On the Dimensionality of Cortical Graphs

Elie Bienenstock

Division of Applied Mathematics, Brown University

and CNRS, Paris

email: elie@dam.brown.edu

J. Physiol., Paris, **90**: 251--256 (1996)

Keywords: Cortex, graph theory, dimensionality, synfire chain, development.

Send reprint requests to:

Elie Bienenstock

Division of Applied Mathematics, Brown University, Providence RI 02912 USA

Summary. We propose to use a random-graph model of cortex as a *tabula-rasa* state, to be contrasted with various types of regular connectivity patterns. Key in our analysis is the notion of graph-theoretic dimensionality, closely linked to that of graph diameter. Our discussion focuses on patterns of synfire type, and on the synfire-superposition model proposed in previous papers

1. Introduction

One of the remarkable features of brains of adult mammals is the precision of their connectivity, a system of complex yet to a large extent predictable wiring patterns. From a graph-theoretic perspective, some insight might be gained from viewing connectivity patterns among cortical pyramidal cells as various types of *departures from randomness*. Consider a hypothetical totally random cortex. In such a cortex—no doubt biologically useless—a synapse from one given pyramidal cell to another would be present with a given constant probability $p \ll 1$, independent of the location and function of these two neurons. The connectivity graphs of real brains differ from such a random graph in many ways, some of which may still be largely unknown.

The most conspicuous and currently best understood departures from randomness in the connectivity of cortex are regularities in its three-dimensional anatomical structure. Chief and foremost, roughly half of all pyramidal-pyramidal connections are local; long-range—white-matter—axons connect specific areas with each other according to well-defined anatomical and functional patterns; on a somewhat finer level, long-range connectivity often exhibits patchiness along the cortical surface, and/or maps one area onto another in a topology-preserving way; connections further obey rules related to laminar organization, etc.

The view to be defended here is that cortex may also contain, in addition to these conspicuous regularities, which we shall refer to as *anatomical*, departures from randomness of a rather different type, which would become apparent only when due attention is given to the *high-dimensional* structure of the cortical graph. These regularities would arise from specific developmental mechanisms, and serve specific purposes in the representation of information in the nervous tissue. We shall term these hypothetical regularities *hidden* patterns. A more precise definition of hidden regularities will be given below. For now, suffice it to say that hidden connectivity patterns are of the same nature as anatomical patterns in that they constitute departures from randomness. However, in contrast with anatomical patterns, they do not live in physical 3-D space.

2. Graphs, random graphs, graph-theoretic dimensionality

Before proceeding any further, some clarification is needed as to the sense in which the cortical graph may be viewed as a high-dimensional structure. This is related to the randomness in this graph, and best seen with the help of a few simple mathematical definitions.

Definition. A *graph* G is a pair $G = (V,E)$, where V is a collection of labeled *vertices*—also called *nodes*—and E is a collection of *edges* between non-oriented pairs of vertices in V. The *order* of the graph is the number of vertices: $N = |V|$. The *size* of the graph is the number of edges: $M = |E|$. Thus, $M \leq N(N\text{-}1)/2$. The *degree* of a node x in G is the number of edges attached to x. The *degree* of the graph, which we shall denote δ, is the average degree over all nodes x in the graph: $\delta = 2M/N$.

Remark: In a *non-directed* graph, as defined above, no distinction is made between edge (x,y) and edge (y,x); this edge shall simply be denoted xy. Although directed graphs provide a more appropriate model for brains, we use the simpler concept; for our purposes this makes little difference.

Example 1. $G_1 = (V,E)$, with $V = \{a,b,c,d,e\}$ and $E = \{ab,bc,bd,cd,ae,ce\}$. G_1 is a graph of order $N = 5$, of size $M = 6$, and of degree $\delta = 2.4$.

Example 2. $G_2 = (V,E)$, with $V = \{a,b,c,d\}$ and $E = \{ab,ac,ad,bc,bd,cd\}$. G_2 has order $N = 4$, size $M = 6$, and degree $\delta = 3$. In fact all nodes of G_2 have degree 3, the highest possible in a graph of order 4. G_2 is the *complete* graph on V.

Example 3. G_3 is a regular square lattice in the plane, with each node x connected to its four nearest neighbors, i.e., the nodes sitting above, below, to the left, and to the right of x. The order and size of G_3 may be arbitrarily large; its degree is $\delta = 4$—ignoring boundary effects, which are small if N is large.

Example 4. $G_4 = (V,E)$, with V a collection of Poisson-distributed points in a given domain of the plane, and E the set of all edges xy such that the distance between x and y in the plane is less than 1. The degree of a node in G_4 is a Poisson random variable whose parameter, the expected number of points in a disk of radius 1 minus 1, is equal to δ, the degree of G_4.

Example 5. $G_5 = (V,E)$, with V a given set of N nodes, and E obtained by including a given edge $xy \in V^2$ with a constant probability p, the decision for xy being made independently of all other edges. G_5 is called a *random graph*; its specification entails $N(N-1)/2$ independent random choices. The

degree of a node in G_5 is a random variable; it obeys a binomial distribution of mean $p(N-1)$, equal to δ, the degree of the graph.

Although a graph such as G_4 is random too, the term "random graph" is usually reserved for graphs such as G_5, which could be called completely random. Random graphs have been studied by Erdös and Rényi in the early 60's (their early papers are most easily found in the collection of Erdös papers edited by Spencer, 1973). There now exists a considerable theory investigating the properties of these apparently innocent-looking objects (Bollobás 1985). As noted above, a random graph such as G_5 would not be a very satisfactory model of the brain. However, random graphs might capture some important aspects of the cortical network (see also Valiant 1994), such as its dimensionality. We define the dimensionality of G, for a large enough graph G, as follows.

Definition. Let a and b be two nodes in a graph $G = (V,E)$. The *distance* $d_G(a,b)$, or simply $d(a,b)$, is the length of the shortest path between a and b in G. The *spreading function* of a node a in G, $f_G(k)$, is the number of nodes at distance k or less from a. If, for k small, the spreading function of any node in G is such that $f_G(k) \propto k^r$ for some $r>0$, the graph is said to be of *dimensionality r:* $\dim G = r$.

To illustrate, in G_1, $d(a,b) = 1$ and $d(a,c) = 2$. In G_2, the distance between any two nodes is 1. G_1 and G_2 are too small for the definition of dimensionality to apply. In the "Euclidean graphs" G_3 and G_4, defined through an embedding of V in the Euclidean space $\mathbf{R}^2$, the $\mathbf{R}^2$ metric agrees, in the intuitive sense, with the graph-theoretic one. Assuming that G_3 and G_4 are large, the two notions of dimensionality also agree with each other. In effect, the spreading function $f_G(k)$ is easily seen to be $1+2k(k+1)$ for G_3, and, approximately, δk^2 for G_4. Thus, in both cases, the graph-theoretic dimensionality is 2. More generally, if G is a Euclidean graph in $\mathbf{R}^d$, defined just like G_4 but through an embedding in $d > 2$ dimensions, $f_G(k) \approx \delta k^d$.

If G is a large random graph of G_5 type with edge probability $p \ll 1$, what is the dimensionality of G? To answer this question, we need to consider the growth of the spreading function of a given node a in G, for small k. The average degree of any node in G is $\delta = p(N-1)$, so a is connected, on average, to δ nodes. Each one of these nodes is connected, on average, to $\delta-1$ "new" nodes, each of which is in turn connected to $\delta-1$ other new nodes, and so on. By this reasoning, $f_G(k)$ grows like $\delta+\delta(\delta-1)+\ldots+\delta(\delta-1)^{k-1} \propto \delta^k$. This however is not quite accurate, for some of the nodes we counted as "new" may actually be "old". For instance, there might exist a path of length 2 between a and a given node x, but also a path of length 1.

In spite of this, our estimate is essentially correct: random graphs do spread exponentially, i.e., $f_G(k) \propto \delta^k$ (see e.g. Bollobás 1985, p. 229). Since this is faster than any finite power k^r, G may be said to be, in effect, infinite-dimensional.

We may remark at this point that large natural systems with interactions taking place on a high-dimensional graph are rare. In the absence of a highly specialized communication device, interactions in a physical or biological system generally take place locally in 3-D space. Neocortex might well be an exception, as a considerable fraction of its volume is occupied by insulated wires which carry information in a partly random way over distances three or four orders of magnitude larger than cell-body size.

How fast does the cortical spreading function $f_{\text{cortex}}(k)$ grow as a function of k? Is there enough randomness in G_{cortex}, the connectivity graph of cortical pyramidal cells, to make its dimensionality infinite, or at least substantially higher than 2 or 3?

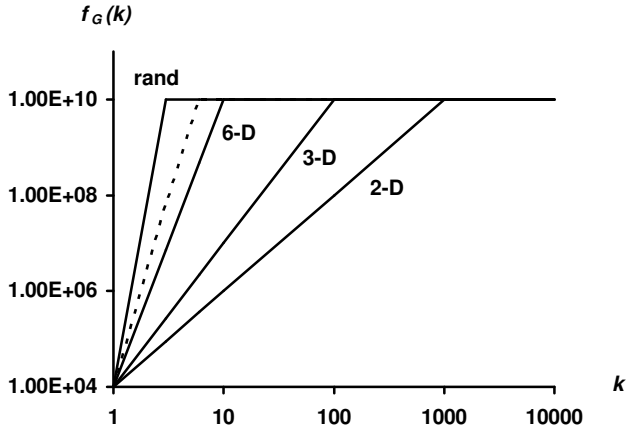

Figure 1. Spreading functions of various graphs of order 10^{10} and degree 10^4. The dotted curve is a rough estimate of the spreading function of cortex, derived from the estimate rad $G_{\text{cortex}} \approx 6$, corresponding to an estimate dim $G_{\text{cortex}} \approx 8$ (see text).

There is no way to directly measure $f_{\text{cortex}}(k)$. Consider however the four solid curves in Figure 1. These are the spreading functions of four different graphs, all of order $N = 10^{10}$ and degree $\delta = 10^4$, roughly the order and degree of G_{cortex}. The curve labeled 2-D is the spreading curve of a 2-D Euclidean graph of G_4 type; as seen above, $f_{2\text{-D}}(k) \approx \delta k^2$, in this case $f_{2\text{-D}}(k) \approx 10^4 k^2$. The curve labeled 3-D is the spreading curve of a 3-D Euclidean graph, defined in just the same way as G_4 except that V is embedded in $\mathbf{R}^3$ rather than $\mathbf{R}^2$; thus, $f_{3\text{-D}}(k) \approx 10^4 k^3$. Similarly, $f_{6\text{-D}}(k) \approx 10^4 k^6$ is the spreading curve of a 6-D Euclidean graph. Finally, f_{rand} is

the spreading curve of a random graph, of G_5 type; thus, $f_{\text{rand}}(k) \approx \delta^k = (10^4)^k = 10^{4k}$. Note that both axes use logarithmic scale.

Clearly, $f_G(k)$ cannot grow beyond 10^{10}, the order of the graph. The value of k where this maximum is reached is essentially the *radius* of G, defined as rad $G = \min_x \max_y d_G(x,y)$. If G is a Euclidean graph of order N and degree δ, we have $N \approx \delta\ (\text{rad } G)^{\dim G}$, since, for $k <$ rad G, $f_G(k) \approx \delta k^{\dim G}$. As shown in Figure 1, if $N = 10^{10}$ and $\delta = 10^4$, we get respectively rad G = 1000, 100, 10 for dim G = 2, 3, 6. If G is a large random graph, from the fact that $f_{\text{rand}}(k) \approx \delta^k$, we get that both the radius of G and its diameter, defined as diam $G = \max_{x,y} d_G(x,y)$, are, with high probability, equal to L+1, where L is the integer part of $\log_\delta N$. (For accurate asymptotic results on the radius and diameter of random graphs, see e.g. Burtin 1974, Bollobás 1985.) Thus, if G is a random graph of order $N = 10^{10}$ and degree $\delta = 10^4$, rad G and diam G are, with high probability, equal to 3.

In short, the notions of radius and dimensionality in large graphs are tightly linked; for practical purposes, high dimensionality is synonymous with low radius. Note that high dimensionality is not a characteristic feature of random graphs. It is easy to construct deterministic graphs of infinite dimensionality. A tree is a simple example. An example that packs

even more nodes within a given radius R given a maximal degree Δ is a de Bruijn graph (Bollobás 1985, p. 227). However, such mathematical constructions are of little interest as models of the brain. In contrast, it is not a far-fetched hypothesis that some developmental mechanisms would act to favor random connectivity in cortex. One mathematical fact of interest concerning random graphs is that their high dimensionality—equivalently low diameter—is, essentially, the best one can possibly achieve: random graphs are close to being the best way to pack as many nodes as possible under the constraint that the degree and the diameter of the graph do not exceed given limits.

3. The cortical connectivity graph

Armed with these observations, let us venture into some speculations as to the graph-theoretic dimensionality of cortex. According to an informal—and admittedly small—survey conducted by this author, the reckoning of most neuroanatomists is that $\operatorname{rad} G_{\text{cortex}}$ is somewhere between 3 and 6. Looking at Figure 1, we see that this imposes a narrow constraint on the cortical spreading function. Assuming for instance that $\operatorname{rad} G_{\text{cortex}} = 6$, we get the dotted curve between $f_{6\text{-D}}(k)$ and $f_{\text{rand}}(k)$; this curve turns out to be the spreading function of an 8-D Euclidean graph. This figure for $\dim G_{\text{cortex}}$ is somewhat arbitrary: if we estimate $\operatorname{rad} G_{\text{cortex}}$

to be equal to 4, we get an estimate of 10 for dim G_{cortex}. We also cannot rule out the possibility that rad $G_{\text{cortex}} = 3$. This would put G_{cortex} in the same category as a random graph, or, alternatively, as a 13-dimensional Euclidean graph. In any case, it appears safe to conclude that from a dimensionality point of view G_{cortex} is more akin to a random graph than to a Euclidean graph in 2 or 3 dimensions.

We can now return to our hypothesis about hidden regularities in cortical connectivity. We define a departure from randomness as any regular *low-dimensional* connectivity pattern, resulting from a development or learning process. As mentioned, the most conspicuous departures from randomness are those apparent in the 2-D and 3-D anatomical structure of cortex. The possibility that these are the *only* departures from randomness in cortex cannot be ruled out. However, we shall now put forward four independent arguments—each one maybe insufficient on its own—in favor of the hypothesis that G_{cortex} contains an altogether different type of regular connectivity patterns, which we propose to call "hidden". Like 2-D and 3-D anatomical patterns, hidden patterns constitute a departure from randomness in the sense that they have low graph-theoretic dimensionality. However, unlike anatomical patterns *stricto sensu*, hidden patterns do not live in 3-D physical space. The existence of hidden patterns is made possible by our observation that the substrate graph of connections in cortex is high-dimensional.

The first argument in favor of hidden connectivity patterns rests on the observation (Abeles *et al.* 1993) that cortical activity, particularly during behavior-related increased firing, is largely made up of accurate temporal patterns in collections of simultaneously recorded neurons. These patterns, dubbed *synfire* events (Abeles 1982, 1991), are defined with millisecond accuracy yet span periods that are considerably longer than a millisecond, sometimes in the order of 100 msec or more. There now exists a small body of literature, both experimental and theoretical, on synfire patterns. Of interest to us is the fact that the *synfire-chain* model, which is the simplest and practically inevitable way to account for the occurrence of synfire events, constitutes a striking departure from randomness. The spreading function of a synfire chain—whether complete or incomplete—is clearly $f_{\text{synfire}}(k) = nk$, where n is the width—or average width—of the chain. (Synfire chains are directed graphs; if edge direction is ignored, $f_{\text{synfire}}(k) = 2nk$ for a node far from either end of the chain.) Thus, if G_{synfire} is the connectivity graph of a synfire chain, dim $G_{\text{synfire}} = 1$.

Synfire chains may serve either as a model for short-range cortical circuitry (Abeles 1982, 1991) or of short-range and long-range connectivity (Bienenstock 1995). In either case, the assumption is that the topology of the chain is largely independent from the embedding of the participating neurons in physical space. Said differently, the graph-

theoretic distance $d_{\text{synfire}}(a,b)$ does not agree with the natural 3-D Euclidean distance in the brain. Rather, the chain is viewed as a 1-D *subgraph* of an essentially random, hence high-dimensional, substrate graph. Thus, although a synfire chain constitutes, just like an anatomical pattern, a strong departure from randomness, it is much less conspicuous than an anatomical regularity—e.g. a topographic mapping between two visual cortical areas—since it does not live in the usual 3-D space. Although other types of hidden connectivity patterns—and their counterparts in firing patterns—may exist in cortex, synfire-type patterns are the only concrete example we shall discuss here.

Note that synfire chains are functional rather than anatomical entities; synfire activity may be reliably carried by chains that are stored as mere synaptic-weight increments within an otherwise random graph. Thus, the existence of synfire chains in cortex is not a statement about cortical connectivity as such, in the sense, say, of retinotopic maps. We nevertheless regard synfire chains as regular patterns, in the sense defined above.

The second argument in favor of the existence of hidden cortical regularities, viz. synfire chains, is that these patterns can be shown, in theory, to develop as a result of Hebbian plasticity on a random substrate graph (Bienenstock 1991, Doursat 1991, Hertz and Prügel-Bennett 1996).

The proposed mechanism of synfire growth is virtually identical to mechanisms of development of anatomical regularities of the topographic type—e.g. retinotopic projections—that have been discussed first in the theoretical literature (Willshaw and von der Malsburg 1976), and then largely confirmed experimentally. The main difference lies in boundary conditions: whereas a 2-D source graph and a 2-D target graph are necessary for the growth of a topographic map, synfire-chain growth takes place when the substrate is a random, i.e., high-dimensional, graph.

The third argument has to do with the distribution of the number of synaptic contacts between pyramidal cells. Based mostly on the assumption that contacts occur at random whenever a candidate presynaptic axon happens to pass close enough to a patch of candidate postsynaptic dendrite, Braitenberg and Schüz (1991) reckon that the probability that an axon collateral of a pyramidal cell will make more than one contact with a neighboring pyramidal cell must be very small. Specifically, if the number of possible targets within a small cube of gray matter of side 100 μm is 1000, the number of contacts on any one of these cells made by a collateral passing through this cube obeys a binomial distribution B(n,p), where $p = 10^{-3}$ and n is the total number of synapses made by the collateral in question in that small region of cortex. A reasonable estimate for n is 20, which results in a 2% probability of 1 contact, and a vanishing probability of more than 1 contact. Only when n

takes on values of the order of 1000 do we get a non-negligible probability for 2, 3 or 4 contacts between two cells.

Recent data appear to be at odds with this prediction of the random connectivity model. Using dual whole-cell recording and light as well as electron microscopy on pairs of thick-tufted layer-V pyramidal cells, Markram (personal communication, see also Lübke *et al.* 1996) reports that (*i*) about 80% of the pairs studied did not contact each other at all, and (*ii*) the number of contacts for the remaining 20% of the pairs was distributed between 2 and 9, with a mean of about 5. (The sample size in such a study is necessarily small. On the other hand, the pairs of neurons were chosen on the basis of anatomical criteria thought to maximize the likelihood of connection between them.) Such a strikingly bimodal distribution, with one peak at 0 and another peak around 5, is inconsistent with the random model. Regardless of parameters, the random model predicts a binomial—or, in the large-n limit, Poisson—distribution, clearly unimodal.

The Markram data are thus suggestive of the existence of a strong departure from randomness in the local connectivity between pyramidal cells. These statistical data leave open the nature of the non-randomness, yet the data are consistent with the synfire-chain model, and, more specifically, with the growth/selection mechanism proposed in

Bienenstock (1991) and Doursat (1991). This mechanism is, in essence, a positive feedback principle, which promotes the growth or stabilization of contacts between cells that already show some propensity to trigger each other's firing. This mechanism will clearly favor the development of multiple contacts between those pairs of cells that happen to win in the competition.

The fourth and last argument is more speculative, as it stems from considerations pertaining to the formats of neural representations that might be used in high-level cognitive functions. One of the striking aspects of cognition is its extensive use of combinatorics. Mental entities of various nature are composed with each other in ever different patterns, a process of *binding* taking place within a second or so, with, possibly, a long-term memory trace. On the biological level, the issues of binding and compositionality have been intensely debated in the last decade. In short, if neural representations are to be composed with each other in useful ways, these representations need to have an appropriate internal structure (Fodor and Pylyshyn 1988; see also von der Malsburg 1987, Bienenstock and Geman 1995). A satisfactory theory of how the nervous system composes still eludes us, yet it would seem that one of the main ingredients in any such theory should be the understanding of the internal structure of neural representations. What is needed is a *common currency*, making it possible for the brain to bind with each other items that had

never been related to each other before. Synfire-type patterns are an attractive substrate for compositional operations, as they are all built along the same, universal, 1-D topology, hence can easily bind with each other (Bienenstock 1991, 1995, 1996).

4. Discussion

In the last two decades there have been a considerable number of theoretical studies addressing various aspects in the development of anatomical regularities (see above) in the brain from the perspective of self-organization. As suggested by von der Malsburg (1987), self-organization mechanisms acting on a much shorter time-scale than that of ontogeny may underlie *logogeny*, i.e., compositional cognitive processes. Composition of neural representations would likely rely upon mechanisms of synchronization and fast synaptic plasticity. We proposed elsewhere (Bienenstock 1991, 1995, 1996) that fast compositional processes may take advantage of a pre-existing structured substrate, laid down in cortical connectivity during development. In this sense, logogeny might be viewed as a seamless continuation of ontogeny.

In the present paper, we emphasize some dimensionality considerations related to this—as of yet largely hypothetical—compositional substrate.

Whereas anatomical patterns are regularities embedded in the physical 2-D or 3-D space, the neural substrate of compositionality would consist of *hidden patterns*, made possible only through the high dimensionality of cortex. The latter is linked to the randomness in G_{cortex}, and the regularities under discussion consist precisely of *departures from randomness*, i.e., of the formation, or reinforcement, of low-dimensional graphs within the high-dimensional G_{cortex}. (A fairly different class of models which also make use of the randomness of G_{cortex} is proposed in Valiant, 1994.)

Not surprisingly, this mechanism of pattern formation in its simplest form produces one-dimensional structures, i.e., synfire chains. However, in the scenario proposed here and in earlier papers, graphs of synfire type are envisioned as a mere first step in the construction of a complex hierarchy of patterns. Not only do synfire chains bind with each other in highly diverse semi-stable configurations (on the long term), they also overlap considerably with each other, in that they share common neurons. Thus, G_{cortex} may consist of a *superposition* of chains in which each neuron is shared by as many as several hundreds distinct chains (Bienenstock 1995, Herrmann *et al.* 1995).

Note that whereas a G_{synfire}-type graph in itself is of dimension 1, a superposition of such graphs may have arbitrary dimensionality. For instance, consider a graph G built on a regular 2-D lattice of neurons by

connecting the neurons into an array of synfire chains running parallel to each other horizontally, say from left to right, as well as vertically, say from top to bottom. G is a superposition of synfire chains, i.e., of graphs of dimensionality 1. Although G is not exactly of the G_3 type defined above, it is clearly of dimensionality 2. In contrast, consider now a random superposition of synfire chains, as studied in Bienenstock (1995) and Herrmann *et al.* (1995). The dimensionality of such a superposition, even for a small degree of overlap, e.g. two chains per neuron, is infinite.

As a final comment, consider the metaphor in which the self-organization synfire-growth process is likened to a process of crystal-growth (Bienenstock 1995). Not only are the two processes similar in that they rely on growth by accretion, they also depend critically on the presence of singularities, or *seeds*, and on the influence of various boundary conditions. In crystal growth, diverse seeds and/or boundary conditions may give rise to a rich diversity of macroscopic forms. However, while crystal growth takes place in three dimensions, synfire growth takes place in a high-dimensional graph. One may therefore expect an even richer behavior, under suitable "environmental" conditions for the development of synfire-based patterns.

References

Abeles, M. (1982) *Local cortical circuits: An electrophysiological study*, Springer-Verlag, Berlin.

Abeles, M. (1991) *Corticonics: Neuronal circuits of the cerebral cortex*. Cambridge University Press, Cambridge UK.

Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993) Spatiotemporal Firing Patterns in the Frontal Cortex of Behaving Monkeys. *J. Neurophysiol*., 70(4):1629-1638.

Bienenstock, E. (1991) Notes on the growth of a composition machine, In *Proceedings of the Royaumont Interdisciplinary Workshop on Compositionality in Cognition and Neural Networks*—I, D. Andler, E. Bienenstock, and B. Laks, Eds., pp. 25-43.

Bienenstock, E. (1995) A model of neocortex, *Network: Computation in Neural Systems*, 6:179-224.

Bienenstock, E. and Geman, S. (1995) Compositionality in Neural Systems. In: *The Handbook of Brain Theory and Neural Networks*. M. Arbib ed., Bradford/MIT, Cambridge, Mass. pp. 223-226.

Bienenstock, E. (1996) Composition. In: *Brain Theory—Biological Basis and Computational Principles*, A. Aertsen and V. Braitenberg eds., Elsevier Science, Amsterdam, pp. 269-300.

Bollobás, B. (1985) *Random Graphs*, Academic Press, London, New York.

Braitenberg, V. and Schüz, A. (1991) *Anatomy of the Cortex: Statistics and Geometry*, Springer-Verlag, Berlin.

Burtin, Yu.D. (1974) On extreme metric characteristics of a random graph, I Asymptotic estimates, *Theory Probab. Applics.* 19:710-725.

Doursat, R. (1991) Contribution à l'étude des représentations dans le système nerveux et dans les réseaux de neurones formels. Unpublished Ph.D. Thesis, Université Paris VI, Paris, France.

Fodor, J.A,, and Pylyshyn, Z. (1988) Connectionism and Cognitive Architecture: A critical Analysis. *Cognition*, 28:3-71.

Herrmann, M., Hertz, J.A., and Prügel-Bennett, A. (1995) Analysis of Synfire Chains, *Network: Computation in Neural Systems*, 6:403-414.

Hertz, J., and Prügel-Bennett, A. (1996) Learning short synfire chains by self-organization. *Network: Computation in Neural Systems*, 7:357-363.

Lübke, J., Markram, H., Frotscher, M., and Sakmann, B. (1996) Frequency and Dendritic Distribution of Autapses Established by Layer 5 Pyramidal Neurons in the Developing Rat Neocortex: Comparison with Synaptic Innervation of Adjacent Neurons of the Same Class, *J. Neurosci.* 16:3209-3218.

Spencer, J. (1973) *Paul Erdös: The Art of Counting, Selected Writings*. MIT Press, Cambridge, Mass.

Valiant, L.G. (1994) *Circuits of the Mind*, Oxford University Press. New York, Oxford.

von der Malsburg, C. (1987) Synaptic Plasticity as Basis of Brain Organization. In: *The Neural and Molecular Bases of Learning*, J.P. Changeux and M. Konishi, Eds., Wiley and Sons, 411-432.

Willshaw, D.J., and von der Malsburg, C. (1976) How patterned neural connections can be set-up by self-organization. *Proc. R. Soc. Lond.* B194, 431-445.