EVOLVING AN INTEGRAL ECOLOGY OF MIND

Chris Lucas

(CALResCo Group, Middleton, Manchester, UK)

Abstract

A deliberation upon the possibility of generating a comprehensive view of 'mind as a whole' by integrating biology, psychology and sociology, and considering 'Mind' as a dynamical interplay between values existing over many levels and scales of complex systems. This view of mind as a coevolutionary whole is related to similar ecological viewpoints being developed in the fields of artificial life and multi-agent systems. A meta-model of mind is proposed which attempts to integrate the various existing views of mind from a time perspective, and positions the field in a way that we think augurs well for the potential of mind studies based upon modern complex systems science concepts.

Key words: embedded mind, higher-order cognition, attractor dynamics, values, multi-agent systems

Introduction

When we attempt to model aspects of the mind, whether within cognitive psychology, artificial intelligence or neurobiology, we must make simplifications. We model, at any one time, only a small subset of the whole, our 'maps' by no means equate to the 'territory', either individually or in aggregate. Yet it is the whole 'Mind' (the capitalisation signifying, in this paper, the wider territory, the lower case specifying a map) which is our chief interest here, not just in itself, but in conjunction with the world and the body of which it forms an integral part. Brains reflect our biological origins and development, mind functions relate to accumulated experience, and consciousness to our awareness of our social and cultural environments. We are not then a form of disjoint 'thought' looking objectively onto a unchanging mechanical world (as thought by Descartes), but autonomous embodied creatures, whose every behaviour echoes the dynamic interplay of all these three levels of influence – 'we' are integrated nature, nurture and culture, considered in the widest sense of each.

Integrating vertically in this way, in other words considering biology, psychology and sociology as a coevolutionary whole, suggests parallels with ecology, that coevolutionary creation of a system of interacting niches at many levels. Such systems however do not comprise static collections of fixed species but are continually fluctuating mixes of populations, and it is here by analogy that a modern dynamical systems approach to intelligence and life, grounded in evolutionary biology, may assist in forming better models of our individual behaviours. Over recent decades much work has taken place in this area, in the study of complex systems (a designation that includes work on artificial life, neural networks, genetic algorithms, self-organization and other linked concepts; for brief definitions see: http://www.calresco.org/glossary.htm). These systems have at least one aspect in common, and that is the presence of multiple interacting components whose interactions are not fixed but can vary dynamically. In this sense these studies are equally applicable to ecologies, societies, bodies, cells and we would suggest here minds.

In this paper we will first outline some problems with current approaches to mind, which suggests that a new, wider ranging, model may be advantageous. The general approach to a multiple level perspective on mind is outlined, we then continue by grounding the mind in relation to the often ignored areas of the subconscious and our autonomous values. The next section introduces the attractor structures central to a complexity science perspective and following on we introduce an ecological analogy to mind operation. We take a necessary detour then to consider those aspects of mind outside the brain and then extend this to consider higher order forms of cognition. A brief outline of a proposed meta-model is then given and comparisons with other proposed models are made. Some possible benefits of our model are then listed, followed by our conclusions regarding the state of the art in such perspectives.

Difficulties with Existing Models

The most striking feature apparent upon consideration of existing mental models is their lack of integration. Each discipline considers an aspect of mind appropriate to their basic methodology, yet these are largely divorced from the theories and conceptual bases of many other related disciplines. Thus, for example, we have approaches from neurobiology which concentrate upon the lower cell

level operations (microscales – neurotransmitters, synaptic structures, etc.), those of neuropsychology (considering mesoscale neuronal groups and brain modularity), of psychology itself (largely macroscale behaviour related) and of social psychology (interpersonally based). Additionally, and often unrelated to 'real world' animal structures (studied in neuroethology), we have approaches to Artificial Intelligence (AI) based upon computer techniques of various kinds, plus the more intellectually abstract approaches taken by philosophers of consciousness. All these perspectives, in theory, relate to the same subject – the operation of Mind, and must conceptually at least be consistent with each other. Such consistency is however often more evident by its absence. Additionally, within each individual discipline, such as AI, strong disagreements about appropriate models and approaches are very evident, even upon a casual reading of the literature. Overall, the disjoint nature of our specialised academic disciplines, each with their own terminology and often hostility to input from other disciplines, leads to a situation in which communication about and comparison between wider ranging models is inhibited rather than encouraged.

By viewing mind from such single perspectives we also generate situations whereby different groups talk past each other, arguing about whose viewpoint is 'correct' whilst ignoring the fact that they are both considering different aspects of a single complex system. Such systems are often held by complexity researchers to be incompressible to any one formal model, due to the presence of multiple interdependent structures – the minimum model of the whole would mathematically then be the whole itself. If this proves to be the case for mind, then we need at least a framework or metaposition that can correlate and compare all the various approaches being taken in relation to studies of Mind, and which can enable us to see the contextual limitations, strengths and possibilities of all the diverse viewpoints around today. From this position we must recognise that none of our models can be complete, and view them instead as simply alternative slices through the whole, looking to understand how collections of such mappings relate to each other and to our objective interests. It is in this area that the interdisciplinary perspective offered by the complex system sciences may assist, not least by supplying a common terminology that seems applicable across all disciplines. Whilst often very abstract (as is necessary perhaps if it is to be widely applicable and not subject specific) this approach has proved useful conceptually in many areas in recent decades, and is the approach upon which this paper is based.

When we look at the functions of mind, at how for example humans actually behave, we find that invariably we have many different needs and interests (see Subconscious Reasoning), along with the ability to choose between them. Yet with rare exceptions academic work ignores this selective multidimensionality, preferring to concentrate at any one time on isolated functions of the whole. Whilst this is understandable, given human limitations, it nethertheless presupposes that all these dimensions of mind are independent (e.g., visual perception doesn't affect hearing; or hunger doesn't affect thinking), an assumption which must nowadays be called into question. Within complex systems, and the mind should be regarded as such, influences between subsystems are common, indeed it is this epistatic aspect that makes the study of such systems a difficult task since the nonlinearity generated by such interactions falsifies the superposition principle central to linear reductionist techniques. Any adequate overall view of mind must therefore take these interdependencies into account, or show explicitly why they are unimportant in any particular context. It should ideally also model the effects of choice itself, in our ability to change our coevolutionary fitness by altering our individual or collective evolutionary trajectories – decisions which may prove positive (synergic) or negative (dysergic) overall, which we can relate to 'true' or 'false' beliefs.

A further aspect concerns the lower and higher level influences on mind, in other words to the biological and cultural constraints upon brain operation. Too often these aspects are completely ignored, leading to viewpoints in which mind operates without body and without context. Yet we have no evidence that mind can do either, such beliefs are philosophical assumptions only, abstractions divorced from scientific reality (in which we find all known minds to be both biologically based and environmentally situated – even AI 'brains' are created by experienced biological lifeforms and have, albeit limited, environments). Additionally, recent evidence has pointed towards the influence of emotions (affects) upon intellectual function, an aspect of mind traditionally excluded from 'serious' academic studies of the subject, which have concentrated largely upon isolated 'rational' behaviours. It seems a rather odd approach however to try to explain mind by artificially ignoring these various known influences on its function, and any such models must be by implication incomplete – forms of simplification understandable perhaps from an historical reductionist perspective, but nethertheless potentially misleading.

Another problem common to most current models of mind is their difficulty in learning new/changing situations or concepts (and in forgetting them). This manifests especially in symbolic AI where the static knowledge bases of current models prove to be extremely fragile even when restricted to highly artificial domains

(the frame problem), but it is also a major difficulty for the connectionist viewpoint where the need for many 'training' trials (typically thousands), and the subsequent cessation of learning capability in the implementation phase, make such (backpropagation) models biologically implausible. Other more recent models, e.g. adaptive resonance theory (ART), can overcome this stability-plasticity dilemma, but still need multiple epochs of training. In the world of our real minds, we very often learn to generalise from single examples (e.g., the Concorde jet), and can completely rewrite our understanding as a result of a single new (e.g., 'near death') experience. Teaching by repetition or imitation (whilst popular in early 'rote' education and for motor skills) is neither efficient nor adequate as a model of concept learning on dynamic landscapes (especially, as Chomsky noted, for the acquisition of language).

Our final problem concerns the lack of attention paid to ontological development, in other words to how those mind functions that we choose to model (e.g., representations) actually came into being. This necessitates a strong evolutionary focus, a view of mind as constantly changing in function – a 'nonstationary' or constructivist perspective (e.g., Quartz and Sejnowski, 1997) – rather than as having some static functionality which can be modelled in isolation – a mode which often leads to tendencies to label people or groups as having fixed 'traits'. The plasticity of neural functionality is well known scientifically, so we have little reason to assume that any higher level functions are unable to evolve, or that we are unable to add new mental functions at a later date. Indeed every new concept that we learn must in some sense involve a change in mind structure (i.e., our cognitive machinery itself is shaped by our environmental interactions), so assumptions of a static structure must prove to be invalid (other than as a simplification). Treating the mind instead as 'open ended', in conjunction with its environment, would seem then more scientifically productive than trying to analyse a imaginary closed system possessing fixed abilities. Just as a tree remains a tree despite continually growing, so we wish to understand how our concept of a tree remains a stable concept of a tree despite the continual growth of our experience of such trees (and their subsequent subdivisions into oak, eucalyptus, deciduous, conifer etc.). In fact the similarities between these two fields do suggest a similar process is at work (e.g., Grobstein, 1988).

Multi-level Dynamics

Let us make a start formulating a more integrated model by first approaching the subject from the opposite direction, by considering

TABLE I

Seven Level Emergent Evolutionary Structure

Ecological (multiple species) → niches
Sociological (multiple cultures) → organizations
Psychological (multiple minds) → specialists
Neurological (multiple modules) → concepts/maps
Developmental (multiple cells) → organs
Biological (multiple genes) → cells
Chemical (multiple molecules) → enzymes

evolutionary theory and the animal kingdom as a way of positioning our psychological model and then using this biological perspective to illustrate recent findings from the complex systems sciences. When we look at any whole organism we can identify three well known timescales of study: firstly the phylogenetic, the variation of the various ecological forms over many lifetimes (modelled by evolutionary biology), secondly the ontogenetic growth of individual organisms over a single lifetime (modelled by developmental biology) and thirdly the epigenetic development of brain with daily experience (modelled by neurobiology). Given the parsimony of nature, it may be thought reasonable to assume that these three very different organic processes are nethertheless based upon the same universal principles, and this is a viewpoint also taken by the systemic sciences, which look for those common principles that apply across all disciplines. Given this possibility, we can perhaps take this analogy a stage or two further and suggest that chemistry, psychology and sociology are also examples of the very same processes, and note that these various areas are all closely interlinked, leading to a proposed 7 level self-similar (fractal) structuring of our living world (Table I), each level of complexity described is presumed to have emerged from the previous level over time, not only by forms of natural selection but also by the processes of self-organization (see Kauffman, 1993).

Many thousands of simulations, investigations and experiments have been carried out by complexity researchers around the world over recent years, covering all these levels of reality, and we now have good agreement as to the typical dynamic features to be expected from any complex system. Given freedom for the parts to interact (in contrast to the fixed interactions of designed human artefacts) the dynamics of the system show forms of self-organization, in other words the structures that appear are not imposed from the outside (natural selection simply chooses between them) but arise as a result of the changing interaction regime of the varying components, emergent forms that cannot be predicted from even full deterministic knowledge of the lower level part behaviours. It should be noted here that this feature of interacting parts does not imply any vitalist influences, and simply relates to the difficulties (possibly soluble) of understanding emergence itself

at *any* level, i.e. how new features (functionality) can appear given constrained part interactions.

In all these forms of organized systems we experience, it seems, a fitness maximising balance between relatively static structural frameworks (the physical 'parts' aspect) and relatively chaotic dynamic interactions (the informational 'processes' aspect), a criticality called in complexity science 'edge-of-chaos', which has similarities with the phase transitions of physics, but in which, surprisingly perhaps, the static and dynamic parts disintegrate and re-integrate into different patterns, swapping places over evolutionary time (e.g., when body structures are dynamically recreated based upon molecular dynamics). At this position complex systems are found to contain a broad mix of attractors (see Attractor Dynamics) of various sizes, giving any individual system many possible alternative stable or meta-stable states. Transitions (bifurcations) between these alternative states are often abrupt and invoke different component or agent groupings (e.g., in biology, cell differentiation). The dynamics of such far-from-equilibrium states under perturbations is complex, typically transients follow a power law distribution (e.g., Bak, 1996), structures occur in space and time at all scales (fractal) and show spontaneous bursts of evolutionary change (e.g., in palaeontology, punctuated equilibria). The groupings or niches that result, given diverse agents (i.e., complex adaptive systems), often involve combinations of generalists (whose behaviour is applicable in most scenarios – implying wide basins of attraction) and specialists (more efficient, but limited in applicability, implying narrow basins), see Holland (1992), and this can be seen in natural world food chains and perhaps in human job specifications. These various features comprise a universality seemingly common to all types of coupled, nonlinear, complex systems.

Subconscious Reasoning

None of these emergent 'intelligent' functions however require the presence of consciousness (animals, and even inorganic systems, exhibit the same dynamics to lesser extents), and we can learn much about the mind if we discard the conventional emphases upon philosophical 'consciousness' (an holistic intellectual mode) or physiological 'perception' (a reductionist animal mode) and look more deeply at the intermediate (subsymbolic) nonconscious or subconscious forms of cognition, e.g. our ability to drive on 'autopilot' whilst thinking consciously of other things – in a sense it is true to say that the more we 'know' the less conscious it becomes, since we adapt to respond automatically to the relevant situations (e.g., Epstein, 2000). Here we will concentrate on such hidden and emergent decision making abilities, viewing mind and intelligence instead as

"*optimizing behaviour in the light of information*", a mode which features even in primitive lifeforms such as bacteria (e.g., see chapter 3 of Cairns-Smith, 1996), and which is prevalent in artificial life studies and our agent models. We will assume that an implicit strategy or 'belief' drives behaviours and this results in evolutionary adaptation of the mind by trial and error (heuristics), such that fitter (more successful) strategies come to prevail. Making decisions however requires both internal values and a context and this brings a teleological (intentionality or self-directedness) perspective to bear (see Lucas, 2000; Polanyi, 1958), implying that facts (distinctions) provide alternative trajectories (choices) through possibility space, which will then affect subsequent needs and resources in different ways, this corresponds to a more directed form of evolution than the 'random mutation' mode assumed in evolutionary biology, i.e. a goal driven exploration of possibilities. Many definitions of overall value or 'fitness' are possible in complex systems, but here we will define fitness in a rather wider sense than that common in evolutionary biology, as "*those global optima available from multiple interacting human needs, evaluated within a full social and environmental dynamical context*", sometimes referred to more informally as maximising 'quality-of-life' and connected to Maslow's (1968) 5 level "hierarchy of needs" (which we will simplify somewhat, relating his model more to evolutionary ideas).

For our human focus here, upon aspects of Mind, we can usefully divide our values into 3 emergent levels of needs: primal, interpersonal, and abstract, roughly corresponding to life/brain (nature – biological), mind (nurture – psychological) and consciousness (culture – sociological). Each level containing a number of separate values which can be roughly placed into three main evolutionary groupings (these groupings are meant to be indicative here rather than definitive).

Primal or Basic Needs

These are those related to lower animal or plant behaviour, essentially concerned with the physical world and our physical existence. They include eating, drinking, respiration, growth, reflexes, shelter/warmth, reproduction, security/survival, sleep, waste disposal and health. They relate to the emergent properties of life itself and correspond to the lower two levels in Maslow's "hierarchy of needs", which he called 'Physiological Needs' and 'Safety Needs'. Many of these needs are completely nonconscious and include all the homeostatic body functions.

Interpersonal or Social Needs

Moving up to a higher level we have the more sophisticated needs associated with the emergence

of mind and community in the middle and higher animals. These add to the list such needs as communication, display, status, belonging, curiosity, stimulation, mobility, work, play, comfort and forward planning (resource stockpiling and simple goals). These correspond to the middle two levels in Maslow's "hierarchy of needs", which he called 'Love, Affection and Belongingness Needs' and 'Esteem Needs'. Our behaviours here are often subconscious, with automatic responses to familiar situations – e.g., habits.

Abstract or Spiritual Needs

At the top of the needs hierarchy are those higher needs attributed to fully developed humans, and these are usually claimed to be applicable only to our species. They encompass, amongst others, art, music, science, mathematics, religion, love, philosophy, justice, ethics, history, beauty, compassion, friendship, creativity, education, enlightenment, and freedom. These levels show the emergence of concepts of a non-material form (no reference to deities is implied by the term 'spiritual' here), the higher levels of consciousness often explicitly excluded from science although they correspond to the same mode of thought as science itself. This level corresponds to the highest level 5 in Maslow's "hierarchy of needs", that of 'Self-Actualization Needs'. Meeting these needs usually requires deliberate and conscious cultivation.

The interdependencies between these needs means that we must also address the relationships between these levels, generating a multi-level scientific methodology where all the levels are regarded as open systems, which take into account the personal, social and environmental effects of our actions as an holistic hypersystem, rather than viewing life in isolated chunks (e.g., sleep, work and lovemaking needs interfere with each other). This implies that multiple goals are simultaneously active, at least subconsciously, a situation more easily modelled by a system of many single-value interrelating agents perhaps than by any single autonomous entity.

Balancing sets of such interacting values and deciding priorities, i.e. which need(s) to pursue at any time, requires a more intuitive integrating perspective, and here recent research suggests that emotions are involved (e.g., Damasio, 1999; Cosmides and Tooby, 2000) an insight now being incorporated also into agent modelling. Whatever the mechanism, in complex systems multidimensional rather than single dimensional approaches are appropriate since we cannot neglect the epistatic (nonlinear) interactions between the 'parts'. This is dependent upon their connectivity which in complex systems can be very large, for example in the brain there are estimated to be about 10,000 synaptic links to 1 neuron on average, with many synaptic (potentially active)

circuits evident (see Shepherd, 1998), although we would predict that few dendrites are actually active at any one time – presumed necessary if concepts are to remain stable. Rather than the single cause, single effect chains of linear treatments we now need instead some form of nonlinear, multiple cause/effect, networks.

These nonlinear networks allow synergy or emergence, the creation of new levels of function from the combinations of interacting parts. This natural modularity, observed throughout the natural world (see Maynard et al., 1995), suggests that a similar model may be applicable for mental functionality. Complex systems methodology takes an autopoietic (self-producing) approach to this, in which the organizations that result and the actions which then take place netherteless depend historically upon coevolutionary interactions (perturbations) with both the environment and with other minds. This necessary 'structural coupling' (see Maturana et al., 1995) precludes treating mind as a 'stand-alone' problem and frames it instead in terms of a dynamically coupled set of cybernetic subsystems, existing at many levels and incorporating multiple nonlinear circular causalities. Indeed a true view of mind, as such a self-producing (homeostatic) system, must include all those arcs external to brain that also form part of feedback loops which can affect the behaviour of the integrated 'Mind'.

ATTRACTOR DYNAMICS

The dynamical systems approach common to many modern complex systems studies considers systems to possess what are called attractors, each of which is a restricted area of possibility (state or phase) space to which the system moves dynamically but then finds it hard to escape (analogous to a gravitational well or a chemical energy barrier). Three types of individual attractor are common: firstly the point attractor in which the system eventually has only one possible state (typified by forms of science in which only single equilibrium solutions are possible); secondly we have the cyclic attractor where the system has a number of solutions which are visited in sequence in a regular order (satellite orbits are of this type), both these types relate to low connectivity networks; and thirdly there is the recently discovered 'strange' attractor where complex and often fractal patterns of behaviour are found. This latter type is generally unpredictable (although deterministic) and relates to 'chaos' which often can be shown to occur if networks are highly connected. The area of possibility space leading to an attractor is called its 'basin of attraction' and in complex systems this can comprise a vast number of individual states. Nonlinear systems will in general have many possible attractors present

(multi-stability), so the attractor in which the system is expected to be found will be indeterminate without some historical knowledge of in which basin of attraction the system starts. This knowledge relates to the control parameters of the dynamical system, parameters whose values determine the phase portrait (the attractor structure within state space) of the overall system and these prove critical in evaluating the possibilities and actualities that we can expect from such process oriented systems [As an aside, from a process metaphysics (as opposed to a substance metaphysics) viewpoint (e.g., Bickhard, 2000), matter itself can be regarded as simply a particularly stable form of emergent attractor, so this dissolves the body-mind problem, converting it instead into consideration of the relative stability of attractors, which seems unproblematic. This idea, compatible with a deeper understanding of quantum theory (e.g., stochastic electrodynamics, and related to standing waves and quantum jumps), seems to have gone largely unnoticed within scientific and philosophical mind circles, i.e. materialism itself dissolves if there is no 'material'].

Given the complex network structure of the brain, it is expected from such a complex systems perspective that there will be many natural attractors present (e.g., our concepts). Since such interconnected (recurrent) networks can support multiple different simultaneous attractors without any structural changes (node or connectivity alterations), then a simple change of input to one component (a control parameter) can be sufficient to flip the entire system to a very different attractor (e.g., the duck-rabbit ambiguity), i.e. minor intentional changes can be expected to drive major behavioural alterations using informational rather than force based switching. We can regard the various influences on our behaviours (either internal needs or external constraints/opportunities) as such control parameters, moving our current state into the basins of attraction of sets of attractors. The system dynamics will then cause the network to quickly relax into a semi-stable equilibrium which we can relate to the combination of concepts or ideas that arise in our mind. Many of these attractors will be short-lived, additional dynamically created forms of attractor, termed 'transient' in Lucas (1997). From a mind perspective this means that not only the structures of our mind (the neurons and their synaptic circuits) but the dynamical processes within those structures can alter under multiple influences. A further consideration concerns cognitive thought directing our actions, this is equivalent to a 'downward causation' or group selection effect, in other words the lower level behaviour of the parts is causally affected by emergent higher level properties of the whole (e.g., Campbell, 1974). Related to this, some recent work in the complex systems specialism of Evolutionary Multiobjective

Optimization (EMO) has concentrated upon the evolution of niches or specialisms, especially the ability of selection, in cooperative coevolution, to optimize multiple resources by generating modularity (a parallel mode) as a more efficient solution to complex problems of interacting values than single solution methods (a serial mode), e.g. Potter and De Jong (2000). Such studies establish the value of small combinatory groupings and diversity in fluid scenarios, whilst at the same time showing that emergent properties of the whole do affect the evolution of the component functions. This work has applicability to the evolution of robot behaviours and by analogy to human and animal equivalents.

Whilst a pure dynamical systems (differential equation) approach to mind cannot directly replace symbolic approaches (e.g., Eliasmith, 1996), we can employ dynamical ideas equally well to discontinuous processes and symbolic systems (e.g., Jaeger, 1999), and here the idea of connectivity producing attractors in discrete dynamical networks proves a powerful one at all levels of reality, including mind [e.g., the attractor-based subdivided neural networks approach outlined in chapters 2 and 3 of Bar-Yam, 1997 provides one feasible model of many brain processes, including psychological malfunctions, necessity for dreaming, creativity and individuality]. Studies of the olfactory system have demonstrated also the presence of chaos and strange attractors (e.g., Freeman, 1991), and these appear additionally in other areas of brain research (e.g., Abraham, 1992). As well as the self-organizing, intentional dynamics aspects of nervous system attractors, studied by Kelso (1997), complex systems theories make much use of natural selection from competing alternatives and this too may have a role to play in theories of the mind (e.g., Edelman, 1992; Goertzel, 1993; Calvin, 1996). Recent approaches to AI that use situated non-symbolic methods (e.g., Brooks, 1990) also show a convergence with the bottom-up methods currently being used in artificial life (AL) and multi-agent systems (MAS) to synthesise lifelike and intelligent behaviours [a view of the immune system as intimately connected to the brain and body in a coevolutionary manner is gaining currency, we can model this effectively as a swarm of agents interacting by signs (e.g., Hoffmeyer, 1994), and perhaps extend this technique to the mind also]. All in all, the signs are positive that such research directions are highly relevant to current mind studies, to which we are now in a position to add our more ecological focus.

ECOLOGICAL MIND

Mind, considered as an ecology (e.g., Bateson, 2000), would consist of a number of levels, an

hierarchy of specialist and generalist niches. At lower levels brain 'organisms' would deal with simple resources (in visual perception these may be 'edges' or 'movement'), different neural components would then 'compete' for survival in successfully utilising the resource (failure would then lead to synaptic 'death' – loss of access to the resource). At progressively higher levels, groupings (e.g., cortical columns) would 'eat' combinations of lower 'organism' resources (patterns, i.e. correlations due to associations) generating categories or classification. These would combine to generate abstract concepts at still higher levels, supervenient upon (but not reducible to) the lower levels. Ultimately consciousness, viewed now not as disjoint to subconscious behaviour but as an extension of it, would be an animal 'browsing' over the whole landscape, attracted by disturbances (incoherences or cognitive dissonances between parts of the brain) within the normal subconscious flow. It is not however being suggested here that populations of neurons 'move' around the brain as animals do, simply that the synaptic circuit activation patterns behave dynamically in an analogous way to ecological encounters in a natural landscape, allowing us to use the same methodologies as those developed within related biological and complexity disciplines (and vice-versa). This builds upon an idea, common in complexity science, that sees many alternative formalisms as isomorphic to each other, some of which are netherthless easier to use in individual application contexts. The idea of 'organisms' at all levels of the brain lends a commonality and parsimony to the whole, and suggests that if we can identify such a structure at one level then we have a potential explanation for all mind levels, i.e. different organism configurations implement different functions (as happens in biology and ecology).

From our perspective here we can relate these brain 'organisms' to our values or needs, which we saw earlier could be placed into three emergent levels. We can now add a number of other levels however, corresponding to the sub-actions needed to ensure the 'survival' of any individual value, e.g. to 'drink' we may need a subsystem to detect water, move towards it etc. – reminiscent of Brooks (1990) 'subsumption architecture'. We can thus establish webs (or 'food chains') of hierarchical dependencies between organisms, each one envisaged as an agent in an artificial life simulation, and having an internal structure appropriate to its role in the functional psychological 'ecosystem'. Evolving this system then would be equivalent to seeing how the mind dynamically met and balanced its needs, establishing over time a balance that can appropriately respond to the environmental opportunities. A further level of modelling however would ultimately be necessary, since each mind is embedded in a landscape comprising other minds and resources, thus we need eventually to include the social or cultural level also, again by using the same interacting agent format.

Natural landscapes vary however in openness and closeness, a closed ecology (e.g., an isolated island) has a relatively fixed set of relationships and species, which we can call autopoietic (energetically open, operationally closed). In contrast, more open landscapes (e.g., part of a rainforest) allow inflow and outflow of species, a fluctuating mix of structural components which we can call sympoietic (see Dempster, 2000). In analogy with the mind, certain individuals (and societies) have closed worldviews and static, dogmatic, behaviours. Other, more open, people are receptive to new ideas and freely replace older viewpoints. These positions relate, in complex systems terms, to being slightly below the edge of chaos (static tendencies) or slightly above (chaotic tendencies) and it is in the dynamic of perturbations across this edge of chaos boundary that we can often identify the evolutionary trajectories (sequences of attractors) that we see around us. Thus we can relate overall human behaviours quite closely with those of known ecological systems, although the detailed modelling of such brain 'ecologies' must be seen as very much of a future aim, given our limited understanding so far of the structural details and dynamics of mind operations.

Extended Mind

The influences driving openness and closeness in the mind however are many and take effect over multiple levels. Our genetically directed neural components provide certain constraints over what can occur structurally, and the allowed processes are also affected by hormonal levels (perhaps driving our 'edge of chaos' boundary position, e.g. sleep/arousal levels). Our available higher level categories also must restrict what we can imagine and these are affected by experience and the cultural norms and variety that we encounter. Thus the triggers or signs that we can respond to (i.e., the distinctions we make) and the behaviours (trajectories) that can then result are always limited (but obviously expandable by learning) and closely tied in with our current value system. This implies that our psychology exists in a transdisciplinary balance between biology (e.g., Gazzaniga, 1992) and sociology, and needs to be understood in that wider context of a dynamically fluctuating mix of intrinsic constraint and possibility (see Mitchell, 1999). In this perspective our effective mind is extended far beyond the brain and exists partly in the actions of our body and partly within society itself. Thus thinking is an embedded activity using and being constrained by much information

external to the brain itself (some of it deriving from other active agents), an open process that has no obvious boundaries and which links constantly, either directly or indirectly, to all other levels of our reality.

Given the large number of biological and social constraints upon our behaviour, we can speculate that these affect the attractor structures at various levels of mind, generating over time relatively stable ways of behaving, worldviews ('structures of consciousness') that operate largely subconsciously, and restrict (canalize) our conscious options in a way that is often unseen and possibly detrimental (i.e., better options or ways of behaving are screened off by the canalization of our search space). An in-depth look at the importance of internal variation, constraint and selection with respect to the evolution of complex systems can be found in Bickhard and Campbell (2003). The idea that multiple alternative modes of thought can evolve across either location or time [thus humans are contingently heterogeneous (e.g., Edelman, 1992), rather than homogeneous as are computers], and that we can flip from one mode to another, can help explain some aspects of our contextual behaviour (e.g., Combs, 1995), for example in the very different sets of cultural beliefs or overall attractors adopted by different societies, and the different 'persona' that we all adopt in different contexts. An evolutionary perspective on such a flow of ideas is found in the concept of 'memes', a field now called memetics. An interesting expansion of this 'ecology of ideas' to social areas and multilevel worldviews in general was pursued by Beck and Cowan (1996), and in this sort of dynamic we see a merging of the bottom up artificial life approaches with the more top down social science research – biology and sociology are to some extent converging upon psychology, within a complex systems perspective.

HIGHER-ORDER COGNITION

In artificial complex system studies we generally discard the idea that there is a qualitative (dualist) difference between non-life and life, between life and intelligence, and between intelligence and consciousness (taking the view that the 'hard problem' is no harder than explaining any other form of emergence, e.g. life itself), and this projected continuum of evolutionary emergence (i.e., a large number of lesser structural steps) applies also to the contrast between 'objective' versus 'subjective' perspectives – the 'intersubjectivity' viewpoint (e.g., Velmans, 1999). Accepting that the mind ultimately does not distinguish between inside and outside (i.e., communication pathways are source irrelevant), e.g. von Foerster (1973), allows us to position instead consciousness as intelligence looking at its own contextual behaviour, i.e. self-referentiality. We sometimes confuse however consciousness with languaging, assuming that these are equivalent, but talking about life (i.e., languaging) is not the same as living it, they are different logical types in Russell's terminology and our consciousness of our languaging adds a further logical level to the mix. Cognition in itself is action selection (using environmental awareness to direct the organism's behaviour), e.g. Franklin (1995), language is then an higher order 'cognition about cognition', a new metalevel, not directly space and time related – a more abstract level. This perspective echoes that in 2nd order cybernetics, where we include the previously detached observer as part of our system, considering how the observer perspective (and the specific language he or she uses) affects what is observed. But we must go further and realise that this higher order viewpoint is grounded in our broader experience of 1st order cognition and directs this in turn, thus (as with modern 3rd order cybernetics) we need a 3rd order cognition that discusses how the previous two coevolve and affect each other – akin to a developmental perspective on mind continued throughout life, i.e. a more timeless cognition about the changes in cognition with cognition. To try to clarify somewhat, first level cognition relates to skills (procedural or machine knowledge) and to the direct dynamical systems views considered previously, second level relates more to connectionist ideas and embodied concepts (valuational or intuitive knowledge), whilst third order relates to symbols (declarative or strategic knowledge) and to the traditional AI approach, and we should not confuse these three philosophically different levels, they are not competing so much as complementary perspectives.

A crucial contrast between them is that first order operates in the 'now', in real space and time, whereas third order adopts a virtual space and time, and thus can reason about any location in the past or future, the not-where and not-when. Thus in abstract thought the tight coupling between organism and environment has been broken, and this new emergent, extended consciousness (allowing 'offline' processing) must be regarded as a major evolutionary step. It is the disjointedness of 3rd order from physical space and time which allows us to generate a fractal set of multi-layered academic meanings or isolated theories occupying various space and time frameworks and levels of reality. Bridging the gap between these two levels (i.e., incorporating second order cognition) needs an approach beyond that of simple first order situated 'real-time' naturalism, but this must also take into account the detached aspect of third order cognition. Such a position needs to recognise that many aspects of this timeless virtual reality exist outside our brains, in the artefacts of our technology (tools, books, etc.) and in our immaterial social ideas (e.g., Clark, 2001). Thus the 'virtual world' overlaps and
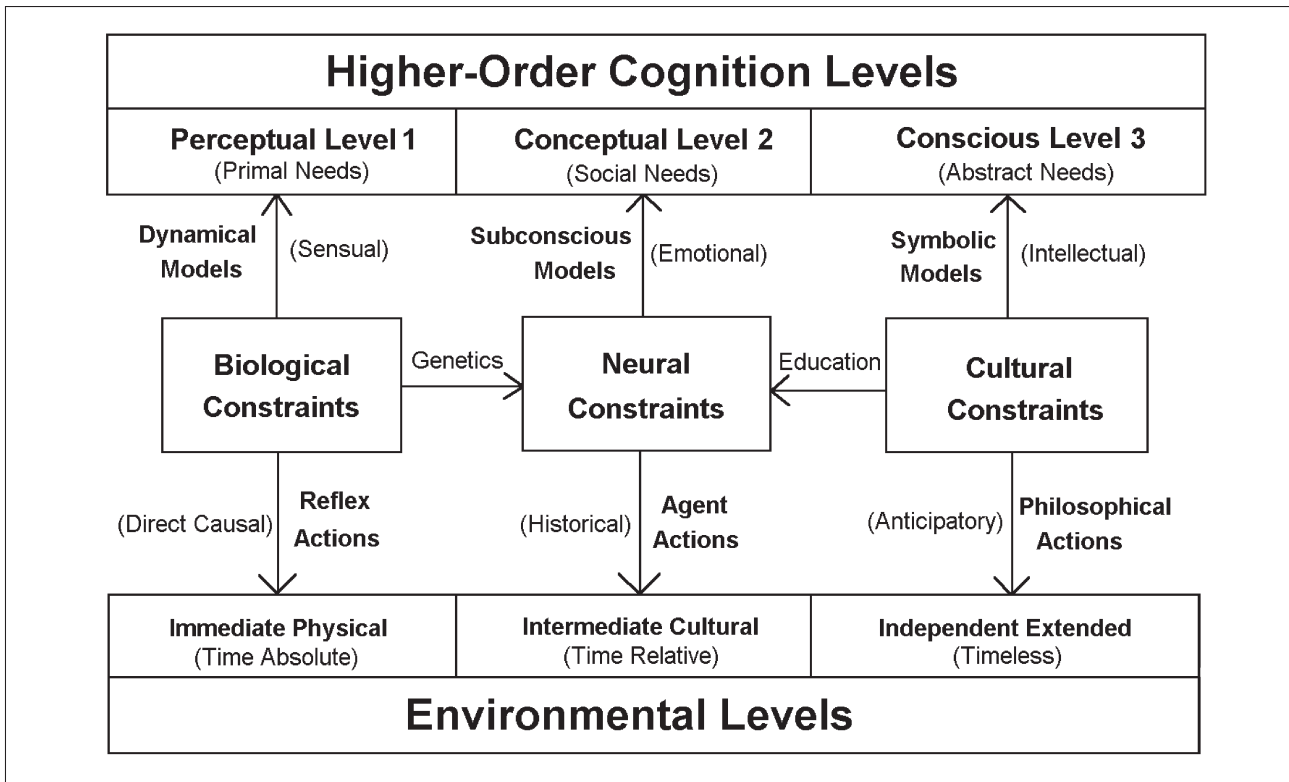
Higher-Order Cognition Levels

Perceptual Level 1 (Primal Needs)	Conceptual Level 2 (Social Needs)	Conscious Level 3 (Abstract Needs)

Dynamical Models   (Sensual)   Subconscious Models   (Emotional)   Symbolic Models   (Intellectual)

Biological Constraints	Neural Constraints	Cultural Constraints

Genetics   Education

(Direct Causal)   Reflex Actions   (Historical)   Agent Actions   (Anticipatory)   Philosophical Actions

Immediate Physical (Time Absolute)	Intermediate Cultural (Time Relative)	Independent Extended (Timeless)

Environmental Levels

Fig. 1 – *A meta-model of embedded mind.*

becomes conflated with the 'real world' (in the merging of 1st and 3rd order cognition in a 2nd order autonomy) with much confusion, especially between 'map' and 'territory'. This is especially the case due to the multiple roles taken by language, which functions both as a prior embodied set of 'subjective' action imperatives (2nd order) which act as (semiotic) informational triggers (included here are indicative and interrogative modes) and also as an abstract 'objective' discussion medium (3rd order) – a possible later exaptation in Gould's terms. In our level 2 perspective we can roughly relate the three semiotic realms of syntactics, semantics, and pragmatics, to the way attractors interact (action possibilities), their basins of attraction (recognition possibilities), and the way they coevolve with the environment to meet our values (fitness possibilities) respectively. This pragmatic coevolution however will occur simultaneously in the interaction of all three of these cognitive dimensions, both with the environment and with each other.

To adequately model this rather complex whole we need recourse to yet a higher level, a fourth order cognition in which we can position the behaviours and interactions of the previous three levels in such a way as to allow us to objectively reason as to how the distinctions we make (first order perceptions – facts) are evaluated (second order judgements – choices) in the context of a whole (third order abstractions – worldviews) within a multi-system environment, a metaview that logically goes somewhat beyond existing

perspectives and may necessitate a new holarchic and paraconsistent form of logic (see Lucas, 2002). The following model makes a conceptual start on this task, in the context of embedded mind (here 'embedded' implies an active environment, including other agents, such that mind behaviours must coevolve dynamically over many timescales and be based upon all levels of human values).

A COGNITIVE META-MODEL PROPOSAL

In an attempt to bring together all these diverse themes we offer a meta-model (Figure 1). It is intended to help position our various approaches and indicate their inter-relationships, without trying to be too comprehensive or dogmatic. The fractal nature of complex systems means that whatever form of analysis we adopt should be equally applicable to any level (if targeted at those details appropriate to that level and those of the adjacent interacting levels) – hence the success of these forms of study in interdisciplinary areas. Our division into 3 forms of cognition within a 3 form environment is thus an illustrative 'coarse-grained' treatment, in practice we would suspect that many overlapping forms of emergence would be necessary (e.g., to model the various levels of generalisation of our concepts). The whole will thus form an effective continuum of complexity, visualizable in many different and complementary ways, and with potentially much mutual influence between the cognition levels (e.g., satisfying primal

hunger may involve social communication and religious taboo aspects also). It should be noted however that the levels operate in parallel, with their own sensory/action pathways, so this does not necessitate the strict hierarchical (pipeline) control structures common to the serial processing (sequential tasking) mode used in many computational models.

In this meta-model, which is based upon a 4th order cognitive perspective, we centralise the neural constraints as our fundamental level, since these must prove to be crucial to the overall mind picture. These constraints comprise in essence the self-organizing dynamics of our neurons themselves and their interconnections. The biological constraints affecting this network comprise those genetically derived chemical and structural components necessary to create and maintain the brain (including hormonal and neurotransmitter influences) along with any physical influences affecting the body's biochemistry and brain operation (e.g., drugs). The cultural constraints on neural operation include our experiences and learning as embedded creatures within a particular culture, along with the ingrained historical social norms (e.g., laws, customs) that affect our behaviours and freedoms (enculturation). The attractors forming at neural level must then self-organize under many parametric influences derived from levels below, above and within our neural architecture. Because of this complexity we must accept that our models can never be assumed to be exact, since we must select out only a subset of factors to consider, i.e. those we take to be particularly causally relevant. As a result we can expect at best only probabilistic predictions within certain restricted contexts (in contrast to our intellectualised abstract models where the domain of discourse is artificially restricted in such a way as to make exact mathematical 'solutions' possible).

Three forms of interaction are indicated in this model, each related to different timescales. In the first, from immediate physical stimuli via direct causal (force based) chains of sensual perception to reflex actions, we emphasise the 'real-time' biological basis of behaviour (common to all animals) which relates mostly to those set of needs we called primal. This is perhaps best modelled by standard (absolute time) dynamical systems theory, although we must not over-simplify this due to the number of simultaneous and interacting needs involved. This interaction type is constrained biologically both by the actual communication channels available (i.e., our sensory modes) and by the physical limitation of our available responses and would normally be regarded as nonconscious or perhaps preconscious. We assume that only one response (attractor) is available to any set of stimuli (e.g., the 'knee jerk' reaction), although different responses could evolve over time (conditioning).

The second form of interaction, from intermediate cultural via concepts to autonomous agent actions, emphasises the aspect of choice amongst multiple attractors that we expect from a complex network such as our neural net. This relates to our social needs and brings in the integrating and directional functions of emotions, which are heavily influenced by social situations and grounded in the historical nature of social norms and evolutionary ability. These behaviours are mostly subconscious in nature, we react generally to social situations without analysis (but often using a great deal of environmental feedback, informational triggers, to co-ordinate and reinforce our responses), but nethertheless we can choose to react or not as the case may be (although generally we must make some choice, e.g. stopping at a red light or jumping it). At this level the same stimulus can generate different responses at different times, in different people, and in different historical contexts. Time here is relative, history is compressed in memories, and needs may be 'postponed' and re-ordered, adopting different (fractal) schedules. This 'elastic time' is nethertheless grounded, i.e. needs will escalate in priority if postponed or ignored (e.g., thirst or deadlines). It is expected that, at this level, the coevolutionary techniques of complex systems modelling will be most applicable, taking into account the intrinsic neural constraints (control parameters) derived from the multiple levels mentioned earlier, which will act to restrict possibility space and evolutionary variation.

The third interaction type, from independent extended environmental influences via conscious thought to philosophically (or academically) derived actions, emphasises the optional nature of these behaviours (i.e., we have no direct biological or social need to pursue such issues, e.g. solving a crossword puzzle, we can put it on 'hold' indefinitely). Here our abstract needs can be discussed in a timeless (but not 'discrete' time as in computational processes, which have been grounded and thus are non-optional) intellectual way via our symbolic languages, although cultural constraints will affect to some extent what is permissible to discuss and how this is done (e.g., academic standards). This aspect relates to our anticipatory plans, and projects our ideas onto a Platonic virtual world, incorporating many remote sources of such timeless knowledge (other people, books, etc.) and relevant tools (forms of external scaffolding) with which to manipulate it. In this area, due to the disconnectedness from real-world issues, we can employ our standard reductionist techniques and scientific methodologies, which are based upon concepts that are removed from their fuzzy contexts (and referents) and discretized as stand-alone symbols. It is important to note however that once we try to implement such plans and theories we must then reconnect to the next

lower level of cognition, and this leads to many potential complications and dynergic effects often overlooked in our academic model simplifications, which tend to divorce concepts and plans from their relevance to other issues and values. In particular, feedback from the environment (circular causality via external loops) tests our actions, along with the assumptions of the theoretical models behind them, in a full multi-agent context, although the distal results of this testing are often only apparent in the long term coevolution of the whole and are easily missed or denied, due to our focus only upon immediate proximal effects (e.g., DDT side-effects).

Our model effectively merges the emotional and connectionist (associative) neural net components of mind, since it seems evident that they do not operate in disjoint ways (i.e., cognitive-affective space is unified). Emotions are regarded, following Cosmides and Tooby (2000), as guiding factors for our decisions, ways of weighting important values and contexts (e.g., by flagging dangerous situations based upon past results in similar contexts) and are thus just as important for effective choice as are the variable strength (experiential) associations between our concepts. We do not attempt here however to specify in detail how emotional and attractor structures relate, simply taking as a working assumption the ability of (limbic system) emotional input to switch (cortical) attractors (and perhaps subsystems) and so direct an organism along different paths, in other words they function as additional (contextually driven) control parameters for the set of neural constraints, perhaps helping to resolve conflicts of values. This does not mean that consciousness (level 3) is emotion free however, our detachment from space and time does not imply a detachment from values, so that many of our conscious behaviours will still have emotional components (see Thagard, 2001). Indeed Damasio (1999) claims that all conscious states have emotional content (he uses the terms 'core consciousness' and 'extended consciousness' in place of what we call subconscious and conscious). However, we should note that our awareness of an emotion is not the same as having such a state (any more than we need to be aware of a pain in the toe, or of perceiving a nearby dog, for these states to exist subconsciously), so we can choose at cognitive level 3 to ignore emotions, just as we can ignore other stimuli that are considered irrelevant to the value we are considering in isolation. Primary (reflex oriented) emotions may also prove to be relevant to cognitive level 1 operations.

Previous Unified Mind Theories

Before we outline the potential benefits of our model, perhaps we should spend a little time looking at some of those other ideas that have been promulgated as 'unified' theories of mind in recent years. We will require from such integration a positioning of the valid work of other researchers, but outright dismissal of such work must be based on scientific rather than axiomatic grounds. We clearly all have reflexes (analogue), and academics do employ symbolism (digital), so axiomatic claims that either are unnecessary to explain Mind (factual denials) seem incoherent scientifically – and we all also have a lot of fuzzy areas in the middle. We have already, of course, hinted at Maslow's humanistic psychology ideas, and in fact Hampden-Turner (1981) includes an additional 59 diverse historical or current mind-map perspectives. He comments: "*This entire book is a plea for the revision of social science, religion and philosophy to stress connectedness, coherence, relationship, organicism and wholeness, as against the fragmenting, reductive and compartmentalising forces of the prevailing orthodoxies*". Each of the models listed however, useful though it might be, targets either a subset of mind, or takes a stance which one might call vaguely 'holistic' and scientifically uninspiring. Only two are worth further mention in our context. Firstly the coevolutionary psychological field theory of Lewin, whose concentration on 'regions', barriers, needs and variable strength beliefs can be mapped quite well onto the attractors in state space and fitness-changing trajectories model we suggest, offering a better formalisation of his core concepts than was available in his time. Secondly the Freudian focus on the subconscious, whose 'id', 'ego' and 'superego' constructs bear a superficial resemblance to our 3 level meta-model. However the oppressively psychopathological attitude Freud took to these has regrettably consigned subsequent discussion and expansion of such issues largely to the rather unscientific world of the psychoanalysts. Our focus is much more positive to all these levels, relating them to fitness changing (if canalized) actions involving instincts, volitions and norms respectively (whether these actions prove to be positive or negative will be an empirical and contextual matter).

A very early claim to an explicit 'unified' theory is Churchland (1986) and this neurophilosophical approach is a necessary and ongoing research direction, with much in common to our central level 2 focus. But by largely ignoring body, emergent causality and social beliefs it was not (in its original form) 'integrating' – the term adopted, 'eliminative materialism', said it all. Perhaps the best known of more modern unified cognitive science perspectives is Newell (1990), based upon the AI production system 'Soar'. However, as Hofstadter (1995) says, this rather superficial approach neglects 'concepts' almost entirely and ignores also the emotional basis of choice (areas we consider vital to any adequate mind theory), such limitations still seem to remain

in current versions. An earlier attempt, which partly inspired (and is complementary to) this paper, was Minsky's (1988) "The Society of Mind", yet this too took a predominantly level 3 symbolic approach, with little neurological or dynamical grounding, and was more a collection of ideas than a joined-up meta-model. Goertzel's (1996) complexity science based 'psynet' model is perhaps closest in spirit to the ideas presented here, but takes a rather more computational and mathematical stance on mind than the fuzzy biological focus that we adopt. It could be regarded as a detailed 'top-down' model of the attractor dynamics of our cognitive level 2, but shows less emphasis on interacting values, emotions and the embeddedness of mind (in both brain and environment) than we think necessary to model human minds. The dynamical systems approach proposed by Beer (1997) emphasises the time-grounded environmental embodiment necessary for adaptive behaviour on levels 1 and 2, but fails to incorporate the symbolic representations and extended environment necessary to adequately model level 3. The latest version of ACT-R, Anderson et al. (2002), another production system, claims integration and biological (subsymbolic) grounding and does show promise, yet by adopting a rigid, largely serial (one goal/model at a time), visual-oriented architecture it neglects the true parallelism of human goal processing, particularly the multi-modality of cognitive levels 1 and 2, nor is it clear how any production system framework could answer the ontological development issues. A more recent attempt at unification by Sloman (2002), using 'virtual machine' computational analogies, also emphasises the complexity of mind, along with the need for multiple levels of architecture and emotions. We regard this work highly, but would emphasise the non-linguistic and subconscious nature of our (parallel) attractor-based central 'deliberative' layer, and include a focus on the time issues and adjacent level constraints missing there.

Moving on to more philosophical integrated accounts of mind, we have the rather poetical (a multiple-meaning thought mode usually ignored in symbolic accounts) treatment of Humphrey (1993) which regards consciousness as "*a special sort of doing*", an embodied naturalistic focus also common to accounts by Flanagan (1992) and Dennett (1995). All these types of accounts have their good points and insights, yet lack (to a scientist's mind) adequate concrete mechanisms, ways of implementing the 'doing', of generating a hard model or an AI system. A recent approach to the philosophy of mind and representation by Christensen and Hooker (forthcoming), highlights the inadequacy of both the dynamical and symbolic approaches, and argues for a similar conceptual paradigm to our own, with an emphasis on the self-organization of embedded, self-directed and multidimensional cognition, but the mesoscale implementation details are missing. An interesting comparison between Buddhist philosophy and complexity ideas by Waldron (2002) however generates a meta-view (discarding the "I" of the 'Cartesian Theater') which proves to be surprisingly similar to the overall structure of our meta-model (but, of course, at a much more spiritual level).

All in all, current accounts claiming some form of 'completeness' seem mostly to remain too vague or to come down to a 'one size fits all' focus, a claim (in various forms) that a single answer covers all the territory (often regarded as exclusively rational and conscious, a very narrow definition of intelligence, or quite the opposite – just the 'tips of the icebergs' to our view). Such models must be regarded as too simplistic to cover a system as complex and diverse as the Mind, although recently many researchers have started to back down from the more 'over-the-top' claims (e.g., van Gelder, 1999). Indeed much irreducible complexity seems implied by the cybernetic 'Law of Requisite Variety', Ashby (1956), if the organism is expected to cope with a complex world. But does that mean that the mind is such a 'hodge podge' that we can never get a focus on it, as some have claimed? We think not, and consider our viewpoint to lead to a perfectly coherent focus on Mind as a whole.

POTENTIAL APPLICATIONS

So, given such a model, is it really of any use? We consider the uses to be many, and identify several new research possibilities. Firstly by conceptually highlighting the whole picture – the Mind as it really exists, we can avoid being misled into thinking that any of our simplified maps are equal to the territory. Thus many of the disagreements about mind and consciousness can be seen for what they are – premature attempts to impose various sorts of one-dimensional simplified models as being adequate to model the integrated whole. We thus are in a much better position to put our restricted disciplinary-based models collectively into a wider context, and to understand their advantages and limitations in terms of the real complexity of embedded mind and the cognitive level(s) being targeted by each.

Secondly, by making explicit the multi-valued nature of our needs, we can see that value interactions (ignored in a reductionist viewpoint) are crucial to understanding the real (as opposed to virtual model) nature of the operation of Mind. One crucial implication of this epistasis is that we cannot optimize each value separately and then merge their solution sets, that approach fails for nonlinear optimization problems, leading to self-contradictory actions. From an interactional

perspective we need to position our models in a way that respects the full context, in other words we need to check if those variables or values taken to be independent in reductionist models are truly independent, as opposed to mere assumptions of being such, and to evaluate the relevance (if any) of those sets of values ignored, within the wider context (which also includes other agents and their, possibly conflicting, sets of values).

Thirdly, by making the neural constraints central to our model, we cut through those disjoint approaches that only consider either neuron electrochemistry or multiply realisable symbols as relevant to modelling mind. The concepts of complexity science show that the intermediate level of self-organized attractors is crucial to a proper understanding of any complex system – we require neither only the parts, nor only the wholes, but the way these both interact by upward (constructivist) and downward (constraintivist) causality, i.e. the 2nd order cognition intermediary, in which a concept's (nondescriptionist) 'context' can be related to the basin of the attractor and the sets of stimuli that enter it. This perspective both allows, and insists upon, inputs from both adjacent causal levels – the biological and the cultural, the parameters of each of these have demonstrable effects on both the static attractor structure and the dynamical behaviour of such complex systems.

Neuron networks are seen to exhibit spontaneous spiking, noise that randomly instigates low firing-rate activity. Given the many parametric inputs in our model, all of which will vary at some rate, we have a situation whereby an unconstrained network is predicted to explore state space, stochastic perturbations causing it to swap between attractors and, due to parametric variations, also dynamically altering the attractor structures themselves by changing system connectivity. This seems just what we require to implement trial and error (variation and selection) learning, with 'success' relating to entering by chance a particularly stable existing attractor. Again this can also implement creativity, whereby new associations between groups of concepts are activated randomly, due to their correlated firings, a particularly stable configuration corresponding then to our new idea, creating a new, perhaps temporary (as in dreams) transient attractor or association, later remembered by Hebbian processes. This mode could also be triggered by environmental novelty, allowing single experiences to generate new ideas and meeting our identified requirement for temporal efficiency in learning. Thus the model could assist in the understanding of child learning and in the understanding of genius – both perhaps related to the level of noise (maybe higher in children) and parametric variation (the more experience, the more parameters, and the more connections, then the greater the evolutionary possibilities and novelty may become).

Another area where the model may give some insights is in the idea of 'small-world' networks (where most nodes have low connectivity but a few more 'critical' nodes have high connectivity – allowing any node to connect to any other within, say, five steps). Generalisation could then be regarded as involving high connectivity attractors, leading to many specialist (low connectivity) 'property' attractors. This 'central node' idea may then be important for analogy, i.e. activating the high level node 'animal' may simultaneously activate both 'dog' and 'toy', leading to a new analogy of the (alive) 'toy poodle'.

The idea of 'elastic time' central to 2nd level cognition, introduces the idea of dynamic priorities. Each need has a strength that varies with time and fluctuates according to the relevant actions (hence optima are dynamic and not static states). Thus, the more time we spend in an action attractor (e.g., drinking) the weaker the need (thirst) becomes and the easier it should be to escape that basin of attraction. We need therefore to consider how basins can vary in depth dynamically, such that barriers to escape lower with 'satisfaction' levels. This is reminiscent of chemical barriers and we may be encouraged to look for enzymatic equivalents (since enzymes lower such chemical reaction barriers) in terms of the effects of parametric inputs on such structure. Here we may find a role for (emotionally derived) hormonal levels altering (neuronal) attractor thresholds (e.g., inhibiting their connectivity and eventually disabling them).

The importance of the ideas of self-organization and downward causality to the brain-mind question has been previously highlighted (e.g., Szentágothai, 1987). That these features apply at all levels of evolution, and not just to our human species is also emphasised. Thus any adequate model of Mind has to cater for this continuum of functionality. The proposed meta-model does this by an explicit evolutionary partition of organism needs into three emergent areas, mapping these onto processes compatible with the physical structures of those organisms possessing such values. In this way we extend the 'three-world concept' proposed by Popper (1972), identifying World 1, his physical world, with dynamical system models (adequate for simple causal cognition – e.g., behaviourism), World 2, his mental world, with the complex systems models necessary for considering interacting subjective values and embodied autonomy, and World 3, his 'objective' world, with the detached symbolic models and the extended virtual world common to academia.

Explicitly separating out the academic world (which operates mostly in World 3) from the world of the passive 'man-in-the-street' (who operates mostly in World 2) makes clear just how different these worlds are and how irrelevant to grounded human issues much academic debate becomes. As

scientists we need to re-evaluate just what it is we are trying to model, is it the Mind or just a intellectualised sub-set of it? If the former, then we need to put much more emphasis on those aspects of mind that involve the mass of humanity in their day-to-day lives, i.e. that takes into account the irrationality, ambiguity, delusions and fallibility of typical human behaviours. This relates to a focus upon the semi-conscious, the way cognitive levels 2 and 3 inter-relate and how we move between them. Given the prevalence of 'edge-of-chaos' balances in complex systems generally, we can expect that our normal behaviours here will prove to inhabit a similar mid-way position, such that roughly half of what an aware person does should prove conscious and half subconscious, with much fractal variation over time. Our model thus suggests that neither the academic nor the passive human viewpoint is balanced with respect to maximising fitness, and this has many implications for our society and cultural behaviours generally.

Our emphasis upon the embedded nature of mind focuses attention more on the environment, and on the ability of organisms to use this as part of their being, either as a form of external information storage or of inter-organism communication (as in a stigmergic perspective). This guiding ability can operate on all three cognitive levels, at level 1 roads (and pheromone trails) for example canalize movement and act as physical route maps, avoiding the need for explicit navigation. At level 2 we have the exchange of signs (e.g., a signpost), indicators of potential fitness and choice, and at level 3 we have the explicit knowledge stored in books and computers (e.g., distances from A to B). Many of our cognitive operations depend just as much upon such external scaffolding for their proper function as do our more autopoietic bodily functions upon their own forms of constant external support (e.g., air, water, food, shelter and mates).

By highlighting the impossibility of any agent creating a full model of its environment (due to the indeterminate number of parametric influences) we help bring to light the sorts of informational filters that we are employing in any context. These can be of several types: limited sampling of the environment (selective blindness), specialisation (niche operations), social rules and norms (taboo actions and fixed roles), institutionalisation (collective solutions), seriality (dealing with only one issue at a time), dogmatism (discretization of continuous variety), imitation (crowd behaviour), delusions (heuristics divorced from 'reality'), withdrawal ('ivory tower' isolation) etc. All relate to attempts to reduce the information processing needs of the agent, thus solutions are adopted that are contextually 'good enough' (i.e., usually adequate – satisficing behaviour) showing that the attempts of computational modellers to obtain exact mathematical solutions in 'toy' domains are

unrealistic in 'real-world' terms (and the Gödel argument largely irrelevant).

A further advantage of this model concerns the possibility of using agent modelling techniques to understand the difficult central level, which to all intents and purposes is analytically intractable (NP-complete). If we take, as an assumption, the idea that brain operation is analogous to that of ecosystems, then we can attempt to model mind as the emergent result of multiple agents (each trying to optimize an individual value – their niche) competing (and/or co-operating) for the use of resources (time, actions, etc.). The patterns of interactions amongst the agents would then form the dynamical attractors of the system in the usual way – we take as an assumption here that agents can have some initial (evolutionarily derived and genetically enabled) functionalities built-in, although what these may be is an open question. Many simulation frameworks are now available within the MAS community that could, with minor modifications, be put to use in the creation of artificial psychology models of this type (e.g., SWARM). Additionally, since such frameworks are also currently being used to model social/organizational systems (artificial sociology) and genetic/biological networks (artificial chemistry and biology) then a good deal of cross-fertilisation can be expected, especially as we need (eventually) to join these three levels of influence on neural function together in order to generate an adequate overall model – a common framework will considerably ease this ideal (see Goldspink, 2000). Indeed, if we add in also the similar studies being undertaken of ecological networks, then this common methodology potentially can incorporate environmental sustainability issues also and how these interact with human beliefs and needs (e.g., Jager, 2000).

Using such a framework allows us to try out experimentally many possible models of the active mind, i.e. many alternative sets of interacting values and associated beliefs (in a way that is impossible using real people, but necessary if we are to test our models). Each agent could consist of a target value, internal states (e.g., priority), and behavioural rules (beliefs). States, rules and connections to other agents would be able to evolve over time as a result of interactional (parametric) or random (noisy) changes. An internal goal, or fitness function, would monitor the agent's performance at its task, triggering connections to other agents in an attempt to find the resources it needs to survive (satisfy its value). Synaptic 'tags' could identify resources, allowing connections to 'die' if an adequate resource is not located. In this way randomly initiated networks could be grown (i.e., their organization evolved rather than designed), and various alternative optima (the evolved 'ecologies of mind') generated by self-organization (recent work on action

selection, based upon negotiating between sub-agents with conflicting desires (e.g., Humphrys, 1996), shows the possibilities of such evolutionary modelling using reinforcement learning). By tuning model parameters (e.g., maximum connectivity distance, maximum number of connections, value consumption rate, and so on) we would have much scope to investigate synthetic neural behaviours (with some commonalties to ongoing work in evolving artificial neural network architectures). This is just one possible scenario, there are various communications languages, interaction protocols, and agent architectures for facilitating the development of multi-agent systems being generated by MAS researchers (e.g., Belew and Mitchell, 1996), but much remains yet to be understood.

Expansions to such models, such as hierarchical structures, interfaces to an environment and the creation of additional personal values following environmental novelty (e.g., a liking for computer gaming), can be envisaged. Additionally we would like to evolve concepts, using fuzzy matching, with groups of associated agent values (properties) comprising the multidimensional attractor basin of the concept, and also labels, viewed as higher level value agents linked to conceptual groups – symbolic identification pointers (e.g., Millikan, 1998) – to attractors that are only fuzzily similar between individuals. This link between the label and an attractor effectively solves both the symbol grounding problem, since the (passive) symbol's 'meaning' is simply all those states (triggered either externally or internally) which enter the (active) attractor to which it has become associated; and also perhaps the framing problem since only 'relevant' attractors would be active in any context, perhaps with a 'strength' proportional to their associative importance. We can also allow that the same label connects to multiple attractors, permitting the multiple meanings associated with poetry and ambiguity. It may prove to be the case that (over time – as in child development) disjoint parts of this label network come into being, areas that are self contained and divorced from the fuzzy values of the main system, this could relate to 'forgetting' part or all of the grounding originally used to create the symbol. This could perhaps be regarded as the evolution of a World 3 capability, the operation of abstract values (e.g., axiomatic mathematical systems, with their associated Aristotelian either/or logics). Rationality then would relate to filling in the missing information (and going beyond experience to create novel ideas) by forms of reason. It would seem a relatively simple matter for parametric inputs to activate and deactivate sets of such attractors in different circumstances, perhaps implementing an 'emotional' control scheme using 'monitor' agents. At a lower level, World 1 features seem also to naturally occur with regard to the causal nature of individual agent interactions (disjoint sets of lower

level input-output agents – habit formation). Thus, potentially at least, such models could generate the full range of activities of our cognitive meta-model, possibly leading to the implementation of a viable artificial mind.

In this way we could, over time, come to understand what sort of belief structures and value adoptions results in a psychologically optimum 'person' (within particular environmental contexts of course) and identify what alternatives (different ways of doing things) are equally viable (i.e., generate the same Pareto-optimal global fitness). This has great potential for revision of our educational systems, which are currently positioned almost exclusively in World 3, divorced from the integration taking place in World 2 and the values and relevances that that level represents. Understanding Mind will allow us to relate our education better to cognitive level 2 (day-to-day behaviours), so that we can then perhaps encourage the learning of more useful, dynamically tested, strategies (obtained via our simulations), rather than imposing educational abstractions and untested 'status-quo' intellectual and cultural norms, as we so often do today – which can either prove irrelevant or socially destructive.

Although this paper does not directly attempt to account for qualia, we should perhaps outline how these would fit in to the model. Emergent properties, of any kind, lead to functions not describable in terms of part properties (so are not then just aggregates). These are supervenient however on the parts (so no substance dualism) and are causally effective (so no epiphenomenalism). The exact status of such evolutionary emergence (at all levels) is an open question in the complexity community, although its scientific existence is unproblematic to most (e.g., Emmeche et al., 1997, 2000; Sloman and Chrisley, 2003), and it follows naturally from a process metaphysics viewpoint (e.g., Bickhard, 2000). Since we take a constructivist view to all mental phenomena (the environment triggering internal attractors, rather than acting instructionally to impose a 'mirror image' map of itself – a coherence view of truth rather than a correspondence one) we regard our perceptions, emotions and qualia in much the same light, how a tree 'looks' as it does and how green 'feels' thus being ontologically (but not epistemologically) equivalent.

A final point relates to the origin of the values which form such a central part of our model and their relationship to human autonomy. We have to stress here the social and cultural origins of many of our higher level values and behaviours, in contrast to the perhaps overemphasised biological origins of the evolutionary psychologists. Evolution operates, in our view, at every level and timescale and, following Dennett (1995), is an unproblematic addition to both psychology and sociology. Coevolution can quite easily generate both generalist and specialist problem solvers in the

same structure (which we can also do artificially in what are called 'learning classifier systems', where specialists override generalists if they compete for the same resource – 'special cases' taking precedence over 'general rules'), and this also applies equally to autonomy or self-control, whether at subconscious or conscious levels (we do, to some extent, engineer the former in planetary-explorer robots). To maximise our, individual or collective, multidimensional fitnesses however (and it is overall 'quality-of-life', not mere 'survival', which drives even animal behaviour in our evolutionary viewpoint) we often need to go beyond the constraints of both genes and memes, and to do this effectively we need to understand their influences on our choices, a point our model emphasises strongly. In practical terms, it is not expected that we will often be able to calculate the global optima necessary to maximise fitness, but we can derive local relative fitnesses quite easily and 'hill climb' from our current position to better ones (although with much propensity for error, given our generally poor understanding of the coevolutionary effects of our actions).

CONCLUSION

To move from an external 'objective' study of passive systems to a coevolving autonomous agent perspective of embodied active systems (whether at a cellular, social or ecological level) has proved beneficial in both artificial simulations and in real world understanding, and this should apply also within the mind. Consciousness requires some form of autonomy, a directed search through possibility space with a view to choosing the best action. This common goal, seen in all forms of life, requires balancing multiple nonlinear objectives or values and is difficult to relate to a serial, disjoint dimensional and linear form of consciousness (a relatively blind mode which fails to account for systemic feedback effects). In complex systems, such global optima as prove possible are more frequently found by distributed (parallel) systems, operating near the edge-of-chaos with dynamically changing interactions, and this suggests that treating mind as a multi-conscious multi-level coevolutionary system of this form may be a more fruitful approach than many current 'single-consciousness' focuses.

This is not to say that such studies are easy, artificial simulations are still many orders of magnitude simpler than real minds and societies, and the artificial neurons or agents employed are similarly simplistic with respect to the real thing. Our understanding of high-dimensional dynamics (where multiple attractors are simultaneously active – as needed to model nested levels of detail) is also very primitive so far. Additionally in Mind studies we have often neglected the mid-level interface between neurological structural detail and conscious symbolism and have much to learn yet here. We have also failed to adequately take into account the effects of emotions on values and on rational states, or to consider how far social and biological structures constrain our inherent cognitive development and abilities. There is still much to do before any reasonably comprehensive metaview of 'mind as a whole', seen as a dynamical interplay between many levels of complex systems, can be said to have been formulated. Our meta-model here is seen as merely a first (biologically plausible) step along this path, a positioning of the whole in a way that better focuses our minds on currently unexplored possibilities.

Some of these possibilities include the effects on attractor structures of dynamical changes to parametric inputs, especially the evolutionary biological influences; the constraints upon attractor development (learning) imposed by the downward causation of social norms; the way in which self-organizational balance operates in multi-valued scenarios; the effects of intentional changes (choices) upon higher level attractor dynamics; and the mechanism whereby new concept attractors can be instantly generated by novel combinations of stimuli. Additionally, by making explicit the three overlapping timescales of action: absolute, relative and timeless, we suggest that different cognitive approaches need to be used to best model each, dependent upon their connectivity and optionality, especially in treating the often neglected subconscious forms of cognition. A common ecological methodology has been suggested with which to approach this neglected middle level of interacting needs, based upon modelling the modularity of mind in the form of a multi-agent system. This approach lends itself well to treating the multi-level nested structures of our world, and also allows for much interdisciplinary cross-fertilisation.

REFERENCES

ABRAHAM F. Chaos, bifurcations and self-organization: Dynamical extensions of neurological positivism and ecological psychology. *Psychoscience, 1:* 85-118, 1992.

ANDERSON JR, BOTHELL D, BYRNE MD and LEBIERE C. An integrated theory of mind. *Psychological Review, 111:* 1036-1060, 2002.

ASHBY WR. *An Introduction to Cybernetics*. London, UK: Chapman and Hall, 1956.

BAK P. *How Nature Works: The Science of Self-Organized Criticality*. Lublin, PL: Copernicus, 1996.

BAR-YAM Y. *Dynamics of Complex Systems*. Redwood, CA: Addison-Wesley, 1997.

BATESON G. *Steps to an Ecology of Mind*. Chicago, IL: University of Chicago Press, New Ed., 2000.

BECK D and COWEN C. *Spiral Dynamics*: *Mastering Values, Leadership and Change*. Oxford, UK: Blackwell, 1996.

BEER R. The dynamics of adaptive behaviour: A research program. *Robotics and Autonomous Systems, 20:* 257-289, 1997.

BELEW RK and MITCHELL M. *Adaptive Individuals in Evolving*

Populations: Models and Algorithms. Redwood, CA: Addison-Wesley, 1996.

BICKHARD MH. Autonomy, function and representation. *Communication and Cognition: Artificial Intelligence, 17:* 111-131, 2000.

BICKHARD MH and CAMPBELL DT. Variations in variation and selection: The ubiquity of the variation-and-selective retention ratchet. *Emergent Organizational Complexity. Foundations of Science, 8:* 215-282, 2003.

BROOKS R. Elephants don't play chess. *Robotics and Autonomous Systems, 6:* 3-15, 1990.

CAIRNS-SMITH G. *Evolving the Mind: On the Nature of Matter and the Origin of Consciousness*. Cambridge, MA: University Press, 1996.

CALVIN W. *The Cerebral Code*. Cambridge, MA: MIT Press, 1996.

CAMPBELL DT. Downward causation in hierarchically organized biological systems. In Ayala FJ and Dobzhansky T (Eds), *Studies in the Philosophy of Biology*. London, UK: Macmillan Press, 1974.

CHRISTENSEN WD and HOOKER CA. Representation and the meaning of life. In Clapin H, Staines P and Slezak P (Eds), *Representation in Mind: New Approaches to Mental Representation*. Oxford: Elsevier, ANNO, pp. 00-00.

CHURCHLAND P. *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press, 1986.

CLARK A. Reasons, robots and the extended mind. *Mind and Language, 16:* 121-145, 2001.

COMBS A. *The Radiance of Being: Complexity, Chaos and the Evolution of Consciousness*. Edimburgh, UK: Floris Books, 1995.

COSMIDES L and TOOBY J. Evolutionary psychology and the emotions. In Lewis M and Haviland-Jones JM (Eds), *Handbook of Emotions,* 2nd Ed. New York: Guilford, 2000.

DAMASIO A. *The Feeling of What Happens: Body, Emotion and the Making of Consciousness*. Orlando, FL: Harcourt Brown and Company, 1999.

DEMPSTER B. Sympoietic and autopoietic systems: A new distinction for self-organizing systems. In Allen JK and Wilby J (Eds), *Proceedings of the World Congress of the Systems Sciences and ISSS,* 2000.

DENNETT D. *Darwin's Dangerous Idea*. London, UK: Penguin, 1995.

EDELMAN G. *Bright Air, Brilliant Fire: On the Matter of the Mind*. New York: Basic Books, 1992.

ELIASMITH C. The third contender: A critical examination of the dynamicist theory of cognition. *Philosophical Psychology, 9:* 441-463, 1996.

EMMECHE C, KØPPE S and STJERNFELT F. Explaining emergence. *Journal for General Philosophy of Science, 28:* 83-119, 1997.

EMMECHE C, KØPPE S and STJERNFELT F. Levels, emergence and three versions of downward causation. In Andersen PB, Emmeche C, Finnemann NO and Christiansen PV (Eds), *Downward Causation. Minds, Bodies and Matter*, Aarhus University Press, 2000, pp. 13-34.

EPSTEIN J. Learning to be thoughtless: Social norms and individual computation. *Santa Fe Working Paper:* 3-22, 2000.

FLANAGAN O. *Consciousness Reconsidered*. Cambridge, MA: MIT Press, 1992.

FRANKLIN S. *Artificial Minds*. Cambridge, MA: MIT Press, 1995.

FREEMAN WJ. The physiology of perception. *Scientific American, 264:* 78-85, 1991.

GAZZANIGA M. *Nature's Mind: The Biological Roots of Thinking, Emotions, Sexuality, Language and Intelligence*. New York: Basic Books, 1992.

GOERTZEL B. *The Evolving Mind*. Gordon and Breach, 1993.

GOERTZEL B. *From Complexity to Creativity*. Plenum Press, 1996.

GOLDSPINK C. Modelling social systems as complex: Towards a social simulation meta-model. *Journal of Artificial Societies and Social Simulation, 3: 2,* 2000.

GROBSTEIN P. From the head to the heart: Some thoughts on similarities between brain function and morphogenesis and on their significance for research methodology and biological theory. *Experientia, 44:* 960-971, 1988.

HAMPDEN-TURNER C. *Maps of the Mind*. Mitchell Beasley, 1981.

HOFFMEYER J. The swarming body. In Rauch I and Carr GF (Eds), *Semiotics Around the World. Proceedings of the Fifth Congress of the International Association for Semiotic Studies*. Berkeley, 1994. Berlin/New York: Mouton de Gruyter, 1997, pp. 937-940.

HOFSTADTER D. *Fluid Concepts and Creative Analogies*. New York: Basic Books, 1995.

HOLLAND J. *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press, 1992.

HUMPHRYS M. Action selection methods using reinforcement learning. In Maes P et al. (Eds), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, 1996, pp. 135-144.

HUMPHREY N. *A History of the Mind*. Vancouver, WA: Vintage, 1993.

JAEGER H. From continuous dynamics to symbols. In Tschacher W and Dauwalder JP. *Dynamics, Synergetics, Autonomous Agents. Studies of Nonlinear Phenomena in Life Science Vol. 8*. World Scientific, 1999, pp. 29-48.

JAGER W. *Modelling consumer behaviour*. Kansas City, MO: Universal Press, 2000.

KAUFFMAN S. *The Origins of Order – Self-Organization and Selection in Evolution*. Cambridge, MA: Oxford University Press, 1993.

KELSO JS. *Dynamic Patterns: The Self-Organization of Brain and Behaviour*. Cambridge, MA: MIT Press, 1997.

LUCAS C. Transient attractors and emergent attractor memory. *CALResCo Group Working Paper*, 1997.

LUCAS C. Value metascience and synergistic choice. In Halloy S and Williams T (Eds), *Applied Complexity: From Neural Nets to Managed Landscapes*. New Zealand Institute for Crop and Food Research Ltd., Christchurch, 2000, pp. 53-87.

LUCAS C. A logic of complex values. In Smarandache F (Ed), *Proceedings of the First International Conference on Neutrosophy, Neutrosophic Logic, Neutrosophic Set, Neutrosophic Probability and Statistics*. Xiquan, Phoenix, 2002, pp. 121-138.

MASLOW A. *Toward a Psychology of Being*. New York: Van Nostrand Reinhold, 1968.

MATURANA H, MPODOZIS J and LETELIER JC. Brain, language and the origin of human mental functions. *Biological Research, 28:* 15-26, 1995.

MAYNARD SMITH J and SZATHMANY E. *The Major Transitions in Evolution*. S. Francisco, CA: Freeman Press, 1995.

MILLIKAN RG. A common structure for concepts of individuals, stuffs, and real kinds: More Mama, more milk, and more mouse. *Behavioral and Brain Sciences, 9:* 55-100, 1998.

MINSKY M. *The Society of Mind*. New York: Simon & Schuster, 1988.

MITCHELL M. Can evolution explain how the mind works? A review of the evolutionary psychology debates. *Complexity, 3:* 17-24, 1999.

NEWELL A. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press, 1990.

POLANYI M. *Personal Knowledge: Towards a Post Critical Philosophy*. Chicago, IL: University of Chicago Press, 1958.

POPPER KR. *Objective Knowledge*. Oxford University Press, 1972.

POTTER MA and DE JONG KA. Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation, 8:* 1-29, 2000.

QUARTZ SR and SEJNOWSKO TJ. The neural basis of development: A constructivist manifesto. *Behavioural and Brain Sciences, 20:* 537-596, 1997.

SHEPHERD GM. *The Synaptic Organization of the Brain*. Oxford, UK: OUP, 1998.

SLOMAN A. Architecture-based conceptions of mind. In Gardenfors P, Kijania-Placek K and Wolenski Kluwe J (Eds), *In the Scope of Logic, Methodology, and Philosophy of Science,* vol II. *Synthese Library 316*, 2002, pp. 403-427.

SLOMAN A and CHRISLEY R. Virtual machines and consciousness. *Journal of Consciousness Studies, 10:* 4-5, 2003.

SZENTÁGOTHAI J. The 'brain-mind' relation: A pseudoproblem?. In Blackmore C and Greenfied S (Eds), *Mindwaves*. Blackwell, 1987, pp. 323-336.

THAGARD P. How to make decisions: Coherence, emotion, and practical inference. In Millgram E (Ed), *Varieties of practical inference*. Cambridge, MA: MIT Press, 2001, pp. 355-371.

VAN GELDER T. Revisiting the dynamical hypothesis. *University of Melbourne, Department of Philosophy Preprints*, No. 2/99, 1999.

VELMANS M. Intersubjective science. *Journal of Consciousness Studies, 6:* 299-306, 1999.

VON FOERSTER E. On constructing a reality. In Preiser FE (Ed), *Environmental Research Design, vol 2. Stroudsburg*. Dowden, Hutchinson and Ross, 1973, pp. 35-46.

WALDRON WS. Buddhist steps to an ecology of mind: Thinking about 'thoughts without a thinker'. *Eastern Buddhist,* vol. XXXIV, No. 1, 2002, pp. 1-52.

Chris Lucas, Ph.D., CALResCo Group, 52 Mount Road, Middleton, Manchester M24 1DZ, UK. e-mail: cl@calresco.org