Belief Propagation and Wiring Length Optimization as Organizing Principles for Cortical Microcircuits.

Dileep George
Electrical Engineering, Stanford University &
Numenta Inc., Menlo Park, CA 94025
dil@stanford.edu

Jeff Hawkins
Redwood Neuroscience Institute &
Numenta Inc., Menlo Park, CA 94025
jhawkins@numenta.com

Abstract

In this paper we explore how functional and anatomical constraints and resource optimization could be combined to obtain a canonical cortical micro-circuit and an explanation for its laminar organization. We start with the assumption that cortical regions are involved in Bayesian Belief Propagation. This imposes a set of constraints on the type of neurons and the connection patterns between neurons in that region. In addition there are anatomical constraints that a region has to adhere to. There are several different configurations of neurons consistent with both these constraints. Among all such configurations, it is reasonable to expect that Nature has chosen the configuration with the minimum wiring length. We cast the problem of finding the optimum configuration as a combinatorial optimization problem. A near-optimal solution to this problem matched anatomical and physiological data. As the result of this investigation, we propose a canonical cortical micro-circuit that will support Bayesian Belief Propagation computation and whose laminar organization is near optimal in its wiring length. We describe how the details of this circuit match many of the anatomical and physiological findings and discuss the implications of these results to experimenters and theorists.

1 Introduction

Perceptual systems have to deal with uncertain information in the world. Thus Bayesian techniques have come to be widely viewed as learning and inference mechanisms employed by the cortex. Bayesian Belief Propagation (BBP) introduced by Pearl [6] is among the most successful inference algorithms in computer vision and machine learning. In [5] Lee and Mumford suggest that cortical regions could actually be doing BBP computations, without giving details of the required mechanisms. Recent work by Rao [7] and Deneve [2] show that Bayesian Belief Propagation can be implemented in spiking neurons. They did not investigate an anatomical connection and treated single neurons as the BBP computation engine there by restricting them to encode binary states. What are the neural and anatomical substrates of the Bayesian computations employed by neo-cortex?

The neo-cortex in mammals is believed by many to have a surprisingly prototypical architecture that remains consistent across different species [3]. In all the examined species, the neurons in the cortical sheet are organized in to 6 layers, with the top layer mostly filled with axons [9]. Several researchers have proposed canonical cortical circuits [3] that are replicated all over cortex. Is there a canonical cortical circuit for Bayesian Inference? If yes, is that circuit related to the prototypical laminar organization of the cortex?

Many researchers have explored the role of wiring length minimization in the organization of neocortex. [10] [1] [8]. The positioning of cortical regions in 3 dimensional space obtained as a result of wiring length optimization matched the positioning of cortical areas on the cortical surface [1]. Does the laminar organization of neurons within the cortical regions, also follow from such a principle?

In this paper we investigate these questions by combining the principles of BBP computations with anatomical constraints and wiring length optimization. The requirement that a cortical region should implement Bayesian Belief Propagation sets a set of constraints on the type of neurons and the connections patterns between neurons in that region (Section 2). Moreover there are anatomical constraints that a region has to adhere to (Section 3). There are several different configurations of neurons consistent with both these constraints (Sections 3). Among all such configurations, it is reasonable to expect that Nature has chosen the configuration with the minimum wiring length. We show how to calculate the wiring lengths for these configurations and explore the solution space (Section 4). As a result of this investigation, we propose a canonical cortical micro-circuit that will support BBP computation and whose laminar organization is near optimal in its wiring length. We discuss how the details of this circuit match many of the anatomical and physiological findings (Section 5). These results have several anatomical and physiological implications (Section 6).

2 Bayesian Belief Propagation, Cell Types and Connections

In this section we describe the assumptions involved in the mapping of a Bayesian Network to the cortical Hierarchy. Every region of the cortex can be thought of as maintaining a set of hypotheses in relation to the concepts encoded by its surrounding regions. The hypotheses at a higher region in the cortex are causally linked to the ones in the lower level. The set of hypotheses encoded by a region can be considered a random variable, with cortical columns encoding its particular values. Each region maintains the association of its hypotheses with the causes in a probability table. Observed information anywhere in the cortex can alter the probability values associated with hypotheses maintained every where else. This is done through Bayesian Belief Propagation. In general, the networks can have loops, and we assume that the inference is done through loopy BP. This does not affect our results.

With these broad assumptions, the inputs and outputs of a region of cortex can be mapped to Belief Propagation messages. A cortical region receives input messages from regions hierarchically above and below it, through feed-forward and feedback connections. The role of the cortical region is to update its Belief based on these messages and to derive the messages to be sent to its parents and children using outgoing connections. These computations are performed using the Bayesian Belief Propagation equations shown below. These equations were adapted from [6] and are for singly connected tree structured networks. We assume this type of topology for the rest of this paper.

$$\lambda(x_k) = \prod_j \lambda_{Y_j}(x_k) \tag{1}$$

$$\lambda_X(u_m) = \sum_x \lambda(x) P(x|u_m) \tag{2}$$

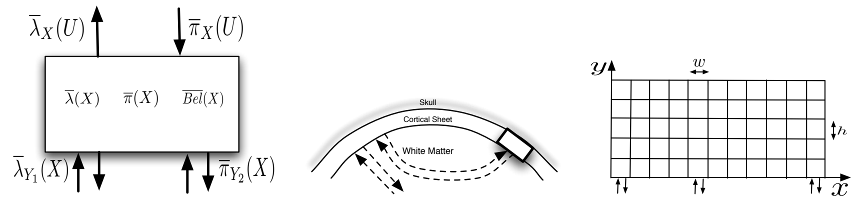

(a) **(b)** **(c)**

Figure 1: (a) Inputs and outputs of a node in Bayesian Belief Propagation. (b) Cortical sheet and its position with restpect to white matter and the skull (c) Idealization of a slice of a cortical region corresponding to the rectangle in (b). At the top is the skull and at the bottom the white matter. Different cortical layers are oriented horizontally. Each square in the grid can hold one neuron(cell).

$$\pi(x_k) = \sum_u P(x_k|u)\pi_X(u) \tag{3}$$

$$BEL(x_k) = \alpha\lambda(x_k)\pi(x_k) \tag{4}$$

$$\pi_{Y_j}(x_k) = \alpha\pi(x_k)\prod_{i\neq j}\lambda_{Y_i}(x_k) \tag{5}$$

These equations are described with respect to a region which encodes X with 2 child regions encoding Y_1 and Y_2 and a parent region encoding U, as illustrated in figure 1(a). The feed-forward input axons and feed-forward output axons carry the λ messages and the feedback axons carry the π messages.

Implementing the Bayesian Belief Propagation Equations in a cortical region will require a diverse set of neurons with different formats of connections. We postulate the existence of 5 types of cells for the implementation of the 5 equations given above. Cell type 1 C_1 is the recepient of the feedforward messages from child regions. A set of such cells do the operation defined in equation 1. This is illustrated in figure 2(a). These cells multiply their inputs together and they do not have weights associated with their synapses.

C_2 is a set of cells which implement equation 2. These cells receive inputs from the C_1 cells defined above. The synapses of these cells implement a sum-product operation. The synapse from cell $C_{1,i}$ to cell $C_{2,j}$ stores the weight $P(x_i|u_j)$. This is illustrated in figure 2(b). C_2 cells send their feed-forward outputs to higher level cortical regions via the message λ_X.

The inputs to C_3 cells are the feed-back messages from higher level cortical areas. These messages are converted to the language of the local cortical region according to equation 3. Similar to C_2 cells, these cells also have synaptic weights. The synapse between cell $C_{3,i}$ and the feedback axon u_j stores the weight $P(x_i|u_j)$. The outputs of cells C_3 are processed internal to the region.

The Belief of a region, according to equation 4 are calculated at the output of cells C_4. These cells receive their inputs from cells C_1 and cells C_3 and combine them multiplicatively to obtain the Belief value according to equation 4. This is illustrated in figure 2(c).

Cells of type 5, C_5 project to child regions and carry the feedback messages to those regions. According to equation 5, the feedback message is specific to child regions. Hence there will be as many type groups of C_5 cells as there are child regions. These cells combine a selected portion of feed-forward messages along with the outputs of cells C_3 to form feedback messages.

As described above, implementing the Belief Propagation Equations in a neuronal circuit

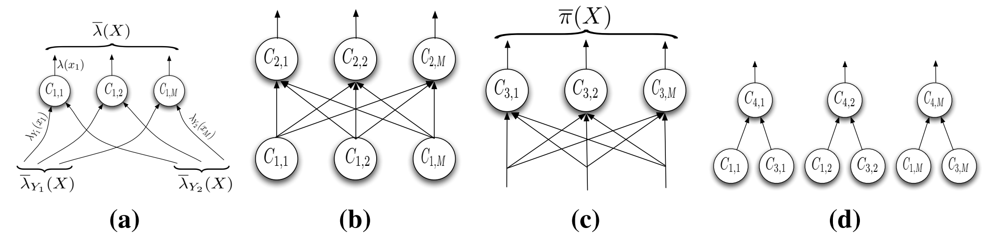

Figure 2: (a) to (d) show the different neuron types and the connections between them arising out of equations 1 to 4. These connection constraints can be encoded as a matrix B (see text).

automatically imposes a constraint on the connectivity between the elements of the circuit. This connectivity constraint can be expressed in the form of a connection matrix B with the neuronal elements and the input and output axon labels along the rows and columns of this matrix.

3 Spatial arrangement of cells within a cortical region

Cortical regions are arranged on a thin sheet covering the white matter. On one side of a cortical sheet is the skull and on the other side the white matter. Hence all the input (output) axons of a cortical region enter(exit) the cortical region via the white-matter side (figure 1b). We can approximate a slice of a region of cortex as a rectangular box with all the inputs and outputs interfacing at the bottom of the rectangle. What would be the best way to arrange the cells within this slice of cortex?

Note that we can achieve the functionality of Bayesian Belief Propagation as long as we maintain the correct connectivity among the neurons. However, different spatial arrangements of these neurons will use different wiring lengths to maintain this connectivity. It is reasonable to assume that among all configurations with the same functionality, Nature would choose the one with the minimum requirement of resources. This enables us to cast the problem of placing the neurons within a cortical region as an optimization problem.

$$\text{Minimize} \quad \sum_{(i,j)\in B} ||x_i - x_j||_1 \tag{6}$$

$$\text{subject to} \quad \text{no overlap between cells} \tag{7}$$

where x_i and x_j correspond to the spatial locations of terminals and B is the connection constraints matrix. Although the objective in this problem is convex, solving this with problem with non-overlap constraints involve a combinatorial optimization.

We can use known facts about cortical organization to reduce the complexity of this problem. The vertical dimension of the cortical rectangle is only a few layers deep . The horizontal dimension is variable. The number of states that a cortical region will have to represent is typically much more than the number of cells that can be accommodated along the vertical dimension of the cortical region. Thus it is reasonable to assume that the states of the region are represented by neurons along the horizontal dimension of the cortical region. We thus divide the horizontal dimension of the cortical region into a number of compartments. We make the simplifying assumption that each compartment corresponds to a particular state of the region.Note that this arrangement corresponds to a columnar organization of the cortex as has been observed using several anatomical and physiological experiments.

This leaves the vertical dimension for cells to support the Belief Propagation operations related to various states of the region.Since we have five different equations and types of cells as given by the Belief Propagation Equations, we divide the vertical dimension of the cortical region into five compartments. This gives us a grid over which we can place cells. Each rectangle in the grid can accommodate one cell.

With no constraints on placement we have $M!(5!^M)$ different arrangements of the cells within a grid on the idealized cortical region with M states. Knowing that the labeling of the states are arbitrary, we can reduce them to number of meaningful arrangements equal to $(5!)^M$. From the pattern of connections illustrated in figure 2 and from the fact that we have exactly 5 compartments in the vertical direction, we can conclude that the optimal solution will have the same type of cell in any particular row of the grid. This insight, combined with the columnar organization constraint helps us to reduce the number of search points from $5!^M$ to $5! = 120$. Thus the approximate optimization problem can be solved by exhaustive search on these 120 possible configurations.

4 Length Function

For each of these configurations, we calculated the length function as follows. In the equations below, we let the symbol for a neuron type to mean its position within the configuration in terms of the number of grid positions from the lower edge of the rectangle (figure 1(c)). Let M be the number of states of the region. We assume that the parent region also has the same number of states. $N_c h$ is the number of children and h and w are the height and with of a grid position. In the calculations below, we assume that the axons branch so as to achieve the least cost wiring. Thus, the equations derived here depend on the order in which different neuron types are placed on the grid. (But the total wiring length does not). We assume that the cells are placed in the order we describe here.

The total length of axons and dendrites required for taking the feedforward inputs from child regions to obtain $\overline{\lambda}(X)$ is calculated as

$$l_1 = M N_{ch} h (C_1 - 0.5) \tag{8}$$

Calculation of the feedforward messages to be sent out to a higher level region requires taking the outputs of C_1 and operating them on according to equation 2 and figure 2 and then sending the outputs to the bottom of the cortical region to be sent out to higher level regions. This gives the total length of axons and dendrites required for this operation as

$$l_2 = M(h\,|C_1 - C_2| + \sum_{i=1}^{M-1}(i-1)w + 2Mh(C_2 - 0.5) \tag{9}$$

Calculating the internal values $\overline{\pi}(X)$ involves taking the feedback messages from a higher level and operating on it according to equation 3. The required length for this operation can be calculated as

$$l_3 = \begin{cases} 2Mh\,|C_3 - C_2| + M\sum_{i=1}^{M-1}(i-1)w & \text{if } C_3 > C_2 > C_1 \\ Mh\,|C_3 - C_1| + Mh\,|C_3 - C_2| + M\sum_{i=1}^{M-1}(i-1)w & \text{if } C_3 > C_1 > C_2 \\ M\sum_{i=1}^{M-1}(i-1)w & \text{if } C_2 > C_3 > C_1 \\ Mh\,|C_3 - C_1| + M\sum_{i=1}^{M-1}(i-1)w & \text{if } C_2 > C_1 > C_3 \\ Mh\,|C_3 - C_2| + M\sum_{i=1}^{M-1}(i-1)w & \text{if } C_1 > C_2 > C_3 \\ Mh\,|C_3 - C_2| + M\sum_{i=1}^{M-1}(i-1)w & \text{if } C_1 > C_3 > C_2 \end{cases} \tag{10}$$

The Belief states of a region are calculated according to equation 4. This requires taking the outputs of cells C_3 and cells C_1 and multiplying them element-wise in cells C_4. The

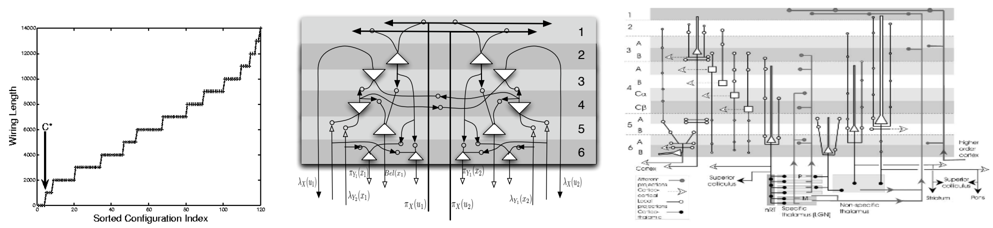

(a) **(b)** **(c)**

Figure 3: (a) Plot of wiring length Vs configurations measured with the optimum configuration at zero base-line . The position of C^* (see text) is marked . (b) The laminar arrangement of neurons and the connections between them corresponding to the C^* configuration. Shown are two cortical columns. The location of cells in this configuration and the connections between them match anatomical data. (c) Laminar organization of cortical micro-circuits - anatomical data adapted from [9] (permission pending). This is included here for the purpose of comparison.

total length of wires required for this operation is

$$l_4 = \begin{cases} 2Mh\,|C_4 - C_3| + Mh(C_4 - 0.5) & \text{if } C_3 \in (C_1, C_4) \\ Mh\,|C_4 - C_1| + Mh\,|C_4 - C_3| + Mh(C_4 - 0.5) & \text{if } C_1 \in (C_3, C_4) \\ Mh\,|C_4 - C_3| + Mh(C_4 - 0.5) & \text{if } C_4 \in (C_1, C_3) \end{cases} \quad (11)$$

Calculation for l_5 is similar, but involves enumeration of 24 different cases.

Finally, the total length of connections for a particular configuration is calculated as

$$L = l_1 + l_2 + l_3 + l_4 + l_5 \quad (12)$$

This is the objective function we attempt to minimize over all configurations.

5 Results: Near-optimal solution Matches Anatomical and Physiological Data

In order to find how the wiring length varies as a function of spatial arrangement of the neurons, we evaluated the objective function described above at all the 120 configurations. We sorted these configurations based on their wiring lengths and examined the configurations starting with the ones with the least wiring lengths. We found that a configuration with near-optimal wiring length matched the anatomical data to a great extent. This configuration, denoted henceforth by C^*, is the second best in wiring length among all wiring lengths, and is only slightly worse (10%) than the best solution when measured as a fraction of the difference between the best and the worst. The spatial arrangement of cells and the anatomical connections resulting from this configuration is shown in figure 3 (b). The position of this configuration among other configurations is shown a wiring-length Vs configurations plot in figure 3(a).

In C^*, the feed-forward inputs cortical regions in a lower level of the hierarchy rise to the layer-4 cells. Layer 4 cells then project to layer 2 cells and also send an axon to layer 5 cells. The feedback axons coming from higher level cortical areas rise to layer 1 and spread laterally. Layer 3 neurons with synapses in layer 1 are the targets of this feedback. These layer 3 neurons project to layer 5 and layer 6. The layer 5 neurons project down subcortically and the layer 6 neurons are the source of feedback to cortical areas hierarchically below. The feed-forward axons that project to layer 4 also synapse in layer 6 to perform the computations in equation 5. These details map on to the anatomical data obtained from [9]. In figure 3 (c) we compare this configuration with a schematic of anatomical data adapted from [9].

The results of the wiring length minimization were independent of the number of states of the region and the number of children. There were three optimal solutions and these configurations shared many properties of the near-optimal solution we chose here for its conformance to anatomical data. These solutions differ from C^* by exactly one interchange operation.

6 Implications

The mapping of Bayesian Belief Propagation on to the cortical microcircuit is of potential implications to experimenters. This framework, can help understand and guide physiological experiments. In this section we explore some of the predictions of this mapping and their potential implications.

Anatomical data describe a class of layer 5 neurons that project to sub-cortical areas [9]. Our results implicate that these neurons carry the current Belief of a region. The current Belief of a region could be be used for making actions or decisions. Emotionally relevant beliefs could be stored for later recall. There are several sub-cortical modules that could make use of the current Belief states of cortical regions. Although the details of the sub-cortical projections are not known, the prediction that these projections carry the Beliefs of a region is potentially significant.

In C^* the cortical columns correspond to states of a cortical region. Typically, these cells are divided into two categories- simple, and complex. However, this mapping tells us that a more sophisticated explanation still consistent with the simple/complex mapping is possible. Layer 4 cells are consistently characterized as simple cells, and layer 2 cells are consistently characterized as complex cells because they pool information from different layer 4 cells. However, layer 3 cells, layer 5 cells and layer 6 cells that are normally characterized as complex cells have a richer meaning. They combine contextual information from higher levels with local information. If higher level context is ignored, and the receptive fields of these cells are mapped using a pure feed-forward technique, they will correspond to the complex cells characterization. However, when contextual effects are taken into account, these cells will have a more sophisticated meaning. For example, In a related study [4], we showed how to explain illusory contours effect [5] and end-stopping effect using Belief Propagation. Results from the current study show that the illusory contour cells and end-stopping cells will be prevalent in layers 2-3 and layer 5. This is consistent with experimental results [5].

This mapping also provides another way to interpret population coding. It is known that cortical columns show a graded response to stimuli. This finding has been largely interpreted as a coarse coding mechanism. C^* too predicts a graded response to stimuli [4]. However in this setting the graded responses correspond to the measure by which the stimuli is likely to belong to the different states of a cortical region. This is applicable at all levels of the cortical hierarchy.

7 Discussion

We derived a cortical micro circuit and its layout within a laminar cortical architecture based on the principles of Bayesian Belief Propagation and wiring length optimization. The discover of a near-optimal solution that matches anatomical data is an encouraging development. Several reasons can be cite for the sub-optimality of the solution. The major reason is our ignorance of the exact constraints and objectives that are involved in cortical organization.

We took into account only the role of excitatory neurons and connections in this study. We think that inhibitory neurons play a very significant role in cortical computations. However, we think that these roles are more in terms of keeping a good operating point for the computationally relevant circuits. It is known that normalizing the messages and intermediate values in BBP is required for numerical stability. Such normalization computations would require inhibitory circuits. We think that inhibitory neurons play a significant role during learning as well. We are currently investigating how to include these as part of the optimization. Missing out the contributions from inhibitory neurons could also be one reason for the sub-optimality of our solution.

Deciphering the functional connectivity of the cortical micro-circuit is a formidable task. Several insights can be drawn by comparing it to reverse-engineering an electronic circuit. Although a single transistor can function as an amplifier, a good amplifier is seldom constructed from a single transistor. A good construction involves a biasing circuitry which makes sure that the amplifier works properly despite changing temperature conditions, different device characteristics, feedback instabilities etc. Its reasonable to expect that a similar situation exists within the cortical sheet where a multitude of neurons are involved in biasing a canonical cortical circuit to function. If the circuit is tested for connectivity when it is not properly biased, one would end up missing some important connections and logging some spurious connections. Hence, deciphering the functional connectivity from an increasing amount of anatomical data would require theories about cortical functions and how they map on to anatomy. We believe that our work is a contribution in that direction.

References

[1] C. Cherniak, Z. Mokhtarzada, R. Rodriguez-Esteban, and K. Changizi. Global optimization of cerebral cortex layout. *Proceedings of the National Academy of Sciences of the United States of America*, 101(4):1081–1086, Jan 27 2004. LR: 20041117; DEP: 20040113; JID: 7505876; 2004/01/13 [aheadofprint]; ppublish.

[2] Sophie Deneve. Bayesian inference in spiking neurons. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 353–360. MIT Press, Cambridge, MA, 2005.

[3] R. J. Douglas and K. A. Martin. Neuronal circuits of the neocortex. *Annual Review of Neuroscience*, 27:419–451, 2004. LR: 20041117; JID: 7804039; RF: 176; ppublish.

[4] Dileep George and Jeff Hawkins. A hierarchical bayesian model of invariant pattern recognition in the visual cortex. In *Proceedings of the International Joint Conference on Neural Networks 2005*, page in press. IEEE Press, 2005.

[5] Tai Sing Lee and David Mumford. Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis*, 20(7):1434–1448, Jul 2003.

[6] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman Publishers, San Francisco, California, 1988.

[7] Rajesh P. N. Rao. Hierarchical bayesian inference in networks of spiking neurons. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 1113–1120. MIT Press, Cambridge, MA, 2005.

[8] O. Sporns, G. Tononi, and G. M. Edelman. Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cerebral cortex (New York, N.Y. : 1991)*, 10(2):127–141, Feb 2000. LR: 20041117; JID: 9110718; ppublish.

[9] Alex M. Thomson and A. Peter Bannister. Interlaminar connections in the neocortex. *Cerebral cortex (New York, N.Y. : 1991)*, 13(1):5–14, 2003.

[10] M. P. Young. Objective analysis of the topological organization of the primate cortical visual system. *Nature*, 358(6382):152–155, Jul 9 1992. LR: 20041117; JID: 0410462; CIN: Nature. 1994 Jun 9;369(6480):448-50. PMID: 8202134; ppublish.