

# Towards an Architecture for A-life Agents

Darryl DAVIS and Barnaby BERBANK-GREEN

*School of Computing, Staffordshire University, Beaconside, Stafford, U.K., ST18 0AD.*

*e-mail: D.N.Davis@soc.staffs.ac.uk*

**Abstract.** This paper is intended as an interim report on aspects of our research on (intelligent) agents. Our (empirical and design) investigations into complete agent architectures is being driven by the development of a theory of (artificial) mind. Rather than channel all our work into one application area, we are considering a number of different applications in order to further the theory (and its possible computational models). This work ranges from experiments with synthetic worlds, through decision support systems, image analysis and adversarial game playing. It is the latter that drives this paper, and builds on a number of concepts currently being addressed in the research areas of agents and artificial life. Our aim in this specific area is to investigate the types of architecture appropriate for developing an emergent game playing intelligence.

## 1. Introduction

Agents can be defined in many ways and there is no one universally correct or acceptable definition [1]. An agent can be defined (simply) as an autonomous entity that can sense and act upon its environment. Such simple agents are analogous but not equivalent to orthodox automata. More sophisticated descriptions draw on concepts such as intentionality, social ability, adaptation, learning, communication etc. Agent environments vary considerably, from synthetic worlds to those which robots inhabit through to more abstract worlds consisting of information and knowledge. The architecture and composition of an agent typically reflects its environment and the role(s) it plays within that environment; i.e. the challenge of an agent's problem or niche space [2]. Our current research relies heavily on the concept of agency across a number of different domains, and categories of processes, agencies and agents. For example, we can consider our ongoing research into decision support systems [3] as that of an investigation into tightly-coupled [4] agent communities making use of modern but relatively orthodox AI techniques. Again we can consider investigation of (co-operative and competitive) emergence in our simulation work [5] as complimentary to that discussed in research related to artificial societies [6]. This current paper provides an interim report on our ongoing (and relatively recent) investigations into the use of artificial life techniques within an agent framework. One of the threads that draws this research together is that of investigating computational architectures that allow or are designed to support computational intelligence.

We suggest that while no one specific agent architecture may be suited to all environments, our concept of agency should be sufficiently flexible and amenable to variations in needs and capabilities, that it allows us to model agents across a spectra of niche and design spaces. The computational techniques used to implement those designs can be varied to suit the requirements of the environment and the agent(s)'s tasks within those environments, and the constraints imposed by the agent's computational space. For example, it is possible to design (and implement) multi-agent communities based on a range of specialised agents that perform specific functions, or on a number of self-similar

agents that are distinguishable only by their environmental stance and, as a consequence, their internal and external states. A comparison can be made with the world of social insects; for example, on the whole, individual ants are identical to one another, excepting the difference between, for example, the fighter, worker, drone and queen classes. By default, each member of a specific insect class performs the same functions and reacts similarly to the same stimuli. Individual differences between instances of any specific insect class (e.g. worker) result from the interaction of that specific insect and its navigation and resulting tasks within the colony's environment, i.e. that insect's environmental stance. These relatively simple biological agents, at least at an individual level, can communicate, using chemical and other type of signals, and seem to work towards common goals. Colonies display complex social behaviours and problem solving abilities, for example the creation of air conditioning systems in termite nests. Biological science has coined a new term, *Super-organism*, for such systems [7]. Using this metaphor it follows that an agent may itself be an environment within which other agents are at work; we will refer to such systems as macro agents. An example analogous to the insect world is the computer simulation of cities with agents (or actors) inhabiting those cities. Each city is itself a macro-agent within a society of cities. Such societies evolve by means of the interactions of the cities, the constraints the cities place on their inhabitants, the redefinition of the parameters of the cities by the actions of agents (or actors) within those cities and the roles the actors play in mitigating resource distribution within and between cities [6,8]. Such modelling strategies may be appropriate in computationally expensive search spaces as a way of deferring computational trajectories through such spaces to higher and higher levels of abstraction. At that point where no further deferment can be obtained, the computational model may be displaying sufficiently intelligent behaviour that the search space problem need no longer be considered. We hope that the computational cost of this is less than traversing the search space in a more orthodox manner!

## 2. Playing the Game of Go

Weiqi was invented by a Chinese Emperor to teach his son, among other things, concepts of war, strategy and patience. Over thousands of years it has developed and gained great popularity in the Far East, particularly in Japan where it got a new name (Go). The rules appear to be deceptively easy. Players take turns to place their stones on a board at the intersections of horizontal and vertical lines. The intersections immediately to the North, East, South and West of a stone, are its freedom points. If all these freedom points are taken up by opposing stones then the stone is removed from the board. However, stones that are placed on freedom points of their own side are linked and share all of their freedom points. Linked stones are referred to as strings or units. Isolated (*complete*) strings that are totally enclosed by opposing stones (or board edges) can also be removed from the board. A string that encloses two separate areas is safe as the opponent cannot occupy all the strings' freedom points without making a suicidal<sup>§</sup> move. The objective of the game is to gain the greatest amount of territory. Territory is loosely defined as the number of freedom points surrounded by safe stones or strings of the same colour.

A full-size 19by19 Go board permits approximately  $10^{172}$  different board positions in the course of a game. In more complex games there may be as many as  $10^{768}$  possible moves. This presents a real and significant challenge for computational intelligence. In order to simplify the game for computer solution, one could use a smaller board; for example there are approximately  $10^{57}$  positions on a 11by11 board. Outside of the strict rule systems of games, state spaces are, to all intents and purposes, infinite. In such scenarios, the human ability to quickly dismiss unimportant aspects and focus on important

---

<sup>§</sup> If a Black move captures no stones, but leaves a black string with no liberties, the move is suicidal.

(and interesting) features and patterns of the environment distinguishes us from our mechanical and technological masterpieces.

### 3. An Artificial Life Approach To Game Playing

Heeding the results from the interesting work going on elsewhere [6,7,8,9], we thought to tackle the challenge of designing a Go playing agent from an A-life perspective. Instead of directly imposing human models of how to play the game of Go, we chose to use simple agencies, integer-valued Cellular Automata (CA), to represent all the available positions (freedom points) on a Go board. These simple agencies communicate to each other about their state and their local environment. By placing a stone on the board, the state of the board as a whole, at a number of levels of abstraction, is changed. This change is a result of the interaction of all the agencies that change state as a result of the placing of a stone. Part of the communication between agencies can be seen as landscaping the board space. For instance, placing a white stone will add 'height' to that freedom point, and subsequently raise the surrounding 'terrain' by a function of the stone's 'height'. This would create a small 'hill' in the model. A black stone will decrease the 'height', creating a local 'basin' or 'dent'. Areas of zero height (i.e. neither increased nor decreased) can be considered neutral. Quite complex terrains ensue from the interaction of freedom points, stone placement and stone adjacency. Our initial aim was to see if the interaction of the low-level (a-life) agencies enabled an emergent behaviour of benefit to the side of a game-playing agent. We hoped to demonstrate how low level organisation, capable of playing a worthy game, can emerge from the properties of the game itself.

### 4. Interim Results

This first implementation was perceived to produce adequate opening moves but failed to play a satisfactory game (i.e. it repeatedly fails to win). In adversarial game playing, there are a number of possible levels of analysis, and types of play, ranging from low-level (almost haphazard) moves to high-level strategic ploys. While we may perceive any sequence of moves as saving, aggressive, defensive or territorial, the emergence of such strategic concepts from simplistic artificial life agencies would be unexpected.

We then considered a more complex agent consisting of multiple terrain models to further explore the limits of this approach. We wondered if a more effective and tactical game-playing agency would emerge through increasing the complexity of the low-level processes within the agent. We devised CA representations that modelled a number of boards to help the agent in determining its next move. One board represented default territory classifications, for example the recognised *fuseki*<sup>§</sup> positions ([10] gives clear overview of this stage of Go). A number of alternative liberty models plus maps that relate solely to the impact of placing a stone on the board (hill or basin forming) were investigated along with hybrid ( $\Delta$ CA) maps that combine some of these models. Figure 1 shows an example of a game that the agent system (playing white) won. In this game five separate boards were kept: the playing surface (top right of figure 1), a fuseki model (not shown), a freedom point or liberty board model (top left of figure 1), a terrain model (middle top of figure 1), and a decision board (bottom of figure 1). The agent system can play white or black (or both), and decides upon its next move using the decision board. If it is playing black it looks for the deepest terrain (i.e. most negative  $\Delta$ CA position), and for the highest terrain (most positive  $\Delta$ CA position) if playing white. As the decision board may have multiple best potential moves, some simple rules are used to resolve any conflict. These rules could be replaced or extended with further A-life mechanisms as discussed in

---

<sup>§</sup> Fuseki describes the opening moves of the game (usually defined as lasting until the first fight begins).

the following paragraph. This relatively naïve approach to modelling the game displays emergent behaviour of benefit to the agent: good opening stone placements, the building of non-empty units, i.e. geometrical arrangements of stones that contain internal freedom points, stone and even string capture.

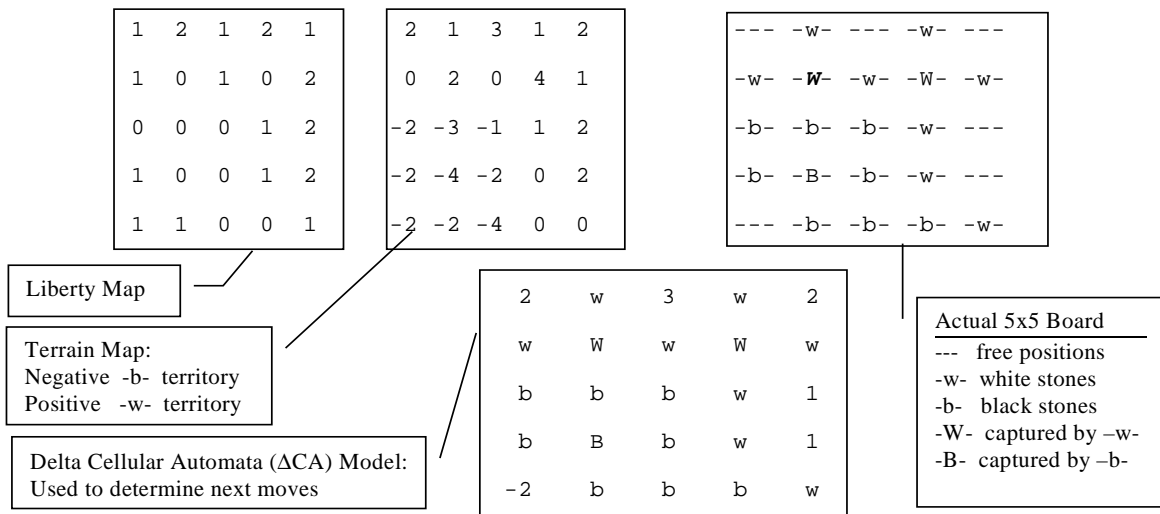


Figure 1. Multiple Terrain CA model– White (the computer) about to win on a highly simplified board.

This system with its multiple CA-based strategies played the game more capably than the first system and even demonstrated some limited success in attacking isolated opposing strings. A number of further improvements could be made to improve its game playing. For example the terrain model should evolve with the game and relate to the recognised different stages of a good game of Go. However the multiple board model mostly failed to demonstrate more sophisticated game playing strategies. Such strategic ploys, while grounded in specific sequences of moves which vary from game to game and within a game as the nature of the current game changes, are qualities that may require modelling as more persistent, global processes. Other work on strategies as emergent properties of tactical behaviours seems to be similarly inconclusive [11]. If we were to continue to adopt a purist a-life approach, more complex computational metaphors are required. For example, we can consider strings of stones to be a kind of multi-cellular agent, for example an irregular coelenterata. Cellular automata representing single stones and empty spaces then become contributory elements of these macro-agents which need to grow (i.e. extend their territory), breathe (extend whilst enclosing empty board positions), eat (extend whilst capturing opposing stones) and meet (extend towards other strings of the same colour). Rather than reproduce in an orthodox sense, such macro-agents become like cities that merge or coalesce.

Another possibility, but again requiring considerable research, is the design sketched in figure 2. This agent architecture combines multiple communities of a-life agents, neural nets and simple reasoning over propositions. In this type of architecture, there exist many competing and co-operating a-life communities. The appropriateness of the actions suggested by any of the base level agents can be measured in terms of the needs of the individual and collective communities; i.e. what is good for the micro and macro-agent. The agencies responsible for determining this can be viewed as a synergistic harness of a-life mechanisms, neural networks and say propositional logic. Some neural networks may be trained to recognise board patterns suited to attacking moves, others to defensive or space utilisation gambits. Such communications paths provide not only a-life driven high level processes, but in a dynamic architecture, where the active combination of agencies shifts as the macro-agent's problem space changes throughout a game, may provide a

means for harnessing (potentially) positive emergent qualities and stabilising disruptive emergent behaviours. The following section discusses some of these issues.

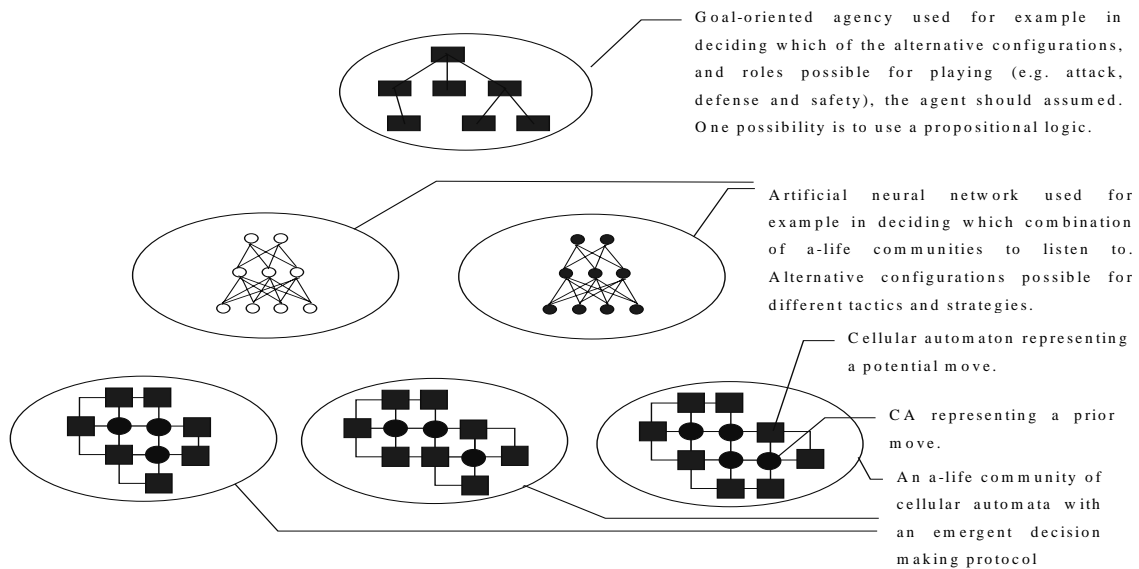


Figure 2. One (of many) possible architectures for a-life based macro agents.

## 5. Emergence, Architectures and Agents

The readings in [6] identify four levels or types of emergence: diachronic emergence, gestalt emergence, representational emergence, and adaptive or functional emergence (the most commonly referenced type); and moreover three categories of problems associated with emergence: sub-cognitive bias, behavioural bias, and individualistic bias. Diachronic emergence relates to evolutionary traits and is inapplicable to our current design and computational models. Gestalt emergence, however, seems to describe our aim of developing multiple a-life models that relate to the different levels of abstraction that we can place on the game of Go. Adaptive emergence is clearly applicable to the current scenario and requires that our game playing agent, modelled using multiple agencies, contains appropriate communication and feedback mechanisms. Of the different types of sub-cognitive bias, our assumption that collective (and co-operative) activity will emerge opportunistically is perhaps our most immediate concern. A consideration of individualistic bias suggests that adaptive emergence needs both feedback mechanisms from the micro (individual) level to the macro (societal) level and from the societal level to the individual level (a kind of focus of attention in a macro-agent). For instance, an individual cell or cell matrix may suggest that a move at a specific location is beneficial. If such information is communicated to other more abstract computational models, it needs to be represented and communicated appropriately at those levels, and then enable feedback to other simpler a-life processes. In such a way, a strategy may emerge from the adoption of a set of tactical moves modelled computationally as the currently attended and therefore active a-life processes.

The architectural and design impact of these sometimes beneficial states is worth considering. As a general framework for computational and cognitive intelligence we [2, 5, 12] have been developing a four layer agent architecture that includes reflexive impulses, reactive processing, deliberative processing and meta-level (*reflective*) processes. Such a complete agent architecture may be able to manage perturbant patterns of processing and/or focus the processes associated with an agent (across the first three layers) to specific categories of tasks. A macro agent's need to monitor emergent patterns of behaviour, to classify them as (potentially) useful or disruptive and manage them is a clearly analogous situation. An underlying assumption here is that perturbant processes do not emerge at the

reflective level. One major challenge, which future work should address, is how does such a macro agent recognise and harness these emergent behaviours, particularly where certain categories of behaviour are perceived to be useful in some but not all scenarios. The architecture sketched in figure 2 contains a-life reflexive processes, connectionist reactive processes and one deliberative agency. As we further develop this architecture, we will need to consider what are appropriate computational models for the reflective level. We have thought about using Hopfield nets but then we may risk undermining the assumption that agencies capable of displaying emergent patterns (in need of further monitoring) are inadvisable at the reflective level. If we consider the work described in [6, 9] and many of the ongoing challenges that AI and computing in general faces, it becomes apparent that no currently known computational process is immune from this charge!

## 6. Future Work

We have made some progress in developing an a-life-based agency capable of playing the game of Go. While our initial approaches show some promise, we have realised that to fully harness the desirable properties of A-life mechanisms we will need to investigate some quite sophisticated agent architectures. The emergent and beneficial complexity of behaviour arising from interacting simply modelled agents can be best utilised in agents whose computational architecture allows them to recognise and utilise the emergent beneficial states and resolve more disruptive emergent patterns. The suggestion is that the utilisation of the social metaphors appropriate to insect and the simulation of artificial societies may be appropriate.

Go presents us with a great computational challenge about which the agent and a-life metaphors offer a fresh perspective. We suggest that by combining such techniques and developing a methodology for creating such agents, we can make progress towards a computational intelligence of use in wider domains.

## 7. References

- [1] S.P. Franklin and A.G. Graesser, Is It an Agent, or Just a Program?: A Taxonomy for Autonomous Agents, In: J.P. Muller, M.J. Wooldridge & N.R. Jennings (Eds.) *Intelligent Agents III*, ISBN: 3-540-62507-0, Springer-Verlag, Heidelberg, 1996.
- [2] A. Sloman, Explorations in design space. In: *Proceedings 11th European Conference on AI*, Amsterdam, 1994.
- [3] D.N. Davis and B. Sharp., An Agent Framework for Decision Support in the Water Industry, *New Review of Applied Expert Systems*, (Under Review). ISSN: 1361-0244, Taylor Graham Publishing, 1998.
- [4] L. R. M edsker, Hybrid Intelligent System s, ISBN 0-7923-9588-3, K luw er A cadem ic Publishers, 1995 .
- [5] D . N . Davis, Reactive and motivational agents, In: Muller, M.J. Wooldridge & N.R. Jennings (Eds.) *Intelligent Agents III*, ISBN: 3-540-62507-0, Springer-V erlag, 1997 .
- [6] N . G ilbert and R . C onte, Artificial Societies: The com puter sim ulation of social life, UCL Press, ISBN 1-85728-305-8, 1995 .
- [7] S. Levy, *Artificial Life*. ISBN: 0224035991, Jonathan Cape, 1992.
- [8] R . H egselm ann and A . F lache, Understanding complex social dynam ics: A plea for cellular autom ata based m odelling. *Journal of Artificial Societies and Social Sim ulation*, vol. 1, ISSN : 1460-7425, 1998 .
- [9] G. Weisbuch, *Complex Systems Dynamics*, ISSN: 0-201-52887-8, Addison-Wesley, 1991.
- [10] O. Hideo 9 dan, *Opening Theory made Easy*, ISBN: 4-87187-036-7, Ishi Press, 1992.
- [11] A . Dragoul, When Ants Play Chess (Or Can Strategies Emerge from Tactical Behaviours). In: Castelfranchi, C , & M uller, J.P, *From Reaction to Cognition*, ISBN : 3-540-60155-4, Verlag, 1995 .
- [12] D.N. Davis, Synthetic Agents: Synthetic Minds?, *Frontiers in Cognitive Agents, IEEE International Conference on Systems, Man, and Cybernetics*, ISBN 0-7803-4781-1, IEEE Press, 1998.