

AI as a Boundary System

By Paul Almond, 17 September 2006

Website: www.paul-almond.com
Email: info@paul-almond.com

© Copyright Paul Almond, 2006. All Rights Reserved.

AI as a Boundary System

By Paul Almond, 17 September 2006

Introduction

“Left to themselves, they had developed a form of intelligence almost devoid of short-term memory ... With the skrode’s mechanical short-term memory they could learn fast enough that their new mobility would not kill them.”

- *A Fire Upon the Deep* (Fiction), Vernor Vinge (b. 1944 CE)

This article will propose a paradigm for dealing with the AI (artificial intelligence) system suggested in my previous article *How AI Would Work [1]* and other articles. I recommend reading the previous article, or some of my other articles about this AI system, before this article: otherwise it may not make sense.

In the previous articles I described an AI system based on a probabilistic *functional hierarchy* of meaning extraction providing statistical constraint on look-ahead searches performed by the *output vectoring system* in conjunction with a *situational evaluation function*.

This article will not change the technical details of the proposed AI system, but will suggest a paradigm for viewing it. The paradigm will not be suitable in all situations and for all kinds of work, but may be useful for some purposes.

The paradigm is that of an AI system as a *boundary system* – a system with limited capabilities, confined to a simple, basic “layer” (I am using that word loosely) of computation between two “worlds” - outside reality and the AI system’s own functional hierarchy.

External Processes and Computation

The association of processes with a thing does not mean that we must regard them as being *part* of that thing. For example, we tend to regard the bark of a tree as part of the tree because of the tree has made it for its own use, but we do not generally regard spacecraft and cars as being part of ourselves, even though we have made them for our own use just as a tree makes bark. The distinction is largely semantic. To some more advanced intelligence, our technology base might simply be viewed as extra parts of themselves that humans have made in the way that trees make bark. Alternatively, the bark of a tree could be viewed as part of the external environment which the tree has manipulated into a form that will protect it – a tool.

Objections could be attempted to this, arguing that tree bark is qualitatively part of a tree in some way that our technology base is not part of us, but it could be hard to decide where the dividing line should be drawn: semantics will ultimately be involved. We can

choose as a matter of convenience whether to regard some arrangement of matter in the world as part of an organism or a tool, external to itself, that it has made.

This also applies to computation. Many people would not regard the computers that we have built as part of our brains or minds, yet to all intents and purposes some of the information processing in our brains is extended into these artefacts. For example, if we are considering some problem requiring the intermediate solution of a further mathematical problem, we may use a computer to solve the mathematical problem, the result being returned to our brains and used to complete the larger problem. It could be argued that this extends our brains in some way. Of course, it could also be argued that our brains are manipulating the environment into doing useful computation, the results of which are returned to our brains.

How This Relates to AI

Using the same semantic convenience, instead of regarding the functional hierarchy - the AI system's modelling and planning system - as part of the AI system itself, we can regard it as something external to the AI system that it changes and manipulates using its outputs, just as it manipulates any other aspect of reality.

If the functional hierarchy is external to the AI system then what is actually left in the AI system? What remains is the output vectoring system and the evaluation function. This leaves a minimal AI system which cannot do much by itself. It needs the help of the functional hierarchy, which it now gets by manipulating an external system rather than from within itself.

How the Boundary System Paradigm Works

I describe this as the "boundary system" paradigm because it relegates the part of the system that we consider the AI to what can be considered to a thin "layer" – the boundary between what would otherwise be considered the rest of the AI system and the outside world.

In the boundary system paradigm the AI system does nothing except perform look-ahead searches to determine what outputs to make. It exists on the boundary between two worlds - external reality (the outside world) and the functional hierarchy (what would usually be considered the rest of the AI system) - and its outputs affect what happens in them.

When outputs affect external reality, changes are caused in it, affecting the system's situation and the scores which will be returned by the situational evaluation function.

When outputs affect the functional hierarchy, changes are caused in it, affecting the sort of constraint that the functional hierarchy provides back to the boundary AI system (the output vectoring system and evaluation function).

Selection of outputs by the output vectoring system still involves simulation. This means that when the output vectoring system is performing its tree search, outputs are only made hypothetically, to see what the results are. The effects of such an output on the functional hierarchy only apply below the node at which it is made.

An attempt to make a physical model of all this, thinking in terms of physical layers, could involve imagining some computer containing an AI system having the functional hierarchy occupying almost all of its volume and the output vectoring system and situational evaluation function – what I call the “boundary” AI system in a thin layer on the outer surface of the machine, between the functional hierarchy and the outside world. It should be pointed out, however, that the paradigm is not really based on physical concepts of where parts of systems are.

Justification for the Boundary System Paradigm

The paradigm treats the functional hierarchy as just another piece of external reality, instead of part of the AI system itself. This is an obvious view to take because the approach taken to the carpet texture problem involves using the system’s outputs to control prioritization in the functional hierarchy - the same method used to manipulate the outside environment. From the point of view of view of the boundary AI system it simply uses outputs to manipulate two kinds of environment – one of which can be used in look-ahead simulations and returns constraint information for its tree searches.

Conclusion

The paradigm of AI system as boundary system has been proposed to complement my earlier articles on AI. The AI system itself can be viewed as being just the output vectoring system and the situational evaluation function, while the functional hierarchy is considered as a part of external reality which the system manipulates by means of its prioritization control outputs. The paradigm is not suitable for all situations, but may be a useful way of considering the AI system when solving some problems.

References

[1] Web Reference: Almond, P. (2006). *How AI Would Work*. Retrieved 4 September 2006 from <http://www.paul-almond.com/HowAIWouldWork.pdf>.