Neural Networks 90 (2017) 29–41

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

Extending the Stabilized Supralinear Network model for binocular image processing

Ben Selby [*], Bryan Tripp

Department of Systems Design Engineering, University of Waterloo, 200 University Ave W., Waterloo, Ontario, Canada N2L 3G1
Centre for Theoretical Neuroscience, University of Waterloo, 200 University Ave W., Waterloo, Ontario, Canada N2L 3G1

ARTICLE INFO

Article history:
Received 20 July 2016
Received in revised form 23 December 2016
Accepted 3 March 2017
Available online 18 March 2017

Keywords:
Stabilized supralinear network
Balanced excitation/inhibition
Interocular transfer of suppression
Primary visual cortex
Convolutional neural networks

ABSTRACT

The visual cortex is both extensive and intricate. Computational models are needed to clarify the relationships between its local mechanisms and high-level functions. The Stabilized Supralinear Network (SSN) model was recently shown to account for many receptive field phenomena in V1, and also to predict subtle receptive field properties that were subsequently confirmed in vivo. In this study, we performed a preliminary exploration of whether the SSN is suitable for incorporation into large, functional models of the visual cortex, considering both its extensibility and computational tractability. First, whereas the SSN receives abstract orientation signals as input, we extended it to receive images (through a linear–nonlinear stage), and found that the extended version behaved similarly. Secondly, whereas the SSN had previously been studied in a monocular context, we found that it could also reproduce data on interocular transfer of surround suppression. Finally, we reformulated the SSN as a convolutional neural network, and found that it scaled well on parallel hardware. These results provide additional support for the plausibility of the SSN as a model of lateral interactions in V1, and suggest that the SSN is well suited as a component of complex vision models. Future work will use the SSN to explore relationships between local network interactions and sophisticated vision processes in large networks.

1. Introduction

The visual cortex has an intricate organization at multiple scales (Cossell et al., 2015; Harris & Mrsic-Flogel, 2013; Markov et al., 2014) and performs many sophisticated functions. Due to its complexity and nonlinearity, computational models are needed to understand interactions between cellular mechanisms, network interactions, and behaviour. However, it has been a challenge to develop cortical models that combine physiological mechanisms with non-trivial visual processing.

Convolutional neural networks (CNNs) are promising in this respect. Their mechanisms are highly abstract, but they have a number of parallels with the cortex in both their structure (Goodfellow, Bengio, & Courville, 2016) and unit activity (Yamins et al., 2014) (see also the Discussion). Importantly, they perform sophisticated visual processing, including state-of-the-art object recognition (Szegedy et al., 2015), stereoscopic depth estimation (Žbontar & LeCun, 2016), segmentation (Chen et al., 2014), and detection of grasp affordances (Lenz, Lee, & Saxena, 2013).

CNNs were invented in the late 1980s (LeCun et al., 1989), and although they have advanced since then, their advances have incorporated little from the last quarter century of neuroscience. Conversely, models that address more recent data (e.g. Bednar, 2012; Mineault, Khawaja, Butts, & Pack, 2012; Rubin, Van Hooser, & Miller, 2015) typically do not perform sophisticated image processing (see Serre et al., 2005, for a counter-example). Therefore, the relevance of recent advances in neuroscience to complex, naturalistic vision has remained largely unaddressed in computational models.

One of the key differences between convolutional networks and the visual cortex is the predominance of recurrent connections in the latter. This recurrence shapes the earliest responses (Somers, Nelson, & Sur, 1995), though perhaps only subtly for simple inputs (Li, Ibrahim, Liu, Zhang, & Tao, 2013). Recently, Miller and colleagues (Ahmadian, Rubin, & Miller, 2013; Rubin et al., 2015) developed the Stabilized Supralinear Network (SSN), a recurrent model with balanced excitation and inhibition that accounts for many receptive field properties (see details in the next section).

The present study is a preliminary exploration of the suitability of the SSN for modelling recurrent connections in large, functional

---

[*] Corresponding author at: Department of Systems Design Engineering, University of Waterloo, 200 University Ave W., Waterloo, Ontario, Canada N2L 3G1.
E-mail addresses: bselby@uwaterloo.ca (B. Selby), bptripp@uwaterloo.ca (B. Tripp).

http://dx.doi.org/10.1016/j.neunet.2017.03.003
0893-6080/© 2017 Elsevier Ltd. All rights reserved.

vision models. The SSN receives an abstract orientation field as input, so we began by developing an extension of the SSN that instead receives images as input, by adding an input stage of linear–nonlinear receptive fields. We found that the extended version behaved similarly to the original SSN. We then developed a binocular extension of the model that incorporated ocular dominance columns. We found that the SSN could reproduce data on interocular transfer of surround suppression. These results provide further support for the biological relevance of the SSN, and its relevance to more complex information processing. Finally, we implemented the SSN as a convolutional network, to determine whether it could run at practical rates for image processing and deep learning. We found that a large SSN model, with 20 million excitatory units, could process an image in about ten seconds on a single graphical processing unit (GPU). This suggests that it is possible to incorporate the SSN into a large convolutional network, but that it should be introduced at a late stage of training. Our general conclusion is that the SSN can be used to incorporate certain modern neuroscience results into highly functional vision systems. This will facilitate future work in exploring the role of lateral connections in a variety of visual processes.

2. The stabilized supralinear network

The stabilized supralinear network, or SSN (Rubin et al., 2015), is a simple recurrent mechanism that accounts for a wide variety of non-linear behaviour without requiring fine tuning of model parameters. It is based on the following circuit properties:

Supralinear (power law) neuron output. All units in the SSN output a power-law function of their input. This is supported by the findings of Priebe, Mechler, Carandini, and Ferster (2004), who found that V1 neurons have a power-law relationship between average membrane potential and firing rate (see also Carandini, 2004; Miller & Troyer, 2002; Priebe & Ferster, 2008).

Recurrent connections with feedback. Neurons in the SSN are sparsely and stochastically connected to one another. Combining the supralinear input–output function of the units with intracortical excitation results in a rapid and indefinite rise in spike rate (for strong enough inputs). Consistent with previous work (e.g. Mariño et al., 2005; Murphy & Miller, 2009), the SSN proposes that this instability is counteracted by simultaneous excitation of inhibitory cells, which then inhibit the activity of both the excitatory and inhibitory units to a point of network stability.

Specific spatial properties of connections between neurons. In the SSN, long-range connections are primarily excitatory, and prefer to connect to inhibitory units. Inhibitory units connect to both excitatory and inhibitory units over shorter distances. This connectivity is consistent with anatomical data (Gilbert & Wiesel, 1989) and is critical to some of the SSN's behaviour.

The model is meant to account for horizontal interactions in layer 2/3. Each unit approximates a group of complex cells that are selective for stimulus orientation. Input to the network is a mathematical function of orientation and retinotopic position—the SSN is not directly pixel-computable.

The SSN consists of a group of excitatory units, E, and a group of inhibitory units, I. Each excitatory unit has a spatial location (either in a ring or a two-dimensional grid, in different versions of the model). Each unit also has an orientation preference. In the two-dimensional version, a unit at two-dimensional grid position x receives feedforward drive $ch(\theta(x))$, where c is stimulus contrast, h is an orientation-tuning curve and $\theta(x)$ is an orientation stimulus at grid position x. There are both excitatory and inhibitory units at each grid position.

The SSN model equations are as follows. Let x be a 2D vector representing the retinotopic position in cortex of an E/I pair of units, and c be the stimulus strength, which is taken to increase monotonically with contrast. W_{ab} is the connection weight from a unit of type b to type a, where a and b can be either excitatory (E) or inhibitory (I). $h(x)$ is the orientation-selective feedforward drive, with range 0 to 1, defined over all cortical positions in the network's grid. Input I to the excitatory and inhibitory units is as follows:

$$I_E(x) = ch(x) + \sum_{x'}(W_{EE}(x,x')r_E(x') + W_{EI}(x,x')r_I(x')) \tag{1}$$

$$I_I(x) = ch(x) + \sum_{x'}(W_{IE}(x,x')r_E(x') + W_{II}(x,x')r_I(x')) \tag{2}$$

where the sum over x' ranges over all other unit positions. $I_E(x)$ is the input to excitatory units, and $I_I(x)$ is the input to inhibitory units at position x. $r_E(x)$ and $r_I(x)$ are firing rates of E and I units at position x.

The steady-state firing of a neuron for fixed input I is,

$$r_E^{SS}(x) = k([I_E(x)]_+)^n \tag{3}$$

$$r_I^{SS}(x) = k([I_I(x)]_+)^n \tag{4}$$

where k is a constant gain, n is an exponent greater than 1, and $[\,]_+$ indicates positive rectification.

Each unit's output spike rate approaches the steady-state value with first-order dynamics,

$$\tau_E \frac{dr_E(x,t)}{dt} = -r_E(x,t) + r_E^{SS}(x,t) \tag{5}$$

$$\tau_I \frac{dr_I(x,t)}{dt} = -r_I(x,t) + r_I^{SS}(x,t). \tag{6}$$

Eqs. (5) and (6) are then solved using Euler's method with a specified time discretization Δt and total steps T.

Connection weights W_{ab} describe the connections between E and I units at each position x. The connections are sparse and computed probabilistically as the product of two Gaussians, one a function of position and the other of orientation preference. Specifically,

$$P(W_{ab}(x,x') \neq 0) = \kappa_b G_{\sigma_{ab}}(x,x') G_{\sigma_{ori}}(\theta(x), \theta(x')) \tag{7}$$

where $\theta(x)$ is the orientation preference, each $G_\sigma(x,x') = e^{\frac{-|x-x'|^2}{2\sigma^2}}$ is a Gaussian function, σ_{ab} and σ_{ori} determine how connection density varies with differences in spatial position and orientation preference, and κ is a scale factor (greater for I connections than E connections). Thus a neuron is more likely to connect to nearby neurons, and also more likely to connect to neurons with similar orientation preference. Rubin et al. (2015) used a 75 × 75-point grid. Their orientation map was generated using a method by Kaschube et al. (2010). Non-zero weights were sampled from normal distributions,

$$W_{ab}(x,x') = N(J_{ab}, (0.25J_{ab})^2) \tag{8}$$

where J_{ab} is the mean strength of connection weight from units of type b to units of type a. Negative weights are set to zero. Using these equations as well as parameter values from Rubin et al. (2015), we confirmed their major simulation results (not shown). The parameter values are reproduced in Table 1.

2.1. Behaviour of the SSN

A major strength of the SSN model is its account for surround interactions. Surround suppression, wherein stimuli outside a neuron's CRF suppress its response to stimuli within the CRF, has been observed in cats and monkeys. When a stimulus extends outside the classical receptive field of a neuron, it stimulates surround neurons that have inhibitory connections to the central