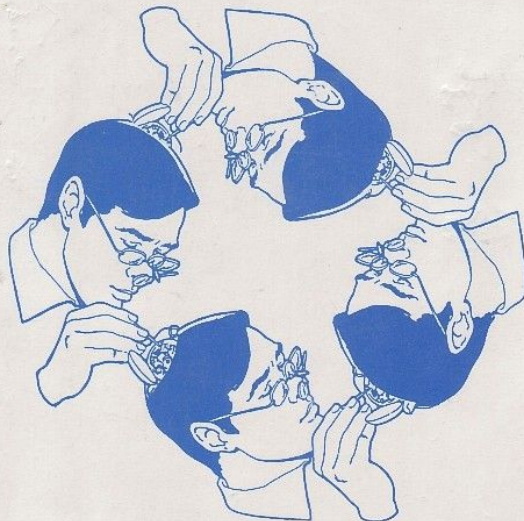# Victor Serebriakoff

THE FUTURE OF
INTELLIGENCE

BIOLOGICAL AND ARTIFICIAL

Victor Serebriakoff

Foreword by ISAAC ASIMOV

TABLE OF CONTENTS

FOREWORD

by Isaac Asimov

When I was very young, I read the myth of Narcissus. Narcissus was a handsome youth who rejected all the young women who were dying for his embraces. One of them prayed that he, too, would feel the pangs of unrequited love, so, of course (since the gods are always eager to answer unkind prayers) he did.

He spied his reflection in the water, thought it was another youth, fell in love with its beauty, attempted futilely to embrace it, and finally drowned when he tried too hard. The myth has made such an impression on people that the word 'narcissism' is a recognised psychiatric term for morbid self-love.

I did not like the myth at all. To my childish self, it seemed incredibly stupid. How could Narcissus mistake his own image for another youth? (I was also more than a little puzzled that he should fall in love with another youth, rather than with a maid, but the problem of the image overrode that.)

No one bothered to explain that point to me. They were only interested in explaining the moral: that if you are unkind to others, then others will be unkind to you; that if you are too fond of yourself, you will find life unpleasant.

I saw the moral, of course; it was obvious. What I wanted, though, was a technological explanation, and I never got it. I had to work it out for myself as I grew older.

The point is that in primitive times, it was perfectly possible for a person to see the faces of those around him with perfect clarity (assuming eyesight was normal.) He (or she) could tell, at sight, the identity of every human being with whom he was acquainted. He would also tell at a glance that some person he encountered was a stranger to him, someone he had never seen before.

There was one exception. Under primitive conditions, no person could see his (or her) own face. If, through some form of magic, his own face were presented to him, he would have no choice but to consider it that of a stranger. To see your own face without magic, you need a smooth, reflecting surface. A piece of smooth, unflawed glass will do, rather dimly, for it lets pass much more light that it reflects. A piece of glass backed by a smooth layer of metal, will do it with near-perfect efficiency. That would be a 'mirror' or a 'looking-glass' (for at what glass would we look more eagerly than at a mirror.) However, in primitive times, smooth unflawed glass, with or without metal backing, was unavailable.

One could simply polish a flat piece of metal. That would do well enough till it tarnished. However, in primitive times, a piece of polished metal large enough to see your face in was not an easy thing to get.

That left the surface of water, which was usually so broken up by waves, foam, and (even in quiet ponds) ripples that a reflected face was too disturbed to make much impression. If, then, Narcissus had come upon a pool so quiet that he could see a clear impression of his own image, you can well imagine that it

was the first time he had seen it and, of course, it would seem like a strange youth to him, one who was hiding under the water.

Once you understand the tale of Narcissus in its true light, then you may come across a sudden analogy. Human beings try to understand the universe little by little. They look at this aspect and at that aspect and learn to analyse the appearance and characteristics and begin to understand it. We can be very proud of the fact that all through our existence, our understanding of the universe has grown enormously. But, then, why not? We are looking at every aspect of the universe with one chief tool, the human brain and the intelligence with which it is associated. Since the human brain is by far the most complexly interrelated piece of matter that we know of, we are using a complex tool to understand the much less complex objects which we are observing. Given enough time and thought, we must understand.

But then comes a point when we wish to contemplate the human brain and the intelligence with which it is associated. Now we are attempting to understand something extremely complex by making use of a tool that is no more complex. The situation is analogous to that of someone using his eyes to see his own eyes. As soon as we face the problem of the human brain and human intelligence we are in Narcissus's case - faced with our own image; and therefore doomed, perhaps, to misunderstanding and death.

But there are solutions. Narcissus's solution would have been familiarity with a mirror.

Our solution is that we are not using a human brain to study the human brain and human intelligence. We are using many human brains to do so.

It is not a scientist who is studying the human brain, it is, rather, the community of science. It has a complex structure of its own, with published papers, with frequent conferences, with communications in which different thinkers present different pictures, different interpretations, different observations. In kaleidoscopic fashion, these all melt together and grow almost without the volition of any individual, so that understanding increases at a speed and to an extent that any one person would find amazing.

What we (who are intelligent but who have done no work on intelligence) need is someone who has followed the work being done on all the aspects of science that impinge on intelligence to present them to us in orderly fashion and, if possible, with his own thoughts and ideas added to the mix.

This is precisely what Victor Serebriakoff has done. Himself a person of monstrous intelligence, he has obviously read, studied, and thought about every aspect of human intelligence, and here it is for us to share with him. He comes to rescue us from Narcissus's fate.

Isaac Asimov

PREFACE 1999.

As I near my eighty seventh birthday this will probably be the last edition I that can alter. It is forty two years since my original 'Eureka' type inspiration, if that is what it was, occurred. It now seems that at long last, my ideas from so long ago are to be tried out in the world of business. Much thanks to an old friend Russell Swanborough, who has created a management consultancy business on the basis of my ideas.

Indirectly the ideas that follow have already led to two important developments in the field of artificial quality control, the first was a simple machine for measuring the strength of building wood and the second was much closer to a pattern recognition program that worked well but probably too well because its introduction was much delayed by the natural resistance of the trade unions of the graders. It was hampered by the fact that the technicians who spent many years designing the machine decided against modern computers and preferred the extant hardwired systems made up of bistables.

Now a group of investors are putting expert programmers on to the job of adapting my ideas about the architecture of the information flow in cells, brains and societies to another practical use in the tricky art of business prediction.

I have written twenty-two books and most have been published. All except three were written to fit a market; I wrote what I could sell. Two slim volumes of light verse were written out of love and joy. The other two were written from deep conviction and a desire to communicate a set of ideas that have always haunted me.

This is a third attempt, in the light of much that has happened in science since my second serious book, (published 1987,) to communicate what is most difficult to communicate, thoughts about thinking, ideas about ideas, intelligence about intelligence. There are a number of loosely bound theses that seem to hang together but not necessarily for mutual support. Some of the package of ideas I present may be overstated or plain wrong without there being a need to reject the entire parcel.

My hope is that the ideas or some others on similar lines may prove to be seminal. I say that because of my grave fear that that central and distinguishing aspect of our race, mankind, is being devalued and, worse, demeaned by a trend in modern thinking that is as genuine and compassionately motivated as it is dangerous. I refer to the trend to reject, despise and devalue excellence and especially the most important kind of excellence, cognitive excellence.

This book is one man's attempt to examine what he sees as the most important phenomenon in the universe, intelligence in all its many forms.

ACKNOWLEDGEMENTS

1. LIFE AND INTELLIGENCE

PREAMBLE

This book aims to present a fresh paradigm for the structure of science. It is based on conjectures in the late fifties that were first published in the books '*Brain*' [Davis Poynter 1976] and '*The Future Of Intelligence*' [Parthenon Publications 1987]. Constant study and reading in the field has revealed much research that fits, much that confirms and none that falsifies the original conjectures. The author has been advised to present the case a new in an updated form. [The word paradigm is used in Kuhn's sense, an approach to a problem from new direction introducing some new concepts and new linkages between them.]

The style is that of popular science rather that of academe, using the simplest and the least technical language which will serve. Unavoidably with a new paradigm there will be some self-explanatory neologisms.

There are a number of linked themes based on a fundamentally new way of looking at the sciences of Things and of those of Creatures, the 'hard' and 'soft' sciences. A new view may help understanding of the areas of incoherence that have emerged as the proliferating scientific disciplines have forked apart into specialist niches. The different new approaches to be developed are separate, only partially coherent. They do not stand or fall as a group. Each theme will be dealt with in a separate chapter and calls for a separate judgement.

THE FIRST TOOL OF SCIENCE.

The exercise starts with a look at the basic tool of philosophical and scientific enquiry, intelligence. Any kind of brain must have its limitations so any enquiry made by a brain will have epistemological limitations. It is not impossible to speculate about what these might be so that they can be allowed for in conclusions or explain lack of coherence.

HOW NEW MODELS APPEAR

Thomas S. Kuhn showed (in '*The Structure Of Scientific Revolutions*') what is really held in common by those tight, inward-facing elites that constitute the scientific communities in each discipline. It is, he argued, the science-teaching textbooks with the carefully selected exemplary models that they contain. He claimed that it is these examples, paradigms, thought-moulds, perspectives, rather than the definitions and laws of the discipline, that underlie and inform the common mind-set and way of thinking of a scientific school or group.

Kuhn showed that the history of science is the story of a succession of revolutionary new paradigms, each of which takes a generation to establish its pre-eminence. New paradigms usually come from the young or from those new to the field. They arise when *ad hoc* adjustments to the accepted paradigm get too messy. The anomalies within the theory or between theory and prediction become too great and, from somewhere, strongly resisted at first, a new paradigm emerges and gradually takes the place of the old one. The science

then proceeds to refine and perfect the new model until the increasing accuracy of observation and new experiments, reveal the inevitable anomalies that are to be the irritant that generates another new paradigm shift.

During the last fifteen decades there has been an accelerating cascade of change and an overwhelming advance in scientific knowledge and understanding. Almost all the previously accepted paradigms have been changed, often repeatedly, in just a few generations. Nothing like it has ever happened before and though much of the change has been improvement, many people do not feel comfortable with a rate of change that seems dangerously exponential.

Professor J. R. Lucas in his Presidential Address to the British Society For The Philosophy of Science [7.6.93], convincingly establishes the inconvenience for the philosophy of science, of the concept '*things*'. But the approach to understanding and prediction via named or nameable things with separate identities is as old as social, sapient, speaking man. And it has been perfectly successful until too many questions began to be asked about the physics of fundamental particles and of relativistic phenomena.

Professor Lucas points to "…an essential difference between organisms and environment which differentiates all the life sciences from the physical ones." Assuming he means the inanimate environment, that is an important proposition. It is possible to show that the difference is so great that, as he says, reductionist science cannot apply to creatures. He also points to the importance of information theory in this context.

Sir Arthur Eddington suggested that philosophy cannot be independent of epistemology. The product is limited by the tools. Why are *things* with an *identity* so well established in human mental processes. Look at the tool, the philosopher's brain. His ancestors up to six or seven millennia ago were, for a million of years, wandering hunter-gatherers. His brain was evolved for that life until a few millennia ago. What sort of data processing system does a sapient, social, hunter-gatherer with symbolic communication need?

[Throughout these chapters the words 'man', 'mankind', are short for *homo sapiens*], both, indeed all, genders.

Where in Nature is man found? Mankind is a set of 'things' in a narrow environmental slot between the *macrocosm* of galaxies and the *microcosm* of fundamental particles. Let us call this central realm the '*mesocosm*'. Mankind was evolved to deal with solid durable *things*, objects, in that realm. 'Things' are made up of trillions of microcosmic entities, stable combinations of fundamental particles. So the statistical laws of large numbers apply to almost all that early Mankind knew of life on earth. The statistical behaviour of *mesocosmic* entities [those in the immediate world around us] seem quite causal, determinate. It has been safe for him to assume strict determinism. But it has been an act of faith. How could it be proved or falsified? Boyle's Law is beautifully predictive though it deals with a chaos of interacting molecules of whose separate paths we need to know nothing. Trillions of uncertainties make statistical certainties and these are what we need to predict and live in this environment.

To anticipate danger and threat mankind had to direct attention to the features of what he observed that were predictable, durable and apparently deterministic, to *things* that had identity and were nameable, that could be handled by a kilogram brain. His brain had to select and model the durable, and predictable, it need not be sensitive to the ephemeral, indeterminate, contingent and unpredictable.

It is necessary to filter out such useless and confusing sense data. Professor Lucas put it like this, "*So if we want permanence, we shall be led to focus on certain general features, certain types of boundary conditions, which can persist over reasonable stretches of time.*" That was the requirement of the brain of man and that is how it evolved. He did not observe the world but that about it which helped survival. Heraclitas was nearly right when he told us 'All is flux', but not quite. We have brains designed to filter out from the mass of flux and manage with what predictable, durable stability remains.

It ought to be no surprise that a brain with that design has difficulty and finds contradictions in the relativistic phenomena of the macrocosm and the quantum phenomena of the microcosm. His brain was not evolved to deal with those realms. What is astonishing is how *well* it copes.

The brain of mankind was designed for a realm even more limited than that of the *mesocosm* [where 'things' are between the size of mountain and that of a dust particle]. The solid, liquid and gaseous environment is highly specific and unusual. Conditions such as the force of gravity, the range of temperatures, the protective atmosphere, and many other features are narrowly constrained and specific, durable and stable compared with most of the Universe. Man lives in a very rare and narrow environmental slot. Most of the Universe is of very hot dense stars or scattered particles, gas and dust. Science has done incredibly well to produce symbolic mental representations that make sense of these radically different realms. We may not be able to resolve the contradictions we find, when we started to ask these questions, until we look at our tool, mind.

Man himself is a dynamically stable system. He needs to detect and react to threats and opportunities, if he is to remain alive, to retain his general form through time, his this-i-ness as Lucas puts it. So he has evolved a perceptual input array [call it a sensorium] influencing a motor output array [his motorium, the full set of his million effectors].

In between the sensorium and the motorium mankind has a brain to filter, classify, codify, record, and process his sensory inputs and transduce them into suitable instructional outputs to organs and muscle fibres. 'Suitable' means appropriate to the survival of the individual, tribe and species. Survival means retention of a creature's form through time. [The paradoxical exception is evolution where the form changes to an even more change resisting variant].

Man's sensorium does not perceive things, it perceives *events* which are extremely complex, muddled and confusing. An event is the total momentary input to 10^8 sensors. Man has not, like most other creatures, a hardwired database adjusted to the environmental slot for which he evolved. He evolved towards versatility, the ability to adapt the individual rather than the species. So he adapts to all environmental slots. To do this he must be born ignorant and

learn in infancy to perform taxonomy, to choose and classify useful features and abstracts from events which are repetitive and predictable within some tolerance limit, often very wide. Mankind also has to be born uncomprehending and incompetent so that he can be free of wired-in obstruction to learning a great amount of cultural information and skills. He has to be able to adjust to Eskimo culture or to that of the Bushmen. He must be free to be an acrobat, a physicist, a conjurer or a professional golfer.

He learns in babyhood to screen out almost all of each event and select features and elements of them that may be relevant to his survival; Things and behaviour. Each thing has an infinite array of possible sensory presentations but they all have recognisable features with more limited variety that he can learn to recognise and label as a cluster. Such repeating features of events, things, are classified and codified in his brain as percepts even though each is unique perceptually.

In mankind alone among mammals each percept or thing is mentally attached to a symbolic counterpart, a noun, for the purpose of social communication. A brain cannot deal with an infinity of unique perceptual inputs except by such classification. He has to know both the classes of Things and the fuzziness of each class. And to deal with the fuzz and variety of his perception of Things and events he uses qualifying words such as adjectives and adverbs. This is how he perceives Professor Lucas's this-i-ness of Things. That means ignoring noise, inaccuracies.

As mentioned above he can only deal with repetitive features of events and combinations of them. Truly contingent features and aspects of events cannot be classified and represented symbolically so they cannot be learned because early learning requires repetition to filter percepts from the fuzz around them. In Information theory Shannon called the meaningless features of events 'noise', the useless and obstructive sensory inputs which have to be filtered out. Our brains are good at that. They have to be, most of perception is noise.

Thus earlier man's brain was designed to detect *Things* because that is all brain can manage. True *Things* are Inanimate. They are largely solid, durable, and have identity. They are therefore nameable objects, stones, hills, valleys, sticks, bones, artefacts, things the hunter-gatherer tribesman had to deal with. Liquids and gases achieve identity and shape from their solid containers and their colour and taste. Lakes are defined by their shores. Gases are imperceptible except by smell and pressure as in wind. It was a long time before they were seen to be substances also.

Creatures are quite different from the inanimate Things, but for the tribesman's purpose the same informational method and mental machinery will serve. Animals, people, plants though in a constant flux of internal change are sufficiently stable and durable in outward appearance to be named, identified and dealt with symbolically, verbally. They have adequate this-i-ness.

But creatures are not really things, they are not substance but form. A stone or a grain of sand is a durable set of molecules taking a certain shape. A living plant or animal is an ever-changing succession of sets of different molecules that remain in the same mutual relationship. The set is like a fountain or a

stream, the durable, therefore nameable form of material substrate of new water. It has no identity of water only that of shape, form, arrangement.

There is a Shona tribal saying, "*You never step into the same river twice*". That tribesman was a physicist. In the same sense you never encounter another person or even yourself, twice. Physically you cannot, biologically you can. The durability is utterly other than that of Things. The questions that physics and biology may ask are quite different.

The physical laws of nature in the mesocosm are statistical and concern very large numbers of equipotent particles [those which have equal effect in their interactions]. Creatures have equally large numbers of particles but they are formed into components that are not equipotent. The neologism 'heteropotent' applies. Creatures are made of long chains and arrays of elaborately arranged molecular and larger components that have highly specific and have quite unequal interactions. The laws of mass action do not apply in the least. Corpuscular theories about creatures make nonsense. You have to look at them like machines. To understand them you have to identify and understand components, signalling systems, sensory and motor. You have to have divined their vastly complex multiple interactions. If you apply heat to a stone you can predict the rise in temperature. Heat a man in the sun and his temperature remains constant.

Animals cannot afford to be determinate, they feed on and are food for other life forms. The best strategy for both predator and prey is to be an unpredictable predictor. For animals, determinism is death.

Parmenides proposed the Theory of Forms, and Plato accorded the central reality to them. For Plato, Things were imperfect examples of Forms. But his Forms were not immanent and objective, they were the learned, chosen, classified, subjective, symbolic representation of features of experience his brain found useful and reliable. They were mental models of the external, variable objective features of events themselves, models that never matched exactly.

He was right in the sense that the forms are all we can know. But he was wrong in thinking they were primary, immutable and perfect. It was the forms which were imperfect, provisional and, like science, subject to constant amendment as we try to improve the match between them and what we perceive. The subjective models are approximate, merely good enough for some purpose.

To understand the enormity of the difference between the biological and the physical we need to look at another phenomenon, information. Information is the basis of all life. What is it? Its source is a genome, a signal string. It functions via strings or arrays of signals. The nature of a signal in creatures is that of a trigger. It is highly heteropotent, catalytic. It is a microcosmic chemical or electrical change in a variable that produces changes or strings or arrays of changes that can have a disproportionate macrocosmic physical effect. A few photons to the retina of a buffalo may stampede a herd.

Life is designed to defeat statistical laws like the second law of thermodynamics for instance. Mesocosmic physics are the laws of the most probable. Biology is

the art of constraining ignorance about unpredictable systems. A later chapter will elaborate this point.

THE SOURCE OF BIOLOGICAL UNCERTAINTY?

Experiment has shown that the brain of man can optically detect and report a few photons. So at the limit, sensor cells are affected by microcosmic phenomena that are quantum-indeterminate. When a photon can trigger the behaviour of a neuron and a neuron can affect a whole brain the amplification factor is vast. Microcosmic uncertainty has the power to overcome statistical laws and affects the mesocosm in sensitive, catalytic devices like cells and brains.

This suggests that creatures can draw on quantum uncertainty, vastly to increase the variety of indeterminate possibilities and outcomes in the mesocosm. Creatures, especially man, could do this using the long highly amplified, triggery signal trains and arrays in cells. Thus the door is open for an enormous extension of the range of variants, choices, from which evolution on one hand and thinking and learning on the other can evolve or select.

This can explain the paradox of free will. The developed brain may use a quantum roulette wheel to generate an infinite array of hypothetical options for consideration and rejection or choice. Human and animal behaviour is not always rational, it is often whimsical, where that means a less than rationally optimum behavioural option is chosen. The social advantage of whimsicality is the same as the biological advantage of mutation. It gets unlikely options tried. A very small proportion of these options prove to unexpectedly advantageous. Whimsy may be a microcosmic event.

The most primitive form of information is the genome, a string of four simple molecular units from an alphabet, [in man, written in this alphabet, is a 'message' four billion 'letters' long]. It is living, a single molecule, its great stability and durability is dynamic, a stable form on a constantly changing, constantly repaired substrate of molecular details. It is a mesocosmically insignificant mass, but a single molecule can trigger the creation of a man or a whale. It is a monstrously heteropotent molecule.

THE HOMOLOGY OF INFORMATION SYSTEMS

Here is something that trans-disciplinary generalists might better see. Specialist scientists seem to have taken something odd for granted. When Crick and Watson found the pattern of the genome and enabled us to look at the very source of our creation in a speck of nucleus, in a microdot cell, they revealed an information system astonishingly like that we have found most convenient for social communication.

Long before men learned what the genetic pattern was, they reinvented it. Speech, writing, and the random access memory that holds information in a computer, all take an unexpected simple form, a one-dimensional string or vector. Complex information is stored as a very long linear array of coded units. The units, symbols, alphabets, are in all cases highly specified units of low variety, nucleotide base pairs, sound phonemes, letters and ASCII symbol

strings. These again are broken into paragraphs, sentences, phrases, words and so on. The genome precognitively plagiarised literacy and computer technology.

'In the beginning was the word.'

The important point is that the information systems that we find at various levels may be found to throw light on one another, that these systems are more than simply homologous, they belong to the same class of phenomenon, information. They have similarities, shared invariance and limitations. We may learn about all by examining those that are easy to probe. We may find models where the explorable will help us to comprehend the unexplorable, even though it works in a different way at a different level.

As has been said, a creature is not substance, it is the form taken and maintained upon it. The reader is not a thing made of molecules but the durable form taken by a continuing series of sets of different molecules that take up the same ordered form. What, against all probability, maintains that form? Genome information strings. But the Information strings themselves have the same nature. The genome is information, form not substance; the radicals that make up the genome are individually ephemeral, constantly replaced. Biological information strings are like human information mono-dimensional strings. They are infinitely replicable and are independent of their material substrate.

Human information strings and arrays can be transduced from one substrate system to any other, from speech to writing, to ASCII or any other code on many different substrate media. They can be vibrations in solids, liquids or gases, magnetic or static differences on tape, optical signals in fibres, electrical signals along wires or electromagnetic waves in space. And they can be potential differences moving along nerves or showers of neuroproteins at a synapse. None of these differences of media or substrate make any difference to what is essential about information, its effect on some specially designed receiver. Signals are nothing like Things. They are independent of substance and physics. Signal strings have no identity. Any one of millions of replicates has the same effect on the appropriate receiving system as any other. Information strings can travel at the speed of light. If they had mass it would be infinite.

THE CONTINUUM OF LIFE AND INTELLIGENCE

Reconsider intelligence. The word is from the Latin *intelligence*, to choose. 'Optimising choice' is shorthand for the basic function of Life. Only living things choose. Intelligence is a biological phenomenon, an aspect of life. The word intelligence is used in this extended sense to establish that it applies to all that lives.

The simplest bacterium makes intelligent choices, it senses the surrounding molecules, and chooses, engulfs or absorbs by infusion exactly the right molecules for its survival. It also excretes the molecules that may disturb its form or that are not further needed and moves in response to sense signals.

Wherever there is behaviour, or activity, that comes within the meaning 'optimising behaviour in the light of information', we can see a form of intelligence. But if we look for beginnings we do so in vain. We find no starting points or borderlines. We see a continuum, a smooth step-less transition from that vastly preponderant mass of lifeless matter which Karl Jung called the Pleroma, the Things, the inanimate material universe, to that infinitesimal smudge in its great array, Jung's Creatura, the Living, the Creatures.

And there is a similar spectrum across the Creatures themselves. Just as there is a smooth slide up to Life there is a smooth slide up to human or cognitive intelligence. We see intelligent behaviour in animals, some intelligent learning in insects. We even see plants choosing which molecules to bind to and which to avoid. We see an optimising, choice-like process in the evolution of life itself. We find no borders where we may pause and say 'Beyond be Creatures' or 'Beyond be Thinkers'.

This book is not concerned with psychometric intelligence, being scalar human cognitive ability. It is a much broader meaning for the concept that is being considered. From the above it is clear that the meaning used for 'intelligence' is not just human intelligence. By broadening the concept, seeing the rest of the continuum, it can be best understood.

Scalar intelligence applies only to humans. It varies positively with the amount, variety and complexity of the informational input, the storage capacity, the number and complexity of the goals, the degree of optimisation achieved therein, the amount and complexity of the output instructions involved in the response and the position on a scale strategic/tactical of the activity involved. These elements are not equiponderous. The last carries the heaviest weight.

There are numerous disputed definitions of intelligence. Most of these attempt to confine the meaning to the small ambit of psychometrics, intelligence testing. A much broader meaning is useful. '*Intelligence is the data-processing activity of entities which respond to information with behaviour that appears to be intended to be optimal with respect to pre-set goals*'. Be it noted that 'information' includes both that which is immediately sensed, and that which is stored, as well as the sensory and the instructional signal arrays.

The weakness of definition as a comprehension tool is revealed. This one calls immediately for several further contextual definitions: data, process, optimal, goal, etc. The useful choice among the range of meanings for all these will appear as our theme develops.

LIFE KNOWS, CHOOSES

Intelligence involves knowing and choosing. Behaviour that is to be labelled as intelligent must have an information store, data bank or memory, an informational input sensing system and instructional action array. The intelligent entity, cell, animal, social assembly, artefact, must have received some coded information about the state of the world, for which it is prepared, and it must have behaved appropriately; acted in some way that it would not have done without that information and without the information store, genome, memory or

other forms of internal data bank. A living entity has choice and chooses on the basis of information.

COGNITIVE KNOWING IS MODEL MAKING

For humans there are two kinds of knowing: knowing like a good billiard player and knowing like Newton. The first sort of knowing, psychomotor knowing, involves the same field, simple mechanics. The laws of motion, acceleration, impact and friction must be understood in some sense by the player but he cannot easily formulate or verbalise them, they are not socially transmissible. Newton created a model, a paradigm, a symbol system, which was socially transmissible and which could therefore be applied more widely. The knowing jumped off the billiard table and exploded into engineering, technology and space travel. Yet the data processing problem in dealing with the laws of mechanics is minute relative to that involved in the billiard player's skill. Robot spacecraft are here to stay, a robot billiard player is not yet feasible.

Intelligence in the extended sense given here is manifest at the level of the cell, again at the level of the animal and thirdly at the social level. Organised communities of animals and people behave intelligently as a whole. But before we consider conceptual or socially communicable intelligence in ourselves we should look at what we usually see as a lower level intelligence, psychomotor intelligence, like that seen in learning skills, tennis, golf, acrobatics, arts and other skills.

Since we have begun to think about making thinking artefacts, artificial intelligence, we have begun to realise there is a lot more to the tacit psychomotor knowing than we thought. In human beings, for instance, an astronomical amount of data has to be received at the 180 million receptor neurons in each retina of the eye. It must be sorted, filtered, processed and transduced into messages to a million or so muscle fibres.

All this happens several times a second. Seen this way it is the psychomotor understanding which is the higher, more miraculous achievement than the comprehension and use of the simple Newtonian equations. There are lots of creatures that excel man in psychomotor intelligence. Watch any bird. Yet we feel that human intelligence, conceptual, transmissible, symbolic knowing, is more important. We are right and we need to explain why we are. We need to revise our models of knowing. Knowing is the business of making useful internal models of the world, symbolic representations that enable us to anticipate and plan.

PERCEPTION, TAXONOMY, CLASSIFICATION

The primary task of an advanced intelligence is perception. The next task appears to be (but is not), classification, putting inputs into classes that require different responses. With the species called 'primitive' the classes of input pattern are hard wired, built in. But even with them there is a process that must have preceded classification, taxonomy, finding the useful way to classify observations, the perceptual inputs or parts of them. In simple creatures, taxonomy must have been acquired genetically, by evolution. There is capital

punishment for all evolutionary mistakes. Mammals are not locked into a hard wired genotype they can adjust by learning, pain and pleasure.

Work on computer pattern recognition is transforming ideas about this. Traditional philosophers and psychologists have taken the perception of complex patterns and objects as given, the starting point of their discipline. We now know that the mere primary act of observation is a highly complex skill that has, in human beings, to be acquired as a baby by trial and error. The mere act of seeing or hearing meaningfully is impossible to the unprepared perceiver. The baby's brain has to be trained to see the patterns that are useful to it.

It is only a highly selected special version of reality that can penetrate the complex, intricate multi-stage sorting, abstracting, filtering and refining network of the brain and get through to the consciousness. We are all drowned in a torrent of data, an incessant cascade of terabytes from many millions of inputs. The primary problem for aspiring comprehenders like babies is taxonomy (I use the word in an extended sense), how to classify experience, how to abstract classify and code the patterns of sensory inputs which help mind or brain to simplify, sort and order data. The baby has to find the most useful, brainable (brain compatible) understanding. The survival tools are predictive, subjective models of the world.

The job is filtering, rejecting the irrelevant without knowing what it is, trying and scrapping models. We have only about one and a half kilograms of the complex spaghetti, the grey fibrous felt called brain, to work with and there is an awful lot of universe out there to be simulated or modelled. What is learned from attempts to simulate simple perceptual and learning skills on today's linear, deterministic (Von Neumann), computers makes it necessary to modify the philosophy of science.

Like the animal, the scientist, too, can now be seen to be swamped with data. His main job is to do systematically, in groups, with instrumental aids, what the individual does in perceiving and predicting his world. Science is sorting needles of relevance from haystacks of facts, condensing, refining, abstracting and then checking the data for congruence with a succession of models to find the most useful. Relevance is the property of being potentially useful to some entity, individual or social.

We now know that messages cannot be sent to an unprepared brain and (from Kuhn) that facts cannot be perceived by unprepared scientists. We have to start from where we are. New understanding has to key into and start from previous understanding. When we have a sudden insight we often feel we have something wholly new. But if we look at it we shall find that it is not; it is a new arrangement of the old. We can perceive new constellations but they are of the old stars. There is an evolutionary aspect to knowledge, as well as the revolutionary one Kuhn suggested. Knowledge is constrained to relate not only to the new data and the new arrangement of it, it has to follow from the track of its own development in the past, like animals and all other life forms. As mentioned above, Eddington warned us that knowledge cannot be cleaned of its epistemological connections, it is designed as much to suit the mind of the knower as it is to suit the data known. We can further say that like a biological

species it must be constrained by its former state and its line of development. There may be sudden leaps but there is a thread of continuity.

ARGUING FROM ANALOGIES

This brings us to the special problem of communication. How to start? Which of the known thought-tracks is the best for our purpose?

In the 'fifties the fashion turned against arguing from analogy, because all analogies breakdown somewhere. No-one explained that there is no other way to argue. To demonstrate something new to B, A has to find the system in B's mind into which the new pattern fits with least anomaly. Analogies never became respectable again; we followed the usual path of reform by euphemism and began to talk of models. Analogies are dead, long live models. We are all the time making and testing models of reality, whatever that is. Seeing an analogy or model is an act of taxonomy, the creation of a new class or set of cohering classes.

Astronomers arranged a model sun and planets in an orrery, because Nicolas Copernicus had conceived the new class 'Copernican planetary systems', of which the solar system and an orrery are both members. In seeing the relationship we have created the new super class, that of systems which behave like that. What is happening in a brain or mind is the building, testing and constant inductive amendment of models of interesting bits of the world. For this purpose a model can be seen as an understood system of some kind that behaves congruently with an external system whose behaviour we need to understand or predict.

GOOD SCIENCE ABOLISHES MIRACLES.

People before the age of science needed the concept 'miracle' to fill in hiatuses where prediction and expectation failed by too wide a margin. The world was governed by gods and demons and anything could happen. Novel, unfamiliar events were unexplained, therefore miracles.

A more recent view would say that a miracle is an extremely improbable event, one which is grossly contrary to experience and expectations. Kuhn's anomalies are miracles, observations which do not fit in with preconceptions. A paradigm, hitherto successful, ceases to serve because in the light of fresh data it challenges credulity. A new paradigm is needed to demote miracles, to reduce them to the status of explicable, expected events.

Where we see miracles we should look for another mental stance, a paradigm shift. Let us look at some miracles that are still left in the world of science. Where is our credulity under most stress?

Miracle One

A friend who is the editor of a biology journal once told me that almost every post brings a paper from some biologist who is disturbed by the astronomic improbability involved in the origin of life, the assembly of the original genome.

Even the simplest living thing, a virus, has a genome, a program, instruction set or prescription, in the form of a string of information consisting of 1300 np (nucleotide pairs) letters from the four-letter genetic alphabet).

This genome is an immensely long, fragile, aperiodic crystal, as Irwin Schroedinger called it, (in his book, '*What is Life?*') The crystal is highly unstable. It has to remain in a very narrow, improbable range of environments and survives only in a temperature slot where its fragile components are under a heavy bombardment by thermal agitation of the molecules. It needs constant repair by wandering enzymes. Yet it is practically immortal.

The probability that such a thing could rattle together from the most favourable environment by chance is so absurdly low that there is a better chance of drawing the right ticket from a hat where the ticket is one particular fundamental particle and the hat is the whole universe of particles. There are around 10^{80} of them in all.

Some biologists have given up as scientists by suggesting that the origin of life is so improbable as to be truly miraculous. The story of divine creation seems as likely as any other origin for life to those who look at the figures that way. Others, like Frances Crick, who first deciphered the genetic code, push the miracle further away by relying on the concept of panspermia. They think the earth was 'seeded' with life from elsewhere.

Further, the paradoxes and conceptual difficulties involved in epigenesis (the science of the development of organisms) are much like the contradictions that began to arise in physics when Albert Einstein, Max Planck, Paul Dirac and Werner Karl Heisenberg came to the field. We do not really understand how a simple linear string code, even if it is billions of characters long, can prescribe organisms as complex as an Einstein, a Planck, a Dirac, or a Heisenberg!

Miracle one is at least partially abolished by the reflection of Dawkins, that mutation and heredity between them allow the Creatura to climb mountains of improbability.

Miracle Two

The miracle of brain. Very detailed, tactical understanding of the function of brain elements, nerves, has been acquired. But no-one has the remotest notion of its functioning at the overall strategic level. More investigation has made perception, taxonomy, classification, memory and cognition seem more, not less, miraculous. Consciousness, mind, memory, attention, learning, sleep, and many of the processes of cognition and data recall are far beyond the realm of explanation, modelling in physical terms or comprehension, as things stand today. The approach we examine can guide the offer of another sort of plausible conjectures as we shall see in later chapters.

We have no credible hypothesis as to how the input from a hundred million receptors is processed and transduced into the appropriate instructions received by a million motor neurones. It is a miracle that has been tucked out of sight in a 'black box'. Attempts to simulate even simple psychomotor-type intelligence artificially have only brought out the enormity of the problems. I

suggest a hypothesis. Or perhaps, as we shall see, this miracle too will surrender to probability mountaineering.

Miracle Three

We have very little idea how complex societies work. Professor Stanislav Andreski is just one frank sociologist who has admitted that sociology is not yet a predictive science (*The Social Sciences as Sorcery*, Pelican). Increasingly we see that economics is largely informed guesswork. If it were a predictive science, economists would own the world.

In '*The Wealth of Nations*' Adam Smith created a successful paradigm for the working of an economic system governed by the concept of the free market. He saw it as kind of autonomous 'engine' (today he would have said computer). His system was not a predictive one. It predicted unpredictability. The model was autonomous and unpredictable in detail since it was driven by billions of unknowable free contracts. However, it remains much more predictive and explanatory than any other in the field, which is saying very little. None-the-less it has informed the work of the most regarded pundits in the field, not excluding Marx.

Recently, determinist theories of economics have been abandoned and there has been a return to Adam Smith's ideas, through thinkers such as Ludwig Von Mises, Friederich Hayek, and the whole school of followers. But Adam Smith described rather than explained the working of the market. From the self-regarding actions of millions of individuals comes a better pattern of the adjustment of supplies to needs and demands than by any other method. Smith had to resort to the idea of a miraculous 'hidden hand' that ensured that the polycentric pooled selfishness of the market produced more general satisfaction than monocentric patrician altruism, as in aristocratic and Leninist communist systems. More recently, Hayek has looked to theories of information for a miracle-abolishing paradigm in economics.

In the book, '*Brain*', (1976), the author proposed a similar idea that is developed in later chapters. Hidden hand, black box, ignorance, anomaly. We need new paradigms. You shall see them here. The third miracle begins to be on unfirm ground.

Miracle Four

Understanding of the universe vastly expanded and improved in the last century and the most important contributions to that improvement were the basic paradigms of modern science. Central were those of causality and determinism, the underlying belief that the universe is an ordered system with comprehensible and unbreakable causal laws. This led to the obstinate disciplined search for such laws and the rejection of the easier options, mysticism, animism, divinity and other such dogma barriers to improved understanding. Determinism and strict causality became the permitted dogmas of the scientist, belief in them their one act of faith. But other philosophers were less happy. However, the incompatibility of strict determinism with the daily human experience of options and free will to choose between them, remained unresolved. If this were an illusion life and science would make no sense.

Many scientists had difficulty with the contradiction between the determinist dogma and their experience of freedom of action and learning. How can any entity learn in a universe which is pre-ordained, which is unrolling, frame-by-frame like a film show? How can any creature whose least action arises from inflexible causes create a working model of things that are remote from it, even of itself? How can a celluloid character in a film create an image of the camera and the photographer?

Here is a paradox. An entity proves that the world is determinist, and by the *way* it did, and the *fact* that it did, shows that it cannot be. The paradigms proposed here indicate that I have options and choices between them need not be false and that my options need be neither illusory nor miraculous. But when we consider that the quantum world is indeterminate and that quantum events marginally affect brains Miracle Four begins to look shakeable.

Miracle Five

Except as a cult-figure, the unfrocked priest, Teilhard de Chardin, carries little weight today but he has pointed to a miracle. He described the overall tendency of evolution, divergence and the competitive elaboration of variety of the species. All was explained by Darwinian theory backed up by modern genetics.

But he also noted a contradictory convergence and symbiotic combination at a higher level. He was puzzled by the apparent symbiotic stability of the biosphere and especially the emergent coalescence of human societies. All these were contrary to expectations from evolutionary theory, where evolutionary divergence is the trend. It is this anomaly with which the newer discipline of ecology seems to be concerned. He also spoke of the 'noosphere' the global knowledge network, as a supervening world of knowledge uniting mankind, which unexpectedly exists and is not accommodated by evolutionary theory.

'*World III*' elaborated by Karl Popper and John C. Eccles is the central database or World Culture that a consensus of educated humanity accepts. Herbert George Wells had a similar thought in his '*World Brain*', while Bernard Shaw's '*Life force*' was similar. There is the inexplicable, emergent phenomenon, our present world culture, which has miraculously appeared within a few hundred years from the informational coalescence of the societies of a species that was bred for a million years to be a network of mutually repelling hunter-gatherer tribes. These are the anomalies that are tentatively addressed in these speculations. An epistemological hypothesis, one of mentality, of brain function, is advanced which it is hoped will advance the understanding of the most important aspect of the universe, understanding itself.

2. ON LIFE.

Imagine a very durable extraterrestrial intelligence that visits Earth once every million years. What changes would it note on successive inspections?

The inspections of the geological and geographical details would cause the intelligence to say with Heraclitas 'Everything flows, nothing abides'. Looking at what we ephemeral creatures see as the durable things, the rocks, stones, mountains, valleys, coastlines, rivers and streams, it would find little that was unchanged between visits. All would be grossly changed.

But if he chanced to look again at certain soft, unstable, ephemeral, moving, constantly changing, excrescences that he finds down in the mud, dust and confusion on the surface of the planet; if he looked at the moving and growing creatures that had been clinging to, swimming, flitting, wriggling, creeping, dodging and just growing on and under the surface of the changing rocks or in the restless waters; if he looked again at the quick, the Creatures; what then? He would find an anomaly, a contradiction, a miracle. Peering again, however closely, at a stickleback, a bee, a pine tree, or the pattern on the wing of a butterfly, he would find that within close margins they had not changed at all, despite scores of thousands of generations of replication. Absurdly, the stable would have proved to be unstable, only the changeable would be unchanged.

About such a broad term as 'life' there is little general agreement among scientists but there is some common ground. They are agreed that there is no agreement about a satisfactory definition. Even the august pages of the *Encyclopedia Britannica* make this complaint (much like 'intelligence', Life is undefined).

Various manifest functions, ageing, metabolising, feeding, excreting and reproducing are mentioned as aspects of life but objections can be raised to any definition in such terms because things that are not alive have some of these aspects, and some things that are, have not. Some things that are undoubtedly living, such as some seeds, remain dormant and without metabolic activity until activated. The tobacco mosaic virus can exist in an immutable crystal form, but applied to a leaf the crystal becomes active and reproductive. Life remains, like many real and useful concepts, perceivable, understandable, understood, but undefined.

The Life Miracle

It is unscientific to use the word 'miracle', which is merely the word for an event so very improbable as to set us seeking new paradigms. Dobzhansky estimated the probability that anything as miraculously unexpected as a genome should appear on this earth. He calculates that the coming together of the human genome with its 3.5×10^9 nucleotide pairs and its 3×10^4 genes is rather unlikely. If any odds are low enough to merit the term miracle these are. He puts them at one in 10^{10000} (plus or minus a few hundred orders of magnitude, I dare say).

Now the distinction between the appearance of something so improbable and a miracle might be hard to establish, so we must not blame Jaques Monod when in the light of similar reflections he calmly talks of the 'miracle' of life. The need to labour the point will emerge. Monod seems to be right. It seems to be true that there are features of living systems that are incompatible with the very methods of science. He feels we need a new paradigm and proposes 'teleonomy', a neologism for an internal autonomous determinism that, he posits, guarantees the invariance of living things.

Karl Popper and John Eccles, in '*The Self and its Brain*', made a similar point when Popper says, 'It is incredibly improbable that life ever emerged but it did emerge. Since it is incredibly improbable, there cannot be an explanation to say how it emerged in probabilistic terms because an explanation in probabilistic terms is always in terms of high probability.' Any explanation based on a very low probability simply is not an explanation in the normal meaning of the term.

We must start from where we are. There are over a million species of Creatures on the earth and each of them is so improbable as to seem unexplainable, miraculous in the normal meaning of the term. Let us make an attempt to abolish this formidable miracle if we can. Let us start boldly with an attempt to define what has been held as impossible to define.

Morphostasis (Formholding)

In my early book '*Brain*', the concept 'morphostasis' was coined. It is defined as the dynamic form-holding exhibited by the Creatures. It was seen as the fundamental essence of Jung's Creatura, the Creatures as distinct from the Things.

Morphostasis is the process by which the form taken by a physical substrate of constantly changing molecules behaves as though it, the immaterial form, ever replenished with new substance, has the purpose to preserve itself through time with minimum change. The continuing essence of a creature is not its material contents, but the immaterial form on an ever-renewed chemical substrate that continues improbably for many years in some cases. Further the replicates, the progeny of morphostats, can continue for millions of years.

Every creature has, as Waddington has said, its basic continuing shape, form, arrangement. This is completely true of single celled creatures, prokaryotes [bacteria, algae]. With Eukaryotes, many celled creatures, to be more accurate, we should say, "The adult form of every creature has its basic continuing shape…" Larger creatures like metazoa [with many cells] usually have, as they enlarge by feeding to reach their final form, to pass through a number of size stages where the adult form would not be viable. Their development takes them through a series of different forms each fitting a different environmental slot. The extreme case is the insect that survives as an egg, a grub, a pupa and finally the adult insect, all grossly different forms. But the succession of those forms is repeated within very narrow margins of variance over the generations. The one way in which the quality of a living thing is distinguishable from that of a river is this; living things are forms, patterns, which behave as if they had the purpose of preserving their adult form with minimum change. Actively and improbably they are homeostatic.

Homeostasis (the tendency for the internal environment of the body to remain constant) is an aspect of all life. All living things, morphostats, are homeostatic on many parameters of possible variability. They are actively homeostatic, using feedback signals and corrective action. They are not stable passively, like a river or a fountain. They strive actively, using trigger like information, to retain their form on a succession of material substrates.

Richard Dawkins, a founder of sociobiology, in his book mentioned above, '*The Selfish Gene*,' says, "Darwin's survival of the fittest is really a special case of the more general law, the survival of the stable". It was just such a thought in the fifties that gave rise to the series of speculations and enquiries that led to what you read.

Life forms, all of them, seem actively to seek selfishly to preserve their form with minimum change through time. They can do this only because they carry coded information stores forward with them through time. All life forms, without exception have some sort of memory, a data bank, a coded, symbolic, representational copy of their survival form. These, technically, are called their phylogenic or species memory. From that is derived their ontogenic or survival and replication form.

Change and Stability

The idea that stability, non-change, is the distinguishing characteristic of life is counter-intuitive at first glance. The living, the quick, the animate, seem to be in a state of constant, restless flux. We see the inanimate as unchanging. But the parable of the infrequent earth visitor warns that things are not that simple.

What the Creatures are restlessly reaching, growing and scurrying about for is to *ensure* the persistence of their particular form through the aeons. My body spends its time transmuting bread, meat, grain and vegetables into ever new physical I things, all of which are called Serebriakoff; all sadly, in the same ageing form. The physical object with that name which was at this keyboard last week is gone forever. It is not, I repeat, my molecules, but the *arrangement* of any old set of them, that I am busy to preserve. My obsessive activity to preserve my form cannot be used as evidence against the fact that I am basically a form preserving entity.

The Secret of Life.

The words 'change' and 'stability' are broad and imprecise in meaning. They are blunt tools for such a profound discussion. But if we look at a world with this crude analysis we see things that are not so clear.

We see that the universe cannot help selecting stability. In a contest between stable systems and unstable ones we can predict without risk that there will be a greater proportion of stable ones left at the end of any substantial time than there were at the beginning of it. Among the infinity of possible changes to unstable systems there must be some towards greater stability and some towards less, changes that would make the object something else or disperse it. But systems that change towards being more stable will become more frequent.

Thus imperfectly stable systems tend to become more stable. In any late sample there will be more of those that have changed towards more stability than there will be in any earlier sample. Those that changed away from stability will be unrecognisable, will have moved beyond some class boundaries (wherever they may be set), or they will have become extinct.

This is the solution to the strongest objection to morphostasis as the key to life. Here is the second strongest.

The False Paradox Of Evolutionary Change

Evolutionary change, adaptive and diversifying changes, have been cited as counter-examples to the concept of morphostasis. There are two kinds of adaptive change. Evolution is an intelligent learning process, it is '*genoplastic*' change, that between generations. Some vertebrate learning is '*ontoplastic*' change, learning-type change in an individual animal. Both can be called 'meta-change', a small class of changes which morphostats accept, seem to seek, those towards higher order stability, those that make them *more* morphostatic. Even diversification into sub-species which are fitted to new ecological niches preserves the basic phylum form better and increases its numbers more than if it did not happen.

The class of second-order (meta-morphostatic) changes is usually so small that it has led to the dogmatic but pragmatic saying of geneticists that '*All mutants are lethals*'. This is the experience dealing with drosophila (fruit flies). The thousand or so known mutations have no examples of advantageous ones, but in the rather larger sample of the biosphere evolution has thrown some up now and again.

So a better definition of morphostats to include this idea could be 'morphostats are a class of entities which are organised to minimise irreversible changes in their form other than those which increase their power to resist changes or allow them to spread to new environments'.

Heraclitas, with support from the Second Law of Thermodynamics, was right when he said that everything is in a state of flux. But as to an infinitesimally small proportion of the world, he was in error. Everything does flow, nothing does abide, with one exception, an exception as infinitesimal in size as it is great in importance, at least to us. *Life forms abide*.

Entropy

The introduction of Boltzman and Maxwell's Second Law of Thermodynamics in the last sentence might seem abrupt. A later chapter will explain, meanwhile we may note that when life forms seek to preserve themselves they are in breach of the very heavy statistical weight of the above Second Law.

States that all ordered or organised systems tend to break down, disperse, scatter, become disordered, randomise with time. Entropy, disorder, increases inexorably in every known, closed, physical system. The Creatura have to struggle upstream against a fierce universal current in order to do their difficult business of survival. They have, constantly and incessantly, to make decisions

that enable them to climb entropy slopes, past and against the pressure of everything else that is sliding down them. They disobey the Second Law and mountains of improbability.

The miraculous improbability of the origin of life forms has been discussed. All life forms, even now, are overwhelmingly improbable by the laws of Things, by among others the Second Law of Thermodynamics.

A paradigm text book in this field is Barton's '*Heat*'. He says, "…so we see that the Second Law of Thermodynamics is true, and so very certainly true, merely because we are dealing with systems of very large numbers of individuals. It owes its validity, to this fact... It is definitely not true for a system consisting of a small number of individuals." (By 'individuals' he means molecules.) In any one sample of a substance, solid, liquid or gas, the molecules are equipotent, they have equal mutual effect. The assumption behind the statistics governing the Second Law is that the elements of the system are equipotent. If they are not equipotent the laws do not apply.

But the molecules of living systems are not equipotent, they have unequal effect on each other, there are catalytic and trigger-like or relay-like effects at all levels. The elements of living systems are, we could say, 'heteropotent', of unequal mutual effect. We must not expect the catalytic elements in a system to obey laws that do not apply to them.

So there are systems (albeit immaterial ones, ones of transduced form) that defy the overwhelming weight of probability that establishes the truth of the Second Law. They subsist on substrates with extremely large numbers of molecules. The systems *do* defy the Second Law. Living things do that, all of them. The monstrous accumulated improbability calls urgently for a paradigm shift.

Death Is Change

It takes practice to see the lively as stable and the dead as changing, mutable. Even the pioneer of this field, the towering Professor Norbert Weiner, of '*Cybernetics*' fame committed the error of saying in that book 'The stable state of any living organism is to be dead'. How wrong! The stillness of death is deceptive. It is the beginning of an irreversible runaway change, decay.

Here is a test for the above hypothesis against Weiner's. Take two horses and kill just one. The living one will act as the control. Keep them in a stable or if you do not have one, in a spare room. Supply feed and water to both. Treat them exactly alike. Check up which changes the most over the next few months. We may predict that the living horse will change its form less than the dead one. It works with mice but it is more odorously and assertively convincing with horses.

Death is resumption of change, the final surrender to the Second Law, the reluctant return to what is overwhelmingly normal in this universe. The decay of death is the end of the rebellion against the universal trend to dispersal, randomisation, disorder. It is the return to probability from the highly erratic and utterly unlikely behaviour of those rebels against the Law, the Creatura. However, the cycle of birth, maturity, reproduction, senility and death is

morphostatic. It has a higher order stability. This brings us to replication, another such higher order stability.

Replication

The Creatura have many techniques for perpetuating form or pattern. The two main ones are survival and replication. Morphostats behave 'as though they had the purpose' of preserving form through time. Teleological explanations are rightly suspected. But meaning has to be conveyed by context and this is the best way to put over the meaning. There is no suggestion that a genome has a conscious purpose. It is just that those creatures that 'appear to be purposive' increase in numbers and spread. It is easier to understand the behaviour of Creatures if, as a tool of understanding, we attribute the 'purpose' to survive to the *form* itself rather than to the individual member of a species. (Richard Dawkins is doing just this. The 'Selfish Gene' by implication, is purposive). But the gene is form, not substance (it is a flux of changing atoms and radicals, it is under constant repair). It is a form and a form prescriber, it is information. Richard Dawkins, Edward Wilson and their colleagues have shown clearly that it is 'as though' the gene itself, the form, the ever changing, mutable complex of unstable molecules with its incredibly immutable *shape*, seem to be motivated to endure and multiply.

ON INFORMATION

The concepts behind the two words 'life' and 'information' have a close semantic connection. There is nothing that lives which is without an information system. The genome itself is clearly one. (No known life is without a genome.) Conversely, there is nothing within the semantic envelope of 'information' that has no biological connections. It can also be argued that no information passes without there being some entity, some purpose (in the loose sense). So information systems and living systems are bound together by a teleological aspect also.

MORPHOSTASIS: THE LIMITS OF THE CONCEPT

The similarity of information systems found in tiny cells and those which serve social communities of mankind leads to my next point. Else, why did I need the neologism 'morphostat'?

We see that protozoa, unitary cells and metazoa / metaphyta (co-operating groups of cells, otherwise known as animals and plants) fall within the semantic envelope 'morphostat'. We may ask what other members of the class there are.

The behaviour of groups of animals that live in communicative co-operation is also morphostatic. Such groups, tribes, nations, firms, institutions all have intercommunication systems, and they use and exchange information. They all appear to be purposive, they act to survive, to preserve form through time. They survive not by welcoming but by resisting all but beneficial change (the rare sort). They are all characterised by conservatism. A cell, an animal, a society is a continuity of form through time. It is a change-resisting machine.

The word 'morphostat', so far, does not describe entities that are passively unchanging but the class of all those entities or systems that actively resist change, the set behaves as if the members had the intention of preserving their form as little changed as possible through time. I may legitimately include human institutions, groups, societies and nations within this class without being accused of drawing false analogies.

An objection has been advanced to the paradigm, morphostasis. There are aspects of living systems which seem to contradict the thesis that they behave as if they seek stability. There are risk-taking propensities in animals, such as curiosity, adventure, exploration. They serve the same purpose as 'mutation' in simpler forms of evolution. There are mechanisms such as sexual reproduction, and the gene-shuffling process of crossover which seem like invitations to change. In a special way Creatures seem to seek change. Creatures with these experimental, probing, risk-taking mechanisms and propensities turn out to be more stable, to survive with less change through more adverse environments and traumas than those without them. Evolutionary mutation, recombination, and exploration behaviour turn out to be better strategies to achieve minimum change than over-conservative ones. Optimum stability of form comes from the offer, at second order, of a limited amount of change for selection by the environment.

This factor has arisen quite separately at the various levels of life: mutation, sex and crossover in the cell, as mentioned, curiosity in the animal, exploration, migration and experiment in social collectivities, adventure, research and experiment in human groups. The approach is indirect but the apparent purpose is the same, to avoid the gross change of extinction by accepting smaller changes and keeping those to a minimum. Death is the ultimate change. Life is minimum adaptation.

Morphostats can be distinguished in many ways. They are very different indeed from the rest of the universe. They are so different that as Professor Lucas says, we shall be at a loss to understand them via the physical laws that prevail among the Things of the inanimate material world. There are good reasons why we should formulate two Realms of natural law. One for the Things and another for the Creatures. Those which apply to the immaterial forms are other. The living are governed by information which has laws other than those statistical laws which apply to the inanimate, as we shall see in Chapter 3 on 'Order'.

There are laws which apply across both Realms, that of the Things, substance and that of the Creatures, the Biota.

Morphostasis: Recapitulation

A morphostat is a miraculously improbable platform of un-change, an ordered functional arrangement organised principally to resist changes in its form, structure, pattern, which would normally result from the effect upon it of the forces and entities in its environment. We see 'the miracle' of life. The highly complex, durable, ordered forms appear to be forbidden or made overwhelmingly unlikely by the physical law which commands order to disperse. So we shall be wise to look more closely at the concept 'order' which is

associated with the pattern or form of our paradigm. The next chapter goes more deeply into that concept.

Protozoa, Metazoa, Sociozoa

Now that we have the concept morphostat, what can we do with it? We can seek to clarify aspects and invariance of the class and we may seek new ones. These may help us in understanding and prediction at all levels, chemical, morphological and social. It is clear that the concept morphostat applies to interactive social combinations. There are social morphostats, societies of creatures. These, too, are morphostats acting to preserve their form, their culture, laws, know-how and practices though time. Different people flow into and out of the roles of a society as different molecules flow into and out of the organs of a cell.

A concept that has come forward several times is that of a social 'meta-creature', one that comprises a collectivity of separate co-operating creatures. A metazoon comprises a number of separate, co-operating cells (which resemble protozoa). Ernst Heinrich Haeckle, the recapitulationist philosopher, talked of this concept as a 'corm'. The philosopher Henri Bergson, one of the last of the animists, talked of the beehive community as being more like a simple dispersed animal than like a colony of them.

There are protozoa, single-cell creatures, and metazoa, co-operating, inter-communicating combinations of cells. There is another stage, co-operating, intercommunicating collectivities of metazoa, of animals. We could call these 'sociozoa'.

There is even a creature that sits neatly and squarely on the semantic borderline between a metazoon and a sociozoon. Slime moulds spend part of their time as a colony of tiny, mobile, worm-like creatures that live and feed independently. But to spread their spores they swarm together and merge to form a single, much larger, worm that crawls to a suitable position, turns vegetative and grows like a small plant so as to produce a sporing pod.

The whole language of human organisation and politics is based on the simplistic Roman paradigm of the State made up of parts related to each other like the parts of a living body: Government came from the Capita, the head. The limbs, the organs, obey and co-operate. The Roman Empire was an organisation. It was an immensely powerful and effective paradigm even if it was oversimplified and increasingly inappropriate after a time. It certainly informed and united a very large empire.

The concept 'morphostat' is to make a more sophisticated paradigm to consolidate the apparent unity that underlies these classes, cells, combinations of cells, (animals or plants) and combinations of combination cells, (social assemblies).

It is very clear that many social colonies and communities of creatures behave in many ways like complete animals. We see that such social collectivities have a social method for passing information between members, speech and writing. They can thus co-operate to preserve their form, their social organisation, their

culture and customs (and thus their individual lives). They have invented another kind of heredity, a non-genetic one, to add to the genetic one. Such collectivities fit perfectly within the meaning envelope of the concept 'morphostat'.

If they do, they open the door to improved understanding of the whole class. Sociozoa have the same problem as metazoa and protozoa, that of preserving and replicating their form through time. We may therefore expect to find ways in which their systems echo each other. Since some are easier to probe than others this may improve understanding of all. For instance, the communication systems of cells and brains are complex and very difficult to unravel, but the simple communication systems of animals and those of human beings, being patent, are easier to probe.

3. ON ORDER.

Order has two distinct but connected meanings. Order of arrangement is an epistemological if not psychological phenomenon. It is a description of mentally economical arrangements of similar things. The arrangement of things in a simple, comprehensible pattern is said to be 'order'. Points which lie more closely in a straight line or a circle are seen to be more orderly than randomly scattered points. The basic idea is that that some clearly mental restraint has been imposed upon the elements of the system, as to their position and arrangement.

Or, quite otherwise, 'order' describes a correctly functioning, purposive creature or system. If it is working smoothly it is said to be 'in order'.

Finding order is inducing a set of precepts and concepts, mental models, which enable comprehension and prediction in our sort of brain. Order is the measure of the comprehensibility, of the coherence, of Things with manageable mental tools. Entropy, disorder, the mentally unmanageable, is shown as the other pole of a continuum.

LIFE PARADIGMS.

Science has been built on a base of mechanistic paradigms since Descartes. But science is a biological activity and a social one. It is an interaction between the tool mind and the world. Modern biology, the growing understanding of perception and of information theory offer a new route to comprehension arising from a better understanding of the philosopher's and the scientist's instrument, intelligence of its abilities and the limitations imposed by its origins. My generalist study in the field ever since shows much research that fits, much that confirms and none that falsifies the original conjectures. I have been advised to present the case anew in an updated form.

The style chosen is that of popular science rather that of academe, using the simplest and the least technical language which will serve. With a new paradigm, unavoidably there will be some self-explanatory neologisms.

There are a number of linked themes based on a fundamentally new way of looking at the sciences of Things and of those of Creatures, the 'hard' and 'soft' sciences. A new view may help understanding of the lack of coherence that has emerged as the proliferating disciplines have forked apart into their specialist niches. The fresh paradigms to be set out are separate but interlinked and only partially coherent. Each theme will be dealt with in a separate chapter and calls for a separate judgement.

ON ORDER.

The Creatures have another feature in common they all seem, in their form and function to be highly ordered systems. But the word has a number of meanings that should be clarified. Things too are often ordered but in a quite different way. In chapter two we looked at the Second Law of Thermodynamics in order to understand Creatures and how they differ from Things. The Second Law makes

use of the use concept 'entropy' that can be thought of as a measure of 'disorder'. To understand disorder we need to understand 'order'.

Comprehension here is not made easier by the fact that James Rudolph, Julius Clausius, James Clark-Maxwell and Ludwig Boltzman have conspired to get it upside down in pointing to and naming the negative property disorder, entropy, rather than the positive one, order. It is hard to visualise perfect order as a state of zero entropy. The reverse semantic polarity would have been better. (Benjamin Franklyn got it the wrong way round when he named that charge 'positive' which turned out to be the absence of electrons.) In both cases the mathematics were made to fit but the ideas are less brainable in that form.

Thermodynamics, as well as life, cannot be understood without the concept 'disorder' or entropy. The curious thing about the concept is that although it first obtruded itself into science in the pragmatic business of engineering heat engines, it seems to be almost entirely an epistemological concept. And so is the concept at the other pole of this continuum, 'order'. Certainly it has invaded physics and engineering but to define it objectively we find ourselves in problems.

The *Oxford Dictionary* contains three pages on the word 'order'. The original meaning was 'rank, grade, class,' the pecking-order[1] of people in a community. The closest to the sense we have in mind comes thirteenth among the array of meanings of a very fuzzy word: 'formal disposition of array'; 'regular methodical or harmonious arrangements of the position of things contained in a space or area, or composing any group or body'. These ideas seem a curious basis for such down-to-earth, pragmatic laws as those governing the engineering of heat engines and refrigerators.

Order as constrained arrangement.

Let us look more closely at this simple sense of order of which I wrote above. We speak of order in this sense in connection with a plurality of similar entities or parts thereof. We could not describe an ordered state or a disordered state of a mountain, a dice, a bicycle, a skyscraper and speck of dust. One point in space cannot be ordered. Two points or particles can be ordered or disordered only in respect of their scatter; the closer, the more ordered, as we shall see later.

But the closer a group of three points are to a straight line, or to the vertices of an equilateral triangle the more ordered we should think them. With sets of four points (or similar things) we might consider the set most ordered that lay closest to being on a single line, plane, or to the vertices of a regular tetrahedron. Obviously if measured accurately nothing real can be *perfectly* ordered.

Order as mental constraint.

---

[1] *Note. Many social animals have a strict dominance ranking system that is seen by biologists as advantageous. It was first noticed in chickens by observing which pecks which. What part can entropy or 'decrease in the degree of regularity, method or harmony arrangements' play in universal laws?

In the real world, entities that are defined as belonging to a class or set have a permissible variation, they are always varied, but individually unique. So in this restricted sense of 'order' we appear to impose order as we impose any (and this is vital) mentally comprehensible constraints on positions or arrangements of similar things. 'Similar' means that there are mental constraints on the position, sizes, aspects, qualities and or features of the 'orderly' entities concerned.

Why do we impose order? Because our mental tool, mind, finds order easy to comprehend and deal with. We can cope better in a tidy, ordered environment than in a chaotic one. We love order. Our idea of beauty lies in order. The pattern of a flower is simple, orderly. We love it because we love to comprehend and our brains had to be designed to select simplicity, comprehensibility, constrained variety, as phenomena we were evolved to deal with. Why? Because ordered systems assist prediction and we live by successful prediction. Ordered systems are informationally economical. The description of a square is easier than that of an irregular polygon or a curved surface.

The entities that we have allocated to a named class or set are specified, their variety is constrained but never perfectly, such is reality. Each element and feature of each class member is none-the-less unique. The tighter the specification the more ordered they are seen to be. Points randomly scattered in space are disordered, those confined to a plane are more ordered, those confined to a line more so, those with a regular spacing on the line even more so, and so on. And a set of objects on a line would be more ordered if they were all pebbles than if they were a rock, a pram, a house and a penny. In the real world they can only be so confined within some tolerance, some permissible deviation, some noisy, fuzziness, but the smaller this is, the more order we should see.

Many ordered Things we encounter in civilised life have been put in order. They have been ordered by selection or manufacture, made uniform. They are further ordered by being arranged and maintained in a constrained, ordered array for convenience. So that we know where to find things. The positional co-ordinates within their 'tolerance' are constraints that confer order. Mind selects order for better comprehension.

'Order of arrangement' can be found in Things, that is what the 'hard' scientists do. Atoms are more ordered than free electrons, neutrons and protons because atoms and their components have less variety of relative position and momentum. The protons, neutrons and electrons bound into an atom are under constraints, easier to deal with mentally. They obey Laws that severely limit their diversity of form. Niels Bohr's atom with its constraint on orbits is more ordered and therefore lawful than its conceptual predecessors. Molecules are more ordered than the free atoms. They are easier to understand and predict, more brainable. The brain of man selects order, loves it, as the way to comprehension, prediction, survival.

Why? Straight lines, planes, fixed orbits and distances, all ordered arrangements, are purely mental abstractions, simplifying ideals from which all cases deviate in practice. Order is never perfect, it is an unreal but very helpful mental model. So what business have these purely mental constructs doing in

the science of objective physical systems outside of a mental context? Is there any such thing as order outside of a mental system? We measure the order of a system simply by the amount of its conformity with a mental abstraction that has no objective reality.

Clausius, Boltzman and Maxwell showed that systems of low entropy, ordered ones, were less probable, less likely to be found. If they are subject to constraints, even if they are mental constraints, that makes sense. Adding specification to a Thing or Creature reduces its frequency. There are more stones than black stones and more black stones than round, black stones. By imposing constraints, by increasing the specification of a class or set, we decrease the probability of encountering a member of that class or set.

Order as Durability, a Paradox.

It seems intuitively obvious that we should see a durable arrangement, in the above sense, as more ordered than an ephemeral one, so the idea of constraint does not apply in the time dimension; the less constrained in time, the more orderly the entity. We are more likely to encounter durable Things than ephemeral ones. Decreasing a constraint, allowing more duration, increases frequency and probability. But the Second Law being what it is, durable entities, ordered within given constraints, are less frequently encountered than short-lived ones within the same constraints. So in that respect the rule holds, more order implies lower probability or frequency.

We should also ascribe more order to repetitive events than we should to less repetitive events. And the more iterations and the less variability, the more the order. Again, more order means lower probability, frequency. Repetitive events aid comprehension, they are more learnable.

Now this concept of order as stability begins to touch closely on the living world to which we belong. The Creatures are more ordered in all the above senses than the Things. They are highly specified, (by the genome) within narrow constraints. Each creature is durable, it actively resists any change in its form by feeding, excreting and avoiding damage. The species is much more durable, by replication. That highly specified form, though in constant internal flux, retains its form for millions of years in some cases. Here the replication increases the frequency of the members of the set but decreases the frequency of such sets themselves. Long-lived Creatures keep up their numbers but there are fewer long lived Creatures than there are of those with short lives. With sets of replicates, the larger they are the less frequently are such sets to be found.

Order As Function

I now come to the second usual sense of the word 'order'. At first glance functional order seems quite different from the order of arrangements or arrays as dealt with above. Yet it is intuitively associated with it. This type of order could be defined as the order that arises in entities that are, or appear to be, 'designed' or evolved for some purpose (in Creatures this is usually for survival and replication). Functional order is truly applicable only to biota and their artefacts.

The functional sense of the word 'order' is inescapably teleological. It has no application to non-purposive systems. Functional order indicates the correct and harmonious functioning of the parts of some purposive interactive system. A clock, car, person, committee or your body, is in order, when it and all its parts are functioning correctly to serve some pre-established purpose even if it is only continuance through time.

In previous usages, we have been dealing with replicates and their arrangements. In this sense the order lies in a much more complex idea, that of the mutual 'fitness' or co-operative suitability (related to some purpose) of many differing and interlocking parts of an entity such as a machine. In his excellent book, '*Genetic Takeover*', Graham Cairns-Smith says that orderliness is not a particularly significant aspect of living things. His concept of 'order' is similar to my first usage: things that are ordered are confined to a small subset of the possibilities of arrangement.

He considers a dewdrop the size of the bacterium *E-coli*, which is about to form from vapour. He shows that the gain in thermodynamic 'order' when it condenses is vastly greater than that which could possibly be specified by the 4,500,000 nucleotide pairs of the *E-coli* genome.

But the highly structured bacterium with millions of highly specified interactive parts is vastly more ordered in the functional sense, than the dewdrop. There are an enormous number of molecular arrangements of the water molecules that would still be classed as a dewdrop, it could be rearranged ad infinitum without ceasing to be one. This does not apply to the *E-coli* bacterium, so each of the much smaller set of viable re-arrangements is much less probable and much more orderly. There seems to have been a semantic shift from the type of order that I first described and the functional order I deal with now.

CONCENTRATION AS ORDER

It might be thought that there is no way to unify, to make sense of, these different senses of the meaning of the concept 'order'. The sense used in the Second Law seems more restricted. But is it? Looking at the supposed beginning, the Big Bang, we seek an underlying sense that links the different aspects of the concept 'order'. The classical view of entropy, if applied to recent consensus cosmology, looks odd. In this view the universe exploded out of unknowable void with a Big Bang around 15 billion years ago. For a short period there were no particles, nothing but a vast exploding ball of super-hot plasma, mixed matter-energy without singularities. It is hard to think of such a chaotic explosion of proto matter-energy as being of low entropy, ordered. But extrapolating backwards through time, this beginning ought to have been the first, the most ordered state of the universe, the one from which the ubiquitous running-down process of rising entropy started. If not, how do we explain the decrease of entropy that led to the highly sorted (ordered) array of galaxies and stars?

From this cosmic view we have to include dispersal, scatter, as a measure of disorder, concentration as an aspect of order. And this seems to be the most important of the various aspects of the concept. The running-down and the expansion of the universe are substantially similar, the more space occupied by

a given amount of matter-energy, the less available the energy, the higher the entropy, the more probable the entity or event. We could fairly designate the Big Bang as an infrequent and infinitely improbable event. Which fits.

Order And Information

The information required to describe the universe after the big bang increases continuously with time from the first femtosecond. The original plasma could be described by simple equations of a few unknowns. When the singularities, particles and atoms appeared on one hand and the galaxies and then the stars condensed on the other, the information needed to describe the universe became continuously greater. Entropy increased, order decreased, information required increased.

Claude Shannon's theory of information takes the same form for the loss of information due to 'noise' as Maxwell-Boltzman's Second Law takes, regarding loss of order in heat engines. The degradation or disordering of information during communications happens in the same way, using the same formula, as the increase of entropy or disorder in steam engines.

There has been a constantly renewed speculation on the theme that there is some relationship between the two formulae. It is often dismissed as mystical but always comes back again. It needs an undogmatic re-examination. When we see that the concept 'order' concerns a mental phenomenon the connection becomes plain.

THE BIOLOGICAL NATURE OF ENERGY AND WORK

To understand the wholly epistemological concept, order, we must look at its peculiar place in thermodynamic theory. As I have said, the universe according to Sir Isaac Newton and classical physics had no real direction in time. Classical mechanics works as well forwards as backwards in time, the fourth dimension. But when Lord Kelvin, Clausius, Maxwell and Boltzman really got down to what happens in steam engines they found that real life physics was irreversible. You will not know if a planetarium is being run backwards but if you film boiling water or breaking crockery you will know if you spool wrongly and run the reel backwards. The smashed crocks have lost order and if they were seen to gain it by flying together instead of apart it would be noticed. There would be something odd about steam that shrank and crawled into a kettle spout.

Norbert Weiner pointed out that there is a one-way arrow in time and we can distinguish later from earlier events. The pioneers of thermodynamics found that time insists on its direction as soon as you begin to do any work or use any energy. In the macrocosmic Newtonian world of classical astronomy we see order and perpetual motion as the planets whirl on, untired, around the sun. The classical microcosmic electron, too, is supposed to whirl on round its orbit without fatigue or creating waves.

But in the 'mesocosm' between the galaxies and the fundamental particles, perpetual motion machines have obstinately been lazier; they run down and refuse to work after a short time. How vexing! As soon as we want to use energy to do something useful we get all these problems.

The teleological nature of thermodynamics.

The Second Law of Thermodynamics is an unashamedly statistical law. It is not a deterministic causal law. It is a 'gamblers' law that tells us about chances. It tells us only what happens to very large assemblies of equivalent, equipotent entities and nothing at all about the behaviour of a system's elemental entities (molecules) that *must*, according to the theory, behave *randomly* and with equal mutual effect.

Yet, assuming only an over-riding but statistical law of large numbers and the randomness rule at the microcosm, it turns out to be the surest law of all. Statistics deals only in probabilities, but a probability that is high enough cannot be distinguished from certainty in the real world. The Second Law was formulated in its first form by Clausius, who invented the term 'entropy'. It arose from the discovery that heat engines and refrigerators were governed by a strange unprovable negative 'law' that limits efficiency. The law governs how much of the available energy can be made use of in any thermal system. But 'use' is a teleological term, it implies some purpose.

The Second Law says that if we have a body that is at the same temperature as its surroundings, although both body and environment have plentiful energy there is no way that it can be used. It is only when energy moves from a hotter body to, and disorders, a colder (more ordered) one that energy is available to do work, to fulfil a biological purpose. The amount of energy thus available depends upon the temperature difference and not upon the total energy contained. Part of the energy can be turned to our useful purposes but part must be given up to lessening the temperature difference.

Everything that we know of what happens in the universe increases the total entropy, the disorder in it, nothing diminishes it. That is, simply stated, the effect of the Second Law. If we extrapolate that rule we see a picture of the universe which will gradually run down from its present fairly sorted, ordered state with much available energy, towards a final Heat Death as everything warms up and matter and energy (which are equivalent) are ever more widely dispersed and climb the entropy slope towards less and less available energy. The final state would be when all particles have decayed into radiant energy of lower and lower frequency, spreading into the void in an eventless Cosmos.

We must question some terms that have been used uncritically in the formulation of thermodynamics. We speak of 'doing work', 'consuming', 'useful' and 'available' energy in these contexts and tend to forget that those words are inescapably anthropomorphic (or biomorphic) and teleological and thus are not respectable in the traditions of science.

These words and phrases all imply purpose. They have no meaning that is not associated with some living thing and its purposes or its artefacts. They are meaningless in the world of fundamental physics. We have to ask, "Who?". Who finds it useful, available? Who consumes things, does work and gains from it? Something living, nothing among the Things! It is only when some purposive creature has to make changes in the universe to suit itself that the base concepts of thermodynamics appear. The terms and the laws that appear are

therefore biomorphic laws relating only to the Creatures and having little meaning apart from them. But the concepts order, and its opposite entropy, have forced their way into Cosmology which is about as far from a connection with any purposive being as it can be. So is it true that the Creatures are, by all their actions, simply hastening the universe towards its eventual heat death? The simple answer is 'yes'. There is, however, a 'but', to which we now come.

Seen in the lower levels of the mesocosm adjacent to the microcosm it may seem that life merely hastens the disordering process of the entropy slide. We said that the total entropy must increase. We cannot reduce that total but we can redistribute it, shift it from here to there, from this system to that so long as we accept the overall increase. Useful work is the process by which the downhill increase of entropy in one place is partly balanced by an uphill climb, a climb towards lower entropy, greater order, lower probability in another. We can dam but we cannot stop the flow of the river. However, we, the Creatures, can divert it and use the energy of its flow to modify our world.

There are levels of order. There is the order of an atom, that of a complex molecule, that of a cell and that of the organs of a body. In each case high-level order is built up at the cost of a low level disorder, increase in entropy. High-level order is of a different and more complex nature than the order that is destroyed to build it. Chemical level disorder buys morphological order in an organism. At the cost of molecular or even atomic organisation here, we buy structural organisation, functional complexity and improbable order there. The entropy stream drives negentropy mills, order-building engines. Things get more disordered and probable here so that unlikely ordered things like cities may happen there. Entropy mill-races work miracle mills.

Ranking entities in order of order, so to speak, we would rate the more durable and the more constrained as the more orderly, the more predictable. We would see particles as more ordered, more stable than radiation, we would count more durable particles as more orderly than the ephemeral ones. Stable combinations of particles, atoms, again are more orderly, the more durable the more so. Next come molecules, transient then durable. Then on the scale of order come gases, liquids, crystals and other solids gaining order as they get colder. High on the 'order' scale come creatures and their artefacts - the more intelligent, the more so. Here we leave the world of Matter and enter the Other Realm, that of the Creatures that seem to be designed to endure in the new sense of preserving their form rather than their substance though time. Some species of protozoa have survived for three billion years. In that time every continent, hill, stream, mountain, every Thing and every detail of any Thing on the earth has been transformed many times.

The fuzzy collection of associated meanings behind the words 'order', 'pattern' and 'form' are not merely a ragbag. I believe they may have hidden understandings in them. I point to what may be seen as a semantic chasm between the meaning 'physical and astronomic order' as described above and the other usages 'order and pattern' as found in the Creatures, the living. And it seems that in concepts such as, 'energy', 'work', 'order', 'disorder',' form', 'time direction', 'the expansion of the Universe' that we may see a road to bridge the gap between the Hard Science of Things and the Soft Science of Creatures. The functional order in a cell, a leaf, an insect or a society may seem to be

fundamentally different from the order arising from natural invariances that Mind seeks to aid in understanding, but they may not be so different as they seem. Our subconscious act of classing the two meanings together may be telling us something not otherwise obvious.

ORDER; A SUMMARY

Firstly 'order' is an epistemological if not psychological phenomenon. Finding order is inducing a set of percepts and concepts, mental models, which enable comprehension and prediction in our sort of brain. Order is the measure of comprehensibility, of the coherence of Things with manageable mental tools. Functional order is the creation of evolution or learning, it is inescapably teleological, evolving.

Figure 1: The morphostasis continuum

**Increasing
(as it goes from top to bottom)**

Order

Epistemological constraints

Durability of form

Improbability

'Useful' energy available

Divergence, variety

Amplification of:
  Microcosmic randomness (creating mesocosmic choices)
  Irrelevance of material substrate
  Complexity
  Organisation

Symbiosis

Intelligence (comprehension of the universe by Creatura)

The realm of the Pleroma
Energy (electromagnetic radiation)
  Low frequency
  High frequency
Matter (stable energy)
Particles
  Ephemeral
  Durable
Atoms (stable assemblies of particles)
  Ephemeral
  Durable
Molecules (stable ensembles of atoms)
Gases
  Stable ensembles of molecules
Liquids
Solids
  Stable crystals, periodic
  Aperiodic crystals (information carrying)
Virus?

The realm of the Creatura
(surviving forms, patterns, on pleroma substrates)
Protozoa   One cell life forms, stable patterns of periodic crystals, DNA, RNA, proteins
  Prokaryotes (bacteria, algae)
  Eukaryotes (nucleate cells, stable assemblies of prokaryotes)
Metazoa   (animals) and metaphyla (plants)
  Stable ensembles of eukaryotic form with internal communication
Sociozoa   Stable ensembles, patterns of metazoa with external communications
  Swarms, herds, flocks
  Colonies of animals
  Human societies
    Stable ensembles, patterns of interconnecting people
  Tribes, gangs, firms, cultures, nations, alliances.
  Earth culture. The interconnecting pattern of international communities, e.g. commerce, industry, the market
  Interplanetary cultures?

4. DETERMINISM

I have to intrude on the ancient dispute between free will and determinism. Determinism was an 'unnecessary hypothesis' in Occam's sense, however, it has been an essential and very successful stage in the development of science. No entity that was really a mind could prove that the universe was predetermined because the conclusion would dismiss the entity from the class 'Mind'. This chapter shows that determinist causality on the one hand and Life and Mind on the other are mutually exclusive.

DETERMINISM OR LIFE?

We cannot explain the miracle of life and intelligence without the notions of options and choices between them. There is no sense in your reading these words in an optionless, determined Universe. If you continue to read, you evidently believe these words have a chance of informing you. So now, in the hope of abolishing miracles we venture on speculation about causality, determinism, and uncertainty, not exactly a new theme. We have to relegate determinism from a scientific dogma to an excellent scientific strategy if we are to understand Creatures.

Evolution, learning, and intelligence in a universe without options is too big a strain on credulity, not to say gullibility. But Descartes, a towering figure in the history of science, is among many sages such as Bishop Berkeley, David Hume and Jeremy Bentham who were strict determinists. Further, despite the failure of determinism in the microcosmic non-laws of fundamental physics to which he contributed so much, even Albert Einstein declared that, "God does not play dice". For him we are witnessing a film show where the future is waiting for us, every frame of it, there on the reel of film waiting to unwind.

Determinism as a paradigm was probably a product of monotheism. Pantheism allows for contention between different Gods, and unsettled and contended outcomes. The Pantheon had an open forking future but one God makes an all-embracing set of rules valid everywhere for all systems. So everything is settled. The celluloid characters have the illusion of choice while the universal movie reel unwinds. A determinist world where we can be sure of the future is comforting. Who wants uncertainty?

Whatever its origin and however hard for wilful, striving creatures like men and women to accept it, deterministic causality is deeply embedded in the roots of modern science and it is so for the best of reasons. As a science strategy it works. "Will it work better?" is the question we must ask of new paradigms before we are forced to abandon established ones. Determinism has been decisively effective and productive in every field of science. Before the causality paradigm, the scientist could sweep all sorts of puzzling anomalous phenomena under the

table of animism or entelechy. Until recently determinism was rightly, fully and firmly established as the central strategy of science, almost its hallmark.

The obstinate, obsessive pursuit of determinist, materialist invariances that explain and predict, has produced almost the whole of modern science. Even where uncertainty and indeterminism cannot be disentangled in the microcosmic world of fundamental particles, to continue the search for causal law is the best course for science that aims to produce predictive models and systems. Try for causality until you fail. Why? Because prediction is needed for human survival. But the determinist dogma leaves us these troublesome miracles. Even some scientists find it difficult to believe that their actions have all been set parts on the universal actors' script. Their thoughts, desires and will were all part of the play! It was a severe trial for the faithful. For a faith is what determinist dogmatism is.

It is fairly common ground since Werner Karl Heisenberg produced the uncertainty principle that indeterminism cannot be distinguished from uncertainty, which is said to be compatible with determinism. This makes dogmatists of both camps in the dispute that has raged ever since. However, our day-to-day certainty that we have freedom to choose between different actions is central and fundamental to our lives. No word you read or hear can convince you that you are a robot with all your choices, every motion, predetermined. You may give verbal assent but your behaviour will show what you really believe, that you can influence outcomes.

VERBAL BELIEFS AND BEHAVIOURAL BELIEFS

There are two kinds of belief, verbal beliefs, which affect only what we say, and behavioural beliefs, those which affect what we do. Words written by Descartes, Hume, Berkeley, Einstein and Bohm can give us the Verbal belief that all our striving, contriving and efforts to change outcomes are an illusion. We will dutifully verbalise this view. But our Behavioural belief that we have options and can choose, will still be seen in all our non-verbal behaviour. In their words, many scientists are determinists. In their actions they are unanimously indeterminist.

There are cults, religions and other belief systems which show that there is no sort of crazy farrago that will not attract passionate, devoted, but inoperative verbal belief. Watch what they do. Do not listen to what they say. On this ground I may say that whatever you may say, you are not really a determinist. If the next page is destined to be turned it will be, so why do you bother? If you really believe in strict determinism (and what other kind can there be?) you take it easy and see what your fingers do.

ANTHROPOCENTRISM AND SOCIOCENTRISM

Scientists often warn us against the error of anthropocentrism, of making the world in the model of man. Perhaps we should be warned against socio-centrism.

The monistic view of the universe and the concept of unbreakable laws may be a reflection not of man, but of society. Human societies were governed by chiefs or kings who ruled through laws. As larger social units, this paradigm can be idealised and 'perfected' into a single god with a universal set of infrangible laws that must be obeyed.

The actual observations of modern science give no reason for belief in a universal system of strict causality. Things are much fuzzier. Without exception, the match between what any law predicts and actual observation is inexact, statistical. Laws can tell us of the variable probabilities that connect observations of events called 'causes' and events called 'effects'. Some of these probabilities approach unity, but they never get there.

Heisenberg established, at base, down in the microcosm, that we could account for what we see only by using probabilistic non-laws, constraints on variance. It follows that the hypothesis of causality was an entity 'multiplied without necessity', as Occam put it. Where there are just two mutually exclusive hypotheses there is a problem in cutting either of their throats with Occam's sharp razor. Causal law, statistical law, which is least necessary?

Uncertainty Indeterminism: The Default Hypothesis

The safest course and the practical one is to scrap the most restrictive law. This will avoid errors from assuming undemonstrated constraints. Since causality is infinitely restrictive it is the one that should go until further and conclusive evidence contradicts.

Fundamental physics tell us that the 'laws of nature', are prohibitions rather than commands. Physics is getting cautious and telling us more what *may not* be than what *must* be. There are invariances, constraints, certain aspects of matter that are conserved, mass/energy, parity spin, etc. We can tell what *cannot* happen, there are no inflexible *musts*.

Man is an animal that knows, at the deepest centre of his being, that he survives only by choosing between real options. Nothing is more vital and obvious to him than that he has the power to make decisions that affect what will happen, that he can change the world. His Behavioural Belief would be strained to the limit to accept, really believe, in deed as well as in word that he has no such power, that he is a puppet dancing on causal strings. Verbally he can still be a loyal determinist. We can only discuss life and intelligence that seem to make choices, on the assumption of a forked world with options to optimise.

A non-scientific but quite decisive analogy against determinism is the argument of this industrialist businessman who is used to making action decisions without enough information or time to get it. It is an argument akin to that advanced by

Blaise Pascal, who proclaimed himself a Christian as a precaution, 'to be on the safe side'.

In business we make fail-safe or least-error-cost decisions whenever we can. Which, I ask, of the hypotheses, determinism and in-determinism, does best in the test? Which carries the most severe penalty for error? In a determinist universe the very ideas of *penalty* and *error* are nonsense. *Nothing* can affect outcome. But in an indeterminate forked world there would be severe penalties for those who acted on the belief that striving and effort were *really* irrelevant to outcome. Determinism has the highest error cost and should be provisionally accepted.

To summarise, nothing is more manifest than that we are intelligent, if only because the fact that something is manifest to a creature implies intelligence in the creature. The concept 'intelligence' makes no sense unless there are options, choices between alternatives and means to optimise outcomes. None of these things can exist in a truly determinist world. So if any intelligence, including our own, informs us that the world is determinist, it undermines its authority to tell us anything by doing so. It sends us a message that cannot be both meaningful and true. Information is either surprise or tautology. If the world is determined there can be no surprises. The message of determinism fails Karl Popper's falsifiability test.

QUANTUM WEIRDNESS

Bohr's lecture at Como in 1927 is still accepted as the last word. In what has come to be called the 'Copenhagen Interpretation', he closed a fruitless discussion in a most unsatisfactory way. Statements of his interpretation differ, but the effect is that while in mesocosmic 'classical' physics particles behave like clockwork, in the microcosm the interaction of the observer and the particle system is such that the particles do not behave independently of the way they are observed, and that it is useless to try to understand this in the normal sense of the word. Mesocosmic minds shrugged away microcosmic realities.

In the macrocosm of astronomy there were other un-brainable discoveries. Einstein showed that the experiments on the speed of light and Heisenberg's work proved that measuring rods, clocks and weights were different for different observers according to their relative speed. He showed that there must be non-simultaneous 'nows' and differing rates of time for different observers. These absurd conclusions were infuriatingly confirmed as being true whenever checked. One remark by Niels Bohr should be etched on the mind of anyone who aspires to be member of the world culture. He said, "Anyone who is not shocked by quantum physics does not understand it". Based on the first of these chapters we can say, "So What?" What would you expect of brains built for a world a score of orders of magnitude bigger or smaller? Is quantum weirdness simply an expression of the unsuitability for our brains of these other realms, to brains

genetically evolved and culturally conditioned for the mesocosm. How well would a molecule-brain or one built of galaxies do in comprehending our world? We are born mesocosmists.

WHAT MUST A FORKED WORLD BE LIKE?

Let us therefore posit a forked world, an indeterminate one where the apparently inflexible causal laws in the mesocosm, where we live, are really statistical, probabilistic, in nature. The inhabitants of the mesocosm in such a world would develop brains based on the laws of inductive trial and error. Such brains must learn by repetitive experience based on the expectation that the future is going to be like the past. How else? In beings that consist of trillions of particles the weight of large numbers makes statistical laws very predictive. So the laws they observe are falsified very rarely. These laws would prove an excellent guide to life preserving prediction. Those creatures who discovered and acted on them would do very well, survive, thrive. They would rightly be very hard to convince that the laws were marginally false in the mesocosm and quite false in the microcosm.

Indeterminism Amplified At Hierarchical Stages

At the Heisenberg level (the microcosm), the laws are not so much laws as constraints, low probabilities. The smaller the assembly of atoms, the greater would be the chance that some rare microcosmic event could have a marginal effect upon it, trigger-wise. So the variety of form of very small assemblies of matter and energy that could happen would become very great indeed, because the envelope of constraints, probabilistic, not law-like, would be opened by the indeterminism amplified from a lower level. Small systems would be under a constant rattling bombardment of randomness just as tiny particles are by agitated molecules in Brownian motion.

This wider variety of forms would have to have a non-equipotent (heteropotent) amplified, trigger like effect on larger systems and some of these again would have a marginal effect on those above. But the chance of such quantum level effects reaching the entities in the mesocosm would be infinitesimal because, relatively, the mesocosm is so huge. However, what would be the effect of immense tracts of time on very large assemblies of such a set of systems? What of that infinitesimally small group of them that fell into self-perpetuating and self-replicating configurations? Those earlier collectivities which are stable in form would increase their frequency simply because the stable are durable and the mutable are not. But all such systems have a problem.

THE ENERGY PROBLEM

The substrate system that takes a form or shape and has the quality of growth and division while retaining that shape has an energy problem and a division

problem. It has to have a ubiquitous supply of energy so that it can climb entropy slopes as it converts environment into self. The substrate system has to be such that its elements bind selectively to those bits of the world around them that increase their size without altering their form. Energy is required; debased fuels have to be disposed of. For a growing assembly full of specialist interacting organs to simply split is not on (you cannot get two watches by splitting one). That is why the two forms, the genotype and the phenotype were needed. It is easier to replicate a symbol string than each of a mass of specialist organs and then reassemble them.

In the enormous variety of things that can happen in the vast variety of an indeterminate universe we know that the energy problem was solved in some array somewhere [e.g. by chlorophyll, the energy collector in plants, and the ATP/ADP cycle]. All the life we know is built on the same four selected of many possible nucleotide pairs and the same small selection of a much larger range of amino acids. No other life schemes are known. It had to be a miraculously rare event and it seems to have happened just once in all of Earth's history. How do you distinguish an event which happens once in four billion years on an Earth sized planet, from a miracle?

When it does happen, something new has come into the universe, a new entity, a self-replicating immaterial form. There are now continuing entities that, not being of matter, are not bound by its mesocosmic laws. (This simply means that they will not be bound to have the same astronomically high probability relationships between succeeding events as those of normal cause and effect in non-biological physics.)

The Advantage Of Access To Randomness

Aspiring 'softsmiths' (and 'hardsmiths') who are playing on the fringes of 'artificial intelligence' know one thing well. Optimisation or trial-and-error systems need a source of randomness, or at least irrelevance as in pseudo-randomness. I suggest that the variety of outcomes of a truly indeterminate system, must be greater, infinitely greater, than any deterministic system such as pseudo-randomness. The former has no constraints the latter is completely constrained and predictable if irrelevant. Has anyone brought in real microcosmic randomness for the provision of irrelevant options? Developing life forms would find that access to an input of randomness is an advantage. The best source for such must be the levels below the mesocosm, those nearer to the quantum world. In a random trial and error system the better the randomness the greater the number of options offered and so the better chance of an advance in evolution or learning.

With time, many-stage amplification systems must evolve as long as each stage remains stable and replicating. Thus we might expect rather than be astonished by the evolution of the primitive Life. Such long trigger chains would promote

microcosmic indeterminism in stages up to the mesocosm. This would offer a wider range of options for selection by proto-life-forms.

Those early life forms that are more various have more options, are less bound by the statistical-type 'determinism' of the mesocosm, must from their greater variety have more chance to throw up metastable forms (entropy stealing forms) than those with less of such options. Indeed, it would be surprising if no such long, dice-throwing, option-selecting, chains were to arise. Such chains can reach down into the quantum world, to evolve the long unlikely chains, the means to defy the statistical laws that seem to forbid the life we know.

PHYSICAL LAWS AND BIOLOGICAL LAWS

Nobel prize winner Philip Anderson, the condensed matter guru, said in an interview with the *Scientific American* in November 1994, that reality has a hierarchical structure, "At each stage entirely new laws, conceptions and generalisations are necessary."

The chains of variety-offering triggers, or relays, posited here, would, among many effects, open another door. They would make it possible for a different kind of triggering chain to be built up by selection over time. Indeed, they could evolve into such a chain. The sort of chain I refer to is the long chain of chemical actions and reactions involved in cell chemistry. This is a story that is gradually emerging as the study of ontogeny proceeds.

These long chains of processes are initiated by the genes. They are immensely long, complicated and miraculously improbable on the hypothesis of random assembly. Their effect is that the reactions we observe in developing cells are as Jaques Monod explained, so improbable as to be almost an infringement of the laws of physics. We can say they conform because they happen, and of course the material substrates are 'legal', but that they should happen is not. It is like any other physical illegality, improbable to the miracle level.

If this speculation is fair we would expect in its later evolution and enlargement, that as this new aspect of existence, form-existence, grows up into life forms it will be found to have a set of laws that are no longer based on the mass-action statistical invariances of the material world. It would be another set of laws, or invariances, those of the biological world, which apply to this new immaterial realm of forms. The laws of forms would be different from those of material because elaborate probability-defying systems (forms) have been imposed upon the material substrates. These have made its material immune from the laws that apply to all other material. Those laws assume that the elements of any system are equipotent, of equal effect, like the molecules of gas that obey Boyle's Law. But biological systems are full of triggers, heteropotent effects. There are many relays and trigger-like amplifiers by which small biological subsystems have unequal effects on larger ones in highly specified, non-statistical ways. What is

law-like, very probable, for material, is not law-like for biological forms; they have their own different legal systems, they are heteronomous (ruled by other laws).

It is against the laws of the world for which our brains were designed that we can observe an electron and know which of two slits it goes through. But we have put together a complex array of instruments that enable us to do so. The existence of such an array in this universe is overwhelmingly unlikely (physically illegal in a purely physical material realm). Why be surprised if we get incomprehensible observations?

Paradoxically, we must expect miracles if it is true that living forms have arisen by amplifying and using microcosmic quantum randomness, amplified at many stages, offering variety for selection. But those are the 'miracles', or disobedience to mesocosmic laws that I want to abolish. They are the miracles of wrong paradigm. They arise from the unwise expectation that systems that can tap the indeterminism of the microcosm and amplify it into effects in the mesocosm will behave like those that cannot. If by many successive amplifications at many levels over aeons of time what has been called 'quantum weirdness' has crept up into our familiar world, the mesocosm, we must distrust our surprise at the apparent improbability and weirdness of the result.

INTERGALACTIC CREATION?

The Russian Professor of Chemical Physics, Vitalli Goldanskii, made a relevant report in the *Scientific American*, Jan. '86, ('Quantum chemical reactions in the deep cold'). His researches show that chemical reactions that are 'forbidden' according to classical (mesocosm) physics are possible at very low temperatures. The 'tunnelling' electron that disobeys laws to climb entropy slopes now has numerous applications.

At temperatures in the range we on Earth are used to, many reactions need an energy input to allow them to climb over the entropy barrier between one stable state and another. In conditions near zero degrees Kelvin, quantum uncertainty seems to apply to whole molecules, not merely to particles, and when this effect comes into play they are able to tunnel through the energy barrier without the energy input required by classical chemistry. The effect can be seen either as due to quantum uncertainty of position or as an aspect of the wave/particle duality of all matter energy.

The significance of this is that it has a demoting effect on the miracle of life. Biologists have thought that the environment favourable to the development of very complex incredible molecules is a very small temperature / pressure / conditions slot found only on Earth-sized planets. If over aeons of time there has been a slow, progressive development of complex molecular chains, anywhere in the endless stretches of cold galactic and intergalactic dust, then a variety door and a time door have been opened to the polypeptides (among countless trillions

of less viable forms). The Creatures seem to have had more tickets in the Great Lottery than we supposed. Life needs only a seed. There were more chances to create one than was thought.

BIO-PHYSICAL DUALISM

The ideas set out here can be seen as a form of dualism which avoids the oddness of psychophysical dualism, which asks us to believe the miracle that mind and world work in unison without connection. Biophysical dualism is the view that there are two realms each with their own differing but compatible laws. There is no mysticism. The laws of nature are brainable conveniences, made to fit the world optimally to the mind. Both sets of laws are obeyed but since one applies only to the ever renewed and ever repaired, physical substrate of the other and not to the surviving non-material essence, there is no call for the two laws to be compatible. Both are obeyed within the appropriate realm. Both produce awkward miracles if misapplied.

The two realms contain entities that are in opposition. They are going opposite ways. One towards order, combination, symbiosis, comprehension, intelligence, the other towards dissolution, random disorder and the final hot or cold death of this universe. We, the choosing Creatures, might be expected to have preferences, prejudices, purposes.

Purpose

Assuming agreement with the proposition that purpose can find no place in a determinist world but that it does appear in both life and intelligence, let us explore the idea. Teleological explanations of phenomena, those implying purpose to some unseen entities such as gods, are unscientific, but the phenomenon 'purpose' does appear to exist in Creatures and it is not unscientific to speculate on its origin. I find it impossible to do without the concept 'purpose', in formulating this paradigm for life. We need to give a better account of the view of how it arose. Entities that behave as if they have purposes such as survival and replication are living. This was done as the simplest way of conveying a meaning, it is a model you know. It did not imply that the purpose was immanent, or that all morphostats have anything like human conscious purpose. The sensible, practical, default hypothesis, is that the future is open, not closed; there are options in it. Perhaps there is a God with purpose. Methinks so.

THE ORIGIN OF PURPOSE

Granted an indeterminate universe, it is inevitable that the phenomenon which we call 'purpose' should arise. Some of the surviving forms will, given time enough, produce mechanisms, systems, which enshrine behavioural preferences, prejudices about what to bind to, how to behave, what general aims to have. Those of such mechanisms that improve stability will survive, replicate

and spread. Unspecified purposiveness has good pay-off in stability because only viable purposiveness survives.

LIFE FORMS ON OTHER SUBSTRATES?

If biophysical dualism, as presented here, has any truth, which is to say if it is biologically useful, it might throw some light on Fermi's question, "Where Are They?" There has been time in the 12 +/- 2 billion years since the beginning to develop at least one intelligent form capable of sending messages across space. There is a lot of space out there. Where are the others? Why have they not made themselves known The one life form we know, our own, is based on a four-letter alphabet genome, is in an extraordinarily confined scale, chemical, space, time, duration, cosmic level, temperature slot. Our sensory equipment, vision, hearing smell, touch sensors and the brains that receive their input, are all designed for the task of time defiance in just and only that slot, the carbohydrate substrate life slot.

But can we be sure that this is the only such environmental slot that will support the meta-morphostatic evolution I described? May there not be, indeed is it not likely that there are other tracks towards morphostatic evolution of surviving forms, on other scales, durations, cosmic levels and based on physical substrates so foreign to our whole cognitive system and equipment that they are unperceived and perhaps not perceivable by creatures like us? Within the apparent chaos of plasma and gases within the sun and stars there may be other kinds of form-preserving and replicating entities. Such entities might build up their kind of preserved order on the substrate of hot plasma, triggering order and growth from solar energy. Or there may be undetected replicating and evolving stable patterns in the frozen dust clouds of space between the galaxies. Maybe the very forms and patterns of the galaxies and galaxy groups involve trigger like control centres that use information sensed and coded in some quite alien way to preserve their form through life spans inconceivable to us. These are unfalsifiable and unnecessary 'if's' and 'maybe's'. And if such other life forms are forever unknowable we can happily forget about them or their possibility.

A last 'if' remains. It is this. If we start actually looking for other types of platform of improbable stability in the universe we may even find out where They are, Some Of Them.

5. ON MORPHOSTATS.

The one common factor to all life forms is the form-holding, the active, energy consuming retention of form, arrangement through time. This applies to cells, assemblies of them, metazoa and to communities of metazoa, creatures. The neologism Morphostat extends the class Creatures to include their symbiotic co-operating assemblies. This Chapter speculates of the origin of Life and of Purpose, one of its functions.

The neologism 'morphostat' or 'form-holder' was introduced to show that the class 'biota', or 'creatures' can advantageously be extended to include all entities which dynamically preserve their form, their functional order, through time. This class of entities are in constant chemical and physical flux yet their form is constant. The members are not of matter, there is a flow of ever different matter through them which, using equally immaterial information systems, preserves their often vastly complex form over periods of time that are improbable based on the laws of physics and chemistry. Morphostats include cells and combinations of them, animals and combinations of them and co-operating social groups such as societies. There may be combinations at even higher levels. Our present World Culture could be an example. Some non-biological, purely informational systems and networks may be shown to fall into the class. The recognition of the unity of the members of this class has explanatory and predictive value. Life is the quality of entities that retain form for long periods with low probability.

Science has been built on a base of mechanistic paradigms since Descartes. But science is a biological activity and a social one. It is an interaction between the tool mind and the world. Modern biology, and the growing understanding of perception and of information theory, offer a new route to comprehension arising from a better understanding of the philosopher's and the scientist's instrument, intelligence; of its abilities and the limitations imposed by its origins. Generalist study in the field shows much research that fits, much that confirms and none that falsifies my original conjectures. I have been advised to present the case anew in an updated form.

The proposition in previous chapters is that there is no borderline at which intelligence as a phenomenon starts to be manifest within the creatura. The proposition is that all living forms have, in some form, options, and preferences about them, at least one purpose [survival], and a tendency to optimise their actions on the basis of their information in support of the purposes.

But the proposition is new and peculiar. The view that the optimising choices of the lower life forms are a manifestation related even distantly, to intelligent behaviour is contrary to usage, it is not contrary to logic. We often have to revise the meaning boundaries of words as we improve understanding.

INTELLIGENCE IN ITS EARLIEST FORM?

What is the simplest creature to which we must ascribe some intelligence? The simplest life form is the virus, which acts like a parasitic nucleus in a host cell. It cannot survive independently. Viruses have a genome or design manual, in the form of a random access memory, a ram, 1300 to 20 000 np (nucleotide pairs) long. A np is a two bit information element. The simplest and numerically most successful independent, self-supporting, life forms known are prokaryote cells (bacteria, some algae). They are the simplest but each is very complex. They are the oldest, they have had longest to adapt. However simple they may be relative to later forms, they have had a much longer period of evolutionary-type learning, adapting, refining than more elaborate biota like ourselves. They are, in Earth terms, immortal. How can this happen without intelligence of some form?

We, the Johnny-come-lately humans, should doff respectful caps to such perfect viability. They were around billions of years before us, they are likely to be around a billion years after us. Home sapiens seems to have been around no more that 200,000 years. Bacteria have a 4 million np prescription string (more information than a book) and their own cytoplasm factory, which acts on these design prescriptions. In bacteria a tighter, more efficient genome is free of the long stretches of 'scribble' that infects eukaryote genomes. Are we intelligent but they not? Eukaryote protozoa have 10^{20} million np (including scribble) and a much more advanced cytoplasm factory including the mitochondria, an elaborate factory where incomprehensibly complex communications systems direct the right signals and products through the system. Metazoa or metaphyta, multi-cell animals and plants, have several billion np, a millibook library of book equivalents. Haeckle's intuitive idea orders the 'higher' and 'lower' life forms and puts the mammals at the top and the viruses at the bottom. It is borne out overall by the simple test of genome riches. The genome affluent are, by and large, at the top, the genome deprived are at the bottom, for all that they are the immortals. So far. But there are always exceptions. Salamanders, some ancient ferns, psilopsida, and one humble amoeba are all in the decabillionaire genome class.

THE EVOLUTIONARY TIME SCALE

Before we express surprise at an event, we ought to know what were the chances of it happening. How long has the miracle of life had to come about, what were the chances that it would? Earth has existed for about four and a half billion years. It took about 2 billion years to get the eukaryote stage and all the nucleate protozoa and metazoa have developed in the last billion years. 2×10^9 years is probably 10^{13} generations (for such simple creatures as prokaryotes). Ten trillion replications when multiplied by the number of bacteria-sized entities that might subsist in the Earth's biosphere comes to an extremely large number of trials (6×10^{26} or so, even if we allow the average of only one cell per m^2 as the mean population.) Many trials, many errors, but many chances of success

too. We would expect a hundred trillion, one in a trillion chances to come off in such circumstances. If selected long shots that come off are replicated by the billion, the population open to experiment changes and the odds change with it. The miracle of life looks less so as the nature of a signal has emerged. It is a low energy event that can trigger more energetic events and thus permit Maxwell demonry, climbing entropy slopes. Triggers, relays, are observed in all manifestations of life from the simplest chemical level. Catalysts are triggers, signals.

SIGNALS DEFY STATISTICAL LAWS

We have little idea what early biochemistry was like. The central problem is the energy problem and there seems to me to be but one way this could be solved. There must be a source of latent energy, there must be a trigger event that can release it, and that trigger must be pulled in a selective way. Morphostatic triggers, demons, are pulled by 'favourable' events, those that favour the morphostasis of the form in question. When this amplification happens there has been vital change that takes the system out of the normal mass action laws of the Things, of Matter. The elements of the system are no longer equipotent, having equal mutual effect. Inequality of influence is present. The elements are heteropotent, of unequal effect. The assembly of material can no longer be expected to obey statistical law because the essential condition of such laws is broken.

Chains Of Triggers And Trigger-Trees

The trigger effect described above with a single heteropotent effect would not be noticed but, as must inevitably happen, such trigger actions can form chains, with one event triggering another and that a further event and so on. Further, they will inevitably form hierarchically diverging triggering reactions, trigger-trees, so that the trigger effect spreads its influence through a system cascade-wise.

The Order Explosion

It is not fanciful to suppose that among the great variety of such hierarchical, heteropotent systems built from quantum indeterminism over the 10^{26} trials some will chance to be morphostatic and autocatalytic, stable and self-replicating. At the end of any period there will be more of these around than at the beginning of it. And this increasing population of 'illegal' or improbable systems is by its nature a cascade reaction in itself, it is an order explosion; there will be more of the environment like this as time passes and the morphostatic systems will get more so as the chains and trees get longer, more ordered, more complex. We can therefore demote the miracle we saw in such effects throughout the Creatura today. Trigger chains and trees make statistical laws inappropriate to heteropotent systems. Very long, complex, catalytic chains of trigger-like

reactions are involved in biochemistry at the atomic level. They show us a world of deepening complexity that threatens to become impenetrable.

An article that gives the flavour of this is that by J. E. Rothman in the *Scientific American* 1985 pp. 88-89. The title is 'The Compartmental Organisation of the Golgi Apparatus'. The diagram shows just seven out of a long chain of serial catalytic reactions. To represent what is going on in a tiny sample of the busy complexity contained within the minute span of a cell nucleus, Rothman shows chains of eighty linked boxes and of the complex linking of one of innumerable such chains. Even in this small section of a long chain there are numerous exits and entrances to other connected chains so that I begin to compare the complexity of the whole process with that of the neural linkage of the mammal brain. Will it ever be open to scientists to comprehend the system as a whole?

A parallel is a vastly larger system; the set of neural connections which mediates the typing of these pages. We can comprehend the functioning of a single component, and painfully we can unravel long complex individual causal nerve chain reactions. But we cannot even begin to guess how the whole engine works when all these things are functioning in parallel, concurrently. Yet one signal event of minute energy can trigger a long complex branched chain or cascade of following events. Comprehension requires an understanding of every link of the process. Other ways of understanding are needed. A creature is an informational input-output system. It might be a simpler method to make electronic simulations. I have done this. It works.

Rothman shows that the Golgi Apparatus, a multi layered structure which lies between the genome and the cell wall, acts as a compartmentalisation system. He explains that this is vital to the functioning of the cell because without it there would be a random mixture of thousands of enzymes, and reactions would be chaotic. Each of the vast number of proteins produced by the ribosomes has to be sorted and delivered to the correct compartment of the sorting office by the macromolecules of the many-layered Golgi Apparatus. In order to be modified correctly each protein must pass through the many stages in correct sequence.

This is all familiar enough to a past Production Director. The modern industrial factory has rediscovered the same system of organisation. We did not know that we were making a system of which we had billions of models within us.

The sociological level echoes the molecular one.

Returning to the theme: whatever the primitive proto-creatures were like they must have been able to do something to 'sense' their surroundings. They must have had means of selectivity, to bind to this kind of neighbouring molecule or radical and avoid binding to that. To engulf what is appropriate and eject what was not so, or no longer so, means selection, which means evidence,

information. This involves climbing entropy slopes instead of sliding down them, it means something as odd as falling upwards.

Maxwell's infamous demon, the doorkeeper between two gas chambers, acted as an energy valve, choosing, letting through energetic molecules and thus making more energy available in defiance of the Second Law. The demon had to be intelligent. He would have to have information and base choices on it to pump order into the system, thus making more energy available, Against The Trend.

It is behaviour something like this that the simplest living systems must have stumbled into, including the way of replicating and preserving basic form. Maxwell's thought-experimental demon, manages to tell us something. The vital essence is the informed trigger-like control the demon had. It would only have worked if he used a very low energy event, perhaps a few photons of light, to trigger a much more energetic event. If the signal and the door action had needed more energy than the fast molecule possessed there would have been no gain of available energy, no entropy pumped. Using signals in this way, it is perfectly possible to make energy available in defiance of the Second Law. Any gleaner picking up seacoal is using a few picowatts of visual information to trigger a few watts of muscular power that bring together the means of making a fire that will drive a locomotive that he could not push with all his strength. The proto-life Creatures who did their Maxwell demonry were doing something similar, using information to climb entropy slopes.

Let me show the kind of language that is used to describe the activities within a cell. The sense of sorting, choosing, packaging, delivering is very clear and is easily comprehensible in terms that serve in human organisations. This is at least *like* intelligent behaviour and this likeness should not be lightly rejected as having no significance.

THE MORPHOSTAT PROBLEM

If my thesis is accepted, if life is morphostasis, actively preserving form, rejecting all changes except those toward greater stability, then certain conclusions follow. They follow even if morphostasis is merely an aspect of life forms. Morphostatic creatures will have to classify events. There will be two broad classes of events, threats and promises. Since the creatures are in revolt against the Second Law there will be many more events that threaten disturbance to their form than those which climb entropy slopes to maintain and replicate it.

Such creatures will have to be sensitive to any changes in their surroundings that falls into the two classes, threat and promise. They will have to have means to detect and predict such change-events. These sensors will have to produce those trigger events called signals that release energetic action within the morphostat which will reverse threat-changes and exploit opportunity ones. With

the rapid advance in microbiology we are beginning to learn a lot about the chemical signalling systems that we see are needed.

Protozoan trade with its environment is optimising, choosing, the selective admission or exclusion of molecules through the cell wall; pocketing some materials from around them and excreting by emptying pockets of waste outside them as they choose. Much of this can go on simultaneously and without co-ordination. But much must happen in time-sequenced phases. This means co-ordination, control. We should be able to find some control based on sequencing and priorities. Their genome has 'homeoboxes' which seem to fulfil this function. They must have these if they are to climb probability slopes.

Consider the definition of 'catalysis': 'the acceleration or retardation of a chemical reaction by a substance which itself undergoes no permanent change'. It lowers the energy of activation. The catalyst has the heteropotency that is required for making negentropic energy economy. An entity has an effect but is not subject to one itself, like any other instruction type signal. It can use energy selectively, cause disorder here to build order there.

It is therefore legitimate to classify the whole genome as a catalyst. It is, to be sure, a very large and complex one, but what it is doing is what a catalyst does, inducing and controlling chemical reactions (by the billion) while retaining its form unchanged. It is an extreme example of catalysis, signalling, heteropotency. This is how it achieves negentropy. It builds morphological order at the price of the chemical disorder in its rejected effluents.

THE MORPHOSTASIS CONTINUUM AT LOWER LEVELS

Let us scan the morphostasis continuum from bottom upwards. The prokaryote genome (its control centre) takes the form of a single ring-shaped chromosome which is so long that it has to be scrunched up into a diffuse, unstructured looking toroid lump. The long, tangled torus pulls apart into two toruses at cell division. There is none of the complex mitosis that is needed for the more advanced departmentalised organisation system of the eukaryotes. There is no membrane separating the genome from the cell. We can call this a first-order brain. All communication is by chemical signals diffused via the cytoplasm.

The eukaryote organisation centre is a distinct walled off nucleus, which divides into departments called chromosomes for the tricky purpose of replication at mitosis. This serves as the brain for the more complex protozoa and its signals and messages are again chemical, catalytic. This is the second-order brain. Plants (metaphyta) are assemblies of millions of separate cells, each group of which specialises in some function. The cells co-operate well but they have no federal brain supervising the whole colony of cells, the plant. However, plants must obviously have and do have means of cellular intercommunication. A cell has to know its place in the plant-form if all the cells are to assume the right form

and produce the right products. The cell gets chemical information from its neighbours that diffuse through the cell wall and activate and deactivate genes, bits of the genome, the right ones for the cell's place and time in the cycle.

Plants are autotrophs, they have the trick of getting energy directly from sunlight which provides the essential energy to bring about chains of chemical changes to create the essential phenotype form. With insignificant exception, their only behaviour is growing, making seeds and the means to have them pollinated. They have no corporate motion. Therefore, the slow-moving diffusion of chemical messengers via the sap is fast enough for them. So metaphyta rely almost entirely on chemical signals. They do without a strategic control centre. But they can sense chemicals and light and respond.

Creatures whose only complex organisation centre is confined to the minute nucleus, where the only sort of brain is a molecule, are very limited in learning power. The genome can only learn the hard way by trial and error. The penalty for error is capital punishment. Evolution has to use life and death as a teaching method. Multicellular plants and protozoa that rely on such slow communications can only be genoplastic; they eliminate failure by destroying those brains (genomes) with wrong thoughts[2]. Such creatures are designed for the centre of broad eco-niches, they cannot deal with short-term problems or exploit ephemeral sub-niches except by the try-and-win-or-die method. The virus is a special case, a parasitic brain; an organisation centre that infiltrates a cell and pulls off a putsch, or coup, taking over the control of a host cell body. (Some say the original eukaryote nucleus is simply a symbiote virus.)

Nerves: Metazoan Signalling

Metazoa developed another life-style, a parasitic one. Animals are heterotrophs, life-form predators. They build and maintain their form by the destruction of other life forms. That life-style is more active, more competitive. For predators and browsers much movement pays off, and the faster and the more skilful the better.

Seeping chemical diffusion came to be too slow a means of communication between cells. A faster, more strategic communications network was needed. The next development was that of nerves. A nerve is an extremely long cell branched at both ends to make many contacts. Impulses of electric charge flick along a chain of relays in microseconds. Nerves talk to each other across synapses via high-speed showers of neuro-transmitter molecules over tiny distances.

The channel of communication, the neural network, is especially designed for fast communication. A new kind of organisation and organisation centre therefore arose, a *third-order* brain, an inter-cellular control centre and communications

---

2 'Wrong thoughts' are of course any genetic patterns that lower probability of survival or reproductiveness.

network. The whole mobile cell-colony called an animal developed strategic control centre to supplement the localised tactical control of the minute genome brains.

Ontoplasticity

Animals had a new possibility. They could become ontoplastic. They could adapt much more closely to smaller tighter, more detailed and specific eco-niches, learning, modifying response to stimuli, *within* a generation. Having a fast telegraph system, they could benefit from many types of sensor to gather more and more elaborate environmental information. They could move when threatened and to get food. They could be viable in larger, wider eco-niches because they had the choices brought by mobility and a better means of information leading to more options and chance of selection between them. They had more choice. They could *find* suitable niches.

I now take a look at the more advanced type of brain, the seat of the highest form of the intelligence that we are examining. This is the outward form of our intelligence tool, the one that is trying, presumptuously, to model itself. The description is of the gross visible features. It has little immediate congruence with the model of information movement I am developing, but is a preliminary reminder of the physical form that the topological (or connective) form must use as substrate.

Mammal Brain: Human Brain

In its size, class and weight, the seat of human intelligence, the brain, is the most powerful, dangerous and promising entity in the universe. That much plutonium can make a small big bang but a single governing, human brain with its heteropotent effect on a social morphostat can trigger many bigger ones, change the biosphere, bust the moon apart, defy the statistical laws of nature or…you name it!

It is very different indeed from anything else of its scale. We have to look at something about 9 orders smaller to find something in the same class of wild incredibility. We have to look at the nucleus of the living cell. Are we not looking at something related? If we probe into the mammal brain we find a level of organisation and incomprehensible complexity which makes it, small as it is, by far the most complex, local, holistically functioning system known in the mesocosm.

The accepted paradigm for the material universe (in the mesocosm) is one where effect follows cause in time. Brains reverse this. The brain anticipates, predicts future events, and the prediction of tomorrow's events governs today's behaviour. The effect, the behaviour, precedes the cause, the future event. The birth of baby today caused its mother to knit little vests three months ago. The

brain reverses the temporal order of causality. It induces negative causality, time reversal.

Filing, sorting and accessing data is now well understood thanks to computers but there are things that the brain can do that defy the comprehension of the computer scientist. Remembering the brain's slow processing speed (a few dozen Hz), how can we explain a classification system that gives us instant access to long complex memories on the basis of a minute irrelevant signal? A whiff of a smell will bring back detailed memories of a whole forgotten childhood sequence of events.

May I remind you what a brain is like physically? Simplifying, a human brain is a Russian doll system of three brains, one within another. The original brain or paleocortex is the deepest layer. This corresponds to the reptile brain, which has been twice overlaid in subsequent evolution. The second enwrapped layer of the parcel is the limbic system or medial cortex, which we share with the other mammals and only partially with the reptiles. This part of the brain is probably over 150 million years old. It appears to be the seat of our emotional life. Above it is the neocortex, which exists in other mammals but which has been developed in the primates and massively so in man. The neocortex is credited with the great difference in cognitive (communicable, symbolic) intelligence between man and the other mammals.

The neocortex is divided into the two massive cerebral hemispheres which have an intricately riffled and infolded surface like a walnut only much more so. The two cortices (surfaces) thus have an area of about a quarter of a square metre. The convoluted surface of the neocortex is 3 mm thick and consists of about ten billion intricately intertwined neurones, the communicative brain cells. They are packed in a vast putty-like mass of supportive glial cells. The neurones are very strange cells. Exceedingly complex and variable, they fall into a number of classes. Some types are much ramified at either end of a long, containing, insulating tube called the axon. The intricately intertwined spaghetti of neurones within this shell of the neocortex are very densely packed indeed so that it is a serious problem to make out their pathways and interconnections. The cell body is usually at the input end, and the hundreds of probing dendrites or thin tendrils reach out from it and contact other cells at thousands of synapses or connecting points on the body or processes of other contacting neurones. A neurone can be over a metre long yet 3000 can be packed into a cubic millimetre. The whole intertwined basket of octopuses is packed in like the myriad fibres in the baize of a billiard table. That is a simplified description of a thinking tool, a tiny entity in the universe that can, in some degree, make a model of it and predict its future.

A good deal is known about functions of gross features of the mammalian brain and even more about the human brain. There are specific functions associated with defined areas of the cortex, visual reception areas at the back of the skull, a memory centre in the amygdala, speech areas in the temporal lobes and there

has been an accurate mapping of the motor centres of the various parts of the musculature, arms, legs, mouth, tongue, hands, fingers and so on. The medial cortex seems to be associated with the appetite system, (instinctive, emotional functions in the animal) It sends streams of excitation to all parts of the brain. Pain and pleasure centres have been found but there is little understanding of how they work. The cerebellum seems to have a very important role in co-ordinating movement. This crude mapping has been done by careful and systematic observation of the effect of lesions in various parts of the brain on many thousands of patients, many of them First and Second World War soldiers with head wounds.

The brain is, as I said, dual. There is a great cleft between the two halves, which are joined by a thick commissure made of packed fibres at the base. By cutting this commissure it has been shown that in many ways the two parts work separately yet without apparent conflict. There are a lot of left brain / right brain theories and discussions at the popular level but little consensus. The duality of the physical brain may be an essential feature of its function or, more likely, I believe, a contingent effect of the basic symmetry ground plan of the vertebrates. All types of learning, including evolution, have to start from where they are and vertebrates are symmetrical about a spine. No duality is essential in the primitive brain model I advance but, if only because redundancy is essential to the model, duality does not conflict with it.

So much for the most complex and advanced sort of brain we know, that is confined to one contiguous mass. We may now look at a brain that is very much more difficult to see as such, a dispersed and/or intermittent one, the brain of a morphostat that is dispersed in space/time. I speak of a brain in which the only essential feature of any brain, intercommunication, passes through much more space and time to connect its elements. Dispersed and intermittent intelligence's are those intelligence's belonging to and working to preserve morphostats of which the elements themselves are the morphostats called animals (which are themselves composed of the morphostats called cells). It is to these third-order morphostats that I now draw attention.

THE SIGNALS OF THE SOCIOZOA

Atoms fell into stable forms, molecules. Molecules fell into form-preserving arrangements, crystals and then information-guided morphostats, cells. Cells fell into stable arrangements, animals, with nervous information systems. Did it end there?

If separate animals can fall into a co-operative symbiotic mode, can develop the essential of any morphostat - an internal communication system - is there any reason why they cannot form and be part of even higher level morphostats? As I have defined and illustrated the paradigm 'morphostat', it would be improper to exclude many intercommunicating social groups of animals from it. Zoology is full

of examples of symbiotic colonies of animals that live, feed, grow, divide and thrive as a result of sociality and social communication. Swarms, herds, schools, flights, colonies, nests - there are scores of names for these third-order morphostats, these assemblies, colonies, which behave as though they had the 'purpose' of preserving, with as little change as possible, their form through time.

Mutual Signalling, The Sign Of A Morphostasis

The visible sign of a morphostat is a mutual signalling system or communications system. I have argued that all mutual signal systems indicate the presence of a morphostat. So to call the class collectives of animals 'morphostats' is not to make an analogy. That act of recognition results from an act of taxonomy, the formulation of a class of entities. We are always entitled to do this. When we do we must try to see if the new class, as a whole, has characteristics which are useful and relevant; characteristics by which what we observe in one subclass can give us predictive hints about another subclass.

It is only in my lifetime that we have gradually seen how very closely plants, animals and micro-organisms are related. Then we found the genome which confirms that all are built on an astonishingly similar plan. It is as though the genome 'program' has a 'menu' (to talk computerese) that allows it to opt for yeast, slime moulds, cabbages or Einsteins at choice.

Important advances in knowledge often come from insights that show that phenomena we thought to be disparate were really different aspects of the same thing. Falling apples and planetary orbits are only one example. (Seeing distinctions that divide a class can also lead to better comprehension and prediction.)

The idea that an interacting social group qualifies as a being is not new. Much of the language of politics derives from the Roman analogy of the State with the body, 'organisation' from 'organs', etc. The paradigm was very powerful and it may be no accident that it was one which led to a large and very influential empire. In William Shakespeare's *Coriolanus* (1.1) Menenius tries using a parable to pacify the citizens who were rioting against the rich patricians:

"There was a time when all the body's members rebell'd against the belly; thus accused it; - That only like a gulf it did remain I'th midst o'th body, idle and inactive, Still cupboarding the viand never bearing Like labour with the rest, where th'other instruments Did see and hear, devise, instruct, walk, feel, And mutually participate, did minister Unto the appetite and affection common Of the whole body."

Menenius then goes on to explain how the stomach does its work and distributes the results via the blood stream. The strength and usefulness of the model of the State as a body is manifest. Some such comprehension tool is essential to a

developing society that is breaking old habits and ways and urging people into new roles so that a much larger and more complex, more viable social creature, sociozoon, can find ways to survive.

However, the analogy with a human body was too crude and had to break down. It is my suggestion that all morphostats have a similar set of problems and though they may be working at very different levels there may be instructive parallels about the solutions that are found. At all levels morphostats will have aspects that are more open to enquiry. By cross fertilisation at several levels we may raise insight, understanding and the object of these, prediction. By understanding how data processing and intelligence work at the various primitive levels we may, in the end, be able to bypass the genome and develop intelligence entities that are not constrained by 4 billion years of biological clutter. Some present attempts at Artificial Intelligence seem like trying to make robot horses to draw carriages instead of making horseless carriages as Herr Mercedes did.

Sociozoan Communication Is Primitive

First, I have to admit that sociozoa, though they may have a great future in the next few billion years, are fairly primitive today. Measured by the actual amount of information measured in bits, which passes between most animals in social groups, it is very small indeed compared with the enormous internal exchanges. It is like the exchanges between countries in the Middle Ages. A few million richly interacting people 'talk' to another such group via a single ambassador. One animal with billions of richly intercommunicating cells and nerves utters a one bit warning cry. But it starts a herd into flight and saves lives.

The signals that pass between social insects within the sociozoon called a colony, swarm, nest or hive are quite elaborate compared with the primitive ones that pass between, say, the birds in a flock or the fish in a school. Precise metric, cognitive (symbolic) information about the locality of food is passed by bee-dance. This is qualitatively different from the imprecise danger squawks of birds, or even the conative (emotional) rather than cognitive cries of many social mammals. Most of the 'higher' animals have little more than the power to express emotions, such as fear, rage, sexual readiness, submission, dominance, and sometimes 'food here' without location information. Yes, the insect sociozoa beat man by mega-years in the race towards cognitive communication, but they had an earlier start. Also, their language is hardware, not software, genoplastic, not ontoplastic. They take many thousands of generations to learn each phrase.

THE LEVELS OF INTELLIGENCE

With interesting exceptions like the above, the continuum of intelligence, of brain complexity, seems to run parallel with the morphostasis continuum outlined earlier. The cell nucleus is one very complex brain-like system, the neural brain is

another on an entirely different scale and level. Yet they work in close concert, their activities mesh together perfectly while each works independently at its own level. We have morphostatic nodes within morphostats.

The effects of physical interference with brain, such as by trauma, drugs, etc., tell us this. The activity we are conscious of, our thoughts, are closely bound up with events that are happening to the ephemeral molecules and the more permanent cells in our brain. But we know nothing of the activity of the ten billion genomes of our neurones. And to expect those genomes to know about the brain of which they are part would be folly of exceptional purity.

Now let us trace intelligence *up* (just in mass) one further step. Concede that the neural brains in a sociozoon morphostat, a herd, jointly create a co-operative behaviour pattern that helps the herd to survive. The signals that pass from one animal to another modify the internal communications system and the result is herd behaviour which is emergent from and includes animal behaviour. The herd behaviour may be directed by one leader but this is not always so. An enormous school of fish will wheel as a body without visible command.

Either way, the total effect, the herd behaviour, is the result of the whole of its set of nervous systems brought into unity by inter-animal communication. This is a primitive version of the system in your body, that by which your cells exhibit unified behaviour based on neural communication. It is a fair parallel of the set of brain cells that are in uncomprehending co-operation to produce your thoughts as you read these words.

HUMAN SOCIOZOA

It is obvious that I have to apply the 'sociozoa' concept to humanity. How well does it fit? If the congruence or fit of the model is good enough it may be useful to see human social assemblies as form-preserving entities.

Certainly my personal experience of many score of organisations and institutions amply confirms the view that they are inherently conservative, none more so than those that label themselves radical or revolutionary. Those that were not conservative seem to have been ephemeral. Apart from personal observation, I have read about and heard of many human collectives. All those that merit the name 'organisation' or 'institution' were the same. For all, the Principal Name of The Game was preserving the form (a set of role relationships) with as little change as possible through time.

Institutions formed for a specific *purpose* often survive long after that purpose has been achieved. Many continue for generations, paying nothing but lip-service to anything but the real, but often tacit, aim of all known morphostats, survival.

Human purposive sociozoa, institutions set up to further a group cause or serve an interest, are, biologically, a very new idea. Until the last few millennia all co-operating human groups were like those in New Guinea and in the Amazon basin today, *purely* survival groups. They had no purpose except preserving the tribal structure, its growth and reproduction, the elementary requirements of morphostasis. For almost all of man's history the successful social pattern was this small hunter-gatherer tribe. Only around one thousandth of the life of the tool-making hominids has been post-tribal.

In a book with this title we must naturally be concerned with the last small sample, the one extraordinary thousandth, the miraculously improbable intelligence explosion that started when agriculture arose and made large viable groups possible. This happened only a few millennia ago!

The subsequent coalescence of tribes into kingdoms and republics, and of these into alliances and empires, the emergence of that strange un-biological coalescence, the present co-operating, intercommunicating and interacting world civilisation and culture, is another uncomfortably miraculous event that fits badly with the rest of biology. It is certainly still true that most human sociozoa such as ethnic, religious and cultural groups, commercial and industrial firms, farming groups, etc, have no real-life purpose other than survival and growth. In this they are the same as any animal or bacterium. The more recent invention the Nation State, is much the same.

However, we also see the astonishing emergence at another level of complex world-scale form-preserving information networks like the commercial and industrial markets, the science, business, legal, banking, maritime, cultural, artistic, entertaining, communications, technological, academic, and many other networks and communities, all interwoven and interacting, nearly all arising spontaneously in the last few centuries. It is mysterious and needs a lot of explanation on accepted paradigms. I say again, "*How come that creatures which slowly adapted for millions of years to life as small groups of hunter-gatherers on plains or in forests have suddenly swarmed together into an enormous network of tightly organised and highly successful, interacting networks of organisations?*" It is odd that I set solid-seeming men and their immaterial associations in the same class? But both are form, not substance. It is hard to hold this in mind. You are not a solid mass of molecules. What you are is the form that may be taken up by any such set. You are a river whose source is your mouth and whose mouth is your anus. You are like a tribe or a limited liability company - immaterial, self-preserving form, not substance.

The miracle of human tribes that suddenly swarm together like slime moulds to form a giant world-scale fourth-order sociozoon, the present world culture, is going to take a lot of abolishing. Perhaps the first clue to the answer was given by the economist Adam Smith. He posited 'engine', already mentioned, a mysterious hidden hand that worked behind the scenes to arrange that all the

autonomous self-interested behaviour of citizens somehow merges and combines to produce a beneficial, stable and autonomous economic distribution system. This happens in any area where individual behaviour is not too much constrained by tribal norms or where law and order are maintained by the evolution of a government tolerant of free traders and artisans.

The Market As A Brain

Smith's intangible 'market' is an emergent primitive societal brain which collates and combines all the signals of need and desire from millions of people in their consumer roles, processes them and transduces them into that enormous set of detailed instructions received by the same citizens in their role as producers at the work place. Result: a closer match between what gets produced and what is required, between supply and demand, than has been observed under any other system. Workshops, mines, distribution chains, farms are started and halted, products changed, batches ordered, people move, voluntarily, each doing his or her own thing, yet the total result is a highly optimised balance of needs and desires with supplies. The whole machine works without monocentric control and planning and the whole vastly complex system arose like Topsy, it was not born but 'it just growed'. Out of nothing?

A billion autonomous selfish people interact, exchange information and co-operate unconsciously in a system to preserve the form of unknown institutions by which they survive. The elements of a cell co-operate and communicate (chemically) so that its form shall survive. The cells of an animal co-operate and communicate (neurally) so that its form shall survive. In no case do the elements seem to be under any straightforward central direction. It was this thought that set me on the long track to the writing of this book. It is this parallel, this congruence, which may tell us more than we knew before. The cell and the brain remain, largely, black boxes. But the communications systems of sociozoa are patent, open to view. We can see how well the pattern of information movement we observe in society fits with the pattern we can barely make out at the molecular and the morphological levels. We can (and are) trying out computer models based on the sociological level to simulate the neural level. It seems to fit, so we may hope to find out more about all three levels of what appears to be a single class of entity.

Readers who deplore the patent imperfections of market methods are invited to look round the world at those economies where the market works freely and those where monohierarchic central control is felt to be a better optimiser for the satisfaction of needs and desires. It seems to be broadly true that, in free economies, living standards tend to be higher, even for the poor, welfare better, and freedom greater, including the much exercised freedom to complain. A greater variety of goods and services is more freely available, but there is more envy, politics and general dissent. If there is an example of a centralist government that does better at getting the right amount of goods and services

available in a more efficient way, then that economy is keeping its success to itself, because I have not heard of it. Perfect the market is not. Something better we cannot find.

The picture of the market as a self-organising system that works in some way independently of the will or thoughts of any one person may seem strange but there are many examples, as Friederich Hayek points out, of similar apparently rational processes which arise unbidden from the interactions of many brains. For a culture, a language is a strange unified communications system, a form that is retained and evolved through time with little or no central intention or planning. It is a morphostat that subsists on the insubstantial substrate of human cognitive intercommunication. There are many human cultural, artistic, religious, business, technological, scholastic and moral traditions that are similar types of entity.

The brain as a control centre?

However much we may be intuitively attracted to the Roman Empire model, the model of the central control system, with a government acting in the role of the control centre or brain, we have to be cautious when we see the above examples of apparently spontaneous undirected self-organisation. The brain itself shows no signs of congruence with a hierarchic control centre paradigm. The cell level organisation is an equally mysterious black box, like the market, that fourth-order brain. This thought is offered as a confirmation of one thing Marx claimed, that detailed central control was inessential. States, he suggested, 'would eventually wither away'. However, the system of social classification by social and occupational class is unpredictive and inadequate. Those vast classes are too large and too fuzzy. We need tighter categories if we are to have hopes of a predictive science of sociology. But institutions are morphostats. That goes even for communist parties and governments. They will be around for a long time like London's ancient Guilds.

THE CONTROL OF COMPLEX SYSTEMS

I want to suggest that there is a fundamental flaw in the traditional model of organisation which bedevils our thinking about it. The organisation chart that we have in our minds, that of a monocentric hierarchy leading down from a chief executive to a ramified structure of departments and sub-departments is muddled and deceiving. The central problem of any organisation, in the light of the morphostasis concept, is to counter change tendencies from many sources. There are always many ways in which a morphostat can lose order, many variables to be kept within some sufficiency boundaries (to be kept sufficiently close to some optimum state). Therefore, for any complex entity to preserve its form it must organise; form organs to specialise in the numerous different homeostatic tasks. Nearly all of these have to work simultaneously and autonomously. Many must have their own sensors, their means of detecting

change threats or opportunity promises. The whole of the world of autotrophs, the world of vegetation, a vast range of complex forms, survives and flourishes with this kind of polycentric control.

The organisation of all the protozoa and nearly all the activity of the metazoa is tactical, not mediated from one centre but polycentric. Most of what is going on, even in an animal, most of the time is autonomic, self-regulating, the body is kept functioning because all its parts detect and respond to local inputs. The strategic level perceptuo-motor behaviour we see in more advanced creatures, the behaviour mediated by brain-controlled muscles, is more obvious and dramatic but it is, in terms of the volume of data-processing, only minor. Most of the data-processing is tactical, specialist, autonomous and at cellular level. It is only in what are classed as the 'higher' animals that there is anything like central strategic control at all. Setting priorities and *doing one thing at a time* is seen only at the top of the morphostasis continuum. The strategic level, deciding priorities on multiple goals that cannot be achieved simultaneously, is essential because of their more active life; first of opportunists like predators, and then defensively of their prey. This development drove them to develop brains with a system of priorities, with flexible attention, a strategic control system. Browsing is vital to herbivore life but fleeing from predators takes priority at some times, so does drinking, sex, and breeding at other times. Central data-processing with means to give priorities and emphases is needed. But that does not necessarily mean CPU^3- type central control, as I shall try to show.

Konrad Lorenz pointed out that human tribes and even nations have two mental modes, the normal gathering mode and the occasional hunter, warrior mode. He suggested two entirely different types of behaviour, two programs on a computer-like 'menu', which can be selected according to circumstances. In the 'military enthusiasm' role the same person responds differently from how he would in the other one. The man who avoids a scuffle in one role will charge a machine-gun nest in the other. Man in the hunter/warrior mode must be more leader-prone than as a gatherer or cultivator. Thus the need for two modes of behaviour. The natural pattern in small groups was that of a single leader for hunting and was an autonomous action for gathering vegetable food.

The same central leader pattern served very well even when tribes began to coalesce, there were warrior leaders who, because of informational constraints grew instructional, (motor) hierarchies with several strata beneath them, commander, officers, sergeants, etc. This leadership pattern is deeply engrained in our traditions and possibly in our genetic nature because it was very suitable for our long stint as small hunter-gatherer groups. So our much larger and more complex institutions are built round a monocentric mental model.

However very large institutions cannot, for data-processing reasons, work in monocentric mode except intermittently in very exceptional circumstances such

---

3 CPU is a computer Central Processing Unit

as war. The problem is simple. In a monohierarchic system over a certain size, where all information has to be channelled to a single apex, the upper strata and especially the apex itself become overloaded and cannot cope with the rate of data movement. Organisation, forming specialist departments (both on the afferent and on the efferent side of the apical leader) becomes necessary. Eventually these 'organs' need and get less and less input from the apical figure. They become autonomous and the system is transformed into another form.

If we examine the actual movement of information and instructions in a large organisation I suggest it would not look remotely like its monocentric authority or organisation chart which is supposed to be the guide to the thinking of those in the network (see Figure 2).

Figure 2: Typical Organisational Chart
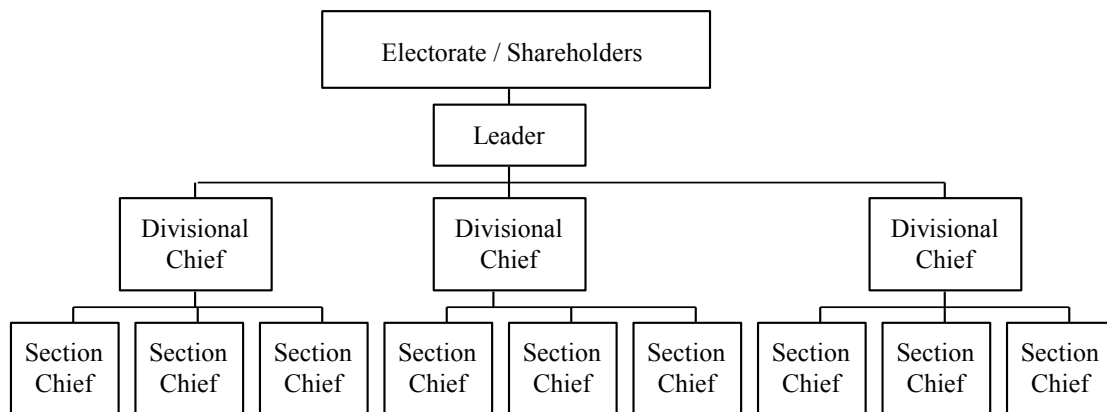
I do not think we are going to be able to understand the various levels of intelligence or create artificial intelligence on a monocentric model. I shall propose another model which has a better congruence with reality and which may be a better ground plan for an artificial brain.

EVOLUTION AND LEARNING REQUIRE RANDOMNESS

Philosophers, so far, have been people. But recent work on computerised pattern recognition shows that in taking perception as their starting point, a simple 'given', they have failed to see its vital importance or its complexity. They have modelled in ignorance of the fact that for people, including philosophers, perception is largely ontoplastic not genoplastic, a learned skill, not a given, hardwired propensity. Learning involves trial and error, trying models and rejecting or improving them over many trials. The very act of seeing is an inductively acquired skill. Adults do not remember that primary baby learning, the hard-won skills in learned taxonomy and pattern recognition, took all of their babyhood, instead they are taken for granted. People, philosophers among them, are too familiar with their adult world to see all that was involved in entering it.

Trial And Error Require Experimental Variety

It is obvious that trial and error in the completely unsophisticated beginner must involve an input of penny-toss randomness. Complete ignorance must try anything. Options for trial must be created, just as genetic mutations must arise before they can be selected. In evolution it is widely accepted that mutations are random. I suggest that at first all and later at least some of the activity of learning involves an input of true randomness. A baby which has very few genoplastic behaviour patterns, almost no preconceptions about what is the right behaviour response, needs to experiment with as little constraint as possible if it is to hit on the right combinations of instructions to muscle fibres.

In this context, randomness could be defined as complete absence of constraints on options. Therefore the desirable source of the options required should be as close to being truly random as possible. To see this clearly we have only to look at the enormous range and variety of human skills. Skaters, acrobats, jugglers, artists, all produce a vast range of complex, intricate, incredible behaviour patterns, many of which seem to be unbiological in that they cannot be necessary to survival. How can all the required option doors be kept open if there are many hardwired constraints?

True Randomness Is Microcosmic

In the heavily statistical world of the mesocosm the only source of pure randomness is the world of quantum weirdness in the microcosm. To draw on this in the mesocosm it must be amplified enormously. Such amplification is possible. A Geiger counter does it, it produces a mesocosmic click for a microcosmic electron or photon. The British Post Office savings lottery system uses a computer called ERNIE which draws on quantum randomness to award prizes. The very great electronic amplification factor is based on complex technology.

Microcosmic Events Can Change A Cell

A single cell such as a receptor is an enormously large system compared with an electron, so the statistical laws of large numbers ought to apply fully. It would seem to be impossible for the minute remote impact of a fundamental particle to be amplified so enormously as to make detectable changes in a human nerve cell.

But they do. Experiments have shown that, marginally, a human eye can score better than chance expectation at detecting a few photons of light. Further, we know that radiation can convert a healthy cell to a cancerous one or cause a genetic mutation. Radiation is a quantum effect, a step function, there must be a step difference between a cancer-causing amount and an amount less than that.

The Amplification Of Quantum Events

Quantum level events, therefore, can indeed trigger changes in cells. Further, we may reasonably posit that cells, being heteropotent and intercommunicating, may, must form trigger-trees and cascades by which an event in a single cell can trigger a much larger action in a brain and then in a musculature. The action thus triggered in one person's brain can be sent by hierarchical social cascades to a society where that brain commands or influences many others. A marginal quantum event in a particle can trigger the movement of an army!

RANDOMNESS AND PSEUDO-RANDOMNESS

Nearly all computers have a means of generating pseudo-random numbers. It is claimed that when these are used in learning programs they are indistinguishable from true randomness. I beg leave to doubt the claim. Firstly, I have too often found that I am in a highly constrained recursive loop when I have used pseudo-random numbers in learning programs. Secondly, as I have said, quantum randomness is without constraints, while pseudo-randomness must, in any physical system, have some constraints if only those of computer time. Take for instance the time intervals between gamma ray emissions from a radioactive substance. This is a truly random magnitude which can have any real number value in femtoseconds between some upper and lower practical limit. An infinite number of such real numbers would be required to represent the possible variety. As I have said, any practical system for producing pseudo-random number must impose some constraints on the real numbers it could produce between these intervals. Merely expressing them in any conceivable number system would be an important constraint. There would be too many decimal places after the zero! Pseudo-randomness is fully constrained in that it is deterministic. The reason it can be used in learning programs at all is that its input is *irrelevant* to the particular process in hand and so can produce previously untried options. However, being determined, pseudo-randomness has limits to the variety of untried options it can offer. True randomness has, by definition, no such constraints and where can it be found, in mechanical terms, except in the microcosm?

RANDOMNESS IS PROMOTED TO HIGHER LEVELS

What I suggest must happen in animal and especially human learning is that as learning advances and becomes more sophisticated the *randomness applies to more and more complex and developed systems*. It moves up into the network, affecting areas at a more strategic level. That may seem strange at first sight but consider this. The learning baby needs trial and error to learn the simplest basic skills, how to move groups of muscle fibres in unison, then a limb to the left, to the right. As these subskills are assembled into skills they become more repeatable, more deterministic, less random at that level but at the higher level where subskills are being combined into skills there is still trial and error and the

need for random input. Later, with skills having gone hardwired as it were, the randomness input needs to move to more senior, more strategic levels still as skills are assembled into behaviour patterns. Until these are learned there is still a need for random input but now it affects penny-toss choices among much more developed behaviour trains. The random input moves distally, inwards, away from the peripheral and tactical zones towards the central and strategic ones as the learning process goes on.

The baby makes tiny behavioural elements deterministic, as it learns, then assemblies of them become so, then assemblies of those and so on. After a time the child has a large and growing repertoire of fully determined skills, abilities, whole temporal trains of behaviour. But where is the random input now? It is now offering options about prolonged periods of skilled activity. Random, experimental trial behaviour now applies to these. Can we falsify this? Is this stage observable in babies? We cannot falsify it. We see such a stage. It is called play. The randomness now applies to what role the child takes and what game is played, no longer to which muscle fibres are moved and how much. Does this process finish with childhood? Not in man. In adults, sport, pastimes, hobbies, games, gambling, whole life schemes with no direct biological benefit continue. At that age the random input is given other names. Free-will, autonomy, whim?

Free-Will, Whim, Perversity

Consider free-will. If our behaviour is the sensible and proper response to the informational input, the data store in memory, and appetitive needs all the time, then it is not free. It is deterministic, the only possible output from the input. It is only when it runs counter to the obvious, where there is an element of self-will, whim, perversity, that it can be called really free. This is unachievable by an automat. Yet we cling obstinately to our power to be whimsical, to behave in spite of indications and not because of them. Is not this precisely the freedom that we are ready to die to preserve?

Do we die to preserve our right to behave illogically? Yes, and rightly so because the whole business of learning is itself illogical, trying options randomly rather than because we know how they will turn out. Lose whim, lose learning, lose progress. Logic is a determined chain. Learning is necessarily indeterminate. The random input moves up the scale and into the network from the tactical level at muscle fibres, to skills, and then longer and longer-term strategic behaviour trains. Where does the process finish? At the limit, whole policies which may last years come into the region of arbitrary choice. Our whole mind set and path through life may have started as random choices. Really free choices. Those that could have been different.

The Perception Of Free-Will

Let me come straight out with it. Let us look at what is so much disputed between philosophers, free will versus predestination. Following the reasoning above our perception of free will is the last vestige of the random function at the end of its journey up the tactical/strategic ladder. Can my treasured free-will be just quantum weirdness borrowed from a senseless particle? This is a hard pill to swallow and I obviously balk at it myself. But I have to show where the argument leads. There is a difficult question. If free-will is not this, how could we tell? What is free-will but the desire be free to do the unobvious? The obvious, a computer can do, must do. A Von Neumann computer is the least free thing of its size and complexity in the mesocosm, designed to be so. Is that what we want to be like?

When we see free will operating it is impossible to say that we are not seeing a microcosmically generated event, like the click of a Geiger counter. Only bigger.

Caveat.

An *apparently* free choice is often much constrained by unseen things below conscious level; there are usually only a few *true* options. To the extent that our will is influenced by these unconscious inputs it is determined, not free, the true freedom remains only at the last most strategic level. Here we plump for one among a small number of highly favoured strongly indicated options. At this point the randomness that used to offer options to muscle fibres as a baby will now operate on decisions that affect the rest of our life. Whom shall I marry? Which job shall I take? Which book or invention idea shall I work on? It would be disturbing if our will free turns out to be free alright, but not ours, just borrowed microcosmic randomness.

The alternative involves something like Julian Huxley's idea. Mind, sapience, he implied, is immanent, it is a hidden aspect of the material of the universe. My argument here shows how it might penetrate the mesocosmic material world. But this is an unfalsifiable and unnecessary hypothesis that offends both Popper and Occam! Without it, if we follow the morphostasis continuum, we have a world in which more and more of the universe is making larger and larger experiments with random options commanded by events at the lowest level. There are known to be genetic factors which partially decide, other things being equal, how tall people are to be. This does not mean that anyone claims there is such an objective reality in the genome as 'tallness'. 'Intelligent' is a description of a position on a rough judgement continuum about thinking power which has been reached by a very widespread consensus.

As I have pointed out in previous books, the sort of intelligence that we are trying to measure in our social judgements (and that formalisation of those judgements, intelligence tests) is social intelligence, verbal or cognitive intelligence, the sort of strategic level conscious intelligence that is more a social than individual phenomenon. It helps everyone more than someone.

It is that very special kind of symbolic coded intelligence that is communicable between the units of that special kind of sociozoon, a human institution. The limitation of intelligence tests is that they depend largely on language. Even the so-called Culture Fair tests involve the comprehension of such artificial conventions as diagrams. These are good when you have been exposed to them but are less predictive or unpredictive when used on subjects who are unfamiliar with the conventions. This means that there are two banal stupidities that are repeated endlessly about IQ tests. One is that because they are not fully valid except on the language/cultural group upon which they were standardised they are of no value at all. The other is the false circularity, that that which we call intelligence is only what intelligence tests measure. It is quite clear that within a culture group, tests do, more fairly and systematically, judge how bright people are. It is equally clear that the further a subject is from the culture and language of the standardisation group the less useful and predictive the test is for him or her.

However, my long experience in this field convinces me that the brighter the subject is the less the cultural differences matter, in psychometric measurement. The higher the flyer, the more you can trust the mental test.

6. THE NATURE OF INFORMATION

The meaning of the word 'information' given in the Oxford Dictionary is 'formation or moulding of the mind or character; training, instruction, teaching; communication of instructive knowledge'. However, the development of modern communications together with what emerges from studies of the brain and the cell have caused a gradual extension of the meaning envelope of the word. I have pointed out that the systems that communications and data processing engineers had to invent for social communication proved to have uncanny echoes in those two entirely different spheres. Brain, cell, computer, telephone network; they are utterly different in every way. Yet in spite of greatly differing origins, scales, dispersions and materials, they fell naturally and immediately into a common linking paradigm. This was the new extended meaning for the word 'information' which has been developing and elaborating all this century. Terms, concepts, algorithms, modalities, from all these disparate levels were exchanged and found to be applicable and predictive. What is nature trying to tell us? Let us see.

Nothing can live without information systems (in this extended meaning). Men know of no information system without biological connections and purposes. Information is a biological phenomenon.

THE COMMON FEATURES OF INFORMATION SYSTEMS

What are the common features that unite these different systems? The first and most obvious point of resemblance between life systems and information systems is that the essence, the message itself and the life itself, is not part of the material universe. In all cases they are form, not substance. Life, at cellular, neurological and social levels and all forms of information, they are all form. They are forms that can occupy a succession of material substrates unchanged.

In the world of data processing, digital computers won out over the analog computers that were their early competitors because the digital pulses get cleaned and squared up at each stage. Computer hardware and software also contain many mechanisms, such as parity checks, to clean up and correct noise. Telephone and cable systems have numerous relay stations where messages are regenerated.

Much less is known about the brain but Lashley's and Penfield's experiments with many others revealed equipotentiality. When brain cells are removed or die others often take over their function.

The fourth feature has already been mentioned. The pattern of coding, symbolism, and data storage is stratified; it has hierarchic levels that are

surprisingly similar, a simple 'string' of 'words' made up of 'letters' from a highly specific and very limited 'alphabet'.

In the cell read 'genome', 'gene', 'np'. In the computer read 'ram', 'address', 'alphanumeric (or code) characters'. In the brain, conscious thought has the same linear sequential presentation, one word after another, and the same hierarchic form, e.g. phonemes/letters, words, phrases, sentences, paragraphs.

Social communication is in sequential linear strings, speech and writing. It also has the same hierarchic structure, a close parallel of conscious thought.

We did not copy the genome when we invented speech and writing. Nor did we go to it for our model when we laid cables, and built computers. We came independently, by trial and error, to the same pattern. Where we find the same solutions we expect to find similar problems.

NON-SEMANTIC INFORMATION THEORY

Before we explore the paradigm I have suggested let us look at the classical information theory. The accepted theory was right for its time. It deals only with the problems of passing signals through channels and avoids semantic considerations. The channel-capacity problems were difficult enough and it was right to solve them first.

Information theory as such probably started with Lord Kelvin's work on transatlantic cables. He was the first to sense that the limit of the information-carrying capacity of a cable was set by the frequency band-width that a channel could accommodate. Between 1934 and 1938 K. Keumfuller, H. Nyquist, and then R.V.L. Hartley gradually developed the formal theory. It was Hartley who first saw the underlying principle that coded information movement via strings of pulse-like signals involved the formalisation of a limited set of signals and that the distinction between the members of the set had to be as sharp as possible. This is the first step towards the reduction of variety and so of ambiguity. He noted that the selection of s pulse signals to represent semantic information, (dots/dashes, low/high pulses) from a repertoire of n possibilities gave a variety of s^n and that the informational 'value' of a string of such signals could best be represented by a logarithmic measure.

His formula was:

Where H_n = *the* quantity *of information.*

$H_n = log2\ s$

Base 2 logarithms were shown to be most convenient for the purpose.

Why was the logarithmic measure important? The essence of the idea of information is surprise. If you are told something you know already, no information has been imparted to you. It is only when what you are told is unexpected that you have learned anything. It follows that the more unexpected, the more surprising a piece of information is, the more its value as information, so improbability is a measure of the value or weight of information.

Claude Shannon and Norbert Weiner brought the theory to its present general form with the formula for a long sequence of signals:

$$H_n = -p_i \log p_i$$

SIGNALS AND NOISE

Philosophers since Plato have speculated about perception and knowledge but it was not until the development of telecommunications that a concept vital to comprehension, noise, was introduced into the field. Even now it seems to be unimportant to most philosophers. 'Noise' was the term chosen by Claude Shannon for the uninformative fuzz that is inseparable from the act of communication in every instance. Whenever a chain of signals passes from one entity to another it has to suffer the fate of everything that is ordered in this particular universe, it has to take its chance of being disordered, of having its degree of entropy increased. Whenever information passes through a channel, it is subject to this process. The remedy is, as I have said, repair, regeneration. The receiving system must have some aspect of expectation, some internal model of what it should receive, and it must be able to recognise a version of this that is imperfect and correct its imperfections before passing it on further. The whole business of perception is bound up with this problem, but our brains are such efficient automatic noise filters that we find it hard to understand that the problem is there at all.

As a simple demonstration of your own power to abstract meaning in spite of distortion I suggest this experiment. First make a guess. Roughly what proportion of a pattern can be obscured before you fail to recognise it? 10 percent, 20 percent?

Experiment 1

Take a piece of paper and place it so that it obscures one line of the type that you are reading. Slide the paper downwards revealing the tops of the letters to find the least amount of type you need to see in order to read them. You will probably find that you have underestimated yourself as noise filter, you need much less than half the line to understand it. You can read through 60 or 70 percent noise.

(True noise is not that systematic and can better be described as fuzz. You would find a similar result if the letters were gradually revealed through a fog of distortion.)

Experiment 2

Look at any caricature. The information you are receiving is far less than 1 percent of what you normally need. Can one explain in terms of information theory how you can recognise the person depicted?

Seeing pattern through fuzz is much more of a problem in the world of classical information theory because the internal models of the receiving system are so much more primitive; pulses high and low, bits, bytes. So it was in that field that the problem was first understood as a result of the collaboration between Claude Shannon and Norbert Weiner.

Shannon's great contribution was to quantify noise in a way suitable to his own non-semantic paradigm. His formula was as follows:

> The greatest quantity H of information that can be transmitted through a channel with a bandwidth W over a period of time T though a mean fuzz of noise D and a mean signal P is given by

$H = WT \log_2 (1 + P/D)$

While this powerful concept is completely correct, predictive and the engineering basis of all telecommunications, I suggest - and I think our programs prove - that Shannon's Law does not apply to semantic information systems. Patterns of great complexity can be detected reliably through very much more noise than would be allowed by the Shannon formula. This was proved by the two experiments above, as will become clear later.

We now have to look at a much more difficult area: the data-capture problem and the fuzz problem of an animal.

INFORMATION NEEDLES IN DATA HAYSTACKS

David Hume said that human knowledge starts with impressions (to him of unknown origin). We would call them sense data and say they are an input from receptor nerve cells to our brain. But what we take in, from 10^8 nerve cells is nothing like the picture we see or the sounds we hear and even less like the symbols and sounds by which we tell each other our impressions. Since we cannot communicate sense impressions, the starting point of science, philosophy and epistemology has this century been taken to be 'propositions'; statements in words or symbols which are supposed to be approximately congruent either with some observation of the world or with some imagined axiom system.

But the words from which speech is made up are symbols for classes of entities, actions, ideas or concepts that have fuzzy, often very fuzzy, boundaries. Words have many overlapping meanings that greatly depend on context.

The early problems of science arose from the fuzziness of the common words they used for communication. Precision was introduced by using the technique called definition. Words were borrowed unceremoniously from the common stock of language and redefined to be less fuzzy in the scientific context. But the words used for the definitions were themselves semantically fuzzy. (Neologisms based on Latin and Greek proved useful because they had an understood semantic background but less of the semantic haze which, I suggest, is necessary in everyday language if social intercourse is not to become prolix and pedantic.)

New paradigms call for new language but comprehension comes more from contextual practice than from definitions. The same words are used in different fields, but they have quite different semantic connectivity networks in different usage. In many cases understanding comes only inductively with constant exposure and cannot be achieved at a stroke.

Although defined and therefore more precise, meanings were an advance, there was a continuing problem because of various subtly different interpretations of the new language. Many of the disagreements and arguments in science are semantic in origin.

Sir Arthur Eddington, in The Nature of the Physical World, introduced an important new idea when he said that the primary informational input of the community of scientists is a vector or profile of meter readings.

Eddington could well have said more. He could have said that the primary input to the brain, perception, was just the same, an array of events like meter readings, a very large one. The accepted view, based on recent neurology, is that all sensory information comes to us as a very large array of signals each of which tells us of change in the intensity of some externally caused phenomenon at some small point in space-time, a sensory neuron. The atom of knowledge and intelligence is a point-intensity in space-time. We learn how strong some tiny action is somewhere, some-when. The informational input element is a magnitude/pixel/micro-second.

A receptor nerve input is a 'meter' reading where the position of the meter needle is represented by the amount of change of frequency of neuronal firing. You read this page because you are receiving many millions of such meter readings in the form of trains of pulses from receptor neurones in the retina at the back of your eyes. Sounds you hear come from the increased firing patterns to the miriad hairs in the whorls of the cochleas in your ears. You observe the sound vibrations as the patterns of intensity of the hair movements as by the rate of firing of many

hair-vibration receptor cells. Touch and all the other senses, including many of which we are not conscious, all work the same way. We get a large array of signals representing various point intensities at own sensor neurones.

MODEL SEARCH

What we need here is a single model that is brainable enough and has enough elementary fit with the facts we know to be useful as a structure upon which a family of developing paradigms can be built. Let me suggest another such model, the one I first suggested in 1958 and which is creating interest now. Take a box and call it a brain. It has many afferent inputs sending in information (in-pointing cones), and many efferent outputs (outward cones) sending out instructions to muscle fibres and organs. These may enter anywhere and emerge anywhere (Figure 3).

Figure 3: The Black-Box brain
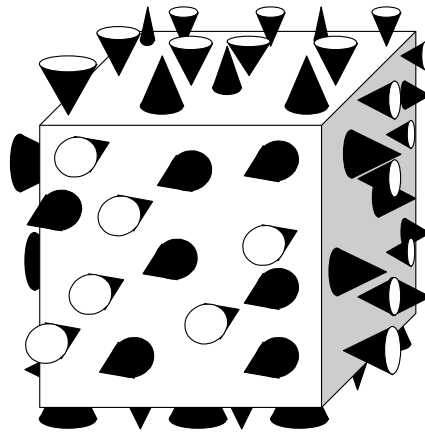
Let us now look at the whole set of inputs. These are the nerves from all the receptor nerves in the body, the receptor rods and cones in your retinas, the receptors at hair roots in your cochleas, the touch sensors in your skin, the proprioceptors on every muscle fibre, and those that are monitoring your bodily functions. Let us give a name to the whole set, the entire input knowledge interface of a brain with the world. Let us call this vast array the *sensorium*. Let us do a permissible topological distortion and draw all these input nerves to one face of our black box.

Now let us take all the efferent, (instruction-giving) nerves to the opposite face (Figure 4). Let us call this the *motorium*; (the entire set of motor neurones). This is a simplification that I believe will aid the sort of strategic level comprehension that we shall need for such a difficult field of speculation.

Figure 4: The Black-Box brain topologically distorted

Motorium

Sensorium
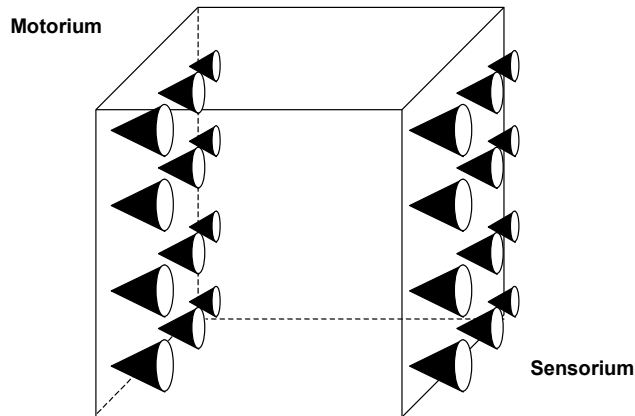
SEEKING FEATURES IN BLACK BOXES

What can be done with brain model based on a 'black box' with input and output interfaces? How does it help our thinking?

We can see that it is going to be a daunting problem but in the light of the more general class morphostat of which our black box is the information centre we can expect to get clues at several levels if the class is a valid one. All morphostats receive and act on information from their environment. The information arrives as a constant array at numerous input ports at an afferent interface. All morphostats transduce this information into complex actions at a very large efferent interface. These actions are purposive, intelligent, and uncausal in the sense given earlier.

To start with the simplest case, the surface of a cell is its sensorium. It has the means to detect the presence of nutrients in its environment and can respond either by forming a bubble to engulf them or by allowing its diaphragm to become selectively pervious. Further, the cell surface is covered with many complex protein 'receptors', vast molecules, scrunched-up balls, made up of strings of amino-acids. The sensor proteins have sites on their surfaces that bind to and thus detect specific chemical messengers that may be signals either from other cells or from the environment. These are used to trigger the cell's response. The signals pass through the cell wall in many ways, then via intermediary proteins to the genome (the cell's central control level). These messages activate or inhibit specific genes or gene combinations on the genome in the correct sequence. The result is that the genome peels off and sends appropriate RNA and mRNA messengers to the layers of mitochondria, and thence to where the ribosomes which are building proteins. Thus the cytoplasm generally and the ribosomes especially represent the motorium in the oversimplified picture I have drawn.

In metazoa the whole chain of ontogeny, cell development, is mediated by the immensely complicated signalling system, by which each cell at a given location learns where it is and by that 'knows' which parts of its program to activate and in

what order. It knows which pages of its instruction book to read and which to ignore. The 'book' is the genome of which every cell has an exact replicate. So the cell's motorium is instructed, what commands to obey and in what order as a result of the perception at the cell surface of quite complex classes of informational input.

Figure 5 shows the information movement in very simplified form

Figure 5: Helpfully distorted view of information movement in a cell
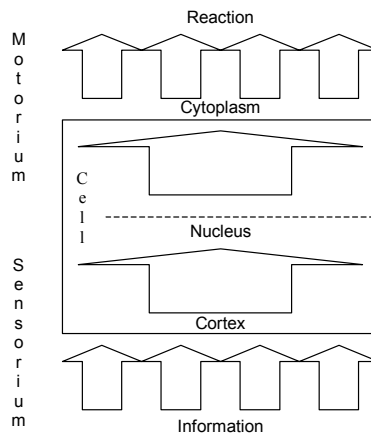
Figure 6 schematically shows the same cell at different times receiving different input patterns of information and causing this to actuate different specific genes. In the 4th diagram we see all systems working concurrently, as they must do in real life.

Figure 6: Sensing and Reacting
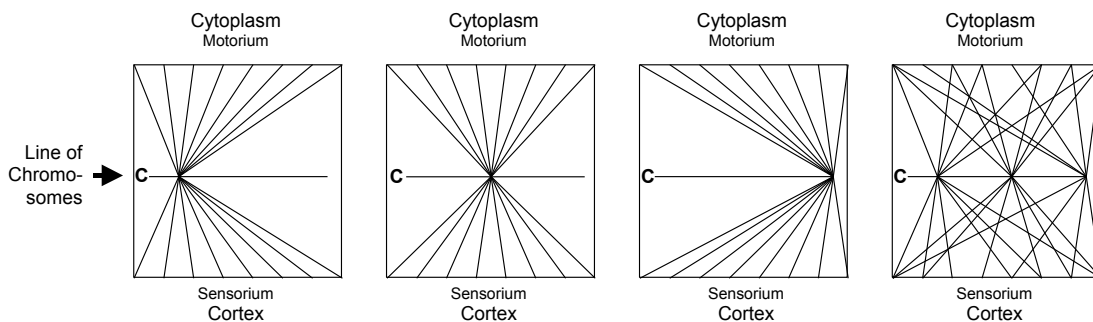
The surface of the cell has been separated conceptually in respect to its two roles: 1) sensing environmental conditions, sensorium and 2) reacting to them, motorium. Information activates different genes (chromosome loci) according to the information from the sensorium, the cell surface.

It is my point that the whole animal, a relatively enormous system of which the cell we have discussed is an infinitely small part (less than a trillionth), has an

echoing system. A neural brain, too, gets signals from its sensorium, deals with them and produces the required behaviour at its musculature.

Figure 7 represents the information movement in a simple metazoon such as an insect. Complex patterns of events in the environment are detected by an array of receptors, patterns of events significant for the creature are detected. Then each triggers its own specific response pattern of behaviour at the muscles and organs instructed by the motorium. (This expositionary picture takes no account of time.)

The case is different but the pattern of information movement is much the same as that we see in every minute cell of the same insect.

Let us now look at the human case, because we have a specially privileged position in this respect. We have an internal inspection system. We have, each of us, some power of introspection. Behaviourists preach that we should ignore such evidence, but I have never quite seen why. So I shall not.

The human brain has a quarter of a billion receptors at the sensorium, the ingoing interface. The movement of information moves into the system and has a great influence on the set of instruction signals that are sent out at the motorium.

Figure 7: Information movement in a simple metazoon
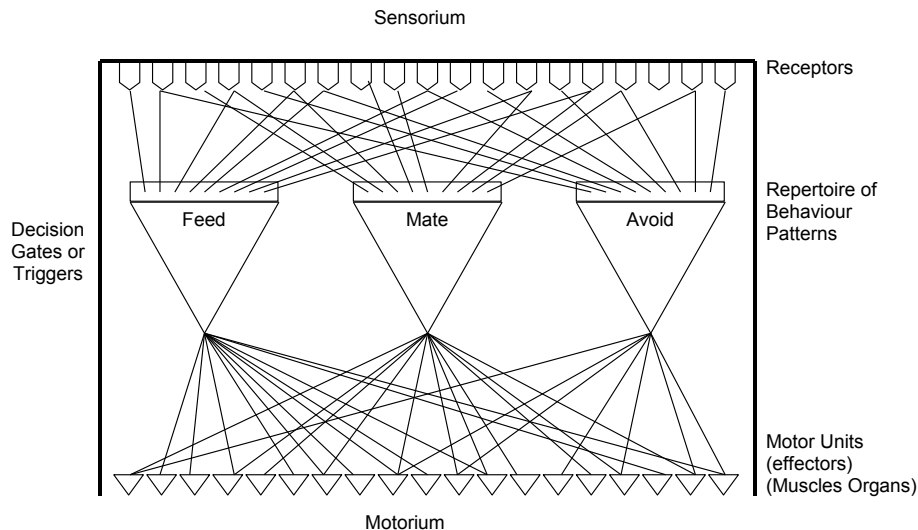
Let us look again at the concept 'sensorium', the entire set of receptor nerves in the body. The axons of these nerves connect these receptors with your nervous system. There is no brainable physiological representative of this sensorium in the physical brain pictured in Figure 8. It is a topological or connective concept. But it is a simplification that I believe is essential if we are to advance the

strategic level comprehension in such a difficult field of speculation. Some physical mapping can be made but it only confuses rather than advances understanding at this point.

Figure 8: The Human Brain
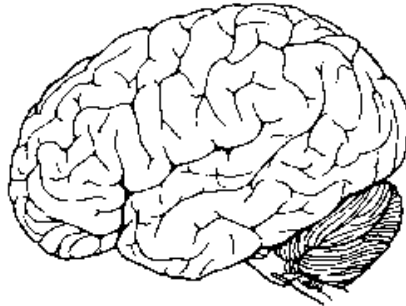
(From Popper, K. and Eccles J. C., *The Self and Its Brain*, p. 287, Springer International)

From this picture of the process of perception we see that the amount and the complexity of your informational input is astronomically huge. If we allow 8 bits for intensity at each point and a view rate of 10 Hz at each of the 180 000 000 receptors, then from one eye we have a data movement rate of 10 terabits/second (10^{10} bit/s)!

There is a popular view that it is impossible to have too much information. Especially in a managerial context, it is taken as obvious that as much information as possible should circulate ubiquitously. Consult everyone on everything is the accepted advice. I have doubts. Here is a point a businessman will see.

The problem your brain faces in dealing with the overwhelming explosive cataract of data described prompts the thought that telling all to all may be the wrong way to deal with information. Raw data is the poorest of ore. It requires massive filtering, sorting, condensing and refining before it can be of the remotest use. The organisation of brains or other order-preserving systems does not require the random spreading of signals ubiquitously but the highly constrained and precise selection, condensation, summation, emphasis and direction of data to specific channels. The incoming problems are *disconnection* and *insulation*, limiting, filtering, selecting, choosing, directing and insulating flows, not elaborating them.

We observe that all that filters through to our consciousness from this great swamping flood of separate signals is a limited selection, a highly abstract condensed set of Hume's impressions. Half a billion nerves are screaming at me ten times a second. Result: I see a part of a green monitor with a few dozen words, less than a kilobyte of information. Within a few microseconds there has been a process of refinement, condensation and relevance selection that no

known computer could begin to tackle. The contingent, the irrelevant peripheral and fortuitous has been blocked and only a many times refined and purified essence of what I have learned, that may be relevant to my life, has been allowed to pass from the tactical level at the periphery and through to the strategic command level of consciousness. So much for the objectivity of knowledge! Terabits are reduced by nine orders of magnitude to kilobits, highly relevant and significant kilobits. This is not an easy task. It is equivalent to finding one needle in a thousand haystacks - ten times a second!

Biological information processing is the business of filtering, sorting and stopping communication rather than making communication. If that sounds wrong consider this. Before learning starts a system does not know which of the incredibly vast number of connections will be required at the end of the process. So all possible connections must be virtually open or potentially so at the start. Learning must therefore be much more concerned with reducing connections than with creating them. You cannot eliminate something you cannot try. Learning is cutting links, not making them. Artifical brains have this problem.

How does the general picture fit our introspective picture of the world. I have described. This vast terabaud, creeping barrage of unique experiences gives rise to a manageable chain of experiences. A limited number of perceptual events composed of a small number of percepts is perceived. We know that we make decisions such 'as walk to the station', which are very simple and brainable relative to the millions per second instructions to muscle fibre groups which are actually involved. There is some type of General Staff translating the strategic level instructions of your consciousness down through many administrative levels to the tactical muscle-pull instructions.

Here is the pattern once again:
- Enormous input array.
- Convergence, recognition of simple patterns.
- Limited variety of central strategic options.
- Divergence.
- Very large array of motor instructions to individual units.

I now probe one that is a much less black box, than the living neural brain. We get a better idea of information movement from this source because it is primitive and therefore relatively simple and open to our view. It is the organisation of an institution such as a business or army, or indeed any organised human group.

THE PATTERN

This is the pattern, the gross architecture, in every case.

1. Sensorium

There is a receptor sensorium that inputs a succession of patterned arrays of afferent signals from the environment. The people as consumers.

2. Conference zone
   There is a layered zone of convergence, summation, filtering. These are many-to-one effects. This is wholesaler factories.

3. Central zone
   There is a line of zone of decision centers from which instruction cascades originate. This is managers.

4. Divergence zone
   There is a layered zone of multiple, diverging, elaborating, one-to-many different instructions. This is departments and such.

5. Motorium
   There is a motorium outputting many efferent detailed instructions. This is the same people in their role as producers in the shopfloor.

Is there a simple conceptual pattern that fits here? Convergence upon a line, a central line, divergence from that line. The convergence and the divergence are both layered, stratified, a poly-hierarchic network.

The pattern is symmetric about the central line; looking at one half we see that being layered and converging (or diverging according to which way we look) it is like a root or tree configuration, but being layered it is hierarchical, organised in successive subordinate grades a normal hierarchic structure does not fit because that is always seen as having a single apex or apical figure, the ruler (hierarch, chief priest). There is a modification of the hierarchical pattern that fits quite well, though, and that is a polyhierarchic network such as shown in Figure 9.

Figure 9: A six stratum polyhierarchic network
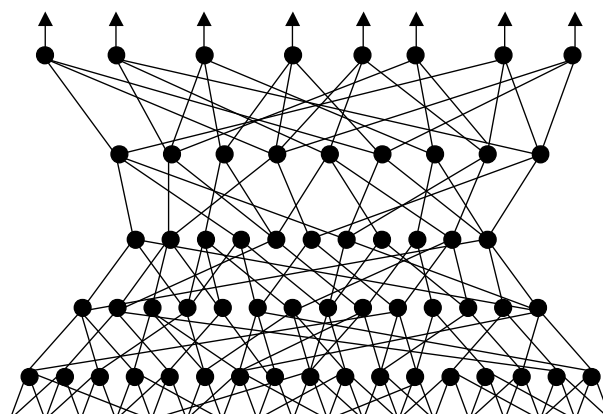
The neologism 'polyhierarchic' was first used by Dr Stanislav Andrzjewski (now Professor S. Andreski) in his book *Military Organisation and Society*. He defined a polyhierarchic society as one where several independent hierarchies exist. I use the word in a slightly different sense and re-define it for our purpose as 'any communications network where the pattern of information movement is such that a plurality of interconnected stratified hierarchies spring from (or move to) the same base'.

Figure 9 shows a network as here defined. Note that there may be (but need not be) cross-interconnections at all levels.

Andreski's concept 'polyhierarchy' provides a hypothetical general architecture for the region from the sensorium to the decision apices. How about the connections to the motorium?

The same pattern, turned upside down fits. We see a converging polyhierarchy from the sensorium periphery to the apical level, and a diverging polyhierarchy from that level to the motorium periphery (Figure 10).

Figure 10: Two polyhierarchic networks joined at the apices
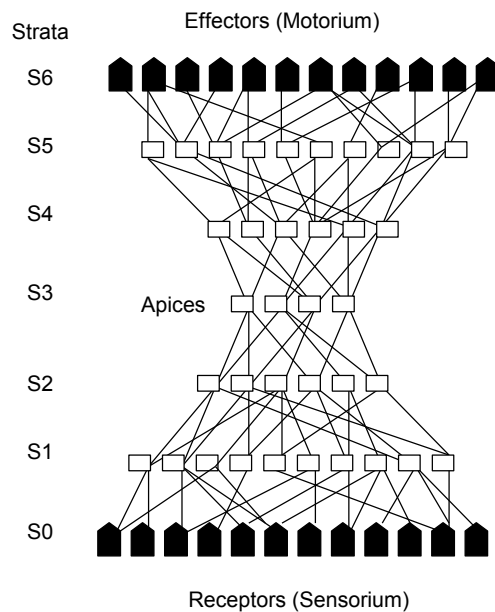
This figure shows a two-dimensional polyhierarchy because paper happens to be flat, but the number of dimensions could be three if the sensorium were an array. Indeed there is no limit to the number of dimensions possible.

Figure 11: A three-dimensional polyhierarchic network
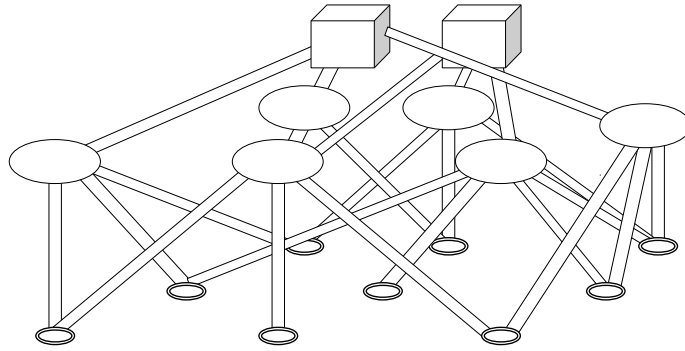
Whatever the tangled pattern of the channels, chemical, neural, inter- animal, the movement of information takes the form of two polyhierarchic networks which are connected at the apices. This fits what we know.

We have a brainable architecture that has a better superficial fit than many others. We have got at least some possible structure into the utter opacity of that dense black box, the brain. We have seen that they may have the same sort of problems. The information movement in them has a similar architecture.

In his or her own experience of the working of a brain each of us has a very good idea of what is going on at the central strategic region of consciousness. We have a very good idea indeed. The best. We know absolutely nothing else than what goes on there. It is all very well for John Broadus Watson to say 'ignore it all'.

The architecture of information movement that comes from introspection fits the pattern already described. Our central perception is nothing like the chaos of coloured dots that our receptors receive and our instructions to muscles are nothing like a million commands, one for each small group of muscle fibres. We make simple strategic level decisions like 'get up', 'go to the station', 'eat the chocolate', 'type this paragraph'. We make policy decisions. Of the administrative detail work at both peripheries we know nothing and care less. The information movement pattern again is convergence from a very large and complex sensorium to a lower variety central strategic region and then massive divergence to the motor region where the details of the movement plan are worked out.

The Patent Case, The Market

In the open case, the market, we can see that the convergence and the divergence happen at discrete stages or strata: retailers, wholesalers, manufacturers. Any human organisation over a certain size there is organisation at several distinct levels. The convergence and the divergence are interrupted, this is subject to a step function. The pattern is hierarchical. This occurs in both directions.

The formal structure, or the authority structure of human organisations. This is often depicted as hierarchic. We are in the context of information and we are talking about the architecture of the information movement. We seek a comprehensible testable pattern which fits in these widely separate fields.

The pattern we see is hierarchic (or rather double hierarchy like a root and tree together). There is root-tree-like convergence and there are strata on both sides of the apical central region. But there is no single apex. There are multiple apices at the centre region.

Our model fits a stable market economy and it fits the same general pattern, but to see it clearly we have to do some unpleasantly severe conceptual surgery. We had to chop the entire population in halves, everyone! Never fear, the dissection is topological and imaginary. It is done to aid comprehension, the motive is not genocidal.

What is the difference between these roles? In the consumer role we are, the whole array of us, a sensorium. We are sending need and desire signals, messages about what we think is required for our survival stability (and pleasure). The signals take many forms, including money and buying orders. We send these into the network. The signals we send, pass into the system just as they did in the other two systems described. The mass of different signals penetrates the system and guides its operation.

Figure 12: Polyhierarchic communication movement in society
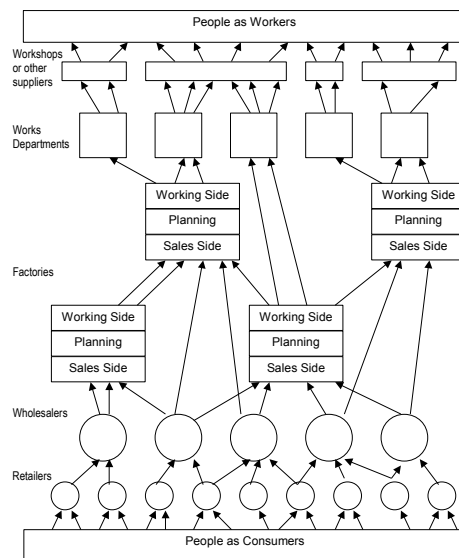
I have now completed my broad brushwork and exposed my bedaubed canvas. Before I venture further to fill in more detail I pause to reassert my claim to

attention by showing that this is a model and hypothesis worth more attention. The model is not falsified by the first few obvious tests we should apply.

7. ORIGIN

No-one is sure where new paradigms, models of reality, come from. Karl Popper rejected the view that they arise from inductive experience and explains them as emergent, creative phenomena, which arise mysteriously. That is to say he does *not* explain them. They might as well be swept under the animist carpet of old, as under the 'emergent' carpet. I offer the view that they arise from the human cognitive activity called taxonomy, perceiving a new valid (useful) class, contriving an advantageous way of classifying inputs. Man and the advanced mammals are constantly engaged in taxonomy, finding useful and convenient ways of bundling together experiences that are each unique but sufficiently similar for some purpose.

Newton saw that the motion of a falling apple and that of a planet were phenomena that could be united. His insight revealed that the two sorts of experience were members of a newly discovered class of experiences, that of objects in a gravitational field. Devising that class was more difficult than perceiving the class 'rattle' from the series of experiences of a rattle to a baby, but it is of the same nature.

The eureka moment, the moment of insight, is that exciting, pleasurable moment when a taxonomic economy is made and two or more whole classes of experience are suddenly united and seen to be one. A relational, or 'inclusive-or' gate is formed. The baby discovers the idea 'man' which unites his brother, uncles and father and separates them from 'woman', the other people who are so similar but distinguishable. Earlier the baby had a similar problem of uniting that set of patterns, all different, which fell on his retina and represented Father.

I came upon the polyhierarchic model for mental activity in 1957 when I had been doing much speculation after an all-night talk with Professor Norbert Weiner and Professor Gray Walter in the Wiltshire village of Urchfont. I was thinking about the similarity of the information-handling problems of cells, animals, and societies, in the context of the class 'morphostat' which I had conceived. As an industrial manager who had observed the growth of a large firm, I had developed the idea of the polyhierarchic network as a paradigm for the shape or architecture of information movement in a firm. It also seemed a way to describe what must happen in the region between input and output in a brain. I was looking at the problem from the professional point of view of an inventor, which I am. I was not asking how it happens, I was asking how I would do it if I had the problem. The seed of the idea may have come from my friend Stanislav Andreski's book *Military Organisation And Society*, which introduced the neologism 'polyhierarchic' in a social organisation context.

The idea came to me with an intense 'eureka' feeling in the middle of the night, with the enormous sense of excitement and wonder that has been described by

others who were or thought themselves originals. I did not remember Andreski's reference until I gave a lecture on the idea under his chairmanship and he reminded me of his earlier use of the term.

From that set of concepts the idea that recognition was analytical and polyhierarchic, that we *learn* to recognise perceptual elements, then assemblies of them and finally whole percepts, arose from the same architecture concept. I began to think about using it for pattern recognition needed for wood grading. This led to much experimentation and finally to a successful product, an automatic machine for grading wood. Wide reading in all the relevant fields has done much to confirm and nothing to falsify the hypothesis ever since.

Influential in confirming my faith in the idea was news such as that of the work described below.

An article, 'Cellular communication', in the *Scientific American* in September 1972 by Gunther S. Stent deals with the various communication systems in the cell and leads up to a description of the work of Steven W. Kuffler and later of David H. Hubel and Torsten W. Weisel on the early stages of perception in cats and monkeys. This work shows that the first few steps in the elaborate process of perception in these animals does indeed conform with the pattern that I proposed. The congruence, the fit between the model and the facts is especially close. There is an informationally economical convergence at several strata towards apices, loci, which 'represent' more and more complex perceptual elements. In the retina itself Kuffler found ganglia that represent small retinal spots. They respond only when there is sufficient (above threshold) stimulus of a small area, and below threshold stimulus to the surrounding area. Thus the 130 million receptor points in the retina are reduced to about a million of these small areas at the retinal ganglia. These connect directly via about a million neurones to the visual cortical area at the dorsal end of the brain. A summational convergence of information and a 130-fold economy is achieved before any signal leaves the retina.

Hubel and Weisel took up the experimentation at this visual cortex itself. They found the same pattern: hierarchical convergence, summation, economy - cells which respond only to assemblies of the Kuffler perceptual elements at one stratum. Cells that respond to more elaborate assemblies of these assemblies at a higher stratum and so on. The gestalt, holistic view of brain function that was the consensus in those days certainly did not apply. It was not in fact until I read of the work of Hubel and Weisel that I really took my idea seriously enough to contemplate writing a book on it. My first paper on the subject was shown to several people at the Philosophy of Science Congress in 1957 and read and signed by Professor Ross Ashby and Professor Kapp, among others. A second paper, '*Cybernetics and factory organisation*', was read to the Theoretical Studies Group of The British Sociological Association on 28 March 1958. They were all well received.

Based on a long chapter on the idea in the context of biological recognition, a highly condensed account was published in The Scientist Speculates, edited by Professor I. J. Good a year or two later. My first book on the theme, Brain, (Davis Poynter) did not appear until 1975. It was followed by a dozen or so lectures to university departments and some experimental work on computers, which I shall describe. I was also invited to, and read a paper at, '*Recognition*', the Cybernetics Conference at Porto Alegre, Brazil, in 1972.

CAN WE FALSIFY THE MODEL?

So much for the unlikely origin of these ideas. Now, before exploring the model further, let me show that it passes the primary tests, it cannot immediately be falsified. We would expect to find evidence of afferent convergence and efferent divergence and of intermediate levels of strata in the three main levels of morphostats posited. I have suggested a model of a social morphostat where the polyhierarchic movement is between nodes that are roles occupied by people. The people in turn have brains that are, in themselves, similar polyhierarchic networks. At the lowest level, the brains too are networks of 'nodes' (neurons), which are again polyhierarchic information systems.

At the top level, in the sociological case it is clear enough that the market system fits the model. Shops summate the need signals. At the wholesaler stratum, and at farm and factory and service and government top level there is further summation and strategic policy scale decision. On the motor, the production supply, delivery side there is a multi-level departmental diverging cascade of instructions to smaller units at more tactical levels. The convergence, divergence and strata are visible. In the Western world visibly, and in the centralist world covertly. A polycentric array of apical control centres, each fed from a wider area of sensorium and each sending diverging instructions to a wide motorium; these can be seen also.

In the animal brain the histological evidence is less clear. We know very little of the pattern of information movement in the tangled mass of nerve fibres, but the primary evidence does not falsify convergence. I have mentioned the convergence from the retina at several stages to the dorsal visual cortex, two great 'cables' of fine nerve fibres from the ganglia at the base of the retina to that cortex. We know the interface between the cortical and the retinal strata. It would falsify my hypothesis if there were more channels in the mass of optical nerve fibres than there are receptors. It is not falsified. The 130 million receptor outputs converge to only a million ganglia and nerve connections.

The evidence that encouraged me most is, as I have said, that provided by the work of Kuffler, Hubel and Weisel already mentioned. To give more detail I draw attention to two articles in "Readings from the Scientific American: Perception: Mechanisms and Models": 1971. The first article is by Charles R. Michael. This

shows that the retina that receives information at the back of the eye has just the convergence and stratification of structure that is required to fit at the sensorium end in the polyhierarchic model of information movement that I posit.

The animal is anaesthetised and the eye fixed. Patterns of light are generated and made to fall on the retina of the cat's eye. Very fine electric probes explore the visual cortex at the back of the cat's brain and detect the actual cells that fire when the stimulus pattern is exposed. Thus the effect on the brain cells of the way the light falls on the receptors is gradually established. Michael's article shows that the vertebrate retina has just the layered structure that is called for and it is of just the type that fits the model I proposed. The connections, the axons, excitatory and inhibitive, of input cells the input end of the receiving cells is a convergent system as is required at the sensorium. There are three distinct strata.

Horizontal cells link the receptor terminal cells that receive inputs from the receptors. (These probably mediate comparison where, for instance what is detected is to be an edge, which involves the detection of a luminance difference). They may mediate mutual inhibition as well. (Our taxonomons will not work without this feature to help the competitive feature search.) At the retina (our sensorium) the receptor ganglion cells detect either "on" or "off" conditions. (Above or below threshold luminance at some spot.) A stratum of bipolar cells connect to the next stratum, the amacrine cells, which are closely mutually intertwined. Behind those two strata is a third stratum, that of the ganglion cells which are also intimately interconnected, mutually and with the amacrine stratum. The converging, summative structure is apparent here even before the information has left the retina.

Having reduced the input variety by a factor of 130 in these three stages the resultant outputs now leave the eyes and travel to the visual cortex. Here the input convergence goes on at a further series of stratifed levels of neurones as can be seen in '*The Visual Cortex of The Brain*' by David Hubel.

He describes how they unravelled the working of the first few layers of this more complex system. The optic nerve fibres 'talk' to several layers of the lateral geniculate body cells. The probes reveal that these cells specialise in detecting particular perceptual elements. They receive messages directly or pass them on to each other. After this sorting, filtering process, they connect to the next stratum, what he calls the 'simple' cortical cells. These are those that recognise the summations of the output of the cortical 'on' and 'off' detectors outputs. It all fits together in just the way called for in the present hypothesis. The next layer is that of the 'complex' cortical cells which detect (fire in the presence of) perceptual pattern details (perceptual elements) that are due to inputs from a converging group of "simple" cells. The first few stages of the process I envisaged are visible. All fits neatly enough. The view that the processes at deeper levels are of basically the same type, convergence, summation, filtering, cannot be dismissed.

Then there is the evidence of the different cell types. Most cells in the areas that are associated with perception seem to be of the type where there are more dendritic, afferent, input connections than axonal, output connections. The perceptual nerve is more funnel-like, convergent, as the model requires. In parts of the brain mapped as motor areas there are many pyramidal cells where matters are reversed. There are more outputs than inputs. If these proportions were reversed it would tend to falsify the hypothesis. It is not falsified. Perceptual areas have more convergent nerves, motor areas have more divergent ones.

In *The Self and its Brain*, p.276, Professor John C. Eccles and Karl Popper say 'In attempting an analysis of voluntary movement and control, it is immediately evident that there are many hierarchical levels'. That these could be mono-hierarchic, hierarchical up to a single apex, is to conceive. We seem to be left with a diverging polyhierarchic shape for the motor information movement, which fits my hypothesis.

In 1958 I suggested that the process of recognition must be analytical. Perceptual elements, spots, lines, edges at various angles, are recognised at any location at the sensorium, then at the next stratum assemblies of these and at the next assemblies of assemblies and so on until at higher strata whole percepts are detected. Above that, events, whole situations calling for response, are detected and trigger motor response cascades. I opposed this polyhierarchic connectionist model to the gestalt model that was ruling at that time. In a holistic paradigm it was posited that patterns were recognised as a whole, not built up as I posited. My hypothesis was seen as connectionist and as such not followed up. Connectionist hypotheses and indeed all 'physiologising' (nerve net thinking) had been ruled out because many animal experiments with mammals as well as observations of human brain injuries showed that a brain could suffer gross damage with surprisingly little loss of function. A paper by Minsky and Pappert was often quoted as a complete refutation of my views, But more recent work has been very much on my lines fully connectionist.

To those with the paradigm of a clockwork world this was like expecting a telephone exchange to work after a bomb had gone off in it. Much was made of equipotentiality, the way in which nerves seem to take over each other's function as they were injured or died. (Thousands die each day apparently.)

However, telephone exchanges and fuzzy logic, biological brains are quite different. My suggested answer to the problem, multiple channels, redundancy and inductive repair, cut little ice then. The importance of the fact that the model worked on an evidential rather than a logical basis and therefore was tolerant of noise and error was not fully seen.

Analytic, Connectionist Recognition Confirmed

How about the motorium? Would it falsify my hypothesis if we failed to find similar loci in the motor regions of the brain at which animal behaviour could be excited?

This question is more complex, the recognition of simple pattern elements is almost instantaneous but complex trains of action have a time dimension. A mammal behaving in a skilled way is getting many feedback signals many times a second during muscular action. Proprioceptors on each muscle fibre are reporting in, with many other situation reports from the eyes and other senses. We cannot expect to be able to excite some centre in the motor areas and get a behaviour train response if these progress-checking cybernetic feedback signals are missing. However, it would falsify my hypothesis if we could not make a muscle fibre twitch at one level and a muscle move at another by exciting various nerves or ganglia in the motor areas.

The work in Toronto by Professor W. Penfield showed just these effects so clearly that he and his followers have mapped out a lot of the lower levels of the human motorium. He could repeatedly excite the same vocalisation as speech, twitches and jerks but not skilled behaviour. He did this by stimulating specific brain cells, loci. This was a vindication of connectionist theory and consistent with the polyhierarchic paradigm.

In areas associated with memory he found that by using electric probes he could trigger long memory sequences. The subject, usually someone having a brain surgery operation, would, while still conscious of his or her surroundings, describe in great detail a lost memory of an often long past event. Memory works on the sensorium side of my posited apical region at the cortex. It is also autonomous. It requires no continued accompanying sensory input. That is its essence. Therefore that discovery neither confirms nor falsifies the hypothesis but it fits very well with the model for memory that emerges from this train of thought, as will be seen later.

And at the level of the cell?

In the cell we have a sensorium; it is the whole of the cell surface and its covering of many special proteins which act as sensors or receptors. These very words are used in that context by those in the field. We have a visible central decision layer, the apical level, the genome itself, the genetic memory where the genoplastic learning called evolution is stored. We are beginning to learn that there is a very complex flow of chemical messengers from the 'sensor molecules' on the surface to the genome and then from that to the motorium. The protein factories themselves are at work making the great variety of complex proteins that build and rebuild and repair the cell's form. The features that take this role are the ribosomes, which crawl along strands of RNA copied from the genome and carry out the instructions delivered by them and the mRNA. Here the many specialist proteins are built up, as long chains of amino-acids which finally

scrunch themselves up into the sheets or balls that have the wonderful array of precise properties required for that highly improbable entity, the living cell.

These proteins are the enormous complex molecules (structural or informational) that carry out the cell's functions. On their way to the ribosomes (the motorium itself) the chemical messengers from the genome (mRNA, etc.) pass through what seems to be a stratified structure in the mitochondria. We are beginning to understand the many complex tree-trigger like processes that occur on the way. We know that there are multiple interactions, inhibitive and exciting, which is a convincing echo of the interactions of nerve cells in brains and people in institutions. The way the double polyhierarchic network applies to the nervous system when it is elaborating behaviour patterns is obviously different from the way it applies in epigenesis, the diverging, specialising development of the great variety of differing cells in the complete animal. But the *pattern* is the same, straightforward, diverging, hierarchical.

In the metazoon the diverging multiple cascades are of *cell* types. Each has an exact duplicate of the whole 'brain' or apical instruction set. The descending motorium hierarchy is the diverging hierarchical development of cell types from multiple gene combinations on the apical genome. In the polyhierarchic *shape* of the *neural* information movement that develops in each animal there is no replication of the whole apical stratum. The central one (the brain) is used and the posited polyhierarchic information movement pattern is built up ontoplastically in each brain by local changes in connectivity such as those which have been detected in learning at synapses. It is the shape of the movement of information that must necessarily be the same in each case. Changes at synapses are known to happen in learning, which is required.

The morphostasis paradigm arises from the congruence between the features of biological information-handling systems at the various levels, cell, animal, society. It predicts this congruence because they are members of the same valid class, a purposive class of entity that has the same goal and problem, preserving form, improving form-preserving power. If the hypothesis were false, it would be surprising to find functionally similar elements and features at all three of these levels, which are so utterly different in scale, time, complexity and traditional function. That is my main point. I have shown that the general features, the sensorium, and the motorium, the central apical decision centres, seem to have echoing versions in each level, cell nucleus, brain, and society. I have also shown that as far as can be seen the movement of information and instruction in each case is afferent convergence, and efferent divergence. I have pointed to other features where the same pattern fits at these markedly different levels. We find something corresponding to *syntax* at all levels and there is broad syntactical congruence. At each level there is language that is *symbolic* in that it is not directly causal in the sense we use the word 'causal' about material systems. Features and events trigger, rather than cause, others. Further, the nature of the language is similar. It is in all cases strings, not arrays or zones, simple linear

strings. In all cases the strings are also hierarchically analytic. Letters, nucleotide pairs; words, codons; sentences, genes. The sociozoon and the protozoon match. The intermediate case is not yet sufficiently understood but computer simulations which ascribe meaning to a class of loci do perform artificial taxonomy, learning unknown patterns.

How about components? We find the nerve, a sort of input/output deciding gate (in computerese), which is matched in the sociozoon by the committee, council, Cabinet, Board. Both deal with many informational inputs and make instructional outputs after some kind of steplike decision has been made. Both have binary inputs. In society they are called commanding and forbidding, speaking for or against. In the brain they are called excitation and inhibition. In the protozoon the language is different but the function the same. There are *activating* or *promoting* catalysts, genes and protein messengers and inhibitive ones. If there were no such suggestive echoes and congruences it would falsify the hypothesis, but we have them. They do not establish the hypothesis but they suggest further study.

The Baby Taxonomist

I have done some experimental work on this and can report the results. I was in my laboratory, the World, the other day, when I turned those sophisticated visual class discriminating instruments, my eyes, on to a baby in the act of learning to know and to know how to deal with that class of entities the generalization loci of which I can excite in your brain by using the words 'chocolate drop'. Learning primitive taxonomy. I produced what was to the baby a strange object and the baby's incessant random exploratory movements stopped and his eyes focused on this strange new entity that had swum into his receptive field. With some apparent difficulty and not without mistakes he brought those unfamiliar instruments, his hands, to bear upon the task of getting the entity into his control. He then held it up and, as well as he could he turned it over, twisted it, moved it around and studied it closely for over a minute. Every part of his retina was receiving many images in rapid succession. Then he transferred it to his mouth and explored it with his lips and tongue before it disappeared. A second chocolate drop produced a similar reaction but the inspection time-span was shorter. On the presentation of the third chocolate drop the exploratory phase was already much diminished. Within ten minutes it was evident that I had, at least temporarily, while the appetite lasted, created a chocolate-eating automaton. Every appearance of the stimulus was rapidly followed by the immediate and identical response. What had happened? One particular exploratory effort for the baby was paid off. The baby has learned to know the chocolate drop as a phenomenon and he learned to know it before he was rewarded with a pleasant taste. In quick succession he presented all parts of his retina with stimulation in the form of those elements and assemblies which make up a chocolate drop in its various presentations. Its brown colour, its curves and planes, its feel, its smell and taste were successively and slowly experienced. Messages from widely separated and very different sensory instruments found

their way together to trigger an automatic response for a period of many minutes. The connection, or set of them, was made, ready to be activated at any time in the future when the appetite for the sweet things should be strong and the attention suitably directed. The child had not only learned about the class 'chocolate drop', its uses and advantages; it had learned to see it. Baby learns taxonomy, distinguishing unknown pattern on its retina.

The Simulation Test

The architecture proposed in the hypothesis could of course be falsified if computer simulations of the described elementary components and pattern of information movement persistently failed to produce acceptable results.

It is clear enough that the primary task of a morphostat at the sensorium, its interface with non-self, otherness, its environment, is to detect and interpret three classes of input, the insignificant, and the two types of significant events, the useful and the dangerous. The significant types of detected event or inputs are the dangers or opportunities, the threats or promises calling for motor choices. If the morphostat makes wrong choices it will turn into a corpse, an entity that changes its form radically with time; it will rejoin the Pleroma. It dies and rots away, resumes change. Some of the signals to which the morphostat will need to respond correctly will be simple analogue ones, one pointer-readings, like temperature. But most significant inputs (those requiring volitional response) are likely, in more advanced animals to be complex: a vector, profile, array or pattern of point intensities, a pattern of light on a retina, or of vibrations on many cochlea hairs.

Can we make computer programs that do what the cell, the animal, and societies do; learn by experience, not programming?

They must do automatic taxonomy; 'learn' to distinguish very noisy unknown significant patterns. I have written a program with algorithms that have been tried and that failed to falsify the paradigm. They are simple and primitive but they show that it is possible for a program that simulates fuzzy logic nerve nets to learn to classify unknown noisy inputs with a low error rate.

Concurrency can be simulated in serial time models on traditional computers and one of the large group of my readers who were interested, James Cherrill, wrote a very successful program which could perform the automatic taxonomy which I posited as the first step towards artificial intelligence. There are now Fortran, and Spectrum Basic versions of this around and being worked on by the growing group of enthusiasts who have become interested in this approach. All models and experiments so far amply fulfill the simple prediction that a system of competitive gates as described below would be able learn to sort out and allocate or capture unknown patterns through up to 85 percent noise. They are available.

The sort of programs needed to test the model is not logical or deterministic. We do not use logic where what is required is judgement. Judgement is evidential-what a nerve, a committee or a scientific consensus does. It is what the semantic gates in fuzzy logic computers do. Chain-logic is what Von Neumann computers have to produce. That is their strength and their limitation. In an indeterminate world we sometimes have to use methods based on indeterminacy.

THE DEFINITION PROBLEM

I have mentioned my distrust of formal definition. In the context of the last section it can be seen that definition, as a method of establishing meaning, simply does not address the vast filtering and abstraction process that is needed to establish both the meanings of, and the fuzziness of the meanings of, the words used.

As a way of establishing exact meaning, definition is a counsel of perfection. We use fuzzy words to reduce word fuzz. We try to capture meaning in a precise net with fuzzy imperfect words.

Definition as a science tool goes against language practice. We learn words at first entirely and later mostly from their repetition in a revealing context. Only rarely do we look them up in a dictionary. When we do we have to use context to select from a confusing variety of meanings. A definition cannot possibly call up the many rich skeins of semantic connections evoked by words presented in different contexts. In contextual learning we have richer, fuller understanding. The receiving brain has to gradually make up its own copy of a whole semantic network with all its subtle connectivities.

To visualise such a network we must think of a three-dimensional connectivity arrangement. Imagine a volume of space containing many scattered boxes that are interconnected by a network of a myriad lines, and in each line a ball showing a positive or negative magnitude (Figure 13).

The boxes are the words of a language, and the numbers in the balls are the semantic connectivity of the words, positive, negative, weak or strong. This is the physical basis of thought, semantics. The connectedness of ideas.

This crudely, is my view of the connectivity within the brain. Models of semantic connectivity, of understanding, of thought, should be tried.

Words connected together in strings can be made to transfer complex ideas from one brain to another if the connectivity system in the receiving brain has sufficient similarity, if the owners of the brains 'speak the same language' in a wider sense than knowing the definitions of the string of words. The receiving brain must have a similar out reaching connectivity network, a thorn bush of semantic tendrils reaching out all along the word tracks.

This is where attempts at machine translation have run into trouble. It is easy to find cognate words in another tongue. It is figuring out what is in the surrounding haze of context words that has been the problem. We have to have a replicate in our own brain of both boxes, (words) and numbers in the balls, (connectedness) before we can understand a language well.

Figure 13: Semantic boxes and connecting balls
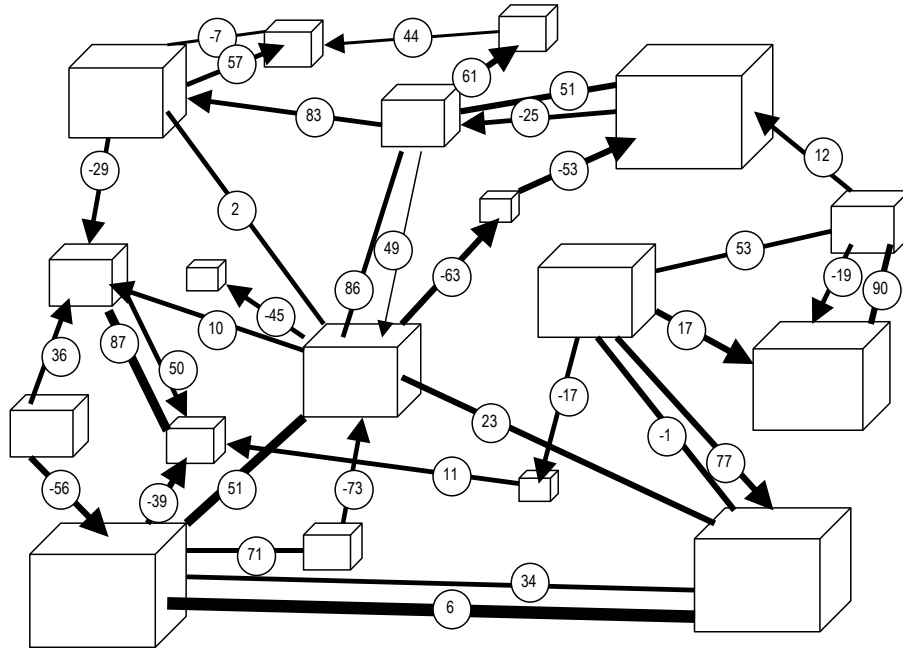
The illustration shows the kind of connectivity network envisaged in the central (or apical) zone where relationships between percepts and concepts are represented by connectivity links (as earlier defined) positive or negative. These central, strategic level nodes (boxes) are seen as connected by links that have changeable weights (balls), plus or minus. The nodes (boxes) have variable sensitivity to each input and variable firing thresholds (sufficiencies).

This is the model truth that lies behind Kuhn's views. We cannot get a paradigm from a definition. We can only get it from inductive contextual exposure.

None of that means that definitions are never useful. But if definition is limited as a way of transferring knowledge of classes how *do* we do it?

CLASSIFICATION

The classification of inputs is, as I have said, the primary morphostat interface problem. Let us start at the beginning with baby learning. How does Baby learn to use nouns to indicate objects? I argue that every perceptual input is unique because of their enormous possible variety. If we allow only 10 states (fire rates) for each of the 200,000,000 receptors in each eye, then the variety or number of possible input patterns would be 10 to the power 200,000,000. This is a largish number. The (*ab initio*) probability of any one input is the reciprocal of the variety

of inputs, which makes the exact repetition of an input a miracle. No retina could possibly receive exactly the same input twice and no visual cortex of a brain ever received two afferent inputs that were alike. In the ear no stirring of cochlea hairs is ever precisely repeated. The other sensory inputs are all the same. And they are not attempting to recognise objects, they are attempting to recognise events. We must start from the axiom that every sensory input is unique, and normal computers do not allow mistakes. Brains have to do better.

In learning to sort out unique bits of pattern in the world that are likely to be significant to it, a baby's brain has to classify, to clump, to treat as alike, inputs which are enough alike for some purpose. The parts of the incoming impressions of the world that represent percepts fall into necessarily fuzzy classes to which words are attached (e.g. persons, houses, chairs, spoons). The members of these classes vary widely but have sufficient mutual resemblance for some limited human purposes. This is the essential, fundamental and unavoidable strategy of baby and of any other intelligence, any brain, in dealing with the stupendous data bombardment and ultra-astronomic variety of its inputs. It is only when one ventures into the tough and tricky world of pattern recognition that this becomes really clear.

For each noun, Baby must conduct the inductive building up of an internal model, in some form with which input patterns, each unique, can be matched. This process of learning to recognise objects is not remotely like instruction from a string of understood words such as a definition.

Recognition

In the presence of a cat turned any way, at any distance, the baby's brain looks at a unique particular configuration of shapes somewhere on the sensorium and asks whether it has sufficient of the features of the percept class 'cat' to be a member of that class or not. If the details, shape, size, colour and other qualities of the features of the impression have enough resemblance, they come within the tolerance bounds, the permissible deviation, the permitted fuzziness (or, as I shall prefer to say, exceed the sufficiency) of their internal models then it is recognised. If *enough* of these '*cattish*' features are present then Baby's brain decides that it is cat; otherwise not. All classes are, must be, fuzzy and so all perception is at base only exact enough for some limited purpose. It can never be complete or exact. (The object comparable, input pattern, can never be completely congruent with subject comparate, the mask, model, template within the brain.) If the brain has *sufficient* confirmatory evidence and *insufficient* contradictory evidence, a binary decision is made. Baby points and says 'pussy cat'. Baby has *recognised*. Millions of local decisions have been at receptors in Baby's pretty blue eye: 'enough stimulus to fire? Yes/No' 'How intense a pulse-train?' These have been transduced into one 'yes/no' decision, 'cat'. In experiments we can simulate learning to see by varying connectivity and the 'sufficiency', which is simply a threshold value which must be exceeded before

an element of the system will actuate, or 'fire', as the neurologist would say of a nerve.

Sufficiency And Tolerance

Sufficiency is the complement of tolerance. If we take the absolute congruence or match of subject and object comparates as **1** and zero congruence as **0**, then where **s** = sufficiency (or acceptable degree of match) and **t** = tolerance (or permissible deviation from match) then

$$t + s = 1.$$

(A tolerance of 0.1 is a sufficiency of 0.9.)

The word tolerance is used in engineering with the meaning 'permissible deviation'. When, as their measurements improved, the early mechanics found that absolute accuracy was an impossible goal, they began to specify the limits of error by fixing a plus and minus tolerance to every measurement on the blueprint. These figures limit errors oversize and undersize. Further progress in engineering could not be made until this was done.

It is quite central to the business of understanding the world that a similar principle should be seen to prevail in the process of cognition. Our kilogram or so of brain can only deal with the immensity of the universe by the strategy of filing events away in classes or sets where the members are sufficiently alike for us to have similar expectations of them. A class that has a sufficiency of I or zero tolerance can have at most one member and is most unlikely to have even one. As we decrease the sufficiency, the number of members grows with the fuzziness and eventually we have a class that is encountered often enough to make it worthwhile for the brain to retain (in some form) the subject comparate (internal model) that is required for the recognition of its members. If we continue to decrease sufficiency we shall eventually reach a class which will be so all-inclusive and have so many members that it will be useless for the vital purpose of helping living things to predict and optimise in the world. Conversely, if we make sufficiency too high we shall fail to recognise the fuzzier members of the class and lose by that failure.

The need is for a good balance, a sufficiency set so as to capture the fuzziest member of the set of inputs that the external object produces; but not inputs produced by non-members of the class. Errors of non-capture and false capture must be kept as low as possible. Neurological research has not revealed how the brain forms its subject comparates and adjusts their sufficiency, but the polyhierarchic paradigm seems fruitful and makes test by simulations possible. Let us look again at the way the baby learns. The mind of the baby starts by delineating very broad general classes with low sufficiency, such as 'face'. This recognition is apparently inbuilt, for the neonate smiles at any crude model.

However the child soon begins to increase sufficiency, becomes more discriminating, and therefore divides the world into more, less fuzzy classes, 'Mother's face', 'Father's face'. Its perception gets more fussy and less fuzzy.

I have suggested that semantic (or evidential) connectivity can be represented by the weight of connections between concepts or percepts. This seems to be a model for the central strategic region of the network. I now suggest that the semantic input to a recognising gate (nerve or nerve-assembly) is the summed 'weights' or intensities of its inputs. This must pass some threshold, exceed some sufficiency, for recognition or classification to occur.

Truth Is Fabrication

In exploring the polyhierarchic model for perception, I start where knowledge does, at the input interface, the sensorium, and we see that there is no pure cognitive truth. There is only fabrication. What we see as truth is simply the sort of fabrication within the brain that has an acceptable degree of congruence with a reality that cannot be directly perceived. The way we know about that congruence is another story. We cannot measure it because we cannot know the objective (external) comparate directly. But we have many crosschecks, tricks. Philosophers have suggested induction, but Karl Popper rejected it. What else is there though? Exact repetition of events is not given to us, but percepts and their elements, within some fuzz boundaries, are repetitive, we and our computers can correct our models for noise if we receive enough exemplary inputs. And there are crosschecks and informative mismatches which enable us to measure congruence by inference. The ground plan is this: that fabrication, that subject comparate system (mask, model) which gets the highest weighted score from a sample series of percept presentations and which also has a threshold which brings errors of exclusion and errors of inclusion lowest and most easily shared is the best, the nearest to truth unattainable.

How do we classify?

We therefore have to classify our incoming impressions to reduce their infinite variety to brainable dimensions. We have to clump (I. J. Good) or chunk (Richard Gregory) analysed bits of experience and treat the members of a class of inputs as if they were what they are not, identical. This is a matter we shall have to clarify a great deal before we can get far with creating non-biological intelligence. Classification in the primitive animals is genoplastic. They are born to classify. Animals can learn by trial and error, with memory records and pain and pleasure as the teachers. Ontoplastic intelligence operates on every individual and it enables each one to snuggle into its own opportunistic sub-niche.

Homo Sapiens is the species which is by far the most gifted with ontoplastic intelligence. He has some of the other sort, more than we used to think, but less than any other animal. Man has softwired tendencies, drives, strategic goals

which set his optimising 'shooing' pleasure/pain system, but few conscious hardwired knee jerk type responses. He has thousands of unconscious ones, of course. Try to stop sweating in a Turkish bath.

At the conscious level, how did you, as a baby, learn to sort out the incessant creeping barrage of sensory impressions with which you were bombarded from the moment you opened your tiny, pretty eyes? Judged by your behaviour, you could make nothing of your input at first and your responses were limited to grasping, sucking, smiling, crying and soiling napkins. Your ontoplastic behaviour, the uncoordinated movements of your little limbs, could have been programmed from a random number table. It was experimental trial-and-error behaviour that was all error at first. You had an almost completely incomprehensible input and a largely random output. It would seem like a pretty bad start for a life it if had not turned out so well long term. Can we illuminate what was happening from our model? Yes.

You had to classify your inputs despite their enormous fuzziness. But before you could even do this there was another task. You cannot recognise patterns unless you have some idea what patterns there are. You had, in some fashion, to elaborate the internal comparates of useful percepts. (Kant's idea that the pattern masks were all there to begin with is too great a strain on our credulity, there are too many recognisable things that our ancestors cannot have known.) Before we allocate perceptual inputs to classes we therefore have to find out what classes there are in the first place, to which is called taxonomy in biology.

Your first task was do-it-yourself taxonomy. You had to sort out your inputs. You had to analyse experience, break it up into useful chunks that enabled you to build up expectations. You had to classify but first you had to decide what classes you needed to know about at that stage. You had only three clues. The first was repetition. You had to latch on to iterative bits of experience. Pain in its various forms from discomfort upwards was your second teacher and your third was pleasure. But every input was unique! There was no repetition! Well then, it must have been repetition within some sufficiency limit. You had to learn those sufficiency limits at the same time as you analysed the input frames and learned which bits were important to you.

Which bits were repetitive? How repetitive? Which bits hurt? Which made you feel good? You started your own personal adventure in d.i.y. taxonomy with only crude stick and carrot teaching, yet by the time you were two you were managing your little world a bit. Enough to get a cognitive input via speech to help you. You did well to sort it all out in the time at your disposal!

Epistemological Classes And Ontological Classes

Epistemology is the science of knowledge. In the influential book, *The Nature of the Physical Universe*, Sir Arthur Eddington emphasised that knowledge could

not be cleaned of its epistemological character and nature. Real objectivity is an illusion. His illustration was this. The view a fisherman takes of the size of fish depends on the mesh of his nets as well as on the properties of the fish. Our sensory organs, our scientific instruments and the hardwiring of our brains between them decide how we shall view nature. Nature then tells us a version of its laws that suits our set of probes. Choosing the method of inspection and primary isolates, classes, subject comparates, was really choosing the laws too. This explains the anomalies where different disciplines interface. It explains inconsistent dualities such as wave/particle duality, and why we have to have a revolutionary revision of analysis and classification when a new paradigm becomes necessary. In a very simple form we have encountered just such revolutions in the programs that are being devised to test the present hypothesis, the automatic taxonomy programs. We create a world with a set of patterns to be distinguished. We present them, as nature does, always in a fuzzy form to the brain simulator program. The program inductively works out which are the most distinguishing elements and assemblies thereof and 'learns' to sort them out, however fuzzy (or noisy) the succession of exemplars. If after many iterations, when the brain has learned its world we now add a new percept 'class' and present its noisy versions, it throws the brain into confusion for a while. It has to revise its scheme throughout to accommodate the new 'discovery'. It often has to give different weights to distinguishing features, in a surprisingly radical revision. The process of learning about the world according to the polyhierarchic paradigm of learning involves analyses into classes that are known, each with *fuzz* that is also known.

The scientist working from definitions can easily communicate the method of analysis and define the classes. What is much more difficult is to pass to others the fuzziness, the sufficiency which is the necessary companion of every class. That only becomes apparent to the insider group within the discipline, as Kuhn's brilliant insight revealed. Fuzz sense frequently comes only with experience, but it often cannot be directly communicated in words. That is why disciplines are mutually insulated and why the work of the venturesome generalist like me is so very difficult. Experimental experience with the computer simulations of the learning paradigm I present produced one important distinction that I must introduce here before we move away from the problems of classification.

There are classes of percept that are natural in that the members of the class tend to cluster about a Gaussian mean or norm. Class divisions fall naturally at the tails of the normal curve. Biological classes are examples. The more the class member deviates from the class median the less frequently the deviant is found. We could call these '*ontological* classes', and their limits *could* be learned merely from their relative frequency without punishment or reward. I believe that some primary human learning, that of perceptual elements (like edges at various orientations, for instance, in the visual field), is of this nature. (Such elements may even be hardwired, genoplastic.)

The other sort of class is the *epistemological* class, whose boundaries are set by a linguistic convention. They are fundamentally arbitrary. The visual class indicated by the word 'yellow' exemplifies. That class is defined by a broad band of light frequencies whose borders are fixed by convention. The class boundaries have to be taught to any intelligence by other intelligences who know the convention. However, before an unsophisticated intelligence has learned the language, before it can even talk to a sophisticated one, how can they communicate? The teaching intelligence has to give pleasure or cause displeasure in some form. The learner intelligence has to have an initial notion of good and bad before the convenient semantic borders and the fuzz limits of epistemological classes can be fixed, but once there is a common language moving between morphostats they can communicate class boundaries and fuzzinesses without too much resort to a pain/pleasure continuum. The built-in pleasure for a child of the mere activity of learning correctly, is very visible. Biologically it has to be.

It seems, parenthetically, that one of the measures of 'progress' which is so urgently denied by many academics like Gould. On the *Scala Naturae* in evolutionary terms is the reduction of the severity of punishment. For simple creatures the sole punishment for 'wrongdoing', un-biological behaviour, is the death penalty. Higher up the sentence may be reduced to torture (pain) and higher still disgrace, disapproval and so on. At the highest level communication can be almost (never quite) free from an incentive element.

THE DYNAMIC NATURE OF RECOGNITION

When I speak of the subject comparate or internal mask or representation of a percept or concept in a human brain I am not suggesting that it is in the form that is in any sense visually like the object comparate that excites it. To fit this paradigm it has to be a dynamic connectivity network in which the pictorial forms have been transduced or coded into locus and excitation magnitudes. The internal representation is, I believe, under constant revision and adjustment both as to its norm and the boundaries of its fuzz. With time it converges asymptotically on improved congruence with the central mode of the class of object comparates that it has specialised in recognising. As excitation moves inwards from the sensorium periphery it transduces arrangement into locus and intensity. For example, one cell or locus above a tiny area of retina represents a left-bright vertical edge. This tiny geometric perceptual element is now represented by a single-point locus in the brain. All the other perceptual elements that may appear at the small area have a point representation. That set will be enough to deal with any class of more elaborate percepts that may be perceived, there is a vast economy. The same thing is repeated at the stratum above but here an assembly of small perceptual elements, a small visible figure like a letter is represented by a locus or set of loci at the senior stratum. Proceeding inwards the subjective comparate (mask model) becomes a pattern of firings at loci that have no visual relation with the object comparate that caused them.

The economy continues at every stratum. There will have to be very large numbers of loci which represent the same element, fewer that represent more central combinations, and so on.

I interpose a caveat. It is just possible that the reader may be making vain attempts to fit my topological map to the physiological map of the human brain. I do not think we are going to get the basic functional ground plan right if we insist on visible physiological-functional concurrence at so early a stage of exposition. Once I have communicated a functional understanding we can take our double polyhierarchic network in our very flexible topological mental fingers and, just as we pulled one kind of connection to one, and another to another, face of a cube, we can pull them all to the known sensory, the motor and the other mappings of particular brains. We know that what I have defined as the sensorium is divided into a number of separate input sensory arrays which are roughly defined by the senses, sight, touch, taste, smell, hearing. Each of these has different organs and modes. We also know that the information from all these organs is pooled or centralised in some way so that any complete experience can have contributions from all these widely separate organs. My model arises from a functional hypothesis, not a physiological one. I ask this, "Seeing what it is and what it is doing, what shape must the information movement have?" When we complete the model we shall have to see how well it can be made to fit the physical thing.

At The Motorium

The polyhierarchic architecture for information movement at the sensorium indicates a stage-by-stage building up of pattern recognition from the base as learning proceeds. What, according to the model, is going on at the motorium side of the central apical region as perceptual learning proceeds? To learn by trial and error you need trials. There must be behaviour of some kind before learning can even start. Ontoplastic learning is impossible if genoplastic behaviour is inflexible (as it is, largely, with insects), so there must be random penny-toss input in learning behaviour. Behaviour must be, at the very beginning, independent of sensory input, pure experiment. There must be whim. Our model is a hierarchic process in which a skill is built up from skill- elements, small elements of behaviour at the periphery. At the next stratum, subskills, combinations of skill elements are connected up, then at successively more central levels more elaborate skills. All these may be employed in more strategic behaviours generated at more central points in the network. As at the sensorium, the learning process proceeds inwards from the motorium periphery towards the apical centre. The information from the sensorium (and the appetite system which can be seen as part of it) has the power as learning proceeds to trigger larger, more complex, longer lasting behaviour sequences.

It is too early to decide what physical aspects of the brain are cognate with the elements of the model. We could interpret the 'nodes' or 'apices' of the model as

represented physiologically by individual cells in the brain. However we could equally represent them as what Dr D.O. Hebb called 'cell assemblies' (*The Organisation of Behaviour*) or the 'modules' of cells that Professor Eccles posited in *The Self and its Brain*.

Now we have a gross model of the brain. How does a computer imitate it?

TRIGGERS, SWITCHES, TRANSDUCERS, RELAYS, GATES, NERVES

Tracing the idea 'information' back to its elementary components, we find as an essential feature an input/threshold/output device which in its earliest artefact form was a trigger, switch or similar device where a small controlled event releases rather than directly causes a more energetic one. Here is the essential idea of regeneration, repair. The transduction of information is completely dependent on the purely biological activity called repair. A transducer is defined as a device that converts the physical magnitude of one form of energy to another form, e.g. a microphone or electric motor. The word is now being used in a very different and more important sense to indicate any device that has an energetic input in one form that creates or triggers a homologous output in some other form. In this sense a meter or biological receptor can be seen as a transducer. The transducer is vital to information systems and we need to explore the range of their functions and possible functions. The simplest form of transducer is one including a step function. It is call a relay. The relay is a device in which a signal operates a switch or trigger receives and responds to a signal.

Figure 14 illustrates the essence of what a signal does in the biological case.

Figure 14: Homeostat
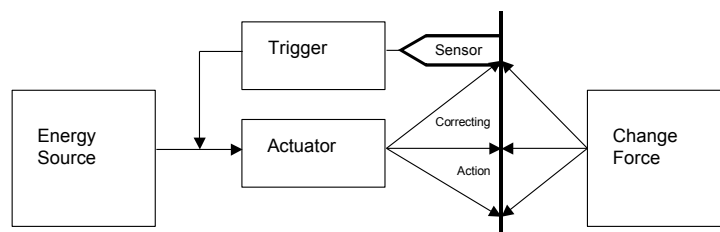
Morphostats and even machines need many governors or homeostats. Homeostats are devices for keeping a variable stable within some sufficiency slot.

The ordinary thermostat is the simplest example. The essence is the *independent source of energy* and a sensor-trigger to release it or hold it back.

In the biological case a signal is usually a representative or symbolic event that is a micro-homologue of some more energetic significant event. It brings information about an event or possible future event that may require purposive

action. It must have a prepared receptor that 'understands' some common 'language'. When you press a switch, you move metal parts. What happens is ruled by classical mechanical and electrical theory, but the switch is 'arranged', it can turn on a light, which is not the natural result of moving a few grams of brass and plastic a centimetre. The switch is a transducer that contains encapsulated know-how or information in its form. Against the normal rules it plays God. It says '*Fiat lux*'. Let there be light! (Genesis)

Returning to the first development of electrical forms of communication, we have to look back at relays. Relays were first used to regenerate weakening direct currents on Morse telegraph lines. The feeble deteriorated signal in a long line was regenerated by being used to open or close another circuit using an *independent source of energy*. This apparently trivial fact is of central importance. The operation of a relay is the simplest form of signal.

A relay is a trigger. It has three distinct features: it has an energetic input that controls an available reservoir of energy and it has an output.

A relay has trigger-like control of the energy in its reservoir. It has the power to regenerate form because locally it can climb an entropy slope. Like Maxwell's Demon it can fiddle the Second Law of Thermodynamics, pumping entropy outside its system, preserving form and order within it.

Another feature of a relay that is of vital importance is that it is a variety reduction tool. It makes decisions. The input of science is pointer readings, as Sir Arthur Eddington said. However, science cannot do much with a forest of flickering meter needles. There are an infinite number of readings or output states of a meter. The simplifier, a relay or bistable, usually has many analogue inputs but by use of a threshold this variety is reduced to two output states (on and off). It is the relay function of transduction from analogue input to binary digital output that constitutes a signal. Variety reduction is the morphostat problem. This helps to solve it. A relay reduces the variety of output states from infinity to two. Above that the number of states can be regulated to any convenient extent

A Gate (Electronic)

A relay is an elementary form of transducer. A gate or logical input/output device in the usual sense allows two outputs and one input. However, gates are used in complexes that may have multiple inputs and outputs, so I suggest that a useful meaning for gate in this context is a device for making decisions that can have any number of inputs and any number of outputs.

Figure 15: A two prong fork gate
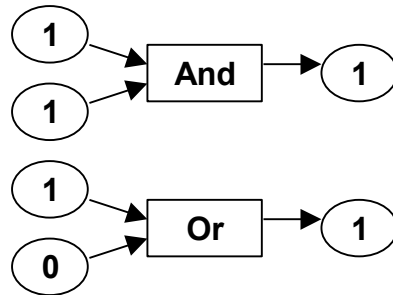
The two-prong fork gate shown in Figure 15 can be an '*and*' gate, an '*or*' gate, a '*nand*' gate or a '*nor*' gate, each with the indicated logical functions (e.g. an '*and*' gate regenerates a signal if both inputs are excited, an '*or*' gate if either is. The prefix 'n' represents 'not' and reverses the action, giving a contradictory output.) The whole repertoire of deterministic computer logic has been built up from these simple components.

Sufficiency gates and semantic gates

In my first version of my book *Brain* I proposed *sufficiency* gates as the component needed in the fuzzy logic, noise tolerant hardware needed for artificial taxonomy. Sufficiency gates are not dissimilar to what are now called 'voting' or 'threshold' gates.

A sufficiency gate (enough gate) has multiple inputs and outputs, the sharp distinction between an 'and' and an 'or' gate is lost. A gate can be anywhere on spectrum between the two. It is not an all-or-none hard logic affair. We could have a ten-input 'andish' gate that responds to (changes output for) any nine of its inputs, or any eight, or any seven. Eventually we have demoted it from a very discriminating, fussy 'and-y' gate with a low tolerance for noise to an easy-going gate, a 'generalising', 'or-ish' gate which will respond to any two or even one of its inputs. We have been changing the threshold downwards. It has lost discrimination and become a relating or generalising gate, an 'or' gate which says 'any of these inputs is equivalent'.

Discrimination and relation or generalisation are the primitive tools of thought. Our new gate is moving away from Shannon's non-semantic paradigm towards a new semantic one. With a semantic gate we can represent or simulate meaning. But we have one more step to take. It may seem like a backward step but it is not. We have to retreat from binary digital expression to analogue expression, to introduce a degree of noise tolerance.

Sufficiency gates can be Boolean/binary as shown in Figure 16, or semantic.

Figure 16: A binary sufficiency gate
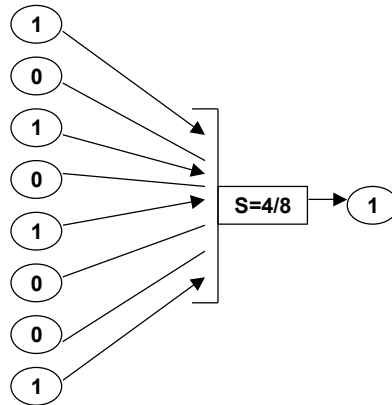
If the sufficency (S=4/8) threshold is set at the total number of inputs it is a full 'and' gate. As we adjust the threshold downwards it becomes less 'and-y' and more 'or-like', finally when the sufficiency is 1/8 it is an 'or' gate.

However, we may also make a gate where the inputs are intensities or magnitudes that may be positive or negative. The sufficiency is now set as a frequency, or magnitude, and the gate responds, sends an output, only if the sum of its active inputs (+ and -) exceeds the threshold. This is illustrated in Figure 17.

Figure 17: A semantic gate
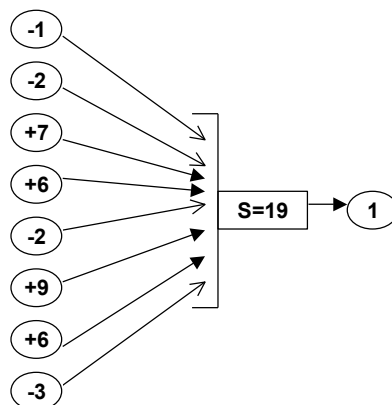
The output of the gate to all its connections can also be a magnitude, positive or negative. We have moved off the Boolean road and we now have a gate that can deal with evidence and meaning. We have a semantic gate. The gate can say more than 'yes' or 'no', it can say 'maybe' (+ 3) or 'I do not know' (0),'I am fairly sure' (+ 7) or 'possibly, but I do not think so' (- 8). Ideally it would have analogue magnitudes, steplessly variable. But in modern computers it can have variety enough in a few bytes.

Why do I say that gates with many inputs and outputs can bring semantic content into information theory? The introduction of fuzzy connections involving magnitudes at junctions between components gives us a calculus for connectivity and that can be used to quantise semantic connectivity that, in the nerve, is the only visible method for dealing with semantic content. The meaning of a word can only lie in the strength, the sign and the direction of its connectivity that nets with the rest of the system or brain. We know that the nerve cell has variable output via its axon, variable receptivity at its dendrites and a variable threshold. Here is a paradigm that shows how this sort of connectivity can represent meaning.

The first program that worked with the concept of this kind of gate in our experiments was used in a program that solved the automatic taxonomy problem of a simulated simple intelligent entity. Such a program is called a taxonomon.

THE FIRST TAXONOMON PROGRAM

In this chapter I have sketched the theoretical approach to present-day information theory and reinterpreted it in the light of the way of looking at that world which I propose. The polyhierarchic information movement architecture that I suggest leads to the suggestion for a possible component and connectivity system for brain simulations, thoughts about possible brain-smith plans will follow in the next chapter. I complete this chapter by republishing (Figure 18) the first taxonomon diagram from my 1974 book *Brain*. This sketch led to James Cherrill's first successful taxonomon programs.

Figure 18: Model of nerve cell programming.
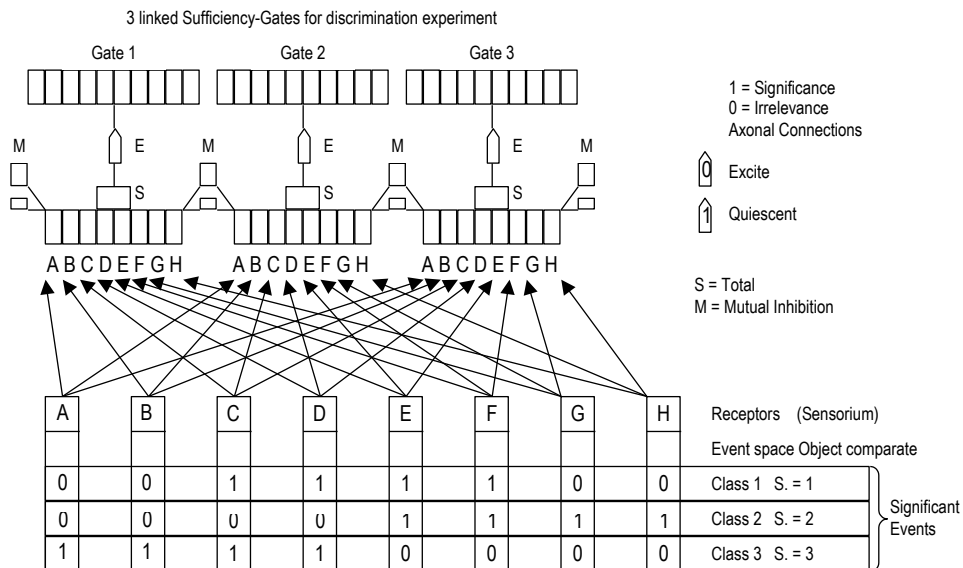
This is the vital first model of a taxonomon.

We see three inter-linked competitive sufficiency gates that are looking at the same 8-input sensorium. Below are the three 'events' or inputs that may be received at that sensorium; they are three binary numbers. The competing gates are randomly set by entering low magnitudes in the input spaces A, B, C, etc. They compete for and 'learn' to specialise, each one responding to one of the input patterns despite the fact that these are noisy or partly erroneous at each presentation. A gate responds or 'excites' when its sufficiency 's' is exceeded by the total of the weights at its excited inputs (those inputting a 1, not a 0). It learns by adjusting its weights and its sufficiency in accordance with the 'learn algorithm' that I arrived at by experiment. In the next chapter we shall explore a primitive range of taxonomons, their limitations and possibilities. I have shown the competing set of nodes that are seeking to distinguish pattern. Now I show a single node and how it competes.

8. MODELLING BRAINS

Programs designed from the taxonomon plan are intended to be experimental and must be seen as several separate programs that interact.

Figure 19: Exemplary taxonomon node

Totaller

It has to calculate the interaction between the output weight and sign from the prenode and its corresponding input weight and sign for all inputs. These it then sums. The result expresses the amount of match of the input pattern with the node setting; the better the match the higher and more positive the score. The result of the interaction is totalled plus and minus to get the node score so as to compete.

Inhibition Input

The inhibition (negative weight) input from other nodes in the stratum group is taken into account in the node total. Inhibition is needed to emphasise the competition for pattern. It has to be a function of output weight.

Node Threshold

The node threshold is another 'learn' variable. For discrimination-learning it is set at first to be a little above the mean unsophisticated (randomly set) node total so that 'good match' nodes only will fire for each pattern class. During learning, thresholds rise side by side with input weights so as to improve discrimination.

Discrimination learning is making a gate more 'and-y', a better responder to a noisy pattern of many inputs.

Node Input

Node input weights are receptivity magnitudes. They are 'learn' variables. They must be many valued, and can be integers or floating point magnitudes. They are set up low and random at first so as to have some initial differences in pattern for selection and improvement.

Trigger

This element simply stores the state of the node, whether it has passed the threshold and fired or not on that pass. If, and only if, it fires it passes its weight and sign to the senior connected nodes whose addresses it has in its link addresses store, and also sends inhibition to the set of nodes to which it is linked for mutual inhibition.

Node Signum And Node Output Weight

These are 'learn' variables, the weight and sign of the output which is sent to all the corresponding inputs of the connected nodes as manifest in the postnode link addresses.

Postnode Link Addresses

This is a list of all the senior (more central stratum) nodes to which the node is connected. These links can change as a part of the learn process.

Objectives

The program must leave plenty of scope for trying different 'learn' schemes and algorithms easily and checking how well they work. What is aimed at, at the sensorium, is to arrange that a top level, apical node fires for each and only one of the input patterns. The higher the tolerated noise and the lower the rates of inclusion and exclusion errors, the better. The aim at the motorium is that the motorium output pattern (vector of magnitudes) shall be as close as possible to that set up as 'appropriate' to the input. The aim is to reward sections of random pattern with a good match by bunching or connecting them by stratum 1 nodes. Stratum 2 nodes bunch these 'good hands (of cards as it were)' or 'genes'. Stratum 3 nodes bunch these.

The motorium behaves randomly at first and gets low-level sections right by trial and error. Then it moves learning towards higher, more senior levels. Excitation that is successful moves towards the apices at the central stratum and simultaneously excited sensorium patterns, while punished excitation is inhibited or directed laterally to skirt towards alternative apical nodes. It is seeking other behaviour patterns looking for the right response to that input.

On both sides of the network the rewarded patterns should probe towards each other and eventually connect so that input of pattern 'A triggers output A' its appropriate response. That is what learning is, finding the right behaviours for whatever is received.

The optimum algorithm system at the motorium is that in which the 'right' output is fired by the 'right' input, in spite of the most noise, with least error, with least iterations, in least time.

Note: Both sides of the network senior nodes are more central, not 'higher'. Junior nodes are more peripheral.

THE LEARN ALGORITHMS

For the discrimination function, the competitive sensorium-observing nodes are set with low random receptivity weights at set-up (before learning commences). Thresholds are also set random but low when the run starts. Threshold setting must be such that some nodes reach the threshold and fire from the first. Once some have fired the aim is that one of these promising nodes will 'capture' that pattern and 'shoo' the other nodes away from it. The need is to discourage the others and make them seek other connections or other patterns on the length of sensorium upon which they are in competitive linkage. This is done by the high score node sending inhibition on the mutual inhibition link and by a 'punish' algorithm applied to the nodes that reach threshold but are not in the highest score group. However, the algorithm cannot 'know' which node is closest to the input pattern norm because each presentation is noisy. The noisier the patterns, the longer the learning has to take so that the noise window can be explored. The more noise, the more time to learn properly. Too fast a learning rate on noisy patterns leads to more wrong capture and failed capture.

Each competitive node learns to be better at discrimination, at capturing just one input class, by changing the weight at each input in the direction that raises the node total when that pattern is sensed. In order to make it specialise the most successful node that fires also has its threshold raised so that it becomes less sensitive to other patterns. It has to raise the threshold enough to stop it firing for other patterns but not so much as to lose fuzzy members of the pattern class it is beginning to capture.

The aim is to open the fuzz window of the node just wide enough to capture the fuzziest example of its chosen pattern but not wide enough to capture others. This can only be done by a number of iterations that depends upon noisiness because the node must be designed so that it learns rather than being instructed how fuzzy the pattern is initially.

To settle competitions where multiple nodes seek the same pattern, total score can be made to relate positively to the weight of the inhibitive signal sent to the other nodes in the competition group. The highest scorers would gradually assert their claim.

A node that passes its threshold but is outvoted or inhibited is therefore 'shooed' away from the current pattern. It has its weights changed so as to lower score for that pattern. Its threshold is lowered to give it a chance at others.

'Over-discrimination' a pattern by constantly inching its fuzz window narrower, it must eventually get too fussy, have too high a threshold (sufficiency), so it loses it again. A warning threshold above the real one is required to slow and eventually stop the threshold-raising process when the 'warn' threshold is reached too often.

To encourage the search for pattern there must be an algorithm that breaks unactivated links and makes new ones. Input weights are set low and changed up or down towards some plus or minus figure. Learning shifts the weight away from zero so those inputs that remain near zero for a long time are not learning and must take a chance of being re-connected elsewhere, usually randomly.

Experiment is needed in designing the relation algorithm. This is that which unites classes of input that are, for the entity, equivalent, require the same response. This is an 'or-like' function which makes a node collect inputs which 'mean the same thing'. The position that must be inductively reached is where the weight on each equivalent input passes threshold and the node fires for any of them. To stop it firing for any plurality of nodes (to make its 'or' exclusive), where required, there has to be a second threshold so the thing works like a 'go/no-go' gauge and fires only for inputs that score between two thresholds. The problem is that this sort of learn algorithm is the opposite of the discrimination algorithm which gets more 'and-like', more fussy, as it learns to be a better discriminator.

WHAT CAN WE DO WITH WHAT WE HAVE?

What can taxonomons, even of the present primitive design, do now? The present design is confined to the sensorium half of the posited brain so they can only seek pattern in data. They can only respond by recognising patterns and classifying them. They can look at any mass of undigested, unfiltered, fuzzy data and sort out any unsuspected repetitive patterns hidden in it.

An example is astronomy. If we got a computer record of the data collected from a large array of stars and presented this as vectors of magnitudes, say Fourier analyses (vectors of intensities by frequency), the taxonomon without previous instruction should develop an apical node that responded only to Red Giants, another for White Dwarfs, another for planets etc. But it might do more, it might find unsuspected clumpings or clustering of qualities and discover hitherto unknown types of star.

This example refers to vectors or strings of magnitudes but the virtue of the taxonomon is that at the higher strata it turns topology into semantics. It can deal as well with information from an array of meters or sensors as with information from a string of them. In fact it can just as easily deal with a three- or four- or multi-dimensional sensorium. As long as the sensorium perceives the input in the same way each time, the number of dimensions is irrelevant. The input is turned into and treated as a string anyway (in present hardware) so the taxonomon is indifferent, it would still pick out patterns that were repetitive because its strings would have their own sort of repetitiveness. Even the present primitive taxonomons can see through a thick haze of fuzz.

A vital point is that there is fuzz-tolerance at each level of the polyhierarchy and the sufficiency, a fraction, is multiplied at each stratum so that the top level noise tolerance can be extraordinarily high. (The sufficiency at a central stratum is the product of the sufficiencies at the intervening strata. If the sufficiency is 0.8 at six strata then the apical sufficiency is 0.26. The taxonomon will discriminate through 74% noise.) (This might explain the miracle that we can recognise caricatures.)

Some applications are already visible. Taxonomy of massed patients' data reports would give known medical diagnosis patterns and might find new syndromes. This would save Doctor's time.

Businesses could have the great cluttering bulk of customer data and other reports filtered for all sorts of useful but unsuspected trends and patterns which were hidden like needles in the fuzzy data haystack.

Geological, meteorological, sociological, commercial, and many other types of mass data could be filtered in numerous new and potentially useful ways. Many unsuspected regularities, classifications, and patterns must surely be found. Economic and even Stock Exchange data can be subjected to automatic number-crunching taxonomic examination to find trends and patterns that may not have been spotted before.

Large farms, factories and industrial plants could have a taxonomon program with a large number of local point sensors measuring all sorts of variables, seeking syndrome-like 'states of the system' which might help management and

production engineering decision-making. (Even before there was motorium that could do its own optimising.)

Law enforcement and national security authorities have to seek patterns in the reports of criminal, espionage and terrorist activity. Taxonomons working on data about offences, suspects, etc., might do what a good detective does seek the revelatory patterns in behaviour. Today they have to rely on bored people inspecting gigabyte data-mountains to spot such trends. A taxonomon would never get bored. It would dredge patiently through the poorest ore to find the predictive patterns lying beneath the sludge and fuzz.

Traffic control is largely a matter of pattern spotting. There are or could be a large number of inputs at a widespread net sensorium and a taxonomon program could spend a time seeking patterns before trying out simulated improvements. Later it could have a motorium activated that would learn by experience, like an animal does. Its 'pain' would be hold-ups, slow flow and accidents, its 'pleasure' would be fast, safe vehicular flow.

It is possible that the very complex problems of aerodynamics and hydro-dynamics can be tackled in this way. Aircraft and ships, as mentioned, could have intelligent surfaces, like those of ultra-fast fish and many birds. These are adjusted minutely and locally to minimise drag. A submarine with an intelligent skin might be less detectable.

Any unknown system where there is a large array of data input can be explored with these programs even as they are now.

Most important, the design and improvement of taxonomons themselves can be undertaken by taxonomons. This is just their kind of problem, and maybe the beginning of what Professor John Good described as the intelligence explosion.

Elaborating The Model

I have thrown up some shaky scaffolding. Yet it might be strong enough to support the harder work of building a more elaborate model.

In the next section I try to tackle some of the gaps in the present model. I have ignored time and delayed reward and behaviour that is spread over time and which needs chains of inputs. I have not dealt with changing priorities, attention, and memory in the ordinary (not the computer), sense. I have not accounted for one-shot learning. There follows some more adventurous modelling to show how future thinkers may move in filling these hiatuses.

DECIDING PRIORITIES

The primitive model of brain drawn so far is a very low level one. It concerns a brain that should be able to deal with the recognition of incoming patterns and the trial-and-error elaboration of instantaneous outgoing instructions.

There is no provision for priorities, for differential urgencies, for memory, or even for time at all. Attention, consciousness, volition, thinking, dreaming, and planning have not been accommodated in the model. It would have been a daunting and perhaps impossible task to make a brainable model with an instantly visible fit with all these observed aspects of our minds. I now want to show how these things can be fitted to the model without destroying its general pattern. I anticipate that these may be called ad hoc additions but the reader must judge whether they clash with or fit with the broad architecture proposed.

Figure 20: A Morphostat
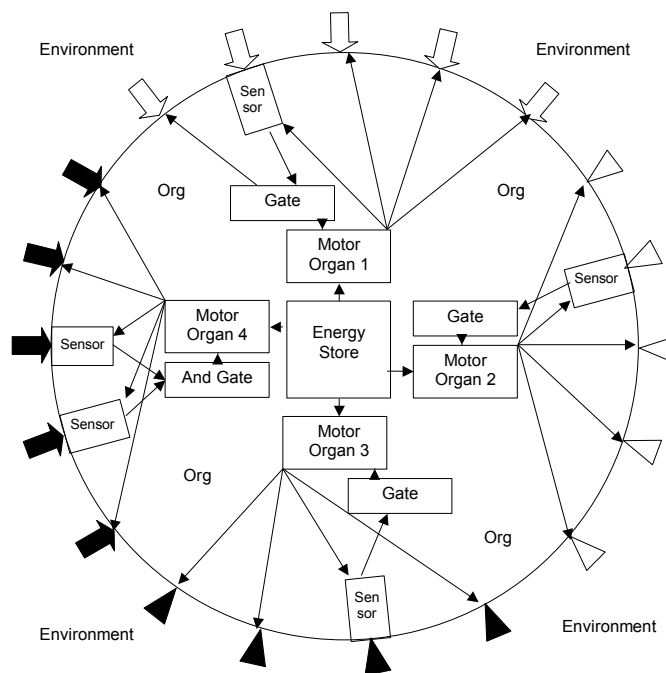
A morphostat threatened by four different disturbance forces has four Motor Organs triggered by gates receiving a signal from receptors (sensors). Motor Organ 4 has two sensors with 'AND' gates.

I am an inventor and you will see that I am looking at the problem like a brainsmith inventor, asking how I would design a brain in view of its functions.

Let us return to the problem of the simple morphostat. If it is a normal plant, a member of the metaphyta, which has neither sensorium nor motorium above the cellular level and only slow chemical signalling, then nearly all its reactions are tactical and local, there is practically no central co-ordination.

It can be represented as in Figure 20 by a series of independent homeostats each dealing with one variable and keeping it within viable limits.

But the metazoa are heterotrophs, they have a different problem from the autotrophs, which can convert ubiquitous and freely available nutrients and gases into self; using energy from the sun. Metazoa have to seek for, find, and maybe struggle for ready-made conveniently packaged nutrients that the autotrophs etc. kindly produce. And that means activity, motion, seeking, judging, ontoplastic learning. The question of priorities arises, the question of central control. Many, if not most, of the problems of preserving form through time have to be solved tactically from second to second, even in metazoa, but there is another class of problem for them. In these problems it pays for the whole animal to have a central strategy of 'one thing at a time'. If it tries to eat, flee and copulate simultaneously it might not work out too well.

Appetites

Once there is any aspect of central co-ordination the question of priorities arises. What if the intelligence perceives two patterns simultaneously, both of which trigger appropriate but contradictory instructions to the same set of muscles or organs ?

We know that the advanced intelligences do have means of deciding to do one thing at a time, a strategic-level way of suppressing or deferring action related to all the inputs but one, the most urgent. The present polyhierarchic model has no such means of decision, no central processing unit to select from a set of action-demanding options that are presented at the sensorium simultaneously. How can this known behaviour be made to fit the root-tree double polyhierarchic network picture? As an inventor, how would I go about that?

Experimental psychologists pursuing determinist models for brain action had great trouble with a feature which was variously called 'attention', 'set', 'tendency' as behaviourists, some of them were trying to treat the animals like black boxes, watch input and output and predict the latter from the former. But the animals were awkward. They refused to be predictable systems input-output systems, like computers. They seemed to have ideas of their own about what they observed at a given moment. It was hard for a behaviourist to deal with the unpredictability of behaviour caused by this aspect. There was a joker in the pack, some hidden force inside the black box seemed to be tossing pennies.

However, as a brain-smith I cannot see how an advanced animal could manage without some simple, even if arbitrary, system of deciding priorities. Penny-toss decisions are often a lot better than no decision or multiple incompatible decisions.

And maybe the decision-making system does not have to be completely arbitrary. It can be guided by another system which acts like the accounts department does in a company - a system which keeps a budget of the whole range of general strategic policy needs of the creature, measures need and satisfaction on each and biases attention. Action choices are then biased accordingly.

Ethologists after the Watsonian behaviourists began to observe that there is an appetitive system in ontoplastic (advanced) animals. There is a set of instinctive appetites, for food, water, sex, to care for young, curiosity, to sleep, to explore and have social contact, and each of these fluctuates cyclically, rising when unsatisfied and falling when satisfied to make a sort of continuous symphony of variation.

This observation fits the polyhierarchic model if we suppose that in advanced animals there is a permeable barrier between the apical regions of the sensorium and the motorium.

Imagine a simple creature which is simultaneously perceiving that whole set of inputs which would cause it to flee a predator, that set that would cause it to seize and copulate with a mate and also that set that would cause it to pursue a prey. Those three sensorium apical nodes are firing, but because of the posited barrier they are sub-threshold for the opposite motorium apical nodes. They are not receiving quite enough excitation to pass the threshold of the motor apical node to which they are allocated or semi-connected. Nothing can happen until an additional input of excitation is received, one that takes the excitation over the barrier threshold.

Where does the additional excitation come from? It comes from the other, the internal sensorium, the one that is linked to and sensing the state of the appetite system. The balance of appetite and satisfaction as well as hard-wired priority system decides which of the presently excited apical motor node gets attention.

Thus the present state of the appetite system together with that system's hardwired priorities decide which of the afferent apical nodes from the sensorium gets the extra input of excitation to trigger the connected motor hierarchy. Thus the animal flees from its predator unless it is too hungry or lovesick. Hardwired, there will be a tendency to put fleeing before feeding for instance.

It can be seen how such a primitive strategic system would develop into features called attention and consciousness.

Attention

The posited appetite system can have two further functions. It could very conveniently act as the success/failure sensing and measuring system and it

could be the system which sends the colder/warmer, pain/pleasure inputs that are needed to work the trial-and-error learning system that our model requires. Another function could be to provide the mechanism attention, the other phenomenon we introspectively observe. We do not pay much attention when we are asleep or unconscious. Consciousness seems to mean paying attention.

Parenthetically, the limbic system in the mammalian brain appears to be a good candidate for the physiological organ that corresponds to the behavioural observations of the instinctive or appetite system. It is observed that this system does send streams of excitation in many directions to many parts of the brain.

Our picture therefore has to accommodate a strategic central control that can set priorities. This can be represented as the central apical region between the two halves of the system where the sensorium apices meet and connect with the motorium ones.

We can see the appetite system as what amounts to a secondary sensorium, one which measures needs and satisfactions and which contributes to the final triggering linkage from the informational side of the brain to the instructional side, the Works Department.

Consciousness? The most difficult question.

This brings me to the complex question of consciousness.

Many thinkers have proposed models of the mind as a great connectivity network. My model is of millions of interconnected loci, balls and boxes with links or junctions of varying strengths. (The connectivity (+ or -) of locus A with respect to locus B is a magnitude which sets a contribution to the probability that the firing of A will cause or prevent the firing of B.)

I suggest a connectivity network where each idea, concept or percept is a idea and the semantic inter-relationships of these elements are the connective states. I am conscious of all that is semantically (connectionally) close to one locus at a time. That is closely connected to my wandering ego and active.

I have the power to move the attention point around the network. I can go to any other point in the great connectivity network of the mind and explore any of the regions of high semantic connectivity around that point. But I have to proceed by a semantic or connectional pathway, though these are often obscure. We can think of connectivity as being the same as semantic relationship. The word 'mother' has a high connectivity with the word 'father' but a lower one with the word 'incandescence').

At the same time I know that the movement of my attention is very much affected by my sensory input at any time. Some attention is involuntary. It can be

attracted despite my will. I cannot withhold attention from sudden emergency. Further, my attention is partially controlled by the state of my appetite system. If I am starving I keep thinking of and notice food.

I venture on to this tricky ground only because the behaviourist assumption is no longer tenable. The speculative model that follows is now testable. Computer simulations are now possible. In all that follows I am outlining very tentative, first guess hypotheses that perhaps within the next hundred years or less may be testable in simulation on the family of concurrent hardware computers that are now possible.

Unlike the behaviourists I feel that our internal introspective experience should be given first status. If any model of reality is contradictory to that then it cannot really convince.

Here, when, I apply a trial first-fit match of the garment of my model both to the dummy, the usual concept of the physical object the brain, and to my mental 'body', my introspective experience of its workings.

Let us suppose that this cortex is topologically or connectionally the central region. Let it be, in fact, the apical region where the two halves of my model brain, sensory and motor, meet. Let its junior connections reach down both to the sensory and to the motor regions in two separate descending polyhierarchies. This would then be the semi-barrier region I posit where sensor-motor links are complete but subliminal, below threshold.

In man, the mammal that specialises in thought, planning, and contemplation before action, there has been an enormous expansion of this neocortical region. I guess that it is the region of planning, option-weighing, cogitation, where whim, indeterminism, and experimentalism are preserved even after sophistication is complete. Here ideas, plans, and imagination happen and decisions are made. This is the zone, I suggest, where all the hierarchic apices meet, the highest, most senior strategic zone.

I see this great region as subliminally interconnected during waking hours, because then something central in us, our attention point or focus, our ego, seems to be able to visit any part of it just by deciding to do so.

Consider interconnectivities measured in the way I have posited (constantly changing excitation, receptiveness to excitation, inhibition and threshold) all over this strategic level system. At any one moment in wakefulness there is some one connectively 'local' zone that has the highest mutual connectivity with all the rest of the system. As connected and active assemblies and thresholds change this centre must move round the network. Let us call that zone of highest general connectivity the ego, the wandering centre of attention.

Appetites Direct Attention

Clearly the excitatory input from the appetite system is especially important in deciding the position and movement of the ego. (Hunger makes us pay attention to food, lust to sex.) Where ego, the centre of maximum excitation, visits, it brings the additional excitation required to turn potential into action, to raise a subliminal connection above threshold and trigger some motor hierarchies or trains of them.

So, simplistically, there are at any one moment many perceptual apices lit up, subliminally active, because the percept combinations that excite them are present to the senses. I see food, drink, a book, my lover. Eating, drinking, reading, lovemaking are my immediate options. It would be difficult, if fascinatingly innovative, to combine them in some simultaneous acrobatic feat. But the relative strengths of my appetite drives, memories of previous experience, are all connected, with different weights, semantic strengths. Between them they decide the location of ego, my attention focus. I decide and, boringly, trigger just one of the motor trains.

Is that, at least, a brainable picture of the processes within us of which we are conscious? We speculate we wander around relevant semantically connected regions and as a result we make decisions. A locus in this network, the point where consciousness is for an instant, will certainly be in mutual excitation with a limited set of cells but it will not have a location in the normal sense because the 'near' and 'far' will mean semantically or connectively near and far. Nearness will be semantic nearness, it will be to do with excitation levels and the permeability of connections and not with spatial millimetres. The ego will be induced to move along high connectivity paths and it will be constrained by low connectivity insulation bounds. But by various routes it will be able to reach any locus. It will be able to explore the highest most abstract strategic regions, or go down the hierarchies to the peripheral tactical levels and concern itself with the feeling in a patch of skin or some small detail of the visual field.

An analogous model of this wandering attention, which moves around the connected apical regions making connections here and there, is an industrial company. A manager moving around a factory (or having people visit him). He is considering all the information, deciding priorities and triggering activity. He is not much concerned with the tactical details, sensory or motor, he triggers strategic-level prepared plans as he moves around the network. He provides the release that causes an informational input to trigger the 'appropriate response' output. Build a new factory or get a new model designed.

MEMORY

The meaning of the word 'memory' has been greatly extended, first by communications science and now by genetics. The one thing that all life forms use in the most extended sense of the word is memory. It is the one essential of

all life forms, a means for carrying the past into the future. It takes these various forms at various levels.

1. The genome
   This is, of course, a memory; it is the encoded form of the phenotype, the creature. This is the primitive and universal form that is present in every living cell. It is the databank that holds the prescription for the living thing.

2. Neural memory
   This has three main divisions:

   1.1. Genoplastic memory
        Hardwired, perceptuo-motor linkage between sensorium input and motorium output. The making of such connections is evolutionary. For example, most insect behaviour is robot-like, automatic, as first demonstrated by Jean Henri Fabre (1823).

   1.2. Ontoplastic memory
        Softwired perceptuo-motor linkage changed by learning within the single generation, e.g., mammal skill, environmental and avoidance learning.

   1.3. Cognitive memory
        The original meaning, the memory track of which we are conscious and which can be communicated between people coded as speech or other symbol systems such as gesture. We can only guess, but cognitive memory probably exists in mammals and maybe other vertebrates such as birds.

2. Social memory
   This is the preserved traditions, mores, arts, skills, with the remembered and written records of an ethnic/language/culture group or more lately of any recruited interacting group, e.g., a commercial company or educational or other institution.

Cognitive Memory

Here is a view of cognitive memory consistent with the present model. We have posited that ego, the attention point, is during waking hours wandering around the connectivity, the relationship network. At each place that it rests there are numerous perceptual apices in close connection, behind them those of the percepts and behind them the elements that make them up. Presence at one such locus is the multiple experience of a whole situation with all its close semantic connections, including perception of self.

The suggestion is that the track is tagged, marked, like that of the runner in a paper chase. Something happens at every point along the ego's track around the network. Like Theseus in the labyrinth, the ego leaves a trail. A silk thread marks the passage of its meanderings round the connectivity network that it has built up from its experience during life. Each point visited has a flag, a marker, a physical change, so that, when the ego returns any time later it finds the flag and the topological locus with all its thorn bush of radiating semantic connections. Here is a part of the memory track. '*Deja vu*' - 'I have been here before,' it says. Remembering and reverie are the processes of exploring that serial track or making leaps along it to search for experiences that will help decision-making. But there is a time when no such track is laid down, we live on but leave no memory trace. We sleep.

Sleep

Does sleep fit the model? There are many hypotheses and little consensus about sleep. Sleep is very, very odd indeed. Competitive animals struggling for survival in a hostile world spend a third of their time blind, deaf and paralysed! It wants a lot of explaining.

Sleep is a condition in which sensory input and motor output are both halted. The survival of many species that go into a state of stupor, of deaf-blind insensitivity to danger or opportunity, is another minor miracle that remains without explanation. Why has not sleep been eradicated by evolution?

All animals live in food chains, fighting to eat their living food and to avoid becoming the food of others. Sleep hampers both predator and prey. Many smaller handicaps have been evolved away. Why not sleep? It is evident that it must confer some enormous but unknown advantage.

Fitting it to the present paradigm we see a state of the network when the wandering, central and singular ego breaks up, dissipates. The centre of neural connectivity fragments, there is no single central attention point or, if there is, it lays down no memory track except the occasional, ephemeral one of a dream.

Experimental psychology has established with some sureness that one function of sleep is to do with the learning process. Not confused by the continuous sensory input and motor output, the learning process seems to be able to advance more surely during sleep. (We say, "Let me sleep on it.")

I see sleep as a mode in which the mind breaks up into autonomous, mutually insulated, regional topological zones in which there is not one locus of maximum connectivity and excitation, but many semi-insulated smaller regions of high mutual interconnectivity. Such analysis would be advantageous if it allowed the learning process to work in parallel instead of serially. Perhaps sleep is the time during which the synapses change their permeability, the thresholds change and

the nodes change their excitability, while free of the stirring, muddling passages of the active, over-responsive ego. In sleep the memory track can be searched for the results of past actions so that delayed. Successes can be found to offer the reward and punishment signals that are needed by my model.

Memory As Learning Amplifier

Work with the taxonomon has shown that a large number of iterations is needed for learning. If an intelligence learns from too few iterations, examples, it cannot get the relation between signal and noise right. Errors of false capture and failed to capture become inevitable. Memory, by storing serial strings of experience, enables the intelligence to get surrogate experience to learn from, by retracking old experience from the data store of memory.

A much more important function of sleep is therefore suggested by these taxonomon programs.

The model insists that learning must be iterative, inductive. But we observe some learning in mammals at least that seems to be from a single instance, especially when the reward or punishment is instant and strong. How would the engineer adapt inductive programs to learn like that? What would be needed, as suggested above, would be time multiplexing, a system in which the shock experience could be stored and then re-run repeatedly, including the inputs of reward and punishment. Further, since the sleep model posits an analytical mode in which perception is broken down and behaviour built up from elementary blocks, it would be greatly advantageous and speed things up if there were many short-cycle re-runs in parallel (concurrent) rather than one long run in series. Now the memory track of which we are all conscious provides us with the data store of tagged strategic-level connectivity states along which the ego may pass in repeated 'ironing in' cycles after the objective experience is over. We are all conscious of the repeated passes our ego takes down a long memory track in the rehearsing again and again of significant incidents we are conscious of. That is conscious learning.

The 'tempting hypothesis' is that more tactical learning would be much better, faster and more detailed if the ego could break up for some of the time into numerous 'egolets', small connectivity centres that made many more concurrent re-run cycles of short, important sections of the track.

Conscious rehearsing, ironing in the lesson, must mean preoccupation, a lower attention to current input. The ego break-up system with its enormous advantage in learning time, detailed analysis, and concurrent, many-track learning has a higher price. Suspension of central strategic control, unconsciousness, sleep is the essential tactical stage of learning.

A parable: In war the General and his staff are in full strategic control, everything moves to his command. The army is awake. The war ends. The troops go back to barracks, into dispersed tactical-level training exercises under many sergeants and junior officers. What do the Generals do in peacetime? They dream. They set up a thousand scenarios and again and again go over the accumulated data from the last and all the previous wars. But the intelligence input and the orders output of wartime are cut off. Sociozoan consciousness, and two kinds of sleep, just what we find in metazoa: rapid eye movement sleep, when we dream and can make a memory track the Generals are doing exercises, mock wars dreams; deep sleep, the troops are in barracks learning to march, to work their equipment, tactical learning at myriads of centres.

As I have said, hypotheses like the above were rightly rejected by psychologists when there was no way of testing them. I suggest that the above ideas are now empirically testable in simulation on the sort of nerve-net models with which we are beginning to be able to play games. I venture to predict that a stage like sleep will be found to be essential when we get to high-level artificial intelligence. A stage will be needed when the input and output are switched off so that important runs of instructive inputs can be iterated many times so as to get the message imprinted and the fuzz explored.

REFLEXIVE BYPASSES

We have to consider reflex responses. They obviously bypass the polyhierarchic system described. Reflex responses (those like the knee-jerk response, or the involuntary withdrawal of a traumatised limb) are involuntary. They evade the conscious control of mind. It has been shown that such reflex arcs, which are mediated through ganglia in the spinal chord, are a bypass to the central nervous system itself. There are many other such direct links that bypass all the strata of the posited polyhierarchic network. Here, signalling from sensorium to motorium is at the very lowest tactical level and is genoplastic not ontoplastic, it cannot be learned or unlearned.

Another sort of reflex type bypass to strategic-level brain, neocortex. All the proprioceptive inputs to the ganglia associated with muscular and other motoneurones. The ganglia close to the muscles where proprioceptive inputs mediate motoneurone are bypasses to the whole system as envisaged. They are not a contradiction of the pattern, they are another part of it. (Proprioceptors are sensors that measure muscle actions.)

This is a first, simple, broad-brush sketch of the two polyhierarchic networks. We have to modify our visual picture of high-level strategic information. To this picture we now add the various hardwired (ontoplastic), tactical-level bypass routes which constitute reflex responses and cybernetic feedback to muscle fibres. These details falsify my model. They are peripherals.

There are almost certainly some learning processes, probably mediated by the paleocortex (the cerebellum), which bypass the whole neocortex system. K. S. Lashley and other early workers were greatly puzzled by the fact that their completely decorticate cats still showed some residual simple learning. The whole of the cerebral cortex was removed, yet the creatures still seemed capable of low level tactical-type learning. This caused Lashley, W. Kohler and others to completely reject the simple connectionist theories of the time. They denied that any one nerve was essential to any mental action and were looking for holistic theories like Kohler's gestalt theory.

The present model is connectionist but does not need the presence of any one component because, at each stratum and at each node there is a margin for error, an allowance for fuzz which is set by the sufficiency, which in turn is set by the threshold of the posited component. Lashley and Kohler were right, no cell is essential, but they were wrong, I believe, to think that this destroyed all possible connectionist theories. (As, to his credit, D.O.Hebb argued at the time.)

A polyhierarchic model with sufficiency gates as elements is a connectionist model that can survive the awkward fact that, as many studies show, there is a steady mortality of brain cells throughout life. Just as societies survive the death of role holders, people, its nodes, so brains survive the death of neurones. In neither case does it mean that individual nodes have no unique role, a connective place in the network, acting upon and responding selectively to other specific nodes. The roles of cells, nerves and people may change with time but at any one time they do exist. The network is connective and functions being interconnected.

THE PROBLEM OF DELAYED OUTCOME

The generalist trying to model the learning process as it appears in people has many daunting problems. Not the least of these is that of finding a credible model for a learning process where behaviour is optimised in spite of the fact that there may be delays of months or years between action and observed outcome, successful or otherwise. How can success and failure signals work if success comes after wards?

When a memory track is laid down there are connectivity loci within the network that were laid down at the time of the original input/output trials of any given problem. If the ego (the attention focus point) returns to that locus on receipt of the corresponding success/failure signal, however much later, then the success/failure input can do its work of making local modifications which are relevant, despite the time lapse.

The memory track explains another hard-to-simulate feature of human behaviour, multi-tasking. We seem to be able to deal with many problems at the same time, switching almost randomly from one to another. This is much harder to

understand than the simple pre-programmed multi-tasking in computers. With people the multi-task behaviour train has itself to be learned by experience. The attention focus can jump around loci in the memory track under the influence of some more strategic-level system.

The Eureka Moment

Karl Popper was frankly mystical on the subject of the eureka moment. I refer, of course, to that joyous exciting moment when what was perplexing and contradictory suddenly falls into place. A consistent model appears as from nowhere and brings all aspects of the problem area into felicitous congruence. 'Eureka', 'I have found it,' cried Archimedes in his bath, when it came to him that density could be accurately measured by water displacement.

I suggest that the eureka moment is that when a valid generalisation emerges. It is that instant when two connectivity systems, which have developed separately, merge and unite, when an 'or' type connection is made which unites the two classes as members of some more general class. It is obviously essential that this should occur. The baby learns about an assembly of face, arms and legs that mothers it, then of another that fathers it. One day, 'Eureka!' a breakthrough. There are a lot of these different assemblies of the same elements; people. A whole set of subclasses has evolved a senior node that responds, 'or' style, to any member of the set. As nodes arise which represent more complex assemblies, finding this sort of valid generalisation becomes more difficult. But is not different from the baby example above?

The act of recognition itself is a pleasure. When you see the face of an old friend there is a moment of pleasure that does not occur with the sight of a familiar, even a loved, one. The recognition, seeing that something unexpectedly belongs to the already known, that it has, in some form, a subject comparate ready within the brain, is itself rewarding. This is biologically advantageous. The pleasure is the success signal that is needed to reinforce the connectivity pattern. Taxonomy itself, even without advantage, is and must be pleasurable. Animals, especially predators, must seek to learn more about what may be relevant. It is biologically advantageous. Why does the pleasure diminish with familiarity? Obviously it must. Once the lesson is learned there are other things to be learned. Pleasure would be devalued if there were no system of rising expectations and parallel lowering of reward.

Boredom has a function. It sets the animal on a search for stimulus so that it can learn more of its environment. In spite of the fact that curiosity kills cats, the survival rate of curious cats must be higher. Predators like felines who get bored and explore must outbreed incurious ones who let their environment force itself upon them. Otherwise, all felines would be politely and supinely incurious.

Boredom is confined to the higher animals, those with some serial tagged-track type memory. Every ontoplastic animal, however simple, has some memory. There is immediate-outcome learning, like learning to walk or play tennis, and there is long-term memory learning where attention backtracks, recapitulates past experience and modifies connections that are no longer in short-term memory. We can call the two styles 'psychomotor learning' and 'cognitive learning' and they are very different. Advanced taxonomons will have to tackle this.

THOUGHT

The perceived wandering of our attention around the network of our mind in the model, is the wanderings of a topological zone of high excitation around the strong linkage paths of the inter-threaded multi-dimensional network at the semantically central zone of the brain. During consciousness, the ego, the '*highest whole-net connectivity centre*', is always on the move. In alert consciousness a single perceptual input from eyes, ears, skin is 'lighting up' a topological region we need. Thus the perceptual input during attentive wakefulness is in large part that which decides the location of the ego. The input from the appetite system and the 'long-term goal' system does have influence, it directs attention, switches eyes to interesting scenes, selects one voice from a hubbub of voices. In reverie, speculation, planning and thinking, the perceptual influence is reduced and the appetite system and its great superstructure, the long-term goal extensions from it, become the principal influence on the location and track of that activating centre.

There is a random-walk element in all speech. Recitation is a fixed phrase-walk. All speech has free choice functions where randomness creeps in. We could identify intelligence in a Turing Test by the unpredictability of the speech functions, the discontinuities, sidetracks, whimsicality. (A Turing test is one which can distinguishe a brain from a computer.)

IMAGINATION AND PLANNING

The model also gives us an insight into imagination and planning. A sophisticated creature needs an internal world model in which to make experimental simulations. The connectivity system proposed can provide such a model, a system built up by each separate being during life as the internal model of its own world. It is a model where it can try out painlessly the results of possible actions and responses and make decisions as a result of planning rather than spontaneous reaction.

Connected with each percept, concept and relationship is a 'word' triggering node across the barrier on the motorium side. Serial thought is a connectivity-walk that is the echo of a random-walk. The ego moves along a high-weight track

connecting nodes triggering 'phrase' nodes, which are triggered by word nodes at lower strata.

Words are the atoms of cognitive thought. In the model, the physical representation of phrases and words lies in the state of the connectivity pattern between the physical nodes (cells or cell assemblies). This connectivity is alive, dynamic, constantly changing, being repaired and regenerated. A simple percept does not have a single locus or node to represent it. Each word has many connective loci, a neurone or a group of them. Percepts and concepts of low specification would have to have a large number of loci to represent them wherever they may be needed. Only very complex highly specified combinations of percepts, concepts, etc. would have a few or only one locus or whole action situations to 'represent' them. This is a vital essence of the model.

Speech, on this model, is the movement of the ego along a phrase-walk on the motorium side of the central zone. The ego is aware of, semantically connected to, all the surrounding ideas to the words used. Thought would be movement along a parallel interconnected track on the other side of the barrier.

A man is picking his way through a wood in the dark with a torch, the words are the stones and rocks and earth of his track, but, as he goes, his lamp illuminates the trees and bushes, all that is near on his track. His mind knows and sees more that there is on the track he follows. From each point, he can see and decide the best next point as his course proceeds. This is how writing and especially poetry is for the prepared brain, a multitrack journey with many simultaneous implicit messages above, beyond and around the simple chain of words. Felicitous sonal echoes too are exited.

9. BIOLOGICAL INTELLIGENCE AND ITS ARTIFICIAL INTELLIGENCE

How were the present entities to which intelligence is ascribed created? By evolution, the selection of random, blind options. It took a long time but it was achieved, *without the help of prior intelligence*. Are we to believe, then, that *with* the help of an advanced intelligence it is impossible? This book is the answer. Contrary to what appears to be common sense, living things, the Creatura, can transcend their creators, pick themselves up by their own bootstraps.

Assume that 'intelligence' can appear on other types of substrate system not confined to the world of organic chemistry. Brains based on silicon-electronic or fibre-optical substrates are possible. I cannot see any evidence for, let alone proof of, the negative hypothesis that non-biological intelligent systems are impossible. The absence of an agreed meaning or definition for the word it will always be possible for people to say that what is put up as an example of artificial intelligence is not so. Our species, *homo-sapiens*, arrogates to itself the pre-eminent position on the scale of known intelligence. Will there be other entities, which excel us in this field, and, if so will they be our evolutionary successors or will they be our own artefacts?

Biology shows no previous example of permanent pre-eminence in any type species so we must not expect too much. Worse, as you will see from what follows, we may already be in the presence of something that is beyond us in intelligent performance.

A RIVAL INTELLIGENCE ON EARTH?

In the early days of science fiction before the full implications of the space- time prison in which earthly intelligence is confined were realised (a 13-year signal response cycle for the nearest neighbour system). There were stories of invasion by alien intelligences from outer space. In many stories they came in spaceships and would have subjugated or destroyed us all if it had not been for a small and clever group of human heroes who thwarted them. Our species on Earth has already been invaded and largely taken over by an intelligence other than that mediated by the intelligence within human brains. Much to our mutual benefit, there is a super-brain of a simple, even an elementary type, operating to co-ordinate and modify our behaviour as individuals.

It is that first discovered by Adam Smith and he called it an 'engine'. It is the world-scale communications network, the money, ideas, knowledge, know- how, market that is the essence of the world culture I dealt with in the last chapter. It is the needs/resources, optimising computer by which we all get daily, hourly, votes with our feet, our money and our behaviour. There are countless millions of signals moving around this vast unexplorable network. It is another intelligence

which is largely independent of the crude switch throwing and lever pulling of political government. A government is like the little boy who thinks he is the pilot of his chairoplane at a fairground. Try as they may, our political masters cannot master this intelligent creature as it evades, dodges and wriggles around the fiats, laws, 'fine tuning' and the rest. Still less can we, the electorate, exert control with our five-yearly, cross-on-paper votes for Jack This or Jill That. We can change the people at the controls of what is not under control. Politicians of all persuasions do have some effect on this other intelligence; effect, not control. They can, with skill, devotion and persistence, corrupt or damage its function of optimising human satisfaction. For a time.

I am in favour of the 'we can sack you' democracy we have in the First World, there is no better way to limit and evict the ever hopeful clever ones who know what no-one does, the non-existent 'predictive laws', the economic system and how to 'control' it.

But these remarks are outside the usual context of artificial intelligence, which is always seen in a master/servant, controller/controlled context. The picture presented here is different. In an intelligent system all the elements are in the dual role of master and servant in a much more intimate and immediate sense than that of a voting democracy. Every element has to survive, to be kept stable, and so has its needs, news of which it sends into the system. Every active element should also have its supply role, its contribution functions. To fulfil these it must get instructions from the system. Not all the functions are vital and essential, but because of the essential nature of redundancy and competition the system would not survive for long without them.

The thing that we human beings should be afraid of is not to be functioning within such an intelligent system. This is the unhappy fate of many people in parts of the world where an ancient, intelligent and stable, if frugal, life in ecological balance with the local biosphere has been broken up and replaced by attempts to govern by systems which have not grown, evolved or spread from successful examples elsewhere but are based on simple-minded cognitive models imposed and directed by an elite of 'wise controllers' who know what to do and induce or force people to conform. Some people stay poor. The hidden, uncomprehended intelligence behind the traditional cultural and behavioural patterns they have evolved is unable to save them from ecological disturbance, famine and disease. They fall into a messy and traumatic muddled cultural pattern that is neither the old nor the new but worse than either. The first thing a politician discovers when he or she attains power is how little power it is. The clumsy central fingers on a great complex of half-understood buttons and levers gets more and more hesitant as meter needles start waggling past red lines and alarms start shrilling at all points in the largely mysterious system. It all seems so wonderfully easy when democratic rules are sternly shouting monitory backseat-driving instructions from the opposition benches. But when you sit at the controls

yourself! You are like poor Alice, with her struggling flamingo stick trying to hit her wilful, animate croquet balls through self-willed, mobile, temperamental hoops.

In the Western World, mankind is part of a of a large primitive network form. It is tentative beginning for a creature of a new type. It may not last, but there seems no chance now of returning to the safer but poorer and harsher world before it happened. The great merger of many different contending human sociozoa occurred when neolithic agricultural man replaced mesolithic man.

We cannot go back, so we might as well go on. We shall do so the better I trying to make better models of it. The chick cannot crawl back into the safety of the egg. It has to go out and face a new world. In our case perhaps an extra-terrestrial one.

IS MANKIND MAKING HIS OWN SUCCESSOR?

In creating artificial intelligence is man making his own successor? Who can tell? successor-making is the stuff of life. We produce personal successors called children, and our societal successor as we evolve new life-styles and cultures. The mesolithic hunter-gatherer sociozoa made the neolithic agricultural ones which in turn made the past civilisation and they the modern world culture. I expect that earth culture sociozoa will create their successor, solar planetary cultures, and so on out. It is the pattern, it is morphostatic, biophilic, and good for things to be so. The morality that arises from the present paradigm says it would be bad if evolution, successor -making, came to a stop?. In other moralities, take your choice. It is an aesthetic matter. Be a specialist culture chauvinist if it makes you happy; it will not make any difference in the long term. Life will choose not us.

My friend Sir Clive Sinclair talks of the time when non-organic intelligence will be so much more intelligent than human beings can be that they keep humans as pets. Even so, I see a long period in which human beings will remain an essential part of something that will grow from the present combination of biological and social intelligence. Already the outcome is affected by a growing input of artefact intelligence. That aspect must grow more influential as time goes on, but a lot of time will pass before humanity descendants disappear if this road is followed. If other roads were followed it would give man a shorter reign as he would be overtaken by a rival somewhere or speciate and divide in the normal biological way.

Species, as well as animals, are mortal. Species are much more durable but none will last forever. If evolution is not to come to a halt something must replace mankind. It might be something evolved from us or from some other animal. It might equally be, or be made by, some human artefacts. The fourth-order, new sort of living thing, the world-culture seems set for a far longer life if it survives the traumas of its first few millennia. While wishing mankind to live forever is a

wasted wish, we do not have to start planning the funeral just yet. Let us try to see that what replaces us, whenever that may be. Will they be better, more durable, more intelligent, more powerful, more exploratory, more peaceful, more orderly and more space-conquering than we are? It seems a moral and an achievable aim. The alternative seems to be a suicidal attitude, '*Apres moi, deluge!*' Deluges have to be and can be survived. Ask Noah.

10. INTELLIGENCE: THE FUTURE

We now move on beyond the realm of speculation to that in bridled imagination. Extrapolation is always dangerous and never more so than in these cloudy heights. Seeing the trend-line I have begun to map, I should be cautious but should not shy away from at least a glance along the road indicated. The implications are startling at first sight, they run counter to cherished ideas, but look at least I must. The reader may or may not accept. At least I shall know the trend oppose.

SOME IMPLICATIONS OF A HETERONOMOUS UNIVERSE

Here is a world model, one of two semi-independent realms. Realm One is the physical world of matter and energy, the Pleroma. Realm Two is the heteronomous world of the Creatura or the biota, a realm of surviving, evolving, negentropic forms living on substrates that obey the laws (statistical) of the physical realm. The forms that live on them do not because they have access to events at a level where there are no strict laws. The creatures in this realm use and amplify quantum indeterminism to draw on options to defeat the statistical weigh of physical laws. They do 'illegal' things. In the realm of the Pleroma, the physical world, there is a trend towards higher probability and increasing disorder. In the counterpoised other realm, that of the Creatura, we see a continuum of morphostasis, an opposite trend towards improbable combinations, negentropy, greater order, complexity, intelligence and the symbiosis of larger and larger combinations of life forms in more and more complex and symbiotic systems. We may have rejected Haeckle's notion of a *Scala Naturae* as anthropocentric because it sets man up as the King of the Castle of Life. But when we see that the trend-line goes on beyond man and life on this earth we may find the trend both more plausible and less attractive.

Can we, after all, get an 'ought' from an 'is'?

The Scottish philosopher David Hume (1711-76) captured a consensus for the idea that we cannot derive a value system from any observation of the world. We cannot get an 'ought' from an 'is'. Societies, he thought, decide priorities, what is good and what is bad, by a social consensus that arises spontaneously. Hume's idea was new and revolutionary; he implied that a society's ethics are its choice, they are not given from nature or even from the versions of God that various religions preach.

There is a later development of this idea that comes originally from the thinking of that other influential Scottish philosopher, Adam Smith. Friederich Hayek, a follower of Smith, and also the biologist Professor Cyril Dean Darlington have suggested that systems of ethics and morality actually evolve in a similar way to that in which biological species evolve.

This fits in perfectly with the ideas developed here. We can see the mores, rules, laws of a tribe or any other society as well as their techniques, skills and general living know-how as non-genetic heredity, a transmissible social information pool. It acts as another sort of heredity side by side with genetic heredity for the survival of the sociozoon, the tribe or other group.

The small tribal group was for some millions of years up to the last Ice Age the only type of social unit of the hominids. It follows that our genetic make-up, our hard-wiring, must have evolved for that kind of hunter-gatherer social life. The last half dozen millennia are far too short a period for there to have been very much genetic adaptation of our species to the much larger social units we now live in. This may account for many of the problems and instabilities of modern urban life. It may account for the need for a written language and morality, for laws. These are a social heredity system to correct the now unsuitable genetic emotional and instinctive tendencies.

It was only when agriculture mysteriously developed a few thousand years after the last Ice Age and the larger sociozoa, like populated cities and empires, were able to form that the simple tribal traditions and mores failed as a tool of social heredity. A better system of trans-generational information movement was needed. More complex skills and techniques had to be passed on. More importantly, different socially necessary values, goals and objectives had to be tried and passed on between generations. This new set of behaviour-modifying rules, laws, traditions, mores, was a new kind of heredity reinvented at a higher level. Only when this became stable, conservative and reliable was it possible for the much-enlarged social units with quite different life-styles to cohere and survive.

A thousand variants of moral systems appeared, a wide variety upon which selection must have operated. These moral and traditional variants (perhaps clothed in the trappings and myths of religions) must have been subject to selection both for their value to the society and also for their transmissibility and durability. The more stable and the more viable of them would survive and spread. Those with the right behavioural content and also the best emotional appeal to the 'conative' or appetite system of the hunter-gatherer substrate would survive because they would be more acceptable and transmissible. This thought might lead to a theory of religion, that myths are set to amalgamate the new within the old form; to use the hunter-gatherer, social norms to embrace the different ethical needs of larger units. Some of these new forms, these ethical systems, must surely have been such as to make the substrate groups upon which they operated more successful, more cohesive and more stable than other groups. Surely the groups with the successful ethical systems succeeded, grew, divided and spread at the expense of those with inferior (meaning less morphostatic) ethical systems. The selection would be for viability and also for

the hereditary grip upon the social unit. If you have a good system you have to be conservative or you will soon lose it.

Can everybody be rich?

'The poor are always with us' is a very old saying which has always been with us. Will it continue to be? A large proportion of the population can work and are working, producing food goods and services, yet poverty is the norm for about a third of the Earth's people. Poverty in that context means living at a bare subsistence level with most dying young and most of the young dying before puberty. This is normal for any creature in the wild. My answer to the heading question is, "Yes". I can suggest how.

In every part of the world there are comparatively rich people. Chiefs and their henchmen for instance. There is a life and death disparity between the rich and the poor in equatorial Africa and Borneo etc. In the rest of the world that kind of poverty hardly exists. In the First World, the richest third, a small minority are poor. No-one starves and the word 'poverty' applies only to those near the bottom of the wealth scale. All scales have a bottom. The poorest in the First World have shelter, adequate clothing, furniture, gas, electricity, lighting, radio, television. The exception is the tiny minority of street sleepers who seem to like that kind of life, freedom and company.

In my young days I lived very close to this level and can vouch that we had no sense of deprivation and were usually happy. There were tramps who begged for a living but no street people, they were sent to an institution, work-house hospital or what were called lunatic asylums. Now only the dangerous are confined. So nowadays poverty means living in a poor and often criminal district in restricted living space with few luxuries, no phone, or vehicle. Invalids are given cars or wheel chairs by the local authority. The affluent are full of guilt about 'poverty'. In the First World it means just adequate provision, no more. In the Third-World that would be 'riches'.

The upper limit and the average in the wealth continuity is constantly rising as science and technology acquire ever increasing know-how and knowledge. Some rich men have more wealth than some countries. More than they can ever consume physically. They can be called satiate. They want money because of the power over others that it brings and because they are interested in the skills of creating wealth and enjoy the challenge. On the whole if their products are not harmful they are doing good, creating more wealth than they can consume. They do it by their skill and brains or by their investment. Without free trade the world would be poorer.

In the past, wealth was always limited by the amount of available land, water and other materials together with respect to local population. Automation reduces these limitations.

For the pioneers of monetary theory, Adam Smith. Ricardo, and so on, the world, was economically a very different place from that in the present and future. The theory of wealth has to be revised. Smith's definition of wealth is now out of date. He said that the essentials were Land and Work. These were what created all wealth then. It was largely about food. What he left out was knowledge, particularly scientific knowledge that was rudimentary in 1766 ('*Wealth Of Nations*').

Later economists, Ricardo and Malthus, were writing in the early part of the eighteenth century when modern automated industrialism was in its infancy. Neo-classical economics started before the century with William Jevons. It was developed by Carl Menger who died in 1921. There followed Pareto, Marshal, etc. But Neo-Classical Economic theory is still accepted by the Oxford Encyclopaedia, as the last word today. Not one economic thinker has take account of modern physics and Einstein's shattering discovery that mass can be converted into energy or that most physical work is now done on machines that are made by machines.

Adam Smith (still surprisingly quoted by policy makers) has also left out energy as an essential of wealth. The sun is a constant supplier of that if you can collect it by growing crops. Anything make-able can be made with knowledge and energy, and a vast amount of energy can be got from very little mass. With more and more automatic manufacture you can produce more and more with less human work. So the future shows no definite limit to wealth. With the mass of the water in the sea alone there is so much energy that we could even manage without sun power that is still hard to collect. With no limit to wealth everybody can be rich. But some will still think of themselves as poor though they are satiate.

Wealth is not the problem. The problem is distribution. That is not the First World's fault. Where there is a despotic ruling class that will not join in the world economy then there is poverty at the bottom of the scale. With democracy and free trade, there will always be entrepreners who will come in and start raising the standard of living. It is undemocratic rulers who want to stay in power that make people poor, not rich nations. However much wealth is poured into a 'poor' undemocratic country, the result will always be that the rich get richer and the poor remain poor. The result of charity is more poverty. I suggest there is no limit to wealth in the rich world that cannot be overcome.

The limit at the moment is the completely unjustified fear of nuclear power. There is a refusal to take the benefit from the vast investment that has been made in nuclear physics for defence purposes. There has been one war-time atomic explosion which killed far less people than many earthquakes, plagues and other disasters. Because of public opinion, the dangers from radiation emitted from nuclear plants are not are allowed to be higher than those from a normal

environment which is full of radiation from many sources anyway. The other ways of producing energy, without exception, kill many more people with toxic substances in fuels and accidents. They cause global warming and many other undesirable effects. It would be a simple matter to put nuclear power stations well out in the deep ocean and send the power to land through cables. There, the absurd precautions that raise the costs so much would not be needed. Such plants could be almost fully automatic. There are hundreds of indifferently run nuclear plants in the Second World and we are still waiting for a second Chernobyl.

The First World only deprives itself of this excellent and wholly beneficial way to get one essential of wealth, energy, because the public are scared and influential in democracy (without which there is little wealth). The extremist environmental groups are responsible for the high cost of nuclear energy. There are very much bigger dangers that are real and they are dangers that nuclear power alone can solve.

The Earth was hit, in the last century, by meteorite that would have destroyed London if it had fallen there. Meteorites that wipe out most species fall onto the Earth every billenia or so. There is even one now whose path is not too well understood that comes dangerously close from time to time.

If we accept the lesser danger (which we have anyway) of developing nuclear power and lowering the expense of it, then why bind it with excessive regulation, especially in the First World which does have the knowledge to develop it safely?

There is a better chance that we can survive a certain final catastrophe to Life on Earth. How. Noah knew. By pouring money into nuclear research and nuclear power so that we have the wealth and knowledge to get as many Arks into space as possible. Take Life knowledge and Intelligence off this single precarious planet and get it, first into the solar system and then - Beyond. The first generation of nuclear physicists is growing old and the incentives for the clever to get into that field are small. I understand that there is a debate going on as to whether a few more million should go on in that direction. We need a counter to the misguided popular prejudice.

Social Heredity Is Lamarckian

This sort of social heredity has an advantageous feature that does not apply to genetic inheritance. The biologist Chevalier Jean Baptiste Lamarck (1744 - 1829) held that acquired characteristics could be transmitted between generations genetically. There was much dispute about this but the view was finally disproved and discredited. There is no way that the somatic experiences of an animal could change the genes beneficially. Mutations are random. But social heredity is Lamarckian. With social heredity, acquired characteristics can be passed on. Man is an imitator, the successful ethical system could be and almost certainly

was copied by other tribes and thus acquired a life of its own as a breeding, developing form, a morphostat which survived on more and more widespread social substrates.

The idea that morality evolved by the evolutionary modification of tribal customs is reinforced by the origin of both the words we use in this context. The word 'ethics' comes from the Greek *ethos*, the word 'morals' comes from the Latin *mores*. Both had the same meaning, 'custom' or 'tradition'. Originally the words obviously applied to the small social unit, the tribe or village. The evolutionary view of morality is a contradiction of the Hume's view because the 'is' of the successful tribe spread and developed the 'ought' of its mores. But Hume would say that consensus must have arisen for the moral code before it had proved successful and thus it was imposed by the society without referring to 'the way things are'.

I now have the task of looking back at the problem of ethics in the light of the world model I am elaborating. We may agree that our present ethical systems are those that survived because they spread of themselves and because the societies that embraced them were stable and successful. They survived to pass them on to us. The tribes with inadequate moralities had no survivors left to pass them on. The morphostat, the ethical system, died. But what of the substrate, the people? They may have died out with the unsuccessful pattern, but that does not follow. People live in societies but they are free to leave them. Successful societies, firms, nations and organisations grow, they acquire more people. Where do they come from? A lot come from unsuccessful firms, nations, organisations. It seems like a very good idea. It is another reduced-punishment-for-failure idea.

Accepting that Hume was basically right, can we still get an 'ought' from the 'is' of the two realms paradigm? I suggest that we can begin to see, if a little dimly, by the faint new light of the model presented, the outline of an *existential* base for ethics. I invite you to make the arbitrary but understandable (*aesthetic*, if you like) decision that we identify with the creatures rather than with the things, that we come within, are citizens of, the realm of the forms, that we have a patriotic/ chauvinistic prejudice in favour of the realm to which we belong, the biota. We are reluctant to see the end of life. If you accept my invitation, and given only that we can get some 'oughts' from the trend and shape of the 'is' of the model, if we accept it, we can embrace those 'oughts' as existential ones.

If we are life-hating Pleromists who see life as a disease of matter and generally a thoroughly bad thing, we would get a clear and definite Pleromist morality from that view. We should know how to behave and start making nice big nuclear bombs. But if we are ready to discriminate, if we are unashamed realists, ready to favour the Creatura, just because it is our own kind, then all we have to do is to set the Pleromist morality on its head and we have a set of 'oughts' which are derived from the 'is' that I am exploring.

To begin with, we would favour the progress along the track of the morphostasis continuum (though we might argue a lot about where the track led and how to go along it). However, there are important aspects of the continuum that are confusing, that do not fit well with the rest. There is a singularity in one of the trends, which was first noticed by Teilhard de Chardin. Like all biologists since Darwin he noted the strong diverging trend of evolution. There is an evolutionary tree-like form. The trend is towards ever-greater diversity. All species complexify and divide again and again towards an explosively increasing variety of ever more specialist forms to fit ever more and tinier ecological niches.

But he also noted the strange, emergent, contradictory and recent convergence trend, the trend towards co-operation and symbiosis. Especially he noted the convergence and combination of human cultures in the last few millennia after millions of years of divergence and mutually competitive repulsion. I am astonished, as the evidence emerges from ethnological and anthropological studies, how little modern thinkers are perplexed by that recent phenomenon - our world civilisation and culture. How, just how, did a creature bred for millions of years to live in small, mutually repelling, territorial, hunter-gatherer groups suddenly become the creature that built the world- wide civilisation of modern man?

There are a few thinkers who followed that track. With his concept 'world brain', Herbert George Wells suggested an echoing thought. Bernard Shaw's 'Life Force', and Karl Popper's 'World III', have the same philosophical root. Contradicting Hume, Hayek and Darlington, there are signs that among educated and informed humanity there is beginning to emerge an expanded or extended view of morality and ethics which does not fit the evolutionary paradigm for cultures. There is a widespread concept, a view of morality that is seen as valid for mankind as a whole, a most extraordinary idea in any biological perspective. There is a widely accepted, passionately believed view of mankind as a single whole, which is contrary to all human traditions and biological precedents. The norm for all highly successful species hitherto is this. Soon after they attain pre-eminence they begin to speciate, divide into different subspecies and then species under population pressure and competition, each singling out some sub-niche for themselves. Mankind is the sole exception. The slight beginnings of speciation visible in the different ethnic groups seem to be coming to a halt as modern travel remixes mankind. I personally rejoice in this as I notice it and it has led me to these views.

The world culture is increasingly a whole-earth community, a noosphere, as Chardin would say. The emergence, growth and success of this earth-scale phenomenon was not predicted and is not congruent with anything in the human evolution before everything began to change about six or seven short millennia ago. It is a surprise from every point of view.

But it does fit one model, it fits the model of the morphostasis continuum I have advanced. It fits it in this way. We may see the world culture as a symbiotic/competitive network of many cultures rather than of many people, a creature in its own right, one that by its existence may preserve a biosphere that might otherwise have been destroyed.

Competition And Co-Operation

There is a problem that has long bedevilled human thought concerning the social value of co-operation and of competition. The two terms are crude and unspecific (they have a very low sufficiency) and perhaps for that reason there is much confusion at the social level as to whether social unity requires the sharp reduction or even the abolition of competition and a society based entirely on the principle of social co-operation. Others emphasise the optimising value of competition and see it as the greatest force for advantageous change and improvement.

The morphostatic point of view seems to help because it makes clearer what the problem is.

Competition is the offering of options. All optimisation, learning, whether genoplastic, ontoplastic or social, requires two things. There must be options, choice among them, and the raising of the probability or frequency of the successful options being chosen in the future. Co-operation is a deeply essential part of the morphostasis continuum. The co-operative, symbiotic assembly of larger and larger systems is the main trend, as I tried to show in Figure 1.

Competition and co-operation are complementary, two essential aspects of all morphostasis. They are not opposites and they are not incompatible. Every morphostat combines the two features in some proportion and in the proportion lies the success or failure of the morphostat concerned.

Meta-learning, the higher, more strategic-level learning, is the business of Finding the best of competing ways to optimise the combination of co-operation and competition. That seems to be the most important human problem today, both socially and in the field of commerce and between nations and life-styles. The taxonomon model of the learning process has a simple built-in system by which the nodes compete to capture noisy patterns. Non-centralist societies have the independent trader, farmer, industrial firm, which works within an environment of social order and restraint, a compulsory co-operative framework of law, order and accepted commercial and legal practice. When it works well the result is the best of both worlds, a system of competition within co-operation that seems to produce increasingly good results for most people.

We should not forget, however, that the very successes of the recent world cultural co-operation that has so suddenly raised and changed the world

population may have its long-term dangers. The usual biological sequence where there are great successes is population explosions, great failures and population crashes. Industrialism in field and factory has called into being an enormous number of people and societies whose ecological fit with each other and the biosphere is new and untested. It may, I fear, prove to be a temporary success only.

However, there is one great hope comprehension. Mankind has already shown that collectively it is able to solve such large problems in utterly new ways, in which the painful trial-and-error, 'die if you fail' method is replaced by soft learning, learning via models. Simulations, analogous worlds can be created by scientists, in which lessons can be learned the soft way, the cognitive way, as they are in an ontoplastic brain or to some extent in governed competitive society.

We are a long way from the central World Government that we dreamed of early in the century after two world wars, but the development of the poly-hierarchic world brain, or world computer, our developing world culture, is at least promising. Now in 1999 we can be hopeful.

Life, death, reward, punishment, pleasure, pain, gain, loss. I wrote above of soft learning and elsewhere of the trend towards reducing the severity of punishments.

Certainly there can be no real learning, optimisation, without options and competition and therefore without winners and losers, rewards and punishments. The visible trend, however, is that as the movement along the morphostasis continuum proceeds the rewards and punishments get less severe. In earlier forms of life there is only one penalty, capital punishment. The more advanced animals have an improvement: pain and pleasure with an appetite system to supply the signals that change behaviour patterns in the learning phase. The animal is modified, not the whole species.

A later stage still is the cognitive brain, which is a simulation system where behaviour patterns can be tried out in an internal model of the world in world simulation experiments called thinking and planning. The pleasure/pain input here is more subtle: pleasure at a good idea or plan, irritation at a bad one or a mistake. But the pleasure/pain spectrum is there, it must be to drive the learning machine.

At the social level the reward/punishment antithesis at the sociozoan level returns to severity. Unsuccessful social groupings, firms, associations, cultures and life-styles, tend to die out, often not by the death of the components, people, but their departure, by their feet voting to join preferred sociozoa other firms, cultures, countries. At the individual level reward and punishment are even further softened to gain and loss, approval and disapproval, honour and disgrace,

rewards and fines, liberty and imprisonment. (Prison is bad but preferable to torture or death.)

What worries me is that in the face of the entire history of life and humanity there are some good and kindly people nowadays who feel that even the softest form of marking failure, to help the learning process, is bad and to be avoided. I see no escape from one fundamental truth. Where there is no success and failure, or where it may not be recognised, where there are no winners and losers, there can be neither learning nor progress – nor in the long run, survival. Each and every creature and combination of creatures on earth today is a winner, selected in a multi-trillion stage low-chance raffle. This is an unpalatable fact to most of us, brought up as we are in the comfortable lap of a protective State. Like it we may not. Ignore it we dare not. As social animals we may and must ameliorate the lot of the losers, ideas-people, planners, scientists, associations and people, but losers and everybody else must face facts and signify failure. There can be no life without options and choice. Where you have options and choice you have losers. Losers are essential to any morphostatic system but we cannot build the future from them and we must not fail to learn the lessons they are there to teach. Losers must be allowed to do their important job and make their contribution.

There are uncomfortable and unpopular conclusions that follow from the present hypothesis. I am not going to pretend that they attract me. I, too, am a product of my time and culture and have all the standard moral preferences of my age. I reassert my warning against untested moralities based on no more than man's simple reasoning powers. However, we ought at least to follow the track in our minds. For pilot trial and possible selection I offer these ideas. Most thinkers in the mainstream of Western thought are appalled by the horrors of the increasingly damaging nature of the old human practice of warfare. Believing that conflicts within nations are often less frequent and less damaging than those between them, they seek the unification of humanity. They want to bring us all together into one political unity, a centrally controlled World State. Most of the world's intelligentsia, including myself, have gone along with this.

Noting the fierce, intractable and bloody enmities between those of ethno-cultural groups, they seek to play down and ignore these differences and move towards a multi-cultural, poly-ethnic goulash society with standardised men and women living in a society where all the cultural styles cohabit the same space at the same time. Again we have all gone along with this idea. Now I have argued that, at least in origin, moralities evolved to serve a function in the survival of small, competing societies. They were based on religions until the so-called 'Age of Enlightenment', when there arose some d.i.y. versions that are still under test. These latter-day 'rational' moralities borrowed a lot from the traditional ones but they contained some very important new elements, among which the idea of individual freedom, which seems to have released a tide of enterprise and beneficial change.

Now we see being painfully born the beginnings of a human-race scale morality that which seems to be aimed at the survival and successes of the species as a whole. This trend appears in the United Nations and the whole panoply of international organisations. That this should appear at all is very much against the history of our species and very strange indeed, but it fits the concept of the morphostasis continuum. Further and even more mysterious, there is a third-order morality, a biophil morality, based on a reverence for life itself. The growing ecology lobby and the Green parties bear witness to the emergence of this species' altruistic meta-morality. This too fits. It is possible that this represents no more than a form of reaction to modern industrialism, the hunter-gatherer within us all protesting at our new life-style. But there does seem to be a real underlying rational concern that deserves serious attention.

I am sad that a few attention-seeking extremists are bringing these ideas into disrepute with selective, simplistic, sentimental campaigns and absurd public antics that do no good. But the fast-growing popularity of biophil ideas among those of the world culture is a sign, and an important one, that biosphere morality is possible and beginning to assert itself.

BIOPHIL ETHICS

What should a biophil ethical system be like? Biophil (life-loving or morphostatic) ethics is defined as that general ethical system that arises from a favourable view of the Creatura and acceptance and approval of the morphostasis continuum. I say 'general', bearing in mind that all existing ethical systems were evolved to serve separate human groups, nations, cultures. These systems have survived and spread because of the success and the symbiotic merging of the separate and previously contending cultures that adopted them. There is much in common between the surviving but different ethical systems around the world but there are many differences also. Many of the differences have a ritualistic, religious aspect that seems, regardless of culture, to have little or no survival value to the group or culture. These mystical elements are much more various as between religions. This helps us to see what is essential and what is contingent. The more variable the feature, the less necessary it is likely to be. So the mythical elements in religion may be their selling gimmicks, the means of spread and propagation between generations. They may be part of the social heredity of morality.

The international governmental community has patched together an eclectic amalgam of the different ethical systems. An attempt to express this is found in the United Nations charter and many other complex treaties and law-like agreements.

Much more important, there is an eclectic but tacit general world morality. This is a selection from the morality systems of the successful cultures. It operates in the world communities that meet and interact in the world culture. This morality

pertains to and informs the commercial, travel, industrial, scientific, artistic, and all other types of international and intercultural communication and interaction. This is a tacit, unwritten, but universal patchwork of understood practice in each field of international contacts that is self-regulating and un-enforced. More than anything else it is what makes the present world civilisation possible. This moral code is observed because people, firms and institutions of all kinds have to acquire and keep a reputation if they are to operate in this milieu at all.

Generally the biophil moralist is humble. He does not arrogate to himself the intelligence to be a moral innovator. He does not trust his intellect to tell him what is good and what is bad for the Creatura. He fears that it is not given to us to understand and predict the outcome of radical social changes in enormously complex modern social assemblies in an age of very rapid technological change. He suspects that what has worked well for a long time and brought a society to enlightenment and prosperity is often quite good and should not be impulsively changed. Innovation is not barred, but radical changes should be tested to destruction in pilot experiments before the risk of general adoption is taken.

Our biophil moralist would favour as many various experimental cultures, life-styles and communities as possible. Only this will offer variety for selection by the social evolution process. But such innovative communities should be seen as what they are, pilot experiments. The biophil sees the working of the commercial market as a symbiotic combination of peaceful, limited competition and co-operation, within an ordered system of social constraint. He might be against the tendency to standardise and equalise cultural life-styles within a nation. Let communities as well as companies run themselves, within some overall constraint, in any way they find good. Let there be variety for selection. Without it there can be no progress along the morphostasis continuum.

The tide of beneficial (at least superficially) change has, as I have said, brought about, among intellectuals and many others, an atmosphere in which it is thought that *all* change is good. We hear pundits talking about social change as though this were an obvious benefit in all cases. Their sense of morality is mainly concerned with speeding up this tardy process.

Biophil ethics would indicate caution about such social change. The wished for effects may be splendidly obvious but how sure, in the present state of sociological prediction, can we be that we fully understand the long-term side-effects? I was one of the generation that rushed headlong into, for instance, rejecting Victorian morality and opting for a permissive society. It seemed positive, an obvious improvement, setting us all free to 'do our own thing'. And it did, for a long time. It may be that the increase in crime, divorce, one-parent families and problem families, and a growing group of disturbed and under-achieving children are other and utterly unrelated problems. But it is easy to be sceptical about that. To the fair question, 'Are we all obviously better off and

happier for these changes?', it is hard for one who has been through it to answer with an unequivocal 'Yes'.

New moral ideas must arise in thought, they cannot yet be tested there. Unlike Jefferson, I do not believe that morality is self-evident. It may be that we can in the future develop artificial intelligence models which will outthink our limited intelligence and be able to know what the long-term effects and side-effects of changes in moral laws will be, but the almost universal failure to develop successful *economic* models shows that that time is a fair way off yet. (If really predictive economic models were possible the discoverer would be visible as the owner of most of the world's goods in a very short time.) I begin to have doubts about some of the moral laws and precepts that became so widely followed in so many cultures since, so recently, the British utilitarians and the French encyclopaedists began to design do-it-yourself moral codes. These thinkers may have been presumptuous to argue that moral codes based only on human cognitive intelligence should be adopted untested.

With that caveat made, I can follow the direction of the morality trend-line indicated by the paradigm we are looking at that of the two realms and the morphostasis continuum.

If we have made the decision that we are favourable to the development and optimising of the Creatura in this way of looking at the world, it follows that what is good and moral is that which favours the continuation of all the trends mentioned in Figure 1. Life and information forms are good and desirable as they reduce in probability, and increase in durability, stability, divergent variety, complexity, organisation, symbiotic convergence, comprehension of and power over the physical universe. We need measures for these properties or trends if they are to be judged rationally.

However, rational judgement is what we must be cautious about trusting. Competitive trials of life-styles, cultural, moral, economic and other systems, within some overall biophil constraints, seem to be the way forward. One day, maybe, there will arise world models that are good and congruent enough and entities, biological or artificial, that are intelligent enough to make confident changes without going through the painful routine of competitive trial and error. That day is not yet.

The principle of least error cost

Paradoxically, thinking about ethics and morality is bedevilled by Gods. If morality is immanent, unchallengable and selectively revealed to an elite of wise men directly by God, it is an impertinence to think about it at all, even if one is one of the elect oneself. Who are you, you small, ephemeral mortal? Take your orders and obey!

But out of the model I present there comes a view about morality, what it is and what its function is. As one of those who questions such a selectively revealed morality I am in a weak position. In the light of my position of biophysical dualism, I do not trust the so-called rational human judgement overmuch either. I agree with the rationalists in questioning revelation but I disagree that they know best and that morality is self-evident.

Pragmatists like myself must judge morality by its results, the kind of life it promotes in the morphostats which have it, not only the people but their institutions, societies and cultures. Do they make progress on the continuum of morphostasis?

There is one thing I mentioned earlier that should be emphasised at this stage. It is Pascal's Policy, the fail-safe principle of least error cost. If we do not have enough information, if we cannot know all the long-term effects of any given moral policy or precept, we can at least be guided by Pascal's Policy. When we stand perplexed between two courses, we can try to calculate which has the least error cost and usually it will be best to follow that road. Generally the moral advice must be to make no changes where you cannot be sure they are beneficial, unless you can experiment first and check your judgement. If you can experiment, subject a pilot sample to the novelty and adopt it only if it is clearly beneficial, then do nothing unless faced by a very severe threat to stability. The greater the severity of the threat, the more radical is the sort of untested change that may be wise.

The feature of human group behaviour called panic seems to be counter-biological but it is not. When the threat of death to a group is dire enough, leader-proneness fails and random individual behaviour asserts itself. It pays off because some deviant, radically new behaviour may be the one way to preserve that form through that crisis, by preserving just one lucky member. 'Every man for himself' is a pay-off policy *in extremis*.

But except in emergencies, the known and tried is the best course, according to this model. Experiment should never cease but really new ways should pass a test to destruction phase before they are widely adopted. I predict that the time has come for return to the traditional suspicion of innovation after the recent incessant tide of it that has been sweeping much of mankind. I predict an age when stability will be more valued than further unceasing 'progress', 'development', 'social change', 'growth' and all the other new goals that have been accepted so lightly and universally with so little thought and almost no prior trial.

Power

Some social scientists are much concerned with power. What is power in this context? It does not mean simply the amount of energy available as measured in

watts (or terawatts). The meaning we seek lies in the idea of a trigger or relay that we explored earlier. The power of a morphostat node lies in the amplification factor more than in the crude wattage. By power in this sense we really mean 'power available to mind', or 'open to purposive or moral choice'. The larger the morphostatic system, the greater will be the amplification factor needed if it is to be able to do what is essential for its survival and further evolution offer variety for selection. Now the only place the indeterminism of real options can come from is the microcosm. It is good for a morphostat to be able to obtain penny-toss type options from the quantum world. But power available to mind and power controlled by chance are not the same things. How do we reconcile the contradiction? There are two ways and one of them involves multiplying entities without necessity.

First: what is offered by the microcosm to the mesocosmic and maybe later to the macrocosmic morphostat are *options* not choices. If the entity is a protozoon it chooses and lives or dies, but if it is a creature of more advanced intelligence it has a world simulator called a brain that has a large datastore of tested, learning patterns, so the new option (called an idea) can be tested in the simulation programs called imagination and cognition before being tried

THE FUTURE OF INTELLIGENCE IN THE REAL WORLD.

The suggestion is that those aspects of the mind we call imagination and creativity have an important contribution from the microcosm. When we seek to be 'free' we are asking for the freedom to behave whimsically, To make experiments which are not justifiable rationally, for access to our own personal penny tosses. It is clear that however large living systems become, those that despite their size preserve this creative aspect best will in the long term, be likely to produce the most viable survivors in the long term, because there will be less restraint on the variants offered for selection. I mentioned an idea which unnecessarily multiplies entities. It is one in which I join many speculations, such as those of Eccles, Julian Huxley and Bergson, in blatant defiance of Occam.

Possibly a purposive mind itself is an aspect of all matter and thus a micro-cosmic phenomenon in itself. Maybe it is not simply options that the microcosm offers but sometimes *selected* options and even sometimes marginal triggerlike decisions. Suppose the universe is indeterminate, as I argue, and that there is a power or force or God, not omnipotent and omniscient, but a purposive origin for the realm of the Creatura that is striving against the other realm, the Pleroma. If we grant it only power at the margin in the region of the indeterminate then it would have to operate at first by building up stable platforms of order and developing them as triggers to mesocosm systems. It would have to build up life forms, create a morphostasis continuum.

The next question is vital. What difference would it make to our view of a biophil moral code, whether or not we accept the above deist hypothesis as dogma?

None, I suggest If it were not so, if there were no original intelligence, mind force, what you will, surely the nature of the two-realm universe (if it is at all congruent with the one out there) wooed itself bring about the emergence of something like it. I gave my reason for this in my remarks on the origin of purpose in the universe, earlier in the book. Thus I am a deist.

The Convergence Of Intelligence

As I have said, there would have been no reason, before it happened, to anticipate the convergence of, the co-operation of, and the communicational integrity of the world culture of men that I have described. Nor would there have been any reason to anticipate the idea of human unity that is implied by the United Nations and the many other international hierarchies and networks that have arisen in this century. That they do not work well and dissatisfy many is not the point, they had no call to arise on an earth dominated by a successful, strongly competitive, intensely territorial, hunter- gatherer predator. That this convergence should happen was unlikely. It is not less likely that, after the trauma of first contact is over, the same kind of symbiotic convergence should continue, if mankind, while he is still around, should make contact with other life forms from beyond the earth. Most earthlings who have thought about it today would expect that, and that fact too is odd and predictive. The life-term of man and the distances in the macrocosm make such contact between men and other types of intelligence most improbable but the scope and duration of artefact intelligence is an unknown that is all potential. The childish science fiction dream that recapitulates the human convergence in imagination among the galaxies may not be so childish or laughable as it seems at first sight.

In forms we cannot possibly imagine but in ways in which we may even now be playing a precursory part, the interstellar and even intergalactic contact, interaction and convergence of intelligence may happen. The primitive hominids never dreamed of our present world culture but they were contributing to it by the way they were and the way they were going.

John Good was the first to see that if man's skills can make an intelligence that can surpass his own, then that is the first step to an intelligence explosion that may not have limits. The advanced intelligence takes a hand in designing its successor, and that successor in the design of its successor. This was John Good's reasoning.

Immanent purpose?

Following this fanciful image a little further, we see a dual universe where the matter/energy realm is vastly dominant and all-pervasive. So far. That realm operates under its own invariances and constraints in the microcosm. Although at base it is indeterminate, in any large aggregate it appears to be determinate because it is governed by the vast weight of statistical laws that apply to very

large numbers. The other realm we posit is that of an intelligence, a Mind or God or Purpose or Life-Force. It exists/existed in some unknown, perhaps non-material form. It had no power in the middle world of the mesocosm. So far. But in the realm of the indeterminate it may have some marginal power.

If that which favours the Creatura is a striving, purposive God who wants Power over events in the mesocosm, the level of our human lives, and later the macrocosm of the stars and galaxies he must build up trigger-chains and cascades so as to amplify his posited, minute power over the quantum world. And to do this he will first need to get a purchase on the cascade of disordering change, the slide to disorder, to entropy, that is universal at mesocosmic levels. The first main objective must surely be to create platforms of stability in the mesocosm. These are needed as a base for operations. Nothing can be done in an anarchy of constant change. The intelligence must retain and develop control of his stable forms via trigger-chains and cascades built into them.

What I have described is precisely the morphostatic model for life forms. A human society that can change itself and the world because of an idea in one man's mind is a model of a system where the microcosm is struggling for control of the macrocosm. At a later stage the God, the Intelligence, the driving force of the Creatura, would need to intervene judiciously via his many-link trigger-cascades, to guide the development of more and larger intelligence-driven creatures, civilisations and cultures. He would have to make them competitive to achieve excellence. He would have to make the competitive forms diverge at first but that would create a dilemma. To have real control he would need a convergence stage at the end. There would have to be symbiosis as well as competition. The convergence of the biosphere as a whole and of what Chardin called the noosphere and I have called world culture is congruent with that thought.

Further, and this is the uneasy thought, the beings and cultures and institutions that were built up might be influenced to experiment with the manufacture of even more advanced forms of mind and intelligence. These artificial forms would not be constrained by the age-long historical clutter of biological systems that had been developed on the crude, slow, genoplastic learning plan. The trial-and-error development of artefact simulations of intelligence could be informed by highly developed intelligence (which the God may have direct access to at a sophisticated level, may be able to *inspire* - to use the common religious phrase). This would be better than the marginal effect on the penny-toss, try-everything mode of genoplastic evolution.

Further such development of Artificial Intelligence is unconstrained by the painful biological evolutionary path by which biological forms came to exist. An example of such is the commercial/scientific community as a whole on earth today. It can invent viable artefacts much more quickly and easily than evolution can do.

Let me propose a moral test. Suppose human technology does, in fact, lead to the emergence of a new form of morphostat, a surviving, organised, self-replicating form of non-carbohydrate life such as intelligent robots. Suppose this is created at first by man and later replicated and improved by itself, perhaps on a base of electronics and silicon components, perhaps on bases and taking forms as yet undreamed of. Suppose this life form, as I think it would be, could live longer, travel further, be more powerful and intelligent than mankind or its associations. Supposing this life form, our creation and successor, should move out into the universe and continue the course and work of the biota there. Here is my moral problem. Would that be morally wrong?

Men are mortal. So are species. All of them. Yet old men think and plan for their successors, the children who will survive them and inherit from them. Is that immoral? Can we as a species and a culture achieve a feat of meta-morality? Can we think and plan for life forms that may one day supersede us? Should we take on the almost impossible intellectual task of opening our morality window wide enough to accept the inevitability of our supersession and try to see that what survives and replaces us is a continuation and development of the universal trend of the Creatura which I see?

This seems to me to be the most important moral question of all.

INTELLIGENCE IN THE FUTURE

In trying to show what an ethics based on the morphostasis continuum would be like I have given an admittedly vague and fuzzy picture of what would be a *desirable* future for a biophil, a life-lover. What will be the future of life and its product, intelligence?

I have claimed that the least-error-cost belief is that we live in a forked world in which there are real options. It follows that prophets are an endangered species, especially those in such heady heights of speculation as this.

Caveats over, I see this. The human race begins to dominate the biosphere of this small, lonely planet and together we are moving towards an important decision fork.

Homo Sapiens and his social institutions has, over the last few thousand years, been growing and changing at an explosive, exponential rate. Even in this tiny time slot there have been several surges and relapses, as various large civilisation patterns have been tried and scrapped. The present world civilisation and culture will either take off and reach the next plateau of stability at a higher morphostatic level, or relapse again for some time before it makes another salmon leap at the falls.

A Fork In The Road

Along one fork, the next plateau will be a stable, durable, biosphere on the earth supporting a world culture which, not at the command of some central control, but because of its autonomous polyhierarchic inner dynamic, pushes on up the improbability and negentropy slopes towards more intelligent people, societies and artefacts. These will, if not this time then at a later attempt, reach out from the earth and start to inhabit and populate the bodies and space of the Solar System and later, one day, galactic space and other worlds. In the very long term I see contact with other types of intelligence, competition, strife, then co-operation and resumption of expansion along the continuum.

But what is it that will reach out from earth's surface? What will the adventurers be like? If they are to be men and women, humanity will have to speciate, breed space folk, because the cost and difficulty of sending earth like environments around the universe will be too high. But we do not like speciation, genetic breeding applied to people. Also it is very slow, generations are too long.

The conclusion is that if future societies believe what we do then people, some people, somewhere at some time, will make the space adventurers. !ntelligence the seeds of sapience and comprehension, which, who, will go out to continue the work out there. And they, unobstructed by the age-long accumulation of hardwired biological decisions on our evolutionary track, will have to learn to make and modify themselves using cognition and its developments to help in offering options for evolutionary choice instead of the blind chance which brought life to where it is (as we believe). On Earth the development of artificial intelligence will bring many changes. The societies that opt for it will be found to be more stable, prosperous and powerful than those that do not. Eventually the feet will vote, if the hands do not, for a system that uses their aid. Or, alas, since we are combative, there will be wars. We can guess where the advantage will lie. You should not put your money on the side without or with less effective AI aid. Either way, the influence of the two other forms of intelligence will grow as they combine to form a single one. The social intelligence of the world culture will be informed and supported by the growing influence of intelligent artefacts.

The Other Fork

That is one fork. What is the other one?

This biosphere will be just one more seed that fell on rock. No problem! There must be plenty more out there.

Humanity will take the soft option, will spurn and reject a rival in the sapience game. We shall spread our earth culture more widely; more and more people will have a longer, more healthy, easier and more amusing life. And eventually there will be a war, a comet, a sun disaster, or disease, and that will be the end. At best, if we survive for a long time there will be the speciation that lies in wait for

the successful. What we now call a mankind will be not so. Mankind will be to the successors what the chimpanzee is to mankind, other.

There are sound biological reasons for believing that the further improvement of human intelligence by evolution is limited. The ground plan of our brain is set, hardwired. Natural selection could increase the average intelligence but it is unlikely to do much to raise the peak. (This has been shown in breeding experiments with rats. They have bred two populations apart so that the maze running-nous of the 'maze-thickest' of one group was better than that of the 'maze-brightest' of the other. But the *brightest*, maze-bright, did *not* improve).

Therefore, the future for humanity if it neglects to provide for succession is that of all such morphostats, it will have none.

BIBLIOGRAPHY

A list of the books that have influenced the text.

Aleksander, I. *The Human Machine.* Georgie Publishing

Alt, F.L. and Rubinoff, M. (editors) (1965) *Advances in Computers,* Vol 6. Academic Press, New York

Andreski, S. *Social Sciences as Sorcery.* Pelican

Andrzejewski, S. (1954) *Military Organisation and Society.* Routledge & Kegan Paul, London

Andrzejewski, S. *Social Sciences as Sorcery.* Pelican

Ardrey, R. *The Territorial Imperative.* Collins

Arkad'ev, A.G. and Braverman, E. M. (1967) *Teaching Computers to Recognise Patterns.* Academic Press, London (1967)

Asimov, I. *The Human Brain.* Mentor

Asirnov, I. *Guide to Science.* Pelican

Ayer, A. L. The *Problem of Knowledge.* Pelican

Barasch, D.P. *Sociobiology and Behaviour.* Elsevier, Oxford

Barton, A.W. *A Textbook of Heat.* Longmans Green & Company

B.B.C. Publications *The Laws of Disorder*

Beer, S. *Cybernetics and Management.* British Steel Corporation

Bell, D. *Decision Trees in Pattern Recognition.* National Physical Laboratory, D.T.I.

Berelson, B. and Steiner, G. A. *Human Behaviour: An Inventory of Scientific Findings.* Harcourt, Brace & World, New York

Bergson, H. (1911) *Matter and Memory.* Alien & Unwin, London

Bono, E., de The *Mechanism of Mind.* Jonathon Cape

Cairns-Smith, A. E. *Genetic Take-over.* Cambridge University Press

Cairns-Smith, A. E. *The Life Puzzle.* Oliver and Boyde

Carter, C.O. *Human Heredity.* Pelican

Cattell, R. (1972) *Beyondism.* Pergamon

Cattell, R. (1971) *Abilities: Their Structure, Growth and Action.* Houghton Mifflin, Boston

Chardin, P.T., de (1965) *The Phenomenon of Man.* Fontana Books, London

Charniak, E. *Introduction to Artificial Intelligence.* Drew

Cherfas, J. *Man Made Life.* Blackwell, Oxford

Cohen, J. *Reproduction.* Pergamon

Cohen, J. (1967) *Living Embryos.* Pergamon, Oxford

Cornell Aeronautic Laboratories Inc. The *Perceptron*

Crick, F. *Life Itself.* Simon & Schuster

Culbertson, J.T. *Consciousness and Behaviour.* W. C. Brown & Co., Dubuque, USA

Darlington, C.D. *The Evolution of Man in Society.* Allen & Unwin

Darlington, C. D. *Genetics and Man.* Pelican

Dawkins, R. *The Selfish Gene.* Palladin

Feigenbaum, E.A. and McCorduck, P. *The Fifth Generation*

Foss, B.M. (editor) *Perception: A Cross Cultural Perspective.* Penguin

Good, I.J. *Speculations Concerning First Ultra-intelligent Machine Advances in Computers.* Academic Press, New York

Good, I.J. (editor) *The Scientist Speculates: An Anthology of Partially Baked Ideas.* Unicorn Books, New York

Gould, S.J. *Ever Since Darwin.* Bumett Books

Gregory, R.L. *Mind in Science.* Peregrine Books, London

Gregory, R.L. (1970) *Eye and Brain.* Weidenfeld & Nicolson, London

Gregory, R.L. (1970) The *Intelligent Eye.* Weidenfeld & Nicolson, London

Gribbin.J. *In Search of Schroedinger's Cat.* Corgi

Guildford, J.P. (1967) The *Nature of Human Intelligence.* McCraw Hill, New York and London

Hayek, F.A. The *Sensory Order.* Routledge & Kegan Paul, Hayek, Earnonn Butler and Smith, London

Hayek, F.A. *Knowledge, Evolution and Society.* Adam Smith Institute, London

Hayek, F.A. *The Road to Serfdom.* Routledge & Kegan Paul

Hebb, D.O. *The Organisation of Behaviour.* Wiley Books

Holland, R.C. *Microcomputers and Their Interfacing.* Pergamon

Hook, S. (editor) (1961) *Dimensions of Mind.* Collier-Macmillan, London

Hume, D. *Treatise on Human Nature.* Pelican

Hunter, M.L. *Memory.* Pelican

Kanal, L.N. *Pattern Recognition.* Thompson Book Co., Washington, D.C.

Katz, D. *Animals and Men.* Pelican

Koestler, A. *The Art of Creation.* Hutchinson

Kuhn, T.S. The *Structure of Scientific Revolutions.* University of Chicago

Lloyd, B. *Perception and Cognition.* Penguin, London

Lorenz, K. On *Aggression.* Alien & Unwin

Lowenstein, O. *The Senses*. Pelican

Luria, A.R. The *Working Brian.* Penguin

MacMillan, R.H. (editor) (1962) *Nonlinear Control Systems Analysis.* Pergamon, Oxford

Macworth, J.E. *Vigilance and Attention.* Pelican

Mair, L. *Primitive Government.* Pelican

Mark, R. *Memory and Nerve Cell Connections.* Clarendon Press, Oxford

Marvin Minsky and Seymour Papert *Perceptrons.* Science Press Inc.

McKenzie, W.J.M. *Biological Ideas in Politics.* Penguin

McLeish,J. *The Science of Behaviour.* Barrie& Rockliff

Michie, D. (editor) (1971) *Machine Intelligence 1.* Edinburgh University Press, edited Collins, N. L.

Michie, D. (editor) *Machine Intelligence II.* Oliver and Boyde, edited Dale, E.

Midgely, M. *Evolution as a Religion.* Methuen, London

Monod, J. *Chance and Necessity.* Collins/Fount

Morris, D. The *Naked Ape.* Corgi

Mottram, V.H. *Physical Basis of Personality.* Pelican

Muses, C.A. (editor) (1962) The -Aspects of the *Theory of Artificial Intelligence.* Plenum Press, New York

Nathan, P. (1969) The *Nervous System.* Pelican, London

Pagels, H.E. The *Cosmic Code.* Pelican

Pask, G. *An Approach to Cybernetics.* Hutchinson

Paul, E.H. *The Physiology of Nerve Cells.* Blackwell Scientific Publications

Penfield, W. and Roberts, L, (1959) *Speech and Brain Mechanisms.* Princeton University Press

Popper, K.R. The *Logic of Scientific Discovery.* Ginsons, London

Popper, K.R. and Eccles, J.C. The *Self and It Brain.* Springer International

Porter, A. *Cybernetics Simplified.* English University Press

Rhodes, F.H. t. *The Evolution of Life,* Pelican

Ross Ashby, W. (1960) *Design for a Brain.* Chapman & Hall, London

Russell, B. *An Inquiry into Meaning and Truth.* Pelican

Sagan, C. *The Dragons of Eden.* Random House, New York

Saparina, Y. *Cybernetics with Us. Peace* Publishers

Sawyer, W.W. *The Search for Pattern.* Pelican

Schade, J.P. *Molecular Neurology.* Roche, Switzerland

Schroedinger, I. *What is Life?* Cambridge University Press

Serebriakoff, V. (1985) *Cybernetics and Factory Organisation.* Paper read to Theoretical Studies Group, British Sociological Association

Serebriakoff, V. *Brain.* Davis-Poynter, London

Sherington, Sir Charles *Man and his Nature.* Cambridge University Press

Slukin, W. *Minds and Machines.* Pelican

Smith, Adam *The Wealth of Nations.* Pelican

Sommerhoff, G. *Analytical Biology.* Oxford University Press

Uhr, L. (editor) *Pattern Recognition.* John Wiley & Sons, New York

Uttley, A.M.(1970) *The Informon Journal of Theoretical Biology,* 27,31–67

Vernon, M.D. *The Psychology of Perception.* Pelican

Waddington, C.H. *The Strategy of the Genes.* Allen & Unwin

Wells, H.G. *The Outline of History.* Cassells

Wiener, N. *Cybernetics.* M.I.T. Press

Wilder, Penfield & Lamar Roberts *Speech and Brain Mechanisms.* Pheneam, New York

Wilkinson, J.F. *Introduction to Microbiology.* Blackwell Scientific Publications

Wright. Q. (1965) *A Study of War.* University of Chicago Press

Woolridge, D.E. (1963) *The Machinery of the Brain.* McGraw-Hill, London

Young, J.F. (1969) *Cybernetics.* Iliffe Press, London

Young, J.F. *Chemical Basis of Life.* Readings from the *Scientific American*

Young, J.F. *Perception* (Mechanism and Models) and *Recent Progress in Perception.* Readings from the *Scientific American*

TABLE OF FIGURES

INDEX

law enforcement, 126
laws, 49
laws of nature, 49
learn, 122, 123
learning, 123
life, 20, 24, 27, 47, 56
life paradigms, 37
life-force, 26, 151, 160
Lorenz, K., 73
lovemaking, 132
Lucas, J.R., 14, 15, 16, 33
lust, 132

M

macrocosm, 14
Malthus, 148
manager, 132
mankind's successor, 143
market, 71, 94
Marshal, P, 148
Marx, 25
Maxwell, 30, 40, 42
Maxwell's Demon, 59, 61, 116
medial cortex, 65
medical, 125
memory, 102, 132, 135
memory track, 137
Menger, C., 148
mental constraint, 39
mesocosm, 14, 17, 159, 161
mesocosmic, 17
mesolithic, 143
meta-learning, 152
meta-morphostatic, 30, 56
metaphyta, 62, 63, 127
metazoa, 32, 34, 57, 63, 87, 128, 136
metazoon, 34, 89, 103
meteorite, 149
meteorology, 125
Michael, C.R., 100

microcosm, 14, 75
microcosmic, 17
military enthusiasm, 73
mind, 45, 47, 160
Minsky, 101
miracle, 23, 24, 25, 26, 27
miracle of life, 34
mitochondria, 103
monocentric patrician altruism, 25
Monod, J., 27, 53
monohierarchic, 71
monotheism, 47
morphostasis, 28, 30, 32, 33, 159
morphostasis continuum, 62, 151
morphostat, 32, 57, 61, 66, 69, 87, 99, 105, 107, 113, 115, 127
morphostatic, 143, 147
Morse, 116
motoneurones, 136
motor, 73
motor information movement, 101
motorium, 15, 86, 88, 92, 99, 101, 114, 123
multi-tasking, 137
mutations, 30

N

Nature of the Physical Universe, 111
negentropy, 44, 62
neo-classical, 148
neocortex, 65, 136
neocortical, 131
neolithic, 143
nerve, 19, 24, 60, 63, 75, 85, 86, 89, 99, 109, 115
nerve cells, 84
nerve endings, 15

neural memory, 133
neuroproteins, 19
neuro-transmitter, 63
New Guinea, 69
Newton, 21, 42, 97
Noah, 144, 149
node input, 122
node signum, 122
node threshold, 121
noise, 83, 105
noise formula, 84
non-carbohydrate life, 162
non-semantic, 82
noosphere, 26, 151, 161
nucleotide pairs, 58, 104
Nyquist, H., 82

O

Occam, W., 47, 49, 78, 159
OED, 38, 81
ontogenic, 29
ontogeny, 87
ontological classes, 112
ontoplastic, 68, 103, 110, 114, 136
ontoplastic intelligence, 111
ontoplastic learning, 128
ontoplastic memory, 133
ontoplasticity, 64
order, 37, 38, 39, 42, 45, 59
order of arrangement, 39
organisation chart, 74
origin, 97
or-ish, 117
orrery, 23
output weight, 122