Scaling self-organizing maps to model large cortical networks

Amol Kelkar, James A. Bednar and Risto Miikkulainen
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712
amol, jbednar, risto@cs.utexas.edu

Abstract

Self-organizing computational models with specific intracortical connections can explain many functional features of visual cortex, such as topographic orientation and ocular dominance maps. However, due to their computational requirements, it is difficult to use such detailed models to study large-scale phenomena like object segmentation and binding, object recognition, tilt illusions, optic flow, and fovea–periphery interaction. This paper introduces two techniques that make large simulations practical. First, a set of general linear scaling equations for the RF-LISSOM self-organizing model is derived and shown to result in quantitatively equivalent maps over a wide range of simulation sizes. This capability makes it possible to debug small simulations and then scale them up to larger simulations only when needed. The scaling equations also facilitate the comparison of biological maps and parameters between individuals and species with different brain region sizes. Second, the equations are combined into a new growing map method called GLISSOM, which dramatically reduces the memory and computational requirements of large self-organizing networks. With GLISSOM it should be possible to simulate all of human V1 at the single-column level using existing supercomputers, making detailed computational study of large-scale phenomena possible.

1 Introduction

Computational models of the self-organization in the visual cortex have shown that input-driven development can explain much of its topographic organization, such as retinotopy, orientation preference, and ocular dominance, as well as many of its functional properties, such as short-range contour segmentation and binding (Grossberg 1976; Kohonen 1989; von der Malsburg 1973; see Erwin et al. 1995; Swindale 1996 for review). However, other important phenomena have remained out of reach because they require too much computation time and memory to simulate. These phenomena, such as orientation interactions between spatially separated stimuli and long-range visual contour and object integration, are thought to arise out of specific lateral interactions between large numbers of neurons over a wide cortical area (Gilbert et al. 1996). Simulating such behavior requires an enormous number of specific, modifiable connections. Currently-practical methods can only model intracortical interactions abstractly (e.g. SOM, Erwin et al. 1992; Kohonen 1989; Obermayer et al. 1990), and thus cannot be used for such investigations.

In this paper we present two interrelated techniques for making detailed large-scale simulations practical. First, we derive a set of linear scaling equations that, when given a small-scale simulation, make it possible to determine the parameter settings necessary to perform a large-scale simulation. The original and scaled simulations have quantitatively-equivalent map-level and neuron-level organization; the larger map will just have more detail. Such a correspondence makes it possible to develop a small-scale simulation first using available hardware, then scale it up to study specific phenomena that require a larger map. The scaling equations can also help tie parameters from small models to experimental measurements in larger systems, help determine simulation sizes needed for realistic simulations, and allow comparison of species or individuals with brain regions of different sizes.

Second, we present a modeling approach called GLISSOM that allows much larger networks to be simulated in a given computation time and in a given amount of memory. The simulations begin with a small network, which is gradually scaled up as it self-organizes. This approach is effective for two reasons: (1) pruning-based self-organizing models tend to have peak computational and memory requirements at the beginning of training, and (2) self-organization tends to proceed in a global-to-local fashion, with large-scale order established first, followed by more detailed local self-organization (as found in experimental animals; Chapman et al. 1996). Thus small maps, which are much quicker to simulate and take less memory, can be used to establish global order, with larger maps used only to achieve more detailed structure.

Although the primary motivation for GLISSOM is computational, the scaling process is also well-motivated biologically, since it represents the integration of new neurons into an existing region during development. Recent experimental results suggest that new neurons continue to be added even in adulthood in many areas of primate cortex (Gould et al. 1999). Moreover, many of the neurons in the immature cortex corresponding to GLISSOM's early stages have not yet begun to make functional connections, having only recently migrated to their final positions (Purves 1988). Thus the scaleup

Technical Report AI-2000-285 August 2001
http://www.cs.utexas.edu/users/nn/

procedure in GLISSOM corresponds to the gradual process of incorporating those neurons into the partially-organized map.

In the next section the model used in these simulations is introduced, and in section 3 scaling equations for it are derived and shown to achieve matching results over a wide range of simulation sizes. In section 4 the GLISSOM scaleup procedure is introduced and shown to greatly reduce simulation time and memory requirements while achieving results similar to the original model. Section 5 shows calculations that suggest that with GLISSOM it should be possible to simulate all of human V1 at the single-column level using existing supercomputers. The remaining sections discuss how the scaling equations relate to biological systems and how they can be used to simulate larger, more realistic systems that would otherwise be intractable.

2 RF-LISSOM model of the visual cortex

The scaling equations and GLISSOM are based on the RF-LISSOM (Receptive-Field Laterally Interconnected Synergetically Self-Organizing Map) computational model of cortical maps. RF-LISSOM has been successfully used to model the development of ocular dominance and orientation maps, as well as low-level visual phenomena in the adult, such as tilt aftereffects and short-range segmentation and binding (Bednar and Miikkulainen 2000; Choe and Miikkulainen 1998; Miikkulainen et al. 1997; Sirosh and Miikkulainen 1994; Sirosh et al. 1996). We will first describe the architecture of the RF-LISSOM model, and then later present our extensions that allow scaling the network.

RF-LISSOM focuses on the two-dimensional organization of the cortex, so each "neuron" in the model cortex corresponds to a vertical column of cells through the six layers of the primate cortex. The cortical network is modeled with a sheet of interconnected neurons and the retina with a sheet of retinal ganglion cells (figure 1). Neurons receive afferent connections from broad overlapping circular patches on the retina. (Since the lateral geniculate nucleus (LGN) accurately reproduces the receptive fields of the retina, it has been bypassed to simplify the model.) The $N \times N$ network is projected on to the $R \times R$ retinal ganglion cells, and each neuron is connected to ganglion cells in an area of radius r_{A} around its projection. Thus, neurons at a particular cortical location receive afferents from a corresponding location on the retina, i.e. its anatomical receptive field (RF). Additional ganglion cells are included around the borders so that every neuron will have a complete set of afferent connections. For an example set of weights, see figure 9*a-c* in section 4.1.

In addition to the afferent connections, each neuron has reciprocal excitatory and inhibitory lateral connections with itself and other neurons. Lateral excitatory connections are short-range, connecting each neuron with itself and its close neighbors. Lateral inhibitory connections run for comparatively long distances, but also include connections to the neuron itself and to its neighbors.[1]

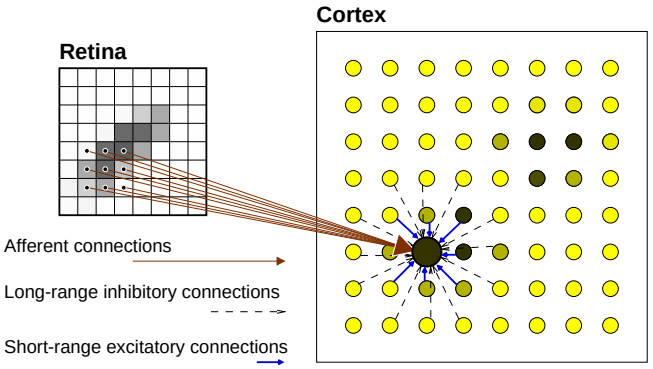

Figure 1: **Architecture of the RF-LISSOM network.** A small RF-LISSOM network and retina are shown, along with connections to a single neuron (shown as the large circle). The input is an oriented Gaussian activity pattern on the retinal ganglion cells (shown by grayscale coding); the LGN is bypassed for simplicity. The afferent connections form a local anatomical receptive field (RF) on the simulated retina. Neighboring neurons have different but highly overlapping RFs. Each neuron computes an initial response as a scalar (dot) product of its receptive field and its afferent weight vector, i.e. a sum of the product of each weight with its associated receptor. The responses then repeatedly propagate within the cortex through the lateral connections and evolve into activity "bubbles". After the activity stabilizes, weights of the active neurons are adapted using a normalized Hebbian rule.

The afferent weights are initially set to random values, and the lateral weights are preset to a smooth Gaussian profile. The connections are then organized through an unsupervised learning process. For an orientation map, the input for the learning process consists of 2-D ellipsoidal Gaussian patterns representing retinal ganglion cell activations (figure 2*a*); each pattern is presented at a random orientation and position. At each training step, neurons start out with zero activity. The initial response η_{ij} of neuron (i, j) is calculated as a weighted sum of the retinal activations:

$$\eta_{ij} = \sigma\left(\sum_{a,b} \xi_{ab}\mu_{ij,ab}\right), \qquad (1)$$

where ξ_{ab} is the activation of retinal ganglion (a, b) within the receptive field of the neuron, $\mu_{ij,ab}$ is the corresponding afferent weight, and σ is a piecewise linear approximation of the sigmoid activation function. The response evolves over

[1] For high-contrast inputs, long-range interactions must be inhibitory for proper self-organization to occur (Sirosh 1995). Optical imaging and electrophysiological studies have indeed shown that long-range column-level interactions in the cortex are inhibitory at high contrasts, even though individual long-range lateral connections between neurons are primarily excitatory (Grinvald et al. 1994; Hirsch and Gilbert 1991; Weliky et al. 1995). The model uses explicit inhibitory connections for simplicity since all inputs used are high-contrast, and since it is such inputs that primarily drive adaptation in a Hebbian model.