Evolution, Adaption, and Behavioural Holism in Artificial Intelligence

T. S. Dahl[1] and C. Giraud-Carrier[2]

[1] Machine Learning Research Group, Department of Computer Science, University of Bristol, United Kingdom, tdahl@cs.bris.ac.uk
[2] ELCA Informatique SA, Lausanne 13, Switzerland, cgc@cs.elca.ch

Abstract. This paper presents work on reproducing complex forms of animal learning in simulated Khepera robots using a behaviour-based approach. The work differs from existing behaviour-based approaches by implementing a path of hypothetical evolutionary steps rather than using automated evolutionary development techniques or directly implementing sophisticated learning. Following a step-wise approach has made us realise the importance of maximising the number of behaviours and activities included on one level of complexity before progressing to more sophisticated solutions. We call this inclusion behavioural holism and argue that successful approaches to complex behaviour based robotics must be both step-wise and holistic.

1 Introduction

After abandoning traditional top-down machine learning (ML) methods for learning novel robot behaviours because of their inherent difficulties with expressing learning biases and background knowledge, we have taken a behaviour-based (BB) [2] approach where complex learning behaviours are implemented through step-wise increases and modifications to already adaptive behavioural foundations.

In order to conduct a number of experiments implementing basic adaptive animal behaviours, we developed a framework of programmable, learning, artificial neural circuits (PLANCS) [8]. It provides a neural circuit class which emulates an independent computational node. The neural circuit abstraction is also a neuron-inspired extension of an object oriented BB architecture called Edmund [5], that supports circuit level cognitive modelling.

During this work we have developed two guidelines that extend traditional BB robotics with respect to developing complex behaviours, in particular learning behaviours. We call the development approach described by The guidelines a *step-wise* and *behaviourally holistic* approach. The terms *step-wise* and *behaviourally holistic* are further described in Section 2. Two of the experiments we conducted on reproducing animal behaviours are presented in Section 3. The experiments are done using the Webots Khepera robot simulator [7]. Section 4 suggests how an analysis of human evolution can be used to provide a road-map for step-wise, holistic, BB robotics. Finally, Section 5 discusses the relationship between the holistic and the ALife approaches to robotics, and answers some common criticisms of our work.

2 A Step-Wise, Holistic Approach to Robotics

2.1 Defining the Approach

Our step-wise, holistic approach to implementing complex behaviours in robots contains two additions to the recommendations of behaviour based robotics:

1. Implement a hypothesised path of evolutionary steps to a desired animal behaviour.
2. Include a maximum number of different behaviours and activities on each evolutionary level.

2.2 A Step-Wise Approach

Brooks' original rules for BB robotics [3], suggested taking inspiration from evolution. We suggest a more extreme approach where rather than just trying to implement a desired set of behaviours inspired by evolution, each step of a hypothesised historical evolutionary path to a desired set of behaviours is implemented.

The motivation for retracing evolution in implementation is that the complex behaviours found in animals and humans are so poorly understood that robust and efficient direct implementations are impossible. In these cases, retracing evolution forces an investigation of the evolutionary history of the behaviour in question.

Undertaking a complete behavioural investigation can be preferable to implementing a brittle or inefficient approximation, something which has often been the result of trying to implement complex behaviours directly.

What we call complex behaviours are behaviours that are not directly implementable. Our failure to implement automated learning of novel robot behaviours directly was originally the inspiration for researching a behaviour-based approach to such learning.

An Evolutionary Theory of Learning One of the main sources of inspiration for our holistic development is the growing amount of knowledge of the physiology and evolutionary history of biological systems that is found in areas such as ethology, neuro-science, and the cognitive sciences. There is currently a wide scope for using this knowledge in AI implementations.

In particular Moore [12] presents a clear theory of how increasingly complex forms of learning might have developed. Below we list the main types of learning as presented by Moore, with contributions from other theories added in italics.

- Imprinting
- Alpha Conditioning

– Pavlovian Conditioning
– Operant Conditioning
– Skill Learning
– Imitation
– *Language Learning*

This list of learning types is not complete, and interesting questions are raised concerning types of learning not included in this hierarchy such as classification and insight.

2.3 Behavioural Holism

The second recommendation reflects the realisation that complex behaviours that involve learning always relate to more basic underlying behaviours supporting a number of activities. Complex behaviours cannot be thoroughly explored without being emerged in a rich behavioural context. The evolutionary histories of behaviours are highly interrelated, and looking at a limited number can not reveal all the details necessary for a comprehensive understanding.

Learning is also a problem that needs to be strictly biased if it is to be successful. The way biases are introduced in biological systems is through a hierarchical structuring of data and control [6]. This kind of structuring is done by pre-existing neural circuitry, and the more effective biases we need, the more underlying circuitry we must provide.

Below we present the three main arguments for the need of behavioural holism.

Conclusions from Our Own Work As we analysed increasingly complex forms of conditioning, it became difficult to design natural learning problems to test the different learning types due to the poverty of the underlying controllers. In designing an experiment for demonstrating alpha conditioning, we needed the robot to recognise that a certain stimulus would regularly occur together with food. The underlying controller could only recognise other robots and food, so we had to invent an artificial pink box sense.

If a holistic approach had been taken, we would have had a number of senses related to other basic behaviours to choose from so that our alpha conditioning experiment would have been more natural and perhaps brought up issues of behaviour integration that were missed because of the artificial nature of our pink box sense.

Our BB analysis of conditioning points out that increasingly complex learning behaviours learning need an increasing number of underlying behaviours to support it.

Arguments from Cognitive Robotics In 1997, Brooks criticised work in BBAI for not having a wide enough *behavioural repertoire* [4]. He recognises the vastly richer set of abilities needed by robots in order to act like a human, and suggests work be done on activity interaction and control.

Brooks explicitly lists *coherence* as an issue to be considered in cognitive robotics. Coherence is a complex and poorly understood problem that involves many different sub-systems. A step-wise, holistic approach provides a study of increasingly complex manifestation of problems spanning many sub-systems, such as coherence, learning and communication. For complex behaviours, this kind of study is necessary to provide solutions of acceptable quality.

In Section 4 we suggest analysing human evolution in order to create a roadmap of human behavioural evolution as a means to support a holistic approach.

Arguments from Evolution Zoologists have provided one of the strongest arguments for a holistic approach to AI:

> No single characteristic could evolve very far toward the mammalian condition unless it was accompanied by appropriate progression of all the other characteristics. However, the likelihood of simultaneous change in all the systems is infinitesimally small. Therefore only a small advance in any one system could occur, after which that system would have to await the accumulation of small changes in all the other systems, before evolving a further step toward the mammalian condition.
> *T.S. Kemp [11]*

This quote was also used in [1], which in addition presents the following example. In order to maintain a constant body temperature and extend their periods of activity, warm blooded animals need to consume an order of magnitude more food than cold-blooded animals. As a result, they have changed the way they chew food, their breathing, their locomotion, their parenting behaviour, their senses, their memory capacity, and their brain size.

In cognitive modelling, we can make simultaneous changes, but we cannot make large changes to some types of behaviour without appropriately advancing others.

3 Experiments on Adaptive Behaviours

3.1 Reproducing Animal Learning

What our experiments show is that it is possible to implement certain types of learning using a step-wise BB approach rather than a direct implementation. The goal of our work is to provide new efficient solutions to learning problems where current solutions are inefficient or brittle solutions, in particular the learning of novel behaviour patterns, but also traditional problems such as natural language acquisition.

As necessary in a step-wise approach, we first looked at low level learning mechanisms. We conducted four experiments on habituation learning, spatial learning, behaviour recognition, and basic association. The experiments were chosen in order to reflect both different types of learning and different types of activities. The habituation and spatial learning experiments are concerned with navigation and feeding, the basic association experiment concerns with feeding and avoiding danger by recognising poison. The behaviour recognition is concerned with fighting and courtship displays. From the four F's of animal behaviour, feeding, fleeing, fighting and procreation, we have touched on all but procreation.

Below we present only the experiment on behaviour recognition, as it best displays the working of the step-wise approach.

3.2 Demonstrating Behaviour Recognition

The Experiment Our first attempt at modelling a more complex form of learning with two interacting adaptive layers, was a courtship display experiment. In this experiment, two robots used a display behaviour to avoid the injuries of physical fighting. These kinds of displays are common in animals and are one of the simplest forms of animal communication [10].

The Environment In order to simulate conflict behaviours, it was necessary to provide a simulated environment with a number of features. To support physical fighting, we gave each robot a certain strength and we simulated physical damage by making energy levels drop noticeably and proportionally to the opponents strength whenever the Kheperas were in physical contact.

3.3 A Step-Wise Solution

The solution to this restricted form of behaviour recognition consisted of three evolutionary steps or layers: the reactive interaction layer, the learning from fighting layer and the display layer.

Reactive Interaction As a basis for more complex interactions, we implemented a reactive behaviour where a robot always tries to get in physical contact with, i.e. attack, its opponent when is sees it close by. In Figure 1 we present the circuits that implement the reactive interaction. These circuits illustrate how we build more complex learning behaviours on top of simpler solutions. The *ApproachFeederPositionController* that the reactive interaction behaviour is put on top of is the solution to the mapping experiment. Up to this level, the robots take no notice of other robots.

A Khepera sense was added which recognises when the opponent is in a threatening position, i.e. near by and facing our robot. In such cases, a touch Khepera drive which inhibits all other behaviour, approaches the other robot in a simulated fighting behaviour. This simple reactive behaviour would in the long term lead to the simulated death of the weakest robot.

Learning from Fighting The first adaptive layer implemented to support behaviour recognition was a layer where the robots learn which one is the strongest by the amount of damage they take. A memory circuit is then used in this behaviour to remember the pain of being the weakest robot. This memory is

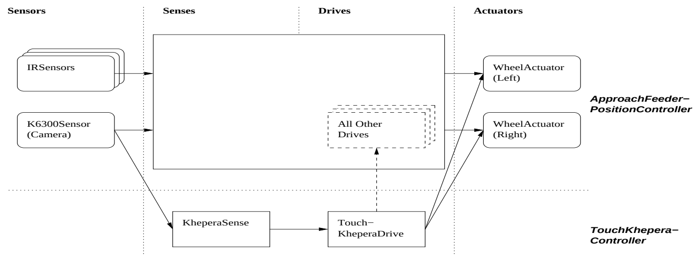

Fig. 1. Reactive Robot Interaction

supported by a pain sense which picks up losses of energy and a fear emotion which is activated by the pain sense. After a fear based memory is established, an avoidance drive ensures that the weakest robot avoids its opponent in the future. The circuits involved in the fighting behaviour are presented in Figure 2.

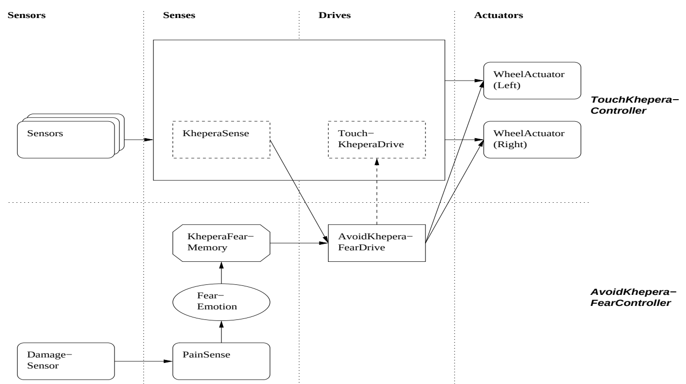

Fig. 2. Learning from Fighting

Learning from Courtship Displays On top of the fighting layer, we implemented a courtship display layer which took the form of a stand-off initiated by the Khepera sense. In a stand off, the robots remain motionless for an amount of time corresponding to their strength. This behaviour needed a strength sense and a memory circuit to keep track of how long the robot had been displaying. These two circuits were used to support a Khepera stronger sense which was activated when it became clear that the other Khepera was stronger.

This use of memory can be described in the habituation type learning framework as increased sensitisation of a yielding behaviour, where the strength sense acts as a threshold.

The stand-off was over when one robot recognised the opponent as stronger. This recognition fired the fear emotion and a basic memory was created using the same circuit that was used in the physical fighting layer.

The final addition was to let the avoid Khepera fear drive inhibit the display drive in order to yield and as a result break up the stand-off by no longer taking a threatening stand. The circuits used to implement the display behaviour on top of the fighting behaviour are displayed in Figure 3.

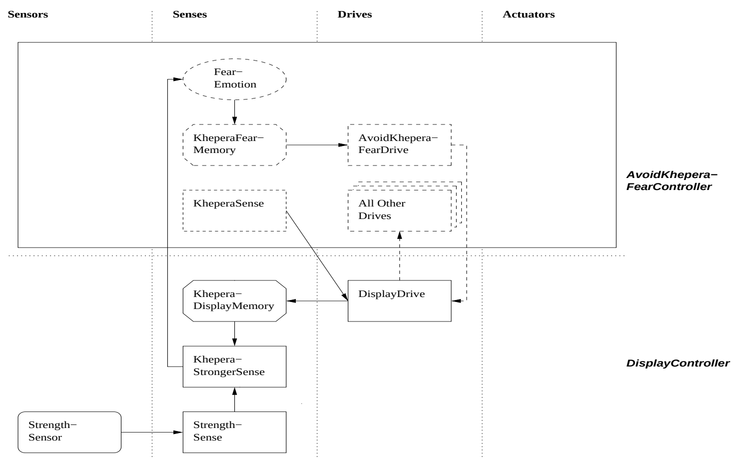

Fig. 3. Learning from Courtship Displays

3.4 Conclusions

Our experiments show that a number of different forms of animal learning can be reproduced in simulated Khepera robots using a step-wise approach. The experiments together with a PLANCS based analysis of the class of learning problems called *conditioning* [15] indicate that a step-wise and behaviourally holistic approach is sufficient for implementing these forms of learning. More complex forms of learning however would need to be studied explicitly before we can evaluate the feasibility of finding a solution using this approach.

4 A Road-Map from Evolution

When hypothesising a path of evolutionary steps to a desired behaviour, it is helpful to have a clear picture of the evolutionary background of that behaviour.

To include appropriate behaviours with sensible levels of sophistication at each step in a holistic manner, it is helpful to have an idea of what types of behaviours and levels of sophistication are likely to have coexisted during evolution.

We suggest that an analysis of human evolution in terms of co-existing behaviours of different evolutionary sophistication as well as bodily complexity can be a helpful road-map for a step-wise, holistic approach to BB robotics.

Figure 4 gives an idea of what such a road map might look like. It presents six numbered evolutionary stages. Between the stages are behaviours and physical attributes that are likely to have coexisted during evolution. Figure 4 is roughly put together from an introductory level text on evolution [16] and is only meant to convey the idea of how a road-map would look. It is not meant to exclude any dimension of behaviour that has an evolutionary history if a better knowledge of that dimension would facilitate development.

	Feeding	Fighting	Fleeing	Procreation	Sensors	Actuators	Adaptability	Habitat	Social Env.
6									
						Speech organ	Symbolic learning		
5									
	Agriculture	Armed fighting	Nesting		Colour vision	Hands	Insight	Settlements	
	Armed hunting								
4									
							Temporal association		Hierarchical Group
				Raise young			Imitation		
3									
		Group fighting	Group protection			Vocal tract			Uniform Group
2									
	Hunters	Displaying	Hiding	Mating	Stereo vision	Legs		Land	Family
			Fleeing pursuer	Release	Stereo hearing				
					Motion compensation	Head			
					Directed vision				
1									
	Grazers	Physical fighting	Moving	Division	Exterio–receptors	Body	Association	Water	Solitary
					Interio–receptors		Habituation		

Fig. 4. Human Evolution as a Road-Map to Robotics

Communication Communication is a dimension of behaviour which was not included in Figure 4. This aspect of behaviour spans all of evolution, has a dramatically increasing complexity and on a human level constitutes the important areas of production and recognition of speech and text.

The holistic approach was developed to study learning, another type of secondary behaviour. It should also provide results in the area of communication since the evolutionary histories of the two behaviours are closely interrelated.

Presumably, there are other dimensions of behaviour or influences on behaviours that could be added to Figure 4. We continue to seek discussion about our analysis and also aim to update it with new results from the relevant sciences.

Currently Unreachable Problems In robotics, research is currently taking place on issues on all the different levels of complexity presented in Figure 4 and very little work has been done on integrating behaviours in holistic frameworks. The behaviours on the lower levels might be studied in isolation with some credibility, while in the levels further up, one-dimensional research gets less and less useful in an AI context.

Existing research in areas like planning, reasoning, and language recognition and production, has produced important scientific results and impressive software engineering tools, but it is not obvious that any of these tools have a place in complex robotic systems.

5 Discussion

Step-Wise, Holistic Approaches and ALife Our reason for not taking an automated search approach to evolutionary robotics the way the ALife community does [14] is primarily because we believe that it is more efficient to implement current theories from the relevant sciences directly, than it is to express those theories in the form of fitness functions and environments in order use automated search to find solutions. It can be argued that the solutions found using automated search are more robust and can take into account parameters that are not known to the developers. Automated search might also provide new knowledge about specific problem domains. We believe that these possible results do not warrant the abandonment of our approach. Our approach has an engineering emphasis rather than an automated search emphasis, but we see the two methods as complimentary and believe that the automated search approach to evolutionary robotics would also benefit from adopting our holistic principles.

Common Criticism It has been suggested to us, and from our results it is sometimes tempting, to try to create a neural circuit based model for high level learning in order to 'solve' the problem of high level learning. We think the utility of such an effort would be limited. Our experiments are an exploration of basic forms of learning and do not test a pre-formulated hypothesis about learning. We do not want to commit to a general theory of learning such as e.g. artificial neural networks or reinforcement learning. Our effort seeks to restrict the search space for common animal learning problems by providing supporting neural structures. The learning problems that remains should be solvable using any learning technology.

The background for this reasoning is that because of the evolutionary cost of adaptability, it is likely that any evolved learning mechanism will be a simple form of learning placed in a complex behavioural context, rather than a more unconstrained complex learning mechanism with a larger probability of learning the wrong things. This belief is reinforced by the way many form of learning, previously thought to demand complex learning frameworks, such as spatial learning and imitation, turn out to be implementable using simple adaptive circuits strategically placed within complex behavioural circuitry. It is further supported by the presence of the necessary contextual neural circuitry in animals [13, 9].

References

1. J. M. Allman. *Evolving Brains.* Scientific American Library, 1999.
2. R. C. Arkin. *Behaviour Based Robotics.* MIT Press, 1998.
3. R. A. Brooks. Intelligence without reason. In *Proceedings of IJCAI 91*, pages 569–595. Morgan Kaufmann, 1991.
4. R. A. Brooks. From Earwigs to Humans. *Robotics and Autonomous Systems*, 20(2–4):291–304, 1997.
5. J. J. Bryson and B. McGonigle. Agent Architecture as Object Oriented Design. In *Intelligent Agents IV, Proceedings of the Fourth International Workshop on Agent Theories, Architectures, and Languages (ATAL'97), LNAI 1365*, pages 15–30. Springer Verlag, 1997.
6. J. J. Bryson and L. A. Stein. Modularity and Specialized Learning: Mapping Between Agent Architectures and Brain Organization. In *Proceedings of the EmerNet International Workshop on Concurrent Computational Architectures Intergrating Neural Networks and Neuroscience*. Springer Verlag, 2000.
7. H. I. Christensen. The WEBOTS Competition. *Robots and Autonomous Systems Journal*, 31:351–353, 2000.
8. T. S. Dahl and C. Giraud-Carrier. PLANCS: Classes for Programming Adaptive Behaviour Based Robots. In *Proceedings of the 2001 Convention on Artificial Intelligence and the Study of Simulated Behaviour (AISB'01), Symposium on Nonconscious Intelligence: From Natural to Artificial*, 2001.
9. V. Gallese, L. Fadiga, L. Fogassi, and G. Rizolatti. Action Recognition in Pmotor Cortex. *Brain*, 119:593–609, 1996.
10. M. D. Hauser. *The Evolution of Communication*, chapter 6 Adaptive Design and Communication, pages 450–470. MIT Press, 1996.
11. T. S. Kemp. *Mammal-like Reptiles and the Origin of Mammals.* Academic Press, 1982.
12. B. R. Moore. The Evolution of Imitative Learning. In C. M. Heyes and B. G. Galef, editors, *Social Learning in Animals: The Roots of Culture*, pages 245–265. Academic Press, 1996.
13. R. U. Muller, J. L. Kubie, E. M. Bostock, J. S. Taube, and G. J. Quirk. Spatial firing correlates of neurons in the hippocampal formation of moving rats. In J. Paillard, editor, *Brain and Space*, chapter 17, pages 296–333. Ocford University Press, 1991.
14. S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines.* MIT Press, 2000.
15. J. M. Pearce. *Animal Learning and Cognition.* Psychology Press, 2nd edition, 1997.
16. M. W. Strickberger. *Evolution.* The Jones and Bartlett Series in Biology. Jones and Bartlett Publishers, Second edition, 1995.