Modularity and Specialized Learning in the Organization of Behaviour

Joanna Bryson and Lynn Andrea Stein[1]

Artificial Intelligence Laboratory, MIT
545 Technology Square, Cambridge MA 02139, USA
joanna@ai.mit.edu and las@ai.mit.edu

Abstract. Research in artificial neural networks (ANN) has provided new insights for psychologists, particularly in the areas of memory, perception, representation and learning. However, the types and levels of psychological modelling possible in artificial neural systems is limited by the current state of the technology. This chapter discusses modularity as illuminated from research in *complete agents*, such as autonomous robots or virtual reality characters. We describe the sorts of modularity that have been found useful in agent research. We then consider the issues involved in modelling such systems neurally, particularly with respect to the implications of this work for learning and development. We conclude that such a system would be highly desirable, but currently poses serious technical challenges to the field of ANN. We propose that in the mean time, psychologists may want to consider modelling learning in specialised hybrid systems which can support both complex behaviour and neural learning.

1. Introduction

Research in artificial neural networks (ANN) has provided new insights for psychologists, particularly in the areas of memory, perception, representation and learning. However, the types and levels of psychological modelling possible in artificial neural systems is limited by the current state of the technology. We are only beginning to be able to model neurally many of the complex structures and interactions we know to exist in animal brains. In particular, modularity is a much-discussed feature of the brain, but we have only rudimentary models of it in neural networks.

This chapter discusses modularity as illuminated by research in *complete agents*, such as autonomous robots or virtual reality characters. We use this research to identify the sorts of modules that have been found useful in our field. We then describe a possible mapping between these established AI modules and the modularity present in mammalian brains. We conclude with a discussion of the implications of this work for learning, development, and the future of ANN and agent models of animal intelligence.

[1] LAS: also Computers and Cognition Group, Franklin W. Olin College of Engineering, 1735 Great Plain Avenue, Needham, MA 02492 las@olin.edu

2. Modularity in AI and Psychology

The extensive use of modularity in complete agents was popularised in the mid 1980's with the establishment of *behaviour based artificial intelligence* (BBAI). Behaviour-based AI refers to an approach inspired by Minsky [22], where many small, relatively simple elements of intelligence act in parallel, each handling its own area of expertise [3, 19]. In theory, these component elements are both easier to design and more plausible to have evolved than a single complex monolithic system to govern all of behaviour. In the modular theory of intelligence, the apparent complexity of intelligent behaviour arises from two sources: the interactions between multiple units running in parallel, and the inherent complexity of the environment the individual units are reacting to.

The behaviour-based approach generated significant advances in mobile robotics [4] and has come to dominate both the fields of robotics and virtual reality [17, 27]. It has lead to a revolution in the way computation is thought about [29]. Nevertheless, it has not been entirely successful. Advances in the development of humanoid agents still come disappointingly slowly. Further, there is no single dominant behaviour-based architecture that is used by even a large percentage, let alone a majority, of complete agent developers.

The engineering advantage of modularity is simple: it decomposes the problem of intelligence into manageable chunks. In essence, it is a design advantage. Modularity is a form of hierarchy. Each module handles a portion of the agent's overall problem space, leaving the complete agent with an exponentially reduced space of behaviour options to consider. However, the modular approach also introduces a number of problems. First, there is the question of how to decompose an apparently coherent intelligence into modules. What state and/or behaviour belongs together, and what apart? Even more problematic, once behaviours have been separated into at least semi-autonomous modules, how can overall behavioral coherence be reestablished? A modular approach is of no advantage if the problem of integrating behaviour leads to a greater problem of design than the decomposition originally avoided.

From a psychological perspective, interest in modularity dates back to Freud [12], and even to Hume [16]. A modular architecture is motivated not only by the inconsistencies of human behaviour, but also by neuroscience, which has shown a diversity of organs in the human brain. Here the same questions emerge: how is the brain modularised? To what extent are the modules encapsulated — that is, how strictly are they separated? Which modules communicate, and at what level of abstraction? What functions do the various organs perform? And how are their parallel operations coordinated into fluid, largely coherent behaviour? In the following sections, we hope to provide answers or at least hypotheses for these questions based on our work in complete agent architectures.

3. Modularity in Complete Agent Architectures

In this section we describe the sorts of modules that are requisite for making a complete agent function. We also show how these required modules relate to animal intelligence. To begin though, we delineate two different sorts of modularity: architectural modules vs. skill modules.

3.1 Types of Modularity

In thinking about the organisation of agent behaviour, we must consider two different sorts of modular decomposition. One is decomposition of generic function, where modules might include planning, vision, motor control, and such-like. The second is decomposition by task, where modules might include walking, sleeping, hunting, grooming and so forth. Initially, the behaviour-based revolution was about moving from the former, generic sort of modularity to the latter, task-oriented sort of modularity. Although the major benefits of BBAI come from the special-purpose, task-oriented sort of modularity, the experience of the last 15 years has shown that some generic modules are also necessary in an architecture in order for that architecture to be useful and usable.

One of our hypotheses (described further below) is that this more generic modularity is analogous to the modularity by organ in the brain, e.g. the roles of the cerebellum or the hippocampus. We call this sort of modules *architectural modules*, both because in the brain it often characterised by different underlying neurological architectures, and because in agent organisation it contributes architectural features. The task-specific modules we hypothesise to be more analogous to the within-organ modularity exhibited in the brain, by neural assemblies differentiated only by physical space and connectivity. We will call this sort of modules *skill modules.* The emphasis in this chapter will be primarily on different sorts of architectural modules, since this is the emphasis of our own research. The reader should be aware, however, that most work in BBAI concentrates on the development of skill modules.

3.2 Architectural Modularity in Complete Agents

Bryson [5] describes the emergence over the last 15 years of three sorts of architectural requirements for complete autonomous agents capable of complex, scalable behaviour. That article reviews the literature in four separate agent traditions: BBAI; the hybrid or multi-layer community, which combines behaviours with more conventional planning; the Procedural Reasoning System (PRS) / Beliefs, Desires and Intentions (BDI) communities, which have adapted conventional planning and representation in response to the success of BBAI and the demands of real-time, embedded systems; and the Soar / ACT-R communities, which have been working with their distributed representation cognitive agent architectures for many years. In this chapter, we will not reiterate that review, but will only describe the results.

Briefly, there are three sorts of architectural modules that seem necessary. First, nearly all autonomous agents consist primarily of a system of skill modules. These are often referred to as *behaviours* due to BBAI. Despite their name, “behaviours” actually *generate* behaviour; they are not a description of expressed behaviour. There is no one-to-one correspondence between a behaviour module and an expressed behaviour. Much of expressed behaviour is supposed to *emerge* from the interaction of two or more skill-module behaviours operating at the same time. The second sort of architectural module performs *action selection*, which provides for coherent behaviour in a distributed system. Finally, because action selection generally works by focusing attention on one subset of possible behaviour, there needs also to be a dedicated *environment monitoring* module. This module switches the action-selection attention in response to salient environmental events.

Skill Modules

Nearly all autonomous agent architectures today use a modular skill structure. They incorporate a set of primitive behaviours that no other part of the system needs to understand the workings of. This reflects the combinatorial advantage of modularity referred to in the introduction. It comes at a cost of less fine-grained control and the inefficiencies of not being able to combine the outputs of related actions or motions. But for a resource-constrained agent working in a dynamic, real-time environment where responses must be quick and appropriate, the advantage of being able to quickly activate pre-compiled skills outweighs these sorts of costs.

Some researchers try to reduce some of the complexity of coordinating skill modules by using homogeneous behaviour representations or coordinated output formats for all the modules (see for examples Arkin [1] and Tyrrell [30] respectively). We feel these strategies overconstrain the sorts of computation and representation in a skill module. Skill modules should contain not only motor actions, but also whatever perceptual skills are necessary in order to support them. Perception is not just sensing. Sensory stimuli are often ambiguous and often require both recent context and longer-term experience and expectations to discriminate. The state needed to learn or tune skills, or to disambiguate or categorise the perception, should also be a part of the skill module. Actions may also vary significantly depending on context, but a designer may want to encapsulate all these expressions as a single behaviour for simplicity. For this reason, we prefer using objects specified in object-oriented programming languages to represent skill modules.

This representation of skill modules has advantages for two reasons. From an engineering perspective, it allows for the behaviour decomposition problem to be addressed with techniques similar to those developed for object decomposition (e.g [9]) and allows us to quickly develop more powerful, complex behaviour [6]. From a cognitive modelling perspective, this level of abstraction is more closely analogous to the level of description usually used for ascribing functions to cortical areas. Although early BBAI eschewed variable state, in general the notion of empowered, perceiving semi-autonomous behaviours is more in keeping with this work than the trend in hybrid and BDI architectures to making the skill modules into simple motor primitives. It is also very like the recent contributions of the multi-agent systems (MAS) community [31].

Action Selection

The next architectural module is hierarchically structured plans for action selection. “Action selection” is the term applied to the ongoing problem of determining exactly what an agent is going to do next. The BBAI literature often refers to this problem as “behaviour arbitration”, in MAS it is called “agent coordination”. When all behaviours are running in parallel, each may have an action it is currently attempting to express. If these conflict, some method of coordination must maintain the coherence of the complete agent’s behaviour.

Plans in the context of agent action selection are structures which indicate, given a particular environmental context and decision history, what to do next. BBAI originally strongly resisted the use of plan-like structures for coordination, because they were felt to lose most of the advantages of modularity by imposing a form of centralised control. However, these systems proved useful not only for coordination, but also as simple memory. A system using plans to coordinate behaviour does not need complex or perfectly ordered episodic memory in order to disambiguate where it is in a multi-step plan, even if the consequences of its previous actions are not present in its environment.

For example, consider, a robot tidying an office. If it has a fixed order for tidying drawers it neither needs to repeatedly inspect closed drawers it had already tidied, nor to remember the complete list of drawers already visited each time it selects the next one. It only needs to remember its agenda, and its current place in its agenda. Notice the agenda does not need to specify any details about how a drawer is cleaned. The drawer cleaning behaviour is free to chose a methodology appropriate to its perceptions, or to determine whether a drawer is already "clean enough."

Although such structures for supporting action selection are referred to as *plans*, this does not necessarily imply that they are created by *planning*, at least not by the agent. Most complete agents exploit plans provided by their designers, rather than engaging in the slow and unreliable constructive planning process. Choosing what to do next this way, without deliberation, is often called *reactive planning*. When plan-like structures are provided to facilitate this process, they are referred to as *reactive plans*.

Environment Monitoring

Hierarchical structures for action selection allow for the focus of attention on a particular set of behaviours that are likely to be applicable in a particular circumstance. The hierarchies are parsed by the recognition of circumstance, and can be reparsed arbitrarily frequently in order to ensure the applicability of current behaviour. This observation leads to the third architectural module found in all successful agent architectures: a mechanism for monitoring the environment and realizing that the agent should attend to new goals. This sort of system is sometimes referred to as an *alarm system* (e.g. [28]). Having a parallel system is necessitated by hierarchical control: any agent in a dynamic environment that may have its attention focussed on a particular task needs a system to guarantee that it notices salient events in the environment. Otherwise, the agent may overlook both dangers and opportunities. The alarm system must necessarily be relatively simple, requiring no cognitive overhead that would distract from the primary task. It is usually a case of pattern matching, of recognising salient indicators in the environment and then switching cognitive attention to analysing and coping with such situations.

3.3 Equivalent Modularity in Animal Brains

In summary of the last section, complete agent architectures have converged on three sorts of architectural modules for supporting complex, dynamic behaviour. These are a system of skill modules, structured action selection, and an environment monitoring system. If such an organisation is necessary or at least very useful for controlling intelligent agents, then it is also likely to be present in animals, since they have evolved to face similar problems of information management. In this section we speculate about what the analogous modules might be.

First, the skill modules we believe are reflected in within-organ modularity in the brain. This is not to imply that they are represented only within one organ or region: many skills necessarily combine vertically a large number of brain areas and organs, for example the retinas, visual cortex, associative cortices, motor pre-planning and motor coordination. To some extent, each of these areas is architecturally specialised, but the combination of a particular ensemble of simultaneous activations across these structures might model the various skill sets modularised in agent architectures.

Next, Prescott et al. [24] postulate that the basal ganglia is the architectural module responsible for action selection (see also [21]). These researchers focus primarily on the problem of action selection as behaviour arbitration, but the basal system is also well integrated with some of the mid-brain systems that have been implicated in species-typical action patterns as well as the cortical systems which might hold the perceptual skill modules required to discriminate context.

Finally, the behaviour of the alert system is analogous to behaviour attributed to the limbic system. In particular, the amygdala has been implicated in learning to recognise and attend to salient situations [8]. This attention takes the form of emotional responses, which are characterised by the selective activation of appropriate behaviours and expectations.

Thus it is plausible, if far from proven, that the sorts of modularity used in complete autonomous agents are also present in animals and humans. This reenforces our proposal that the agent platform could be useful for psychological modelling.

4. Learning in Behaviour Oriented Architectures

A particular modular organisation provides structure not only for how an intelligence operates in particular situations, but also for how it can learn. We assume that learning happens *within* modules, not across or outside them. So, for the organisation of modules described in the previous section, there are only certain places learning can take place. There can be specialised or perceptual learning within the skill modules. There can be learning of new skill modules within the architectural module of skill modules. There can be the learning of new reactive plans within the action selection module, and there can be learning of selection rules for priorities and alarms. Notice that the sorts of perceptual learning that takes place in skill-modules can also affect the execution of reactive plans and of attention switching. In our model, learning new categories, both new discriminations and new generalisations, is the sort of perceptual learning that should take place in a skill module.

Not all complete agent architectures perform learning in this way. Some which are more closely aligned traditional planning maintain a single database of “beliefs” [15]. Production-based agents [18] often have both a database of beliefs and a database of *productions*: if-then rules governing intelligent control. These productions, analogous to behaviours in their modularity, are not as complex as the sorts of skill modules we have been describing. For the remainder of this chapter, we will be emphasising complete agent architectures that are *behaviour oriented*. Behaviour-oriented systems differ from being strictly behaviour-based because they have structured action selection, but are

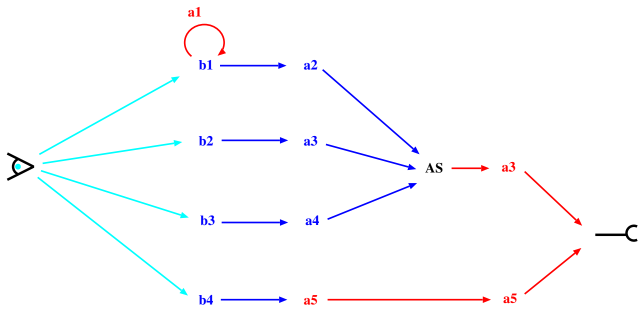

Fig. 1. Behaviour-oriented systems have multiple, semi-autonomous skill modules or *behaviours* ($b_1 \ldots$) which generate actions ($a_1 \ldots$) based on their own perception (indicated by the eye icon). Actions which affect state outside their generating behaviour, whether internal to the agent or external, are generally subject to arbitration by an action selection (AS) system.

otherwise near to traditional BBAI systems in their emphasis on relatively autonomous behaviours capable of perception and action (see Figure 1).

Agents tend to be created for particular environments and tasks; in other words they are niche specific. Consequently, the kinds of things they are likely to need to learn can generally be determined in advance. This is a currently popular view of animal learning [13, 25], and we have made it central to our methodology for developing behaviour oriented agents, which we call behaviour oriented design (BOD) [6]. We determine the modular decomposition for a set of skills around the kinds of perception and memory those skills need to operate appropriately. Thus each skill module contains specialised representations and perception and action routines for maintaining those representations, as well as control for the skilled actions that the agent applies to its environment.

The second most common form of learning in complete agents is the learning of new reactive plans. This is usually done in one of two ways, either by reasoning or planning a new plan (e.g. [2, 14]), or more recently by social learning such as imitation or receiving instruction (e.g. [10, 26]). However, in practice, a surprisingly large number of agents use plans programmed by their designers. So far, this is still the fastest and most reliable way to get appropriate behaviour from an agent.

5. Learning New Behaviours

In our earlier discussion of skill modules, we claimed that behaviour oriented design requires the use of complex algorithms and specialised representations, and that modules are therefore better represented in object oriented languages than in current ANN. However, there is at least one reason to favour an ANN representation of skill modules. That is the problem of designing an agent that can learn or develop *new* skill modules. This is clearly desirable, and has been the focus of significant research (see [10] for a recent example and review.) However, to date, most efforts on these lines would qualify as specialised learning *within* a single skill module / representation system from the perspective of behaviour oriented design.

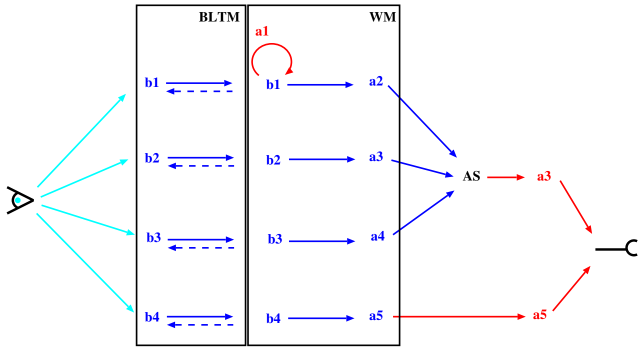

Fig. 2. A system capable of learning behaviours must 1) represent them on a common substrate and 2) allow them to be modified. Here behaviours are represented in a special long term memory (BLTM) *and* in a plastic working memory (WM) where they can be modified. During consolidation (dashed lines) modifications may either alter the original behaviours or create new ones.

The reason that we would like to be able to represent behaviours in terms of ANN is as follows. Consider Figure 2. In this figure, representation of the skill modules has been split into two functional modules: the Behaviour Long Term Memory (BLTM) and the Working Memory (WM). The working memory allows for rapid, short term changes not only for perceptual memory, but also in the representation of the behaviours. The BLTM provides a relatively stable reference source for how these modules should appear when activated. Skill representations might be modified due to particular circumstances, such as compensating for tiredness or high wind, or responding to a novel situation such as using chopsticks on slippery rice noodles for the first time.

In this model, the adjustments made in plastic, short term memory also affect the long term memory. This sort of dual- or multi-rate learning is receiving a good deal of attention in ANN currently (see [7, 11, 20]). Depending on long term experience, we would like this consolidation to have two possible effects. Let's imagine that b_2 has been modified in working memory in order to provide an appropriate expression of a_2. If the same modifications of b_2 prove useful in the near future, then they will be present for consolidation for a protracted period, and likely to effect the permanent representation of b_2. However, if the modifications are only sometimes applicable, we would like a new behaviour $b_2{}'$ to become established. This process should also trigger perceptual learning, so that the two behaviours can discriminate their appropriate context for the purpose of action selection. Also b_2 and $b_2{}'$ would now be free to further specialise away from each other.

6. Conclusions and Future Directions

In this chapter we have shown that complete agents can provide higher-level AI models for psychologically important concepts such as modularity and specialised learning. We have also described the sorts of systems that are the current state-of-the-art in complete agent architectures. We have shown that using ANN to represent at least some parts of a complete agent might be highly desirable, but unfortunately we have also argued that the complexity of the algorithms and the specialised representations are not yet met in current models of ANN modularity (e.g. [23]).

On the other hand, we believe that combining research in these two fields might be a highly useful direction. Complete agent researchers are already experimenting with ANN for learning and simple control, but most such systems are not yet advanced enough to be interesting to psychologists. However, the psychologically interesting ANN systems currently under development such as those described in this volume might be furthered by embedding them in a complete agent. In this way, conventional programming can be used to provide for appropriate experimental platforms to test such systems. A complete agent can provide realistic inputs, and test outputs in realistic settings. Thus, we hope to see a further uniting of these two fields in the future.

Acknowledgements

Thanks to David Glasspool for his comments and suggestions.

References

[1] Ronald C. Arkin. *Behavior-Based Robotics*. MIT Press, Cambridge, MA, 1998.

[2] R. P. Bonasso, R. J. Firby, E. Gat, D. Kortenkamp, D. P. Miller, and M. G. Slack. Experiences with an architecture for intelligent, reactive agents. *Journal of Experimental & Theoretical Artificial Intelligence*, 9(2/3):237–256, 1997.

[3] Rodney A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47: 139–159, 1991.

[4] Rodney A. Brooks and Lynn Andrea Stein. Building brains for bodies. Memo 1439, Massachusetts Institute of TechnologyArtificial Intelligence Lab, Cambridge, MA, August 1993.

[5] Joanna Bryson. Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, 12(2):165–190, 2000.

[6] Joanna Bryson. Making modularity work: Combining memory systems and intelligent processes in a dialog agent. In Aaron Sloman, editor, *AISB'00 Symposium on Designing a Functioning Mind*, 2000.

[7] John Bullinaria. Exploring the Baldwin effect in evolving adaptable control systems. In this volume. Springer, 2000.

[8] Niel R. Carlson. *Physiology of Behavior*. Allyn and Bacon, Boston, 2000.

[9] Peter Coad, David North, and Mark Mayfield. *Object Models: Strategies, Patterns and Applications*. Prentice Hall, 2nd edition, 1997.

[10] John Demiris and Gillian Hayes. Active and passive routes to imitation. In *Proceedings of the AISB'99 Symposium on Imitation in Animals and Artifacts*, Edinburgh, April 1999. AISB.

[11] Robert French, Bernard Ans, and Stephane Rousset. Pseudopatterns and dual-network memory models: Advantages and shortcomings. In this volume. Springer, 2000.

[12] Sigmund Freud. *The Interpretation of Dreams*. Avon, New York, 1900.

[13] C. R. Gallistel. *The Organization of Learning*. MIT Press / Bradford Books, Cambridge, MA, 1990.

[14] Erann Gat. *Reliable Goal-Directed Reactive Control of Autonomous Mobile Robots*. PhD thesis, Virginia Polytechnic Institute and State University, 1991.

[15] M. P. Georgeff and A. L. Lansky. Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, pages 677–682, Seattle, WA, 1987.

[16] David Hume. *Philisophical Essays Concerning Human Understanding*. 1748.

[17] David Kortenkamp, R. Peter Bonasso, and Robin Murphy, editors. *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*. MIT Press, Cambridge, MA, 1998.

[18] John E. Laird and Paul S. Rosenbloom. The evolution of the Soar cognitive architecture. Technical Report CSE-TR-219-94, Department of EE & CS, University of Michigan, Ann Arbor, September 1994. also in *Mind Matters*, Steier and Mitchell, eds.

[19] Maja J. Matarić. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental & Theoretical Artificial Intelligence*, 9 (2/3):323–336, 1997.

[20] James L. McClelland, Bruce L. McNaughton, and Randall C. O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, 1995.

[21] Jonathon W. Mink. The basal ganglia: focused selection and inhibition of competing motor programs. *Progress In Neurobiology*, 50(4):381–425, 1996.

[22] Marvin Minsky. *The Society of Mind*. Simon and Schuster, New York, NY, 1986.

[23] David C. Plaut. Systematicity and specialization in semantics. In D. Heinke, G. W. Humphreys, and A. Olson, editors, *Connectionist models in cognitive neuroscience: Proceedings of the Fifth Annual Neural Computation and Psychology Workshop*, New York, 1999. Springer-Verlag.

[24] Tony J. Prescott, Kevin Gurney, F. Montes Gonzalez, and Peter Redgrave. The evolution of action selection. In David McFarland and O. Holland, editors, *Towards the Whole Iguana*. MIT Press, Cambridge, MA, to appear.

[25] T. J. Roper. Learning as a biological phenomena. In T. R. Halliday and P. J. B. Slater, editors, *Genes, Development and Learning*, volume 3 of *Animal Behaviour*, chapter 6, pages 178–212. Blackwell Scientific Publications, Oxford, 1983.

[26] Stefan Schaal. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(6):233–242, 1999.

[27] Phoebe Sengers. *Anti-Boxology: Agent Design in Cultural Context*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1999.

[28] Aaron Sloman. Models of models of mind. In Aaron Sloman, editor, *AISB'00 Symposium on Designing a Functioning Mind*, 2000.

[29] L. A. Stein. Challenging the computational metaphor: Implications for how we think. *Cybernetics and Systems*, 30(6):473–507, 1999.

[30] Toby Tyrrell. *Computational Mechanisms for Action Selection*. PhD thesis, University of Edinburgh, 1993. Centre for Cognitive Science.

[31] Gerhard Weiss, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, Cambridge, MA, 1999.