bio-X for funding the project last summer.

References

[1] Saxe, A. "Deep linear neural networks: A theory of learning in brain and mind." Stanford University, PhD Dissertation Defense, June, 2014.

[2] Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. Nature, 412(6846), 549–53

[3] Raiguel, S., Vogels, R., Mysore, S. G., & Orban, G. A. (2006). Learning to see the difference specifically alters the most informative V4 neurons. The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 26(24), 6589–602.