

A bird's-eye view of PetaVision, the world's first petaflop/s neural simulation.

Daniel R. Coates¹, Craig Rasmussen², Garrett T. Kenyon²

¹Computer Science Department, Portland State University.

²Los Alamos National Laboratory.

Contact: dcoates@pdx.edu

1 Introduction

Despite decades of intense research, fundamental aspects of visual processing in the cortex are still poorly understood[1]. Even the mechanisms for pre-attentive early vision remain elusive. To address such questions, the PetaVision project at Los Alamos National Laboratory aims to demonstrate synthetic visual cognition using a simulation of mammalian cortex running on the Roadrunner supercomputer. Visual computation is realized through the temporal dynamics of a massive anatomically-derived network of spiking neurons responding to images presented to an artificial retina.

2 Motivation

Several results from neuroscience inform the approach. First, it is now well-established that the connectivity of neurons in the early visual path is highly specific[2]. Lateral connections in V1 have been rigorously quantified[3], and recent theories hypothesize a role in contour integration[3, 4, 5] and other perceptual processing[5]. Second, it is becoming clear that temporal dynamics are important in neural processing. Thus the system employs the spiking neural model of Worgotter and Koch[6], an integrate-and-fire-like neuron augmented with separate excitatory and inhibitory channels.

The benchmark application involves detection of a closed contour amongst a field of distractors. The difficulty traditional computer vision algorithms have with such a task, versus the ease with which a human can detect such 'pop-out' targets, typifies the gap that the PetaVision project is trying to close.

3 Methods and Results

The neural connectivity patterns suggest a natural implementation on a grid of parallel processing nodes. Designed with this target platform in mind, the code scales using both shared-memory multiprocessing and MPI, running efficiently on commodity multicore PCs and on the 12,000-processor Roadrunner. The idiosyncratic hybrid architecture of Roadrunner[7] presents significant programming challenges, but the layered, columnar structure of visual cortex has an obvious topographic mapping.

Initial experiments utilize a 64x64 monochrome image and more than 100,000 neurons, with simulations corresponding to 300-500 milliseconds of neural response. Using standard spike analysis techniques, cortical activity is post-processed and object detection is implicated in self-organized temporal correlations in the spike train data.

References

- [1] B. A. Olshausen and D. J. Field. How close are we to understanding v1? *Neural Computation*, 17:1665–1699, 2005.
- [2] D. H. Hubel T. N. Wiesel. Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)*, 195:215–43, 1968.
- [3] W. H. Bosking, Y. Zhang, B. Schofield, and D. Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *The Journal of Neuroscience*, 17:2112–2127, 1997.
- [4] D. J. Field, A. Hayes, and R.F. Hess. Contour integration by the human visual system: evidence for a local association field. *Vision Research*, 33:173–193, 1993.
- [5] O. Ben-Shahar and S. W. Zucker. Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Computation*, 16:445–476, 2004.
- [6] F. Worgotter and C. Koch. A detailed model of the primary visual pathway in the cat: Comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity. *The Journal of Neuroscience*, 11:1959–1979, 1991.
- [7] <http://www.lanl.gov/roadrunner/>.