Coordinate-free representation of sensorimotor spaces

Francesco Frisone, Pietro G. Morasso, Vittorio Sanguineti

Department of Informatics, Systems and Telecommunications
University of Genova, Via Opera Pia 13, 16145 Genova, Italy
E: friso@dist.unige.it F: +39 10 3532154

Abstract

The paper overviews a computational framework for characterizing the cortical representations and processes which underly the kinematic invariances of movements, taking into account the new understanding of the cortex as a continuously adapting dynamical system, shaped by competitive and cooperative lateral connections. We show how a coordinate-free representation of sensorimotor spaces can emerge from self-organized learning which builds a topological representing structure. This implies a mixture of local and long-range lateral connections, consistent with known anatomical facts, thus allowing the representation of high-dimensional spaces in an apparently flat anatomy. The dynamics of cortical maps is analyzed taking into account the excitatory nature of the majority of cortical synapses and the puzzling presence of long-range competition without long-range inhibition. A model is proposed which combines a process of diffusion (via the excitatory topologically organized connections) and a process of competitive distribution of activation which tends to sharpen the active map region.

1 Lateral interactions in the cortex

It is commonly assumed that the kinematic invariances of movements [8] are somehow related to the organization of sensorimotor cortical areas. The plasticity of such areas, revealed for example by experiments of sensory deprivation in young animals, traditionally has been attributed to adaptable thalamo-cortical synapses, underestimating or neglecting the role of lateral connections. However, a new understanding of the cortex is emerging as a continuously adapting dynamical system, shaped by competitive and cooperative lateral connections. Among the exciting recent data about intracortical connectivity and dynamic processes in the brain, we may quote the evidence that lateral connections seem to be *topology preserving* in the sense that cross-connected cortical columns tend to have similar properties [4, 7] and such organization is not static but changes with ontogenetic development together with patterns of thalamocortical connections [5, 1]. Shortly, it has been suggested that cortical areas can be seen as a massively interconnected set of elementary processing elements, i.e. as a computational map [6].

The cortex is a thin layer (2 mm thick) of neurons arranged on a surface of a little bit more than 2000 cm^2 with a density of about 75000 neurons/mm^3. It appears to be structured in (mini)columns, as revealed by Nissl stain which only colors cell bodies thus showing cylindrical clusters (with a diameter of about 30 μm and a number of neurons of the order of 100) separated by gaps of vertically-running bundles of axons and dendrites. Several hundred millions of columns cover the cortex, arranged in a number of functionally identifiable larger cortical modules. The structure of lateral connections is not genetically determined but, indeed, depends mostly on electrical activity during development. More precisely, they have been observed to grow exuberantly after birth and reach their full extent within a short period; during the subsequent development, a pruning process takes place so that the mature cortex is characterized by a well defined pattern of connectivity.

Many models of brain cortex (wrongly) assume that lateral connections are *local* and *radially symmetric*, with a strength which varies according to a mexican-hat profile. In general, the assumption of locality of lateral connections implies that, from the point of view of circuitry, the cortex is basically a 2-D lattice. However, the observed patterns of lateral connectivity are much more complicated and include a large amount of non-local connections, thus ruling out all the models limited to a purely 2-D circuitry. It can be argued that non-local lateral connections might simply follow pseudo-random patterns, without any topological organization, but our claim is that, even if this were true at birth, it would be inevitably changed during development by the interaction with a structured environment - a typically *ecological* argument. In other words, our basic assumption is that Hebbian learning applied to an arbitrary initial connectivity has the power to *carve* in the chaotic jungle of connections a structure with a robust topological order and this is the essence of the self-organization of sensorimotor neocortex.

The greatest number of neurons in the cortex (pyramidal and stellate) have excitatory synapses. Recurrent excitation is especially prominent in the superficial layers of the neocortex: in fact, superficial pyramids send myelinated axons to distant columns (both in superficial and middle layers) as well as unmyelinated collaterals to neighboring pyramids.

The superficial connections to non-neighboring columns are organized into characteristic patterns: a collateral of a pyramidal axon typically travels a *characteristic lateral distance* without giving off terminal branches and then it produces tightly terminal clusters (possibly repeating the process several times over a total distance of several millimiters). Such characteristic distance is not a universal cortical parameter and is not distributed in a purely random fashion but is different in different cortical areas: 0.43 mm in the primary visual area, 0.65 mm in the secondary visual area, 0.73 mm in the primary somatosensory cortex, 0.85 mm in the primary motor cortex, and up to several mm. in the infero-temporal cortex [2]. Horizontal connections also occur in the deep layers but without the regular spacing pattern. The formation of the network of lateral connections is known to be characterized by a process of growth & pruning immediately after birth [5] and the process is likely to be activity-dependent, as suggested by experiments on sensory deprivation. Thus, the development of lateral and afferent connections depends on the cortical activity caused by the external inflow, in such a way to capture and represent the (hidden) correlations in the input channels. Each individual lateral connection is "weak" enough to go virtually unnoticed while mapping the receptive fields of cortical neurons but the total effect on the overall dynamics of cortical maps can be substantial, as is revealed by cross-correlation studies. Although a possible traditional interpretation of the pattern of characteristic distances is simply that hierarchically "higher" cortical areas have "longer" lateral connections and "larger" receptive fields, we propose a more radical alternative according to which the observed pattern is a side-effect of the need to allocate m-D grids onto a 2-D substrate and areas with longer lateral connections represent "higher-dimensional" spaces.

Lateral connections from superficial pyramids tend to be recurrent and excitatory [9] and recurrent excitation is likely to be the underlying mechanism which produces the synchronized firing which has been observed in distant columns [12]. The existence (and preponderance) of massive recurrent excitation in the cortex is in contrast with what could be expected, at least in primary sensory areas, considering the ubiquitous presence of peristimulus competition (or "mexican-hat pattern") which has been observed time ago in many pathways and has been confirmed by direct excitation of cortical areas. In other words, in the cortex there is a significantly larger amount of long-range inhibition than expected from the density of inhibitory synapses. In general, "recurrent competition" has been assumed as synonimous of "recurrent inhibition", for providing an antagonistic organization that sharpens responsiveness to an area smaller than would be predicted from the anatomical funneling of inputs. Thus, a still open question is how long-range competition can arise without long-range inhibition and this has attracted the attention of several people [10, 11]. But apart from the question of *how* this is obtained, we also must address the question of *why* it emerged: this is necessary in general for any computational theory, as argumented by David Marr, but is particularly significant in our case because there is no apparently cogent reason for such paucity of cortical inhibitory synapses, except a computational one.

The answer to the *why* question comes natural in the framework of a theory of field computation in the cortex: if the basic computational paradigm is the propagation/manipulation/interaction of fields then one needs lateral excitatory connections and stabilizing local self-inhibition, exactly in agreement with cortical data, not recurrent inhibition which would freeze the activation patterns. On the other hand, field computation is apparently absent in lower computational layers and thus simple recurrent inhibition is the natural solution for sharpening the data in extra-cortical centers. Neocortical machinery, on the other hand, may have evolved under the pressure of two somehow contradictory requirements: (i) *arborization* of massive excitatory connections, as the basic feature of the field-computing architecture, and (ii) *preservation* of a sufficient degree of peristimulus competition in order to sharpen the response function of individual units [11].

2 Cortical dynamics

Let us now consider a family of simplified dynamical systems which are compatible with many of the facts and notions about cortical dynamics summarized above. For simplicity, we lump the generic i-th cortical column into a single processing element, characterized by an activity level V_i and two kinds of inputs:

$$\frac{dV_i^x}{dt} + \gamma_i V_i^x = f(h_i^{lat} + h_i^{ext}) \tag{1}$$

The equation simply says that V_i evolves under the action of three competing influences: (1) a self-inhibition (weighted by γ_i), (2) a net input h_i^{lat} coming from the set of lateral connections inside the same cortical area, (3) a net input h_i^{ext} determined by thalamo-cortical or cortico-cortical connections across different cortical areas. A non-linearity $f(\ )$ is applied to the total input for generality.

The h_i^{lat} term is intended to express the massive lateral excitatory connections mentioned before and thus can have the following form: $h_i^{lat} = \sum_{j \in \mathcal{N}_i} C_{ij}\ V_j$ where $\mathcal{N}_i$ is the set of columns laterally connected to the given element and the connection weights C_{ij} are positive and symmetric ($C_{ij} = C_{ji}$). This is a recurrent input which triggers the coupled dynamics of the whole network, according to global processes as diffusion, dispersion, convection, etc.[13].

As regards the external input, we may use a similar model $h_i^{ext} = \sum_{k \in \mathcal{M}_i} W_{ik}\ X_k$ where $\mathcal{M}_i$ refers here to the set of input neurons projecting to the given column (e.g. thalamic neurons or output neurons from other cortical areas). These connections are uni-directional, which means that are asymmetric and cannot support recurrent dynamics: as a consequence, what h_i^{ext} expresses is a smoothed version of the pattern of activity present in (a small portion of) the projecting region. Considering the somatotopic organization of such thalamic or cortical regions, if we denote with $\vec{x}$ the sensorimotor vector collectively represented by the given region, then the previous equation is approximately equivalent to the following more synthetic form $h_i^{ext} = G_i(\vec{x})$ with $G_i(\vec{x})$ a positive unimodal function of $\vec{x}$, such as a broad Gaussian, centered around a vector prototype or receptive field center $\vec{\pi}_i^x$.

We assume, in general, that both the input synapses W_{ik} and the intra-cortical synapses C_{ij} are modified according to the classic Hebbian paradigm, starting from an initial highly redundant set of connections. This will yield, automatically, a *robust pruning* of such set of connections. In the case of input synapses, in particular, this process is virtually equivalent to a *smooth tuning* of the receptive field centers $\vec{\pi}_i^x$: $\Delta\vec{\pi}_i^x = \eta(\vec{x} - \vec{\pi}_i^x)U_i$, according to the soft competitive learning paradigm. In any case, the same Hebb's rule serves two different but complementary purposes: (i) to smoothly tune the receptive field centers of the cortical processors and (ii) to build the topological structure of the cortical map. The two connection patterns jointly identify the coordinate-free representation of the sensorimotor vector $\vec{x}$.

2.1 Combining diffusion and competitive activation

Let us now consider the cortical dynamic equation (1) and let us forget, for simplicity, the non-linearity $f(\ )$:

$$\frac{dV_i^x}{dt} + \gamma_i V_i^x = \sum_j C_{ij}\ V_j + G_i(\vec{x}). \tag{2}$$

We can observe that this is a diffusion equation which spreads the initial Gaussian pattern, progressively flattening it. If we wish to use an electrical analogy, the equation above is equivalent to a network of leaky integrators ($1/\gamma_i$ is the time-constant of each integrator) with a resistive lattice structure (C_{ij} are are lateral conductances) driven by a set of current sources $\{G_i(\vec{x}), i = 1, 2, \ldots\}$, modulated by a common input signal. This kind of dynamical behavior, although important from the point of view of field computation, is not particularly interesting per se because it lacks the necessary capability of sharpening the edges of the broad input pattern $G_i(\vec{x})$. Therefore, we need a mechanism which combines diffusion with sharpening and, in particular, can obtain the latter without using recurrent inhibition: *competitive distribution of activation* can be the answer.

Originally it has been proposed by Reggia et al. [10], as a class of simple mechanisms which are based on the following idea for achieving competition without inhibition: it is assumed that any columnar processing element has a finite amount of excitation to distribute at any moment to the set of columns which are the targets of its axon collaterals and such distribution is competitive in the sense that the more active is a column the larger is its excitation share. In principle, there is no reason for arguing that this kind of operation might not be carried out by the columnar circuitry. Thus, if a column is very active it will tend to syphon a larger share of excitation than neighboring columns, thus functionally inhibiting them by a sort of of multiplicative (or gating) rather than a subtractive mechanism, which is somehow reminiscent of Grossberg's shunting. From the modeling point of view, this concept can be implemented in a variety of ways, such as the following one:

$$\frac{dV_i^x}{dt} + \gamma_i V_i^x = \sum_{j \in \mathcal{N}_i} C_{ij} \frac{V_j}{\sum_{k \in \mathcal{N}_j} V_k} + c_e G_i(\vec{x}) V_i \tag{3}$$

where the multiplicative action is applied both to the lateral and external inputs. and, for simplicity, there is no global non-linearity $f(\ )$. In this way we can combine in the same non-linear dynamics two complementary and somehow contradictory dynamic effects: sharpening a broad pattern and diffusing a sharp pattern. This is illustrated in figure 1 for the very simple case of a 1-D cortical map of 60 neurons. The input x is a scalar with a variation range uniformly distributed between 0 and 3. The Hebbian learning paradigm mentioned above yields a distribution of receptive field centers visualized as small circles on the abscissas of the top part of the figure. As a consequence of the 1-dimensionality of the input manifold, each neuron develops only two lateral connections, thus acquiring the structure of a 1-D lattice. In the figure, the network is initially at equilibrium for an input distribution $\{G_i(x); i = 1, 2, \ldots | t < t_0\}$ peaking at $x = 0.5$. The dotted pattern at the initial time $\{V_i; i = 1, 2, \ldots | t_0\}$ is a narrow island of activity around the neuron whose receptive field center π_i^x is closest to $x = 0.5$ and this is the consequence of the competitive distribution of activation mentioned above. But what is more interesting is the transient behavior which is determined by the sudden shift of the input variable x from 0.5 to 2.0 which yields the dashed distribution of the external input $\{G_i(x); i = 1, 2, \ldots | t \geq t_0\}$: we may interpret it as the selection of a new target. As can be seen from the figure, which shows the distribution of activity over the map for different time instants, the transient is a combination of a *diffusion process* (which initially flattens the population code spreading the activity pattern over the whole network) and a *re-sharpening process* (which builds up faster and faster as soon as the diffused wavefront reaches the new designated target). The combination of the two processes is the *propagation* of the population code toward the new target: remarkably, the time course of the peak position (bottom part of the figure) is smooth and with a bell-shaped profile, in spite of the sudden/discrete selection of the new target.

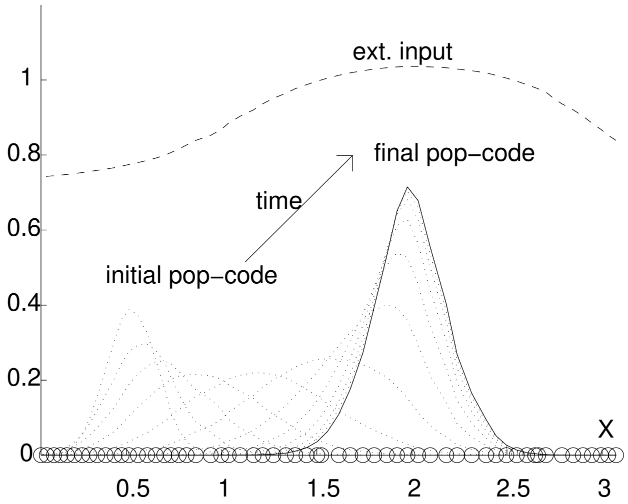

Figure 1: Dynamic behavior of a 1-D cortical map when the input vector suddenly shifts from an initial value ($x = 0.5$) to a final value ($x = 2$).

2.2 Flow-line tracking

In general, competitive distribution of activation can be seen as a mechanism which transforms an initially broad activity pattern into a tightly focused island of activations. Which is the putative relation between such islands of activation and the diffuse patterns revealed by the population coding of Georgopoulos et al. [3]? And why is such focusing computationally necessary? The answer to the former question is that the neural activity could be related to the input element h_i^{ext} or $G_i(\vec{x})$ of the posited columnar mechanism. In practical experimental condition it could be easier to observe it rather than V_i, whose rapid transients would probably be masked by G_i. On the other hand, the sharpening mechanism is absolutely necessary for allowing the map to operate as an effective component in a field-computing architecture.

In the example of figure 1, although trivially, the smooth shift of the population code toward the target can be interpreted as the process of tracking the flow-line of a field which is implicitly associated with the gradient of the interpolated distribution of activity levels $\{V_i, i = 1, 2, \ldots\}$. The same computational process is much more interesting in the case of m-D grids, although is still described by the same equation. The dimensionality is hidden in the pattern of lateral connections and in particular in the size of the set $\mathcal{N}_i$: the larger the dimensionality, the larger the size. Also in an m-D map the population code is an island of activation and the two dynamic processes of diffusion and sharpening evolve in the same qualitative way. In particular, the peak of the propagating waveform will travel in an approximately straight line to the target, coded by the G_i pattern, provided that the lattice of connections is approximately uniform (as regards topology of connections and distribution of connection weights) thus inducing a quasi-isotropic diffusion process. Moreover, what is straight or round in the underlying m-D space might appear to be quite irregular on the cortical 2-D surface as a consequence of the mixture of short and long-range connections and such irregularity is indeed a common and apparently puzzling empirical finding in a variety of bioimaging investigations.

More complex behaviors can occur if the input pattern is not fixed but varies over time: for example, the external input could be projected on the cortical area, via cross-connections, from another cortical area which is undergoing, itself, a diffusion/relaxation process. Since the two areas might represent, in a coordinate-free way, sensorimotor vectors of different dimensionalities, the combined dynamics of two or more cortical areas can yield an extremely powerful computational engine [11]. Moreover, through the external input pattern it is possible to code obstacles as well as targets: a target is identified by a local maximum in the input distribution and an obstacle by a local minimum, for example a negative value. In this case the network dynamics will generate a generally curved trajectory, which reaches the target while avoiding the obstacles. The asymptotic stability of the cortical map dynamics is assured by the fact that, as a consequence of the symmetry of lateral connections, we can see it as a Hopfield network for which an explicit Liapunov function is known.

As a final remark, we wish to observe that although the proposed computational architecture has a direct link with the many forms of self-organized maps, it differs, among other things, for its homogeneous treatment of lateral and input connections. In the original SOM the lateral connections are fixed and in other models like the "neural gas" are learned during training but they are treated in a different way. This is not biologically plausible and the proposed model is a contribution in this direction.

References

[1] A. Burkhalter, K. Bernardo, and V. Charles. Development of local circuits in human visual cortex. *J. of Neuroscience*, 13:1916–1931, 1993.

[2] W. Calvin. Cortical columns, modules, and hebbian cell assembles. In M. Arbib, editor, *The handbook of brain theory and neural networks*, pages 269–272. MIT Press, Cambridge, MA, 1995.

[3] A. Georgopoulos, A. Schwartz, and R. Kettner. Neuronal population coding of a movement direction. *Science*, 233:1416–1419, 1986.

[4] C. Gilbert and T. Wiesel. Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J. of Neuroscience*, 9:2432–2442, 1989.

[5] L. Katz and E. Callaway. Development of local circuits in mammalian visual cortex. *Annual Review of Neuroscience*, 15:31–56, 1992.

[6] E. I. Knudsen, S. du Lac, and S. Esterly. Computational maps in the brain. *Annual Review of Neuroscience*, 10:41–65, 1987.

[7] S. Lowell and W. Singer. Selection of intrinsic horizontal connections in the visual cortex by correlated neural activity. *Science*, 255:209–212, 1992.

[8] P. Morasso. Spatial control of arm movements. *Exp. Brain Res.*, 42:223–227, 1981.

[9] A. Nicoll and C. Blakemore. Patterns of local connectivity in the neocortex. *Neural Computation*, 5:665–680, 1993.

[10] J. Reggia, C. D'Autrechy, G. Sutton III, and M. Weinrich. A competitive distribution theory of neocortical dynamics. *Neural Computation*, 4:287–317, 1992.

[11] V. Sanguineti, P. Morasso, and F. Frisone. Cortical maps of sensorimotor spaces. In P. Morasso and V. Sanguineti, editors, *Self-organization, Computational Maps, and Motor Control*, pages 1–36. Elsevier Science Publishers, Amsterdam, 1997.

[12] W. Singer. Development and plasticity of cortical processing architectures. *Science*, 270:758–764, 1995.

[13] M. Zak. Neurodynamics with spatial self-organization. *Biological Cybernetics*, 65:121–127, 1991.