Foresight Cognitive Systems Project

Research Review

Self-Organisation in the Nervous System

David Willshaw

Institute for Adaptive and Neural Computation,
School of Informatics, University of Edinburgh

This discussion document is part of a series of reviews of the state of the art in cognitive systems. It was prepared as background material for the Foresight Cognitive Systems Project. Technical Co-ordinators, Professor Richard Morris, Professor Lionel Tarassenko; Series Editor, Michael Kenward

The objective of the project is to examine recent progress in two major areas of research – computer science and neuroscience, and their related fields – and to understand whether progress in understanding cognition in living systems has new insights to offer those researching the construction of artificial cognitive systems.

Further details are available at the Foresight web site: http://www.foresight.gov.uk/

Contact the Foresight Cognitive Systems team at:

Foresight Directorate
Office of Science and Technology
1 Victoria Street
London
SW1H 0ET

Fax: 020 7215 6715
E-mail: cognitivesystemteam@dti.gsi.gov.uk

Contents

Self-Organisation in the Nervous System

1 Introduction

The term self-organisation is commonly held to describe the process by which individuals organise their communal behaviour to create global order by interactions amongst themselves rather than through external intervention or instruction. Despite this term receiving only scant mention in dictionaries, it has been used to describe many different types of activities. The clouds formed by birds in the sky, the coordinated movement of schools of fish or the paths formed by ants, as well as the intricate patterns seen in snowflakes are all the results of self-organisation. Other complex examples of spatial patterns are the many man-made or natural crystal structures.

In physics, the simplest examples are closed systems, where the system acts independently of external influences. The future state of the system is then controlled by its constitutive elements. Crucially, the emergence of a global pattern of order requires interactions between elements. Cooperative interactions will iron out local variations whereas competitive interactions will exaggerate them.

In the visually stunning Belouzov-Zhabotinsky reaction, two chemicals inhibit each other's autocatalysis, resulting in striking periodic changes in colour, as indicated by an appropriate dye. In magnetic materials, it is energetically favourable for the dipoles of neighbouring atoms to coalign, resulting in a global magnetic field. An example where there is a simple external influence is found in a laser. At low levels of excitation, individual atoms emit their light independently to produce incoherent light; at higher levels, the emission of light from all the atoms becomes highly coordinated through local interactions, producing coherent light.

Many systems exhibit both competition and cooperation. A well-analysed example of a temporal pattern of self-organisation in biology is found in the statistics of the populations of hares and lynxes, their predators, as recorded by the Hudson Bay Trading Company in Canada between 1849 and 1930 (Murray, 1993).

The lynx is a predator of the hare. Analysis of the number of pelts collected suggests the following pattern of events: a large fluctuation in one population can upset equilibrium states, in which the rates of reproduction and death of both species balance out. For example, a decrease in the number of prey will cause a corresponding decrease in the number of predators, who will have less food. The presence of fewer predators will then increase the number of prey and consequently will increase the number of predators, until finally the preys will decrease in number again. This pattern of events will repeat over and over again, yielding the cyclical variation in both prey and predator numbers over time that is seen in the records. Clearly this behaviour emerges from interactions between lynxes and hares and thus is an example of self-organisation.

1.1 Self-organisation in the nervous system

As the words suggest, order in a self-organising system emerges through local interactions between individuals in the absence of any external influence. As a highly complex and dynamic system involving many different elements interacting with each other, the nervous system displays many features of self-organisation. However, there will be very few, if any, examples of true self-organisation within the nervous system.

It is very likely that the organisation of regions of the nervous system depends on external influences, either from other regions of the nervous system of the body or under the influence of external stimuli, such as sensory stimulation from the outside world. The resulting organisation will be the result of interactions between the elements of the system itself as constrained by the particular boundary conditions that are in force, together with ongoing external influences.

1.2 Outline of this essay

I take the term self-organisation to refer to those aspects of organisation that result from interactions between the elements of the system as well as with external influences that do not themselves provide ordering information. I identify three forms of neural self-organisation, which I shall discuss in turn. These are:

- **Self-organisation in development** Since a key challenge in our understanding of the nervous system is to comprehend how such a highly structured yet complex system can emerge from a single fertilised egg, many phenomena displaying self-organisation are concerned with how the nervous system develops. Many of these developmental processes are a result of interactions within the system itself. External influences exist but they can be regarded as initial constraints or boundary conditions acting on the system.
- **Self-organisation as a complement to experiential changes** This refers to later stages in development, when self-organisation plays a role along with other mechanisms such as those involving external signals arising from the sensory environment. I examine the effects of external influences only when these do not contain any patterning information. Therefore I do not discuss the neurobiology of learning and memory, where specific patterns of activity are required to be stored in or recalled from the system.
- **Self-organisation as a complement to damage** The adult nervous system can respond to surgical or accidental damage. The facility for damaged brain to regenerate is either minimal or non-existent, which implies that the brain can self-organise, allowing healthy regions to take over functions previously carried out by other regions.

Section 2 considers development. I introduce some concepts of development at the genetic and molecular level. I then describe self-organisation in the formation of pattern within collections of cells (Section 2.1), in producing the correct numbers of cells (Section 2.2) and in the formation of ordered nerve connections (Section 2.3).

In Section 3, I look at the role of self-organisation in experiential changes. Section 3.1 describes the self-organisation of patterns of feature selectivity in the cortex and Section 3.2 provides a brief introduction to the self-organisation of cognitive function.

Section 4 is concerned with self-organisation as a response to injury, principally in the adult.

Finally, in Section 5, I discuss some open questions that are relevant to the subject of this essay. Section 6 gives a short reading list.

2 Self-organisation in development

Generating nerve cells of the right type, in the right numbers, in the right places and with the right connections is a formidable task. It involves cell division, cell migration, cell death and the formation and withdrawal of synapses. The essential steps of embryonic development are reviewed in many books. Wolpert (1991) provides a simple readable introduction: Price and Willshaw (2000) discuss mammalian neural development.

Every organism is defined by the sets of genes in its genome. This contains the initial instructions from which development proceeds. The set of three-letter 'words' obtained by reading the sequence of bases along the DNA defines a sequence of amino acids. Proteins are made out of amino acids and cells are made out of proteins.

There has been considerable progress in our understanding of how genes control development. The fruit fly, *Drosophila melanogaster*, has been used intensively in genetic research for many decades. It is small, has a short life cycle of two weeks, and large numbers of mutants have been identified and studied. The combination of the extensive knowledge of mutants and experimental embryological and mo-

lecular biological techniques has provided a profound understanding of the genetic regulation (i.e. control) of development in this species.

Remarkably, not only have many of the control mechanisms that operate in *Drosophila* been conserved in mammals, but so have many of the genes themselves. It is now commonplace to use information obtained from studies of *Drosophila* to search for specific regulatory genes in higher species and to formulate hypotheses regarding the general principles that underlie development in all organisms. In particular, work on *Drosophila* has provided a comprehensive understanding of how different regions of a developing organism can develop regional specificity. For example, certain morphogens – molecules that control the development of form, or morphogenesis, a term coined by Turing (1952) – are distributed in gradients in the early *Drosophila* embryo. They evoke different cellular responses at different concentrations, specifying the expression patterns of other genes that themselves regulate later-expressed genes. In this way, complex patterns of later-expressed genes emerge to confer positional identity on cells at each position in the embryo. The combined action of the specific cocktail of regulatory genes that each cell expresses is essential for conferring on each cell a particular phenotype appropriate for its position.

Many groups have shown that vertebrates have genes that are similar to those of *Drosophila*. Researchers have found vertebrate homologues for *Drosophila* genes that act within cells to regulate the expression of other genes (transcription factors) or that signal between cells to control processes such as axonal guidance. A good example of transcription factors is the large family of Hox genes (members of the homeotic clusters of genes) in mouse which have homology to the genes of the Antennapedia complex of *Drosophila* and which regulate the identity of segments of the *Drosophila* body. Another good example of conservation of developmental mechanisms is in the guidance of axons. Many of the receptor systems that have been implicated in this process are highly conserved between *Drosophila*, other invertebrate species and mammals.

It might be thought therefore that the genome could hold a coordinate-by-coordinate blueprint of the nervous system, specifying where each nerve cell is to be situated, what its functional properties are to be, and which other cells are to be contacted. If homeotic genes control the production of gross anatomical structure and cell differentiation, is it not possible that subordinate gene families subsequently control the remaining development processes? This is extremely unlikely given the large numbers of nerve cells and the many more connections that they make compared to the relatively small size of the genome.

If the 10^{14} connections between the 10^{10} neurons of the human neocortex were made at random, this would require at least 10^{15} bits of information compared to the 10^{9} bits in the human genome. It is more likely that the genome contains the 'rules of development'. For example, it is well known that the connectivity of the brain is highly structured, with topographic maps found between many sensory structures and neocortex. The rules of development would specify the general features of the mapping and the fine details could be arranged through interactions between the constituent parts. Many of the features of the system are, so to speak, arranged by the nervous system itself, or self-organised. The next three subsections describe different parts of the developmental process where mechanisms of self-organisation make an important contribution.

2.1 Self-organisation and pattern formation

A central aim of developmental biology is to understand how cells in different positions develop differently; i.e. how regional specification comes about. This is as true for the development of a structure as complex as the cerebral cortex – where each point in the dorsal telencephalic wall (the precursor of the cortex) acquires a unique functional property, with relative invariance in the layout of these properties between the individuals of the same species – as it is for the development of the five distinct digits of the hand or the patterns of markings on sea-shells, zebras or leopards. Many of the principles that govern the ways in which regional specification arises during development have been identified in studies of early embryogenesis. We are relatively ignorant of the mechanisms that control brain regionalisation, which makes it all the more important to generate hypotheses with knowledge of principles deduced from studies of earlier developing systems.

The key questions are:

- How does a mass of developing cells acquire differences one from another?
- How is this information used to determine their different fates?

These two questions are interlinked. For example, does each cell acquire its own identity independently from the instructions in the genome? Or do certain cells act as organisers and instruct their neighbours, which happens in limb morphogenesis?

Turing (1952) investigated one possibility theoretically. He analysed the emergence of pattern in a collection of cells and showed that, starting from a uniform concentration of morphogens, which interact with each other and diffuse between cells (giving rise to the term reaction-diffusion), the patterns of molecular concentration produced over the cell population had peaks and troughs defining characteristic periodicities. Turing argued that in a developing organism, a set of chemicals could be used in this way to set up a prepattern that will determine specific features of the organism. The patterns are not preprogrammed and emerge through self-organisation.

In the brain, generation of regional differences has to occur at different levels. How are particular areas or regions within a brain nucleus or system distinguished one from another? How is cellular identity within a given region determined? I now discuss regional specification within the forebrain and within the neocortex. I discuss specification within a brain region in Section 2.3.

2.1.1 Pattern formation in the specification of forebrain

Much of our understanding derives from work on the amphibian *Xenopus laevis*. Two important terms used to describe complementary mechanisms that can generate regional specification are mosaicism and regulation, i.e. whether the development of an individual cell is independent of or dependent on the development of other cells.

Regulative behaviour is commonly found among cells undergoing regional specification in the developing mammalian nervous system, which indicates that the mechanism of specification involves intercellular signalling. In early embryogenesis, major sources of such signals are well defined, and include the so-called Spemann organiser (Spemann, 1938). The signals that affect the developmental pathway are termed inductive signals as the process involves transfer of information from mesoderm to ectoderm, two of the three germ layers formed very early in development. Although inductive signalling is almost certainly a widespread mechanism in the later stages of cortical regionalisation, its clearest roles are in the early stages of forebrain development.

The very early regionalisation of the developing forebrain can be detected by morphological criteria and by analysis of the discrete domains of expression of regulatory genes. It is possible that the many genes known to be involved give each region of the developing forebrain a unique identity, probably through combinatorial actions. They may do this by controlling the expression of numerous other genes required for the characteristic morphological differentiation of that region. Amongst the molecules known to be involved are the diffusible proteins notch and delta, the wnt family of glycoproteins and the hedgehog family of proteins first identified in *Drosophila*.

Regional specificity of gene expression in the telencephalon, from which the forebrain develops, is likely to control regional differences in morphological characteristics, through actions on the cellular processes of proliferation, migration and differentiation. How different regions come to express different genes in the first place is a subject of speculation. One simple possibility is that a small number of genes distributed over the neural plate and very early neural tube, the forerunners of the nervous system, generate gradients of molecules.

Through transport and inter-cellular exchange, molecules at different levels of concentration would become localised in different cells. This can create domains of gene expression with sharp boundaries. This type of process is known to generate regionalised domains of gene expression in the early embryo of *Drosophila*. Although there are homologues of these genes in the mammalian forebrain, drawing close parallels between mammalian forebrain and *Drosophila* development may be dangerous given

the differences between them at a cellular level. Nonetheless, the principle that continuous molecular gradients may be read out to create domains of expression of other genes distributed with discrete levels is well established. There are various models for how this can be done. These models are usually formulated according to the concept of positional information and are constrained by the regulatory phenomena often seen in embryogenesis.

Positional information Evidence from classical embryological experiments on a mass of cells after the removal of some cells or the transposition of cells to a new position resulted in the proposal that the fate of a cell is determined by its position within the morphogenetic field of cells. A particular set of cells that makes a single field can form its own organ when transplanted to a foreign site and cells within the field can regulate to take over the function of other cells that are removed from it. How is each cell within the field instructed or, as expressed by Wolpert (1969), how does the cell acquire its positional information? As already discussed, one fundamental way in which information is supplied in development is through inducing signals supplied through extracellular means and various ways of assigning differences amongst cells by means of morphogens have been considered. In the simplest case of a one-dimensional field of cells that specifies the digits of the hand, for example, a gradient of morphogen would enable different parts of the field to be distinguished; specifying particular threshold values of morphogen would determine which cells would develop into which digit.

Simple source/sink models Various different ways of producing spatially varying profiles of a putative morphogen have been considered. For a single dimension, morphogen flows from a single source to a single sink to set up a graded variant of morphogen down the line of cells. Alternatively there could be a single source and all cells acts as sinks through leakage and other forms of loss. These models have been found to be unsuitable. In particular, they do not adapt in the required fashion following perturbations such as the removal of a substantial number of cells.

Reaction-diffusion model Gierer and Meinhardt proposed a model of the reaction-diffusion type in which there are two molecules with different properties: an activator, which stimulates its own production, and an inhibitor, which diffuses at a faster rate (for a review, see Meinhardt, 1982). The activator stimulates production of the inhibitor but the inhibitor represses the production of the activator. A small local increase in the amount of activator will result in more activator being produced, thus giving rise to a local source of this molecule. The inhibitor produced as a result will spread out more quickly than the activator and so a sink for activator will be established nearby. In this way, spatial patterns of activator and inhibitor become distributed across the array of cells. In these reaction-diffusion models, a crucial parameter is the size of the morphogenetic field over which the pattern is being formed compared with the diffusion lengths of the two molecules. If the field is very much smaller than the diffusion lengths, periodically repeating patterns will be produced; if the field is comparable in size to these diffusion lengths, a single gradient of morphogen results. Imposing a weak gradient of activator production to determine polarity yields a single gradient. Diminution of field size causes the full gradient to be restored, up to a limit. This is important as an explanation of the findings in developmental biology that in some animals structures can regenerate from partial structures.

Reaction-diffusion mechanisms have been applied to the generation of many different patterns, such as stripes, spots and other markings that appear on animal coats, and to other naturally occurring patterns, such as those on butterfly wings or those on sea-shells (Meinhardt, 1982; Murray, 1993). There is a close relationship between these mechanisms, involving different types of non-linear interactions, and the self-organising systems studied in physics.

Role of gradients The primary role to be fulfilled by systems of gradients is to provide a way for cells to be distinguished from one another. The reaction-diffusion scheme at least provides a way of doing this which is resistant (within limits) to changes in morphogenetic field size. It is assumed that a separate mechanism translates an amount of morphogen into an instruction to build a cellular structure. In some cases, patterns of morphogens are required to specify the coordinate systems of developing organs. It is natural, although not necessary, to assume that the axes of the morphogens will match those of the required coordinate system. For example, a rectangular coordinate system might be provided by two morphogens, each identified with one axis. In cases where there is no such requirement, as long as cells can be distinguished one from another the pattern of morphogens can be arbitrary.

2.1.2 Pattern formation in the specification of neocortex

The neocortex, the uniquely mammalian structure which has evolved rapidly and extensively in primates, is thought to be the source of our highly developed cognitive functions. We can subdivide it into distinct areas, according to anatomical and functional criteria. There has been much discussion of how these distinct areas of neocortex develop from the early cortical plate, with a relatively homogeneous appearance. It has been suggested that neocortical organisation is determined by: its afferents (-inputs); the significant amount of information preprogrammed into the neocortex; interactions within the developing neocortex, independently its inputs.

In the adult mammal, the cerebral cortex has been divided into areas according to their histological appearance. This is illustrated in Figure 1, which shows the cytoarchitectonic fields of the human brain as defined almost a century ago by Brodmann (1909). These distinct fields were defined according to the relative thickness of cell and fibre layers and Brodmann was able to delineate sharp borders between neighbouring areas. This analysis of the human brain has been extended to other species, showing that their cortices can also be subdivided into cytoarchitectonic fields, with equivalent fields occupying relatively similar positions. Most importantly, these anatomically defined regions have different functional specialisations.

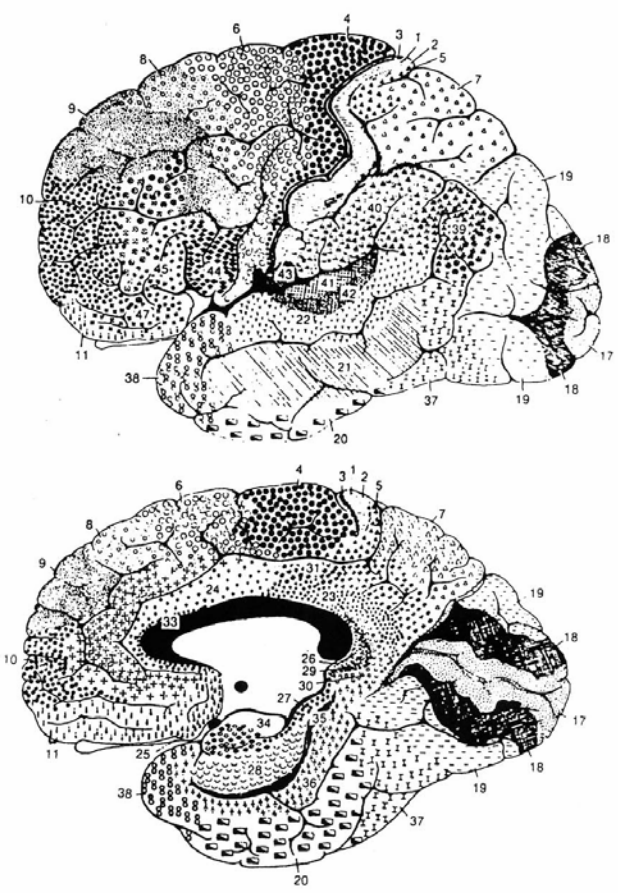

Figure 1: Brodmann's cytoarchitectonic fields of the human brain marked on the left cerebral hemisphere, seen from the lateral (top) and medial (bottom) aspects. After Brodmann (1909).

The question of how much of the regionalisation of the cortex is imposed on it by the order of its afferents (i.e. inputs) and how much is specified before innervation has become a major preoccupation. One extreme view is that the cortex is a naive sheet of cells whose identities are determined by the nature, the order and the detailed connectivity of the afferents that they receive. Another view is that the cells of the cortex do have at least some regional identity before they become innervated.

There is now considerable electrophysiological evidence from the results of embryonic and neonatal transplantations that differences between distinct areas in the neocortex are induced by the afferent thalamocortical axons that innervate the cortex.

However, some area-specific differences are detectable before any cortical innervation has taken place. There is evidence for region-specific differences in the rates of proliferation and expression of molecules in the cerebral cortex prior to innervation. In some cases these molecular differences are not altered when expressing regions are transplanted to non-expressing sites, suggesting that region-specific differences may be determined (i.e., irreversible) prior to innervation. The abolition of the activity in cortical afferents does not prevent the development in the cortex of characteristic region-specific distributions of molecules. In addition, there are several reports of cortical area-specific gene expression that begins before or is independent of afferent innervation.

In mouse, a specific gene is expressed in the somatosensory cortex from before the time of afferent innervation. The gene is still expressed if the developing somatosensory cortex is transplanted to an ectopic location but is not expressed by other regions of developing cortex even if they are transplanted to the somatosensory cortex. Most recently, the arrangement of cortical areas in two different mouse mutants that lack thalamocortical connections have been found to be abnormal. These results argue in favour of cortical cells having some positional identity without innervation, although how much is far from clear.

Other experiments have addressed this issue by studying the properties of different cortical regions either in culture or after transplantation. Tissue culture experiments to investigate the specificity of axons from different thalamic regions for different cortical areas showed that axons from thalamic explants exhibit no preference for the area of neocortex with which they were cultured. They grew equally well on their normal target areas as on other non-target areas of the neocortex.

Other experiments have involved the transplantation of regions of the developing cortex to abnormal sites. In transplants between neocortical regions, the donor tissue was found to develop attributes of the new host region rather than retaining its normal attributes. Pieces of visual cortex grafted into motor cortex developed persistent projections to the spinal cord, as does normal motor cortex but unlike normal visual cortex. When pieces of motor cortex were grafted into the visual cortex, they developed persistent projections to the superior colliculus, as does normal visual cortex, but unlike normal motor cortex.

Sur et al. (1988) carried out experiments on the regeneration of connections in ferrets where target structures of sensory fibres were removed, leading to the fibres being diverted to other cortical structures. Removal of lateral geniculate nucleus (a relay station between retina and cortex) and visual cortex led to visual afferents innervating medial geniculate nucleus (the destination of auditory fibres). The result was that cells of the auditory cortex became responsive to visual stimuli and acquired the functional characteristics of cells of the visual cortex.

All of these experiments indicate that different regions of the embryonic and neonatal neocortex have a low level of commitment to their specific regional fates. Although it is widely accepted that the fates of embryonic cortical regions are not determined by birth (i.e. they are not irreversible), the degree to which they are specified remains uncertain. The results of the experiments described above suggest that it is easy to deflect developing neocortical regions from their normal developmental pathways. However, recent transplant experiments similar to those outlined above have come up with opposite results with no clear explanation for the difference. Furthermore, it may be harder to alter the fates of embryonic tissue transplanted between the neocortex and other cortical areas, such as the limbic cortex.

2.1.3 Self-organisation in forebrain and neocortical development

To summarise Sections 2.1.1 and 2.1.2, there is much evidence at the genetic, molecular and neural levels for the roles of self-organising influences in the development of regional specificity in these systems. We know most about the early stages of development of the forebrain, where inductive effects are important. In the development of cortical regionalisation, the influences of specification prior to innervation and cortical afferents contribute. Several mathematical and computer models have been developed but at present they lack specific application.

2.1.4 Pattern formation in the positioning of cells

For nerve cells to function correctly, they have to be placed in the correct position. To investigate how this is done, we need to look at systems where it is clear what it means to place cells correctly. This is most easily accomplished in parts of the nervous system where there is a high degree of order.

Nerve cells within invertebrates are well ordered, while the vertebrate nervous system is less highly ordered. The degree of order varies greatly between different parts of the nervous system. In mammals, hippocampus, olfactory cortex, cerebellar cortex and retina are examples of relatively ordered brain structures.

In the cerebellar cortex, the Purkinje cells (the main output cells) form, with other cell types, a regular three-dimensional lattice. These cells define regular, typically hexagonal, neighbourhoods with inter-cellular spacings that grow steadily in size during the first few postnatal weeks.

In the retina, neurons form regular arrays, called mosaics. These exist in many species, being more regular in invertebrates than vertebrates. Several cell types are distributed regularly in the two dimensional plane of the retina, giving a constant cell-spacing between adjacent cells. In insects, receptor cells are highly ordered, forming precise hexagonal patterns.

It is important for cells to be arranged regularly across the retina. The vertebrate retina is organised in such a way that many local circuits can analyse each part of the image. This allows it to handle the vast information flux of a complex, ever-changing visual scene using only slow, noisy neurons. The circuits work in parallel to assess different static or dynamic aspects of colour, brightness and contrast. As it is the regular repetition of these circuits across the retina that gives rise to mosaics, each mosaic can be assumed to embody a unique function in sensory processing.

There are different types of mosaics, such as those involving cholinergic cells (amacrine cells with acetylcholine as transmitter), horizontal cells, different classes of ganglion cells and the different types of photoreceptor cells. The mosaics appear early on in development. They form while the cell membership is still forming, by the processes of neurogenesis, cell death and cell migration. For example, in developing cholinergic mosaics, as new cells enter the array, neurons move sideways to preserve a constant inter-cell spacing.

It is thought that the rudiments of the pattern are determined at a very coarse scale by molecular markers, derived from genes such as the pax family of homeotic genes, in the embryonic neuroepithelium, which forms the eye. The ordered spacing can be simulated by a very simple local exclusion rule, applied to cells of the same type, which specifies the minimum distance between cells.

Recent research suggests how this rule can be implemented at the molecular level. A computer simulation study has shown that regular, advancing arrays such as cone mosaics can emerge from simple cellular-automaton rules that are applied to initially random arrangements, but also that many different sets of these rules converge on the same simple patterns.

Although information must be supplied to position cells in the correct general region, this solution has the advantage that there is no need for a means of pre-specifying the positioning of nerve cells precisely. There are other advantages. Since interactions are short range, neither local errors nor the introduction of new cells at the periphery of the array disturb the pattern. If the local interactions are restricted to cells of the same type, introduction of an array of cells of a different type will not perturb the arrangements in the preexisting arrays. Finally, the ability of cells to recognise others within a limited range could lay the basis for the formation of the topographically ordered maps of connections that exist in the retina and which subserve visual processing (see Section 2.3).

2.2 Making the correct numbers of cells: cell death

Of all the mechanisms involved in the formation and maintenance of the nervous system, cell death is the best understood, especially at the level of how genetic instructions can bring about cellular self-destruction. It has long been realised that cell death can be a physiological as well as a pathological process, i.e., that many cells die even during normal brain development.

Cell death during animal development was first observed by Vogt in 1842, in amphibia. The basic findings are: there is substantial motor neuron death in normal development; removal of a developing limb bud from chick embryos causes increased death of the motor neurons; some of the motor neurons that would have died during normal development can be rescued by grafting in an extra limb bud.

Researchers have reported similar findings in *Xenopus*. Subsequent studies have shown that cell death is a normal occurrence amongst many neuronal populations in the developing vertebrate nervous system, taking place when axons begin to reach and activate their targets. The number of neurons in a given population first rises then declines towards the constant adult number. The proportion of cells that dies

varies among different neuronal populations, ranging from the removal of only a small number of the neurons in some regions to more than half the population in others. In the retina, results from a range of mammalian species indicate that 50-90% of the retinal ganglion cell population will die during development.

That so much cell death occurs during normal development is counterintuitive. Thinking anthropomorphically, it appears wasteful. It is not clear why evolution should have selected such a developmental process.

Blocking normal neuronal death – by making transgenic mice with mutations of genes that regulate cell death, thereby increasing the numbers of neurons – does not affect lifespan. Nonetheless, cell death is a significant developmental process that demands an explanation both in terms of its role in development and the molecular mechanisms that control it.

Researchers have advanced various explanations for nerve cell death, amongst them being:

- failure of neurons to find their target;
- failure to make the correct connections;
- the elimination of entire structures that may act as transient scaffolds; examples of this are the subplate, a layer of cells that is critical for the development of cerebral cortex, and the Rohan-Beard cells in amphibian embryos, which are temporary sensory neurons;
- removal of transient branches of the tree of lineage. This seems to be the case in invertebrates. In the nematode *C. elegans*, around 20% of the 300 nerve cells generated by cell division are pre-programmed to die;
- lack of adequate innervation, as is the case in insect optic lobe.

All these explanations seem to apply in special cases.

A commonly held view is that in many cases this substantial amount of nerve-cell death results from the action of a mechanism that matches the number of presynaptic cells to the number of postsynaptic cells. One hypothesis for this is that neurons compete for a supply of one or more so-called 'neurotrophic' factors that are known to be produced by their target cells. According to this neurotrophic hypothesis, insufficient neurotrophic factor is produced to support the excessive numbers of neurons generated and those that are unsuccessful in the competition die.

Additional support for the idea of matching presynaptic and postsynaptic cell numbers has come from experiments in which chick lumbosacral cords were transplanted into quails and vice versa before the limbs became innervated. Chicks are larger than quails and have bigger muscles with more muscle fibres. More quail motor neurons survived in the chick than in the native quail, and fewer chick motor neurons survived in the quail than in their own environment. There was a correlation between the number of motor neurons surviving and the number of muscle fibres available for innervation.

The nervous system contains control mechanisms by which the numbers of presynaptic and postsynaptic cells are matched apparently automatically. This involves the apparently destructive side-effect of cell death. This self-organising mechanism has great benefit, making it unnecessary to have ultra-precise controls over the generation of neuron numbers. Whilst qualitatively the effect is well established, the underlying mechanisms and their quantitative implications represent exciting challenges for the future.

2.3 Development of connections

Once nerve cell axons have found their correct target within the nervous system, neuron numbers have been adjusted to their adult levels and neurons have been correctly positioned, the appropriate connections have to be made. It is generally thought that this happens in two stages. Initially, the axonal terminals are distributed across the target relatively diffusely. Subsequently, there is a rearrangement or refinement of connections. This picture has emerged from many different animal preparations, principally

from experiments on the innervation of skeletal muscles and autonomic ganglia of neonatal rats, and on the visual pathways of *Xenopus*, kittens, ferrets and infant monkeys.

There are good reasons why some sort of refinement of connections is essential. Firstly, far too many neurons exist for the positioning of each to be controlled by a genetically determined programme. Secondly, it is difficult to see how such a programme can determine the fine details of the connections between independently specified sets of neurons. Finally, as a consequence of the continual growth of some animals there has to be a continual remapping of connections during development to accommodate the generation of new cells and consequently new connections.

The two stages are thought to involve mechanisms which guide axons to their initial, approximate destination to generate an initial pattern of connections, followed by a stage during which connections are remodelled to form the adult configuration, which involves the loss of existing connections and the generation of new ones. Many people think of the first stage as being programmed genetically and the second one driven by neural activity so that the refinement of connections is sculpted to fit the uses to which the neural system has to be put. The precise division between these mechanisms is not clear, particularly in that it is not clear how much specificity of connection is imparted during each stage.

I will describe the role of self-organisation at two levels. I will first discuss the network level, as exemplified by the development of ordered maps of connections between the vertebrate retina and optic tectum in lower vertebrates (equivalent to the superior colliculus in mammals). I will then consider the single-cell level, as exemplified by the elimination of connections from the developing neuromuscular junction.

2.3.1 Map formation

A striking feature of many of the connection patterns between collections of nerve cells is that they are highly ordered. Evidence for this comes mainly from two types of experiment.

Firstly, in electrophysiological experiments, stimulation of a small region in one structure, such as a sensory surface or a nucleus, leads to activation of cells in a small region of its target. As the stimulus is moved systematically across the structure, the region that responds shifts in a corresponding fashion. The region in stimulus space that produces a response at a particular target position is called the 'receptive field'.

Secondly, the mapping between two points in different structures can often be established in anatomical experiments using axonal tracers (molecules that can be injected at discrete points to label axons running to and from those points). Tracers placed at one point in one structure typically label a small, circumscribed area in the target, the spatial layout of points of administration (in different animals) being reflected in the layout of points to which the tracers go in the target. Such ordered anatomical layouts of connections provide the substrates for

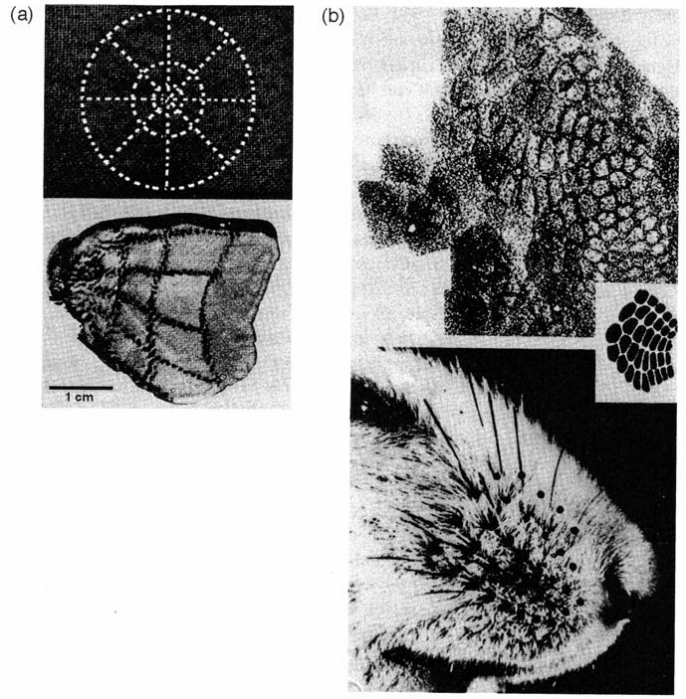

Figure 2: (a) Showing the ordered projection of the retina onto the visual cortex of monkey. The animal was exposed to the image (top photograph) and activity in the visual cortex was revealed by processing for deoxyglucose uptake (bottom photograph). Note the ordered projection and non-linearities. (Reproduced from Tootell et al., 1982. (b)) The pattern of whiskers on the adult rodent snout (lower photograph) and the pattern of barrels in somatosensory cortex (upper photograph, schematised in inset) to which the whiskers project in a one-to-one fashion. The pattern of whiskers is almost invariant from one animal to another. (Copyright acknowledged to Woolsey and van der Loos, 1970).

tions provide the substrates for the ordering observed in electrophysiological experiments.

Many neural maps are effectively projections of one two-dimensional surface onto another. For example, axons from each small cluster of ganglion cells in the mammalian retina project, via the lateral geniculate nucleus (LGN) of the thalamus, onto a small area of visual cortex, with the result that a map of the retina is spread over the surface of its target structure (Figure 2a). In amphibia and fish the retina projects directly to optic tectum where, once again, an orderly map of the retina is found. Auditory cortex contains an example of a one-dimensional map of frequency (Figure 3a).

Another striking example is the existence of precise maps of connections in somatosensory cortex. Rodents make much more extensive use of tactile information than of visual information. Correspondingly, their somatosensory cortex is relatively large whereas their visual cortex is relatively small and simple. A large area of primary somatosensory cortex is occupied by an ordered representation of the facial whisker pad (Figure 2b).

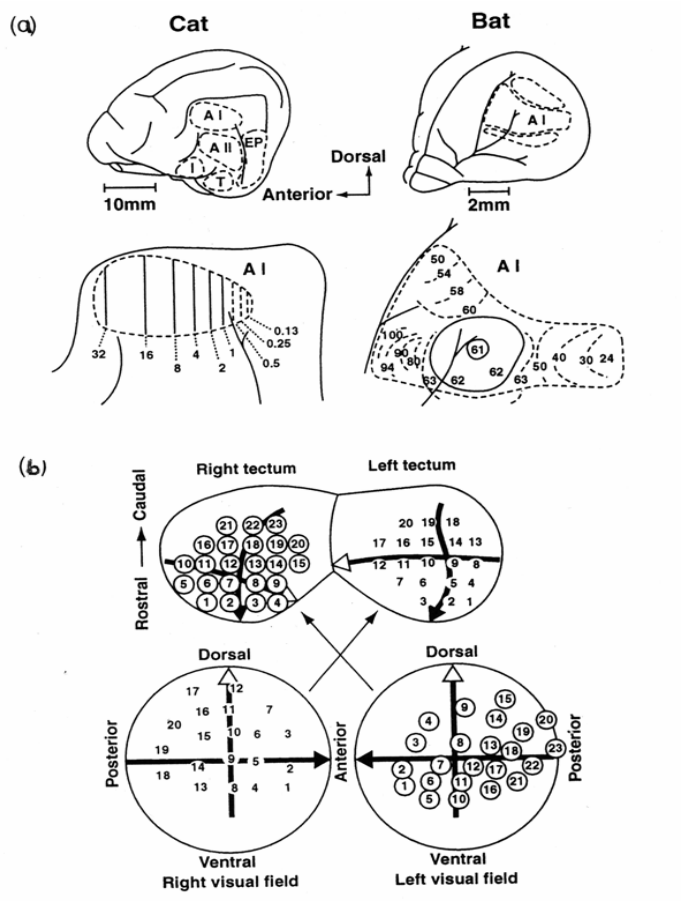

Figure 3: (a) The one-dimensional map of frequency in auditory cortex. (AI and AII, primary and secondary auditory cortex; EP, posterior ectosylvian gyrus; I, insula; T, temporal field; numbers 0.13-100, kHz). There is a regular tonotopic representation in cat but a distortion of this regularity in bat by a large representation of 61-62 kHz, the frequency of its echolocating signal (based on Sugar, 1978 and Shepherd, 1994).

(b) The ordered projection from retina to contralateral tectum in adult Xenopus loevis. The numbers indicate where in the visual field a small point stimulus evoked maximal response at the correspondingly numbered tectal position. Reproduced from Jacobson (1967).

Anatomical investigations have revealed a set of barrel-like structures, with a one-to-one relationship between the arrangement of whiskers and the arrangement of barrels: where the muzzle contains an extra whisker, there is an extra barrel in the topographically equivalent place and vice versa. Neurophysiological recordings have established that activation of an individual whisker excites the cells in the corresponding barrel. There is also a topographical representation of the whiskers in each of the two nuclei which form the relay stations linking the sensors with the somatosensory cortex. This ordered map is not preserved throughout the length of the pathway from sensorium to cortex but rather is recreated at each individual relay station.

There is much evidence for plasticity in the whisker-to-barrel pathway in rodents. An intact sensory periphery is required during a certain critical period of development for the normal map to develop. When rows of follicles are injured at birth, before barrels form, the corresponding row of barrels is absent in the adult, and is replaced by a small barrel-less territory. Barrels develop in the first postnatal week and their morphology can be manipulated by the selective lesioning of the whisker follicles. The earlier the follicles are removed, the more extensive the resulting morphological aberration.

Generally, where axonal projections are through intermediate structures, the intermediate maps are themselves well-ordered. Ordered projections are the rule, although there are a few apparently disor-

dered projections such as that of visual space onto the pyramidal cells of mammalian hippocampus, which in rats respond to specific locations in the animal's environment, and in the direct projection between cortex and striatum.

The mapping of the elements in one structure onto another is studied at the coarse-grained level, as in the case of the mapping of thalamic nuclei onto the cortex; or at a fine-grained level, as in the case of the mapping of cells within a particular thalamic nucleus onto a particular cortical region. Such maps are readily understood on the basis of zone-to-zone or point-to-point connections between the structures, mediated either by bundles of axons (in the case of coarse-grained mapping) or by individual axons (in the case of fine-grained mapping).

In a complex structure such as the cerebral cortex, more complex properties of the sensory environment than position in the field are detected by specific nerve cells. The properties of the stimulus that produces maximal excitation varies over the cortex, defining 'feature maps' (see Section 3.1).

2.3.2 Topographic map formation in the visual system

Most work on topographic maps has been carried out on the vertebrate visual system. Here I describe the direct projection of retina onto optic tectum (the analogue of the superior colliculus in mammals) in amphibia and fish. The first crude maps were constructed from the results of axon degeneration studies but the first maps with any precision were made by extracellular recording from the optic tecta of goldfish, and of the frog, *Rana*, and *Xenopus* (Figure 3b). The connections are not precise at the cell-to-cell level (as in invertebrates) but in the retinotectal map in *Xenopus*, for example, at least 50 recording positions can be distinguished, all arranged in topographic order. The other important attribute of such maps is that they always have a specific orientation. All retinotectal maps are arranged so that temporal retina projects to rostral tectum and dorsal retina to medial tectum (Figure 3b).

The problem of understanding how maps are formed was originally formulated in terms of the establishment of a functional relationship between the cells in the two sets of cells. With the advent of powerful methods for tracing patterns of connections, this problem has become that of specifying the connections themselves

One powerful set of results, generated from electrophysiological recordings carried out in the 1970s, involves experiments on the regeneration of connections in the retinotectal system of adult goldfish. The paradigm experiment involves removing part of the retina, allowing connections to regenerate and then making a map of the remaining part of the retina onto the tectum. After several months, retinal fibres had expanded to innervate the entire tectum. Complementary experiments revealed that after removal of half the tectum from goldfish, the projection from the entire retina eventually becomes compressed, in order, on to the surviving half-tectum. The basic result in these so-called 'mismatch' experiments is that the retina and the tectum match together as a system, independently of the actual sizes of both structures.

Although these results were found in the regeneration of connections, exactly the same type of phenomenon occurs during development. In *Xenopus laevis*, the retina and tectum grow in different ways. New retinal cells are added on in rings to the outside of the retina whereas tectal growth predominantly involves addition of cells to the back. Nonetheless, there is an ordered projection of retina onto tectum from a very early stage. This implies a gradual shifting of connections. For example, throughout development, central retina always projects to central tectum, the position of which moves progressively backwards as more tectal cells are added. This inference was first made from extracellular recordings and then confirmed by electron microscopy studies which demonstrated the degeneration of synapses as axons shifted their positions during development. Later work on frog tadpoles that combined electrophysiological and electron microscopy confirmed this by showing that retinal ganglion cell axons move across the tectum during development, continually changing their tectal partners as they do so.

These results demonstrate that connections cannot be made by means of a simple set of instructions specifying which cell connects to which other cell; more likely, the two populations of cells self-organise their connections so as to ensure the correct overall pattern. Several different hypotheses for the formation of nerve connections in this paradigm system have been made. The two main contenders are:

1. As first proposed by the Nobel laureate Roger Sperry (1963) in his doctrine of chemospecificity, the two populations of nerve cells, one in the retina and one in the tectum, are labelled separately by sets of molecular markers. By some means, the correspondence between the two sets of markers is communicated to the participating cells. Each retinal axon then uses this information to find its correct tectal partner. In addition, there is some means of regulating the constitution of the molecular labels, when required, to account for the lability of connections during development and regeneration. This proposal has received renewed interest recently due to the discovery of a type of receptor located in the retina, the Eph receptor, and the associated neurotrophins located in the tectum, the ephrins, which bind to these receptors. The Ephs and the ephrins could be the markers that label the two sets of cells. The origin of the markers themselves has not yet been linked to the generation of regional specificity discussed in Section 2.1.

2. Initially, a roughly ordered map of connections is made with subsequent refinement of the map through electrical activity.

At present, the status of both contending proposals is unclear. They lack experimental verification at the mechanistic level. For (1), it could be that the tectum acquires its markers from the retina (Willshaw and von der Malsburg, 1979), thereby ensuring that the two sets of markers are properly coordinated. For (2), a Hebbian type synaptic modification mechanism[1] might operate. By reinforcing the contacts made by neighbouring retinal cells to neighbouring tectal cells at the expense of the connections made by non-neighbouring cells (Willshaw and von der Malsburg, 1976), each pair of neighbouring presynaptic cells comes to connect to postsynaptic cells that also are neighbours, resulting in a topographically ordered map.

2.3.3 The elimination of superinnervation from developing muscle

The second part of the two-stage process thought to underlie the development of nerve connections is illustrated by cases where the development of the connections on individual targets has been monitored, as in many vertebrate skeletal muscles. In the adult each muscle fibre is innervated at its endplate by a single motor neuron. This pattern of innervation arises from an initial state in which there is innervation from a number of different axons at a single endplate.

During early development, contacts are withdrawn until the adult configuration is reached (Figure 4). The same pattern of events takes place in the adult after transection of the motor nerve. In the initial stages of reinnervation, muscle fibres are superinnervated and this pattern is transformed into one of single innervation after a few weeks. Muscles vary in size. The soleus muscle is one of the larger muscles in rat with around 3500 fibres innervated by some 25 motor neurons, so that in the adult each motor neuron innervates on

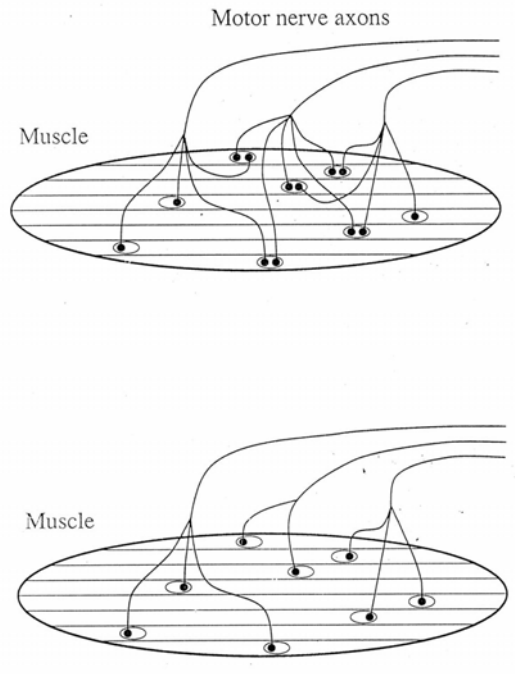

Figure 4: Schematic of the states of innervation of mammalian skeletal muscle.

Top: the initial pattern of superinnervation;

Bottom: the adult state where each muscle fibre has contact from a single axon.

After Rasmussen and Willshaw (1993).

[1] Referring to the hypothesis due to Hebb (1949) that synapses are strengthened by conjoint activity in the presynaptic and the postsynaptic cells.

average 140 muscle fibres; the rat lumbrical muscle has about 100 muscle fibres and 10 motor neurons.

This developmental loss of synaptic contacts has been observed at both central and peripheral sites, in systems as diverse as the neuromuscular junction of invertebrates and the cerebral cortex of primates. The precise time course over which synapse elimination occurs and the proportion of afferents lost varies greatly between areas of the nervous system, even within a single species. In neonatal rat skeletal muscles, there are on average four to five contacts per fibre initially. This reduces to exactly one per muscle fibre over the next two weeks.

In rat cerebellum, the elimination of climbing fibre synapses onto Purkinje cells occurs during the second postnatal week, about the same time course as at the neuromuscular junction. In contrast, The elimination of preganglionic synapses onto neurons of the rat submandibular ganglion occurs over at least five postnatal weeks, far longer than is required for elimination at the neuromuscular junction.

The number of cochlear nerve synapses on neurons of the chick cochlear nucleus declines rapidly, from about four to two afferents and reaches a mature state even before hatching. Therefore, synapse elimination appears to be a widespread phenomenon, although there are no general rules about the percentage of afferents that is lost or the duration of time required.

There is a long and established history of the role of neural activity in the development and regeneration of nerve connections in the neuromuscular system. Tenotomy (cutting the tendon) delays the withdrawal of superinnervation by a moderate amount. Muscle paralysis results in increased, long-lasting levels of polyneuronal innervation, as does application of tetrodotoxin (TTX), a neural activity blocker, to the motor nerve. Chronic muscle stimulation accelerates the elimination of synapses during development.

Researchers can compare the effects of activity with inactivity by applying nerve blocking agents during the reinnervation of neuromuscular connections to a rat muscle with two separate branches of the motor nerve that can be manipulated independently. After cutting or crushing both nerve branches, motor axons regenerate, they superinnervate the muscle fibres and then gradually the pattern of single innervation is reestablished. By blocking one of the nerves with TTX during reinnervation, a competitive advantage can be given to the active synapses as against the inactive synapses (i.e., those made by the axons from the nerve blocked by TTX). It turns out that active synapses have an advantage over inactive synapses. For example, the ability of a regenerating nerve to regain its territory is enhanced if the other nerve is blocked and is diminished if its own nerve is blocked.

The final pattern of one contact per muscle fibre is both clear and unequivocal and several possible causes of the elimination process have been considered. Here is an assessment of them.

- It is unlikely that withdrawal of connections is random as this would leave many fibres uninnervated, contrary to observation.
- Nerve-cell death cannot provide the appropriate reduction in contacts as there is no cell death during this stage of development.
- Some terminals might withdraw if they were misdirected to the wrong region or the wrong fibre type. This is also unlikely as the muscles are almost homogeneous in fibre type and the somatotopic ordering of motor neurons across the muscle is very low.
- Another possibility is that synapses are pre-programmed to die. This is also unlikely as when a proportion of the motor neurons are removed before connections have been established, the surviving motor neurons make more contacts than normal.

This last point provides strong evidence that the identity of the surviving contacts depends on which other contacts are present; i.e., there is competition amongst synaptic contacts for survival.

There are various formal models for this competitive process. According to some models, there is competition for a fixed amount of synaptic strength possessed by the motor axons and shared amongst its terminals; in other models, the making of synapses is thought of in terms of the binding of neurotrophins onto the receptors on different axons, with competition amongst the receptors for the neurotro-

phins. How the various effects of activity can be interpreted is not yet clear. It is interesting that the mathematics underlying these models bears a strong family resemblance to the mathematics underlying many self-organising phenomena studied in physics and the competitive interactions of the predator-prey type studied in population biology and other disciplines.

In summary, the production of the precise connection pattern of one contact per muscle fibre seems to be a task that is not fulfilled by the genome but instead is the responsibility of a competitive, 'self-organising', mechanism involving interactions at the synaptic level.

3 The role of self-organisation in experiential change

I now examine situations where the putative self-organising mechanisms are more strongly influenced by the signals impinging on the system from outside.

3.1 Feature maps

In mammalian cerebral cortex, convergence of inputs onto neocortex causes cortical cells to respond to complex properties of the input. For example, in certain areas of the visual cortex, there are cells that are sensitive to the orientation of a stimulus or its direction of movement as well as its position in the visual field. Such attributes of the external environment can be detected by means of the neural circuitry and the connectivity of the central nervous system. The properties of the stimulus that produces maximal excitation in each small area changes over the cortical surface, defining 'feature maps'.

Figure 5: Computer reconstruction of the pattern of ocular dominance columns in layer IVc of area 17 of a macaque monkey, produced by reduced silver staining, translated into the visual field (Figure from Hubel and Wiesel, 1977).

3.1.1 Ocular dominance, orientation and direction selectivity

Visual cortical cells receive innervation from both eyes, but with differences in the strength of the innervation from each eye that vary systematically across the surface of the cortex (Figure 5). The Nobel laureates Hubel and Wiesel (Hubel et al., 1977) discovered that cells in monkey binocular visual cortex vary in their responsiveness to the two eyes. Similar ocularity preferences extend down to layer IV of neocortex[2], where axons from the lateral geniculate nucleus terminate, and thus the concept of 'ocular dominance columns' arose.

These systematic variations in ocular dominance are superimposed on the basic retinotopic map (Hubel et al., 1977). Subsequently, existence of such columns was confirmed anatomically; the map of ocularity specificity across the entire surface of binocular cortex resembles a pattern of zebra stripes.

[2] One of the six layers in neocortex defined on anatomical criteria.

In cat and monkey, segregation begins at or around birth and is complete about six weeks later. Ocular dominance columns seem to be the result of a competitive process between the axons from the two eyes. They result from an initially overlapping distribution of innervation originating from the left and right eyes. In another preparation, a similar pattern of stripes can be produced through competition within the axons from a single eye. In *Xenoptis*, so-called 'compound eyes,' constructed from two matching half eye rudiments, develop a single optic nerve which produces a striped projection on the optic tectum.

Cells in certain visual areas of mammalian neocortex are orientation selective; i.e., each cell is responsive to a small bar of light when presented in a particular orientation at a particular position in the visual field (Figure 6). The existence of orientation maps was established by extracellular recording and, more recently, the method of optical recording has been used to produce detailed orientation maps over the entire surface of the visual cortex. The maps produced are complex and have a number of features, such as periodically repeating patterns, and more complicated features, such as saddle points and singularities (points on the cortex around which orientation domains are clustered in a pinwheel fashion). This type of data has provided an irresistible challenge to modellers.

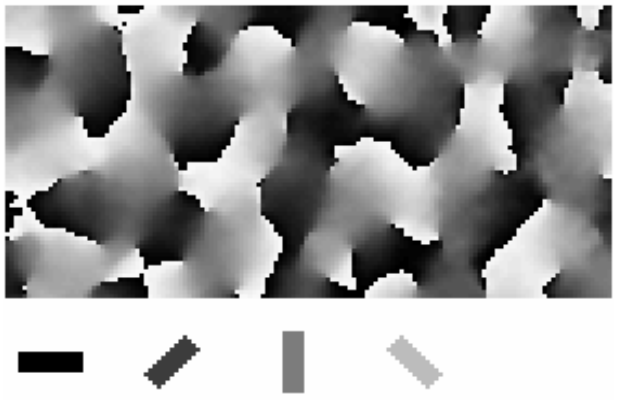

Figure 6: Orientation preference map for area 17 of the cat visual cortex, constructed by optical recording. The orientation at each point is indicated on a grey-scale. The key below shows how orientations are assigned on the scale. The length of the bars corresponds to 0.5 mm. This type of picture is best seen in colour (e.g. Hubener et al., 1997).

The magnitude of the response from some cells elicited by a moving bar stimulus in a particular orientation may depend on the direction of movement (at right angles to the orientation of the bar). This forms a directionally selective map.

3.1.2 Relations between the different feature maps

The different types of map are interrelated. The ocularity map effectively interrupts the retinotopic map; i.e., if all the pieces of cortex innervated by one of the eyes were removed and the remaining pieces, innervated by the other eye, were pushed together, then a completely ordered retinotopic map would result. In cat visual cortex, pinwheel centres in orientation maps are mainly located in the middle of ocular dominance columns. In addition, according to recent optical recording experiments in cat, the orientation domains tend to intersect at right angles the borders of ocular dominance stripes. In the classic model of Hubel et al. (1977), developed for the cat, iso-orientation columns run in straight lines at right angles to the inter digitating ocular dominance columns.

3.1.3 The effect of neural activity

It is well established that lack of normal visual experience early in postnatal life prevents the normal development of the visual cortex. This suggests that some information required to generate the normal visual system derives from interactions of the developing system with the external environment. Clearly, many aspects of cortical development occur prenatally and are immune to the effects of postnatal functional deprivation.

Studies in both cats and primates have shown that before the neonate has received any visual experience, geniculocortical fibres find their main target. In layer IV, they converge to generate immature orientation selective cells clustered into rudimentary orientation columns, form a retinotopic map and, at least in the primate, begin to segregate to form ocular dominance columns.

Visual deprivation early in life, during the so-called critical period, does not abolish these features of early organisation. The continued refinement of cortical connections is strongly influenced by patterned neural activity. Visual deprivation can have a devastating effect on the development of the detailed circuitry required for the normal functional properties of visual cortical neurons. For example, in areas 17 and 18 of the cat[3], the normal appearance of a large proportion of orientation selective cells is prevented by dark-rearing or by binocular eyelid suture.

A variety of results shows how experimental interference can affect the development of ocular dominance columns. Most of these experiments involve the manipulation of activity levels by: monocular deprivation; rearing animals in the dark; distorting the retinal input by artificially inducing strabismus, and removal of spontaneous retinal activity by administering the neural activity blocker TTX. The results of most of these experiments indicate the important role of neural activity in the formation of ocular dominance columns. Thus, deprivation prevents the emergence of the full richness of functional architecture and receptive field properties of the normal adult visual cortex.

The most striking effects of deprivation occur when vision through only one eye is impaired during the critical period. This results in expansion of the cortical territory of the projection serving the normal eye relative to that of the projection serving the deprived eye. Neuronal activity plays a crucial role in this organisation, and it appears that monocular deprivation places the geniculocortical afferents from the deprived eye at a competitive disadvantage. The effects of monocular deprivation can be reversed by opening the deprived eye and closing the other before the end of the critical period (reverse suture). These changes involve the sprouting and/or trimming of geniculocortical arbors. It appears that, as the developing geniculocortical fibres elaborate on their initial framework of immature inputs, it is especially important that the projections serving one eye should be as active as those serving the other eye. A balance of activity is required to ensure that the growth of terminals from each eye is restricted within its own cortical territory.

3.1.4 Self-organisation and the formation of feature maps

Self-organisation plays a role, in combination with external signals, in determining the response properties of the individual cells in feature maps as well as the pattern of response properties distributed over the map itself. In the development of ocularity maps, determination of the ocular preference of an individual cell is influenced heavily by the nature of the afferent activity, whereas the pattern of ocular dominance is the result of an interaction between activity and mechanisms of self-organisation.

In the development of orientation maps, self-organisation may be involved in specifying both single cell properties and the overall pattern. von der Malsburg (1973) published the first paper demonstrating that the pattern of orientation specificity in visual cortex could be developed in a self-organising model under the instruction of simulated orientated bar stimuli. Since then, research has established that patterned stimuli are not needed to develop the individual cell's response properties. Radially symmetric patterns of activity drive the production of orientated receptive fields by a process of symmetry breaking. There has been much recent theoretical work developing models of all types of feature map and the relation between them. In these models, changes in the external conditions (simulating, for example visual deprivation) lead to the model successfully self-organising to adapt to the new conditions.

3.2 Self-organisation and the acquisition of cognitive function

All the examples of self-organisation in the nervous system we have discussed so far concerned the development of the nervous system in cases where we can relate the results of the self-organisation to structural and functional changes at the cellular and sub-cellular level. I now discuss self-organisation and the acquisition of cognitive function.

Perhaps because of the lack of a strong structural basis, less can be said about the precise form of the self-organising mechanisms than in the examples of neural self-organisation discussed already. It will

[3] These are the primary (18) and secondary (17) areas of cat visual cortex.

be seen that many of the themes already discussed are echoed, but with different vocabulary to place an emphasis on the cognitive and psychological levels.

3.2.1 Cognitive self-organisation

The term self-organisation has been used to understand how new cognitive behaviours are acquired. As noted earlier, these behaviours could be derived from built-in knowledge (in cognitive science terms, nativism, attributed in this context to Fodor) or through learning (attributed in this context to Piaget). However, it has been pointed out that through self-organisation, order can emerge from interactions within the system rather than by explicit instruction.

According to Karmiloff-Smith (1992), children's brains are genetically prewired to contain modules (or domains) of knowledge which grow and interact during development. She proposes that in each domain, the child's development is subject to constraints that initially shape the way that information in different domains is processed and represented. The child's initial knowledge in each domain is then progressively redescribed according to existing knowledge and the child's experiences. The level of redescription in each domain thus depends on a complex interaction between the current state of knowledge in the domain and the child's experience. Along with mastering new behaviours, the child also learns to introspect about what he/she has done and ultimately constructs his/her own theories about how the world works.

While accounting for a large body of knowledge, these theories of re-representation are difficult to interpret insofar that no details are given about the mechanistic basis of re-representation. An attempt to make a bridge between the cognitive and neural levels of brain function has been supplied under the name of neural constructivism. This has been developed from Piaget's constructivism, a term that reflects his view that there is an active interaction between the developing system and the environment in which it develops. Neural constructivism emphasises the dynamic interaction between the developing system and the developmental' task' to be accomplished and in this respect it can be regarded as a form of self-organisation.

Neural constructivism lies between the extreme versions of the theories of chemospecificity (the entire blueprint of the nervous system is specified in the genome) and a tabula rasa theory (nothing is prespecified). It has been contrasted with the doctrine of selectionism which, according to some authors, is the idea that neural development proceeds by initial overproduction of neural structure followed by the selective pruning away of the inappropriate parts. We can view selectionism as an extreme form of the two-stage process thought to underlie the development of connectivity, mentioned in Section 2.3. According to selectionism, the first stage is involved in the formation of connections and the second stage in the breaking of connections.

Neural constructivism is concerned with how neural activity guides the development of connections, the patterns of activity being generated from the external environment rather than through, for example, spontaneous activity. It is suggested that, as neocortex evolved in mammals, there was a progression toward more flexible representational structures, rather than there being an increase in the number and complexity of innate, specialised circuits.

Different types of evidence are cited in favour of neural constructivism, some of which were reviewed earlier in this essay.

Changes in synaptic number During development in primates the number of synapses rises and eventually falls but there is no decrease until post-puberty.

Axonal growth In the visual system during the period of activity-dependent development there is evidence for axonal growth, contrary to the doctrine of selectionism.

Dendritic growth Dendrites grow at a much slower rate than axons, over a longer time scale which stretches over the period of time prescribed by selectionism. Dendritic morphology can be moulded by the patterned neural activity incident upon it.

The extent of cortical development Human cortical postnatal development is also more extensive and protracted than generally supposed, suggesting that the neocortex has evolved so as to maximise the capacity of environmental structure to shape its structure and function through constructive learning.

The power of constructive neural networks Finally, arguments based on the computational power of artificial neural networks, used to simulate the phenomena observed, are used to support neural constructivism, particularly the view that networks which develop their structure whilst they are being trained are more powerful than fixed architecture neural networks.

3.2.2 Neural constructivism and neural networks

Neural networks are collections of highly interconnected simple computing units modelled loosely on the architecture of the brain. In such a system, a computing unit is a stylised representation of a nerve cell. The networks are required to learn specific input/output relationships (Hertz et al., 1991; Bishop, 1995) by selective adjustment of connection strengths between neurons. Most applications of neutral networks are to problems in pattern recognition, classification and prediction. Their behaviour is usually investigated through computer simulation.

Neural networks can be trained by two main methods. In supervised learning, many examples of the input/output pairings to be learnt are presented to the network, in the forms of patterns of activity over the computing units, which can be likened to patterns of neural activity. As a result, the strengths of the connections (weights) between individual computing units changes. Ultimately, the weights in the network become set so that the network, when tested, gives the required output for any input presented during learning. Hopefully, it generalises its behaviour to give the appropriate response to an input that it has not seen before.

In unsupervised learning, there is no teaching signal, in the form of the required output being supplied for each input. Presentation of every input generates an output and the network modifies its weight strengths 'by itself,' with the result that in the initial stages the output generated for each input may change. After many presentations, a stable output comes to be associated with each input.

Neural networks have been said to be self-organising in that, in both learning paradigms, learning depends critically on the structure of the network and the interactions between computing units. In supervised learning, both the pattern of weight strengths that emerge in learning a given mapping, and the ability of the network to respond to novel inputs, is self-organised by the network itself. In unsupervised learning, the nature of the input/output mapping produced depends on the interaction between network structure and the nature of the input patterns.

How a given task is learnt, or whether it will be learnt at all, depends on the structure of the network. Using recently developed methods, the structure of the neural network can be built up whilst the task is being learnt, rather than training a network with a fixed architecture. The network structure becomes tailored to the specific problem and in many cases it is claimed to yield better performance than networks with a fixed architecture. By analogy, neural constructivism describes the idea that the development of structures such as the neocortex involves a dynamic interaction between the mechanisms of growth and environmentally driven neural activity. How information is represented in the cortex depends on the particular computational task being learnt.

Neural network models have been used to simulate human performance on specific cognitive tasks. This area of research is known as connectionism. Typically, the modeller prescribes a specific network structure. The strengths (weights) of the connections between 'neurons' are chosen randomly, so that the model does not contain any preprogrammed knowledge.

One early and influential example of the use of neural network models in this context was Rumelhart and McLelland's model of past tense learning (Rumelhart and McClelland, 1986a). They showed that a neural network could learn the mapping between the stem form of a verb and the actual past tense without the need for the rule that prescribes such transformations to be specified externally. In addition, the U-shaped pattern of learning characteristic of a child's learning is produced automatically by application of this model. Whilst this model has been subject to much criticism, its development and applica-

tion did demonstrate that there are alternative ways to viewing language acquisition than the application of preprogrammed rules.

4 Self-organisation as a response to damage

It has long been known that the effects of brain injury early in life are much less severe than similar effects in the adult. Until relatively recently, the commonly held view has been that this is because the brain has a capacity for plasticity during development that can be brought into play as a response to injury. It is now clear that the adult brain has a much higher capability for reorganisation than previously thought. Recent experimental studies illustrate that the adult mammalian nervous system does have substantial capacity to reorganise itself functionally. Plasticity should be thought as a property of both developing and adult systems.

4.1 Self-reorganisation

A landmark study was made by Raisman in the early 1970s on experiments in rats. This involved partial denervation of the septal nuclei, which are part of the limbic system that includes the hippocampal formation and the amygdala. The lateral septum has two main inputs. When either one of these was cut, there was a large temporary reduction in the number of synapses on the septal cells. Over the following two weeks, the number of synapses returned to normal levels but all the synapses were now the type characteristic of the intact input. The same result was obtained whichever septal input was cut. This study is important as it provided the first solid evidence of plasticity at the cellular level in the adult mammalian central nervous system.

Since that time, there have been several studies showing substantial plasticity in mammalian neocortex. Many sensory and motor modalities have multiple representations on the neocortex. Originally it was thought that these maps would be relatively permanent once they have formed. However, it is now established that the neocortex retains a large degree of plasticity throughout adult life.

In the 1980s, Merzenich and colleagues found in monkeys that, following peripheral nerve injury, the ordered mapping of the body surface onto primary somatosensory cortex becomes reorganised substantially. Electrophysiological recordings were made to examine the responsiveness of the areas that had responded to stimulation of the cut sensory nerve.

When recordings were made soon after surgery, it was found that a large part of each of these areas was responsive to cells from neighbouring areas of skin, which normally projected to neighbouring cortical areas. If the area was large, there was a region in the middle from which no response to sensory stimulation could be obtained. Over the next month or two, the cortical representation of nearby somatic areas gradually expanded into the unresponsive area until the entire somatosensory region of the cortex responded to sensory stimulation, making a map that was a reorganised and distorted version of the original one.

Another manipulation that has been tried is removing part of the body instead of denervating it. For example, removal of whole digits caused the affected area of cortex to come under the control of the digits from either side of the ablated ones.

Most likely, this reorganisation involved two different mechanisms. The gradual spreading of the somatosensory projection seen in the long term is likely to involve an anatomical rearrangement involving the regrowth of connections. However, this type of change is relatively slow and so cannot also account for the changes seen after surgery. The most likely explanation is that surgery triggers an unmasking of a population of silent synapses, agreeing with inferences drawn from earlier experiments on the cat spinal cord.

Similar effects have been found in both primary auditory and visual cortex in monkey. The auditory cortex contains a one-dimensional map of frequency, with high frequency tones represented in caudal regions and low frequencies more rostrally. Destruction of nerve fibres in the cochlear which are respon-

sive to high frequencies caused, a few months later, a reorganisation of this tonotopic map with low frequencies being represented rostrally and mid-range frequencies more caudally. In primary visual cortex, small retinal lesions initially produce an area of unresponsive visual cortex. Over the next few months, this cortical region gradually acquires new receptive fields from places in the retina near the site of the lesion.

These effects are also seen in humans. In the 1990s, Ramachandran and colleagues carried out a series of studies on sensory reorganisation following limb amputation. On examination a few weeks after amputation of an arm, patients reported sensations from their phantom limb which were referred to regions of the face and the intact arm. In several cases it was possible to define fields of sensation across regions of the face in which the normal somatotopic organisation of the digits of the hand was maintained. In other cases these same types of map were reported in the region of the operated arm above the level of the amputation.

The most likely explanation of these results is again that the sensory fields in cortex that were adjacent to the part of somatosensory cortex that has suffered deafferentation had invaded the deafferented region. It might be noted that in the normal somatosensory map, representation of the face is near to representations of the hand and the arm.

4.2 Can the nervous system regenerate after all?

The commonly held view is that damaged axons in the mammalian nervous system will regenerate in certain cases but nerve cells will not. Damage to major axon tracts or large areas of nervous tissue leads to permanent loss of function at the neuronal level as neither damaged nor killed axons will regenerate. In the mammalian peripheral nervous system, limited repair is possible as axons can regenerate, leading to the restoration of functional connections with other nerve cells and muscle fibres. In contrast, invertebrates and non-mammalian vertebrates have the capacity to regenerate axons and thereby nerve connections throughout their nervous system.

The view that nerve cells cannot regenerate is being challenged. Recent research shows that an important class of cells, called stem cells, exists in many parts of the adult mammalian nervous system. These cells can differentiate into all types of cell, including nerve cells.

Stem cells have been found in the dentate gyrus of the hippocampus and in the olfactory lobe. It may be that nerve cells can be generated from these cells, even during adult life, possibly to replace damaged nerve cells. One series of experiments observed continuous formation of cells (not identified as nerve cells) in adult mouse neocortex. Targeted destruction of nerve cells in a small region of neocortex then led to a population of new cells being generated, a small proportion being nerve cells. These new nerve cells formed connection with other neural structures suggesting that indeed they were substituting for the set of destroyed cells.

Research into the potentialities of stem cells is a current hot topic and new results are reported frequently. For example, a very recent paper (Mezey et al., 2003) reports post mortem analysis of humans who had received bone marrow transplants as treatment for leukaemia. In the four patients studied, there was evidence for newly generated nerve cells in the brain, which had derived from the bone marrow transplants.

These findings are still preliminary and controversial but may yet change the view that individual nerve cells cannot regenerate.

5 Open questions

In attempting to understand a complex system such as the nervous system, it is important to identify important general principles of operation. ‘Self-organisation’ is one such principle. It manifests itself during both development and the functioning of the nervous system.

This section suggests key questions for activities that could arise out of the work reviewed here, both for the furtherance of research into neuroscience and for the construction of new types of computational ('cognitive') systems. It is worth restating the obvious point that these two activities have different goals. One is concerned with understanding a given system: the other is concerned with designing a system to achieve a particular computational task, where the nature of the internal workings of the system are dictated by the task to be accomplished rather than having to reflect any biological plausibility.

5.1 Questions for the neurosciences

The term self-organisation refers to a postulated mechanism rather than a collection of phenomena such as those embraced by, for example, perception or learning and memory. The nature and scope of other examples of self-organisation remain to be determined. Therefore, questions such as "What is the future for research into neural self-organisation?" are premature. Instead I focus on those aspects of the life sciences which are related to the areas of research described in this essay.

5.1.1 Levels of analysis

This essay describes how self-organisation operates at the synaptic, cellular and network levels. At what other levels may self-organisation apply?

- **The cognitive level** The role of self-organisation within developmental cognition was described briefly in Section 3.2. To make sense of self-organisation at the cognitive level, it is important to be able to identify the nature of the elements that do self-organise. As this type of self-organisation involves extensive regions of the brain, this will require the use of powerful methods for assaying whole brain activity to identify these elements. As mentioned already, modelling is a crucial experimental tool here. To apply computer modelling successfully, models must contain the correct degree of neurobiological realism.
- **The sub-cellular level** Although we now have at our disposal several completely sequenced genomes, we are still for the most part remarkably ignorant about how genes interact and regulate each other to develop the nervous system. A plausible picture of cellular regulation would involve networks of multiple interactions, with feedback between all layers. There is great scope for exploiting the parallels between maps of metabolic pathways and those of gene expression, both of which involve processes of global control by self-organisation. More widely, the similarities between the patterns of connectivity over ecological, neural and biochemical networks are being compared (Lee et al., 2002; Milo et al., 2002). Arguably, perhaps the only system of many interacting elements that has been characterised in detail is the artificial neural network. There must be a parallel, so far unexploited, between neural networks and sub-cellular networks.
- **Crossing the levels** Quite often, testing a postulated mechanism requires observation of phenomena at one level and the identification of mechanism at lower levels. More links between levels are required to achieve this. For example, there is still no solid evidence linking the nature of the nerve connections making up the topographically ordered maps and the underlying molecular mechanisms. This will require a cross-disciplinary approach. In this case, what is needed is: experimental evidence of the nature of the connections made at both the electrophysiological and anatomical levels; evidence for the distribution of signalling molecules or patterns of neural activity amongst nerve cells that carry the information enabling the correct connections to be made; an explanatory framework to link the two levels. New imaging methods that obtain information at many different levels, particularly the protein, synaptic, cellular, network and whole brain levels will be crucial.

5.1.2 The use of mathematical and computer models

Often the consequence of any given hypothesis involving a large number of interacting elements can only be obtained by constructing and analysing the properties of computer and mathematical models. This approach is now recognised within neuroscience as an important means of formalising ideas and

concepts and testing them out for their self-consistency and against the large amount of neuroscientific data that is becoming available, and forms part of the field of neuroinformatics. Many areas of neuroscience described in this essay have profited from the application of computer and mathematical models.

- **Neural simulators** There is an increasing trend towards standard powerful neural simulators that will enable researchers to share data, models and results.
- **New types of model** Completely new types of model will be required in some cases, such as those required to model the regeneration of nerve cells (if indeed this occurs) as described in Section 4.2.
- **Modelling of real nervous systems** The power of the present generation of computers is now sufficient to enable simulation of the complex geometry of the nervous system, which is an important constraint on function. For example, up till now, nerve cells have been often modelled as abstract entities rather than occupying particular positions in a complex three-dimensional environment.

5.1.3 Evidence from invertebrates

This essay has been restricted almost entirely to vertebrates, with very little discussion about the organisation of invertebrate nervous systems. Nonetheless, two general comments can be made.

- **Lessons to be learnt from the evidence** Observation of the same type of phenomena in invertebrates may suggest completely different types of mechanism. For example, much of the evidence at the synaptic and cellular level suggests a very precise and inflexible organisation. For example, in the neuromuscular system of the fruit fly, *Drosophila*, there is a precise and fixed relation between a motor neuron and the muscle fibre that it innervates. There are various possible reasons for this. In this case, for example, it may be that a more subtle form of self-organisation operates; or it may be that for the smaller nervous system the genome can afford to specify precisely all the parameters values needed which have a smaller number of neurons? Knowledge of invertebrate nervous systems can provide another source of inspiration for the physical sciences.
- **A complete explanation** To have a complete knowledge of how the underlying neural substrate generates and controls the animal's behaviour, it is necessary to understand it at many different levels, from genes through proteins, synapses, cells, networks to behaviour. In the essay, emphasis has been on the properties of mammals and other vertebrates. However, it is more likely that the first complete understanding will be obtained for an invertebrate system. One prime example is *Drosophila*. It has a known genome, containing around 15,000 genes, with a variety of complex behaviours including those involving learning and memory. In addition, newly developed imaging techniques can be applied that enable visualisation of nervous system activity at the sub-synaptic, synaptic and cellular and network levels.

5.2 Inspiration for other sciences –'cognitive systems'

Replication or Inspiration? Whereas living and artificial systems are made out of the same basic set of elements, namely atoms, clearly their structures are different – silicon chips are different from pieces of brain. The nature of the substrate constrains the properties of the system. This means that whereas it may be possible to build a system that mimics the input/output relations of a living system, at the mechanistic level the systems will be different and will operate in different ways. As a consequence, results from neuroscience can at best only inspire the construction of new types of computing device rather than lead to the construction of exact replicas of living systems. To take a well known example, neuroscience has inspired the growth of the field of artificial neural networks. This can offer many interesting ideas for the construction of new computing devices themselves rather than give a precise blueprint for such a system.

Three inspirations:

- As emphasised throughout this essay, the mathematics developed to describe neural self-organisation, itself influenced from studies of physical systems, has wide application. Applications of self-organisation to social systems and to economic systems are those that have not been mentioned so far.
- Using principles of self-organisation to solve problems of coordination among autonomously interacting agents, such those that as occur within e-communities, is an obvious specific application.
- Neural self-organisation will serve as an inspiration to self-repair technology as an example of application to a hardware problem.

Acknowledgments

My very grateful thanks to Fiona Jamieson, Stephen Eglen, Kit Longden, David Price and Peter Dayan for reading and commenting on this essay. The errors and omissions are all mine.

6 Further reading

This essay draws on a large body of literature. Instead of providing a long list of papers, I give below a number of texts that cover most but probably not all of the work I have discussed, together with descriptions of their various scopes. In addition, a few important papers are included which are referenced in the text.

Self-organisation There is no book specifically about self-organisation and the nervous system. Kauffman (1993) discusses theoretical aspects of self-organisation in evolution and Camazine et al. (2003) is a very recent book discussing self-organisation in biological communities, typical examples being aggregates of bacteria, communities of fireflies, fish and ants.

Mathematical basis The book by Murray (1993) is a classic, concentrating on the mathematical basis of theoretical developmental biology; Edelstein-Keshet (1987) is a very good alternative.

Development Alberts et al. (1994) is a text at the molecular biology level and Wolpert (1991) is a readable introduction to embryonic development for non-specialists.

Nervous system Nicholls et al. (2001) is a classic treatise on the nervous system, first published in 1977 which has undergone continual revision since then. Shepherd (1994); Bear (2001) are recent research level text books.

Development of the nervous system Purves and Lichtman (1985); Sanes et al. (2000); Brown et al. (2001) are research level texts on the development of the nervous system, the latter two being more recent. The recent research monograph by Price and Willshaw (2000) describes genetic, molecular, systems and modelling approaches to understanding neocortical development. Elman (1996) takes a connectionist approach to development.

Neural networks Excellent texts amongst the plethora available are those by Hertz et al. (1991) and Bishop (1995). Rumelhart and McClelland (1986a,b) are two collections of classic papers. Arbib (2003) is an encyclopaedic collection of short papers written by experts on many different aspects of theories of brain function, connectionism and artificial neural networks. As the name suggests, the treatment of the nervous system is generally at the cellular and network level.

Brain damage and repair Fawcett et al. (2001) is a collection of papers reviewing recent research in this field.

7 References

Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., and Watson, J. D. (1994). Molecular biology of the cell. Garland Publishing Inc, New York, USA.

Arbib, M. A., editor (2003). The handbook of brain theory and neural networks. MIT Press, 2nd edition.

Bear, M. F. (2001). Neuroscience : exploring the brain. Lippincott Williams and Wilkins, Baltimore, 2nd edition.

Bishop, C. M. (1995). Neural networks for pattern recognition. Clarendon Press, Oxford.

Brodmann, K. (1909). Vergleichende Lokalisationslehre der Groshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. J. A. Barth, Leipzig.

Brown, M., Keynes, R., and Lumsden, A. (2001), The developing brain, Oxford University Press.

Camazine, S., Deneuborg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., and Bonebeau, E. (2003). Self-Organisation in Biological Systems. Princeton University Press.

Edelstein-Keshet, L. (1987). Mathematical Models in Biology, Birkhauser.

Elman, J. L. (1996). Rethinking innateness: a connectionist perspective on development, MIT Press.

Fawcett, J. W., Rosser, A. E., and Dunnet, S. B. (2001). Brain damage, brain repair. Oxford University Press.

Hebb, D. (1949). The organisation of behavior. Wiley, New York.

Hertz, J., Krogh, A., and Palmer, R. G. (1991). Introduction to the theory of neural computation. Addison-Wesley.

Hubel, D. H., Wiesel, T. N., and LeVay, S. (1977). Plasticity of ocular dominance columns in monkey striate cortex. *Phil. Trans. R. Soc. London Ser. B*, **278**:377-409.

Karmiloff-Smith, A. (1992). Beyond modularity: a developmental perspective on cognitive science. MIT Press.

Kauffman, S. A. (1993). The Origins of Order. Oxford University Press.

Lee, T. I., Rinaldi, N. J., Robert, F., Odom, D. T., et al. (2002). Transcriptional regulatory networks in saccharomyces cerevisiae. *Science*, **298**:799-804.

Meinhardt, H. (1982). Models of Biological Pattern Formation. Academic Press.

Mezey, E., Key, S., Vogelsang, G., Szalayova, I., Lange, G., and Crain, B. (2003). Transplanted bone marrow generates new neurons in human brains. *Proc. Natl. Acad. Sci. USA*, **100**:1364-1369.

Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., and Alon, U. (2002). Network motifs: Simple building blocks of complex networks. *Science*, **298**:824-827.

Murray, J. D. (1993). Mathematical biology. Springer-Verlag, 2nd edition.

Nicholls, J. G., Fuchs, P. A., Martin, A. R., and Wallace, B. G. (2001). From Neuron to Brain. Sinauer Associates, Sunderland, Mass., 4th edition.

Price, D. J. and Willshaw, D. J. (2000). Mechanisms of cortical development. Oxford University Press.

Purves, D. and Lichtman, J. W. (1985). Principles of Neural Development. Sinauer Associates, Sunderland, MA.

Rasmussen, C. E. and Willshaw, D. (1993). Presynaptic and postsynaptic competition in models for the development of neuromuscular connections. *Biol. Cybern.*, **68**:409-419.

Rumelhart, D. E. and McClelland, J. L. (1986a). Parallel distributed processing: explorations in the microstructure of cognition. Vol. 1, Foundations. MIT Press.

Rumelhart, D. E. and McClelland, J. L. (1986b). Parallel distributed processing: explorations in the microstructure of cognition. Vol. 2, Psychological and biological models. MIT Press.

Sanes, D. H., Reh, T. A., and Harris, W. A. (2000). Development of the nervous system. Academic Press.

Shepherd, G. M. (1994). Neurobiology. Oxford University Press, 3rd edition.

Spemann, H. (1938). Embryonic development and induction. Yale University Press, New Haven. Reprinted, New York: Garland, 1988.

Sperry, R. W. (1963). Chemoaffinity in the orderly growth of nerve fiber patterns and connections. *Proc. Nat. Acad. Sci., USA*, **50**:703-710.

Sur, M., Garraghty, P E., and Roe, A. W. (1988). Experimentally induced visual projections in auditory thalamus and cortex. *Science*, **242**:1434-1441.

Tootell, R. B. H., Silverman, M. S., Switkes, E., and De Valois, R. L. (1982). Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science*, **218**:902-904.

Turing, A. M. (1952). The chemical basis of morphogenesis. *Phil. Trans. R. Soc. London Ser. B*, **237**:37-72.

von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex, *Kybernetik*, **14**:85-100.

Willshaw, D. J. and von der Malsburg, C. (1976). How patterned neural connexions can be set up by self-organisation. *Proc. Roy. Soc. London B*, **194**:431-445.

Willshaw, D. J. and von der Malsburg, C. (1979). A marker induction mechanism for the establishment of ordered neural mappings: Its application to the retinotectal problem. *Proc. Roy. Soc. London B*, **287**:203-234.

Wolpert, L. (1969). Positional information and the spatial pattern of cellular differentiation. *J. Theor. Biol.*, **25**:1-47.

Wolpert, L. (1991). The Triumph of the Embryo. Oxford University Press.

Woolsey, T. A. and van der Loos, H. (1970). The description of a cortical field composed of discrete cytoarchitectonic units. *Brain Res.*, **17**:205-242.