Preface

> *And the world was One vast cellular automaton (or `CA'), calculating out the instants - and each of the world's diverse objects was but a subcalculation, a simulation in the One great parallel process.*
> Rudy Rucker, *Wetware*

This thesis takes the form of a critique of what I shall refer to as computationalism, the metaphysical view that the world at its most fundamental level is computational in nature. A number of issues immediately arise in connection with any such critical study including the factors motivating the study, the existence of precedents supporting the critique, the approach adopted and how it differs from other approaches, and finally, the implications of the investigation for the subject under study. These issues require consideration.

Why should such a critique be necessary ?

Physics is the study of the most primitive phenomena of the universe. Metaphysics, on the other hand, is the study of that which is beyond what physics *can* study; it is the study of our assumptions about the ultimate nature of reality. But is it reasonable to study computationalism as a *metaphysical* position ? I would argue, following Cariani (1989), that computationalism has *become* a metaphysical position. No longer is it the guiding hypothesis uniquely identified with its founding disciplines, cognitive science and the philosophy of mind. Computationalism has transcended its traditional confines and is now being investigated in connection with disciplines associated with the study of other phenomenal domains, for example, the philosophy of life and the philosophy of matter. According to Bunge (1977), "metaphysics only recently has undergone a revolution so deep that nobody has noticed it: indeed ontology has gone mathematical and is being cultivated by engineers and computer scientists." (p.7) As a consequence of the trend towards viewing computation as ubiquitous in the phenomenal world, and supported by the conceivability of a unification of phenomena under an integrated scheme grounded in computation, metaphysical computationalism has become an increasingly attractive position. Unification is a fundamental goal within science and philosophy, drawing inspiration from Ockham's Razor, which is captured in the maxim "simplicity is the key", and exemplified by the unending search for Grand Unified Theories or/and Theories of Everything. Computationalism, like materialism, idealism and other antecedent philosophical positions, attempts to provide a basis upon which such a phenomenal unification might be achieved. However, computationalism, like materialism, idealism and other positions, *is* a metaphysical position. It makes claims about reality, that which, in Kantian terms, is beyond mere phenomenal appearance. Hence, any serious study of computationalism must, of necessity, involve a consideration of metaphysical issues thereby justifying the examination of such concerns as reported herein.

What precedents exist in support of such a study ?

The idea that the ontology (that is, being or nature) of the universe is computational is not original to this study. Nor is the attempt at critically evaluating this metaphysical position. A number of popular works could be enlisted in support of such a critique; for example, Theodore Roszak's classic *The Cult of Information* (1986), Stephen Talbott's

The Future Does Not Compute (1995) and Mark Slouka's *War of The Worlds* (1995). Such works are usually written in an attempt to counter what the authors perceive to be the polemical excesses of adherents of computationalism and its related creeds. Although global in scope, the style of such works is often rhetorical and hence, generally tends to be dismissed by the academic mainstream, rather unfairly and somewhat prematurely in my opinion, as the "neo-luddite ramblings" of "misinformed intellectual light-weights". However, indirect support for a critique of computationalism from a more scholarly perspective has been provided by the various scientific and philosophical investigations of computational models used in connection with phenomenon-specific studies of matter, life and mind.

What approach is to be adopted and why is it different from others ?

As implied above, and following precedents set by Dreyfus (1979), Searle (1980) and others, previous approaches to investigating computationalism have mainly involved evaluating computational analogues of certain natural phenomenal kinds such as matter, life and mind[1]. There is an intimate connection between computationalism and another notion, viz. *artificiality*, which makes it possible to use the latter to investigate the former when suitably conceived; hence, the investigations of artificial intelligence, life and reality reported in the literature. Such studies have, however, tended to ignore the underlying generic concept of artificiality itself, the one exception being Herbert Simon's *The Sciences of the Artificial* (1969). Unfortunately, this work, which introduced the concept of artificiality, did not adequately clarify the various distinctions holding between this term and opposing concepts such as naturality (or the natural which can be contrasted with the artifactual, man-made or synthetic) and reality (or the real which can be contrasted with Kant's *as-if* or apparent). It also failed to address the ontological and epistemological issues associated with the concept from the perspective of that entity (human being) responsible for asserting that such a thing as artificiality exists. Consequently, a generic view of artificiality was not readily conceived, thereby rendering difficult the translation of arguments associated with one kind or instance of artificiality to other phenomenal kinds. (Most critiques based on argument-translation, when attempted, have usually been of limited scope, a result of their being grounded in a relatively superficial and non-generic analysis of what is common between the various kinds of artificiality).

This study attempts to resolve some of the genericity problems associated with Simon's concept of artificiality. This is accomplished using a series of conceptual frameworks that describe the ontological and epistemological relations holding between naturals, artificials and that which is making (or can make) the distinction, viz. human beings[2]. A new concept of artificiality is derived by identifying features common to all instances of the class of artificiality using these frameworks. Artificiality is defined as a category or class with artificialities such as artificial intelligence, artificial life and artificial reality as instances or kinds. A particular form of artificiality based on an emergentist

---

1 For the sake of argument, it is assumed at this preliminary stage in the discussion that matter, life and mind *are* natural phenomena.

2 This should not be taken as implying that human beings are the only beings capable of making the distinction; rather, that it is only human beings that are *known* to be capable of making such a distinction.

metaphysics is investigated. This *emergent-artificiality* enables individual artificial kinds to be hierarchically unified such that an isomorphism between evolutionary nature (or naturality) and artificiality is established. This both permits and simplifies the translation of arguments between the various disciplines and kinds of artificiality, making possible the application of phenomenological arguments due to Searle, Dreyfus and others in the context of a computationally-grounded unifying framework of emergent artificiality. Using such a framework, computationalism as a metaphysical thesis can be evaluated.

What are the implications ?

I believe that the implications of this study will lead to two main developments: First, the supplanting of computationalism, either by an alternative metaphysical tradition[3] or by a post-metaphysical *Weltanschauung* (world-view)[4]; second, and independently, an interpretative or *hermeneutic* shift in favour of reinterpreting computation in anthropocentric (or human-centred) terms. This would be consistent with the pragmatist (or instrumentalist) attitude towards technology and would contribute greatly to the appreciation of computer science as an engineering discipline incorporating existential or humanistic concerns[5].

Personal development of the problem

To understand and appreciate the numerous underlying factors motivating any long-term study, whether of a philosophical nature or otherwise, it is necessary to trace its historical development back to the point of its inception when the foundations for the study were first being laid. As Theodore Roszak states, "understanding an idea means understanding the lives of those who created and championed it."

My study began in opposition to the argument presented in this thesis with an attempt to realize what would later be identified as a variant of the Leibnizean dream of a universal language for intellectual discourse. According to MacDonald Ross (1984),

> Leibniz's approach was to try and reconcile the logical, rhetorical and geometrical traditions by blending their three distinct emphases (on formalism, on linguistic proprietary and on mathematicisation) into the single vision of a formal language notated mathematically. (p.50)

I was motivated by the goal of facilitating unambiguous communication between members of diverse intellectual communities involved in modelling natural phenomena with a view to designing artificial analogues of such phenomena. I held that what Winograd and Flores (1986) describe as "the rationalistic orientation" was the correct view on the basis of which that goal might be brought to fruition. Inspired by the then

---

3 For example, Whiteheadian panexperientialism.

4 For example, Heidegger's ontology of Being.

5 This would go some way to alleviating the concerns about neo-Luddism raised by technologists such as Florman (1976).

recent speculation on the possibility of a "post-symbolic" form of communication[6] associated with the emerging field of virtual reality, I embarked upon the ambitious project of developing what came to be called a generic interactive modelling paradigm (or GIMP) for modelling systems at the behavioural level. The term `generic' indicated the universality of the modelling framework and its application in disciplines as diverse as physics, biology, psychology, sociology, politics and even art. `Interactive' marked the connection to virtual reality and the possibility of a post-symbolic form of dialogue based on the gestural play of agents embedded within graphical environments or `virtual worlds'; in this respect, I was particularly inspired by the application of the Aristotelian theory of dramatic action to human-computer interaction described in *Computers as Theatre* (Laurel,91b). Why post-symbolic communication ? On the basis of what I now understand to be an implicit Heideggerian interpretation of the world, I argued that perhaps the post-symbolic world of `artificial reality' might mirror the pre-symbolic world of natural reality. By `modelling paradigm' was understood the amalgamation of a particular modelling methodology and an associated modelling technology; hence, the GIMP project was to have two deliverables, one conceptual and the other, technological.

I began to investigate the concept of a model and search for candidates that might meet the `genericity' requirement. The *McGraw-Hill Dictionary of Scientific and Technical Terms* (Third Edition, 1984) defines a model as

> a mathematical or physical system, obeying certain *specified conditions*, whose behaviour is used to understand a physical, biological, or social system to which it is analogous in some way [emphasis added].

The requirement that a system obey certain specified conditions indicates the imposition of constraints and the definition of boundaries. The problem of how these constraints are imposed, that is, whether via internal (endosystemic) or external (exosystemic) agency, is a notion that would become pivotal to the framework I was to later develop for differentiating between artificial and natural phenomena. Two ideas related to this problem were briefly considered at this stage of the study: First, that a model is only an abstraction of the phenomenon it is being used to help understand; hence, a model is, of necessity, an incomplete representation of reality. This idea was to form the basis for investigating the difference between the simulation of a natural phenomenon and the realization of that phenomenon; second, that there is more than one way to abstract away from a phenomenon, and hence, that modelling is, of necessity, a relativistic activity. However, at this stage it was still unclear whether such relativism was ontological, thereby supporting a subjective-idealist position, or merely epistemological, in which case it would be consistent with variants of objective-realism. The framework for differentiating between natural and artificial phenomena described herein makes it possible to resolve this issue.

The next stage of work involved concurrent exploration of the two aspects to the GIMP,

---

6 Lanier (1992) defines *post-symbolic* communication as follows: "In the physical world, you can't make physical changes to your world very quickly. The only thing you can do is use your tongue to form words that refer to all the possible changes you might make if you could. For example, you could say the worlds, `let's go flying on a giant squid', even though you can't actually do that in the real world. But in a good shared virtual reality system, you can just *directly* make up the objective world instead of using symbols to refer to it [emphasis added]." (p.69)

viz. its methodology and associated technology. The former necessitated a more detailed investigation of the notion of interaction within dramatic contexts and the idea of *agent-based* causation was introduced. Agents are entities with the capacity to initiate action (Laurel,91b). A *semiotic*[7] approach was adopted because it was seen as providing the necessary foundation for analysing the syntax (structure) and semantics (meaning) of agents and agent-based interaction. Technological issues motivated consideration of the notion of *behaviour*. The result was a first attempt at classifying systems according to their behaviours based on the binary oppositions tangible-abstract and animate-inanimate: Tangible systems were perceptual while abstract systems were wholly conceptual; animate systems were `living' whereas inanimate systems were not. At this stage, I was already starting to establish links between systems modelling and the emerging discipline of artificial life. I began to evaluate candidate technologies such as George Cherry's Stimulus-Response Machines (Cherry,91) and was then introduced to John Holland's work on genetic algorithms (Holland,92). It appeared that I might find the modelling primitive I was seeking in a synthesis of the two approaches. For this reason, I spent the next year looking more closely at the connections between genetic algorithms and evolutionary biology. It was during this time that I came into contact with the work of Ted Steele (1979). Steele had been reconsidering the ideas of Jean-Baptiste Lamarck, a predecessor of Darwin, who held that characteristics acquired during ontogeny (development) could be inherited by an organism and thereby pass into the phylogenetic (evolutionary) lineage of the species. Steele's work was inspired by the idea that retroviruses might provide the means by which to transmit somatic (or bodily) characteristics across Weismann's barrier and into the germline. Such a scheme clearly contradicted the `central dogma' of modern genetics which states that the passage of biological `information' is always in the direction DNA to proteins and never the reverse. Irrespective of the biological validity of such work, I made the connection between genetic algorithms, Steele's Lamarckian thesis, and the interactive modelling framework I was attempting to develop. Adopting a form of what might be described as "computational organicism", I embarked on the development of a biologically-inspired modelling primitive based on Steele's neo-Lamarckian ideas, viz. the CyberCell (Ali,92).

Ironically, it was during this time that I began to have serious doubts about the underlying assumptions of the project. Two independent sources were responsible for placing me in this situation. The first can be traced to the critiques of neo-Darwinism that I had encountered when studying evolutionary theory. Works ranging from Michael Denton's *Evolution: A Theory in Crisis* (1985) to the collection of papers in *Evolutionary Processes and Metaphors* (Ho,88) forced me to look beyond contemporary Darwinian and Lamarckian concepts of evolution. I discovered that the difference between the two interpretations was not so significant once both schemes were reconceptualized in the language of modern genetics with its metaphysically dualistic assumptions of chance and necessity. Implicit arguments against the adoption of a genetically-based behavioural modelling primitive such as the CyberCell were to be found in critiques of genetic determinism and biological reductionism such as *Not in Our Genes* (Rose,84) and *The Doctrine of DNA* (Lewontin,91), both of which were popular works written by leading geneticists. The second source of doubt lay in more direct criticism of the CyberCell concept by a colleague which forced me to reconsider its status as a modelling primitive and as the ontological substrate of artificial analogues of natural phenomena: The project

---

[7] Semiotics refers to the scientific study of systems of signs and sign-users.

was an attempt to develop a *generic* modelling paradigm; yet, it did not seem right that a biologically-inspired entity such as a cell should be regarded as more primitive than, say, an atom. After all, I reasoned, what are organisms composed from other than atoms ?

It was at this point that I was introduced to John Conway's "Life", a mathematical `game' (originally) played using tokens placed in some initial configuration on a potentially infinite grid of squares (with a maximum of one token per square). Life was a strange sort of game in that it required no players, merely the repeated global and concurrent application of an update rule to each square in the grid. Yet, given specific initial configurations of tokens, it could generate `structures' (or token-patterns), both static and dynamic, of potentially unbounded complexity. Although based on an analogy with the growth, reproduction and death of simple biological organisms such as bacteria, the ideas underlying the Game of Life clearly had implications transcending biology. In particular, there was the possibility that *our* universe might itself be computational in nature, not in the rather unrealistic sense of it being a gigantic Turing machine as had been previously conjectured, but as the physical instantiation of an object belonging to the class of mathematical formalisms to which the Game of Life belonged, viz. cellular automata (CAs). This neo-Laplacian thesis was explored in William Poundstone's *The Recursive Universe* (1985) and has increasingly assumed the status of a literal truth, reappearing in numerous popular works such as *The Mind of God* (Davies,92), *Complexification* (Casti,94) and *Darwin's Dangerous Idea* (Dennett,95).

I had been interested in the problem of the origin of life for a long time and during my reading in evolutionary biology, had stumbled across the work of Graham Cairns- Smith whose "life from clay" hypothesis (Cairns-Smith,71) fascinated me. An abstract reinterpretation of his theory of proto-biotic evolution in terms of self-organization provided me with a suitable context in which to investigate a modification to the Game of Life which I believed would make CAs more realistic when viewed as models of natural phenomena (Ali,94a).

Although cellular automata might have provided me with the behavioural modelling primitive for which I had been searching in the context of the GIMP project, I had at this stage in the project become much more interested in investigating the implications of the view that the universe was a cellular automaton, that is, an atomistic computer. Consequently, I redefined my thesis objectives: I would now attempt to integrate the `sciences of the artificial', disciplines such as artificial intelligence, artificial life, artificial physics and artificial reality, into a unified framework of `artificiality'. This would provide me with the means by which to investigate the limits of computationalism, the metaphysical view that the underlying reality of the world is computational in nature (or being). Inspired by Stuart Kauffman's (1993) argument that natural selection was really only a secondary mechanism operating on the products of self-organizing processes, I began to look for a broader, more encompassing notion than evolution. This I found in an idea related to self-organization, viz. emergence, which provided me with the conceptual basis needed to integrate the various disciplines within artificiality such that the resulting framework was isomorphic with the natural phenomenal hierarchy[8]. The first version of this framework was described in (Ali,94b), which had as its main theme the argument that conventional interpretations of functionalism, of which computationalism is a variant, were flawed. A number of philosophical reasons were advanced in support of the contention that the idea of a function necessitated the existence of an observer, a position related to notions within second-order cybernetics for which I was later to discover support in the work of Searle (1992,1995).

I had been working on integrating the various artificial disciplines using a computationalist interpretation of the emergentist framework presented by Samuel Alexander in *Space, Time and Deity* (1920) when I began teaching on the Foundations of Intelligent Systems course with Mike Elstob. During the lectures, I found it extremely difficult to communicate the phenomenological (specifically, Heideggerian) critique of artificial intelligence (and computationalism in general) that I had been assigned to certain `hard science' students. I also noticed the way students with opposing views often talked past each other rather than to each other and it was then that I became convinced of the importance of clarifying precisely what computationalism meant as well as exposing the philosophical assumptions underlying this concept. Consequently, my study broadened to include concerns of a more metaphysical nature as I began to examine computationalism from an ontological perspective.

Casti (1989) has identified at least four perspectives from which critiques of artificial intelligence have been attempted: (i) logical and mathematical arguments against the view that the mind is a computer program based on interpretations of Gödel's theorems, (ii) aesthetic and ethical arguments based on the implications of such a view for human self-regard, (iii) philosophical arguments based on the irreducibility of mentalistic intentionality, and (iv) philosophical arguments based on the irreducibility of *non*-mentalistic intentionality. Arguments based on Gödel's theorems have been and continue to be the subject of endless controversy (Lucas,61) (Benacerraf,67), most of which is either a continuous restatement of themes or simply polemical in nature. (One need only investigate the subtlety of the points made in discussions about Penrose's reinterpretation of Lucas within the `comp.ai.philosophy` newsgroup to appreciate this fact.) On the other hand, while I generally agreed with the ethical arguments made in works such as Joseph Weizenbaum's *Computer Power and Human Reason* (1984), I was not convinced they provided the foundation for constructing an alternative to the computationalist programme[9]. For this reason, I concentrated on developing a critique of a more philosophical nature, applying arguments previously made by philosophers such as Searle and Dreyfus across the various disciplines within artificiality. I discovered that many of the philosophers were connected to a common tradition within philosophy, viz. phenomenology; Searle with the early work of Edmund Husserl and Dreyfus with the existential interpretation of phenomenology due to Martin Heidegger. Although I gained from both the transcendental and the existential phenomenological perspectives, it was the latter that had by far the greater impact on my thinking. There are at least two reasons for this: firstly, transcendental phenomenology can be shown to provide support

---

8 Assuming, of course, that natural phenomena are hierarchically-structured, irrespective of whether such a hierarchy is ontic (real) or merely epistemic (apparent).

9 Assuming, of course, that an alternative *is* possible.

for the computationalist view, at least in the context of its application within cognitive science and "good old-fashioned" artificial intelligence (Dreyfus,82) (Fodor,80) and this was precisely the position that I was attempting to critique; secondly, it was in Heidegger's writings that I discovered the concept of *poiēsis* (becoming, coming-forth, bringing-forth), a concept that was to provide the basis upon which to establish a phenomenological framework for differentiating natural from artificial phenomena.

I attempted to identify what I believed to be the main contributions of existentialist thought to a critique of artificiality and in this respect was greatly benefitted by Dreyfus' (1991) commentary on Division One of Heidegger's *Being and Time* (1927). As a result of reading this particular work and also Heidegger's *An Introduction to Metaphysics* (1959), I became increasingly aware of the role that categories play in metaphysics. Whether viewed as objectively discernible features of the physical world as in the realist tradition or as features of a transcendental subject as in variants of idealism, categories are the means by which reality is approached. All metaphysics since Aristotle, Heidegger argued, has been concerned with the epistemological use of abstract categories to understand the world and in the process has ignored that which is prior to epistemology, viz. the ontological fact that the world *is*[10]. His response to this condition within philosophy was to develop a set of *concrete* categories - or existent*ials* - for ontological analysis, that is, for investigating what he called the question of Being. However, it should immediately be apparent that irrespective of whether the study is of an ontological or epistemological nature, in both cases the approach is made through the use of categories. They are the means by which we `cut' what Heidegger called Being (exist*ence*) into beings (exist*ents*). Categories are lines of demarcation, coming into existence during what Spencer-Brown (1969) has referred to as the `act of severance' or the `drawing of a distinction', and enable us to bound, constrain and identify as significant that which we hold to be significant. This last point, viz. that significance, and by implication meaning, only belongs to that which is capable of appreciating significance, has deep implications for computationalism in particular and functionalist philosophy in general.

During my investigations into artificial life, I came across a work which discussed what was called computational emergence (Cariani,91). It was argued that as a consequence of the finite nature of the physical substrates in which computational formalisms such as CAs are implemented, the extent to which emergence in such systems is possible is itself bounded; hence, computational systems were only capable of supporting emergence up to some limit. Cariani encapsulated this view in the idea of emergence-relative-to-a-model. In full agreement with this position, I began to investigate extensions to CA that would solve this problem. Whereas Cariani had attempted to solve the problem of infinite or `open-ended' emergence by postulating an external analog continuum from which the discrete symbolic primitives in computational systems were generated by measuring devices (sensors), I chose to remain within the formal computationalist framework. My investigations resulted in the development of an extension to the standard CA formalism incorporating the notion of an epistemological "cut" of reality

---

10 In short, *that* the world exists entails that there *is* something rather than nothing. As will be seen in chapters 1 and 6, Heidegger's approach to addressing the problem of explaining *why* there is something rather than nothing (assuming that this constitutes a legitimate problem) involves *questioning* from the perspective of a historically-embedded existential interpreter. Thus, Heideggerian phenomenology is *essentially hermeneutic*.

and supporting the potential for infinite or open-ended emergence via an infinite hierarchy of rules governing the transitions between CA states (Ali,98a). Although the scheme was well received in session on emergence at the *Tucson II: Toward a Science of Consciousness* conference, I had been having serious doubts about its ultimate significance both during the development of the formalism prior to the conference and upon my return. After all, it was a formalism and hence, an instance of a categorial `cut' itself ! This point requires clarification. It was not the use of representations or models *per se* that was being criticized, although I certainly held certain modelling schemes to be superior to others. Rather, I was attempting to critique (1) the tendency towards *identifying* models with the phenomenon being modelled, (2) the disregard for the modelling *activity* which brought the model into being, and most importantly, (3) that which was responsible for *initiating* the modelling activity in the first instance.

At this stage, I became aware of the implications of a tacit assumption made in the original formulation of the GIMP project, viz. that natural phenomena are *systemic*. Very early on in the project I had been exposed to two opinions regarding the ontological status of systems. One of my supervisors, Robert Zimmer, viewed them as abstractions dependent on the perspective of an external observer, a relativistic position explored in (Ali,94b). However, my other supervisor, Mike Elstob, viewed systems as ontologically primitive, thereby adopting a systems-realist perspective. It became clear to me that irrespective of which position was adopted, a position *had been* and *had to be* adopted, that is, realism and relativism were themselves instances of categorial `cuts'. Inspired by this view, I began to build upon ideas that had first been discussed in the Tucson conference paper. Exposure to numerous ideas presented at this conference had left a deep impact on me. In particular there was what David Chalmers (1996) called the `hard problem' of consciousness, viz. the relation of subjective experience to brain neurophysiology. Consideration of this issue and my interest in the problem of emergence and the `cut' hypothesis led me to investigate Whiteheadian organicism[11] and panexperientialism (Griffin,88), viz. the view that the phenomenal world is internally-related and hierarchically-structured with experiential events (or actual occasions) as ontologically primitive. On this basis, a radical reinterpretation of the notion of emergence grounded in a synthesis of Heideggerian ontology and Whiteheadian metaphysics was formulated in an attempt to solve the mind-body (or `hard') problem (Ali,98b).

At the risk of overstating a point, I would again like to draw attention to Heidegger who has been singularly instrumental in providing me with the basis upon which to develop the ideas presented in this thesis. His account of the relation between the Classical Greek notion of *poiēsis* (coming-forth or bringing-forth) and modern technological `Enframing' (*Gestellen*) in *The Question Concerning Technology* (1977) led to a shift of attention away from semiotics, the science of signs and sign-users, and toward the development of a complementary discipline of *poiētics*, the phenomenological science of Being, beings and becoming. In this thesis I have concentrated on developing a *poiētic* ontological critique of computationalism in contrast to the semiotic epistemological critiques of others such as Cariani (1989), Harnad (1990) and Fetzer (1990). My reasons for doing so are connected with the Heideggerian stance I have adopted and an attempt

11 Ironically, it seems that I was being forced to reconsider a variant of my original biologically-inspired GIMP primitive based on an organicist metaphysics, albeit in post-computational form.

to tip the philosophical balance in favour of considering ontological over epistemological issues.

Intellectual Debts

A number of people have influenced the development of the ideas presented in this thesis and none more significantly than my two doctorate supervisors. Robert Zimmer is responsible for pointing me in the direction of continental philosophy. If it had not been for his advice to read Husserl, I would never have discovered the difficult, yet highly inspiring ideas of Heidegger and this thesis would never have got off the ground. Mike Elstob has been singularly supportive throughout the last four years and to him I owe a great intellectual debt. His relatively few papers on emergence, indeterminism and downwards causation, which attempt to establish the reality of volitional autonomy, made a deep impact on me. Although our views on these issues began to diverge quite rapidly as my study progressed, our exchanges have always been mutually enlightening.

The work of Peter Cariani (1989) deserves special mention for a number of reasons. Cariani was probably the first to draw attention to the constrained nature of emergence within CAs, a consequence of the "closure" of the computational substrates in which such dynamical systems are embedded. CAs are examples of formal systems and can be reduced to essentially two components: (i) a set of primitive symbols and (ii) a set of rules for combining these primitive symbols into symbol structures. Cariani's approach to the problem of how to design devices supporting `open-ended' emergence was to abandon the computationalist approach of attempting to generate novelty through permutations of rules and investigate methods for constructing new symbolic primitives. On his scheme, "open" emergence was supported by introducing analog components into the system in the form of measurement (or sensor) and control (or motor) devices. Through a system of feedback with the external world by which the performance of the system was pragmatically evaluated, the hybrid analog-digital system `evolved' new semantic functions, and, as a corollary, new semantic primitives.

Cariani addressed the problem of semantic emergence from the perspective of a biologist; hence, his endorsement of an evolutionary approach. In his doctorate thesis, he argued in favour of developing a "theory of biological semiotics", the need for which had already been anticipated in the writings of semiotically-inspired biologists such as von Uexküll (Emmeche,94). However, extending semiotics beyond its immediate domain of application, viz. the study of *human* communication, into other domains such as biology and "evolutionary robotics" rests on the assumption that semiotic processes, functionalities, and categories are ubiquitous - or *perfusive* - in nature, an assumption which sets American semiotics in opposition to Anglo-Continental semiology.

Additionally, there is a contradiction in the idea of a "designed emergence" if the emergence supported by such devices (or `artifacts') is understood as an instance of the same phenomenon that (supposedly) occurs in natural systems which are, on the Darwinian view, certainly not products of an intentionalistic design process. Yet, this must be the case if emergent-artificiality is to be interpreted on the "strong" view, that is, as a realization rather than a mere simulation of the phenomenon since this is necessitated by the functionalist basis of the "strong" artificiality programme. For reasons such as these, I strongly disagree with Cariani's view that phenomenology constitutes a

"cognitively impenetrable" realm which does not "explicitly point the way to an alternative research program." On the contrary, I would argue that existential phenomenology, taking its lead from Heidegger, provides us with the possibility of a far more realistic research programme, one which attempts to go beyond categorial thought and point to the source of the categories. Nonetheless, I am indebted to Peter Cariani, whom I have met in `cyberspace' (via the internet), and who I regard as a veritable `comrade in arms' in the intellectual struggle against what we both consider to be the polemical excesses of philosophically naive adherents of computationalism.

I must also recognize the debt that I owe to the authors of four works: Stephen Pepper, author of *World Hypotheses* (1942), for the `root metaphor method' which provided me with an important early tool for philosophical analysis; Samuel Alexander, author of *Space, Time and Deity* (1920), for providing me with the basis of a unifying framework of computationally emergent artificiality; Herbert Simon, author of *The Sciences of The Artificial* (1981), for the first real study of the notion of artificiality *as such*; and Edward Fredkin, author of "Digital Mechanics" (1990) for presenting the notion of a cellular automaton universe as a serious proposition, thereby ensuring that the arguments in this thesis were not directed at a "straw man".

Finally, I would like to acknowledge the support and encouragement that I have derived from precedents set by members of what can only be described at this juncture as `the philosophical underground', the vanguard of what Mike Elstob likes to refer to as "the new paradigm". In particular, David Bohm for his implicate order thesis and process philosophers such as David Ray Griffin and Christian de Quincey for their concise and precise presentation of neo-Whiteheadian panexperientialism.

I should like to end by stating that I think the reader will consider this study to be a rather ambitious undertaking given my former training as an electronic engineer and computer scientist. However, despite a lack of formal philosophical training prior to this study, I believe that critiques such as the one I have attempted are necessary in order to prevent thinking from descending into dogmatic obscurantism. As Mike Elstob perceptively informed me, an attack made from within a tradition is so much more devastating than one made from without. One need only see the difference in impact that the writings of Hubert Dreyfus (a philosopher) and Terry Winograd (a former practitioner within the discipline of AI) have had on members of the latter's research community to appreciate the wisdom and validity of this remark. For this reason, accusations of naive sophistry notwithstanding, I believe the critique presented in this thesis to be *justifiable*. Whether or not it is *justified* is left to the reader to decide. However, in my defense, I should like to cite the following Heideggerian (1959) maxim:

> *When the creators vanish from the nation, when they are barely tolerated as an irrelevant curiousity, an ornament, as eccentrics having nothing to do with real life; when authentic conflict ceases, converted into mere polemics, into the machinations and intrigues of man within the realm of the given, then the decline has set in.*