# Mind as Creative Organization:

## The nature and development of human cognition as a creative process

Matthew Fisher

Thesis for Master of Arts (Research)

Philosophy
School of Humanities
Faculty of Humanities and Social Sciences
Adelaide University

March 2004

**MIND AS CREATIVE ORGANISATION:**

**The nature and development of human cognition as a creative process**

Matthew Fisher - March 2004

Thesis for M.A. (Research), Philosophy
School of Humanities, Faculty of Humanities and Social Sciences,
Adelaide University

# ABSTRACT

The central subject of this thesis is the nature and development of Mind and cognition in human beings. It argues that the Mind, as it develops through the pre and post-natal periods of life, displays a number of key characteristics that together amount to a form of *creative organisation*.

In addition to this main theme, the work attempts to locate and frame this argument in two ways:

- Firstly, to ask how the issues raised, within the scope of cognitive science and philosophy of Mind, might be relevant to wider social and political concerns.
- Secondly, to examine relevant contextual ideas in ontology and biology, such as: emergence, complexity and organization, and the features of simple organisms

The progress of the discussion is as follows: after an introductory chapter, the thesis examines some background issues in ontology; in particular the idea of emergence. It moves on to consider some features of simple biological systems, and outlines a basic model of creative organisation on that basis. The main body of argument then proceeds, focusing on several areas directly related to the nature and development of cognition; suggesting that the key features of creative organisation continue to be displayed throughout:

- Morphogenesis and developmental processes that underpin and enable cognition
- Connectionist models of cognition and how these inform an understanding of cognitive development
- The nature of moment-by-moment as cognition.

In the final chapter, the thesis orients these central themes towards consideration of a number of wider concerns; questions of how we interpret issues of human nature and Mind in our social or political lives, and of what science offers into these socially mediated domains.

**Matthew Fisher** • 5 Bennett St Brighton 5048 •
(h) 8296 6214 •

# MIND AS CREATIVE ORGANISATION:

## The nature and development of human cognition as a creative process

Matthew Fisher - March 2004

Thesis for M.A. (Research), Philosophy
School of Humanities, Faculty of Humanities and Social Sciences,
Adelaide University

SUPPORTING STATEMENT

I attest that this thesis contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of the thesis, when deposited in the University Library, being available for loan and photocopying.

Yours sincerely,

Signed ... ..

Date ..........24-3-04...............

# CONTENTS

**"MIND AS CREATIVE ORGANISATION"**

# CHAPTER 1: BACKGROUND TO THE STUDY

## 1.1 INTRODUCTION

The central parts of this study focus on questions of our evolved human nature, in particular the nature of *Mind* in human beings. I will argue that the biological and evolutionary antecedents of human life manifest a broad dynamic of *creative organisation*; and that the phenomenon of the human Mind, emerging from this background, continues and extends upon this creative character.

Secondarily, however, the work is also intended to orient these central themes towards consideration of a number of wider concerns; questions of how we *interpret* issues of human nature and Mind in terms of our social or political lives, and of what contemporary science offers into these socially mediated domains. Here are located the overarching motivations that have led to the work being undertaken at all.

Because these motivations are, perhaps, somewhat unusual as a pathway into the particular subject matter of Mind, it is appropriate to identify them straight away. Although they will not figure prominently in the body of the study, they will continue to be present as background, informing the content, style and direction of the argument, and then coming out more explicitly in the final chapter.

### 1.1.1 Initial Aims & Motivations

The background of experience brought to this study does not include a long record in academic philosophy. Rather, it is, amongst other things, a history of active personal engagement with contemporary social and political issues. On the basis of that experience I would assert that contemporary Australian society, as part of an increasingly 'globalised' world, faces a number of significant problems; circumstances that, if allowed to continue unchecked, will seriously undermine our long-term future health and well-being, both as individuals and as a social body. Allow me to simply indicate some main headings:

- **Ecology**: global warming; land degradation; population growth; loss of biodiversity; over-consumption of natural resources; in general, problems with the current health and long-term stability of global ecosystems, and how these impact on human welfare.
- **Social and political stability**: an on-going polarisation of wealth, resources and political power on a global scale; fundamentalist ideologies and extremist politics; rapid cultural and technological change; war; refugee movement; family breakdown; erosion of shared values.
- **Public health**: AIDS and other diseases; food poverty; depression and other mental illness; domestic violence; social disaffection.

(E.g. Brown 1990; Ehrlich &Ehrlich 1991; Shearman & Sauer-Thompson 1997; Shiva 1991; Suzuki 1993; Trainer 1996)

In relation to such problems, and at risk of stating the obvious, I would make two further observations. Firstly, it is clear that most if not all of the significant problems we face as a society are, in one sense or another, circumstances in which we ourselves, human beings, are deeply implicated. In some cases they are a direct result of human activity, in others, they are exacerbated by human actions (or culpable inactions). In every case, I would suggest, the only real likelihood of constructive change to reduce or resolve significant social problems lies in our own hands. In the main, it will happen (and does sometimes happen) through deliberative human action, or not at all. The general truth of this I think is self-evident.

Secondly, I would assert that the real-world actions we take are often *shaped by individually and/or collectively held beliefs about who and what we human beings are*. To unpack this assertion a little, this could include beliefs formulated (roughly) in such terms as:

- **The kind of being we are –** we *are*: competitive, violent primates; God's creation; intelligent rationalists; acquisitive consumers; etc. (or some combination of such).
- **The kinds of values we ought to hold –** we *should*: obey religious laws; compete in free market relations; respect human rights; Recognise the inherent value of all living things; do whatever we feel like; keep white-skinned and black-skinned people separate; etc. (or, again, a combination of multiple options).
- **The kind of potential we embody –** to realise our *individual* potential we *should*: be creative; exercise economic power; receive a formal, Western-style education; or what have you. To realise our *collective* potential we *ought to*: value the nuclear family; support free market capitalism; foster democratic institutions; grow the economy faster; and so on.

In other words, I would assert, beliefs of this kind are a significant influence on our real-world behaviours. For example, if I strongly believed in the picture of humanity (and of myself) drawn by the doctrines of a Christian church, my behaviour would in all likelihood be significantly different from what it is now, when I do not.

If this is so, then it would seem to follow that such beliefs may also be a significant influence on our response to the kinds of problems indicated above. Indeed, one need not look very far to find contemporary examples where different beliefs about our nature and associated values seem, on reflection, clearly linked to different positions on such current issues. In this way, fundamental beliefs about our nature and values as human beings take on significant political[1] relevance. For example, the political dynamics of the Israeli-Palestinian conflict would seem very clearly to be partially driven by underlying, deeply held beliefs on both sides about 'who *we* are' (and 'who *they* are').

In seeking to confront some of the explicit challenges facing our society, therefore, one might reasonably look beyond the overt facts of any particular 'problem' as such, and come to examine what is, broadly defined, a political relationship, between our own various *beliefs* about the world and ourselves, and the real world *effects* of these beliefs on our behaviour and forms of social organisation. To what extent, one may ask, are the various beliefs and associated values we hold – in so far as they do shape our actions – either implicated in the problems we make or, conversely, adequate to any putative solutions? And if they are indeed both implicated and, perhaps, currently less than wholly adequate, are they amenable to change and what might be the alternatives?

Such questions are certainly appropriate for consideration by philosophers, and on the whole it would seem a part of the whole *raison d'etre* of philosophy to believe that better, deeper, more truthful ideas will in the main serve rather than hinder positive, life-enhancing social capacities.

It is this kind of reasoning that has led me, from an initial set of sociopolitical concerns and aspirations, to consider more directly our beliefs about 'human nature' and the way these (positively or negatively) inform our social dynamics. It has led me to conclude, indeed, that without explicit consideration of these issues, our long-term prospects for social development beyond current challenges are seriously reduced. Even leaving aside the relative merits of various

[1] 'Politics' in this context is conceived broadly, as individual or collective action taken to effect matters of social organisation and resource use; particularly actions undertaken by governments, political parties or other influential groups or institutions.

beliefs; if our more explicitly political or ethical debates are being shaped by underlying beliefs that are *not* in themselves brought to light, or subject to direct enquiry, then surely the long-term potential for a static polarization of views is increased.

### 1.1.2 The Role of Contemporary Philosophy of Science

On the whole it is appropriate, I believe, for the business of science to be conducted, and allowed to occur, as far as is possible within a climate of political disinterest. In support of such a value, it also seems reasonable to assert that, as a general rule, the information offered by science and the socio-political uses of that information, can be treated as distinct domains – and that the distinction is important.

Nevertheless, it is also equally true that the information, insights and capacities generated by science do lead, sometimes, to profound social consequences. One could not reasonably argue, in my view, that the particular physicists involved in the United States' nuclear weapons development program during WWII are wholly or directly responsible for the existence and use of these weapons. There can be no doubt, however, that their work informed the development of these weapons, and in this sense did in fact allow for their use.

Equally, we would be naive to think that scientists are somehow immune to political beliefs and biases, or that these never enter directly or indirectly into the ways individual scientists interpret the world, or how they communicate their ideas and theories to others. (Rose et al 1984, Ch. 1) And, of course, a lively debate between science and scientists and other social players about the interpretation and use of scientific knowledge is an appropriate and visible part of everyday social discourse.

It is important therefore that those engaged in science, while valuing a stance of disinterested neutrality, should – like any good citizens – consider the actual or potential social implications of their work, both in itself and in the ways it is communicated. Indeed, we do increasingly see formal mechanisms in place to do things in this vein, such as the ethics committees functioning in many university science departments.

Clearly Western science has over time informed a number of more widely held beliefs and understandings in the vein of those mentioned above – broadly speaking 'beliefs about human nature'. For example, it is a widely canvassed view of modern history that the growing influence of science-based ideas played a significant role in the so-called 'Age of Enlightenment' in 17th and 18th century Western Europe, whereby the influence of religious doctrine over the popular mind was significantly unsettled, and human reason was put forward as a source in its own right of ideas, values and understanding of the world. (E.g. Saul 1993)

Or one could consider the role of Darwinian theories of evolution in fuelling the development of so-called social-Darwinist views of human beings as fundamentally competitive creatures, and the extended economic and social values aligned with that basic belief. (Hofstadter 1955) While it might be argued that such politicized interpretations are ultimately inadequate to their source material, one could hardly argue that there is no influential relationship between the two.

So, in recognizing the reality of such exchanges between scientific insights and social beliefs, it is appropriate and useful for philosophers of the sciences to ask how scientific ideas, and science-based philosophies, currently might inform popular beliefs about 'human nature' – and then how they may also shape our real-world actions guided by those beliefs – and how, if in any way, they might better inform those beliefs? The main content of this study is intended (for my own part) as a ground from which such questions can be considered. The philosopher, social commentator and linguist, Noam Chomsky, also alludes to the potential social import of such inquiries:

> To command attention and respect, a social theory should be grounded on some conception of human needs and human rights, and in turn, the human nature that must be presupposed in any serious account of the origin and character of these needs and rights.
>
> Chomsky 1988, p. 195

In my view Chomsky has the order of ideas right here, locating an account of human nature at the foundation of subsequent social theory. I would only assert again that this economy of ideas is not limited to formal social theorizing, but is far more widespread; so that, beneath the surface of most if not all seriously held socio-political ideas, one is likely to find underpinning beliefs that are properly characterised in the same way; as about human nature – about who or what we are in this more fundamental sense. These ideas may at times be poorly formed, partial, or accepted uncritically (or sub-consciously) as 'just the way things are', and still be exercising considerable influence on more day-to-day attitudes and ideas. (Fromm 1991, Ch. 2)

Furthermore, in initially contemplating the role of the sciences, one might conclude that the prominence of examples like evolution theory is not surprising, simply because the physical or biological sciences, whether intentionally or not, can have the effect of locating the human experience within a wider physical and/or biological world and history. Thus science, in a sense, offers its own versions of what religion has typically provided, a wider worldview within which we may contemplate our own place and meaning. As will become clear as the argument unfolds, I would also generally recognise the importance of locating the human within the physical and biological realms, for the purposes of clarifying a 'base-line' view of our nature, and therefore also as a sensibly prior task to more specific discussion of Mind and other matters.

Finally, it bears saying what is probably becoming clear to the reader anyway. The argument presented is also motivated by a view that *some* of the more popular, socially mediated beliefs about human nature – informed by ideas derived from science – are lacking in some respects, and that better alternatives are possible. The lacks are not all necessarily due to the state of scientific knowledge as such; but may be in part simply a result of poor (or slow) translation of ideas between science and the popular understanding. However, be that as it may, if science is implicated – either actively or by omission – in the perpetuation of less than adequate ideas which in turn contribute to (or fail to assist much with) the perpetuation of significant social problems, and implicated in fact because better alternatives *are* indeed possible, then this ought to be a matter of concern to scientists and philosophers alike.

## 1.2 THE SIS SUMMARY

As stated above, the main goal of this study is to develop a view of Mind as creative organisation. However, the wider social concerns expressed – about our social interpretations of 'human nature'; their real world effects; and the influences of science in this regard – all act to shape the approach taken towards this goal. More specifically, I have found it useful to preface an argument about Mind with a survey of some relevant background issues in the physical and biological sciences. It is to say, in other words, that if these areas are fundamental to a scientific worldview, then they must also be basic to what we 'get out of' science by way of self-understanding and social interpretation.

It is also the case that Mind can be seen – with scientific eyes – through different timescales: that Mind can be understood as an outcome of events on the timescales of: universal history, or terrestrial, evolutionary history, or the ontogeny of an individual person, or the more immediate nature of real-time cognition. Indeed one might say it is an outcome of all of those histories simultaneously. The progress of discussion through the body of the study thus also aims to read matters through several of these lenses, and to track Mind down, as it were, through progressively shorter timescales; beginning with a universal, physical view, and moving through a broadly biological view, before coming to focus on some of the specific morphogenetic and developmental processes that underpin and enable human cognition, and finally considering the moment-by-moment nature of Mind as cognition.

The specific view of Mind to be developed also, I hope, grows by degrees through the stages of the discussion. Initially, some of the primary concepts underpinning this view are located in our understanding of the physical world. Proceeding from that, it is suggested that biological life in all forms manifests a creative dynamic of sorts, and that this also extends into the nature of human biology and morphogenesis. Finally, this creative dynamic is seen to enter into the character of Mind itself.

In the final analysis (or synthesis, perhaps), the discussion as a whole offers a view of Mind as something that has come about through events and processes with similar creative characteristics: initially – through the lens of a universal timescale – these characteristics 'appear' in partial ways, and in 'developmental' changes highly attenuated in time; subsequently they manifest in processes over shorter 'biological' timescales; and in Mind they are seen to become more 'condensed' still, into the moment-by-moment timescales of working cognition. Thus contemporary science and philosophy is seen to offer a perspective on Mind in human beings as a particular realisation of the nature of complex, emergent, living systems in a material universe; and as a process displaying a character of *creative organisation*.

It is then only after this view of Mind has been put forward, that the study links back to the issues raised in this introductory chapter, and concludes with a brief discussion of how the view of Mind presented might begin to be extended into an explicitly social and political discourse.

Lastly, the structure of the study also reflects the development of my own understanding over the last several years. It reflects the scope of issues that I have been interested in, and found it necessary to grapple with; all shaped by a desire for a founding perspective on *us*, the human animal, from which some more coherent social understanding might begin to arise.

### 1.2.1 Chapter Summary

- Chapter 1 – as above.
- Chapter 2 surveys some basic, contemporary science-based perspectives on the physical world, including a pluralist ontology, the key concept of emergence, and aspects of the reductionist tendency in science.
- Chapter 3 considers some characteristics of simple biological organisms, and proposes an initial description of creative organisation in that context.
- Chapter 4 sets the terms for the discussion of Mind as cognition, and outlines a range of background issues in cognitive science
- Chapter 5 puts a model of creative organisation in more detail, and discusses processes of human morphogenesis in those terms.
- Chapter 6 considers connectionist models of cognition and the nature of representation. These ideas are then applied to the realisation of representational capacity in humans, again as a process of creative organisation.
- Chapter 7 draws upon the argument of Chapter 6 in putting a case for moment-by-moment cognition as embodied, embedded and creative.
- Chapter 8 briefly surveys issues and prospects for the socio-political interpretation of Mind-as-creative-organisation.

### 1.2.2 A Note on Method

The structure and content of this thesis requires it to range across a number of major areas in the philosophy of science. It becomes necessary therefore to only touch upon selected ideas within these areas, rather than go too deeply into topics each of which could warrant detailed study and argument in its own right. Given this necessity, I will try at appropriate points simple to indicate the assumptions made for the purposes of the discussion. In general terms, one might say that it is not a densely argued work on a specific area of thought, but an attempt to draw concepts and ideas from a range of areas together within a particular view of Mind; and then to ask what this might mean from a social perspective.

**"MIND AS CREATIVE ORGANISATION"**

# CHAPTER 2: BACKGROUND PERSPECTIVES: ONTOLOGY, EMERGENCE & REDUCTIONISM

## 2.1 PRELIMINARY PERSPECTIVES

The main purpose of this chapter is to provide an overview of some essential ideas and debates shaping contemporary science, ideas that also underpin the view of Mind to be considered in Chapters 5-7. Amongst the ideas surveyed will be the notion of *emergence*, as part of a naturalistic and *pluralist* ontology. The discussion will also briefly consider some issues in relation to so-called *reductionist* approaches to science, insofar as they relate to the wider goals of the work.

In general terms, the perspectives canvassed here are grounded in a pragmatic science-based naturalism. On the whole I would accept a view of the world that a) generally takes things to be composed of natural entities as studied by the physical sciences; and b) endorses methods of argument and explanation that are consistent, broadly speaking, with those in science. I also accept that modern scientific inquiry, generally speaking, does offer access to real (but not necessarily complete) understanding of a real world; a world that exists independently of our individual experience of it.

With regard to methodology, I have a general preference for the spirit of pragmatism, open enquiry and diversity that much of Western science has displayed; at least in the sense of simply taking the world to be real and consistent, and proceeding to investigate it on that basis. As others have argued, one may take the historical record of success for science in offering reliable understanding of the world, and ameliorating the human condition in some substantial ways, as support for such a choice. (E.g. Popper 1974, pp. 260-269)

On the other hand, I would tend to disagree with any tendency to over-inflate science (or quasi-science) into ideological claims, or in any other way to believe uncritically that science in itself can solve all our problems.

### 2.1.1 Selection and Classification

If we assume that the world is both real and intelligible to us; we may ask whether there are any terms that might define and further clarify the notion of intelligibility in this context.

In the first instance, one might suggest that our knowledge of the world is subject to some inherent limits; limits represented most clearly, perhaps, by the very facts of our own individual and collective limits in time and space. So, for example, while it may be true that an individual human being can in some way adequately contemplate 'the universe' as one thing, it would seem clear that he or she is not capable of 'knowing', or being aware of, every actual thing existing or every actual event occurring in the universe at any particular time. Indeed, contemporary understanding of complexity in the universe would suggest that even on a far more modest scale similar limits apply. In a particular volume of gas, for example, we may understand some things about the type of molecular 'behaviours' and interactions occurring within it; we quite clearly cannot (and do not find it necessary to try to) 'know' or measure or describe the position and velocity of each and every molecule making up this volume of gas at any particular moment. (Auyang 1998, p. 50)

So, it would seem clear that the individual and collective (or shared) understandings of the world we generate are clearly subject to some limits, even though in a specific instance the precise boundaries of those limits might be difficult to define. Rather than contemplating limits in this sense – as definitions of boundaries – what such hypothetical cases indicate is simply that in knowing the world we necessarily *select* from the notional totality in which we are immersed. In identifying this necessary act of selection we may thus draw a distinction between our knowledge of the world and the world in itself. We may come to recognise that 'no knowledge can be established about the *whole* of nature. Given our finite powers, every investigation is and must be partial and perspectival. Objective knowledge works through abstraction.' (Grene 1971, p. 25)

Bringing out the notion of selection does two things that are relevant to this study. Firstly, it introduces an element of 'subjectivity' to the concept of knowledge, to the extent at least of suggesting that our knowledge of the world is never entirely detached from our own human participation in its realisation. This is not to deny either an objective reality or the validity of shared knowledge. It is to say that the genuine knowledge we do have, and share, is nevertheless a human knowledge, and it is formed and carried within the frameworks of individual and shared human experience. We do *actively* select from the totality, and 'narrow the view' as it were, in order to know the world; and therefore, in this sense, genuine knowledge of the world will always carry with it both a subjective and objective aspect.

As physicists, Bohm and Peat, comment:

> We suggest that there is indeed a meaning to reality that lies outside ourselves but that it is necessary that we, too, should be included in an essential way as participators in this reality. Our knowledge of the universe is derived from this act of participation which involves ourselves, our senses, the instruments used in experiments, and the ways we communicate and choose to describe nature. This knowledge is therefore both subjective and objective in nature. ...
>
> The fact that this knowledge can bring order to experience and even predict new kinds of experience shows that it must in some way be directly related to a reality beyond this knowledge alone.
>
> 1989, pp. 55-56

This conceptualisation of knowledge is regarded here as generally useful and appropriate, portraying scientific knowledge not as an idealised absolute which it is our business to progressively conquer, or as a key to limitless manipulation of the world; but as something limited and perspectival, but nonetheless real, about real things, and useful.

Recognition of selection in knowledge also introduces a further important idea, that the very act of knowing in turn requires a primary level of *classification*. If the notion of selection applies then it would seem inevitably to follow that the selection of any one thing as the subject of one's attention or interest implies some criteria by which the distinction of that thing from the general field of being occurs. It would seem quite implicit in our whole construction of knowledge about the world, in other words, that we distinguish *things* to know about, and not only that but *types of things*. To see an animal in front of one and say (correctly), 'That is a dog' is not only to identify that particular dog as a distinct thing in itself, but also to classify it as a type of thing.

This essentially classificatory approach is so much bound into our knowledge structures that, ultimately, it becomes meaningless to conceive of the two as separate. As philosopher and physicist, Sunny Auyang puts it:

> The categorical framework contains our most basic and general presuppositions about the intelligible world and our status in it. It is not imposed externally but is already embodied in our objective thoughts as oxygen is integrated in the blood of breathing organisms.
>
> 1998, p.ix

(Also: Maturana and Varela 1987, p.40; Bateson 1972, p. 456)

If this is so, then we may certainly ponder, debate, refine or change the basis of a particular classification of one type of things vis a vis another, but to reject classification itself would be to deny the very structures of knowledge. This study accepts such an argument and thus recognises classification in this general sense as one of the essential terms on which our knowledge of the world proceeds.

To put this another way, it is accepted here, as a generally necessary feature of our knowing the world, that we identify things in the world, and types of things, and that these identifications (often) correspond to (but do not entirely encompass) characteristics of an independently real world. There are real things in the world that I distinguish as 'that cup on my table', or 'that dog in the corner', and the distinctness of them (to me) corresponds with certain facts about their real existence, independent of my awareness of them.

To take this as a provisional and pragmatic view of reality from which to begin does not yet imply anything much about the more specific terms on which such distinctions of things are made, and classes of things are identified, and distinguished in turn from other classes. It is to recognise that, in general terms, this is indeed what we do, so it may as well be initially accepted, so that we may move on to ask what these primary distinctions mean, and what collectively they offer by way of a general understanding of the world. The discussion that follows may indeed be thought of as about contemporary science's classifications of the world, and the specific terms on which that takes place.

## 2.2 Ontological Foundations

### 2.2.1 A 'Developing' Universe

It is now the most common view within science that the universe we inhabit had its origins in a profound cosmological event, usually referred to as the 'Big Bang'; an event which brought into being not only all the material substance of the universe, all the physical matter, but also space and time. Following on from this theory of origin, it is further understood that the material universe in its early life after the initial event consisted mainly of the most basic physical materials, spread diffusely through space.

If we take this description of the 'early' universe as generally true, and compare such a state of affairs with the world we see around us now – the specific forms of cosmological being, planets, stars, galaxies and so on; the diversity of life; and our own existence – then one broad conclusion would seem inescapable. Somehow, within the life of the universe, change through time has occurred, and the initial diffuse simplicity has given rise to the present diversity of form of which we are a part. In the most basic of terms, by making such a comparison we are led to consider some form of cosmological 'development'[2], a process of change in the physical universe, and a *coming into being* by some means of larger, more complex material forms *out of* the prior elemental diffusion. (Davies 1995)

Furthermore, of course, alongside this preliminary sketch-view of cosmology, we have the whole raft of insights from science about the biological evolution of life on Earth. Here again, as a 'minor' theme within a larger cosmological history, we have a description of beginnings in a diffusion of the simplest water-borne life forms, moving through processes of change into greater diversity of size and form and behaviour. The oldest fossil records of life on earth, dating back some 3.5 billion years, are of non-nucleated, single-cell life forms, or prokaryotes. It is only in the last .5 billion years that large, multi-celled organisms have appeared, including, of course, ourselves. (Margulis, 1987, pp. 98-99)

Taken in themselves, neither of these most basic descriptions of the contemporary scientific world-view is regarded as controversial. More open to debate and discussion, of course, are the various descriptive, classificatory and explanatory frameworks brought to bear within such a generally agreed understanding. In overall terms, one might say that, the bare evidences of these 'developmental' processes of origin and change are now amongst the major conceptual foundations of contemporary scientific investigation.

On the basis of the discussion above – about the terms on which we know the world – such bodies of theory and knowledge clearly imply classificatory distinctions between things and types of things. If we provisionally accept such accounts as these, along with their history of associated supporting arguments, evidences and experimental substantiation, then we may consider some of the specific terms on which they classify and codify the world.

[2] Although the term 'development' might imply one self-consistent process, its use here only as a convenient description of processes of cosmological change, whereby a greater diversity and complexity of forms has come to be.

### 2.2.2 Materialism, Complexity and Composition

It should be noted, first of all, that the overall picture of a 'developed' and naturalistic universe begins with the positing of an essentially simple and universal 'substrate' of physical materiality; that 'all known stable matter in the universe is made up of three kinds of elementary particle coupled via four kinds of fundamental interaction'. 'The ... particles are the electron and two light quarks. ... The four kinds of fundamental interaction are gravity, electromagnetism, and the two nuclear interactions.' (Auyang 1998, p.1 & p. 347) Assuming this sub-atomic foundation of the universe's material being, and of its 'development' through time, it generally follows that all of the larger structures and systems we observe in the world are, in some sense, '*built*' from this physical material. Here two things must be recognised: firstly, this is, in simple terms, the *materialist* assumption that is basic to most contemporary science-based ontology. Subject to the discussion to follow, this assumption is accepted here. Secondly, on the basis of this assumption we are confronted with the notion of relative differences in what I shall refer to as *aggregational complexity*.

If something we regard as real, a specific dog for instance, is indeed 'built' in this sense (somehow) then it would seem obvious that it is a more 'complex' thing than one of its sub-atomic parts, at least to the extent that it is (if nothing else) an aggregation of many billions of such parts. In this rudimentary sense, a pile of bricks is also more complex that one brick by itself. 'Aggregational complexity', then, can be seen as a relative term denoting a basic physical difference between one thing and a larger *aggregation* (definable in some sense as a thing in itself) of those things. Although a fairly crude criterion on its own, even such a basic definition offers a basis for types of things identified in the world to be compared and related to each other. On these terms, a galaxy is more complex than a single star; a human body is more complex than one of its cells. In the task of articulating an ontological position around these ideas, aggregational complexity is a useful starting point. It offers two things.

Firstly, it provides one or two initial substantive criteria of *difference* between things – relative degrees of aggregation; a difference between parts and wholes. Secondly it links these differences together under a common theme; extending from the recognition of a basic (and original) materiality to the world, the various types of things we see about us (in the present) are portrayed as different *elaborations* of the same essential substrate[3]. Different things are thus also linked, by being viewed as *instances of composition* – a 'building up' of the basic physical elements of the

---

[3] The term 'substrate' here is used to generally describe the physical constituent elements that are commonly assumed to be shared amongst the larger structures and forms of the world.

universe into larger and more elaborate structures. (Bunge 1977, p. R82) We might say, then, that an idea of 'composition' may begin to be defined in these terms, as a 'coming together' (or aggregation). Clearly, some *process* (or processes) of composition is also implicit in the notion of temporal 'development' in the universe, from the 'simple' and universal to more complex and particular kinds of things.

On this basis we may move on to a range of more specific questions that science is concerned to ask about the nature of things that manifest different degrees of aggregational complexity, and the nature of the compositional mechanisms, processes and/or relationships (or what have you), that may account for those differences.

The next question here is, what are the further classificatory strategies to be applied from this position? And what do these differences of complexity, and instances of composition, then tell us about the general nature of the world, and things in the world; that is, what more fully developed ontological scheme do they propose?

### 2.2.3 Emergent Properties

The investigations carried out by scientists within particular fields are commonly understood to describe, specify or quantify the *properties* of their subjects. Chemistry will tell us that water has the properties of a fluid which freezes at such and such a temperature: biology will tell us that a dolphin has the properties of a mammal, in that it has mammary glands to feed its young, and so on. These, and many other identifiable, experimentally verifiable properties have conventionally been amongst science's main criteria of distinction and thus of classification.

In considering questions of complexity and composition, the properties that are understood to be displayed by notionally less complex 'parts' vis a vis those of a more complex 'whole' can serve to further refine and substantiate a classificatory distinction between them. For example, one might compare the properties of single $H_2O$ molecules with the properties of bodies of water. Here we can begin to see that to make such comparisons only in terms of differences in basic aggregational complexity, in other words simply as a difference between a single 'part' and an aggregation of those parts, will fail to account for the different properties displayed. Here it would seem that, in addition to aggregational complexity as a criteria of difference, sometimes the properties of the more complex (relative to those of their parts) suggest an 'ordering' occurring between those parts, as interacting *participants* within the compositional whole.

As Auyang describes it:

> The homogeneity and simplicity at the elementary level imply that the infinite diversity and complexity of things we see around us can only be the result of that makeup. *Composition is not merely congregation*; the constituents of a compound *interact* and the interaction generates complicated structures. ... Composition is as important to our understanding of the universe as the laws of elementary particles ...
>
> Auyang 1998, p. 2 emphasis added

Auyang argues, in other words, a view of 'mere congregation' (my aggregation) is too limited and that, further to an understanding of the physical foundations of the universe, equally essential is an understanding of composition, and the compositional relationships between things manifesting different levels of complexity.

To clarify this point we may return briefly to the analogy of the bricks. Previously I likened a difference of 'basic' complexity between the elements of the physical substrate and a more complex entity to the difference between a single brick and a pile of bricks. Let us now extend the analogy, and compare this case to the difference between a single brick as a house made of bricks. Intuitively it would seem clear that the house has additional properties of arrangement over and above the simple pile; that some sort of order pertains to the house that is not present in the pile. Is it possible (by assessing relative *properties*) to say in any more specific terms what it is, if anything, that is new and significant about this 'extra difference'?

In considering the pile as an instance of composition it is straightforward, I would suggest, to view the substantive properties of the more complex whole (the pile) as a simple extension of the properties of its constituent elements considered separately (the one brick). The weight of the pile, for example, is simply the weight of one brick multiplied $x$ number of times. In this way one can fairly easily conceive of the properties of the whole as a simple, aggregated result of the properties of a single part; the pile is pretty much just a lot of single bricks together in one place. As a process of composition *from* the separate bricks *to* the pile, one could readily conceive an answer in terms of simple addition; 'take a brick, add another brick, etc...' To put it another way, the difference between the brick and the pile of bricks would seem to be (mostly) a quantitative difference.

In the case of the house, however, such simple addition would not work. There is an overall order, the design of the house, to which the individual elements must conform. The properties of the house therefore are not so simply extrapolated from those of a single brick. It is this new element of relational 'order' pertaining between the parts that seems to be required for a proper account of the difference.

This analogy is not in all ways applicable and should not be taken as a definitive explication, but it does serve to suggest why in some cases differences in basic physical complexity might be further defined as different forms of 'ordering'. *Properties* that pertain to the more complex and not to their less complex constituent parts (considered separately) seem to require that we speak not only of instances of aggregational complexity but of *complex order*: because there would seem to be something more than the simple congregation of parts going on. The properties of the whole are not explicable as a simple, quantitative extrapolation of the properties of an individual part.

In this case, the concept of 'order' is used to refer to some kind of apparent arrangement that pertains between 'the parts together', and not to a single 'part' considered separately; such that the properties of 'the parts together' are in some way new (or qualitatively distinct) relative to properties of the single 'part'. In applying such a concept of further difference, we shall see later that the classification of things in general begins to resolve into a hierarchy of *levels* of complex order.

This matter is now widely understood in science under the general heading of *emergence* or *emergent properties*. These concepts have a broadly 'developmental' / temporal flavour in so far as we could say; if we generally take more complex things as having 'come out of' – or evolved from, or been assembled from – a less complex substrate of individual 'parts'; then to the extent that the properties of the more complex are new, and qualitatively different from those of the less complex, these properties are seen to come into being, or to *emerge*, along with the more complex form itself.

On this basis, we may distinguish between two categories of properties belonging to more aggregation ally complex things, as related to the properties of their constituent parts considered separately; properties that are either *resultant* or *emergent.*

Mahnen and Bunge offer a basic distinction between the two categories as follows:

> A property of a whole which is also possessed by some of its parts is said to be *resultant*. If, on the other hand, a property of a whole is not possessed by any of its components, it is called *emergent*. For example, the property of being alive is an emergent property of cells, but a resultant property of multi-cellular organisms.
>
> 1997, p. 29

Again, these definitions reflect the derivation of a concept of emergence from within the broadly 'bottom-up' world-view of science; that is of a world where complex, compound things are 'built up', through processes of composition, from a basic material substrate. Resultant properties generally are tied to the constituents as the *material content* of the whole. They represent composition as the quantitative, material aggregation of these constituents. So, as with the analogy of the bricks, the mass of a complex whole, say a long chain carbon molecule, may be understood to represent the cumulative mass of its atomic parts.

Emergent properties, by contrast, generally arise in the structural nature of systemic wholes; or, in other words, the complex of *relationships between* the constituent parts. Furthermore, they are understood to be of a *qualitatively different character* to the properties of the constituent parts considered separately. For this reason, emergent properties are also often described as 'novel', 'new' or 'unexpected' in relation to the constituent properties. In a sense, the identification of emergent properties may begin in the negative, as those properties of a systemic whole that fall outside, or seem to transcend, a purely constituent-based explanation.

To use a popular example, the property of *temperature* is one that pertains to water but not to single $H_2O$ molecules. A body of water clearly is understood to *have* a temperature, while a single molecule within it cannot be understood so. The property of temperature *x* in a body of water in therefore emergent relative to the properties of its constituent parts considered separately. Get enough $H_2O$ molecules together in one place and the property of temperature will *emerge* as the molecules combine to form a body of water (along with other emergent properties, such as fluidity). It would seem to result, therefore, not only from the separate molecules in and of themselves, but also from the 'order' of relations that pertains to them when together in large numbers. In this sense, water may be described as an instance of *complex order* relative to a $H_2O$ molecule, and the two may therefore be at least provisionally regarded as instances of two distinct levels of order.

The working definition of order in use here is therefore closely linked to the recognition of emergent properties. If the idea of aggregational complexity, on its own, directs us to consider a complex thing as a collection of constituent parts; then emergent properties, and the idea of order, suggest a necessity (for adequate consideration of that thing) to look not only at the parts as such but also, specifically and explicitly, at the compositional *relationships* between the parts, and the cumulative effects of those relationships under certain conditions. References here to some thing as an instance of 'complex order' (or 'ordered complexity') are intended to encapsulate the combination of (differences in) aggregational complexity, resultant and emergent properties as criteria of distinction. Where a more complex aggregation of constituent parts is understood to possess its own definite, coherent structure and to display emergent properties, it may also be referred to as a *system*.

The above quoted definition of emergence is accepted for the moment; the concept will be further refined later in relation to complex systems and Mind.
(Also: Bertalanffy 1971; Clark 2001)

### 2.2.4 Levels and Hierarchy

If these criteria are applied to the manifest diversity of our universe, and that diversity is regarded as the result of a 'developmental' outgrowth from a simple, physical substrate into more complex forms, then a more general classificatory schema may be constructed. Here things may be broadly categorised into conceptual *sets* or *levels* according to a relative degree of complex order they are understood to share. The sub-atomic physical substrate becomes, on these terms, the base level; the notional set of all sub-atomic particles. 'Up' from there, other levels commonly distinguished include: the atomic level (the set of all atoms as things in themselves and as composed of sub-atomic parts); the molecular level (the set of all molecules as things in themselves and as composed of atomic parts; the cellular level (the set of all cells as things in themselves and as composed of molecular parts); and so on. (Bunge 1977, p. R78)

The first thing to note about this, as Bunge does (1977, p. R78), is that the level as such is to be understood as a *class*, and thus as conceptual. The material entities that are understood to belong to a class, however, are real, not conceptual. It is real properties and characteristics embodied in these entities, when understood as shared between like entities, which provide the main criteria by which such classes may be defined. (This is an important distinction to be kept in mind in discussing emergence, and other related matters. For example, one might employ a shorthand description of properties being emergent at one level vis a vis another 'lower' level.

Strictly speaking this is incorrect; emergence does not occur between levels as such at all, but is a phenomenon occurring in material entities having particular properties and concrete, compositional relationships. Provided this is made explicit from time to time, however, some shorthand description of 'relationships (or emergence) between levels' can be convenient and useful.)

How then do we move from the definition of such classes – of atoms, molecules, cells, etc – to a linking of them together as levels within a hierarchy? The concept of the hierarchy in this context has two related aspects; what I will call a temporal aspect and a compositional aspect.

The temporal aspect arises specifically in relation to the broadly 'developmental' character of the universe – the universe in so far as we understand it having begun, in its earliest life, with the most essential, sub-atomic materiality, and then having changed through time, to display other more complex and ordered structures. In then applying the classificatory terms mentioned above to this history, as a form of 'mapping' some of the general characteristics of change, the transitions from more simple/essential to more complex/ordered/structured appear also as a progressive realisation of levels; each, as it were sub-structing the next. Thus the universe in its temporal aspects may be understood to have embodied a developing hierarchy of levels; from the sub-atomic to the atomic, to the molecular, the cellular, and so on. In all cases, of course, the generalised development is embodied in the concrete entities understood to populate each level, and the specific relationships between them. The notion of hierarchy, then, is a description of this progressive realisation.

The compositional aspect of the concept of hierarchy may be understood by considering an entity at the more complex end of the spectrum; an individual human being, for example. From the perspective of contemporary science, and a view of the world 'built up' via processes of composition from a basic materiality, it follows that the complexity of person 'x' remains a material phenomenon. It remains, in other words, that whatever our unique characteristics as a species may be, each of us is made of the same basic materials as any other physical thing. We assume that the sub-atomic particles making up person 'x' at any particular moment, notwithstanding their 'participation' in his being, remain subject to the same laws of physics as all sub-atomic particles.

Furthermore, it is commonly accepted that the sub-atomic particles making up person 'x' at moment 'y' are also structured into atoms, and that (many of) these atoms are further structured into complex molecules, and so on.

In other words, person 'x', as an instance of complex order and of material composition, may be generally seen to embody (or contain) parts that, as distinct structures in themselves, also represent many of the classes defined above. They may be seen, again, as embodying levels in a hierarchy in so far as they are understood to represent degrees of increasing order and structure, simultaneously present within the one articulated, complex whole. (It bears noting also at this point that the description of something as a 'whole' or 'part' within such a hierarchy of complex order will always be a relative description. A human cell may be appropriately considered as a unitary whole, a thing in itself, relative to its own constituent parts; it may equally be regarded as a part of a 'greater' whole; the person in question.)

The concept of hierarchy as employed here, then brings together these two aspects in a similar way to that described by Bunge:

> The scala naturae or "hierarchy of life" is no longer static: it has become part and parcel of an ontology that is not only pluralistic but also evolutionary. The levels are not static layers that happen to be piled atop one another. They succeed each other in time (metaphorically speaking) and they do so by virtue of a definite and pervasive mechanism, namely the self-assembly of things. And they are not rungs in a hierarchy leading from lowly atom through middling man to the Supreme Being: they are but stages in a natural evolutionary process that may have occurred and indeed may be occurring at various places and epochs in the history of the world, though not twice in the same manner.
>
> 1977, p R79

So, it is both the temporal and compositional aspects described – together with the criteria of aggregational complexity and emergence – that here gives meaning to the definition of levels of complex order within a hierarchy. However, as Bunge notes, it is composition (his 'self-assembly') that links together the concrete entities of each level, and thus makes sense of the many levels as parts of one hierarchy. The designation of levels as 'higher' and 'lower' that the notion of hierarchy inevitably introduces, are again generalisations. What they are describing is the cumulative dimension of composition; i.e. that the entities of the higher level are seen (temporally) to *arise out of* the entities of the lower; and continue (compositionally) to embody the entities of the lower level as parts within a new, systemic whole.

In any case, the description of a level as 'higher' or 'lower' than another is not intended to convey any connotations of relative ethical value or significance. (The argument might be taken in that direction, but that is not my present purpose.)

### 2.2.5 The Task of Ontology

The task of ontology, we might say, is to make a claim about the essential and substantial nature of *being*, about the 'stuff' of which the world is made. What, then, can be said about the ontological questions and issues that arise from within a contemporary, secular and scientific worldview, insofar as it addresses the matters under discussion?

It would seem, as has already been suggested, that the physical nature of the universe – insofar as we understand and describe it in the language of physics – is primary. There is a widespread consensus in science that the essential 'beingness' of the things around us lies, if nowhere else, in their shared, physical *materiality*, in that they have in common basic, physical constituents. Such a supposition, of course, is the very thing that casts more complex combinations of material 'stuff' as 'instances of composition' in the first place.

There is little doubt, therefore, from the point of view of science, that the entities described by physics are properly regarded as independently real things. Given the commitment here to a naturalistic and science-based view of the world this, as a grounding ontological claim, is not regarded as controversial.

Furthermore, there is also general recognition of the reality of specific 'things', as larger aggregations of this basic materiality. For example, there is a widespread recognition that planets are real features of the universe; that they are made up of physical material; and that as such they are subject to physical laws and so on. The same would be said of other distinct and definable things or systems; a dog, a galaxy, a human being, at least so far as their substantial materiality is concerned.

Some of the more contentious ontological questions, then, arise beyond this general consensus, and come into play around the matter of 'levels'; how they are defined, what their ontological status is in relation to the underlying materiality, and how (as instances of composition) they relate to each other. In other words, for the ontology of contemporary science, it is often questions and debates about the nature and significance of 'levels', emergent properties and the like, which form the background against which particular questions are raised and positions are adopted. (And as shall be discussed shortly, it is also in the relations between *theories* developed and applied in relation to notionally different levels that many contemporary epistemological debates have their source.)

In the broadest terms, and within the bounds of a generally materialist naturalism, one might, for example, counterpose two basic options. One the one hand, we may recognise that some real properties of real things are genuinely emergent (that is, qualitatively distinct from the properties of their individual, constituent parts); and that, if we take the properties of something as very much part of what defines the specific, concrete nature of that thing – a part of its beingness – this would seem to demand recognition of multiple kinds of things in the world, albeit sharing the same kind of underlying physical substance. Thus we would tend towards a generally materialist and what I would call a *pluralist* ontology. On the other hand, one might claim that any apparent plurality of properties and things is epiphenomenal, and in fact arises only and directly from the nature, properties and laws of one kind of elementary materiality, the purely physical substrate of the universe. Thus we would assert a strongly and narrowly *monist* ontology; that *really* there is only one kind (or level) of real thing(s).

Narrowing the focus from such encompassing metaphysical positions, one may recognise also a range of ontological questions and claims that arise in the investigation of specific complex entities, their properties as notional wholes vis a vis the properties of their constituent parts, the concrete causal or compositional relations that pertain between these two notional 'levels' and so on. Again, we may note, however, that, whatever the specific conclusions drawn, it remains a conceptual framework of 'levels', emergence, complexity, etc (or something akin to it) within which these questions are generally asked, and ontological claims are made. It remains, in other words, that for the contemporary science, these would seem to be some of the most basic terms around which a variety of ontological arguments are constructed, and that the position one adopts in relation to the compositional or causal relations across 'levels' (or, to be more precise, between the material entities which populate and define these levels), will largely define one's (explicit or tacit) ontological commitments.

The intention in this work is to recognise emergent properties in some complex structures, and so to adopt a broadly pluralist ontology as described above, as a basis for the work of later chapters. In that later discussion it will also be useful to have at least an outline of some other related issues here, under the general heading of *reductionism*.

## 2.3 REDUCTIONISM

The terms *reduction* or *reductionism* are ones widely used in relation to aspects of both contemporary and historical science practice and theory. The particular meanings attached to the terms often vary. Both ontological and epistemological theories or claims have been labeled 'reductionist'. The term might be applied to an interpretation of specific experimental results, or to the positing of a supposed 'theory of everything'. It will be useful, therefore, to survey some of the different uses and how the term is to be employed in this context.

To begin with, the general notion of 'reduction' in this context has been used to describe circumstances where some particular science-based interpretation of the world becomes subsumed within another (seen to be) more encompassing, more adequate understanding. In other words the former is seen to be generically 'reduced' to the latter. Historically, of course, it is part of the record of modern science that particular theories have been overtaken or radically revised as knowledge has changed and advanced. An example often cited is where the separate studies and theories of electricity and magnetism gradually converged and, early in the twentieth century, were combined under a unified theory of electro-magnetism (Koestler 1975, p.230). On we might consider the case of 'Newtonian mechanics [being] superseded by relativistic mechanics, the domain of which covers Newtonian mechanics and more.' (Auyang 1998, p. 52) Common sense would suggest that reduction in this sense, as a reinvention or consolidation of scientific knowledge, is an entirely appropriate matter.

Another related usage of the idea of reduction is in the characterisation of scientific enquiry as often having a broadly 'reductive' orientation; which is to say that it is common practice in science to investigate the properties and causes of more complex structures and systems by reference 'down the chain' of complexity, as it were; that is, in relation to their constituents: and that science has often achieved tangible successes in so doing. (E.g. Auyang 1988; Bunge 1977; Popper, 1974 p.260)

The more specific term 'reductionism' however, is generally used to describe stronger, or more doctrinaire approaches to the relationships between different domains of scientific enquiry, and/or between the apparently different levels of ordering within complex systems. It can refer to claims in either an epistemological or ontological vein. In the epistemological terms this would generally amount to a view that theories/concepts specifically identified with a notionally higher level can be effectively *eliminated*, by being entirely translated into the terms of a lower level theory.

For example, let us say (in rough terms) that we have a body of concepts and theory describing the properties and behaviour of single cell organisms. We also have a body of biochemical theory describing the properties of the organic molecules that constitute these organisms. The general intention of a reductionist approach would be to translate the concepts and theories of the first into those of the second (perhaps modifying and expanding the constituent theory in the process; and with appropriate 'bridging rules' for the translation), to the point where the 'higher level' theory becomes *dispensable.* As philosopher of science, William Bechtel, describes it, the aim would thus be to arrive at a situation where 'laws about the behaviour of cells (now stated in the language of the lower level and so using terms referring to molecular structures) are to be derived from laws about the behaviour of molecules (e.g., the laws of biochemistry).' (1988, p. 73)

In ontological terms, reductionist claims would generally be suggesting that properties (and particularly emergent properties) or characteristics seen to be specifically generated at a (supposedly) higher level of complex order are *nothing but* epiphenomena arising from the concrete properties and lawful relations of entities and processes at a lower level. One may note that such contemporary claims as are made in this vein more commonly arise as derivative interpretations of more substantial debates and ideas about reductionist epistemology and methodology; in so far as 'the elimination of the reduced concepts suggests that the things represented by these concepts are *nothing but* the things represented by the reducing concepts'. (Auyang 1998, p. 51) For example, from the sociobiologist, E.O. Wilson:

> The central idea of the consilience worldview is that all tangible phenomena, from the birth of stars to the workings of social institutions, are based on material processes that are ultimately reducible, however long and tortuous the sequences, to the laws of physics.
>
> 1998 p.266

Now, it is not within the scope of this study to argue in any substantive way the relative merits of reductionist or non-reductionist approaches. We may note however, that the main thrust of contemporary argument against reductionist epistemology is simply that, on the whole, it has *not* demonstrably succeeded in eliminating concepts and theory which describe complex material systems as such (particularly with reference to their emergent properties) in favour of constituent-based concepts and theory. (Auyang 1998, pp. 51-60; Fodor 1974)

What will be useful for later purposes, however, is to note that these (roughly) pro and anti-reductionist debates within science present a number of points of contemporary tension: as to what the world is *really* like; how science explains the world and the human condition; what can be achieved by way of explanation and (perhaps) what can't; and, by extension, what wider social or political interpretations arise out of science. In the areas of scientific enquiry of most interest here it is generally the case, I would suggest, that these tensions arise between two broad 'schools'.

On the one hand there are approaches that generally acknowledge the idea of emergence and the particular challenges it presents in terms of how highly complex structures and systems are studied, and the terms on which explanations about them are generated. Particularly with the profound material and behavioural complexity of biological systems, it is recognised that such systems are appropriately considered through *multiple levels of description,* and that levels will often require some descriptive and explanatory concepts and principles tailored to them, so to speak. (Clark 2001, pp. 160-161; Montalenti, 1974 p. 17) It would seem therefore, on the whole, that such an approach may indeed be broadly considered as scientifically pluralist, allowing for different modes of explanation to suit different domains of enquiry. On the other hand, the general tendency of reductionist approaches is to seek to 'explain away' apparently higher level system properties by reference to the properties of their constituent elements at a lower level (Rosen 2000, pp. 85-86).

Related 'tensions' are also apparent in other areas such as the identification of causes in explaining complex system behaviours. On the whole, I would suggest, it is often seen to be appropriate in a more pluralist approach to highly complex systems to recognise multiple causal factors. (E.g. Clark 1997, p. 107; Auyang 1998, p.65-66; Rosen 2000, p. 280-281). Some of the reasoning behind such an approach will be explored further in later discussion. Conversely, the general reductionist tendency is to narrow explanation down to single casual factors.

For the purposes of this study it is useful to be aware of these tensions, and the general epistemological character of the different approaches. These issues will be returned to in Chapters 7 and 8 where some of the social interpretations arising from science will be the subject of discussion. It is recognised, of course, that the description here would probably not do justice to the many of the specific debates that arise, or to the detailed positions taken in relation to them.

In the next sections, it is the generally pluralist 'pathway' that continues to be followed – including a recognition of both multiple levels of description and multiple causal factors – in considering some of what this approach has to offer for an understanding of complex biological systems.

## 2.4 Conclusion

The ideas and issues surveyed in this chapter are relevant to the ways that contemporary science, and the philosophy of science, approaches questions about the nature of human Mind and cognition. As shall be discussed in Chapter 4, it is reasonably typical to consider the human organism and the brain as physical structures, as instances of highly complex composition, and as embodying multiple levels of ordering; all in keeping with an underpinning materialism. From that point of departure, as it were, some of the central questions are then concerned with *how* various notional properties or capabilities of human mentality arise from these highly complex physical structures; and allied to that, what investigatory or explanatory strategies and concepts may best serve our understanding of this complexity. Here the concept of emergence will continue to be important.

By initially locating issues of Mind within these wider perspectives we are also led to view it as something which has itself arisen out of compositional processes occurring over the timescale of billions of years.

 

# CHAPTER 3: THE CHARACTERISTICS OF BIOLOGICAL SYSTEMS & AN INITIAL MODEL OF CREATIVE ORGANISATION

## 3.1 SOME GENERAL CHARACTERISTICS OF BIOLOGICAL SYSTEMS

In the following chapters the study will move to the subject of Mind as an emergent property of biological systems. As a further preparatory step for that discussion, this chapter will employ ideas and perspectives developed so far to briefly consider some of the characteristics of biological systems. The focus of examples will generally be on some of the most simple and (it is commonly assumed) the original forms of life on Earth.

Through this focus on the most basic of living organisms, the aim is to pick out certain key characteristics seen to be common to biological systems in general. Concepts capturing these several characteristics will then be carried forward into Chapters 5-7, and form essential ingredients for the view of Mind-as-creative-organisation to be proposed.

### 3.1.1 The Single-Celled Organism

It is a matter of some agreement in biology that the earliest forms of life on earth were probably basic single-cell organisms resembling modern bacteria, thought to have appeared roughly 3.5 billion years ago. Subsequently, so it is thought from fossil records, it took a further 3 billion years for the first large multi-celled organisms to appear. As some have argued, it was within that long interval that the prevalent single-cell organisms evolved many of the adaptations that enabled the subsequent outgrowth of greater size and complexity of form. (E.g. Margulis 1987, pp. 98-100)

As the earliest form of life, these simple organisms by their presence pose questions about their own origins, and about the special combination of characteristics that distinguishes them from the non-living 'primordial soup' of chemical elements from which they are thought somehow to have sprung. Although it is not necessarily easy to point to one particular property that marks our the living from the non-living, it is commonly recognised that a number of

essential characteristics combine in simple life forms that do serve to mark them out from other bio-chemical structures. As is considered below, some of these include the presence of a boundary structure (the outer cell membrane), metabolic processes, and the capacity for reproduction. Having this in mind there are numbers of thinkers who would say that this combination of characteristics in simple organisms is to be understood as emergent – in a way consistent with previous definition and discussion – relative to the properties and characteristics of their constituent molecular parts.
(E.g. Rosen, 2000; Auyang, 1998; Maturana & Varela, 1987; Ayala & Dobzhansky, 1974).

Let us consider, then, a basic description of a single-celled organism as a particular kind of systemic whole. The first part of the description points to the existence of a membrane, the outer surface of which is generally regarded as the defining boundary of the organism in itself. The second characteristic identified is the operation of a dynamic process internal to the organism, which serves to sustain and reproduce the components of the organism – the cell metabolism. Part of the workings of the metabolism is an exchange of material with the world outside of the membranous boundary – the intake of material to feed the system and the expulsion of 'waste' material.

As biologists, Maturana and Varela, argue, the highly inter-dependent combination of these two characteristics may serve to capture something of the emergent character of the single-celled organism, relative to the physicochemical world from which it has sprung.

> What we have, then, is a unique situation as regards relations of chemical transformations: on the one hand, we see a network of dynamic transformations that produces its own components and is essential for a boundary; on the other hand, we see a boundary that is essential for the operation of the network of transformation which produced it as a unity
>
> 1987, p. 46

Maturana and Varela define this form of organic organisation as *autopoietic*; that is, as 'continually self-reproducing' (p.43). Although the term will not be extensively relied upon here, the basic description of organism that it summarises is endorsed.

### 3.1.2 Stability & Dynamism

For the purposes of this study, one of the most significant aspects of this initial description lies in the notion of dynamism and change as essential to the basic character of life. (Also: Weiss 1962, p.3) It is clear that the existence of a single-cell organism relies upon the continuation of its dynamic, metabolic processes of self-reproduction. One might say that the organism, in this sense, is in an on-going state of *flux*, as its material content is continuously renewed and remade, and is continuously exchanged with the wider environment. If this process ceases, then the disintegration of the organism will soon follow.

Alongside this dynamic, changeable aspect, however, is what would seem an equally important *stability* in the overall order and integrity of the organism. This we might say is the consistent pattern of relationships pertaining between the many material elements of the organism; the pattern within which these elements function as parts of a coherent whole. The persistence of this ordering through time is again critical to its continued existence. As Maturana and Varela state the matter:

> Thus, autopoietic unities specify biological phenomenology [ie require a specifically 'biological' understanding] … because the phenomena they generate in functioning as autopoietic unities *depends on their organisation* and the way this organisation comes about, and not on the physical nature of their components (which only determine their space of existence).
>
> 1987, p. 51, emphasis added

In other words, they suggest, the emergent, stable character of the biological system must be seen as *vested in its overall organisation*, not simply the collective presence of its physical constituents per se.

Recognition of this systemic, relational order in a simple organism, *as a characteristic in itself*, is consistent with the recognition of emergence properties in biological systems (relative to the properties of their physicochemical constituent elements). (Rosen, 1991 p. 117) It is also consistent with the recognition of the specificity of some of the explanatory concepts and strategies required by the fact of emergence. Out of this recognition, a particular term is introduced which will be important throughout the remainder of the study.

### 3.1.3 Organisation

Sometimes, in literature relevant to this study, the term 'organisation' is used in a way roughly equivalent to the way I have hitherto used 'order' or 'complex order'. For particular purposes here, however, it will be used to refer specifically to the type of complex order displayed by living organisms. Therefore it is proposed: that an organism, as an autopoietic entity, insofar as it displays this dual aspect of both an emergent, stable, systemic order, and dynamic processes of flux and change in its constituent elements; may be considered as having an overall character of *organisation*. In other words, the idea of organisation is here intended to capture the way a living organism displays emergent, complex, stable order, but that it also does so in a way dependant upon a continual process of dynamic change; and that through a combination of these two aspects the organism functions as an active, self-reproducing, bounded system within an environment.

So, whereas 'complex order' might be seen to be manifest in the emergent properties of a body of water vis a vis its constituent molecular parts; water does not display the same kind of 'dynamic ordering' characteristic of an organism. Assuming stable conditions, a body of water will remain stable also, and maintain its properties, without, one might say, needing to *actively do* anything. It is seen to be appropriate and useful here, therefore, to differentiate and summarise the particularly dynamic and active ordering of the organism with a specific term, and 'organisation' is the term chosen for that purpose.

The term should be understood to describe what is in reality – in each specific instance – one unified, living *process in motion*, as it were, from which the two inter-dependent aspects of stable, systemic order and dynamic changeability are differentiated.

### 3.1.4 Organism-in-Environment

The specific characteristics of the organism-system encapsulated in the notion of organisation, beyond the nature of the organism in and of itself, require also a particular perspective on *environment*.

As suggested earlier, the basic definition of a single cell organism may begin with the identification of a membranous boundary, which is also part of the cell's dynamic self-maintenance processes. The identification of a boundary simultaneously specifies an 'inside' *and* an 'outside', and across that boundary an active exchange occurs. The actions of metabolism, taking in and assimilating material from the environment and expelling material into it, is one part of this exchange. Already, therefore, a characteristic process of the organism actively incorporates the environment, and could not be adequately understood as a process in isolation from that wider context.

Also of relevance to the wider issues of this study, is the capacity of even the simplest organisms to respond to certain environmental conditions. For example, a particular variety of the single-cell organisms called myxomycetes, *Physarum*, will in humid conditions, 'grow a flagellum and become motile. If the environment is dry, the cell develops pseudopods and becomes an amoeba.' (Maturana and Varela, p. 77) Without dealing with the detail of how and why these capacities are enacted, what such examples indicate is that at least some of the specific behavioural properties of the organism must also be understood *in context*; one might say, within the wider set of relations that is organism-in-environment. The difference between myxomycetes' development is either of two directions is not fully understandable without recognition of the active engagement of the organism with its environment.

In terms of explanatory strategies applied to such cases, clearly the environment external to the system itself must in some sense be taken into account as a domain of potential causal influences on system behaviours or system-internal states and events; alongside factors identified within the system itself. To follow the same example, in the case of *Physarum*, if we are concerned to ask why the 'flagellum growing procedure' was initiated, then it would seem that there will be a number of factors to potentially be considered including: that the environmental conditions were humid; that a particular precursor bio-chemical change occurred within the cell; that the organism was in such-and-such a point in its lifecycle, and so on. A range of such factors may be implicated in certain system events occurring, and this is a useful example of the way in which the consideration of 'multiple causal factors' comes into play when living systems are concerned.

As a general principle, also, it is proposed that at least some of the emergent properties of organisms are appropriately understood as *embedded*; that is, as realised within a wider set of relations of organism-in-environment. 'Environment' becomes more than an easily separable set of 'external conditions or surroundings' (Collins English Dictionary 1993); it is something 'implicated by association' with defining characteristics of the organism itself. This is a point of understanding that will also important to later discussion of Mind, where the notion of embeddedness will be further extended.

Having identified the significance of this point, we can now clarify the idea of an organism's capacity for 'responsiveness' within an environment that examples such as that above clearly suggest.

### 3.1.5 Responsiveness

Although 'responsiveness' is chosen as a more neutral term, it nevertheless raises the spectre of teleology; in this case of construing from examples such as that above that a simple organism is somehow capable of *purposefully* directing its own activity towards a *future* goal beneficial to itself. Here the notion of organisms as responsive to their environment is one more significant 'stepping stone' for the discussion to follow; but is used in a way that avoids teleological inferences. Here, saying that an organism is responsive is only to say that there is a congruence between (some) systemic capacities (internal to the organism) and (some) environmental conditions. Certain environmental conditions will trigger changes in the organism's structure that in general are adaptively effective for the organism-in-environment. Thus, in total effect, the organism has a capacity to respond to certain changes in environmental conditions as they occur. In general, as the complexity of organisms increases, so too does the range and variety of responses of which it is capable. (Maturana and Varela, 1987 pp. 147-150)

One should note that this is another distinct aspect of the quality of dynamism in the organism. Whereas the dynamism of metabolism involves the 'exchange' of material with the environment, responsiveness is better understood as environmental conditions triggering the dynamical operation of more exclusively *internal* sub-systems, which is subsequently effective for the organism-in-environment.

The business of describing and explaining the development and workings of these responsive capacities we may also note as another place where multiple causal factors may be seen to shape and influence system behaviours; possibly factors identified in the environment, in system-internal, constituent states and events, and/or in the disposition of the system as a whole. (Rosen, 2000 p. 281)

### 3.1.6 Further Emergence – From Cellular to Multi-Cellular

Some varieties of myxomycete cells show another interesting responsive capacity, and that is to join with other cells to form a larger, multi-cellular organism. With *Dycostelium*, for example, individual cells will join together to form a single, fructiferous body as part of an overall lifecycle. The originally single type of cells will, within that body, adapt their form in order to meet specific functions; to make spore, to support the stem wall, and so on. The important point of interest here is that the multi-cellular organism displays emergent properties in its own right, properties of different order to those of its individual and formerly separate cellular components.

The larger ordered structure within which the cells diversify in function is the signature of a new level of overall *organisation*.

In terms of the new capacities for responsiveness to environment shown by the multi-cellular *Dycostelium*, (such as a capacity to form a gripping 'foot' onto a surface) it is most useful to note that these can only be regarded as emergent also; and as likely to require some explanatory concepts and strategies beyond those appropriate to its cellular components. These concepts and strategies are again also most likely to address themselves in some sense to the wider relationship between the organic system as such and its environment. (Nicolis and Prigogine 1989, pp. 31-36)

In wider terms, this is one example of a general capacity in living systems to develop within individual ontogenies, or over generations to evolve, through 'stages' of organisation, such that new emergent properties arise. These instances will generally fit into the terms of classification in levels and hierarchies of levels as per previous discussion. There is a demonstrable capacity, one might say, for evolved organisms to manifest *multiple levels of organisation*. In the business of describing and explaining these as hierarchical, it will sometimes be appropriate to recognise that an organic system which has functioned (or been understood) as an individual, emergent whole relative to its own parts – and embodied the higher level relative to the lower – can become (or be understood as, from a different perspective) a constituent part in a larger emergent whole – and thus embody the lower level relative to the higher. It becomes in turn, we might say, the substantial content of a higher order organisation. It is in this general way that the multiple levels of organisation may be recognised within larger and more complex forms of life (like ourselves).

## 3.2 Key Concepts for a Model of Creative Organisation

The goal of the above discussion has been to briefly outline a number of general characteristics of biological systems. Having a conceptual 'tool-kit' that captures these characteristics will be important for the next chapter, and the main discussion of the thesis. Below a number of key concepts are further spelt out, and comments made about these characteristics combining as *creative organisation* in biological systems.

- **Organisation**: firstly, I have suggested that the organism-system embodies – as dual aspects of one, dynamic process – a combination of two general capacities: one for stable, systemic ordering, and the other for variability or change in the way that order is instantiated at any particular time. This particularly dynamic, changeable-yet-coherent form of systemic order is described as the system's organisation.

- **Stability**: one side of the dual aspect captured in the idea of organisation is the systems capacity for *stability* in the sense of maintaining a set of stable system features through time.

- **Variability**: within that overall stability, however, the constituent elements of the systemic process, on a moment-by-moment basis, are liable to change under the influence of a range of causal factors, including environmental factors.

- **Emergence**: firstly, the systemic organisation displayed by a single-cell organism displays certain emergent properties or functional characteristics relative to its constituent parts considered separately. Secondly, it is also apparent that organisms can embody emergent organisation on multiple levels; for example, *Dycostelium* displays within its lifecycle a process whereby initially multiple and separate instances of cellular organisation become the parts of larger organised whole, one with systemic properties of its own that emerge via the combinatorial process.

- **Constraints**: in all of these dynamic processes of system functioning or development, the idea of constraints is also an important one. It is to say that these dynamic processes occur within a range of stable and/or limiting factors, potentially either environmental or system-internal, that shape and 'constrain' the outcome.

- **Multiple causal factors**: finally, we may note the assertion that these various dynamic, changing processes are subject to multiple causal influences.

It is in the combination of these several characteristics that the notion of creativity enters the picture. Specifically, creativity is seen to occur in biological systems when organisation at one level moves through and completes a process whereby emergent organisation at a new, 'higher' level is generated; the example of *Dycostelium* being a case in point. Within the view taken here, this generative process is creative in that it occurs via a concatenation of multiple factors including: stable organisation at the constituent level; influences between constituent elements;

and wider environmental influences. Due to the complexities of the biological systems involved, and the influence of these multiple factors entering into the process, (which engender a certain degree of randomness) each instance of emergent organisation – let us say, each time *Dycostelium* cells engage in the generative process whereby the larger, emergent structure is produced – will be somewhat different to each other instance. The combination of constituent elements will be 'put together' is a slightly different way each time. Thus we have a process where the characteristic stable ordering and the variability displayed at one level also enter into this 'production' of new, emergent organisation, and are embodied in what is produced.

The description of 'creativity' is thus seen to be warranted, to the extent that the combination of factors entering into an instance of emergent organisation has this 'ad hoc' nature: the precise combination of relevant factors that pertain in each instance will be somewhat different. In many of the examples to be discussed here, particularly those relating to human morphogenesis, the result is then an emergent structure of a particular general type (and the processes are reliable and stable enough to achieve this consistently), but also a unique instance of that type.

Complex biological systems that develop multiple levels of organisation are thus regarded as commonly (but not always) displaying a general capacity for *creative organisation*. This claim, however, must be qualified by saying that the degrees of difference between instances of relatively simple systems, like *Dycostelium*, coming about through the ad hoc combination of factors within their emergent development, does not necessarily signify anything of great import for our understanding of the type. The degrees of resulting difference between one *Dycostelium* fructiferous body and another may produce, as it were, only marginal differences between instances, in terms of the properties of the higher-level structure itself.

Thus the 'creative' part of the description of these processes as creative organisation is put, at this stage, only as 'small-c' creative. However, the significance of the characteristics identified is that, in the view taken here, they remain consistently present from the relatively simple to more complex biological processes such as human morphogenesis, and that it is there, at these higher levels, that they then begin to combine in processes that are creative is a more meaningful and important sense.

"MIND AS CREATIVE ORGANISATION"

# CHAPTER 4: COGNITIVE SCIENCE & MIND

## 4.1 INTRODUCTION

Over the course of the next three chapters I will extend upon the initial perspectives on biological systems outlined in the previous chapter; to consider aspects of our human biology, and of Mind as a specific feature of this active, living system[4]. The broad aims here are two: the first is to suggest that processes of human morphogenetic development display the same characteristics of biological organisation as previously mentioned. The second aim is to argue that Mind is an emergent feature of these processes, and then to propose a view of Mind also as a process of *creative organisation* – occurring, as it were, in the medium of an environmentally embedded, highly complex, biological system. One might say that the proposed key characteristics of Mind-as-creative-organisation are seen to be already present in the evolved structures and processes that enable the flowering of Mind to occur. An understanding of this continuity of character is therefore seen as useful and appropriate, even perhaps necessary, to the better understanding of Mind itself.

Broadly speaking, then, the case presented here may also contribute to a wider view about Mind that would seem to be becoming more prevalent – in cognitive science circles and elsewhere – that knowledge of our structure and particularly our *development* as biological organisms may be far more than tangential to understanding ourselves as Mind-possessing, intelligent, social beings. (E.g. Clark 1993) Although it is not my intention to argue the point directly or in detail, some aspects of this view, and of those contrary to it, will be raised as the discussion proceeds.

[4] Although this view would accept the possibility of Mind-like characteristics in other forms of animal life; in keeping with the wider purposes of this study, the focus here will be exclusively on human Mind. And, of course, when it comes to the nature of Mind, it would seem on the face of it that human experience and behaviour – from our own perspective, anyway – are the domains where we can most confidently identify mental properties and abilities in action (Auyang, 2000c p. 93; Edelman 1992 p. 115).

The ordering of ideas in what follows will also be shaped by the on-going goal of considering Mind, and the ground out of which Mind arises, at various timescales. Chapter 2 looked at some of the features of complex order in the universe as it has developed over a universal timescale. Chapter 3 considered some features of simple biological systems and, although evolutionary processes are not discussed directly, the idea of a continuity of character from simple systems to the human organism locates the coming into being of human Mind within an evolutionary timescale.

From this point forward, however, the distinctions become somewhat more pointed insofar as we are now looking specifically at the domain of the human, and thus must recognise that the nature of Mind can be, and is, considered within this domain (by science) from various, timescale-related perspectives. Or, to put it another way, if we are to ask (roughly), 'How is Mind *realised*?' then it becomes clear that the question might take us in one of several more specific directions. It might lead to a focus on human morphogenesis and the 'realisation of Mind' in the sense of properties of Mind coming about via the physiological development of the organism as a whole and/or the brain in particular. In this case, perhaps, the most relevant span of time will be that period of ante and post-natal development where the main phases of morphogenesis and neurogenesis are occurring. Equally, however, one could focus on the 'realisation of Mind' in a much more immediate sense – perhaps in terms of what the brain is doing to realise cognitive abilities on a moment-by-moment basis, and where it would seem the relevant timescales are much shorter.

The intention here is to recognise the sense of such distinctions in how the discussion is structured, and to continue to 'track' the realisation of Mind at progressively shorter timescales. Thus (following a discussion of background cognitive science issues in this chapter), Chapter 5 will focus on aspects of morphogenesis, and in particular neurogenesis – examining the nature of the processes by which body and brain are 'constructed'. In Chapter 6, the discussion will consider areas of the developmental trajectory where morphogenesis overlaps with the (mostly post-natal) 'take-off' of cognition – where the developing structure of the brain and a range of other influences begin to 'come together' in the establishment and growth of cognitive processes. In Chapter 7, I will then discuss the nature of Mind as something realised on a moment-by moment basis; through a nexus of brain, body and environment.

The logic of the approach remains throughout to: a) draw on the wider science-based case for seeing Mind as very much bound up in an environmentally embedded biology, and b) to thus make a case for seeing Mind as a particular emergent feature of a system which, from its genesis in conception onward, maintains a consistent character of creative organisation. The field of *cognitive science* will be an area of specific focus; mostly via cognitive science-based views of Mind as *computation* and in particular, the so-called *connectionist* view. The latter model will be further pursued as (in some limited ways) consistent with a biological perspective, and supportive of the notion of creative organisation.

We may note also that some of the issues in this area will be significant in terms of the wider goals of the study; to develop a view of Mind with one eye towards its wider social and political interpretation. For example, some of the contrasting views about the developmental character of Mind, also – not surprisingly – inform views of *human behaviour*, and 'behaviour' we might say is one of the ways we describe the domain of interaction between the individual, Mind-possessing organism and the social realm. Different ways of understanding the underlying causes of behaviour are thus very likely to inform some quite different social interpretations. The discussion of Mind here will also inevitably reflect somewhat on the nature and origins of behaviour, and some points about the significance of these reflections for social interpretation will be carried forward into Chapter 8.

In summary, the next three chapters will put forward a case for Mind as creative organisation via a number of steps, with the following main themes:

1. Background issues and concepts of contemporary cognitive science
2. Aspects of human morphogenesis and neurogenesis
3. The connectionist view of representation in computational systems
4. The realisation of representational capacities in the developing human Mind, in the light of connectionist ideas
5. The nature of moment-by-moment cognition

## 4.2 Cognitive Science: some Background Issues

So far as philosophy is concerned, then, the approach taken to the nature of Mind in this chapter relates most directly to contemporary philosophy of cognitive science, and then, to a lesser extent, to some longer-standing issues in the philosophy of Mind. For these reasons it will be useful to preface the chapter with a brief overview of the field of cognitive science and philosophy of Mind.

Firstly we may note that cognitive science – the scientific study of Mind – is commonly recognised as a relatively young domain, and as one that bridges across a number of more specific disciplines. Philosopher, Jerry Fodor, for example, describes a field covering 'artificial intelligence, computational theory, linguistics, cybernetics and psychology.' (1981, p. 124) Varela, Thompson and Rosch include 'neuroscience … sometimes anthropology, and the philosophy of mind.' (1996, p. 4-5) As shall be relevant later, one could also include the application of dynamic systems theory to questions of Mind and intelligence. How then may we begin to 'get a grip' on some of the central questions that link these various areas of enquiry together?

### 4.2.1 Some Basic Issues & Questions

A number of thinkers (e.g. Fodor 1981; Clark 1993, Ch 1) locate the precursors of a specifically *cognitive* science, and science-based philosophy, in our common, 'folk' psychological descriptions and explanations of our own mental experiences.

In folk psychological terms, we seem – both to our own introspection and, by inference, to each other – to have mental 'things' such as beliefs, desires, ideas and intentions. We also seem to be able, with our Minds, to imagine, to think abstractly about things not actually present, to plan, to remember, to intelligently control our own actions toward certain goals for certain reasons, and so on. In general terms, we could say (roughly) that it is in these kinds of descriptions that a number of apparently 'mind-ful' characteristics of ourselves as human beings are *marked out* for attention: and thus 'the Mind' begins to be identified as a subject of interest (to itself) in the first place. (Clark 1993, p. 4; Varela et al 1996, Ch. 1)

Furthermore, we also commonly *explain* our own *behaviours* in folk psychological terms: 'I did such and such *because* of this or that belief or desire.' In other words, we implicate our own mental states as being causally effective in shaping our actions and, indirectly, shaping the world we act upon. (And science, of course, has always displayed a strong interest in cause-effect kinds of relationships of any kind.)

Thirdly, in considering how our psychological descriptions mark out a notional domain of Mind, it is also important to note a broad distinction between Mind as something that seems to intelligently shape and guide behaviour, and Mind as the locus of conscious, subjective experience. In other words, beyond the folk explanations of our behaviours, we also have folk description (and direct experience) of ourselves as feeling, conscious *subjects*. This is the realm of what is sometimes called *phenomenal consciousness* (e.g. Clarke 2001, p. 174) – the subjective 'feel' of mental experiences. This too must be considered a part of the putative domain of Mind.

In simple terms, one could say that cognitive science questions begin when the methods of modern scientific enquiry are applied at these points – when one asks what in fact such 'naive' folk-psychological descriptions, ascriptions and explanations imply about our own real nature, and/or the nature of the world; or about the relationship between the (seemingly) mental and physical domains. (Fodor 1981, p. 124)

These basic questions are also reflected in the history of Western philosophy of Mind, and in particular we may note the contrasting approaches taken in a *Cartesian dualist* interpretation of Mind, and a broadly *materialist* approach.

The former position was of course most notably advocated by French philosopher, Rene Descartes, who theorized a Mind-body *dualism*; in essence proposing that the nature of Mind is non-material and therefore fundamentally different to all the rest of a purely material, machine-like world. (Clark 2001, p. 163) In terms of then explaining how the mental domain was able to influence the physical – how, for example, the intention to grasp a pen could initiate the required movements of arm and hand – Descartes proposed the pineal gland as the conduit of these influences.

Generally speaking the development of philosophy of Mind into the present represents a vote against Cartesian dualism and its separation of Mind out of the physical realm. Instead the more common choice has been to pursue a materialist worldview as outlined in previous discussion. In the making of that choice, however, the questions of causality, and of the relationship between the notionally mental – the things and explanations invoked by folk psychological description – and the physical recur in other forms.

### 4.2.2 The Mind-Body Problem

On the basis of materialist assumptions about the world, we take it that most of the physical stuff of the universe does not have mental properties or abilities. A piece of rock, we assume, does not have experience; does not think, dream, imagine, or do anything else remotely Mind-like. In temporal terms, we also assume that the early eras of the physical universe – prior to any possible appearance of life – also did not anywhere display such mental properties. And yet in this world apparently grounded in the material (and seemingly non-mental) we are now confronted with our own experience. Our capacities for intelligent reflection, ideas, desires and all the rest seem to suggest, to our own raw intuition at least, that we are in some way fundamentally different from the apparently unthinking, material things we see about us.

Thus the materialist approach to Mind has confronted what is commonly known as the 'mind-body problem'. If, as good materialists, we take the physical domain as a starting point, then how are we to adequately account for these signs of a distinctively mental domain (and explanations couched in mental terms) in this physical world; and how to use the tools of scientific enquiry and analysis – which otherwise would seem most useful in rendering the world transparent to shared understanding – in formulating that account. Cognitive science and philosophy, as we shall discuss shortly, offer a range of responses to that question, but nevertheless these are typically premised on an underpinning materialism.

Fodor in fact notes two sub-themes of the mind-body issue: one relating to questions of how a physical system could come to display intelligent behaviour; the other to questions of how a physical system could come to produce phenomenal consciousness. He also puts the general view that the former question is far more tractable to scientific enquiry that the latter, insofar as science would seem to offer some promising strategies for addressing it. (1981) The distinction is an important one here, given the commitment throughout to tie the discourse into a science-based, naturalistic worldview, and the intention in this chapter to focus on information about human morphogenesis, and the computational modeling of intelligent systems. In both cases, insofar as these areas of science inform what we refer to as Mind, it is Mind as intelligent, behaviour-shaping *cognition* that is the main focus. The intention here, therefore, is to conduct the discussion mostly within these parameters, those of the first mind-body problem as Fodor portrays it, and to leave matters of phenomenal consciousness aside.

### 4.2.3 What is Cognition?

The idea of 'cognition' may not in itself necessarily be an easy one to define precisely. Let us say, however, that in relation to the human domain, it generally refers to a notional range of *capacities* for intelligence, in the vein of those things mentioned above: a capacity for memory, for perception and discrimination of objects in the world, for language use, abstract thought, or forward planning. Barring mental illness, disability or other gross impediments, these are also capacities demonstrated by most adult human beings.

Furthermore, these capacities are very much tied into our functioning, as discrete individuals, *in an environment*; they all in their own ways imply or directly refer to this interface. We say cognitive 'capacities' I would suggest for this very reason, because, by and large, we are interpreting in ourselves various kinds of features that are *enacted* through processes of intelligent functioning and *behaviour* in an environment.

Thus, although not all cognitive capacities will necessarily 'show' in some linked, overt behaviour; we may again note how, when considering cognition, behaviour must be taken into account. Typically, the 'third-party' observation of behaviour is a crucial tool for the subsequent interpretations we make about other beings (human or otherwise) having 'internal' cognitive capacities at all, and for finding out what those capacities are, how they function in particular circumstances, and so on.

Regarding cognition as something typically enacted (in some sense) in an environment, and as notionally implicated in behaviour, now also leads us to a further point about what 'cognition' means. Let us say that we observe a person A walking along a path, when person B comes into their view. At this point person A smiles, waves and moves towards person B with his hand extended. In such a case, we would not (at least in terms of a folk psychological description) take it that A's behaviour is simply some kind of direct mechanical response – like a light coming on when the switch is turned. Instead, coming to some kind of interpretation of events in cognitive terms, we would very likely infer a *process within* A that, as it were, has come between the 'input' (to A) – seeing person B approaching – and the 'output' (of A) – smiling, waving and so on. We might interpret this internal process in terms of 'person A recognised person B as someone he knows and likes'.

In other words, to notionally regard human beings as having cognitive capacities, and as enacting these capacities through their functioning in an environment (at least at this folk psychological level), would also seem to imply some internal states or events (indeed the putatively cognitive states or events) that mediate in some effective way between their perceptions of the world and their behavioural responses to it[5].

[5] Although, of course, radical versions of the behaviourist school have argued against the positing of specifically mental causes of behaviour. (Fodor 1981, pp. 124-125)

### 4.2.4 Propositional Attitudes

Here again, therefore, we come up against the folk psychological interpretation and explanation of behaviour, as one that invokes mental 'things' – beliefs, desires and so on – as causally effective. One of the specific approaches to these kinds of explanation lies in the formal description of these notionally mental things (and causal processes) as *propositional attitudes.*

To designate a particular propositional attitude is, in simple terms, to ascribe to a person a particular item of mental *content*; for example, a particular belief such as 'Mary believed it was too hot in the sun'. In the same way notional explanations of behaviour can also be arrived at: 'Mary went into the shade *because*; she believed it was too hot in the sun, and she wanted to be cooler, and she believed it would be cooler in the shade.' In general terms, propositional attitudes, as explanations of behaviour, are seen as mental contents with this kind of belief/desire form. As Clark describes:

> The mental contents ascribed … are structures of concepts and relations … embedded in propositions governed by attitude verbs (believing and wanting).
>
> 1993, p. 3

As philosopher Jerry Fodor quite rightly observes, we do indeed, in our everyday lives, commonly ascribe particular beliefs and/or desires to ourselves and other people, and we explain or predict behaviour in these terms, and often these explanations or predictions prove to be reliable. For example, 'John will come into the kitchen soon because he believes there is a cake here, and he loves to eat cake.' (And, as it turns out, he did come into the kitchen, etc.) If nothing else, Fodor argues, our theory of Mind ought to be able to account for this explanatory reliability. (1987, ch. 1)

Further, it is by reference to propositional attitudes that folk psychology is also suggestive of a basic *causal economy* of perception, thought and action; in other words that, influenced by perceptual 'input' from the world, we form particular mental states, that in turn shape other mental states, and that these (sometimes) effect some appropriate behavioural 'output'. What is it, we may ask, about the (presumed) intervening, effective mental links in this causal economy that ensures, when confronted with some salient environmental feature or event, our cognitive capacities are usually able to initiate an *appropriate* environmental response?

### 4.2.5 Cognition as Computation

Again, a key issue from a cognitive science perspective is what such explanations and questions, couched in folk psychological terms, may have to tell us about the concrete nature and functioning of Mind in human beings. If we adopt a materialist perspective, and if the brain is taken to be a primary locus of our cognitive capacities, we begin to get a notional view of the brain as some kind of *device*, where the states of that device somehow correspond to the thought-ful, belief/desire elements of this causal economy.

In other words, we may begin to form a basic hypothesis about human cognition as computation; and here a number of key propositions are to be considered:

- Firstly, that mental events are in fact physical events in the brain (notwithstanding the involvement of perceptual organs, the nervous system, and so on)
- Secondly, that these brain events notionally feature *intentionality*; that is, not only are they brain states and events as such, but that by virtue of some particular features they are also *about* things external to the cognitive system (the person in question). Philosopher, John Searle, defines intentionality as 'that feature of certain mental states by which they are directed at or about objects and states of affairs in the world'. (1982, p. 358)
- Thirdly, that somehow the physical events going on in the brain, insofar as they act as part of the basic causal economy, are *semantically well behaved* with respect to the external domains of everyday human activities. In other words the kinds of cognitive abilities and behaviour (assumed to be) directly or indirectly affected by the workings of the cognitive device are (usually) sensible and appropriate to activities 'as disparate as trout breeding and subatomic physics.' (Clark 1993, p. 4-5)

We might say that these are a set of base-line hypotheses about human cognition, ones that would seem to follow from certain assumptions; and that much of the enquiry, debate and philosophy of cognitive science is about a more substantive interrogation of this basic picture. In terms of the specifically human capacity we may ask, if this is what our cognition (roughly) consists in, then how in the real workings of brain and body does it actually take place?

### 4.2.6 The Role of Computer Science

One of the contemporary areas of science that has now very much come to the fore in strategies to explain the nature and abilities of intelligent physical systems in general is computer science; including the development of digital computers, as an outgrowth of work in formal logic systems.

Of particular historical note in this regard, of course, is the work of Alan Turing, whose theoretical work in computational structures – crystallized in the concept of the 'Turing machine' – contributed to the construction of the first digital computers. (Clark 2001, pp. 7-11)

What these developments suggest in relation to Mind (amongst other things) is a way to arrive at a useful intermediate level of description between a folk psychological-type description of intelligence, and a purely mechanical account of physical states and events going on in an intelligent device (such as the brain). Inspired by the abilities of digital computers and the distinction between *software* and *hardware*, it is to say, in other words, that mental events might be described and understood in terms of their causal relations to each other within some internally consistent logic, rather than as a very specific kind of physical (brain) state or event. Just as with the computer programs now familiar to many, one can fairly readily grasp the idea that a particular program could be run on different kinds of computer hardware, and, if it still displays the same suite of *functional* relationships (e.g. it still produces 'fish' on the monitor when you type f-i-s-h on the keyboard), then one could conclude it is indeed still the *same program.*

In relation to human cognition, therefore, this distinction from within computer science has offered a way of conceptualizing cognition and giving some substance to the basic computational story. It proposes, in simple terms, that the mental states and events implicated in cognition might be understood and described in relation to physical brain anatomy in something like the same way as software is understood in relation to hardware.

Computer science, in other words, offers us a way to say *how* it is that a physical state within a cognizing device can also be *about* something external to the device; or to say how linked events in these internal states might come to be semantically well behaved in relation to some domain of activity. After all, well-programmed computers can be counted on thereafter to be semantically well behaved in various domains, like playing chess.

### 4.2.7 Symbols

In order to delve a little further into the conceptual 'space' between the physical and the computational (or the cognitive), and to better come to grips with such claims from computer science, we also need to introduce the notion of a *symbol.*

As Clark explains:

> [A] symbol is just a physical state which can be both nonsemantically individuated (i.e., recognized by its weight, its shape, its location, or any other property which can be specified without mentioning the content of the state) and reliably interpreted (i.e., assigned a meaning).
>
> 1993, p. 5

Thus the symbol is at the center of the 'overlap' between the physical and the computational. It is apparent within a well programmed digital computer, for example, that specific states of the computer hardware can also be assigned a meaning as specific (semantic) elements of a software domain: thus they qualify as symbols.

At a more general level, then, one of the basic ideas about how a physical device might achieve a capacity to affect semantic good behaviour is to suggest that it does so by virtue of *manipulating symbols*. It is to say, in other words, that that part of the basic causal economy attributed to the cognizing structures (whether a brain, a computer, or some entirely hypothetical device), within a semantically well-behaved system, is a matter of *certain kinds of symbols giving way to certain others is some reliable way*, such that the semantically appropriate behaviour is achieved.

For example, if we have a computer programmed to accurately translate German words into English words, then in order to be semantically well behaved in that domain, there are certain things the system would have to do. If we put into the system the word 'zehn', it would be required to produce the output, 'ten'. If this in fact occurs (along with many other successful translations) then we might reasonably conclude that the internal workings of the system through that transaction have been such as to capture the input as an individuated state (the symbol for 'zehn'), and to undergo certain state transitions so as to produce another individuated state, the symbol for 'ten', and to 'offer' this as output.

These symbols, and the ways in which they are reliably manipulated in a cognising structure, may also be said to be effectively *representing* certain salient features of an external domain; in this example, word-to-word meaning correspondences between English and German. Symbols, we might say, are items of representation.

The notion of representation will be gone into in more detail below in sub-section 2.9. For the moment, we may note that the ideas of symbols and representation offer an approach to the general idea of intentionality. They are part of a broad hypothesis, one might say, that intentionality is achieved by virtue of physical systems containing and manipulating symbols in ways that reliably represent salient features of some external domain.

### 4.2.8 Introducing the Classical & Connectionist Models

The investigation of human cognition-as-computation within contemporary cognitive science revolves around two main theoretical models; the *classical* and the *connectionist*. Each, according to its own fashion, extends upon some of the basic views about cognition as canvassed above. Each might be regarded as a theoretical 'school' encompassing a spectrum of more specific theories; each makes use of modeling with artificial systems. Although the major focus here will be on the connectionist view, it is important to have a reasonably clear picture of both. Occasionally the characteristics of the one are best understood by contrast with those of the other.

*The Classical View of Mind:*

The classical view about human cognition (and the various notional capacities marked out in folk psychological descriptions) takes up the view that the brain, as the 'central processor' of the human structure, conforms to the broad description of cognition as computation.

Building on that view, and taking some of its cues from the functional nature of digital computers, it also asserts a view about the way brain states and events instantiate symbols, and the ways those symbols are manipulated, combined, or give way to other symbols: theorizing that these computational processes have a broadly *linguaform* structure.

In order to grasp what this means we may note that language may be treated (by some degree of abstraction from everyday language use), as a formal, logical structure in which there is a finite stock of basic concepts (words) and a set of formal rules for combining those words into more complex conceptual structures (sentences). Out of this suite of tools it is possible to construct vast numbers of novel structures by combining and recombining elements according to the rules. Furthermore, it can be argued that, so long as the rules are properly applied then logical truths can be preserved through such combinatorial processes.

To say, then, that human cognition employs a linguaform 'computational economy' is to assert that the suite of symbols represented in those structures are something like the suite of basic concepts in language, and that the rules for combining (manipulating, etc) symbols in more complex symbolic structures, or through certain transitions are somewhat akin to the logical rules of (formalized) language.

This view about cognition seems to have a certain explanatory potency in extending the basic computational story in a way that also addresses some of the particular capacities picked out in folk description. It seems to offer a way, firstly, to explain the *productivity* of cognition; that is, the way the human Mind is able to link more basic ideas together to form more complex ones. Secondly, in the idea of preserving truths *through* certain computational manipulations, it offers a way to explain the capacity of the human Mind to enable our individual, semantically well-behaved engagements with all sorts of complex domains of shared activity and information. (Clark 1993, Ch. 1)

Thus in the classical view it is the physical realisation and manipulation of symbols within a linguaform computational economy that is understood to be the causally effective reality underlying folk-psychological ascriptions of one thought leading to another, or of thoughts directing behaviour. This one might say, is the broad (cognitivist) answer proposed to the Mind-body problem.

There is a range of experiments in artificial intelligence which attempt to model this kind of conception of intelligence, and how it might function within various domains. The SOAR project, for example, is one attempt to generate robust, intelligent negotiations through a range of 'general knowledge' domains, using linguaform, symbol processing computational architecture. (Clark 2001, Ch.2)

As an added note to this initial introduction to the classical view, we may also note a particular interpretation most notably espoused by philosopher, Jerry Fodor. Fodor, in putting forward his 'Representational Theory of Mind' is generally regarded as a realist about propositional attitudes and Mind. (E.g. Fodor 1987) What this means, in simple terms, is that he believes that the computational processes of the brain engaged in the realisation of propositional attitudes will (in due course) be shown to mirror fairly closely the causal relationships picked out in folk-psychological descriptions. In other words, he expects that there will be a reliable correspondence between the relevant representational brain states and processes and the kinds of belief/desire concepts and 'causal' linkages picked out in folk psychological terms. As Fodor describes it:

> To a first approximation, to think 'It's going to rain; so I'll go indoors' is to have a tokening of a mental representation that means *I'll go indoors* caused, in a certain way, by a tokening of a mental representation that means *It's going to rain*.
>
> (1987 p. 17)

It proceeds also from this view that the computational economy of Mind is theorized not simply as having a broadly linguaform structure, but as conforming quite closely to the actual structures of 'real-world' language. Thus Fodor asserts a view about the computational Mind employing a (representational) *Language of Thought*.

A significant point to note here is that the idea of complex cognitive events 'built up' in such a directly linguaform manner, and as rule-conforming combinations of 'smaller' symbolic elements, as a means of enabling semantic good behaviour, then requires certain things of the symbols in question. It requires, in short, that symbols at some basic level must retain the same semantically interpretable role, independent of any combinatorial context in which they're employed. So, as a simplified hypothetical example, it might require that a specific symbol semantically interpretable to mean 'dog' should bring the same, stable meaning to bear in cognitive events interpretable as, 'Our dog is in the yard' and 'I saw a dog yesterday'. One can see fairly readily that the capacity of a Language of Thought-using cognitive system to reliably 'build up' more complex cognitive items out of more simple elements, and thus to function in a semantically well-behaved manner, would be in doubt if the simple elements – the symbols – were not stable at some basic level; if they were liable to change or fluidity of semantically interpretable meaning.

*The Connectionist View:*

The other main computational model currently vying with the classical view is the *connectionist* view of Mind. Again in introductory terms, the connectionist approach is informed by a model of computational architecture, commonly known as *parallel distributed processing* (PDP). Rather than pursuing the formal logic-based computational style of digital processing, PDP architecture is (in a highly simplified fashion) modeled on the general neural structures of the brain.

At a very basic level of description the human brain is, of course, made up of a special type of cell, the neuron, and each one of these is connected, via axons and synapses to a large number of other neurons. In essence the physical workings of this network consists in the passing of electro-chemical impulses between neurons, such that the activity of each neuron at any time may be both affected by the impulses it is receiving, and exerting effects on other neurons by the 'signals' it is sending.

An artificial PDP network reproduces the basic architecture, insofar as simple individual processors – typically referred to as 'units' – are linked to multiple numbers of other units by 'connections': thus, of course, the 'connectionist' description. It is also this simplified modeling of the brain that begins to immediately suggest the more deliberately 'biological' approach that connectionism takes to understanding Mind. It is part of a wider picture that will emerge as the discussion proceeds of an underpinning intention expressed in connectionist approaches: to study, model and explain Mind in a way that is consistent with what is know about our evolved biology and the processes of pre and post-natal development. (Clark 2001, Ch. 4; Elman et al. 1998, pp. 104-106)

There are now many examples of such connectionist networks, used to investigate certain aspects of cognition. Typically, these systems are set up to undertake learning towards some kind of simple task competence. Beyond some basic rules 'written in' to the system in terms of the form of inputs and outputs, or learning rules for how units and connections interact, the systems display a (apparent) capacity to actively learn, and *develop* competencies in a task domain. Of particular interest here is that, in undertaking learning, these systems often employ a mode of representation quite different to that proposed by the classical model. (A detailed discussion of this view of representation will follow in Chapter 6.)

In this sense we might say that connectionism offers an alternative view about how physical systems might be intentional, and one that will prove to be important for the themes of this study. The connectionist emphasis on learning also has some engaging consistencies with the broader developmental approach to Mind and cognition as emergent aspects of human biology and morphogenesis.

However, in terms of teasing out these issues, discussing particular examples of PDP systems, the detail of connectionist theories, or the specifically connectionist approach to representation; these will be the subject of a far more detailed discussion in later chapters.

### 4.2.9 Representation

We should note, to begin with, that the notion of representation can be a tricky one to pin down precisely. It is an idea closely linked to a broadly computational view of Mind, and as Clark notes, an unequivocal definition of a computational system can itself be problematic (2001, p. 17). There are extant debates in philosophical circles about how anything can really be said to 'represent' something else at all. There is no intention here to engage deeply in such debates about what does or does not, could or could not, count as representation; however, there is value in clarifying the use of the term here, and capturing some of its typical meaning in a cognitive science context, and in reference to the human Mind.

Firstly, it is clear that the notion of representation implies some kind of relation between a thing or domain that is the 'object' of representation, and a physical system that is doing the representing. It generally follows, therefore, that representation (as something a representing system is understood to be doing) is to be understood as such by reference to a *context*; salient objects or events in the system's wider environment.

Secondly, in considering that relationship, the designation of a system as a representing one is commonly done on the basis of the system manifesting some kind of behaviour; and that this behaviour achieving what I referred to earlier as semantic good behaviour in a wider (environmental) context.

Thirdly, in terms of beginning to interpret the causal links between system and environment, one can crudely describe environmental effects on the system as inputs (to the system), and systemic effects on the environment as outputs.

For example, it is now fairly well understood that, with certain species of frog, a fly passing across a frog's visual field in a certain way will generally trigger a chain of events, including brain events, that lead to the frog shooting out its tongue in an attempt to catch the fly (e.g. as discussed in Maturana & Varela, 1987. p 125 –126). Here we have a system displaying a certain fly-catching intelligence; in an environment with small black flies flying about. Given that the fly catching behaviour reliably occurs when a fly passes the frog (and generally *not at other times*), we might fairly interpret a correspondence of some kind between visible (to the frog) fly-passing input, and tongue-shooting output. (Furthermore, it has been shown experimentally that frogs will also respond to a small, fly-like black dot in a similar way as they do to real flies.)

Other knowledge about frog physiology would not support an explanation of this ability in terms of any direct, mechanical link between the frog's eye and its tongue. On the contrary, actual neural pathways and corresponding brain events all point to a causal 'chain' routed via the brain.

Given all of this, we have some reason to believe that the frog's brain is doing something that effectively constitutes discrimination and recognition of a certain kind of pattern in the environment (one salient to the frog's evolved nature), and the subsequent triggering of a bodily response to that pattern. To invoke brain structures and events as the crucial factors in this capacity for discriminatory recognition is to assert nothing more or less than that the frog brain *represents* flies (or small black, fly-like patterns) in a certain way.

This representational story draws together the elements mentioned about: the relationship between system and environment, behaviour of the system in the environment, and some notional links between input, system-internal states and events, and output. We should note that there is nothing in this interpretation requiring that the frog actually thinks, 'There's a fly', or that it has a 'picture of a fly in it's head'. The representational view only says that, on the behavioural evidence and for other reasons, we conclude a *reliable correspondence between a salient environmental pattern, and some kind of brain states and events*. In context, the behaviour reliably induced via these brain states is effective as 'fly-catching behaviour'.

As outlined in earlier discussion, folk-psychological descriptions of ourselves also claim that we have certain mental states – in the form of beliefs, desires and so on. If we take the brain to be the seat of mentality, and as physical, then these claims would also seem to imply similar correspondences between human brain states and events and salient environmental patterns, such that the system as a whole (the person) is able to be semantically well behaved in a wide variety of domains of activity. The folk view, on these terms, is thus strongly suggestive of the brain as a representing device in some way.

The computational story inspired by folk psychology puts forward a basic view about what *representations* (notionally) *are*: they are physical states and events in the brain: ones that (in the sense we have been using) 'correspond' to certain salient patterns in the environment; that lead to other states and events; and that, sometimes, trigger certain behaviour in the system as a whole (Clark 2001, p. 15).

Not any environmental pattern, we surmise, will lead to any (representing) brain state, or to any behaviour. On the contrary, by the folk psychological account, the common occurrence of contextualised behavioural appropriateness – it *was* hot in the sun; she *did* go into the shade – leads us to interpret a representational process that enables semantic good-behaviour, insofar as it can reliably 'pick out' salient patterns in the environment, engage in some mediating internal activity, and reliably initiate behavioural responses that are appropriate in some way to those patterns.

As discussed above, the classical computational theory of Mind then goes on from that basic position to put a view about *how* the human brain represents; in short, that brain states and events represent information based in symbols, and the way representations are operated on mirrors the formal properties of language, and of digital computation. When operated on in these ways, these symbols act as the 'building blocks' of more complex representations.

Within Fodor's 'Language of Thought' model, the whole system is theorized as operating in a way that mirrors the structures of real language and folk-psychological descriptions of mental events. Such a close lining up of the symbolic and representational dispositions of a system with our own, language-based conceptions would mean that such a system could be considered *semantically transparent*. (Clark 2001 p. 29) In other words, if we had some means of identifying and reading (semantically interpreting) the symbolic content instantiated by such a system, we would read a body of semantic items with some close correspondence to the words and structures of our everyday spoken language.

### 4.2.10 The Approach to Cognitive Science Issues Here

With this background information about cognitive science in place, the discussion can now move on to consider more directly the view of biological systems as put forward in the previous chapter, how it may apply to human morphogenetic processes, and then its applicability to the nature of Mind. It must be understood, however, that when the discussion returns to consideration of cognitive science ideas as such, the aims will not be to enter directly into most of the issues and debates around the competing claims of these two broad models.

Instead, the purpose will be to reflect upon ideas and information generated within cognitive science about Mind in relation to the perspective of previous chapters, and in terms of the model of creative organisation to be proposed. In so doing, the bulk of the discussion about computational models will focus on the connectionist model because it seems to be the one more in keeping with these perspectives. The goal is *not* to otherwise argue in a more general sense either for the connectionist model per se, or against the classical model.

## "Mind as Creative Organisation"

# Chapter 5: Morphogenesis & Mind

### 5.1 Introduction

The first major elements of the case for Mind as creative organisation will focus on some aspects of human morphogenesis, and neurogenesis in particular. The broad goals here are, firstly, to discuss these processes in terms of the kinds of characteristics identified at the conclusion of Chapter 3; secondly, to suggest how it is that morphogenesis underpins the emergence of characteristic features of Mind; and thirdly, to say why thereby morphogenesis 'shading into' Mind should be understood as a process of creative organisation.

It is suggested, in this sense, that our cognitive capacities are largely prefaced and proceed out of a period of pre and post-natal development where the 'construction' of the major features of human anatomy takes place. The discussion that follows, therefore, is about the realisation of Mind in this sense, and over the timescale of morphogenesis. I will initially consider some more generic aspects of morphogenesis at the level of the cellular 'building blocks' and the formation of multi-cellular structures. Following that, the focus will increasingly fall on morphogenetic processes seen to be more directly implicated in the eventual realisation of cognition. It will become clear, one hopes, that there are no neat points of division between these processes occurring over a timescale of months and years, and the 'launching' and progressive realisation of cognition.

Looking further forward again, we may note again that discussion in later chapters will eventually enter areas touching on matters of human behaviour; specifically the relationship between behavioural characteristics and cognition. For the moment what I would wish to note is only that behaviour is a factor taken here to signal a shift in orientation in the developmental pathway, from predominantly system-internal morphogenetic processes to processes that increasingly must be seen as shaped by or responsive to the system-external environment.

To this extent, then, some of the study will come to refer, if indirectly, to questions about the links between morphogenesis, cognition and behaviour; and thus will touch upon questions of *nature* versus *nurture*: that is, to what extent do behaviours, and/or properties of Mind, and/or the workings of the brain occur as inherent characteristics of the individual/organism in question; or to what extent do they arise as a product of the interactions between individual and environment (including of course, other people)? Is intelligence (or aspects of it) better understood as a product of a 'mechanical' process of development – simply the implementation of a 'pre-set' genetic blueprint – or is it more a matter of environmentally situated learning, or something that incorporates both these aspects?

Here (as discussed below) a distinction between genotype and phenotype will be relevant, and it will become clear, one hopes, that the approach favored here is one which cuts across any simple dichotomy of nature *or* nurture; recognizing instead a complex, developmental interaction between genotypic constraints, environmental factors, and the particular pathway of phenotypic instantiation. The connectionist model of Mind, along with associated ideas about morphology, is directly relevant to substantiating this approach; and, by implication, it also argues against suggestions of simple, linear causation from genes to Mind or behaviour. In the next chapter these issues and tensions will also recur in terms of how they might inform social interpretation.

We may note also that some of the ideas presented here about morphogenesis will also be relevant later to the idea of connectionist models in cognitive science having a level of 'biological credibility'. The seemingly consistent character of the processes implicated from morphogenesis to cognition itself is broadly regarded as lending some support to this assertion.

### 5.1.1 Morphogenesis as Creative Organisation: An Initial Overview

There can be no doubt that, in some sense of the word, human beings are *developmental* creatures. We each begin life as a zygote, the union of male and female sex cells, too small to be seen with the naked eye. All being well, this *develops* into an embryo, which develops into a newborn infant, which develops into an individual capable of language, and on it goes.
So too, in this general sense, would we assume the nature of Mind to be developmental. We would not regard the zygote as yet possessing any of the abilities or properties suggested above as key markers of the mental (Edelman 1992, p. 52). Again the human organism must be seen to *develop* its mental capacities through a morphogenetic and maturational process, and there one may seek to pick out the key factors that enter into the realisation of Mind.

In the previous chapter, it was suggested that living organisms embody a dynamic 'balance' between overall systemic ordering, and the flux and variability of the material instantiation of that ordering at any particular time. This I described in terms of two aspects: on the one hand, that organisms display a capacity for *stability* and continuity of overall systemic ordering – and that these could also be regarded as forms of *constraint*. On the other hand, it was suggested that within those stable constraints, the constituent elements of the systemic process, on a moment-by-moment basis, are also subject to changes under the influence of a range of causal factors – and these aspects were characterised more generally as a capacity for *variability*. In what follows this description will be substantiated with examples relating to human morphogenesis. Here it will be argued that both aspects of this dual capacity are necessary to these processes. (In particular, variability is regarded as a significant factor in its own right, and not merely as 'noise' that is present but incidental to the outcomes.)

Finally, I noted that living systems can display developmental transitions, where 'higher' levels of organisation emerge from 'lower' ones, and I suggested that it was here – in this multi-leveled 'building up' of complexity and ordering – that we could identify a process of *creative organisation*.

The consideration of processes of morphogenesis, now allow for an expanded description of these various aspects, and of their significance within far more complex biological processes. As suggested above, the concepts of *genotype* and *phenotype* are central; where 'genotype' refers to the evolved genetic constitution of an organism as a member of a species; and 'phenotype' refers to the specific physical constitution of an individual organism, shaped by both its genetic inheritance and its individual history in an environment. Their general significance in beginning to develop a view of morphogenesis-as-creative-organisation is as follows:

- Firstly, the genotype is seen to encode and enable features that provide *stability* and consistency of systemic ordering in complex organisms, and thus to provide certain *constraints* on morphogenesis;
- Secondly, the phenotype is seen to realise the developmental potential encapsulated in the genotype via processes and structures that are open to *multiple causal influences*, and thus also to embody some degree of inherent *variability*.

In the previous chapter I suggested the notion of biological systems embodying in their organisation a general dynamic 'tension' between stability and variability. The mutually reinforcing roles of genotype and phenotype are now being put forward as a concrete expression of this general character within morphogenesis. As such, morphogenesis will generally be seen to also constitute a process where 'Genetic *constraints* interact with internal and external environmental influences, and they *jointly* give rise to the phenotype.' (Elman et al. 1998, p. xii; emphasis added.)

Furthermore, the morphogenetic processes to be considered will also be seen to display at times *emergent* properties, and to constitute a process of *epigenesis*: that is, a developmental process where certain developmental 'steps' necessarily precede others, and thus the timing of processes in relation to each other becomes a factor to consider.

Assuming for a moment that morphogenesis does display this range of features, what then is it about the combination of factors that might lead one to identify this process as 'creative'? A number of key arguments will be developed in what follows, but in simple terms the description of *creative* organisation rests on several points:

- Firstly, because multiple levels of organisation are seen to be realised;
- Secondly, because these systems display a characteristic variability, and because they are subject to multiple causal influences, it is seen that effectively the phenotype over time instantiates an 'ad hoc' combination of variables. I will suggest, in other words, that these various factors coming together in each developing organism enter into the formation of a *unique phenotypic history* – one that embodies the species-consistent constraints provided via the genotype, but also represents an idiosyncratic, individual, 'case-history' instantiation of those constraints. Thus, over and above the differences between individuals arising because of their different genetic make-up, there will be differences that arise through the development of the individual phenotype, generating greater divergence between cases and diversity across cases. The idea of variability will thus be regarded as a significant factor in itself in shaping the unique phenotypic history both within and beyond morphogenesis.
- Thirdly, I would argue that essentially the same combination of characteristics continues to be displayed as morphogenesis dovetails into the occurrent nature of cognition, socially embedded human learning, and the maturation of the individual personality. There, I think, creative organisation as characterised here extends into that which counts as 'creativity' in a more conventional, socially mediated sense. Thus there is a continuity of character, and this is seen as potentially very useful in developing our understanding of creativity in the

latter sense (although the scope to explore these 'further dimensions' in this study is very limited).

I would also add as a further important point that an understanding of creativity in the sense offered here is seen as tool with the potential to be transposed into a more explicitly social or political discourse, and to inform the relationships between socially expressed values or norms, and individual development.

The discussion to follow will be concerned to substantiate and further inform this outline description.

As a final preliminary point on description, we may note that the complexity of human morphology seems to demand of science generally descriptions and analyses at multiple scales – from individual cell structure to whole-of-system behaviours – and multiple levels identifiable within whole articulated structures. It is not un-typical of cognitive science literature, for example, to range from neuronal cell structure and function, to micro-structural, cellular networks or 'modules', to macro-structures of the brain, to a nexus of interaction involving brain, body and environment. As we have already seen, the computational models also display this tendency: for example, in the classical vein, with distinctions between the physical structures, the individuation of symbols within those structures, and the building up/causing of more complex symbol types through the application of combinatorial rules.

## 5.2 Creative Organisation & Morphogenesis

In discussing morphology, and the morphogenesis of the human organism, a basic 'unit' of development to consider is the cell. We each begin life as a single eukaryote cell; a cell with a nucleus in which the genetic material is carried. In the developmental processes that proceed to occur, the cell continues to feature as an essential constitutive element in the formation of larger structures, including of course brain structures. This part-whole relationship will continue to feature strongly as one where epigenetic developmental processes, and a character of creative organisation, are in evidence.

In reference to some of the relevant evidences about these and other processes, it must also be borne in mind that the scientific study of brain morphology is often not, for obvious ethical reasons, always appropriate via direct examination of human subjects. There is some necessity therefore to make inferences about human brain development from experiments on animals. (Elman et al. 1998, p. 287)

This section will focus on the development of brain structures, initially considering matters at the cell-micro structure level. The key features to be considered in terms of the above model are:

- That the causal factors seen to be at work within developmental processes are multiple, not single and (only) genetic.
- That the morphogenetic processes are stable and consistent within genotypic constraints; while also being variable insofar as the specific realisation of those patterns through an ensemble of constitutive elements is not genetically pre-determined, but follows a unique pathway of phenotypic instantiation, within overall constraints.
- That the properties of complex structures arising through morphogenesis are (often) appropriately understood as having an emergent character;
- And that the processes whereby these structures develop may be regarded as *self-organising*, insofar as there would seem to be no single, prior factors in evidence acting to determine the outcome of processes that produce emergent organisation.

### 5.2.1 The general role of the genotype

It is generally clear that the genetic material contained in each cell provides powerful constraints on the pattern of development that ensues from the moment the zygote is formed. After all, cells containing human DNA can be consistently relied upon to develop into the species-typical human form, in all its complexity: just as turtle DNA reliably turns out turtles. The history of genetic science and heredity since the work of Gregor Mendel in the nineteenth century now offers a broad consensus about genetic material as the essential factor of inheritance. There is no sense in which that consensus is being questioned here.

What is more at issue, however, is the question of how the genome carries out this role within the phenotype – and, subsequently, how the exercising of this role is understood to 'carry through' into the nature and behaviour of the mature organism. Some views about genetics appear, at least on the surface, to portray the genome as an encoded 'blueprint' for both the morphogenetic process itself, and the 'finished' organism. They suggest, in other words, *linear* processes of development, where a pre-set, genetically expressed design exists, and development itself may be understood simply as the construction of that design. It is to suggest forms of 'direct causality' from genotype to key aspects of the mature structure and functioning of the phenotype. Of some interest here, from the point of view of social interpretation, are suggestions that this 'direct causality' enters into the nature of Mind (and issues of representation). (See Chapter 8)

As a place to begin to weigh such claims, we may consider the (over) simplified picture of the genome coding directly for every single micro-structure in a human body consisting of roughly 100 trillion cells. As Elman et al. point out, these are instances of small organisms where such direct coding does seem to occur. They cite the case of the nematode, *C. Elegans*, as an example, referring to this mode of development as 'mosaic development'. (A primary piece of evidence for this conclusion is that genetically identically individual nematodes display essentially the same details of cell morphology.)

However, as they go on to argue, such direct coding would only seem to be possible in more simply structured organisms; where there is some direct correspondence between the informational capacity of the genome and the structure of the organism. As far as human beings are concerned, they suggest that such 'direct specification of the human brain alone, for example, could plausibly require something in the order of 10 trillion base pairs of DNA, which is far in excess of what is structurally feasible'.

'Mosaic development' in some simple organisms, therefore, is contrasted with 'regulatory development'; described as:

> Regulatory systems rely heavily on cellular level interactions. The orchestration of cell differentiation and the final outcomes are under broad genetic control, but the precise pathway to adulthood reflects numerous interactions at the cellular level that occur *during development*.
>
> 1998, p.13-15

Without wanting to anticipate discussion of such developmental processes, such issues of informational capacity suggest against a too simplistic picture of direct, coded specification of phenotype by genotype, and lead us to consider other options.

### 5.2.2 'Place & Time' Factors

In general terms, the alternative to a view of direct coding lies in a recognition of other causal factors (other than genetic factors) that bear upon the actual development of cell structures that occur. Above I referred to the influence of environmental factors, and these may be considered both in terms of *place* and *time* related factors. As far as 'place' environmental factors are concerned, generally what this means in terms of the earlier, morphogenetic processes of cellular differentiation and formation of micro-structures, is an environment of other cells. What is often identified as a domain of influence on the developmental processes at this level, is the responsiveness of cells *in relation to* other cells around them. In other words, these relational factors are understood to make a difference in terms of how precisely the process proceeds, and what precisely individual cells are doing as 'participants' in that process.

One area where this has been shown is in the plasticity of neural cells in the still-developing fetal brain. There are now numbers of studies with vertebrate animals showing that, in the development of specific brain structures with particular functions; fetal neurons already developing within one structure can be transplanted into another structure, and will begin to develop in accordance with that new placement. For example brain tissue from the visual cortex of rats has been transplanted into the sensorimotor cortex. The neurons ceased to function in a way characteristic of visual cortex cells, and began to function as sensorimotor cells. (Schlagger and O'Leary 1991)

Such results suggest that these cells, at least in this phase of development, are indeed sensitive to their context, and that the specific nature of their development is therefore shaped by this context; and seemingly is not determined a priori by the genome in this way at least. Again, this is not to suggest that genetic factors are not vital. It is to say that the influence of genes can be (and is) indirect, and when other factors (such as location within a developing brain structure) can be relied upon to provide sources of constraint in themselves, there is no need (from an evolutionary point of view) for the genome to code this directly, in order to generate more complex structures. It is this combination of genetic and environmental factors that is one marker of epigenetic processes; it is the capacity of cell ensembles to self-organize under these multiple factors that places them in the category of Elman et al's 'regulatory systems'.

Alongside such examples of environmental factors couched in terms of placement – a spatial context, so to speak – it is also important to recognise the importance of timing, and the sequence of events as further influencing factors on the formation of brain structures. Edelman discusses the example of the mapping that occurs between the retina of a frog and a region of its brain called the optic tectum. Mapping occurs such that, in a developed frog, there is a correspondence between the configuration of light sensitive cells in the retina, and neural cells in the tectum.

Although the mapping area of brain tissue itself first forms independently, its full functional development, and the establishment of the mapping relationship, relies on iterative feedback between retina and tectum as 'differential growth occurs' in both areas. (1992, pp. 22-24) The full development on the tectum, in other words, depends upon the corresponding development of the retina and optic nerve, and would not occur if the feedback were not available. This is a matter of timing in the general sense, and is the key aspect of epigenetic processes; that certain events need to occur before other events can occur. We may note also that in this example the focus of 'environmental influences' on cellular function and growth has shifted somewhat, from the influence of other cells with a developing structure, to the influence between structures.

### 5.2.3 Phenotypic Variability

Under the influence of such multiple factors within the development of brain structures, cells are moving, dividing, dying, forming synaptic connections and so on. Through these dynamic processes, the specific neural micro-structures that result in each individual phenotype will be somewhat different, *even between genetically identical cases*, whether in water fleas (Edleman 1992, p. 26) or human identical twins (Elman et al. p. 298). The multiple factors, in other words, begin to constitute a unique phenotypic history, and through that history we see a definitively individual realisation of the genotypic constraints. So then, this is how the aforementioned organisational 'dynamic' manifests at this morphogenetic level: between the stable continuity of a developmental pattern according to the genotype, and the variability that comes about through the morphogenetic instantiation of that pattern in each phenotype.

However, although there is such evidence for multiple causal factors influencing morphogenetic processes and structures at this level, it may be argued nevertheless that the genetic factors are pre-eminent insofar as they (albeit more indirectly) still constrain and effect the overall pattern and function of the larger brain structures that develop along with the organism itself; and, for the purposes of understanding Mind as a computational process and brain as a computational device, it may be these macro-structural consistencies which matter. It is to say, in other words, that the variability apparent in the formation of micro-structures (under the influence of multiple factors) becomes insignificant within a genotypically consistent global formation of macro structures. The variability, on these terms, may be considered as analogous to the background level of static 'noise' that comes with a radio signal. By this argument, it is the translation of the formal elements of the signal into the message that is important; the noise can be ignored – it doesn't signify anything meaningful.

Now, if it is the case that the variability 'side' of morphogenesis wanes in significance as the developmental processes proceed, then there might indeed be a case for such an argument. The point about the above examples, however, is not simply to indicate the presence of both stable constraints and variability at the cellular-micro-structural level, it is also to form part of an argument for saying that essentially the same character continues to be in evidence 'all the way up'. Variability (or 'noise') and multiple causal factors, on these terms are important in themselves, precisely because they are seen not only to enter into micro-structural formations, but to continue to be present and significant in the development of the brain as a whole, and into the emergence and maturation of Mind. As Clark suggests:

> A multi-factor perspective leads rather naturally to an increased respect for, and theoretical interest in, what might be termed the historical idiosyncrasies of individual development. What needs to be explained here is the delicate balance between individual variation and developmentally robust achievements. ... Soft assembly out of multiple, largely independent components yields a characteristic mix of robustness and variability. The solutions that emerge are tailored to the idiosyncrasies of context, yet they satisfy some general goal. This mix, *pervasive throughout development*, persists in mature problem solving and action. Individual variability should thus not be dismissed as "bad data" or "noise" that somehow obscures essential developmental patterns.
>
> Clark 1997 pp. 42-44

### 5.2.4 Emergence & Self-Organising Processes

The idea of emergence appears in much of the literature relevant to this discussion and, as Clark observes, with some variety of interpretation. (2001, pp. 112 – 117) For purposes of this study, the above characterization of morphogenetic processes can also offer a more refined and specific grasp on the linked notions of emergence and self-organising process in this context.

On the whole it seems difficult to evade at least a general view that complex, composite structures of brain morphology do indeed have properties, or display systemic behaviours, that are qualitatively different from the properties or behaviours of their component parts considered separately (individual neurons, for example), and thus are appropriately seen as instances of emergence as per my previous discussion in Chapter 2. For example, certain brain structures are empirically shown to 'map' (that is have a configuarational correspondence to) certain distinct sources of sensory 'input'. It is not clear that the individual neurons making up these structures are able, in and of themselves, to 'map' in this way, nor that simply having a bunch of neurons in the same place enables mapping to occur. Edelman argues that the interconnected ordering of neurons as functional structures at various scales is crucial to the mapping capacity. (Edelman 1992, Ch. 3))

However, liberally interpreted in this context, the initial definition may be in danger of becoming too broad. The further refining of it is closely related to the above discussion insofar as it proceeds also from a broadly developmental perspective – that is the view of developmental *processes* as epigenetic, and as *self-organising*.

As discussed above, a key to a description of processes as self-organising lies in the recognition of multiple causal factors, and that the *combination* of these factors presages and enters into the development of a larger morphological structure. It would seem to follow from this view that no one factor could be fairly represented as *the* single and only causally effective factor within the process; prescribing the outcome by causally directing the participant elements. It is the (apparent) absence of such a 'central controller' that leads to a view of processes as self-organising: instead it is a process of interactions amongst the constituent elements, or between elements and environmental factors, which is understood to generate the outcome.

Where the 'outcomes' of these complex, interactional, self-organising processes include 'salient patterns of systemic behaviour' (Clark 2001, p. 114) one would generally expect that these behaviour patterns will also be appropriately regarded as emergent features.

For example, the visual cortex of a number of species (including ourselves) is known to display areas of tissue that primarily serve one or the other eye. These areas are referred to as *ocular dominance columns*. It is known also that these configurations of 'preference' on the cortex are not present from birth, but develop thereafter. (E.g. a study by LeVay, Stryker & Shatz 1978 on the visual cortex of cats; cited in Elman et al 1998 p. 87.) Relevant studies also show that the development of the columns requires a normal level of environmental 'input' to both eyes. (Abnormal 'input' conditions show abnormalities in, or under-development of, column formation.) Subsequent simulation studies – showing strong predictive success for both normal and abnormal outcomes in animals – (Miller, Keller & Stryker, 1989) suggest that the 'columns emerge as the result of naturally occurring differences in the degree of correlation between neighbouring cells in each eye versus across eyes, along with competitive interactions among … cells.' (Elman et al. 1998 p. 87)

In other words, the development of columns is understood to *emerge* from the coincidence of a number of factors, including system-environment interactions, without the need for of presence of any single pre-determining (perhaps, genetic) specification for column formation. This is an example of a self-organising process, and the columnar pattern that results is also regarded here as a instance of emergence under the above description. Again, the process displays the characteristics of variability within certain system constraints; the precise pattern of columns will be unique in each case.

According to the perspectives taken here, the formation of these columns therefore counts as a (relatively minor) instance of creative organisation. The 'completed' columns are seen as manifesting a salient pattern of organisation that does not appear in the compositional elements, and so the developmental 'step' between the elements and the overall pattern is understood as creative organisation in action.

### 5.2.5 Brain Architecture & Connectivity

In light of such examples as those mentioned above, and through the idea of epigenesis, we now have some reasons to re-state the role of the genotype in phenotypic development. If such examples are capturing something of the general nature of human morphogenesis, then it would seem clear that over-simplified notions of direct genotype encoding of detailed, prescriptive 'instructions' are inadequate to the reality. Instead, a view is emerging of the genotype specifying some key *constraints* that strongly shape the character of multiple developmental processes, without 'having' to specify the precise way those processes 'play out' in each instance. In more concrete terms, this might mean that the genotype 'directs' the kind of cells that may be present together in a particular location, and the types of extra-cellular molecules that are present, which act on the cells in particular ways. What then makes the process an epigenetic one (at this level), is that these cells may provide 'feedback' influences on each other within a developmental process, or that the instantiation of 'higher' structures is shaped by factors external to a particular ensemble of cells, molecules and so forth. (The example of ocular dominance columns being a case in point.)

In terms more specific to brain development, then, the general notion of genotypic constraints can be recognised in a number of aspects of brain *architecture*; that is, the actual deployment of neurons into composite structures, the formation of synaptic connections within and between these structures, and ultimately the overall structuring of the brain itself into a number of highly inter-connected macro-structures (e.g. the cortex, hippocampus or cerebellum). Such constraints might include, for example: (at the cellular level or micro-structural level)) the types and response characteristics of neurons found in particular areas of the brain; or (at the global level) the disposition of major structures within the brain, or the general pattern of connectivity between major structures. (Elman et al. 1998 pp. 28-30)

The formation of these structures may (and should, in the view taken here) still be understood as epigenetic and subject to multiple causal factors, but in a way where a 'normal' mix of factors in embryonic development and early childhood development seem adequate to account for the result. The more extended notion of genotypic constraint allows for this, insofar as we can say that, given 'normal' developmental conditions, the phenotype will exhibit these genotype-typical features.

Notwithstanding these consistencies, however, the parallel character of phenotypic variability continues to hold also. Of particular relevance here is the profoundly dense and complex networks of synapses and/or nerves linking ensembles of neurons together, linking brain structures to each other, and linking the brain to the nervous system and sensory organs. Within the early development of the brain, for example, neurons develop a staggering level of synaptic inter-connectivity at the local scale. At this level, the variable aspect of neurogenesis is reflected in the stochastic nature of these formations. Multiple causal influences mean that the instantiation of these connections will not be repeated in precisely the same way between any two cases. (Edelman 1992, p. 64).

It is now well understood from neuroscience that the development of brain connectivity continues well into early childhood. Some evidence here also suggests that the formation of connections also occurs in a wider context of multiple influences, implicating to various extents brain, body and environment as significant, interacting factors. (E.g. Quartz 1999)

In order to go more deeply into the implications of all this connectivity, however, and to begin to look 'outwards' (at what the brain *does* as one part of the bodily structure, and within the relationship between organism and environment) it will be useful now to return to the connectionist modeling of brain structures and learning capacities. In doing so we now begin to canvas ideas that will take the discussion beyond morphogenesis and the initial formation of brain *structure* (and the timescale germane to those processes), into questions of what that formed structure is *doing* over shorter periods to launch and to develop cognition itself.

"MIND AS CREATIVE ORGANISATION"

# CHAPTER 6: CONNECTIONISM, CREATIVITY & COGNITION

## 6.1 CONNECTIONISM

We now come back to one of the key issues identified in the earlier outline of cognitive science models; that of intentionality, and questions of how the brain (seemingly) represents aspects of the external world. It seems quite clear, on the face of it that I, while sitting at home, can conjure up a 'mental image' – probably with a fair degree of accuracy – of what is on my desk at work. Such abilities of Mind suggest that, in some sense, my brain has somehow stored this information, to the extent that I can deliberately 'access' and use it without actually having to have my desk there in front of me. This is one example of what *intentionality* can mean to contemporary cognitive science and philosophy of Mind; that mental states are seemingly *about* other things, things external to the experiencing individual. The issues of representation of particular interest here arise when one asks what, in more concrete terms, are these putative mental states, what form do they take in (assumed to be) physical structures of the brain, and how, if in any way, do they enter into more complex 'chains' of thought – where, perhaps, thinking about my desk as work (seemingly) leads me to remember that I need to call Fred on Monday – or into behaviour?

As outlined earlier, the classical and connectionist models of intelligent structures offer some answers to these questions. In sub-section 2.9 we considered some general ideas about what representation entails. We may now more on to consider in more detail the connectionist 'take' on representation.

### 6.1.1 The Connectionist View of Representation.

As described earlier, the architecture of an artificial connectionist-style system is loosely based on the basics of neural architecture in biological brains, and consists mainly of individual unit processors (analogous to individual neurons) and multiple connections between these units in various configurations (analogous to axons and synapses).

When these systems are functioning then, again in a way loosely analogous to brain neuronal function, each unit will be receiving input, either as information from outside the system, or from other units to which it is connected. Each unit will be set up to operate on those inputs according to some mathematical function, and then – depending on what these operations entail – producing output signals to other units. The operations conducted on input by each unit can vary widely, and have a non-linear character. It is not simply a matter, in other words, of a direct linear correspondence between input and output – 'receive a signal; send one out'. For example, a unit might do nothing, until the combined strength of input signals from other units reaches a particular level.

Added to the intricacies of unit function, the connections between units are assigned a numerical weighting, one that can either inhibit or amplify the signals passing through it. The signals arriving at a particular unit, therefore, will vary according to combinations of output from other 'sending' units, and the intervening *connection weights*.
(Clark 2001, Ch. 4)

When a connectionist system is engaged on a learning task, it is these connection weights that will be changed, by automatic procedures known as 'learning algorithms'. For example, one of the better-known procedures is known as *back-propagation.* In simple terms, when a system is engaged on a learning task, it begins with randomly assigned connection weights. Some input is then provided to the system (via input units), and the output value (from output units) is compared to the target value it 'should' have to achieve success in the task in question. Adjusted values are then 'back-propagated' through the system from the output 'end', in order to decrease the level of error, by changing the connection weights involved in producing the first value. As this process goes through a number of iterations, the output values generally move closer to the target values.
(Elman et al. 1998, pp. 66-70)

To get a somewhat better feel for what all this means in more practical terms, let us consider an example. The NETtalk system (Sejnowski and Rosenberg 1986) is one now well-known example of a connectionist system in action: in this case focused on a task of translating written English words into the appropriate 'pronounced' sound of the word. The system was constructed in three layers: an input level consisting of 29 units specifying particular letter inputs, 80 'hidden' units (neither directly input or output related), and 26 output units, producing coded versions of phonemes (word-sounds).

The training of the system takes place with a body of training data; in the case of NETtalk a number of 7 letter words. Initially, again, the connection weights of the system were randomly assigned. As each word was put into the system, and an output produced (in the form of coded phonemes), a learning algorithm was applied to make small adjustments to each connection weight, towards a target of reduced overall output error. When the error level was reduced to a certain point, then the system was considered (to that extent) to have learned the task. The output, in other words, reliably coded, with a fair degree of accuracy, the 'correct' English phoneme-combination of the input word.

To begin to interpret what connectionist modeling has to offer, we may begin by noting that the NETtalk system, along with many other examples, has demonstrated a basic level of success within the task domain set for it: outputs do indeed move closer to the target standard over time. What then might we understand by representation occurring in such a system?

To follow the same example, given that NETtalk, having been trained, continues to exhibit consistent success in matching an appropriate (coded) pronunciation to a number of word inputs; clearly we may conclude that the translating processes have somehow been encoded into the system, and therefore it may be said (in this sense at least) to *represent* these processes. They are *available* in the disposition of the system, to be applied to an input, and to reliably produce an appropriate output. If this is so, then we may ask how, in more detailed terms, is representation understood to be instantiated in this connectionist system, and how does it compare to the classical or 'Language of Thought' models of encoding and manipulating information?

### 6.1.2 Distributed Representation

Although some of the early connectionist systems use a form of representation similar to the classical, digital form, most of the more recent and advanced systems display what is known as *distributed representation*. In order to understand this term we must first note, in a preliminary way, a distinction between two aspects of representation in connectionist systems. Firstly, to follow the example used above, if we considered the trained NETtalk system *at rest*, so to speak – not actively doing any text to phoneme translation tasks – we could say nevertheless that the array of (trained) connection weights is 'holding' the system's potential to reliably engage in a whole range of specific representational processes within its particular domain of competence. The *connection weights* in themselves may therefore be legitimately understood as a form of representation in this sense.

When NETtalk is actively doing what it is trained to do, however, then for each item of input (a piece of text to be 'translated') a specific pattern of activation will occur within the system; one we may think of as representing at that time, that particular instance of task competence. In other words, the *activation pattern* is a more specific, activity-linked form of representation arising out of the 'background' array of connection weights.

Both of these aspects of connectionist-style representation will need to be borne in mind in the discussion to follow.

To say that a system displays distributed representation is to say, firstly, that the activation patterns arising in relation to specific inputs are typically *distributed* across 'large' areas of the system array; and, relative to the overall size of the system, many units and connections are activated in the representational process.

Secondly, given this distributed character for all or many of the activation patterns occurring from time to time within one system, it is not surprising that individual units and connections will often be, at different times, activated 'participants' in a number of representations. It is for this reason, indeed, that the disposition of connection weights in a typically trained system may be seen to carry the 'overlapping' potential for a whole number of representations. As Elman et al. point out this means that:

> [In] order to know which concept is being considered at a point in time, one needs to look across the entire pattern of activation (since any single unit might have the same activation value when it participates in different patterns). The information needed to decide which concept is being represented is *distributed* across multiple units.
>
> 1998 p. 91 emphasis added

Such a description is of central interest here, because it brings to the fore a main point of difference between the classical computational view of representation and the connectionist view. By contrast with the classical notions of representation built around individuated symbols, distributed representations cannot be understood to be present in a connectionist system as discretely coded items. Instead they must be conceived more as dynamic processes, which occur across multiple elements of a system, in response to a particular kind of input. They exist not as a suite of fixed, individuated symbols within a computational structure, but as a one of multiple representational potentialities encoded within the array as a whole.

We may see then that distributed representation systems also offer a view of representation quite different to the Fodorian, Language of Thought model, in *not* featuring a semantically transparent matching up of distinct patterns of coding with the language-style symbols and structures that we might consciously use to conceptualize a task. In this sense, distributed representation is sometimes referred to as 'sub-symbolic'. As Clark describes it:

> The activation of a given unit (in a given context) thus signals a semantic fact: but it may be a fact that defies easy description using the words and phrases of daily language.
>
> 2001, p. 67

### 6.1.3 Features of Connectionist Systems

It will now be useful to examine in more detail some of the interesting features of connectionist systems; features that lend some further weight to the supposed 'biological plausibility' of connectionism; and which are also of interest in terms of some of the key features of notionally creative biological processes mentioned earlier[6], that is:

- Multiple causal factors
- Self-organising processes, yielding emergent properties
- Unique patterns of phenotypic instantiation within certain systemic, genotypic constraints

Such features are of interest also because, at first glance, the abilities of a trained system like NETtalk might be seen merely as outcomes contrived by the system designer in setting the parameters of system architecture, input and output coding styles, learning algorithms, and so on. The following discussion is intended also to fill in the notion of connectionist systems as genuine (although still comparatively very simple) learning systems.

To begin with, we may note that the actual patterns of coding in a distributed representation system of any complexity – where that system has learnt to do things within a particular task domain – are themselves extremely complex. The system is demonstrably representing certain information, and information transformations, in its task domain; but it is apparent that the actual

[6] Note, however, that there is no intention to argue as such that artificial connectionist systems are creative in the same way or to anything like the same extent as the 'human system' and Mind is considered to be. It is only that particular features of these systems are seen to be significant to a developing picture of the brain as part of a creative, developmental, biological process *and* as a representing device.

coding patterns of these multiple and overlapping representations have not been directly 'programmed' into the system (by a human programmer). Instead, they have come out of an iterative process between the designed system, the corpus of training data, and the application of the learning algorithms. Indeed, generally speaking it would probably not be possible to prefigure a specific setting of unit activation levels and connection weights to achieve success as a representational array in an untried task domain of any complexity. As with the notion of non-semantic transparency, in this sense at least the connection between system design/programming and behavioural success is non-obvious; and the iterative learning process undergone by the system seems necessary and non-trivial. The 'final' pattern of representation would not seem to be prefigured in any of the elements of the process in themselves.

Secondly, we may note again that, in a 'trained' system, the representation of a specific item of information occurs as a distributed activation pattern within a connectionist system. The *realisation* of particular activation patterns continues to be linked to the presence of particular corresponding inputs. When a trained system is 'at rest' one might say that it's representational content is still 'there', but not as a number of discretely coded items. Specific representational patterns only 'come out' of the overall disposition of the system, as it were, when the system itself is engaged on carrying out the relevant tasks, prompted by the presence of specific, external inputs.

Again, some of the significance of this can best be appreciated in contrast to the classical model, and its 'take' on success within task domain. Without going into detail, it is commonly recognised that a classical approach in artificial systems will rely on the (programmed) building up of a highly detailed 'inner symbolic model' of the task domain, and that task success then depends on a detailed search of this world-model to marry inputs with appropriate outputs. (Clark 2001, Ch. 2)

The learning character of connectionist systems, and the links between the representational process and the *presence* of the external item (as input), leads thinkers in the field (e.g. Clarke 2001, Elman et al 1998) to consider instead the extent to which intelligent systems may actively make use of information *present in the environment* in achieving their behavioural success. Here the on-going dynamics of system-environment interactions begin to come more into the foreground, and the classical division of the causal economy of cognition comes into question (being a view of the external 'world' only as an initial, passive source of input, subsequently interpreted within the detailed inner model, and then again as the passive arena for a behavioural response). We will return to this matter in more detail in later sections.

This notion of active use of information present in the environment extends also into other interesting features of advanced connectionist systems, where the organisation of information and the behavioural capacities of systems both seem to extend beyond the 'confines' of their training information. Of particular note here is the capacity of some systems to generalise, in generating conceptual categories, and/or applying task domain skills to instances outside their training corpus.

For example, a PDP network developed by Elman (1991) was trained to predict the next word in English sentences. The training data consisted on thousands of sentences, with words input to the system one at a time, the system making a prediction of the next word, and back-propagation used to adjust connection weights after each prediction.

As Elman et al. explain:

> The network eventually learned to carry out the task, although not precisely as trained. For example, given the sequence of vectors corresponding to the words "the girl ate the …," the network activated all the output units which corresponded to the various words representing edible things, rather than predicting exactly the specific word which followed.
>
> 1998, p. 94

In would seem, then, that the network had generated a category of 'edible', although there was no such information explicit in the training data itself. Indeed, the system displayed capacity for a number of such categorisations, distinguishing between words designating animate and inanimate things (nouns), and (as sub-categories of the animates) between words designating animals and humans.

These categories were apparent both in the output of the system, and in the patterns of activation occurring within its units and connections. In simple terms, the activation patterns are closer to each other within the total 'state space' of the system, than the patterns of concepts not categorised together. (The state space is a conceptual space representing all possible states of the system.) (Clark 1993, pp. 60-67)

This is a useful example of a more general point, that connectionist systems are (sometimes) understood to display emergent properties of organisation and behavioural capacity. In this case the emergent character is most clearly apparent in the development of generalised categories of concepts from a combination of the plasticity of the connectionist network, the training data of sentences, and the back-propagation learning process. Given the way the system and data were structured, it would seem that distributional statistics within the language structures themselves provide the informational cues for the system to group certain inputs together, both syntactically and in a way that make semantic (contextual) sense.

Similarly, the NETtalk system discussed earlier was, after training, able to successfully apply its skills of text to phoneme translation to input words it had not encountered before. Again, it would seem that some more generalised capacity had emerged; the system was doing more than merely reproducing 'rote' routines within the limitations of its training data. Thus such a case would seem to qualify as an instance of creative organisation.

Finally, we may note that representation in connectionist system also leads us to consider the temporal dynamics of the representational processes – that the 'location' of a representational process *in time*, that is in relation to prior states of the system, is also significant in shaping the precise pattern of the process. To continue with the example of Elman's word predicting system, the system has been tested both in terms of how it is 'clustering' concepts (as per the above), and in terms of its information-processing transitions from one state to the next. As one might readily understand, the particular grammatical context in which a word appears in a sentence may make a considerable difference to what might follow. Elman (1991) uses the example of three sentences:

1. The boy broke the window.
2. The rock broke the window.
3. The window broke.

Clearly, the word 'window' in the first two sentences functions as the final word of the clause or sentence; whereas in the third case it could not play this role. Thus the matter of what is able to grammatically follow the word 'window' is quite different in each case.

As the system is able to demonstrate success in distinguishing between such cases, it must be assumed that dispositions of the trained system encode not only the lexical 'knowledge' enabling the formation of categories of concepts, but also grammatical 'knowledge' of this sort, where the relation of information before or after other information is significant. In relation to the example, what is shown is that the disposition of the system at the occurrence of the word 'window' between, say sentence 1 and sentence 3 is different, because it allows different potentialities for what follows.

This is of particular interest here, as an instance of connectionist systems showing a degree of moment-by-moment variability in the precise ways that information is being processed, depending on the previous state of the system. The precise pattern of activation in representing, in this example, 'window' will thus reflect this prior state in some way. It is not, as in the classical conception, a fixed and static syntactic item stored in the system, available to be retrieved and manipulated according to higher level processing rules. As Elman describes it:

> There is no stage of lexical retrieval. There are no representations of words in isolation. The representations of words (the internal states following the input of a word) always reflect the input taken together with the prior state … the representations are not propositional and their information content changes constantly over time … words serve as guideposts which help establish mental states that support (desired) behaviour.
>
> 1991, p. 378

Again, the capacity of this system to demonstrate such context sensitivity and generalise grammatical strategies seems to strongly indicate that it is doing more than merely reproducing its training data. As Clark argues, the distributed character of representation, and the 'non-linear' processing strategies, in most connectionist systems would seem to be necessary to the generation of these 'extended' and emergent abilities. (Clark 1993, pp. 59-60)
(Discussion of the word predicting system also draws upon Clark 2001, pp. 70-72; and Clark 1993, pp. 60-67.)

In the terms of previous discussion, this example would seem to show that a connectionist system can display something analogous to the genotype/phenotype distinction insofar as the general architecture of the system itself or patterns of (trained) connection weights can be taken to form certain stable constraints on system abilities and performance, while the dynamics of representation in the moment-by-moment patterns of activation begins to reflect a specific, 'individual' history in a (highly simplified) environment.

For the wider purposes of this study, connectionism (and the key notion of distributed representation) offers another perspective on what it might mean for a physical system to represent aspects of the world, how those representations might come to be 'in' the system, and how they contribute to the behavioural capacities of the system. Although connectionism is of considerable interest in terms of what artificial systems – computers – might do for us; it is also very much shaped as a methodology by a desire to better understand biological intelligence, and the human Mind in particular. Within cognitive science, connectionism provides an alternative to the classical computational view of Mind; the brain as a digital computational device, and representation as the encoding of a formal language-like collection of symbol-tokens and processing rules.

What, then, are the key elements of the view of representation (within a computational structure) that connectionism offers; those things which can now inform our thinking about representation in the biological Mind? Connectionism shows that representation can occur in ways that do not necessarily match our adult, language-based 'break down' of a task domain into basic symbols and rules. It shows that robust, behaviourally potent, complex representations can *come about* within a plastic, learning structure exposed to information in a particular task domain. This organisation of the computational space that develops over time, is understood to *emerge* from a non-linear, self-organising, iterative exchange between a system architecture, a learning process and the environment. Representation in a 'educated' system is able to remain dynamic, constantly being 're-cast' through a flexible combination of the prior states of the system and present-time inputs. Out of these characteristics, connectionist systems show a capacity, not only to achieve, as it were, first-level representations within a task domain, but to go on to self-generate higher level categories (such as the categorisation demonstrated by the word-predicting system discussed above) (Elman 1991).

It would seem, therefore, that the basic elements of creative organisation as proposed above are displayed by artificial connectionist systems in a way that is at least analogous to its occurrence in biological systems. The artificial systems too are functioning within a range of constraints, including the architecture of the system, the settings of input and output modes, and the nature of the training data. Within those constraints, the iterative combination of multiple causal factors, in advanced systems, begins to come together as an individuated, dynamic representational profile; one, in other words, analogous to the unique phenotypic realisation of a biological system within its particular developmental history. The key elements outlined earlier would all seem to be at least nascently present:

- Multiple causal factors
- Self-organising emergent processes
- Unique patterns of 'phenotypic' instantiation within certain systemic, 'genotypic' constraints

All of these features of artificial connectionist systems provide a way of thinking about how representation might be achieved in the human brain, and how it might come about as part of the development of the human organism.

## 6.2 CONNECTIONIST-TYPE REPRESENTATION AND THE BIOLOGICAL BRIAN

Armed with this range of ideas from connectionist modeling of cognitive systems, we now return to the consideration of the human organism and Mind, through a perspective on the process of development, considered on various physical and temporal scales. The purpose now is to pick up the thread of previous discussion about human morphogenesis, still through the lens of the creativity model as proposed. In terms of ontogenetic history and the scale of the processes under view, the discussion will now be focusing mainly on the representational capacities of the formed post-natal organism. We may note therefore that the relevant environmental factors at this whole-of-system level are now generally external to the organism itself and, of course, might begin to include the influence of other people.

Continuing the themes of the previous discussion, the point here is to argue that representation in the human system occurs as a part of the organism's developmental trajectory, and as with processes charted at other 'lower' scales, within a creative (organisation-building) 'tension' between various sources of constraint (including genetic constraints) and the variabilities associated with a unique pathway of phenotypic realisation.

In terms of the on-going reflection about timescales, this section focuses on processes within the early post-natal period, where we might say the morphogenesis has put the major structural elements in place (over the timescale of gestation for the most part), most notably the brain, and attention now goes to shorter-time scale processes occurring within that 'established' structure. However, we are not yet at the point of directly considering the nature of moment-by-moment realisation of cognition. The processes considered here still mainly refer to developmental changes in brain morphology as part of the early 'take-off' phase of the representing brain.

### 6.2.1 Representation in the Computational Brain: Some Basic Questions

In order to conduct this part of the discussion, and taking a cue from Elman et al (Ch. 5), it will be useful to firstly present some basic outlines in terms of: *where* computational (and connectionist) representation is regarded as occurring in the biological brain; and *what* form representation is understood to take there. From that point we will then be able to consider what is, for the main purposes here, the more pertinent question of *how* representation comes to be 'in' the biological brain.

*Where is it?*

So then, where (if anywhere) and at what scale may representation be regarded as occurring in the brain, from a broadly computational perspective? In the most general terms, all cognitive scientists would probably agree that the capacity of the brain to function as a computational device lies in the massive interconnectivity of its internal components and structures, and its links to both the central nervous system and to various sensory organs. In terms of components, we are of course talking about specialized neural cells of various kinds. In the formed brain these cells are the basic parts of various micro and macro structures; and at all scales, from the 'local' relations between cells, to small ensembles of cells, and through various levels of structure, the parts are functionally related via neural connections. All of this and more, the whole variety of anatomical features of the brain, is known to us through medical science.

Focusing only on the brain for the moment: if we take the view, as cognitive science generally does, that the cognitive capacities of the brain arise by virtue of its workings as a physical device; then we may assume, as a working hypothesis, that representation occurs (somehow) *in* the physical structures of this profoundly interconnected biological structure, conceived of as a multitude of *states* and *events*. The computational view, we might say, is about *the nature of computational operations* that are (or happen to be) instantiated by states, events and changes in the physical structure; at least as much, if not more than, the anatomical structure in and of itself.

This point is a significant one for cognitive science generally. In order to clarify it I will refer to the anatomical structures as such as brain *architecture*, as a means of distinguishing them from computation (and representation). We may note that there is no dualism involved here. At any moment there isn't anything physically present except the architecture. The point is to understand that, for the living, working normal brain (in situ), we may use 'architecture' and 'computation' (or 'representation') to describe distinctly different aspects of the same unified whole.

Both the classical and connectionist perspectives on the role of the brain in cognition then propose more detailed ideas that fill this basic computational perspective, and make claims about what it means for the physical device (the brain), as a functioning part of a whole person, to instantiate cognition and other aspects of Mind. Ideas about representation are central to these claims, and while the two approaches obviously diverge in many ways (as do many 'sub-schools' within them), they share a general view that human, 'biological' representation is *in* the specific patterns of states and events instantiated in the cells, structures and neural connections of the brain. The various theoretical and artificial models of computation (as already discussed) are, in essence, attempts to investigate and describe those patterns of activity.

To some extent, the question of scale (or of, 'where exactly?') will depend on the generality or specificity of representational content or process that one is discussing. Much of the literature naturally pays particular attention to the role of the fundamental building blocks of the brain, specialized neural cells, and the ways these connect to other cells. Inevitably and correctly, if one considers the physical make up of the brain, then one comes 'down to' these basic structures. It is to be generally understood, therefore, that representation occurs most concretely in states and events taking place in and amongst neural cells, axon and synapses. In terms of a 'basic unit' of representation, although there are extant debates about the computational capacity of individual neurons, this is more commonly taken to occur at the level of relatively small, highly interconnected ensembles of cells, rather than individual cells as such. (E.g. Edelman Ch.2)

At a larger scale, the folded outer layers of the brain, the structures of the *cortex*, are widely considered to be crucial in most of the higher functions of cognition and their associated representational processes; e.g. 'processing' information from the sensory organs.

*What is it?*

Without going over previous ground, the general view of the classical approach to the computational brain is to assert that representation is encoded into the dispositions of cells and cell connections in a way at least directly comparable to the way that information is encoded into a digital computer.

In connectionist terms, the basic elements of representation are understood to occur in a way again comparable to artificial connectionist systems; that is; a) encoded in the distribution of relative strengths of synaptic connections between cells and within cell ensembles, and b) realised through the patterns of activity taking place within those structures. In other words, they are seen to lie in the neural factors most closely corresponding to, firstly, the distribution of *connection weights* in artificial systems (Elman et al. 1998, Ch. 5), and secondly to the *activation patterns* that realise and specify the representational potential encoded in the weights. At higher scales, representational processes are also interpreted in processes occurring within larger brain structures, and (importantly) in linked activities *between* brain structures also, in a way that suggests overlapping, influence-exchanging activities between *multiple* connectionist-type, representational structures (or sub-systems).

A more detailed picture of this multi-layered conception will emerge from the following discussion, where the third question, the question of *how* representation come to be 'in' the brain will be considered.

### 6.2.2 How Representation Comes About Within Development: Some Background Issues

Discussion here inevitably touches upon a number of nature/nurture-type debates about the degree to which representation might come about in the brain under the influence of specifically genetic factors – and so be regarded as *innate* – or, conversely, come about via learning processes within an individual, developmental history.

Bearing in mind previous discussion here, it is important to recognise firstly that these debates are best not regarded as a choice between 'all nature' or 'all nurture'. One way or another, all substantial views would likely recognise some kind of spectrum with three main categories:

1. Outcomes seen to be subject to strong and fairly direct genetic control;
2. Outcomes that may still be seen as innate, but in an 'extended sense' – where developmental processes are seen to occur through certain genetically shaped 'pre-conditions' or dispositions that subsequently are triggered, and/or develop a mature form, with environmental factors often implicated in the 'triggering' or maturational processes. Often the environmental factors considered here will be those that are commonly present in early stages of individual development.
3. Outcomes that are seen to be more the result of specific environmental influences on individual development and/or learning.

(Elman et al 1998)

The substantial debates therefore are more about the weighting that one might give to one or other of these broad categories, in terms of the degree to which representation is seen to be innate (or not). For example, there are those who would argue that quite specific representational outcomes are indeed subject to quite direct genetic control (shading between the first and second categories above). (E.g. Carey and Spelke 1994; Crain 1991)

We may note also that these questions are significant to the two main models of computation within cognitive science. As Churchland has argued, one of the issues for the classical view of cognition is the question of how exactly a linguaform type of representation is seen to come about through the process of ontogenetic development. Churchland notes that one response to this question is to argue that 'the cognitive activity of animals and infants is linguaform in its elements, structures, and processing right from birth.' (Churchland 1999, p. 132 [footnotes])

These debates are noted here for three reasons only: one, as background information, against which one may better portray the relevance of a connectionist-type view of human representation; two, because they will be significant later, in relation to the social interpretation of views about

Mind; and thirdly, because they pertain to the issue of *variability* as a characteristic (or not) of ontogeny. If representation is seen to be substantially determined by the genome then on the face of it this would seem to suggest that any variability intrinsic to epigenetic development at 'lower' levels is merely 'noise', and only marginally relevant, in relation to representational outcomes in the 'fully formed' individual. Again, I note that my approach to this issue here is to set out the general case *for* a connectionist-inspired account of representation, and assert that variability is indeed a necessary factor within that account.

### 6.2.3 A Connectionist-based Account of Human Representation

Perhaps by contrast with classical approaches, it is very much part of the claims of 'biological plausibility' for connectionism that it *does* offer some cogent ideas about how complex representations can come about through the ontogenic development of the individual organism. This doesn't need to imply that explicitly genotypic factors are not relevant (an 'all nurture' type of position); and in fact proponents of connectionism cited in this study are often at pains to rule out such a view.

The general attitude of connectionist-based views about human representation is that it is does mostly come about within an individual ontogeny, and that: it is emergent, it is shaped by multiple causal factors, including a range of genetic factors, and the process displays the same characteristic tension between genotypic constraints – and other sources of constraint – and various factors of variability. These elements come together in a creative realisation of representational organisation in multiple levels, and a unique pathway of representational development and activity.

### 6.2.4 Genotypic Constraints

The idea of a 'constraint' in this context – one of biological systems – refers to a source of stability and consistency that is typical to the ontogeny of members of a species. Constraints act as limiting control parameters within which a particular individual ontogeny can be played out. As we have already considered, the genotype of a species is a potent and necessary source of constraint on the phenotype. The genome passes on, as it were, fundamental structures of adaptive success; the cumulative results of evolutionary development.

However, it is not the only source of constraint. Above it was mentioned that the relationship between representing system and environment, as part of the causal picture, would increasingly come into consideration; and indeed it will later be suggested that the environment also offers sources of constraint within ontogenetic development.

For the moment we will consider constraints that are internal to the organism, and specifically to the brain. Generally these are appropriately regarded as innate in at least the extended sense above, in that they are broadly species-typical, and thus likely subjects of a significant degree of genetic determination (albeit *via* the epigenetic processes of morphogenesis discussed earlier). Elman et al mention a number of likely candidates (1998, p. 27-35), and for the most part these are to do with specifics of neuron type or response characteristics, and larger scale aspects of brain *architecture*. For example, the types of cells which develop within particular areas, the number and density of layers of brain tissue with larger brain structures, and at the global level, the characteristic macro features of brain structure and connectivity. In general terms, we could assume with some confidence that such consistent features of brain anatomy are innate, and that they are highly significant in (at least) 'setting the parameters' for the individual person's realisation of representation, and mentality generally. In other words, these reliable features of biological organisation are (at least) a necessary precondition for any human realisation of complex representation at all; and also, in general terms, for the relatively large representational capacity we obviously display (by comparison, say, with frogs). For the moment, then, let us accept this range of features as sources of constraint that enter into the typical human representational 'profile' and capacity in at least these general ways.

### 6.2.5 Sources of Variability

If such features of brain physiology as mentioned above are taken to be innate, insofar as they are reliable structural features that emerge through morphogenesis; and if they are taken to enter into the representing world of post-natal life; then where do we identify sources of variability within those structures, and how might these factors 'enter into' representation?

Some of the relevant factors have already been discussed above. The purpose here is not to go over the same ground, but to group those with others not yet mentioned, in this critical part of the argument for Mind as creative organisation. In general terms, the first point here is to recognise that *within* these features of structural constraint, there are indeed a number of ways in which variability in some sense can and does occur. Secondly, it is to say how these factors may be regarded (within a connectionist view) as central to a process whereby representation comes about through and within the specific ontogenetic history of the individual. Together, these factors suggest that variability is as necessary to representational success as are the constraint factors mentioned above.

Firstly, then, the connectionist view of representation (together with other related views) consistently points to the fine-grained scale of neuronal function and connectivity as the core area, in biological terms, where the instantiation of computational representation may be understood to occur. At that scale, and within the constraints of large structural consistencies, a number of sources of variability present themselves. We have already noted, as one key factor, the plasticity of neuronal function in the developing brain, such that neurons from one area may be transplanted to another area, and are able to respond with functional changes appropriate to their new location[7]. (Schlagger and O'Leary 1991).

Alongside that we must consider the massive development in synaptic and axonal connectivity that occur within and between localized areas of the brain in the post-natal period. There is evidence offered by neuroscience that virtually all of the neural cells one will possess are present at birth, together making up the gross structural elements of brain anatomy (E.g. Rodier, 1980 cited in Elman et al. p 288). In several crucial phases within the first several years of life, however, the brain undergoes considerable structural change in its patterns of *connectivity*, in terms of both 'progressive events' (adding new connective structures or strengthening existing ones) and 'regressive events' (subtraction or destruction of structure). (Elman et al. 1998 p. 288-298) There are now numbers of studies which link the timing of such neuroanatomical events, particularly key phases of 'progressive' growth in connectivity, with key phases of childhood development. (E.g. Bates et al 1992)

The connectionist case here argues that significant degrees of variability enter into the specific patterns of connectivity that arise through these events, because they have a flexible, self-organising character, and because they are subject to multiple influences as they occur. It also argues that, in a way analogous to the learning processes displayed by artificial connectionist models, that fundamental elements of the system's representational profile are being 'built' within these connective structures (as forms of emergent organisation), through a nexus of multiple causal factors.

In basic terms, the way this may be understood to occur is through processes analogous to the learning of artificial connectionist systems; where, through an iterative learning process (with

[7] I argued that, within morphogenesis, the capacity of neural cells to respond to their locale introduces a variability at that level, in the way neural configurations are instantiated in each case.

multiple elements), the configuration of connection weights changes over time (some connections are strengthened, some are weakened) in generating a systemic capacity for robust representation.

In terms of biological brains, then, it can be argued that a similar (but far more complex) learning process takes place in the (progressive) development and strengthening of some neural cells, cell ensembles and connections, and the (regressive) weakening or elimination of others. Several writers portray this a competitive process where, as a young individual explores the world, the flow of input effectively is acting as a *selective* influence, assisting to launch and continuing to reinforce areas of developing representational success and 'starving' areas of lesser success. Elman et al describe the process as one where:

> [The] overabundance of cells axons and synapses .. can be viewed as the raw material for development through competition – including competition from the cell next door, competition between regions, and competition among alternative solutions to the problems the child encounters in her physical and social world. Subtractive events (particularly synaptic degeneration) can be construed as the processes by which experience "sculpts" the raw material into its final configuration. Thus human brain development may represent a particularly strong example of the intricate codependance of maturation and experience.
>
> 1998, p. 296-297

(Also Edelman 1992, pp. 83-85)

Clearly the authors are here extending on the wider themes of their work, which portray the instantiation of representation in each phenotype as an *emergent* and somewhat 'idiosyncratic' result of multiple causal factors including: experiential input, the plastic and epigenetic facets of neurogenesis, the genotypic constraints of brain architecture, cross-influences between brain structures, and others.

Neuroscientist, Gerald Edelman, is noteworthy in taking the themes of neurogenetic competition and selection further in his Theory of Neuronal Group Selection (TNGS). (1992) Dynamicists, Thelan and Smith, also lend support to these views, and make note of what they regard as the vital role of variability in neurogenesis and representation.

> [Local] fluctuations and the resulting sensitivity to initial conditions mean that the exact nature of the resulting local diversity [of neural patterns] cannot be predicted. This "noise," however, is also the ultimate source of new and emergent patterns because it provides the variability and flexibility necessary to generate new combinations upon which selection can work. … The

result of early neurogenesis, therefore, is both a modal and species-constant neural architecture and, at a more microscopic level, a system of enormous three-dimensional variability in the number and connectivity of neurons.

1994, p 158

In relation to claims about cognitive representation being subject to significant levels of direct genetic specification, Elman et al sum up their position as follows:

> In principle it should be possible to inherit specific and detailed patterns of cortical connectivity … [However] the evidence that we have reviewed here suggests that the cerebral cortex is not built this way. There is some evidence for primitive innate representations in the midbrain …, but the cortex appears to be an organ of plasticity, a *self-organising and experience-sensitive network of representations that emerge progressively* across the course of development.
>
> 1998 p. 315, emphasis added

In other words, while they recognise several aspects of brain architecture as innate, on the whole they regard representation as not so.

### 6.3 Summary

On the basis of such arguments (some specifically put forward as connectionist arguments) it is concluded here that representation – for the moment particularly in the context of early childhood development – extends upon the creative character of earlier morphogenetic development. All the key features of the creativity model would seem to be indicated:

- The multiple causal factors
- The unique pathway of phenotypic instantiation within stable constraints, including genotypic constraints; the combination of organisational consistency/stability and variability
- Self-organising processes with emergent results

The 'fulfillment' of the creative model as proposed lies in the generation of new levels of organisation, and here it is the 'successful' representational forms themselves, enacted within growing levels of behavioural competency, that are the case in point.

Clearly to fully substantiate the above portrayal of representation would require more than what is put forward here. However, the point of the exercise for my purposes in not to do that as such, but to survey a line of argument, supported by connectionist thinking and modeling, which is credible and fits with the terms of the creativity model.

The particular significance of representation, and the connectionist perspective, is that in an approach to Mind seeking to combine developmental biology with cognitive science, it is a 'building block' that bridges the 'gap' between morphogenesis and cognition. Cognitive intelligence, by the computational story, *is* an intelligence founded on a flexible systemic capacity for representation. Connectionism offers some key explanatory strategies for how that capacity is generated.

As such, the above argument about the ontogenetic development of representation as a creative process is the initial piece in the culminating point of this study; a description of Mind as creative organisation. There remain, however, a few further steps. These are required to move our notion of representation (and related matters) beyond the primary 'take-off' phase of brain development and into a basic description of Mind *as the creative moment-by-moment realisation of 'mature' cognition*. In the context of this work, that will be a summary point. In the next chapter we will also, to some extent, be extending beyond the specifically connectionist-based perspectives developed above.

# "Mind as Creative Organisation"

# Chapter 7: Embodied, Embedded, Creative Cognition

## 7.1 Moment-by-Moment Cognition

At each moment of our waking lives we are typically cognizant of a whole range of salient objects and events in our surrounding environment. We are constantly engaged in processes of recognition, negotiation through physical space, manipulation of objects and exchanges with other people. Often the relevant intervals between events in the environment – a dog approaches – and cognitive responses to those events – the processes engaged in the thought, 'There's my dog, Benji.' – are relatively short. In terms of our 'tracking' of Mind over progressively shorter timescales of realisation, we are now considering the smallest scale, and processes that occur over seconds and parts of seconds.

The overall view taken here is that cognition in this sense is appropriately understood as both *embodied* and *embedded*. The former term encapsulates a view of human cognition as inextricably 'tied in' with the nature and functioning of the developed, biological system; body, brain architecture, bodily movement and so on. The latter term asserts that the nature of cognition is also deeply engaged with the system-external environment – in short, the concrete settings in which our development occurs and our lives are conducted.

Again, connectionist models of cognitive processes and the nature of representation continue to offer ideas about what is going on in this more immediate sense. The reader will recall in section 6.1.2 I noted the important distinction between connectionist-modeled representation as connection weights – task competencies encoded in the relative strengths of connections between units – and as activation patterns – the specific patterns of activity occurring in a PDP network in response to the presence of particular stimuli. In the previous chapter I suggested that the establishment of certain stable patterns of neural connectivity during post-natal development might be considered a process analogous to the 'setting' of connection weights via the iterative learning processes of PDP systems. Now the idea of representation as activation pattern comes more to the fore, insofar as it is, in connectionist terms, the characteristic short-term response to present environmental stimuli.

We must also recall, however, that both aspects of representation figure in the real-time realisation of task competencies by connectionist systems. This is important because, in considering moment-by-moment cognition, and the example used above, the recognition of the dog would also seem on the face of it to imply both some relatively stable representational capacities (sufficient, say, for reliable discrimination of all dogs from all other environmental objects), and the more immediate recognition of this dog, Benji, at this specific time.

### 7.1.1 Recurrent PDP Networks

One of the important distinctions to note with connectionist models is that between *feed forward* networks and *recurrent networks*. Previously I discussed some of the basic architectural features of a simple connectionist system; consisting of a layer of input units, a layer or layers of hidden units and a layer of output units. In basic terms, with a feed forward network the flow of signaling prompted by a particular excitation of the input units proceeds, via the connections between layers, in a unidirectional manner: from the input layer, through the hidden layer/s and to the output layer. The NETtalk system previously discussed is a feed forward system.

With a recurrent network, connections and the signals they carry are set so as to provide a mechanism of feedback so that, we might say, signals circulate within the system. The Elman word-predicting model cited (1990), for example, is a recurrent network where, in addition to the basic features above, it also employs a set of 'context units' that do not receive external input, but receive signals from and send signals back to the hidden unit layer. In this way 'at time 2, the hidden layer processes both the input of time 2 and, from the context layer, the results of its processing at time 1. And so on recursively.' In other words, it achieves a certain form of dynamic short-term memory, allowing it to 'capture the sequential nature of the input.' (Elman et al 1998, p. 120)

In considering connectionist systems as models of moment-by-moment cognition it is generally the recurrent networks that are of greater interest here, because they display a combination of characteristics that seems usefully analogous, in simplified form, to aspects of human cognition in action. A number of relevant matters have already been raised in sub-section 5.4, but it will be useful within this discussion to reiterate to some extent, and identify some further points.

Firstly, then, it would seem intuitively clear that real-time cognitive processes (for brevity, for the moment, 'cognition') are fast; the interval between the dog coming into my field of view, and the corresponding thought of 'There's my dog, Benji' is only a matter of parts of a second. To be convincing, a model of human cognition must display at least a comparable capacity to deal quickly with events.

When a trained recurrent network is presented with a particular input the system will go through a process where the input excites signaling from the input layer into the system, and thus throughout the system until a particular optimal state is reached, one which satisfies the constraints exercised by the nature of the input, along with the inhibitory or reinforcing influences of the connection weights and perhaps other factors as well. This stable state is in effect, the system's short-term (typically task-oriented) response to that particular input. This process is known as *relaxation* or *relaxation search* (Rumelhart et al 1986, pp 8-10). The pattern of positive excitation emerging via relaxation search is what we have discussed as the (representing) activation pattern.

The immediate point about the process of relaxation search is that it is the characteristic *short-term* response of a recurrent system. The actual time taken for the process of relaxation search will depend on the nature of the system hardware; but the point is that it occurs over a relatively short time, typically in the order of seconds or parts of seconds. Given the demonstrated representational success of trained systems, such a capacity for generating short-term representation would seem reasonably convincing, in at least this respect, as a simplified model of human cognitive responsiveness to immediate events.

Another related point is that this process of relaxation search, and its results, is an instance of self-organisation, where multiple factors enter into the process and the activation pattern in an emergent result. Another way to put the above point would be to say that one of the strong points of recurrent networks – in terms of being biologically plausible – is their capacity to self-organize over very short timescales. We may now move on to discuss in more detail what some of the 'multiple factors' impinging on this self-organising process are and, again, how they relate to information or ideas about human cognition.

### 7.1.2 'Feedback' Connections and Context Sensitivity

As mentioned above, recurrent PDP networks feature mechanisms to generate feedback between units, layers of units or multiple structures. As Clark describes, in reference to the Elman system previously discussed, what these mechanisms allow is a certain form of short-term memory, and thus the context sensitivity displayed by recurrent networks:

> [What] is preserved is some kind of on-going trace of the network's last activity. Such traces act as a kind of short-term memory enabling the network to generate new responses that depend both on the current input and on the previous activity of the network...
> {Thus we can] see beyond the classical image of static symbols that persist as stored syntactic items ...Instead, we confront an image of a fluid inner economy in which representations are constructed on the spot and in light of the prevailing context and in which much of the information-processing power resides in the way current states constrain the future temporal unfolding of the system.
>
> 2001, p. 70

One might observe again a certain on-the-face-of it plausibility in this, insofar as context sensitivity would appear to be a common feature of our day-to-day mental life. Consider, for example, a situation where I encounter a group of dogs together, and the short-term representational response is shaped by a particular over-arching goal. If for example, my goal was to locate my dog, Benji, amongst the group, my moment-by-moment representational response might be focused on picking out Benji's particular white and tan coloring. On the other hand, if I had the goal of counting all the big, black dogs in the group, the representational response would be different.
On a wider level, it is this feature amongst others that has led me to regard recurrent networks as (simple) instances of creative representational organisation; by virtue of this context sensitivity, as Clark suggest, each instance of representation is 'made up on the spot' – a unique concatenation of the multiple factors, including the prior activation state of the system, that enter into that instance.

This idea also relates suggestively to some studies of biological brains in action. For example, Thelan and Smith describe the studies of Freeman (e.g. Freeman 1991) on brain activity in rabbits, in response to olfactory stimuli. Electroencephalogram (EEG) measurements were recorded simultaneously from over 60 sites across the surface of the olfactory bulb. Responses to specific scents were shown to manifest not as excitation in localized specific (representing) areas of tissue, but as dynamic, oscillating patterns of neural excitation across relatively large areas.

The representation of individual scents, therefore, was shown to occur 'not in any single group of neurons, nor even in the shape of the EEG waves, but in the spatial pattern of the amplitude of waves across the entire olfactory bulb.' (Thelan and Smith 1994, p. 132)

We should note that this study is looking at patterns of neural activity across a far larger scale of connectivity than that modeled in PDP systems. However, the pertinent point here is that the response patterns for particular smells were shown to have altered, *depending on the prior (contextual) presence or absence of other smells.* In other words, the neural response to a particular stimulus was not one stable pattern of excitation, recurring the same way each time the stimulus was present, but a dynamic, 'new' pattern constructed out of multiple factors, including the presence of an external stimuli, and the prior state of the system.

Such evidence would seem to add to the picture of moment-by-moment representational organisation as self-organising and emergent, and thus of human cognition as the unfolding of a unique 'pathway' of representation (and action) shaped by multiple causes. (The example is also most useful in informing the idea of *embodied* cognition. It is suggesting that, similarly to the features of recurrent connectionist systems, real-time cognition is taking account of prior states across whole systems, and between sub-systems; and in this sense is very much 'bound in' to the specific architecture and on-going dynamics of the developed biological brain.)

### 7.1.3 'Feedback' and Categorisation

As a further look at relevance of recurrent connectivity, and the notion of embodiment, we may also consider that huge amounts of 'feedback' neural connections within and between structures are a key feature of human brain architecture. Gerald Edelman takes up this point within his Theory of Neuronal Group Selection (TNGS) theoretical framework (as discussed in sub-section 6.5), proposing interaction between localized structures as a potential source of higher-level representational organisation.

Specifically, Edelman portrays a capacity to perceptually *categorise* - that is, to engage sensory input in discriminating objects in the world from other objects – as primary to, and more basic than language-based categorisation. To do this he firstly describes localized neuronal repertoires as *mapping* (representing) input from certain sensory modalities, by the selectional process outlined in the previous section; ie by the strengthening of some neural connections and the weakening of others. He cites, for example, the visual system of the monkey, as having 'over thirty different maps, each with a certain degree of functional separation (for orientation, color, movement, and so forth)' (1992, p. 85).

Categorisation is then seen to be achieved through essentially the same selectional process, but at a higher level, that is in the connectivity between maps. In other words, environmental factors (as 'input') act as a selectional mechanism within development that strengthens some of these connections, and weakens others.

Edelman then suggests that specific patterns of 'reentrant' feedback between multiple maps and linkages with other areas of the brain (simultaneously active in motor and sensory modes), triggered by an animal's movements within an environment, participate in the realisation of moment-by-moment categorisation:

> [A] global mapping ensures the creation of a dynamic loop that continually matches an animal's gestures and posture to an independent sampling of several kinds of sensory signals. Selection of neuronal groups within the local maps of a global mapping then results in particular categorical responses.
>
> 1992, p. 89

Edelman cites a range of empirical evidence from studies with animals and with artificial systems in support of these claims.

The key points of interest here lie not in the detailed verifiability of this theory as such (although I find it convincing in many respects), but as an illustrative example of what the notion of *embodied* cognition can mean. We may note in particular how, within this view, the real-time realisation of a cognitive capacity is bound into constant processes of interaction, both with the environment, and between multiple areas of brain function at both micro and macro scale. I would particularly emphasize the proposed role of bodily *movement*, and that of an *iterative feedback between brain areas* that are representing in different ways, e.g. sensory and motor-function ways, as (notionally) two key influences entering into cognition.

Equally, the argument is suggestive of the way cognition is also *embedded* in the system-external environment in which it occurs. Bodily movement is shaped by factors as basic as gravity; sensory activity is continually shaped by the properties of the natural world.

### 7.1.4 Embedded Cognition

The notion of feedback, interpreted more generally, offers a picture of human cognition as very much a dynamic process, and the patterns of activity in the cognizing brain as constantly evolving. It is also a concept which, in some senses, can even be extended into the arena of relations between system and environment.

As Rumelhart et al (1986) discuss, one of the characteristic features of PDP systems is that their dynamic activities are essentially responsive to (or dependant on) input, and indeed changing input over time. Although, they suggest, some would find connectionism less convincing as a model of biological cognition on this basis; they further reflect that our own cognition takes place, in effect, in an environment that is constantly changing – or, perhaps more to the point, where the nature of the system-environment interaction is changing. One source of this change is our own acting upon the world.

> [Our] interpretation of an event often dictates an action which, in turn, changes the environment. The environmental change can then feed back into the system and lead to another interpretation and another action. … this feedback loop can continuously drive the system from state to state.
>
> Rumelhart et al 1986, p. 40

This picture of constant interplay between movement and action in the world and cognition is a constant theme in connectionist-inspired (and other) views of Mind – in terms of development, learning and real-time cognition. It can also be seen as a view that takes us further away from the more classically-inspired view of the causal economy of cognition; where input, cognition and behavioural output appear as more easily separable stages in a linear process, and the emphasis falls on the internal domain holding all the required information – a detailed, symbolically structured inner world model – to be operated on for the production of an appropriate response.

With the idea of cognition as embedded, this neat division of the causal economy falls away and, furthermore, the environment is recognised not simply as a passive arena in which the cognizing system operates, but a source of information, feedback and constraint in itself *for the business of moment-by-moment* cognition. As Clark suggests in discussing the nature of perception, 'Perception and action, in this view, form a deeply interanimated unity.' (2001, Ch. 5)

### 7.1.5 Representational Resources & Levels of Organisation

We should note in relation to the above, then, that the ideas of embodied-ness and embedded-ness in relation to cognition are themselves intertwined. They present a cumulative picture, as it were, of cognition and representation as self-organising on a moment-by-moment basis, shaped by high levels of connectivity, including feedback structures, within and between brain systems; and these systems themselves both influencing and being influenced by body movements and sensory systems; and all going on via a highly interactive engagement with the external environment. In the debates of cognitive science and philosophy there would seem to be a certain level of recognition about this kind of picture as a portrayal of cognition through its formative stages of post-natal development and early learning. In contemplating the exploratory learning of babies there would certainly appear to be some close links between manipulating (holding, hitting, tasting, etc) the world and perceiving it, or learning about it. However, critics raise questions about how this view might account for seemingly 'higher order' cognitive processes such as abstract thought, or conceptual planning of complex tasks.

While the views we have been discussing may seem well suited to explanation of how the biological brain launches into basic, 'ground floor' cognition, do they do less well in suggesting how this might go further, into these 'higher' processes? Clark puts forward a view to suggest that we are able to do this because the biological brain gets us 'far enough along the way' to be able to take advantage of a range of further external, and/or artificial factors; in his terms, 'cognitive technology'. (2001, Ch 8). These factors Clark describes as providing 'scaffolding' to extend the capacities of the evolved brain into these other areas.

In particular he considers the role of language as a tool to increase our powers of categorisation – perhaps building on the kind of primary categorisations capacity described above. In simple terms, the idea here is that language provides symbols (words) that designate a *class* of object: the word 'dog' does not simply describe that one hairy animal in my lounge room, but a whole class of similar animals. Thus language provides a means in which these higher order levels of similarity and difference can be represented in the cognitive domain. As Clark describes it:

> Learning a set of tags and labels (which we do when we learn a language) is, we may speculate, rather closely akin to acquiring a new perceptual modality. For like a perceptual modality, it renders certain features of our world concrete and salient, and allows us to target our thoughts … on a new domain of basic objects. This new domain compresses what were previously complex and unruly sensory patterns into simple objects. These simple objects can then be attended to in ways that quickly reveal further (otherwise hidden) patterns, as in the case of relations between relations.
>
> 2001 p. 145

Of course, the area of language acquisition is a whole vast area of cognitive science and developmental psychology in its own right. Again, for our purposes, the question is what such a view adds to the picture of real-time cognition.

Firstly, I would suggest it offers another perspective on the notion of embeddedness and the interaction between cognitive processes and the environment. Acquired language in the Clark view becomes a set of stable representational tools available to the biological brain in its moment-by-moment work, extending the basic capacities of pattern recognition, and in this sense enabling the rapid exchange of language-based information that plays such a large part in our everyday adult activities.

This notion of stable representational features also harks back to an earlier comment, that of keeping in mind the interaction between representation in the form of connection weights and activation patterns in trained PDP systems. Here we touch on the analogous view, that human cognitive processes, over the longer timeframes of development and learning, take on more stable representational forms (and associated structures of neural connectivity).

Consideration of this dimension of cognition is important to the goals of this study in that, so far as cognition is concerned, it fills in that more general notion of the creativity model; that of complex biological systems generating and sustaining multiple levels of organisation. I would hope that my case for creativity in this sense, and in relation to morphogenesis is already made. However, the question remains as to what, if anything, we might make of the idea of multiple levels of organisation in the domain of representation, and embodied, embedded cognition.

To explore this idea let us recall the view that some stable representational capacities come about through the developmental phase; where brain events at the micro level (particularly through progressive synaptic growth or regressive decay) are seen to be subject to selective pressures from other brain events, or from 'incoming' information from the senses, and so on. As Elman et al discuss such process may be implicated in a normal 'burst in learning activity' between 2-4 years of age, a phase immediately preceded by a massive growth in cortical connectivity. They also suggest, however, that, under the 'competitive' pressures of representation-building, that regressive events also may be crucial, and can be regarded as 'processes by which experience "sculpts" the raw material into its final configuration': and that a gradual thinning of the cortex between 8 years and adulthood may in fact be just these processes at work. (1998, pp. 295-298)

Let us now add to this the idea that processes of learning, on various timescales (remembering a phone number; learning a new language) may engage somewhat similar processes – that is, processes of growth or regression in relatively stable patterns of synaptic connectivity – although perhaps more in terms of 'refining' the gross patterns of connectivity generated through development. Drawing on contemporary neuroscience, there is now a level of recognition within cognitive science that the precise patterns of connectivity in the brain continue to change, develop and regress over lifetimes, and in response to environmental circumstances. (Elman et al. 1998 Ch. 5) In the terms of the view of representation presented here, it is these changes that would be invoked as the likely physical manifestation of learning.

In more general terms, then, perhaps we can say that the instantiation of specific representational patterns is in effect a *limiting* of both the initially high level of plasticity and representational potential, and of some on-going responsive flexibility and variability. If, within the constraints of brain architecture, the brain then builds and sustains a number of stable representational features, then might these – as factors which simultaneously limit (in the sense we are using) and organize – also be considered as *additional sources of constraint,* ones which open up a potential for further representational capacities.

Language might be notionally regarded as a prime example of such a process. And as Clark suggests above, what it seems to involve, beyond the new configurations of 'wetware' involved, is also a new level of *informational* organisation. Such a new level, in other words, exists in the specifically cognitive domain; in the real-time processes of cognition; the ways these concatenate multiple sources of influence and constraint; and in the 'entry' of new forms of constraint (stable representational forms acquired through development or learning) into that equation.

The idea of certain forms of basic representational consistency acting to underpin others also comes into the frame when we consider questions about the seeming *motivation* of cognition, so visible in young children, to actively explore the world of sensation, to explore relationships between objects, and so on. What, we may ask is the source of this *appetite*? One possibility discussed by Thelan and Smith (1994, Ch. 11), following Edleman's ideas about cross mapping and reentry canvassed above, is to look to reentrant 'feedback' and cross stimulation between the actively representing, behaviour shaping, cortical areas of the brain and other structures engaged more directly in the body's physiological and emotional responses. The combination, in the terms of dynamics, creates a stable *motivational attractor*. Perhaps again we may speculate that such stable features, which themselves have come about within the developmental process, may provide sources of constraint (and motive 'force') for further representation. As they describe it, 'certain movements and perceptions form higher order categories, these categories become associated with particular motivational attractors.' (p. 319) And again this also adds some depth to the notion of cognition being, and continuing to be, embodied; in this case tied in with physiological and emotional responses.

## 7.2 Creative Cognition

Taken together, and in light of earlier discussion, the ideas outlined above would seem to support a view of moment-by-moment cognition consistent with the broader theme of creative organisation.

The connectionist 'take' on cognition, broadly considered, suggests that these moment-by-moment processes mirror the dual aspect of representation in artificial connectionist systems; representation as both connection weights and as activation patterns. It suggests, in other words, that moment-by-moment cognition: a) draws on (or engages) stable representational features – features instantiated in relatively stable patterns of neural connectivity arising via either development or learning; and b) is a process where representational features emerge and evolve through self-organising patterns of neural excitation over very short timescales.

In introducing the creativity model in Chapter 5, the genotype was put forward as a fundamental source of constraint and stability during morphogenesis. Genes also continue to play a role in the ways that physiology, including that of the brain, changes over a lifetime. On the whole, however, they are probably less significant in terms of constraints acting on real-time cognition. Here the emphasis is on matters such as the stable representational features as mentioned, the immediately prior states of the cognizing system, bodily sensation and movement, and features present in the external environment.

And again it is this very multiplicity of influences and constraints, now as entering into embodied, embedded, moment-by-moment cognition that – in a way also consistent with the characterisation of morphogenesis and development – is also regarded as the source of a fundamental *variability*. The processes in question, in other words, are seen to be subject to an ad hoc concatenation of these multiple factors in each individual, at each moment; a situation that would seem to rule out reductive explanation in terms of simple, linear causal processes.

If we accept the earlier view offered of variability entering into the establishment of some foundational representational features during development (and – although I've not argued it in any detail – in stable features established through learning), then in this sense variability in real-time cognition also 'builds on' features where, as it were, variability has already entered into the phenotypic 'pathway' and history.

Thus we might not be surprised to learn that the moment-by-moment patterns of representation, even under controlled conditions of identical stimuli, or compared between genetically identical individuals, can be highly individualistic. As Elman et al. conclude in relation to several studies in the area, 'it becomes increasingly evident that there is considerable variation in structure and function [of the cognizing brain] in normal adult subjects.' (1998, p. 298)

In relation to this it should be recognised that there are debates about certain localized brain structures which would seem, in numbers of cases, active in relation to certain quite specific representational 'jobs'; that is, they are notionally both 'domain specific' and typical across cases. If so, then this might be seen as prima facie evidence for innate representation: in other words, that these features and their representational 'job' are typical *because* they are genetically specified.

While there is no intention to enter into these issues here, it is worth noting them as matters of 'live' debate in cognitive science, and a potential source of argument against the notion of variability embodied in a unique phenotypic history. However, we may also note the review of this debate by Elman et al, where they argue that domain specificity is often addressed too simplistically, and in any case can itself be (and the cumulative evidence suggests it is) an *outcome* of epigenetic development, in the same manner (so they argue) as representation generally. (1998 p. 240-250)

In more general terms, then, we would seem to have a view of real-time cognition that fits with the overall model of creative organisation proposed. There is the interplay of features providing stability and consistency through time, and acting as sources of constraint, with a characteristic variability driven by the ad hoc combination of multiple causal influences. Alongside relatively stable features of neural connectivity arising through development and learning, and these as emergent levels of organisation; the *emergent activation patterns of neural excitation seen to instantiate cognition in each moment are also rightly regarded, in my view, as new levels of organisation too.* The activation patterns occurrent in a particular moment of real-time individual cognition, in other words, are an event of organisation; unique to that moment by virtue of being born out of a unique concatenation of the specific external events pertaining and the prior states of the system – the latter bringing to bear anything from sensory information to contextualising goal orientations to stable representation resources of various kinds: indeed, in a certain sense, the whole prior history of the individual.

Thus the basic conditions of a creative biological process as I've defined it are fulfilled.

- The multiple causal factors entering into a unique pathway of phenotypic instantiation
- A range of stable constraints, including genotypic and neural-architectural constraints *and* stable representational features
- Thus the combination of variability and organisational consistency/stability
- Self-organising processes with emergent results
- New levels of organisation

In more descriptive language, there is a basis here for what I would regard as a highly stimulating picture of cognition; not as a mathematical codification of an objective world into formal symbolic structures, but of an always developing and evolving collaboration between body, brain and environment. We are, I would suggest, engaged in a constant, real-time cognitive process consisting in re-creating (representationally) the world about us, and/or otherwise 'manipulating', 'sampling', connecting elements of our representational profile. At each moment cognition is being created at the 'leading edge' of a dynamic yet stable organisational *process*.

In, around and through that 'bass beat' of cognition, we are also, for much of our time, engaged in actively exploring elements of the world around us (using representational tools like stable perception and language), capturing the *relationships* between them; the larger *patterns* of which they are part. When these patterns are captured and consolidated as simplified representational items in themselves, then we are able to use those items for further exploration; in this sense creatively building representation on multiple levels.

The creativity view of Mind portrays the human being as a deeply 'developmental' creature. We exist as the 'product' of a long evolutionary process, where forces of natural selection acting on populations (embedded in environments) have 'developed' homo sapiens as an organism of high material and behavioural complexity. Each new organism embodies this evolved genetic inheritance but also, through morphogenesis and then early childhood stages, enters into a uniquely instantiated process of physiological and cognitive development, shaped by multiple factors. Thirdly, we develop by acquisition, as it were; using and assimilating 'cognitive technologies' such as language.

The significance of describing developmental and cognitive processes as 'creative' is then also particularly vested in the view that they are part of what constitutes a unique phenotypic and cognitive history. The variabilities inherent in the confluence of multiple influencing factors in each individual case, in other words, opens a space for *divergence* between cases; a greater divergence than that which would arise through the influence of genetic differences alone.

In considering the significance of such putative divergence, it should be recognised that the main focus of the arguments in this chapter (apart from the section on connectionist models) has been on processes occurring during ante-natal development and the earlier years of post-natal life. The matter of how the views put forward might then apply (or not) to the complex, language-fluent, sometimes highly specialized nature of mature adult cognition remains a largely open question for this study, and one that will not be addressed here. However, I would say that on the basis of the work undertaken I can see no particular reason to believe that the general nature of mature cognition is in some way fundamentally different from the embodied, embedded view offered here. (Thelan and Smith 1994, Ch. 11)

Nevertheless, with this proviso in mind, the notion of divergence, and the whole designation of cognition as 'creative' organisation, also opens some ground for further interpretation in relation to 'creativity' in the more conventional and social mediated sense of the term. Let us briefly speculate about where this line of thought might go.

Firstly, the view taken here clearly locates Mind at the nexus of an individual phenotype, a specific developmental trajectory, and a cognitive history within particular external circumstances. There is no reason within this view why a particular person shouldn't come to be, as it were, at a point of coincidence between their own natural and developed capacities and some external, historically contingent ideas or influences; and within this coincidence of capacity and circumstance then be 'struck' by one of those moments of original invention or inspiration – the classic case of human creativity. (E.g. Vernon 1970, Ch. 2) There is in this way an allowance for creativity in the sense of historically new ideas coming into being, while suggesting that they will not occur in a vacuum but require, as it were, some confluence of prior factors.

Perhaps of more interest for my own further work is the idea that activities particularly marked out by contemporary society as creative, such as the making of art, might be understood to engage – in a particularly potent way perhaps – both the immediate (embodied, embedded) nature of moment-by-moment cognition, and the longer term process of learning. (In connectionist terms, as discussed, this is the dual character of representation as connection weights and activation patterns.)

Consider for example a skilled artist involved in the process of painting a picture. On the basis of the view of Mind developed here, we might assume initially that she has 'built up' a range of body-brain-world abilities such as manipulating the physical tools (brushes, paint, etc), discriminating subtle differences of color, and so on. We may speculate that, by virtue of practice, these are consolidated to the extent of becoming largely 'automatic'; that is, not requiring a great deal of direct, cognitive 'attention'. What this make possible, I think, is for the more immediate, in-the-moment cognitive process (the immediacy, say, of her seeing the landscape scene in front of her) to more directly enter into the sensorimotor process of applying the paint; without the 'interruptions' and cognitive effort that might be required to work out, say, how to manipulate the brush to get *that* particular textural effect.

Similarly, a skilled pianist, having 'internalised' a set of skills to move his fingers along a keyboard and depress keys according to different patterns and so on, might then have the cognitive 'space' to pay attention to the emotive interpretation of a piece of music; while a new student might only cognitively 'engage' with the mechanical task of getting one's fingers to play *those* two chords consecutively.

In both these examples the creative process may again be understood, in the terms developed here, as an iterative, emergent combination of certain stable constraints – and here I refer in particular to stable representational features in the form of learned skills – and the variabilities inherent in the multi-faceted confluence of internal and external factors entering into moment-by-moment cognition.

My interpretation of both these examples would also, in my view, also fit reasonably well with some of the contemporary cognitive science perspectives focusing more on creativity as it is, perhaps, more 'conventionally' understood: for example, in the work of Margaret Boden (1994, pp. 75-116) and Philip Johnson-Laird (1988, Ch. 14)

Unlike the approach taken here, both these writers consider 'creative' events as a specific class of psychological / cognitive events including things such as invention, artistic creativity, new scientific theories and the like. Both regard the creative 'product' as something that should be more than merely novel to its producer, but must be new in a deeper sense. Boden describes it in terms of a process that works through exploration of or changes in a 'conceptual space' (i.e set of understood rules or conventions within a particular domain), rather than simply a trivially new instance of use of an unchanged set of rules (such as a sentence with novel combination of concepts). (1994, p. 78) Both also note that some form of valuing of the product by other people is also relevant to what is ultimately regarded as 'truly' or importantly creative within a society.

Of particular interest here is the fact that both these writers also discuss creativity in this sense as something that arises within certain pre-conditions (it does not spring out of nothing), and that occurs 'in the moment', as it were, (as a psychological event) but within certain stable constraints, either internal or external – e.g. the rules of a 'conceptual space'. This I would see as broadly consistent with the approach I have taken, and my characterization of the two examples above.

From a wider perspective, however, I would also take issue with the basic notion of limiting creativity to a special class of events in the first place. My own view would be that it is ultimately more accurate (and potentially more useful from a sociological perspective) to understand that all normally endowed, conscious human beings are, in a genuine sense, continuously engaged in the realisation of a creative (cognitive) process. (Dartnall 2002) It is only subsequently that this same essential process – in particular people in particular circumstances – happens to become engaged in those events and products that as a society we tend to mark out as 'specially' creative.

Finally, if cognition is understood to develop and continue to function in a way that is influenced and shaped by environmental conditions, then it would seem that interactions with other people and other socially mediated conditions will necessarily enter the equation in some way. Precisely what might come to be said about the extent or significance of any putative effects, or how they might be weighed against other sources of influence (such as the genotype), is again an open question: but on the basis of the views put here, one could not, in my view, reasonably take the view that socially mediated environmental conditions are likely to have no significant effects at all. Again, then, this might be regarded as 'opening space' for further issues and questions; in general terms a space where a particular view of Mind can also be extended towards a specifically social, or indeed a socio-political interpretation.

The next and final chapter briefly considers some of the further issues in this direction; one that are relevant to the original motivations and interests of this study, as outlined in the first chapter. In more specific terms, given the overall view of Mind developed here, what issues might be raised about the ways contemporary science informs socio-political interpretation, discussion and debate?

## "MIND AS CREATIVE ORGANISATION"

# CHAPTER 8: MIND, SCIENCE & SOCIETY

### 8.1 INTRODUCTION

The conclusions of the previous chapter represent the primary goals of this study, to develop a view of Mind as creative organisation. However, as outlined in the introduction, the underlying motivations of the study are also to consider this view in relation to matters of wider socio-political concern. In Chapter 1, I asserted that 'views/beliefs about human nature' could be (and are) potent forces in shaping our social and political domains, including the ways we respond to serious social problems. I also suggested that modern science has introduced a number of such defining ideas into the public domain, and (sometimes) to profound effect; for example, Darwin's theory of evolution through natural selection, and its implications for our understanding of our place in the world.

In recognition of such social effects, and other forms of exchange between science and society, I identified a value in science being able to function and to present findings in a way that is free from partisan political or ideological influences. I also asserted that, notwithstanding this general value, science (and philosophy of science) should also be open to critical scrutiny and debate in terms of its actual or potential role within wider social processes. (Rose et al. 1984, Ch 1)

The aim here, therefore, is to briefly consider some ways in which a creative-organisation view of Mind might be relevant to issues of the social or political interpretation of science. The discussion will fall into two main parts.

- The first will point to a particular area of extant public discourse where, in my view, some flawed 'ideas about human nature', drawing upon certain 'source' ideas derived from science, underpin certain socio-political views. The aims here will be to: a) locate a creative-organisation view as, potentially, a critical perspective on these ideas, and the ways they are interpreted; and b) to ask what the implications of this discourse might be for other related science-based ideas.

- The second part will briefly outline some potential directions of social interpretation the creative organisation view might offer in its own right.

In general terms, however, it should be understood that the goal of this chapter is *only to offer a 'statement of intent' about potential future areas of work, in light of the view of Mind proposed and given the initial motivations as indicated.* It is not to substantively argue any of the issues raised.

## 8.2 Science-based Approaches to Explanation: Some General Observations

Prior to the first part of the main discussion as outlined above, however, it will be useful to firstly identify two contrasting tendencies I see in the broad approaches science takes to the business of *explanation*, particularly in relation to cognitive science and the nature of Mind.

In Chapter 2, I outlined a schematic contrast between certain *reductionist* tendencies within science and other, let's call them *emergentist*, approaches to understanding complex systems. I suggested that reductionist approaches:

- Focus their attention on the constituent level within whole systems and aim to account for the nature of the whole at that level; in terms of the properties of the parts and/or constituent-based causes of system properties or behaviours.
- Favor narrow, *single-level* explanations, and seek to eliminate 'higher level' explanations in favor of 'lower level' ones.

By contrast, emergentist approaches generally:

- Recognise emergence as a genuine phenomenon
- Recognise multiple causal factors, sometimes acting in non-linear ways
- Accept multiple levels of description and explanation

Clearly, on these terms, it is the emergentist path that has been argued for here, as having a basis in ontology, and as an approach suitable to an understanding of complex biological systems and the nature of human cognition. The more 'radical' reductionist approach has been avoided.

We should also add, however, that the emergentist approach taken here may equally be regarded as avoiding extremes in the 'other direction'; that is, of radical anti-reductionism – for example, in claims that system-level properties of complex systems are somehow immune to even a moderate kind of reductive analysis. Again in chapter 2, I recognised the value in methods that are *reductive* in orientation, seeking an understanding of complex structures through the *analysis* of things and the identification of constituent-based causal factors, without being strongly reductionist (as defined above).

At a more general level then, we might say that an emergentist view can take a 'middle way' approach to analysis and explanation: one avoiding the extremes of either strong reductionism or anti-reductionism; one recognizing the benefits of moderate reductive analysis (as Auyang does with her arguments for 'synthetic analysis' [1998]), and also taking a pluralistic approach; acknowledging that explanation of complex systems can be pitched at many levels, depending on the perspective taken and the goals of the exercise.

More specifically in relation to the business of cognitive science and philosophy of Mind, such an attitude would seem to be adopted by Andy Clark in his conclusion to *Mindware* where he suggests that the human Mind, 'is a system that resists any single *approach* ..., that resists any single *level of analysis*, ... and that resists any single *disciplinary perspective*'. (2001, p. 160) Based on the work undertaken here, I would concur with this view. Thus, in the wider arena of science-based claims about causes of human mental states (or linked behaviours) my general disposition would be, again, to treat with caution any claims that invoke causal factors *at a single level* as acting to strongly shape or determine subsequent outcomes for the developed system *as a whole* in areas like mental representation, intelligence or behaviour.

Now, rather than frame a contrast around the above portrayal of reductionist and emergentist approaches as such, it is this somewhat broader distinction I would like to bring out; contrasting an emergentist/pluralist approach to explanation in cognitive science and philosophy of Mind, with other tendencies towards more 'single-level' approaches.

Is the contrast a valid one? It would seem that there are some claims made within contemporary cognitive science that could be interpreted as at least leaning in the latter, 'single-level' direction. For example, Elman et al (1998, pp. 367-391) cite a number of works they would regard as taking a too-strongly 'representational nativist' position. This critique is based precisely on a view that these other approaches, in the context of making claims about the 'innateness' of certain notional features of mental representation, do indeed rely too much on genetic factors as *causing* these features to occur (and consequently overlook other relevant factors). Amongst the authors mentioned are Carey and Spelke, from an essay of 1994 entitled 'Domain-specific knowledge and conceptual change', who say:

> We argue that human reasoning is guided by a collection of *innate domain-specific systems of knowledge*. Each system is characterised by a set of core principles that define the entities covered by the domain and support reasoning about those entities.
>
> (cited in Elman et al 1998, p. 367, emphasis added)

They also quote linguist Stephen Crain (1991), who claims that:

> [Syntactic] knowledge is in large part innately specified. What is *innately given is knowledge of certain restrictions on the meanings that can be mapped onto sentences* as well as restrictions on the sentences that can be used to express meanings.
>
> (cited in Elman et al 1998, p. 369, emphasis added)

It is the view of Elman et al that both these claims (presumably taken in context) are too strongly 'nativist', because each would seem to require '*genetic specification* of cortical wiring at the synaptic level' (1998, p. 370, emphasis added) – a very unlikely prospect in their view. In other words, this is where an over-emphasis on putative causal factors at the level of genetic information is seen to lie. (The position of Elman et al in this respect of course is that a combination of causal factors enters into the development of representational features at the synaptic level. The same basic view is also central to my characterization of these processes as having a certain level of intrinsic *variability*.)

More generally Elman et al argue that, if nothing else, such views should be presented with full cognisance that the term 'innate' can mean some quite different things (note the discussion about 'innateness' as either 'direct' or 'extended' in section 6.2.2). When 'innateness' is 'unpacked', they suggest, a more adequate picture of the complexity of factors entering into mental representation may well result.

By way of contrast, we may note another view that in its own way has also tended towards claims of systemic outcomes in human behaviour being caused by factors pertaining at a single level. The radically empiricist 'behaviourist' school of psychology, most exemplified by B.F. Skinner, has been widely discussed for its theories of operant conditioning, and the view that the bulk of human behaviours are a direct product of positive or negative reinforcement from environmental events. (E.g. Fodor 1981) Here, of course, it is the environmental factors that are 'picked out' as the preeminent causal factors in shaping subsequent behavioural outcomes.

It is not my purpose here to engage directly in any argument about the relative merits of these views, but simply to suggest that there are views around which (at least partially) reflect the distinction between a pluralistic, multi-level approach to explanation in cognitive science, and other more 'single-level' approaches. I will return to the significance of these observations later.

## 8.3 Science and Social Interpretation

Having made mention of the above points, let us for the moment put to one side these points of debate or tension that are, in a sense, more internal to science itself, and turn to the first part of the main discussion; issues arising in the relationship between science and socio-political ideas.

My aim in this section is to sketch out what I regard as an area of socio-political discourse that fits the rough criteria mentioned above – it encapsulates certain views about 'human nature', and these views are potentially significant for how we see ourselves and how we act in social or political contexts. Having done that I will consider a number of further issues regarding the relationship between these views and science-based ideas, and the particular perspectives that might be offered by a creative organisation view.

### 8.3.1 'Biological Determinism', Social Darwinism and Competition

In order to outline the area of discourse I have in mind, and do so succinctly, I will rely on the work of Rose, Kamin and Lewontin in their book, *Not in Our Genes: Biology, Ideology and Human Nature* (1984). These writers, referring particularly to currents of public and political discourse in the United States, propose a nexus of consistent socio-political beliefs, grounded in science-based ideas, under the description of *biological determinism.*

Citing a range of evidence, including the published work of a range of writers (both scientists and non-scientists), they identify a stance of biological determinism in three basic claims:

- That inequality in (contemporary, free-market, Western democratic) society is primarily a consequence of differences in intrinsic merit and ability in (and between) individuals
- That these differences in individual's intrinsic merit and ability are in large part a result of differences in their inherited genetic make-up

- That the inherent differences in people lead to hierarchical societies because it is 'part of biologically determined human nature to form hierarchies of status, wealth and power'.

(Rose et al 1984, pp. 68-69)

My intention for the purposes of this chapter is simply to take this portrayal at face value; to accept Rose et al's case that a set of such views about human society exists (or has recently existed), and that the ideas which constitute it have (or have had) some level of social currency. From this point, then, what might be further observed about the biological determinist position?

Firstly it would seem clear that amongst these points one may readily discern a broad 'view about human nature', together with some further ideas of what this view means in a wider social context.

Secondly, it would also seem clear that this view refers to issues and ideas within the specific purview of scientific enquiry; e.g. the genetic heritability of phenotypic characteristics, or the derivation of certain contemporary human characteristics in the 'forces' of natural selection. Rose et al argue that, in a more general sense, the biological determinist position is very much influenced by – or, one might say, is a version of – 'social Darwinism': that is, taking a cue from Darwin's theory of natural selection, a view which extrapolates ideas like 'survival of the fittest' into a specifically human, social domain. (Also: Hofstadter 1955) This is not to argue, of course, that such references or influences are necessarily adequate to their scientific sources, but only to note that the link is there.

Thirdly, we may observe that such ideas, if adopted, would seemingly have significant impacts on one's attitudes and responses to a range of social issues. At a general level, as Rose et al again are concerned to argue, such ideas may readily be used to *justify social inequality*; the argument being (roughly): if the tenets of biological determinism are true, then any socially mediated or political efforts to create a level of social equality substantially different to the prevailing conditions will be doomed to fail because it would have to contend with the genetic factors that are, in effect, creating those conditions. (1984, p. 70)

Amongst the many examples cited by Rose et al is Harvard academic, Richard Hernstein, who says, in his book, *IQ in the Meritocracy*:

> The privileged classes of the past were probably not much superior biologically to the downtrodden, which is why revolution has a fair chance of success. By removing artificial barriers between classes, society has encouraged the creation of biological barriers. When

people can take their natural level in society, the upper classes will, by definition, have greater capacity than the lower.

Cited in Rose et al 1984, p. 69

Hernstein, to my reading, would indeed seem to be claiming that 'natural' biological differences between people are the dominant cause of class-based social inequality in contemporary society. We can only presume that the 'biological barriers' he refers to are the social working out of genetic differences between individuals.

Hernstein, I would suggest, is also an example of someone setting out a socio-political view against a background of more direct involvement in academia and science. There are several others mentioned by Rose et al who again would seem to be expressing views consistent with the biological determinist position, again drawing on ideas derived from science, but doing so on the basis of work or reputation in other fields. For example, businessman, John D. Rockefeller, who said, 'The growth of a large business is merely a *survival of the fittest*. ... This is not an evil tendency in business. It is merely the working out of a law of nature.' (Cited from Hofstadter 1955, emphasis added)

So we may have some initial sense of how a biological determinist-type view might be expressed in different domains.

Given that Rose et al's book is now some twenty years old one might be concerned to suggest that times have moved on, and comments such as those above would be less likely to be heard today. I would concur with such a view in a limited way, but also point out that other (not entirely unrelated) ideas about inherent differences in the value and worth of individuals have been at the roots of many of the modern world's most shameful episodes – and particularly one might think of the eugenics research programs undertaken by the Nazi regime in the 1930's and 40's. In historical terms, these events are not so very far removed from the present.

Secondly, I would suggest that, although times have changed, that certain views about 'human nature' remain current and influential, which are also still at least partially or implicitly consistent with a more 'fully expressed' biological determinist or social Darwinist view. I refer in particular to (overly simplified) beliefs that human beings are inherently 'competitive' creatures, that the competitive society expresses this defining human characteristic, and that our success or failure in society depends, for the most part, on our individual, intrinsic capacity to compete in that social domain.

For example, the view of the Australian Prime Minister, John Howard:

> [Australians believe that] the individual is more important than the state, … that competitive free enterprise is the ultimate foundation of national wealth, and that *the worth of a person is determined by that person's character* and hard work, *not by* their religion or race or creed or *social background.*
>
> Federal Hansard, 2003, emphasis added

In my opinion such words, understood in a wider historical context, strongly indicate (as sub-text, we might say) a larger set of socio-political beliefs and ideas at least partially consistent with those of biological determinism. I would particularly note the rhetorical linking of notions of competition, individualism and individual worth based on 'character'.

My limited goal here, however, is not to substantiate the veracity of this interpretation, or to otherwise verify a continued current of biological determinist-type beliefs. It is only to say, *if* such views continue to exist, or exist partially, and continue to have a genuine socio-political import; and *if* they explicitly or tacitly maintain some points of reference or justification in ideas derived, however distantly, *from science*; then this is a matter in which science and scientists would have some interest and some stake.

### 8.3.2 The Role of Science

The immediate point I would make about a biological determinist-type view, of course, again taking Rose et al's characterisation of it at face value, is that it appears – as a socio-political interpretation of science-based ideas – to be following (in its own fashion) the single-level explanatory approach identified as a tendency within science. In this case, seemingly, the genotype is being invoked as *the* pre-eminent causal factor involved in generating the subsequent 'intrinsic merit and ability' of the individual phenotype; and it is differences in genotypes that are invoked as sufficient to 'explain' differences in 'merit and ability' between individuals (and thus, by extension, different levels of social status within a hierarchical society). In other terms, one might suggest that the Darwinian-derived idea of 'genetic fitness' as the key to individual survival is a competition for resources has been narrowly interpreted, and inflated to an encompassing social principle.

If this is so, then what might one say about any implications for science itself? For one thing, it is important to guard against any assumptions, on the basis of (perhaps) disagreeing with current levels of social inequality, or with beliefs justifying that inequality, that *therefore* any scientific ideas or theories invoked by those beliefs are necessarily wrong. Furthermore, there is no assertion being made here that scientists working in genetics or evolutionary theory, or what have you, are necessarily in any way deliberately complicit in promulgating such views.

Having made those qualifications, however, there are two main issues I would raise. The first regards the creative organisation view (or indeed other related views as canvassed in this study that are consistent with the emergentist/pluralist approach outlined earlier). In broad terms it would seem clear to me that the creative organisation view put here, at least in the areas of early human development and cognition it has been concerned to address in detail, is one that would *not* sit easily with the central tenets of biological determinism, or other similar views. The creative organisation view should, if nothing else, give rise to considerable caution about any claims to the effect that the development of various cognitive abilities and capacities (which surely must fall somewhere within what is understood to constitute an individual's 'intrinsic merit and ability' within the biological determinist view), are the simple or direct result only of prior factors encoded in the genotype (or indeed that they are wholly and simply *intrinsic* to each individual in any sense). The creative organisation view would for its part say, if nothing else, 'Hold on! It's more complex and variable than that, *because* … etc etc.'

Now that on its own, having 'come all this way', might not seem (yet) very much, and if so I would in a certain sense agree. However, the point of this exercise for me has also been to explore and set out a *foundation* of ideas about some fundamentals of 'human nature' – in particular of course, ideas about our nature as intelligent, responsive and creative beings. It is hoped that that foundation may, in due course, provide a fresh, scientifically-informed ground for other, further ideas about our social development. On the basis of other work (Fisher 1997), it is my view that such ideas may eventually offer a contrasting view about such matters as social inequality; its roots in the dynamics of socially embedded human development, and its open-ness to socially constructed strategies for change.

The significance and value, in my view, of a case to support the critical 'Hold on!' perspective above lies in two basic matters. Firstly, regardless of saying 'Hold on!' to any other view, I see it as vitally important that any kind of social theorizing be built on a strong foundation of more fundamental ideas about the 'kind of beings we are'.

This is indeed the sense of the observation by Noam Chomsky cited in the first chapter:

> To command attention and respect, a social theory should be grounded on some conception of human needs and human rights, and in turn, the human nature that must be presupposed in any serious account of the origin and character of these needs and rights.
>
> Chomsky 1988, p. 195

It is also potentially important politically, simply because other socio-political perspectives, such as biological determinism, will very likely lay claim, rightly or wrongly, to some equally fundamental ideas; and these may need to be contested directly.

The second issue to consider then concerns the approach taken by scientists at large in putting certain views into the public domain. If we can recognize that there are certain currents of ideas extant in society – such as biological determinist, or social Darwinist-type ideas – which can 'pick up on', interpret and use science-based ideas for their own purposes; then the way ideas are put can take account of that, and perhaps seek to avoid what might well amount to a misinterpretation and/or a misuse of ideas.

In particular, I would suggest, one might recognize that, today, particular ideas which are suggestive of simple or direct causal links between genetic factors and certain systemic outcomes of human development and functioning – even if not intended to be so, or if various caveats and qualifications about complexity or other factors are omitted for ease of communication –may inadvertently provide grist to the mill of such political beliefs.

As Elman et al put the issue in relation both to their own work on the innateness (or otherwise) of cognitive development and representation, and the kind of critiques mentioned earlier:

> The words we use to explain out interim findings still have important consequences. The word "innate" means different things to different people … to some people "innateness" means that the outcome in question cannot and should not be changed. … If scientists use words like "instinct" and "innateness" in reference to human abilities, then we have a moral responsibility to be very clear and explicit about what we mean, to avoid our conclusions being interpreted *in rigid nativist ways by political institutions.*
>
> 1998, p. 391, emphasis added

Thus we come back to the potential significance of the kinds of general scientific approaches and attitudes to the business of explanation mentioned earlier. Scientists may choose for many reasons to put into the public domain ideas that purport to explain aspects of our complex human intelligence, behaviour or what have you, and these ideas may (according to the descriptions I have offered) be reductionist in approach and/or invoke single-level causes – particularly direct genetic causes – or be *liable to interpretation in that vein.* If such claims are rigorously argued and seriously intended, then of course they should be subject to scientific scrutiny and debate in the usual ways. However, in any case one would hope that public discourse is not entered into glibly, or in ignorance of any unintended socio-political consequences that might follow.

## 8.4 Prospects for Future Work

Finally, as noted above, on the basis of previous work (E.g. Fisher 1997) I already hold the working view that the ideas developed here do have some potential as a ground for further exploration and exposition in some of the areas typically addressed by psychology, sociology, social or political philosophy. In very rough terms the line of enquiry and argument I would like to pursue is as follows:

- That, by virtue of the creative-organisational, embodied, embedded character of cognition we are, as Minded beings, indeed highly responsive to the character of our relationships with other people, and to the particular social or cultural norms we inhabit.
- That the creative character of cognition, as a developmental potential (in the broadest sense of the term), can be developed in different ways, or to different *extents*.
- That the fulsome development of this creative potential in individuals can be shown, on the whole to also be of social benefit
- That certain kinds of socially mediated conditions may be shown to be conducive to the fulsome development of our creative potential, and certain others to distort or inhibit it.
- That the above distinction is significant in terms of understanding and constructively responding at a socio-political level to a range of negative trends within our current social dynamics.

Now that may be an overly ambitious agenda, however, I would not resile from it, for one main reason. I believe that for any science-based understanding to 'flow on' into socio-political interpretation; to provide insight into our current, in some ways problematic, situation (as outlined in Chapter1); and to inform alternative pathways of social development, it *must* do one particular thing. It must, at some point, articulate some kind of critical distinctions between forms of individual-social interaction, and the differing results for individual or social health and development that generally follow from them (as per the fourth point above). This is not something foreign to science. For example, there are numbers of studies in the epidemiological area that do this kind of thinking; about differential health effects that stem from different kinds of social situation. (E.g. Marmot & Wilkinson 1999) There are also a wide range of existing understandings developed within fields such as child psychology, and many other places, that do similar things and will no doubt also inform and constrain such a line of enquiry. However, based on my current limited knowledge, it seems to me that a creative organisation perspective, pursued as philosophy, *may* have something fresh to offer in these areas.

~

# BIBLIOGRAPHY


- Atlan, H. 'Uncommon Finalities' in Thompson, W.I. (ed.) (1987) *Gaia: A Way of Knowing* Massachusetts, Lindisfarne Press
- Auyang, S. (1998) *Foundations of Complex-system Theories; in Economics, Evolutionary Biology, and Statistical Physics* Cambridge U.K., Cambridge University Press
- Auyang, S. (2000a) *Synthetic Analysis of Complex Systems I - Theories* University of Sydney http://www.usyd.edu.au/su/hps/newevents/Auyang1.html
- Auyang, S. (2000b) *Synthetic Analysis of Complex Systems II - The Genome and the Brain* University of Sydney http://www.usyd.edu.au/su/hps/newevents/Auyang2.html
- Auyang, S. (2000c) *Mind in Everyday Life & Cognitive Science* Cambridge Massachusetts, Bradford
- Ayala, F.J. and Dobzhansky, T. (eds.) (1974) *Studies in the Philosophy of Biology* London, Macmillan
- Bates, E., Thal, D. and Janowsky, J. (1992) 'Early language development and its neural correlates.' In Rapin, I. And Segalowitz, S. (eds.) *Handbook of Neuropsychology, Vol. 7: Childhood Neuropsychology* (pp. 69-110) Amsterdam, Elsevier
- Bateson, G. (1972) *Steps to an Ecology of Mind* New York, Ballentine Books
- Bateson, G. (1979) *Mind and Nature* New York, E.P. Dutton
- Bechtel, W. (1988) *Philosophy of Science; An Overview for Cognitive Science* New Jersey, Lawrence Erlbaum Associates
- Bergson, H. (1912) *Creative Evolution* London, Macmillan
- Bertalanffy, L.V. (1971) *General System Theory* London, Allen Lane The Penguin Press
- Birch, C. (1990) *On Purpose* Sydney, University of New South Wales Press
- Birch, C. (1993) revised edition *Confronting the Future* Ringwood, Victoria, Penguin Books
- Birch, C. (1999) *Biology and the Riddle of Life* Sydney, University of New South Wales Press
- Boden, M. (1996) *The Philosophy of Artificial Life* Oxford, Oxford University Press
- Boden, M. (1994) 'What is Creativity?' in Boden, M. (ed) *Dimensions of Creativity* London, Bradford
- Bohm, D. 'Some Remarks on the Notion of Order' in Waddington, C.H. (1969) *Towards a Theoretical Biology* Vol. 2 of four, Edinburgh U.K., Edinbugh University Press
- Bohm, D. & Peat, F.D. (1989) *Science, Order & Creativity* London, Routledge

- Braddon-Mitchell, D. and Jackson, F. editors (1996) *The Philosophy of Mind and Cognition* Massachusetts, Blackwell Publishers
- Broad, C.D. (1925) *The Mind and its Place in Nature* London, Routledge & Kegan Paul
- Brown, L.R. (ed. 1990) *State of the World* Sydney, Allen & Unwin
- Bunge, M. (1979) 'Ontology II: A World of Systems' vol. 4 of *Treatise on Basic Philosophy* Holland, D. Reidel Publishing Company
- Bunge, M. (1977) 'Levels and Reduction' in *American Journal of Physiology* 233(3), pp. R75-R82.
- Capra, F. (1992) *The Turning Point* London, Flamingo
- Capra, F. (1996) *The Web of Life; A New Synthesis of Mind and Matter* London, Harper Collins
- Carey, S. and Spelke, E. (1994) 'Domain-specific knowledge and conceptual change' in Hirschfeld, L.A. and Gelman, S.A. (eds) *Mapping the Mind: Domain Specificity in Cognition and Culture* Cambridge UK, Cambridge University Press
- Chomsky, N. (1988) *The Chomsky Reader* edited by James Peck, London, Serpent's Tail
- Churchland, P.M. (1999) 'Eliminative Materialism and the Propositional Attitudes' in Lycan, W.G. (1999) *Mind and Cognition: An Anthology*, second edition, Massachusetts, Blackwell Publishers
- Clark, A. (1993) *Associative Engines: Connectionism, Concepts and Representational Change* London, MIT Press
- Clark, A. (1997) *Being There: Putting Brain, Body and World Together Again* Massachusetts, Bradford
- Clark, A. (2001) *Mindware: An Introduction to the Philosophy of Cognitive Science* Oxford University Press, New York
- Cole, D.J, Fetzer, J.H. and Rankin, T.L. editors (1990) *Philosophy, Mind and Cognitive Enquiry* Dordrecht, Netherlands, Kluwer Academic Publishers
- Crain, S. (1991) 'Language acquisition is the absence of experience.' *Behavioural and Brain Sciences, 14,* 597-611
- Darwin, C. (1950) *On the Origin of Species by Means of Natural Selection* London, Watts
- Davies, P. (1995) *The Cosmic Blueprint; Order & Complexity at the Edge of Chaos* London, Penguin Books
- de Chardin, T. (1964) *The Future of Man* London, William Collins Sons & Co. (first published 1959)

- Dobzhansky, T. 'Chance and Creativity in Evolution' in Ayala, F.J. and Dobzhansky, T. (eds.) (1974) *Studies in the Philosophy of Biology* London, Macmillan
- Edelman, G. (1992) *Bright Air, Brilliant Fire: On the Matter of Mind* United States, HarperCollins (Basic Books)
- Ehrlich, P. R. & Ehrlich A.H. (1991) *Healing the Planet; Strategies for Resolving the Environmental Crisis* Chiping Norton, NSW, Surrey Beatty &Sons
- Elman, J.L. (1991) 'Representation and structure in connectionist models.' In G. Altman (ed.) *Cognitive Models of Speech Processing* Massachusetts, MIT Press
- Elman, J.L.; Bates, E.A.; Johnson, M.H.; Karmiloff-Smith, A.; Parisi, D.; and Plunkett, K. (1998) *Rethinking Innateness: A connectionist perspective on development* London, MIT Press
- Fisher, I.M. (1997) *The Journey Home: Holism, Human Nature and Social Change* Adelaide, Seaview Press (self published)
- Fodor, J.A. (1974) 'Special Sciences' *Synthese 28* pp. 97-115
- Fodor, J.A. (1981) 'The Mind-Body Problem' from *Scientific American* Vol. 244, pp. 124-132
- Fodor, J.A. (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* London, the MIT Press
- Freeman, W.J. (1991) 'The physiology of perception.' *Scientific American, 264,* pp. 78-85
- Fromm, E. (1991) *The Sane Society* London, Routledge (Originally published 1956)
- Goodfield, J. 'Changing Strategies: A Comparison of Reductionist Attitudes in Biological and Medical Research in the Nineteenth and Twentieth Centuries' in Ayala, F.J. and Dobzhansky, T. (eds.) (1974) *Studies in the Philosophy of Biology* London, Macmillan
- Grene, M. 'Reducibility: Another Side Issue' in Grene, M. (ed.) (1971) *Interpretations of Life and Mind* London, Routledge & Kegan Paul
- Grene, M. (ed.) (1971) *Interpretations of Life and Mind* London, Routledge & Kegan Paul
- Harman, G. editor (1993) *Conceptions of the Human Mind: Essays in Honour of George A. Miller* New Jersey, Lawrence Erlbaum Associates
- Hofstadter, R. (1955) *Social Darwinism in American Thought* (revised edition) New York, George Brazillier Inc
- Howard, J. (2003) Welcoming speech to the President of the United States *Federal Hansard, Thursday 23 October 2003* Australia
- Johnson-Laird, P.N. (1988) *The Computer and the Mind: An Introduction to Cognitive Science* London, Fontana Press
- Koestler, A. (1975) *The Act of Creation* London, Pan Books

- Laplace, P.S. (1951) *A Philosophical Essay on Probabilities* New York, Dover Publications (first published 1819)
- Lycan, W.G. 'The Continuity of Levels of Nature' in Lycan, W.G, (ed.) (1999) *Mind and Cognition: An Anthology* Massachusetts, Blackwell Publishers
- Mahnen, M. & Bunge, M. (1997) *Foundations of Biophilosophy* Springer.
- Margulis, L. (1987) 'Early Life' in Thompson, W.I. (ed.) (1987) *Gaia A Way of Knowing; Political Implications of the New Biology* Massachussetts, Lindisfarne Press
- Marmot, M.G. and Wilkinson, R.G. (eds) (1999) *Social Determinants of Health* Oxford, Oxford University Press
- Maturana, H.R. & Varela, F.J. (1987) *The Tree of Knowledge* Massachusetts, Shambhala Publications
- Mayer, E. 'Cause and Effect in Biology' in Waddington, C.H. (1969) *Towards a Theoretical Biology* Vol. 1 of four, Edinburgh U.K., Edinbugh University Press
- Miller, K., Keller, J. and Stryker, M. (1989) 'Ocular dominance column development: Analysis and simulation.' *Science* 245, 605-615.
- Montalenti, G. 'From Aristotle to Democritus via Darwin: A Short Survey of a Long Historical and Logical Journey' in Ayala, F.J. and Dobzhansky, T. (eds.) (1974) *Studies in the Philosophy of Biology* London, Macmillan
- Nicolis, G. and Prigogine, I. (1989) *Exploring Complexity: An Introduction* New York, W.H. Freeman & Co.
- Pattee, H. 'The Problem of Biological Heirarchy' in Waddington, C.H. (1970) *Towards a Theoretical Biology* Vol. 3 of four, Edinburgh U.K., Edinbugh University Press
- Peacocke, A. (1986) *God & the New Biology* London, J.M. Dent
- Popper, K. 'Scientific Reduction and the Essential Incompleteness of All Science' in in Ayala, F.J. and Dobzhansky, T. (eds.) (1974) *Studies in the Philosophy of Biology* London, Macmillan
- Popper, K. and Eccles, J. (1977) *The Self and its Brain* Berlin, Springer International
- Prigogene, I. 'Unity of Physical Laws and Levels of Description' in Grene, M. (ed.) (1971) *Interpretations of Life and Mind* London, Routledge & Kegan Paul
- Prigogene, I. (1980) *From Being to Becoming: Time and Complexity in the Physical Sciences* San Francisco, W.H. Freeman & Co.
- Prigogine, I. & Stengers, I. (1984) *Order Out of Chaos; Man's New Dialogue with Nature* London, Shambhala Publications
- Quartz, S. (1999) 'The constructivist brain.' *Trends in Cognitive Sciences*, 3(2), 48-57

- Rosak, T. (1993) *The Voice of the Earth* New York, Simon and Schuster
- Rose, S., Kamin, L.J., and Lewontin, R.C. (1984) *Not in Our Genes; Biology, Ideology and Human Nature* Middlesex U.K., Penguin
- Rosen, R. (1991) *Life Itself; A Comprehensive Inquiry into the Nature, Origin and Fabrication of Life* New York, Columbia University Press
- Rosen, R. (2000) *Essays on Life Itself* New York, Columbia University Press
- Rosenthal, D.M, editor (1991) *The Nature of Mind* New York, Oxford University Press
- Rumelhart, D.E., Smolensky, P., McClelland, J.L. and Hinton, G,E. (1986) 'Schemata and Sequential Thought Processes in PDP Models' in Rumelhart, D.E., and McClelland, J.L (eds.) *Parallel Distributed Processing: Volume 2* Massachusetts, MIT Press
- Sachs, O. (1990) *Awakenings* (revised edition) New York, HarperCollins
- Saul, J.R. (1993) *Voltaire's Bastards: The Dictatorship of Reason in the West* London, Penguin
- Schlagger, B. and O'Leary, D. (1991) 'Potential of visual cortex to develop an array of functional units unique to somatosensory cortex.' *Science 252*, 1556 - 1560
- Searle, J. R. (1982) 'Minds, Brains, and Programs' from *The Mind's Eye* Hofstadter, D.R. and Dennett, D.C. (eds) New York, Bantam Books
- Sejnowski, T. and Rosenburg, C. (1986) *NETtalk: A Parallel Network That Learns to Read Aloud.* (Technical Report JHU/EEC-86/01) Baltimore, MD. Johns Hopkins University.
- Shearman, D. with Sauer-Thompson, G. (1997) *Green or Gone; Health, Ecology, Plagues, Greed and our Future* Adelaide, Wakefield Press
- Shiva, V. (1991) *Ecology and the Politics of Survival* New Delhi, Sage Publications
- Skinner, B.F. (1972) *Beyond Freedom and Dignity* London, Jonathan Cape
- Spelke, E.S. (1994) 'Initial knowledge: Six suggestions.' *Cognition, 50*, 431-445
- Spretnak, C. & Capra, F. (1985) *Green Politics; The Global Promise* London, Paladin
- Suzuki, D. (1993) *Time to Change* St Leonards, NSW, Allen & Unwin
- Teeple, G. (1995) *Globalisation and the Decline of Social Reform* Toronto, Garamond Press
- Thelan, E. & Smith, B. (1994) *a Dynamic Systems Approach to the Development of Cognition and Action* London, MIT Press/Bradford
- Thompson, W.I. (ed.) (1987) *Gaia A Way of Knowing; Political Implications of the New Biology* Massachussetts, Lindisfarne Press
- Thorpe, W.H. 'Reductionism in Biology' in Ayala, F.J. and Dobzhansky, T. (eds.) (1974) *Studies in the Philosophy of Biology* London, Macmillan.

- Trainer, T. (1996) *Towards a Sustainable Economy; The Need for Fundamental Change* Sydney, Envirobook and Oxford, Jon Carpenter Publishing
- Varela, F., Thompson, E., and Rosch E. (1996) *The Embodied Mind* London, MIT Press
- Vernon, P.E. (1970) *Creativity*, Middlesex UK, Penguin
- Waddington, C.H. (ed.) (1968-72) *Towards a Theoretical Biology* in four Volumes, Edinburgh U.K., Edinbugh University Press
- Weiss, P. 'From Cell to Molecule' in Allen, J.M. (ed) (1962) *The Molecular Control of Cellular Activity* New York, Mcgraw-Hill
- Westphal, D. & Westphal, F. (eds.) (1994) *Planet in Peril; Essays in Environmental Ethics* Orlando Florida, Harcourt Brace & Co.
- Whitehead, A.N. (1926) *Science and the Modern World* Cambridge, Cambridge University Press
- Whitehead, A.N. (1978) *Process and Reality* (corrected edition, eds. Griffin, D.R. and Sherburne, D.W.) New York, Macmillan (first published in original edition 1929)
- Wilson, E.O. (1998) *Consilience; The Unity of Knowledge* New York, Alfred A. Knopf Inc.
- Wilson, E.O. (1975) *Sociobiology: The New Synthesis* Cambridge, Harvard Press