Directions for Bio-Inspired Artificial Intelligence

William J. Teahan*

School of Computer Science, Bangor University, Bangor, Wales, UK

Introduction

Bio-inspired Artificial Intelligence [1] attempts to synthetically produce systems that exhibit intelligence by taking inspiration from the processes of natural life systems. Examples of bio-inspired artificial intelligence include behaviour-based robotics, artificial neural networks, evolutionary algorithms, particle swarm optimisation and ant colony optimisation.

Biology is a branch of science that seeks to determine the laws of nature that lie behind the structure and behaviour of living organisms. A biologist will postulate a theory or model in order to explain certain natural phenomena, and will then experimentally verify how well the model predicts what is observed in nature. If the experimental results do not match with what occurs in nature, then in the words of the acclaimed physicist Richard Feynman, "it's wrong". Although there is no design constraint in the field of Artificial Intelligence to restrict oneself to applying exactly the same mechanisms to those used by natural life systems, the "wrongness" of the particular model in reproducing natural intelligence can be a useful guide in measuring the success of a bio-inspired AI model. Furthermore, in attempting to reproduce human level intelligence (since few AI systems to date has come close to human capabilities), insight can be gained that can potentially lead to further breakthroughs. The purpose of this article is to emphasize the importance of this process of validation against real-world systems in evaluating bio-inspired AI models. We can gain insights from the comparison with observations of naturally occurring processes in order to inform the design of our intelligent artificial life systems.

Many of the models that have been developed to date are only loosely related to their real-life equivalent (such as artificial neural networks to name the most obvious example). Many bio-inspired AI systems have overlooked the importance of maintaining the real-world comparison. Once a technique has been deemed effective, the approach often diverges from the real-world equivalent, and improvements to the method (such as back-propagation for artificial neural networks, for example) will often concentrate more on improving machine performance and less on how to reflect better observed natural phenomena in order to improve the realism.

It is often important to revisit the design objectives of the original research by drawing further inspiration from real-life systems. This can be done by examining those phenomena which are not covered well by the model and exploring alternative solutions, or by trying to extend the model to improve the realism and analyse where the model is wrong.

With this in mind, three research directions that are currently being explored by the author of this article are described below: firstly, the problem of concept representation and learning, with specific focus on colour recognition; secondly, improvements in existing evolutionary algorithms which take their inspiration from natural evolutionary processes; and thirdly, real-time behavioural animation

of virtual creatures and virtual humans. The first area is an example of using the validation approach advocated in this article; the second area shows how significant improvements in performance can be made by making modifications inspired by biological mechanisms not included in the original model; and the third area in real-time animation uses agents that are embodied and situated in virtual environments that can be visually analysed to assess the realism of the models.

Representing Colours

Some important areas in AI research include the problems of knowledge representation, learning and categorisation. The 'symbolic' approach to AI posits that it is sufficient to represent knowledge symbolically. An opposing approach, called 'connectionism', rejects the classical symbolic approach, and instead advocates that intelligent behaviour is the result of sub-symbolic processing – that is, the processing of stimuli rather than symbols; for example, using artificial neural networks that comprise interconnected net-works of simple processing units. Here the knowledge is stored as a pattern of weights of neuron connections. A number of other approaches have been devised, some proposing a combination of symbolic and sub-symbolic processing, such as conceptual spaces [2]. Gärdenfors postulates that what is fundamental to our human cognitive abilities is our capacity for processing concepts and these emerge from a distributed connectionist representation at the lowest level where stimuli from receptors are processed, and combine to form symbolic structures at the highest level. In Gärdenfor's approach, concepts are represented as regions in n-dimensional space, with similar concepts represented in geometric regions that are spatially located near to each other. The multi-dimensional geometry for a conceptual space is constructed from what Gärdenfors calls 'quality dimensions' which correspond to the different ways an agent's stimuli are judged to be similar or different. The primary function of the quality dimensions is to represent various qualities or features of objects.

The colour naming problem involves the question of how to categorise colours. Berlin and Kay [3], for example, identified eleven basic colours – white, black, red, green, yellow, blue, brown, purple, pink, orange and gray. The colour naming problem provides an ideal vehicle for testing out cognitive models for AI. The problem is easily understood but presents difficulties because colours occur across a continuous spectrum (for example, try separating cars that occur

*Corresponding author: Dr. William J. Teahan, School of Computer Science, Bangor University, Bangor, Wales, UK, E-mail: w.j.teahan@bangor.ac.uk

along a highway into discrete categories). There is also a wealth of experimental evidence which we can use to validate our models against. For example, experiments with human perception of the colour domain show that the colour conceptual space is described using three quality dimensions – brightness, hue and saturation. Other studies show that humans have the ability to create prototypes for specific colours, distinguish the most similar colours and distinguish complementary colours. With the conceptual spaces model, these capabilities require the application of a simple distance-based metric (e.g. Euclidean distance) in three dimensional space. The colour naming problem can also help investigate the relationship between language and thought – the conjecture that the structure of language can determine the way we perceive the world (called the Sapir-Whorf hypothesis). And colour blindness raises further issues that require investigation.

Surprisingly, in the field of AI, there has been relatively little investigation of the colour naming problem. Although the problem has been highlighted in robotics, the importance of understanding the problem from the knowledge representation or machine learning perspectives has largely been ignored. Gärdenfors conceptual spaces model also raises many questions. For example, how are concepts (to be) represented in the neural networks (both human and machine)? By exploring different models for representing the colours and validating them against naturally observed phenomena, we can potentially make discoveries that might help explain mechanisms behind human intelligence and/or lead to improved capabilities for artificial intelligence.

An initial investigation by the author into this problem has shown promise. An implementation based on the conceptual space model leads to highly accurate colour categorisation from surprisingly little training data.

Bio-inspired Improvements to Evolutionary Algorithms

The Grammatical Evolution (GE) algorithm [4] is an evolutionary algorithm that takes inspiration from nature and can evolve complete programs in an arbitrary language using populations of variable-length binary strings. Grammatical Evolution's unique features compared to other evolutionary algorithms are the degenerate genetic code which facilitates the occurrence of neutral mutations (various genotypes can represent the same phenotype), and the wrapping of the genotype during the mapping process which enables the reuse of the same genotype for the production of different phenotypes.

Georgiou and Teahan [5] have developed a modification to GE called Constituent Grammatical Evolution (CGE) that takes inspiration from two very elementary concepts: genes and conditional behaviour switching. It augments the standard Grammatical Evolution algorithm by incorporating and utilising these concepts in order to improve the performance of the later in agent problems.

The new evolutionary algorithm constitutes a conditional behaviour switching evolutionary algorithm, based on the Grammatical Evolution algorithm, and which incorporates the notion of genes in the individual's genotype. CGE builds from elementary and more complex building blocks a control program which dictates the behaviour of an agent. The modified algorithm significantly outperforms GE and other evolutionary algorithms, such as Genetic Algorithms (GA) and Genetic Programming (GP). The next step in the research is to apply

the modifications to alternative types of evolutionary algorithms, such as GA and GP, to see if the improvements apply to all classes of algorithms and not just to GE in particular.

Real-Time Animation of Virtual Creatures

Real-time virtual environments are becoming increasingly complex, realistic and physically simulated. Similarly, computer game characters continue to advance towards the elusive goal of photo-realism. However, virtual characters are still largely animated using pre-created data, such as motion capture and key-frame animation. Whilst perfectly adequate for simple or specifically designed environments, the interaction between agent and environment is still primarily an illusion. Removing or at least limiting this discrepancy between agent and environment potentially allows for much more interesting, complex and realistic animations to form. Furthermore, by embodying the agent in such a way leads to behavioural realism – provided that the simulation model remains relatively realistic, any animation that forms is inherently realistic for the specific situation, however complex.

ap Cenydd and Teahan [6] have developed a real-time simulation of walking arthropods, capable of autonomous exploration over an arbitrarily complex environment (see Figure 1). The physical grounding of the creatures allows for complex, emergent animations to form that are specific and unique to the situation. Not only can the arthropod interact with physically simulated objects in the virtual world, the actual interaction in turn affects the behaviour and resultant animation. This constant creature-behaviour feedback elevates the vividness of the animation beyond all but the most contrived or carefully designed agent-environment interactions found in other simulations. Recent advances in the simulation include the ability to fully physically simulate the articulated creature, allowing for an unprecedented degree of interaction with the environment, including full-body collision detection, rag-doll recovery, rigid-body interaction, friction and inertia.

A primary goal of this research is to simulate these creatures with enough autonomy that complex animations emerge due to embodied circumstance, not predetermined direction. The next step will be to explore the problems associated with the addition of higher cognitive functions to the artificial creatures, by linking these behaviours to the creature's physical embodiment in the virtual environment, and also to proceed up the evolutionary ladder, with the ultimate goal being realistic animation of virtual humans.

Conclusion

Many of the important advances in the field of Artificial Intelligence (AI) such as artificial neural networks and evolutionary algorithms have taken their inspiration from biology. However, these AI systems are only very loosely related to their real-life equivalents. Further research
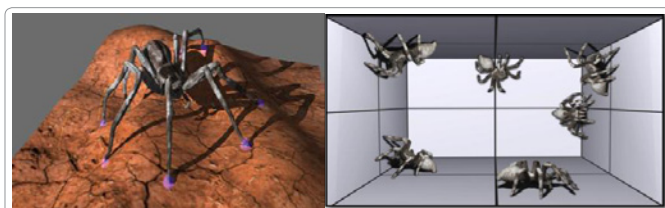
Figure 1: Screenshots of virtual arachnids that have been procedurally animated in real-time.

is essential into how biological mechanisms can be used to enhance existing AI systems and develop new models. It is crucial that these systems are validated against real observations. From understanding gained through the development of artificial life systems of behaviours that reproduce the patterns we see in real-life, we can go forward to uncover the as yet undiscovered mechanisms that lie behind intelligent life. The hope is that the insights gained through this understanding will lead to the crucial breakthroughs needed to produce intelligent artificial life.

References

1. Floreano D, Mattiussi C (2008) Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies. The MIT Press.

2. Gärdenfors P (2004) Conceptual Spaces: The geometry of thought. The MIT Press.

3. Berlin B, Kay P (1969) Basic Color Terms: Their Universality and Evolution. University of California Press, Berkeley.

4. O'Neill M, Ryan C (2003) Grammatical Evolution: Evolutionary Automatic Programming in an Arbitrary Language. Springer, USA 4.

5. Georgiou L, Teahan WJ (2011) Constituent Grammatical Evolution. Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence *(IJCAI)*, Barcelona.

6. ap Cenydd L, Teahan B (2012) An embodied approach to arthropod animation. J Computer Animation and Virtual Worlds.