BRAIN THEORY

SPATIO-TEMPORAL ASPECTS OF BRAIN FUNCTION

A. AERTSEN / EDITOR

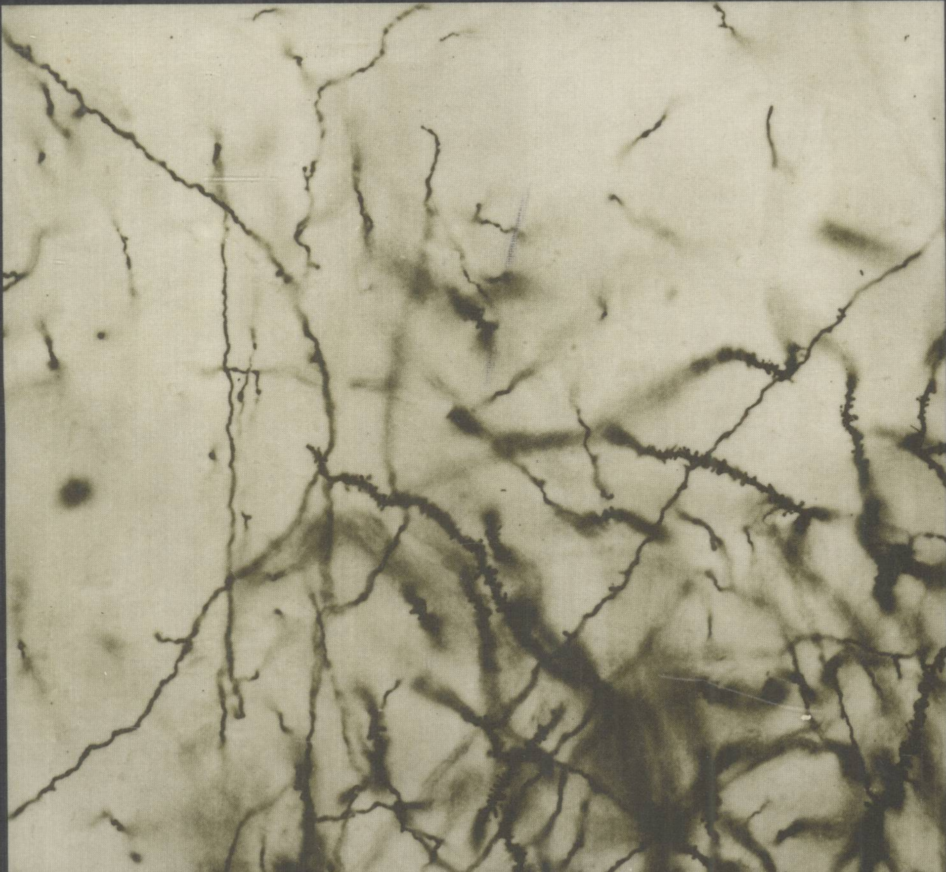

ELSEVIER

BRAIN THEORY

SPATIO-TEMPORAL ASPECTS OF BRAIN FUNCTION

Edited by

A. AERTSEN
Institut für Neuroinformatik
Rühr-Universität Bochum
Bochum, Germany

1993
ELSEVIER
AMSTERDAM - LONDON - NEW YORK - TOKYO

ELSEVIER SCIENCE PUBLISHERS B.V.
Sara Burgerhartstraat 25
P.O. Box 211, 1000 AE Amsterdam, The Netherlands

ISBN: 0 444 89839 5

This book is printed on acid-free paper.

Printed in The Netherlands.

PREFACE

Modern theories of brain function are increasingly concerned with dynamics. The task of organizing perception and behaviour in a meaningful interaction with the external world prompts the brain to recruit its various resources in a properly coordinated manner. Vis-a-vis the complexity and the multitude of the dynamics involved, a careful orchestration of the various processing components, distributed over space and time, is essential. Hence, it should come as no surprise that a number of recent developments in brain science have emphasized the aspect of spatio-temporal coordination. This holds for experimental and theoretical investigations alike. Experimental studies of the brain at different levels of resolution have established a multitude of distributed physiological mechanisms, ranging from interactions at the level of individual ionic channels, via the integrative processes governing the behaviour of single neurons, all the way up to the dynamics within and across populations of neurons or even entire regions of the brain. Similarly, theoretical investigations of the behaviour of model neurons and computational networks of neuron-like elements have revealed a number of intriguing dynamical phenomena which acquire their meaning only if the entire spatio-temporal complexity is considered. Thus, much of recent theoretical work in brain science is, in fact, concerned with the development of an adequate conceptual framework for the space-time dynamics exhibited by natural as well as by artificial brains.

The present collection of papers intends to capture these various developments in the brain sciences. It brings together new insights and concepts from various branches of experimental and theoretical neuroscience, partly in the form of review chapters, partly in short, focussed contributions, or critical essays. Starting point of this enterprise were the presentations and discussions at the Fourth International Meeting on Brain Theory held at the Istituto per la Ricerca Scientifica e Tecnologica (IRST) in Trento (Italy) on April 11-13, 1992. This meeting, organized by Valentino Braitenberg, Werner von Seelen, Luigi Stringa and Ad Aertsen, set out to explore "the problems of the processing of the temporal dimension of sensory input and of the generation of space-time patterns in the motor output, as well as the intervening storage and transformation of temporal patterns in nerve nets". The meeting was the fourth in a series, starting in 1984 at the International Center for Theoretical Physics in

Trieste [1], and continuing in 1986 in Bad Homburg [2] and in 1990 at Schloss Ringberg [3].

The meeting lasted three full days, providing a natural segmentation into three subheadings. On the first day, spatio-temporal aspects of brain function were discussed in the context of processing of sensory input and perception. Similarly, the third day focussed on spatio-temporal aspects of brain function at the output end: planning and control of movement. These two sessions flanked the middle one, which was dedicated to the intervening level of neuronal activity in the working brain, and the various dynamics observed at different levels of resolution in space and time. We adopted essentially the same tripartition to organize the present book. A fourth part combines contributions that transcend this scheme. A declared goal of these various meetings was "to raise an interest in theoretical models that actively seek confrontation with experimental data from the functioning brain, and by a didactic effort aimed at experimentalists to present their data in a format that makes them more amenable to theory". The same goal was pursued in the composition of the present book. The proof of the pudding is left as an exercise to the reader. The Fourth International Meeting on Brain Theory was jointly sponsored by the European Commission, the Gertrud Reemtsma Stiftung, the Istituto per la Ricerca Scientifica e Tecnologica (Trento), the Accademia degli Agiati, as well as by contributions from the various institutions who financed the participation of their delegates. Splendid hospitality, together with most efficient organization was provided by IRST, Trento. The last afternoon was hosted by the Accademia degli Agiati in the neighbouring town of Rovereto. The generous support by all these institutions is most gratefully acknowledged.

Ad Aertsen

References

[1] Palm G, Aertsen A (eds): Brain Theory. Berlin: Springer, 1986.
[2] von Seelen W, Leinhos U, Shaw G (eds): Organization of Neural Networks: Structures and Models. Weinheim: VCH Verlag, 1987.
[3] Aertsen A, Braitenberg V (eds): Information Processing in the Cortex: Experiments and Theory. Berlin: Springer, 1992.

CONTENTS

III. Space-Time Dynamics of Neuronal Activity in the Working Brain

IV. The Brain: The Organization of Perception and Action

Part One
Space-Time Aspects of Brain Function: Sensory Perception

Brain Theory / A. Aertsen (Editor)

Embodiments of geometry

J.J. Koenderink

Utrecht Biophysics Research Institute (UBI), Buys Ballot Laboratory, Princetonplein 5, 3584 CC Utrecht, the Netherlands.

Abstract This tutorial is meant to offer a broad overview of the importance of geometrical language for the description of the structures encountered in front end visual systems. Since the paper is of a tutorial nature, technicalities are suppressed and quite a few liberties with mathematical rigour are taken.

The brain of any optically guided agent has to be able to deal with geometrical problems as a matter of routine. Yet the notion that formal geometries might be powerful and apt tools in the description of functional brain structure is an alien one to neurophysiologists and psychophysicists alike. Yet I will argue for this very possibility.

The discussion will focus on the structure of the "front end" visual system in neurophysiology and the "visual field" in the psychology of perception. Several striking examples of empirically disclosed structures that map almost immediately upon certain formal geometries are discussed. Of course many research problems remain, those relevant to the issue of the brain as a "geometry engine" are summarily discussed.

1 GEOMETRY AND THE BRAIN

For the neuroanatomist the brain exists primarily *in space*, it is an incredibly complicated three dimensional shape. Good neuroanatomists have to be excellent intuitive geometers. Yet the notion that parts of the brain can advantageously be described as *embodiments of space* is probably in need of some explanation. The problem is that most nonmathematicians tend to hold an extremely limited view of geometry. My first aim in this introduction is to set the scene by trying to broaden this view.

One may distinguish (at least) three different views of "geometry":

logician's view: a geometry is a categorical system defined via an arbitrary set of axioms, stating all relations that may hold among a set of arbitrary items[1]. Some of such systems are (for historical reasons) known as "geometries";

algebraist's view: a geometry is the set of algebraic relations admitted by some fiducial structure, *e.g.*, all relations invariant under the action of some arbitrary fiducial group of transformations[2]. Various examples of such systems are conventionally known as "geometries";

(synthetic) geometer's view: geometries are about certain primitive objects ("point", "line") that are abstractions of certain physical entities (pin pricks, rulers) and carry intuive connotations[3]. In a geometry one studies the properties of the configurations that can be constructed using a predetermined set of ideal operations (*e.g.*, connecting two points with a line).

We may also distinguish the *physicist's*, or *engineer's* view of geometry[4,5], which is totally different from all of the above, because the physicist's geometry is *about* certain entities that exist as material systems (pin pricks, rulers).

The "geometry" one learns at school (in the tradition of Euclid's Elements) is (ideally) of the synthetic variety, although usually with a physicist's flavour. If this is your background you are bound to think of the notion of describing functional brain structures *as* geometry as preposterous.

The abstract logician's view of geometries as sets endowed with additional structure at first sight appears to be "about nothing". However, when you realize that a description of the functional order of parts of the brain can be little else but a description of the formal relations that exist or can exist among possible brain states, it becomes clear that abstractions such as the logics, geometries, *etc.*, are exactly the formal structures one needs. Of course one needs at least one thing in addition to the pure logic: Some way to interprete or relate the abstract entities to actual observations (neuroanatomical, physiological or psychophysical).

Geometrical descriptions of functions of brain and mind are indeed common. In psychophysics one describes the internal structure of sensory modalities via multidimensional Riemann spaces (*e.g.*, color space[6,7] or binocular space[8]), in physiology one often describes the empirically determined correlation between signals recorded from different structures in geometrical terms ("common input", "overlapping receptive fields"), in neuroanatomy one quite naturally describes the common "projections" geometrically. Indeed, one often uses the geometrical description as an "explanation" of certain phenomena. For instance, "local sign" is most often explained away (though quite mistakingly I think) by pointing out the retinotopical organization of certain cortical areas[9]. Attempts to interprete brain structures as "embodiments" of geometrical entities have been comparatively rare though. Perhaps the most common instance is that of an identification of mechano– or photo–receptors or cutanous or visual receptive fields as "points" in a rather literal sense. A more sophisticated (and rather extreme) example is the identification of certain dendritic tree structures with the Lie group generators of movements in the visual field[10].

In most cases the identification is essentially *ad hoc* and very teneous. In order to arrive at something more objective one has to provide *operational definitions* of geometrical entities for which physiologically plausible implementations can be pointed out. Such need not necessarily be very complicated at all. Consider a—argueably the simplest—example. According to Euclid a point is "that which has no parts", *i.e.*, which has either no discernable inner structure or whose inner structure we choose to ignore. (*E.g.*, as an astronomer may treat our galaxy as a "point" in a somewhat more global context.) How can one operationalize such a definition? An example is to consider the photoreceptors as points, but such an identification doesn't lead to exciting consequences, it is hardly more than a metaphorical way of speaking. A more viable approach[11] is to define a point as an *operator* that returns a minimal description of what "is at it". An example would be an operator that operates on retinal illuminance patterns and returns the illuminance at some location. By recording the output of such an operator for various inputs one characterizes the point. (For instance, the operator will only be responsive to a restricted set of punctate stimuli, or by noting the response to linear illuminance gradients in two orthogonal directions we obtain the "Cartesian coordinates" of the point and so forth.) For such an operator one easily imagines a physical embodiment, *e.g.*, a flat photovoltaic cell, a receptive field. There is no need for the physical operator to have "zero size" (if that were possible at all!), but one should turn a blind eye to its inner structure. (Thus the photovoltaic cell returns only a number, the incident flux per integration period.)

Of course the definition of geometrical objects is only a start. One will attempt to define *relations* between the objects in an operational manner. If succesful, one may use geometry as a *language* that describes the functional brain structure[12]. This has obvious advantages: because geometries are among the best understood mathematical objects one is then in a position to predict as yet unnoticed relations, and so forth. This is exactly the manner in which many natural phenomena have been formalized (*e.g.*, symplectic geometry is often used as a model of classical mechanics). It is just that brain science has largely ignored such obvious possibilities.

An additional issue I haven't mentioned is that some sensory modalities (haptics, optics, to a lesser extent acoustics) and motorical competences (grasping, manipulation, locomotion, ...) are most naturally described in geometrical language in the first place, this holds for both the stimulus as well as the response domains. The full process of optically, ...guided behavior can often be treated as a sequence of mappings between geometries.

2 GENERAL ISSUES

It will not be possible to discuss all (or even many) of the geometrical aspects of optically, ...guided behavior and the related brain processes and structures. I will focus on the optical domain ("vision") and merely pick out a few goodies. Because I can't assume that all of you will be thoroughly conversant with issues of abstract geometry I need first to spend some time on explaining various important distinctions and domains.

2.1 Local sign

The problem of "local sign" has been latently present in philosophical discussions since the classical period, but was first formulated precisely by Hermann Lotze[13] in the second half of the nineteenth century. His reasoning is fairly subtle. Lotze is concerned with conscious perception. Of course "perception" is by definition an aspect of consciousness, but I want to stress the fact that local sign does not apply to mere sensorimotor coordination. The latter can always be explained as an—be it incredibly complicated—map from the receptors to the skeletomuscular system for which one may think of network implementations, that is essentially a triviality. The point is that the perceiver *knows* where the photoreceptors (or rather retinal receptive fields) are in the visual field. That such is not trivial, but is somehow established during ontogenesis, is clear from the fact that cases exist in which this process fails (cases of amblyopia sometimes termed "tarachopia", or "scrambled visual field"). That the brain has a knowledge of *where* parts of the body are is not a trivial feat, as is well illustrated by the familiar phenomenon of "referred pain".

That somatotopy doesn't provide an "explanation" may become clear when you perform the *Gedanken* experiment of spatially permuting two cortical pyramidal cells leaving all the wiring intact: will this permute the locations in the visual field? If you believe so you'd better skip the remainder. Hermann Lotze had no need for homunculi and neither have I.

2.2 Symmetries

Another important point has to do with the issue of symmetries, or the "universality" of the human visual system. When someone explains a theory of the visual system usually a person from the audience will rise and object that the human visual system is not at all like the theory because .... Such a remark is almost certainly factually correct, but most often somewhat besides the point. I'm not claiming here that theories should be immune from empirical verification! It is just that a healthy pragmatic attitude increases potential pay offs. A theory that doesn't explain the universe may still be useful, as long as one is clear about its (always present) limits. For instance, strictly speaking all theories of physics are false except for some that are too hard to verify (superstrings?). For instance, geometrical optics or classical mechanics have been proven "false" a century ago, yet everybody uses them in daily laboratory situations: it would be sheer madness to discard anything as useful merely because one perceives obvious limits to the domain of applicability. Yet a similar practice is customary among brain scientists.

Most theories of vision start with symmetry assumptions, such as homogeneity, isotropy and selfsimilarity of the visual field. All these assumptions are obviously false[14,15]: the peripheral visual field differs from the central part, the vertical differs from the horizontal (not to speak of "oblique effects"), size invariance doubtless fails at very small or extremely large sizes. Does this doom the theories from the very start? I don't think so. For such a practice would lead one to abandon theory altogether, and one would be left with phenomenological descriptions of specific systems. Then "science" would become mere butterfly collecting and effectively cease to exist. This can't be tolerated.

Suppose one had to design a visual system for a robot[16] fit for an unspecified environment. Such a system cannot be committed to any specific world. One has to assume that one place will (on the average) be like the other, that there are no special directions, or any preferred size. The system would have to be designed on the assumptions of full symmetry: homogeneity, stationarity (in time), isotropy and selfsimilarity. Once one knows something about the prospective environment, one would of course use these constraints. (For instance, if all objects the robot ever had to deal with would be known to be of a single size, one would abandon selfsimilarity.) It is the same in the animal kingdom, where evolution has fitted species for specific (very generalized) biotopes. As these biotopes differ, so will the visual systems. From a theoretical point of view it is advantageous to view all visual systems as *specializations* of the fiducial totally uncommitted (tailored for an extremely varied biotope) system. Then one can build a general theory[17] and study the hierarchy of possible specializations of it. Although *no particular* visual system will be exactly like the fiducial universal prototype, yet they are more easily described as (in itself lawful) deviations from the ideal than as completely independent entities. Nontrivial comparisons between the visual systems of different varieties depend on the existence of such a genetic tree of theories.

In this paper I have no room to discuss many possible specializations and will largely concentrate on the fiducial, totally uncommitted system. Such a system is only constrained by the expectation that the structure of the world is basically equally important at all places, in all directions, at all sizes, at any epoch or timescale. Evolution will weed out systems that overly specialize in such a world. This principle turns out to be much more powerful than one might expect at first blush.

2.3 Local, multilocal and global structure

The concepts "local", "multilocal" (or neigborhood) and "global" are often used in a loose sense. In mathematics these terms have obtained a precise meaning[18], and are sometimes extended by further terms of the same nature (*e.g.*, "punctal"). I will use a fairly strict convention in this paper[19]. A property will be denoted as

punctal if it applies to a single point. *Example*: barometric pressure at Trento;

local if it applies to an infinitesimal neighborhood of a point. By that I mean that such geometrical objects as vectors, or tensors, which may be defined by operators ("receptive fields") at a single location, may enter the description. *Example*: direction into which the water will flow at some point;

multilocal if it applies to more than one distinct (though usually close) point. *Example*: height difference between Trento and Amsterdam;

global if it applies to the space as a whole. *Example*: the surface of the globe has no boundary.

Punctal and global properties are typically of little interest for our pursuits. It is very important to understand the vital differences between local and multilocal fully though.

The all important difference between the two is that *locally* all geometrical expertise can be encoded into the structure of the receptive fields. There is no need to be able to find the distance between receptive fields (they're all at the same location anyway) for instance. In contradistinction, *multilocally* the geometrical expertise has to be encoded into the relations between the receptive fields. For instance, when are three simple cells collinear (an extended local sign problem)? When are two simple cells at different locations parallel (problem of "parallel transport", or "connection")? Whereas local processing does not even need local sign, for multilocal processing one has to postulate some kind of calibration procedure.

3 TOPOLOGY

"Topology" is a branch of geometry that abstracts from all metrical (distance, angle, ...) and projective (collinearity, convexity, ...) properties. It is merely concerned with coincidence, overlap, boundary of regions, *etc.* Although this may seem very insignificant, exactly because of these restrictions topology is actually the all important backbone (or "glue") common to all geometries. Without topology little is left to do. For our purposes we practically equate "topology" with "local sign".

In order to start with *something*, I assume that we have a collection of "points". These will be operationally defined in the way indicated before. Let's consider relations between the points. Such relations will have to be operationalized too of course. A very simple example will be *coincidence*, or equality of points. I consider two points to be coincident if and only if they never fail to produce identical responses to all stimuli I try. The physiologist can distinguish the locations where he puts his probe, but *functionally* such points are equal, and one cannot distinguish between them. (Notice that the location where the physiologist puts the probe never enters the operational definition of the point. The brain *itself* cannot know this location, or one would have another local sign problem!).

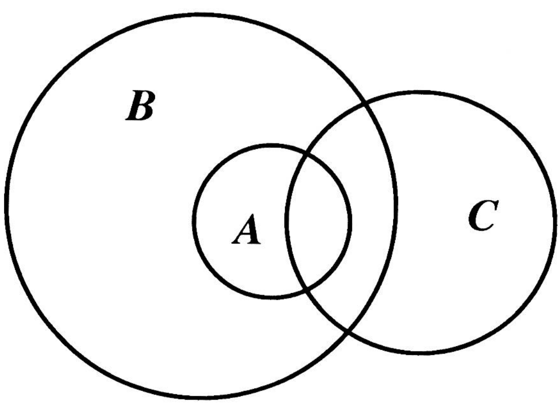

Figure 1: Functional definition of "inclusion": a receptive field (A say) is included in another (B say) if and only if any receptive field (C say) that correlates with it (A with C) also correlates with the other (B with C).

That correlated neural activity leads to the notion of "being at the same place" in the mind is vividly illustrated by a curious phenomenon first described (and explained in our manner) by Helmholtz[20]: a patient with toothache usually can't indicate whether the troubling spot is in the upper or lower jaw! Helmholtz explained this by noting that

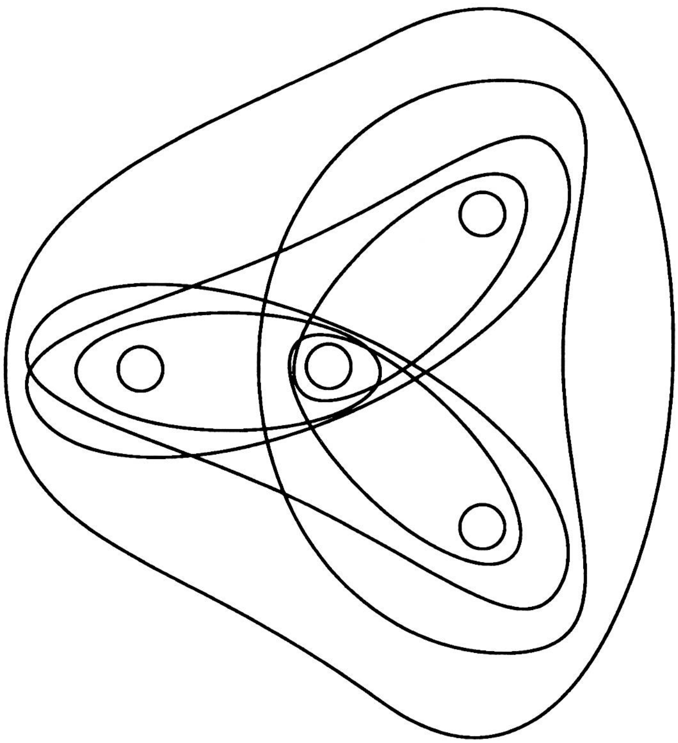

Figure 2: A simple configuration of eleven heavily overlapping receptive fields.

in chewing our food the neurons at the upper and lower jaw are always stimulated in synchrony. For the brain they are "coincident" in our definition, though physically they are well separated.

A somewhat less trivial example would be the case in which the responses where highly correlated, but not identical. In such a case the points are *related* without being coincident, I will call them "overlapping" (or "incident"). I will use the notation $\mathcal{A} \sim \mathcal{B}$ for the relation "$\mathcal{A}$ is incident with $\mathcal{B}$". Notice that this information is in principle available to the brain itself, whereas the position of my probe isn't. The relation is obviously symmetrical, that is, the relations $\mathcal{A} \sim \mathcal{B}$ and $\mathcal{B} \sim \mathcal{A}$ imply each other. Moreover, you (trivially) have $\mathcal{A} \sim \mathcal{A}$. From the relation of incidence I can derive many other useful relations via logical operations, for instance, I define "$\mathcal{A}$ is included in $\mathcal{B}$ (or $\mathcal{A} \subset \mathcal{B}$)"