CONSTRUCTING COMPLEX BRAINS

Building minds using sub-minds from biotechnology authors.

[1]Ray Walshe, Mark Humphrys,[2] Ciarán Ó Leary
[1]*School of Computing, Dublin City University,IRELAND.* [2]*School of Computing, Dublin Institute of Technology,IRELAND.*[1]*computing.dcu.ie/~ray computing.dcu.ie/~humphrys*[2] .*www.comp.dit.ie/coleary*

Abstract: The World-Wide-Mind (WWM) is an architecture that facilitates the construction of complex "minds" (software solutions) for virtual agents in online virtual environments (software problems). These complex minds are constructed by re-using third party sub-minds remotely through the world wide web by hosting both the mind and environment online. Until now there has been no successful language independent framework for reusing complete or partial solutions by third party authors in the construction of more diverse and complex solutions. The WWM architecture attempts to provide such a framework. This paper describes an ongoing project that uses WWM technology to develop a novel set of minds that incorporate a range of algorithms and techniques. We demonstrate that it is possible to evolve superior minds by the artificial selection and combination of existing online minds. The project that is described here involved 8 authors who developed 8 different minds for the virtual animal in a re-implementation of Tyrrell's Simulated Environment (Tyrrell, 1993). These 8 authors were from a biotechnology background and were enrolled on a postgraduate masters course in bioinformatics. The minds created by these authors were subsequently selected based on their performance in the virtual world, and integrated into larger minds. Since a huge number of combinations of minds are possible, it is important that the work is distributed among a community of researchers. The architecture of the World-Wide-Mind makes this possible

Key words: Distributed AI Algorithms, Distributed AI Systems and Architectures, AI Languages, Programming Techniques and Tools

1. INTRODUCTION

The World-Wide-Mind (WWM) (Humphrys and O'Leary, 2002) is an architecture that facilitates the construction of complex "minds" (software solutions) for virtual agents in online virtual environments (software problems). These complex minds are constructed by re-using third party sub-minds remotely through the world wide web by hosting both the mind and environment online. Many people have discussed how little re-use of other authors' agent software and agent architectures takes place (Guillot and Meyer, 2000). The WWM attempts to provide a language independent framework by which communities of remote authors can construct agent minds online.

This paper describes an ongoing project that uses WWM technology to develop a novel set of minds that incorporate a range of algorithms and techniques. We show that the WWM architecture provides a framework that significantly reduces the need for knowledge about operating systems and networking (Walshe and Humphrys 2001). We demonstrate that it is possible to develop superior minds by the artificial selection and combination of existing online minds. The project that is described here involved 8 authors enrolled on a postgraduate masters course in bioinformatics, who developed 8 different minds for the virtual animal in a re-implementation of Tyrrell's Simulated Environment (Tyrrell, 1993). These authors constitute a sub group of researchers on the project, as the other 242 authors are undergraduate students and the results of that Tyrrell experiment are described in a forthcoming paper (O' Leary et al forthcoming). The minds created by these authors were subsequently selected based on their performance in the virtual world, and integrated into larger minds. The architecture of the World-Wide-Mind makes this possible

1.1 Structure of Project

The ongoing project that is described here has so far involved 250 authors who each developed minds (or action selection mechanisms) for the virtual animal in a re-implementation of Tyrrell's Simulated Environment (Tyrrell, 1993). This simulated environment models a complex environment of a virtual animal world inhabited by a creature with multiple conflicting goals. Each author (the majority of whom were undergraduate Computer Science students) is permitted to develop any number of minds, which are then made available online, as web services. Since the minds are online, third parties can select a specific mind and run it in the virtual world, located at w2m.comp.dit.ie. This effectively means plugging the mind into the animal in the virtual world, and letting it select the actions at each point in time. The score that is achieved by each mind in each run is recorded in an online scoreboard that can be viewed by all other authors. The scoreboard had two effects (i) it introduces an element of competition between authors that encourages them to write improved versions of their minds, and (ii) it identifies for authors the best minds to incorporate into their own minds as sub-minds or modules.

1.2 Mind$_M$ Services

The manner in which existing minds can be integrated into new minds takes a number of forms. For example an author can select the two minds at the top of the scoreboard and wrap a "high level mind" around them. The high level mind will make the decision on which sub-mind to call at which point in time, resulting in a division of the problem state-space between the two sub-minds. From the point of view of a student project, this use of other people's work in order to select actions for the virtual animal, could be seen as "cheating". We encourage this "cheating" by students (so long as it was written up properly) since such re-use is the whole point of the WWM. Using the terminology introduced in (Humphrys, 2001), this is a Mind$_M$ service.

These better minds are then more likely to get artificially selected for inclusion in a Mind$_M$ service. The WWM opens up the possibility of large numbers of people addressing a problem over long periods of time with the best sub-minds being selected on their fitness to create the next generation of Mind$_M$ servers. The results outlined here are based on three Mind$_M$ services developed by the third author. They are accompanied by a description of the strategy used for developing each Mind$_M$ service based on the scores available in the scoreboard. In addition we describe how these Mind$_M$ services, which are also available online, will be subject to the same sort of evolution by artificial selection that the original minds are.

2. WORLD WIDE MIND

The World-Wide-Mind (WWM) is a scheme for facilitating the development of large minds through publishing virtual worlds and sub-minds online as web services. A web service is effectively a programming language object whose methods are invoked over HTTP, the language of the web, rather than being invoked in a local environment.

A simplified version of the well known web service architecture was developed. A large body of published work has demonstrated how protocols such as SOAP, WSDL and UDDI (w3.org/2002/ws) facilitate the discovery, description and invocation of remote services. Work on the Semantic Web, in particular OWL-S (formerly DAML-S) (www.daml.org/services/) has shown how semantic markup of services can further automate the composition and monitoring of services as composite web services, where multiple tasks are automatically integrated into a single service. While each of these architectures promise much, as do standard architectures such as CORBA (corba.org), and the Java model (java.sun.com), the authors found the level of complexity to be a significant barrier to participation to those interested in the creation of agent minds, without an interest in networking and distributed systems. To address this, the WWM architecture is web service based, but a single XML based protocol named the Society of Mind Markup Language (SOML) (O'Leary, 2003) is used as a substitute for the much more complex existing protocols. This point is expanded upon in (O'Leary and Humphrys, 2002).

The WWM specification (Humphrys, 2001) defines three types of entity, a world service, a mind service and a client. The world service supports a set of SOML messages for retrieving the state of the world and for instructing the agent in the world to execute a given action. The mind service represents an action taking mechanism developed for a given type of world service. The client entity is responsible for plugging a mind into an agent in a world. It does this by selecting a world service at a given URL and selecting a mind service at another URL. It then informs both the world service and the mind service that it intends to start a run involving them. The world service algorithm implements the problem (e.g. Tyrrell's SE) and the mind service algorithm implements a solution (one author's mind).

3. PROBLEM DOMAIN

At present there are two world services online that conform to the latest version of SOML. The first of these is a simple blocks world implementation, for which twelve mind services have been developed, including two $Mind_M$ services. A description of the problem and the design of the minds is discussed in earlier work (O'Leary and Humphrys, 2003).

A second available world service, which is the problem world at the centre of the work described here, is a re-implementation in Java of Tyrrell's Simulated Environment (SE) (Tyrrell, 1993). Tyrrell's PhD thesis examined a number of different action selection mechanisms when applied to a complex, multi-goal problem. In order to do this he wrote a virtual environment that modeled an animal in a heavily populated, dangerous environment (Figure 1). The primary goal of the animal was to mate, as this would ultimately determine how likely it was to pass on its genes to subsequent generations. However, in order to mate the animal needed to survive in the world long enough to be presented with sufficient opportunities to court and mate with other members of its own species.

The world that the animal occupied had fifteen different features, including food and water (both of which could be toxic), a den, cover and shade. It also featured several different types of animal; predators and non-predatory animals that needed to be avoided, prey that could be consumed and animals of the same species that could be mated with. The world experienced changes in weather and also cycled through different times of day.

At each time-step the animal was required to choose one of 35 actions to execute. The animal should choose the action that is most likely to increase its expected future genetic fitness i.e. the number of times it mates in this case. In order to do this the animal should maintain good health by eating and cleaning. It should also avoid an early death by hiding from predators and looking out for other dangerous places such as cliffs and marshes. In addition it should try to make its way home at nighttime to sleep.

Tyrrell implemented six different algorithms for the animal's action selection in the environment. His own algorithm which he named the Extended Rosenblatt and Payton algorithm performed best in the tests he conducted. His results over 6600 trials [mating average 8.09] were later improved upon by Bryson with her Edmund algorithm [mating average 9.12] (Bryson, 2000).

4. TYRRELL'S SE ONLINE

An implementation of Tyrrell's SE is available online as a WWM world service at w2m.comp.dit.ie. When queried for its state it will return a vector of numbers representing the animal's perceptions and memories and well as some additional information about the environment. It will accept as an action any number from 0 to 34, each representing one of the 35 actions that can be selected by the action selection mechanism. For a detailed explanation of state and action see w2m.comp.dit.ie.

Any mind service that will be created for this implementation of the SE should accept the state, process it and return an action i.e. it will serve as an action selection mechanism, except it will be located somewhere else on the Internet.

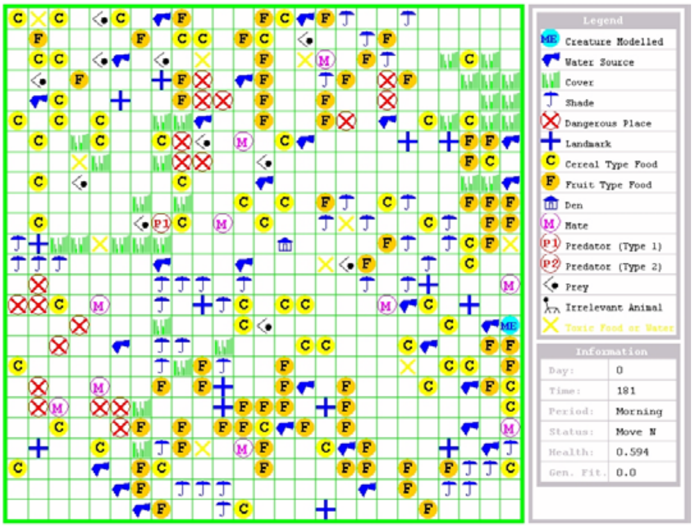

Figure 1. Graphical display of reimplementation of Tyrrell's Simulated Environment

5. FORMAT OF EXPERIMENT

The 8 biotechnology authors were given a tool kit that they could use to develop the mind and put it online. Once a mind was put online, it was automatically run in the world to record some account of its performance. Many minds were then run again several times by their authors or others who wanted to observe the performance of the mind, with a view to integrating the mind in their own Mind_M services. As the experiment continued, the better minds rose to the top of the scoreboard. Over a prolonged period of time, the artificial selections made by authors and users will result in a larger number of runs in the better minds, which will in turn result in their being selected more often for inclusion in Mind_M services, and consequently the creation of better minds.

6. INTEGRATION STRATEGIES

Using the large set of mind services that had been put online, we hand-designed a further two Mind_M services based on the biotechnology students' sub-minds. The first problem that we are faced with in developing a Mind_M service is in identifying at least two existing minds that could work together.

There are two main issues involved in the creation of a Mind_M service:

(i) Profiling each mind based on its performance in the world. The profile should provide sufficient information to be able to identify what goals are handled well by each mind. As more runs are performed using this sub-mind the profile should become a more accurate reflection of the strategy employed by the mind.

(ii) Deciding when to switch between minds. In the implementations above, minds were selected and deselected based entirely on an action that could be related to only one goal. When appetitive actions can be associated with more than one goal (e.g. move towards direction of nearest mate) it is difficult to infer that the sub-mind should be given control of the agent.

6.1 Mind_M 11 and Mind_M 12

By profiling the minds that scored best on the scoreboard we selected the top three biotechnology sub-minds on performance and ranked them as predominant maters or predominant survivors. As the top two sub-minds were predominant maters they were paired with the third best sub-mind which was a predominant survivor to produce two new minds which were Mind_M M11 and Mind_M M12

M11 was built using the *kb2* sub-mind as the survival algorithm and the *jm* sub-mind as the mating algorithm. *jm* ranked position 1 of the 8 sub-minds and *kb2* ranked position 3.

M12 was built using the *kb2* sub-mind as the survival algorithm and the *t* sub-mind as the mating algorithm. *t* ranked position 1 of the 8 sub-minds and *kb2* ranked position 3.

For both M11 and M12 the algorithm used was to listen to the survivor at all timesteps except when the mater suggested COURTING or MATING. This basic algorithm then provided the switching between the sub-minds based on their perceived strengths. See next section for summary of results of running the new Mind_M servers over multiple instances.

7. RESULTS

The results of the biotechnology sub-minds are shown below (Table 1).

Table 1. Biotech Mind Results

URL *	Average Lived	Average Mated	Best Lived	Best Mated
jm	1043	0.85	2850	3
t	1355	0.20	3426	1
kb2	2458	0.00	2660	0
ms	1742	0.00	2816	0

URL *	Average Lived	Average Mated	Best Lived	Best Mated
al	1637	0.00	3470	0
noc	638	0.00	1716	0
fm	560	0.00	1024	0
md	481	0.00	945	0

The results of the combined two $Mind_M$ servers M11 and M12 are shown below (Table 2) [** preceed initials in this column with* http://w2m.comp.dit.ie/services/]

Table 2. Biotech $Mind_M$ Server Results

URL *	Average Lived	Average Mated	Best Lived	Best Mated
M11	2202	0.00	2571	0
M12	2252	0.00	2607	0

The top two biotech minds were *jm* and *t* on the scoreboard. An analysis of these minds showed that the students from a biotechnology background developed diverse agents with different policies being implemented for animal cleanliness, animal temperature, carbohydrate deficiency and animal cell temperature. Unfortunately when these sub-minds were combined based on the two primary goals (*mating* and *survival longevity*) the results show that no significant improvement was made. As a result the $Mind_M$ servers M11 and M12 were significantly better than the other five sub-minds but only ranked alongside the third placed sub-mind. This division of labour between the sub-minds based on profiling is a complex task (Bryson 2002) and was not successful within the biotechnology sub-group. This leads us to believe that profiling information from the sub-minds may be inaccurate. Treating the sub-minds as a “black box” may have no future and if online re-use does take off then people may choose the online minds that are documented clearly as to how they work.

Details of the results of the undergraduate computer science students minds showed greater promise as these minds concentrated on the primary goals of mating and survival longevity and as a result could be easily combined as $Mind_M$ servers and reducing the search sub-space by 50% for each sub-mind.

The results for two such $Mind_M$ servers M7 and M4 are shown below (Table 3). The M7 and M4 servers outperformed all 513 sub-minds that were created during the project and are explained in (O’ Leary et al 2004).

Table 3. Computer Science $Mind_M$ Server Results

URL *	Average Lived	Average Mated	Best Lived	Best Mated
M7	1539	3.70	3956	20
M4	1276	3.67	2583	9

8. DISCUSSION

From the results that have been shown it is clear that it is possible to create superior minds from existing minds if the correct minds are chosen, and if the integration strategy suits the algorithms that are being executed by each of the minds. When the two worked together they appeared to get the best of both worlds. The mating mind sought out opportunities to mate whenever possible, whereas the remainder of the time the other part of the mind was focussed on ensuring that the creature lived long enough to be presented with sufficient opportunities to mate. The main reason that better results were not achieved is that the sub-minds were not dedicated to one specific goal. Selection of a sub-mind such as “mating” or “survival” was based on the test results on the individual minds performance in the simulated environment. Both algorithms used within the sub-minds could not be automatically tested for fitness as “best mater” or “best survivor” without better profiling techniques being utilised. Selections for combination into superior minds was based on the best results achieved by a sub-mind in only ONE of the goal categories, mater OR survivor even though these are competitive goals.

Tyrrell (1993) and Bryson (2000) both built action selection mechanisms that outperformed our collection of minds. However, as mentioned earlier, our sub-minds require much more testing, as does the re-implementation of Tyrrell’s SE. Also, the action selection needs to be more sophisticated. Minds need to be designed to co-operate with the types of action selection that is taking place in other minds. For example, minds may need to flag appetitive and consummatory actions, perhaps by returning strengths or W-values (Humphrys, 1997) with each action.

9. FURTHER WORK

The assignment that was given to the undergraduate students only ended recently, thus limiting the time available to develop more sophisticated Mind_M services. However, it is clear that if minds are available online as services, they effectively exist forever, and can be used by anyone with access to the World-Wide-Web. This should mean that over time, as more runs get executed and more scores are recorded, it will be possible to get a clearer image of the profile of each mind. Further work on multiple minds co-operating socially in a common simulated environment is also on going.

10. SUMMARY

Complex minds are necessarily composed of a multitude of different algorithms and structures (Minsky, 1986). Such minds are difficult to build because the required expertise is rarely available to a single project or lab. The World-Wide-Mind project aims to make it easier to build these types of minds, by proposing the online reuse of existing mind components. The integration of minds from various authors and from various and diverse backgrounds have shown us that distribution of intelligence is possible.

REFERENCES

Bryson, Joanna (2001), *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*, PhD Thesis, MIT

Bryson, Joanna. (2002) *Modularity and specialized learning: Reexamining behavior based artificial intelligence*. In Martin V. Butz, Pierre G´erard, and Olivier Sigaud, editors, Adaptive Behavior in Anticipatory Learning Systems, Springer.

Guillot, A. and Meyer, J.-A. (2000), *From SAB94 to SAB2000: What's New, Animat?*, Proc. 6th Int. Conf. on Simulation of Adaptive Behavior (SAB-00).

Humphrys, Mark (2001), *The World-Wide-Mind: Draft Proposal*, Dublin City University, School of Computing, Tech Report CA-0301 computing.dcu.ie/~humphrys/WWM/

Humphrys, Mark and O'Leary, Ciarán (2002), *Constructing complex minds through multiple authors,* Proc. 7th Int. Conf. on Simulation of Adaptive Behavior (SAB-02)

Minsky, Marvin (1985), *The Society of Mind*, Simon and Schuster

O'Leary, Ciarán (2003), *SOML Specification*,http://w2m.comp.dit.ie/services/documentation/

O'Leary, Ciarán and Humphrys, Mark (2002), *Lowering the entry level: Lessons from the Web and the Semantic Web for the World-Wide-Mind*, 1st Int. Semantic Web Conf. (ISWC-02)

O'Leary, Ciarán and Humphrys, Mark (2003), *Building a hybrid Society of Mind using components from ten different authors*, in Proceedings of Seventh European Conference on Artificial Life.

O' Leary, Ciaran, Humphrys, Mark and Walshe, Ray (2004 to appear) *Constructing an animat mind using 505 sub-minds from 234 different authors*, Proceedings of 8th International. Conference. on Simulation of Adaptive Behavior (SAB-04).

Tyrrell, Toby (1993), *Computational Mechanisms for Action Selection*, PhD Thesis, University of Edinburgh, Centre for Cognitive Science.

Walshe, Ray and Humphrys, Mark (2001), *First Implementation of the World-Wide-Mind, Advances in Artificial Life*: Proceedings of the 6th European Conference on Artificial Life (ECAL-01)