Study of Same-Species Emergent Agent Behaviours in a Controlled Genetic Environment

Daniel Barry
School of Computer Engineering Science
University of Hertfordshire
Hatfield, Hertfordshire AL10 9AB
Email: danbarry16@googlemail.com

***Abstract*—Agents have been observed both naturally and artificially displaying cooperative behaviours within a species. This study aims to understand how and why these behaviours naturally occur and how we may begin to classify them. This is tested by simulating agents in a 3D environment where agent data can be analysed. This study shows that it is possible to distinguish behaviours of agents and begins to show how they can be more optimal working together. There is good grounds to believe agents working together can benefit from having different behaviours in an emergent system.**

Keywords—Artificial Life, Genetic, Emergence, Simulation.

I. Introduction

In nature and in particular the realms of the small, it's seen that simple agents can evolve to show emergent behaviours to better optimise for their environment. Delicate systems that both rely and thrive on cooperative behaviour can interact to create even more complex systems [1] [2] [3].

Of particular interest are agents that depend on one another, where neither is hunter or hunted and the coexistence is sustainable. Agents can be of the same species but yet display different behaviours to benefit both themselves and the greater population. The purpose of this study is to show that a common relationship can be demonstrated that doesn't rely on the usual predator and prey examples seen more commonly in some other studies [6]. With such a relationship, it should be possible to show that cooperative agents can out-compete with non-cooperative counter parts [9].

The agent in this simulation will be a Braitenberg vehicle at its centre it represents both a simple and iconic agent in the field of artificial life. In this sense, the simulation should be kept relatively simple whilst complex behaviours are still possible [4]. My addition to these simple agents will be the ability to mate and therefore evolve [7], coupled with the ability to have a weighting attached from sensors to actuators instead of binary mappings which differs greatly from the original design of Brainteburg vehicles. There'll be no strict concept of generation programmed into the bots as to mimic what is seen in more natural systems.

This effectively gives each agent a very simple neural network with training happening over generations as opposed to on the same network. The reason for using these agents will be to maintain as much simplicity as possible in the already well-proven design, allowing for easier classification of perceived behaviour.

To achieve this, the following questions need to be answered:

- Is it possible to measure agents behaving in similar fashions and categorise not reliant on their "genetic" code, in this case the weights between sensors and actuators? Does other factors like the agent's external environment also make a noticeable difference to the overall behaviour?
- Can we provide evidence of emergence of emergent behaviours working together?
- What are the conditions which emergent behaviours become more optimal by working together?
- If a group of emergent agents were to be removed, if they truly work together, would there be an impact on their collective performance, e.g. their ability to breed?

II. Problem Description

This study tackles the problem of knowing when a behaviour is dependant on another behaviour and whether that relationship is optimal for all sets of agents involved. To do this the agent behaviour first needs to be classified, followed by the recognition of an effect when an agent group is removed from the environment and a measure of success there after.

III. Review of Literature

It has been shown that a gendered genetic algorithm can prove optimal when taking into consideration theory of sexual selection, where agents mate depending on sexual selection, mutation variations between gender and competitive/cooperative outcomes of effectively having two species [13]. This study intends to take this a step further by looking at the cooperative behaviours of agents without the need for effectively independent species.

Another study shows how cooperation may occur if agents are arranged in a predator-prey situation, where predator agents must work together in order to capture prey agents [12]. Whilst the agents do again cooperate to achieve the task, there is a direct reward for being close to their prey in their genetic fitness algorithm, so agents achieve the cooperative behaviour regardless as it's the most optimal thing for a single agent to achieve. This project differs from this idea as an agent can still do well in the environment without achieving cooperative behaviour, so cooperative behaviour is not immediately obvious.

There is another interesting study of giving agents the ability to evolve code instead of genes to achieve cooperative behaviour [11]. This is a nice approach as it effectively encodes agents with a greater intelligence but makes each individual agent much harder to understand when analysing them. It's a non-trivial problem to say how closely related two agents may behave in the environment. Behaviours become much more difficult to classify which is why this experiment uses the much simpler genetic code.

Using a co-evolutionary algorithm was yet another consideration if the world complexity is to increase [10]. In order for a cooperative behaviour to emerge, some agents may have to take important actions that they may not directly benefit from. As the agents aren't very complex and complexity is difficulty to simulate, the other option to take would be to make agents aware of other agents success through use of shared fitness.

IV. Experiment Methodology

A. High Level Concept

This will involve a simulator that can handle multiple agents where there is the possibility to breed, die and be rewarded for a given action. In our case the reward also encourages breeding as many agents will meet where the reward exists. The environment should be simple enough that key concepts may be seen but complex enough that key concepts may occur.

B. Agents Definition

1) Diagrams

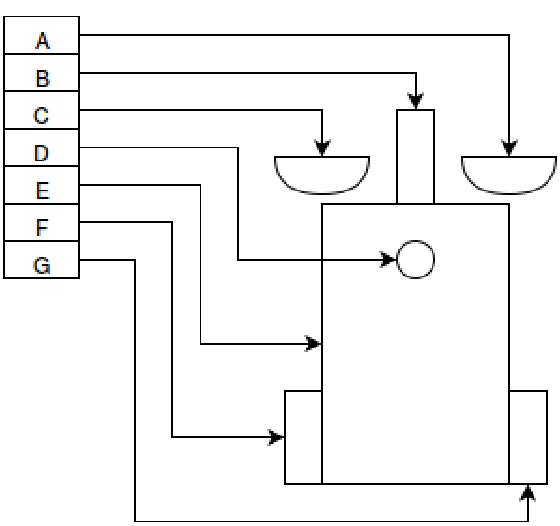

Fig. 1: Vehicle layout and general agent representation

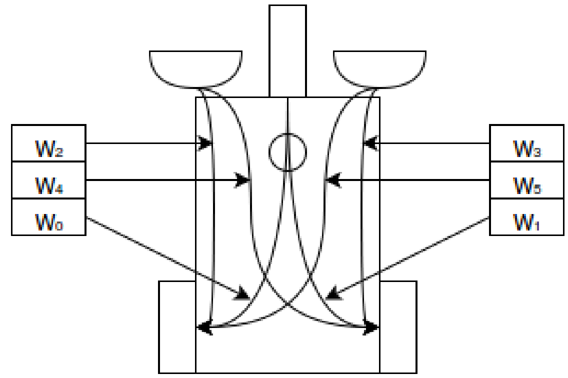

Fig. 2: Vehicle sensor to actuator weighting

(1) A – The right light/temperature sensor. For the purpose of this they are interchangeable, one could think of it as an infra-red sensor where both are satisfied.

(2) B – The mating sensor used to both detect a mate and initiate the making of offspring. In this simulation there is no gender and the offspring appears into existence.

(3) C – The left light/temperature sensor. (See the right sensor for detail).

(4) D – Zero friction "caster-wheel" support object. This is to be under the agent and to allow the agent to move. This is more for the purpose of the simulator than for purpose of the results.

(5) E – The agent's body.

(6) F – The left wheel at fixed velocity, defined by

$$F_{VEL} = F_{MAX\_VEL} * \frac{(W_2 * C) + (W_0 * B) + (W_5 * A)}{3}$$

(7) G – The right wheel at fixed velocity, defined by

$$G_{VEL} = F_{MAX\_VEL} * \frac{(W_3 * A) + (W_1 * B) + (W_4 * C)}{3}$$

Where $W_i >= -1$, $W_i <= 1$, $S_n = \{A, B, C\}$, $S_n >= 0$, $B <= 1$, and $\{A, C\} <= MAX(A_{Reward})$.

The weights in *figure 2* almost resemble a simple neural network, where a genetic algorithm through breeding is how the agent "learns". Knowledge and/or learning is passed from agent to agent through generations such that inter-agent breeding effectively is sharing this knowledge.

2) Overview of Simulator Model

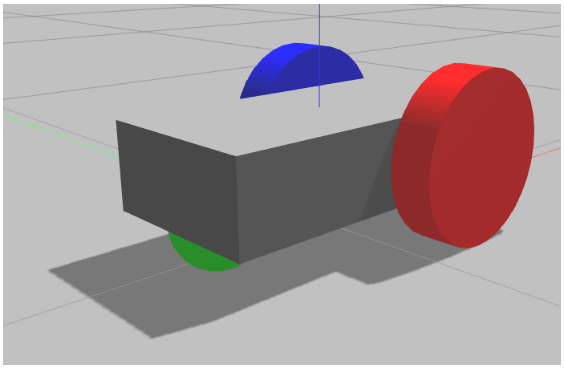

Fig. 3: The vehicle built in the Gazebo environment. *Note that the sensing and mating parts have not been added for the purpose of collision and visualisation.*

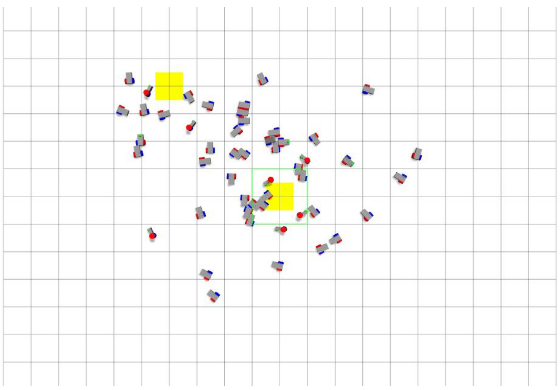

Fig. 4: The agents running during the simulation.

In the two figures we can clearly see an example of an agent and multiple agents interacting with one another. The design of the vehicles has been kept simple for the purpose of simulation speed and visualisation of many agents. This means that with 40 agents the simulation can be run at 3.5 times reality.

There are several note worthy states the agents have been seen in as simptoms of using a 3D environment, including agents that move on their side with the flat edge of a wheel to the ground, agents that move on their back or agents that throw themselves rapidly between many of these states. Whilst these were not designed to occur, they serve as an interesting artefact that could help to counter natural occurrences in nature we also don't count on occurring.

The positioning of the plates isn't particularly important in the environment, they have just been designed such that they encourage collisions between them, one of the factors that help such a small population to continue to breed for so long.

C. Environment Definition

Below is a brief description of the simulation set-up.

Starting Locations – Agents are spawned in a grid like manner around the centre, each facing a random direction. Given that the world is complex, it doesn't take too long before any evidence of a grid disappears within the randomness and complexity of the simulation.

Agent Starting Weights – The starting weights of the agents are $0 <= W_i <= 1$ to ensure that all agents are moving in a forwards direction at the beginning of the simulation. This is to ensure that agents are working correctly. If agents evolve to go backwards, this makes the action look more deliberate. *Mating* – Mating will require energy from a pair of agents after a collision from the frontward direction, where an agent will be created. This will only occur if an agent has previously died as to prevent over population and will require a sacrifice of health from both agents to produce another randomly placed agent within the environment.

When mating, one of the parents is chosen at random to be the dominant parent gene with a small probability of 10% to include the other parent's genes. There is then a further 10% random genetic mutation to accompany this, meaning on average the child agent will have 20% of its genes not belonging to its dominant parent.

A mating collision occurs when each agent is within half the agents width away (measured from the centre of each agent) and ± 0.5236 radians or ± 30 degrees. These numbers only serve to allow relatively successful mating when an agent is in the correct direction and agents facing one another.

Reward – There will be multiple rewards in the simulation with the same benefits for each, offering health to the agent. There will be two static plates in the world, where health is added with respect to distance to the centre of the plate.

$$d = \sqrt{(TileX_i - A_x)^2 + (TileY_i - A_y)^2}$$

$$A_{Reward} = \sum_{i=0}^{N} \frac{TileReward_i}{(1+d)^2}$$

Where $TileReward = 4$ and $N = 2$ for the purpose of simulation, updated once a second.

Health Decline – A consistent and constant health decline ensures that agents will eventually die even in perfect settings. This means they have to pass on their genes in order for their behaviours to survive. Agents lose 0.5 from their starting 100 units of health every second.

Death – This is simply defined by there being no health left ($< 0Units$) for one of the previously mentioned reasons.

D. Environment Measurements

1) Generic Measurements

The experiment should be able to show the following:

Emergence – The point at which agents turn from a random collection and start to show emergent behaviour. This will be done using the measurements below.

Multiple Emergent Behaviours – This will be showing multiple emergent behaviours co-existing in the same environment. This will be measured by looking at the genetic code and the time survival.

Cooperative Multiple Emergent Behaviours – Experiments should show that emergent agents can work together for near mutual benefit and therefore express an even more complex behaviour. This is planned to be done by removing agents from the environment that are identified as of a particular classification to see whether the other agents are affected by their lack of presence.

Stability – Agent populations should be considered stable if they are able to maintain their count over several generations without sign of diminishing. This amount or percentage will be decided at a later date when the simulator is built - the likely deciding factor on the number of agents that can be reasonably supported.

Genetic variation – The genetic variation will be measured visually in terms of entropy in comparison with other agents. Each weight can be compared individually, with variation also being a way of telling whether a gene is redundant as redundant genes will be random or locked out. They will then be measured by taking the average of all significant weights as to look for behaviours.

Emergent behaviour – This will be partially done through visual inspection, but behaviours will also be modelled and classified on their role in the larger group. An attempt will then be made to generally classify behaviours.

2) Agent Measurements

To categorise agents, there are several methods that are planned to be used:

Weights – The weights of the agents will be a good indication of how the agent is supposed to behave in it's environment.

Existence Time – If an agent survives a long time compared to it's closely genetic relatives, this tells us about the performance of the agent and something about its circumstance in the environment.

It's also worth noting that these may not be adequate during the simulation as it depends on what situations arise. There will be an element of classification by observation before a more formal description of an agents behaviour is generated.

V. Results

A. Recorded Data Description

Data from an agent is recorded upon an agents death as the agent collects data about itself throughout it's life. Below is a description of the recorded data from agents.

Agent Number (AN) – This simply describes the order in which the agent died. An agent for example may be the 100^{th} agent to have died, but it may have taken a million years for that to have occurred as agents don't perceive time in that way. The idea if a generation is meaningless in this setting as agents continuously breed independent of one another and are not influenced by external factors.

Breeding Number (BN) – This number represents how many times an agent was able to breed before it died. This may be measured as success as it is very likely that large numbers mean the agent genetics were passed on more for future use.

Time Lived (TL) – This is the time lived by the agent, as it is not obvious whether living long is the best strategy for passing on genetics. This may also be used as a measure for different types of agents.

Genetic Variation (GV) – This is measured by looking at the agents weight and comparing overall connections to indicate whether variation increases or decreases. Individual weights are then used to highlight important genes. This is done by visual inspection for patterns as data is relatively complex.

B. Graphs

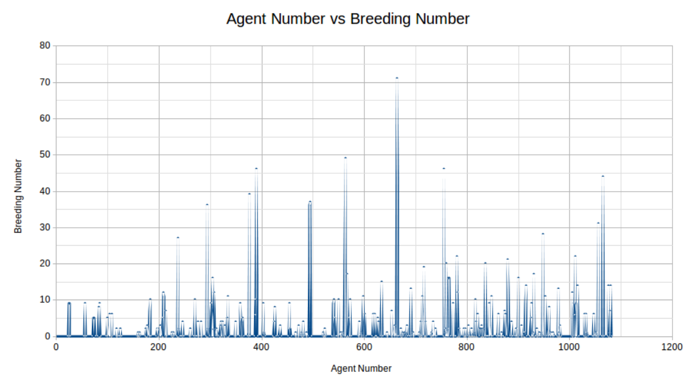

Fig. 5: A comparison in Agent Number and Breeding Number to see whether the number of times an agent breeds before death increases when agents have bred multiple times as part of their general self improvement.

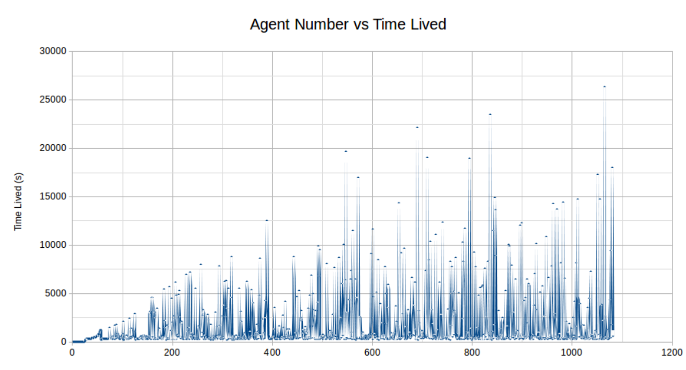

Fig. 6: A view of how long an agent could be expected to live after much breeding has occurred.

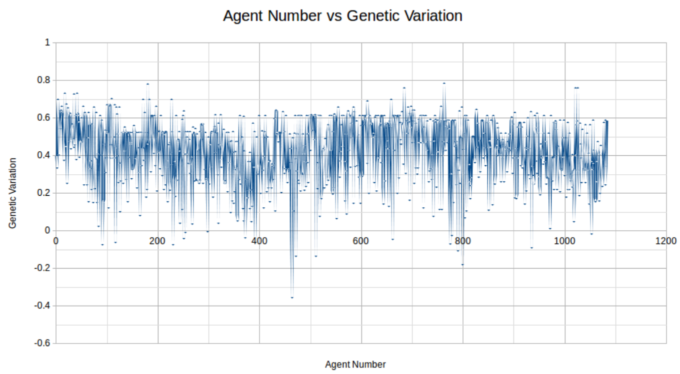

Fig. 7: An extremely simple look at whether the agent weights change by taking the average weight, $\frac{1}{n}\sum_{i=0}^{n+1} W_i$. In this graph we can see fluctuations and therefore have reasonable grounds to investigate further.

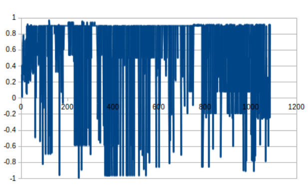

(a) $\overline{A_{W_0}}$

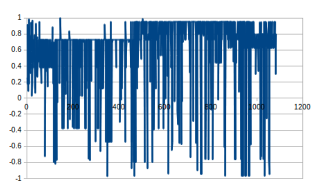

(b) $\overline{A_{W_1}}$

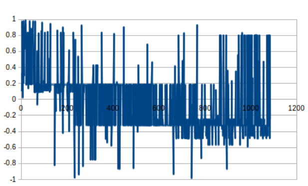

(c) $\overline{A_{W_2}}$

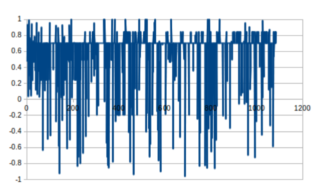

(d) $\overline{A_{W_3}}$

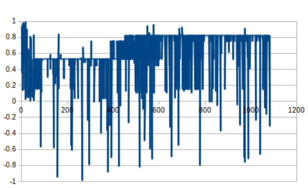

(e) $\overline{A_{W_4}}$

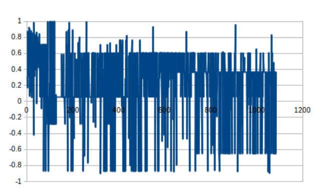

(f) $\overline{A_{W_5}}$

Fig. 8: Average of individual genetic weights where we can see that there is a stepping where a group of agents find a local optima for a weight.

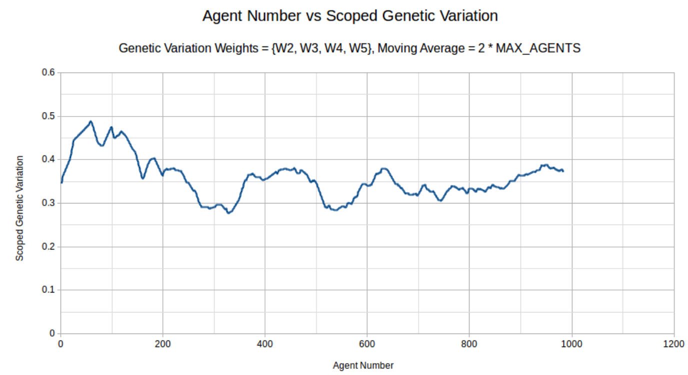

Fig. 9: A look at the average weights W_2, W_3, W_4 and W_5, showing genetic deviance over time. Agents that bred zero times before death are ignored and there is a moving average of $2 * A_{MAX} = 100$.

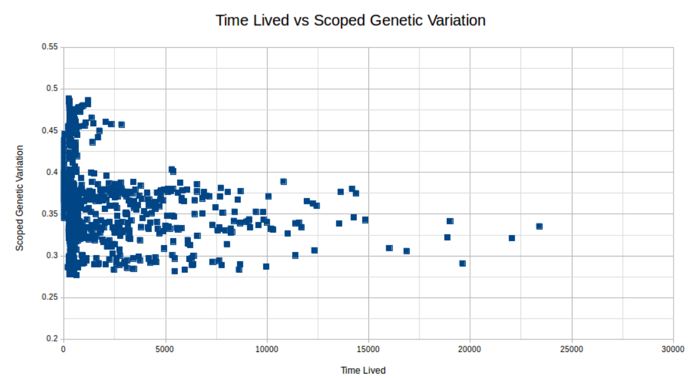

Fig. 10: Another scoped view of the weights, this time being compared to the time the agents spent alive before dying. In this graph we can clearly see two peaks.

VI. Analysis

A review of *figure 5* shows no interesting relationship between breeding over time. Occasionally some agents achieve a higher breeding number, but ultimately the number of times an agent breeds throughout it's life remains the same on average.

Figure 6 shows that agents tend to live longer, which makes sense as this increases their ability to pass on their genetic code. *Figure 7* shows that agents tend change their genetic variance together, meaning there is a relationship that requires further analysis.

In *figure 8* we see that the attraction to a mate, W_0 and W_1 quickly lock out to near enough to a maximum attraction of 1 without too much opposition. Agents that survive value breeding as a high priority concern, so this result is expected. We can then exclude this value from any calculation of genetic variation as these values don't appear to be interesting.

A look at W_2, W_3, W_4 and W_5 reveal these areas of what appears to be stepping, where it seems two or more values were optimal at the same time. Looking a these values more closely in *figure 9* we see that they were becoming less diverse and steadying. *Figure 10* reveals that there were two sets of agents evolving at the same time, one shorter living pair and another longer living pair. On closer inspection in the simulator the longer living pair would sit near to the centre of the reward tile and spin. The shorter lived agents would then take a much larger circling approach around the tiles, colliding with many spinning agents before dying as they were not consistently close enough to the reward tiles to gain much benefit.

The benefit from this behaviour is that the fast spinning agents were more likely to be at the correct angle to mate when the circling agents headed towards them. It seems there is also an optimal number of agents to maximise this relationship, a value that can't be pulled out with such little data. It can be said that the circling agents are a lot less than those that spin waiting to be hit. This is likely from two factors, one being that they are favoured in terms of likelihood to die and another being that fewer agents are required to search for spinning agents.

VII. Limitations

Time – Due the short time allocated to this study, amount of simulation has been limited against producing a report of work done and programming the simulator.

Simulator – The simulator, although having a lot of support, has a reasonable amount of overhead to learning as every simulator has it's own ways of handling each feature. This should have been taken into consideration with the proposal of this project and expected work.

Agents – There are many aspects of the agents that may limit dependant behaviour, including complexity, their ability to mate and choose whether to mate and the randomness in child agent spawning.

Environment – There are many limiting factors to the environment, but to name the most prevalent: Unlimited reward system that doesn't run down or regenerate in some way as this is highly un-natural and how the reward is static in the world. More complexity in the environment would most likely require a more complex agent to overcome it, or simple agents working together.

VIII. Conclusion

It seems as though it is possible to measure agent's behaviour using more than just their genetic code and in fact this is actually an important aspect to classifying them. The external environment certainly makes a difference to the performance of an agent as some agents were "born" disadvantaged far away from the reward tiles. Some agents also appeared facing away from the tiles, ultimately leading to them breeding zero times before mating.

Emergent behaviours have been shown working together which is one of the main aims of this project. Whilst that idea is not unique in itself, having multiple emergent behaviours from agents that share both the same sex and individual fitness measure is interesting. Child agents don't need to have a specified role, e.g. male or female, in order to be able to take on one of many roles within it's environment.

Although not directly tackled, the question of what conditions are required to see cooperative agents is something that the study can start to understand. There needs to be some method in which child agents can retain dominant behaviour from either parent. Early testing indicated that this was extremely important in the development of consistent and meaningful behaviours.

Due to resource constraints outlined by the limitations section, removal of agents within the environment could not be directly demonstrated although it is relatively trivial to see how the two behaviours, spinning and circling, would be very worse off without one another's existence. Not having on of the groups of agents would clearly be detrimental to the mating process.

IX. Critical Evaluation

It would have been good to have been able to run more simulations and do further experimentation but due to limitations in time these have not been covered. It may have also been good to have derived the equations on the graphs also, but again due to time constraints this has not been possible.

There could have also been more use of entropy and information theory in describing genetic variance for the agents as this may have allowed for a more accurate description of behaviours.

X. Further Discussion

One area that should be further explored is whether it is possible for more than two emergent behaviours to coexist in a similar setting, as this would really differentiate the importance of emergent behaviours from a system involving male and female sexes for example. The implications of this could be agents with specialist knowledge to achieve a task where each individual agent doesn't need to be aware of all possible behaviours in order to be able to share knowledge of how to perform a particular behaviour.

Further experimentation should almost certainly involve more complexity and agent intelligence. This will mean that agents will not be able to achieve a task themselves and therefore encourage more than two individual unique behaviours to emerge. This should be thought about carefully as analysis complexity is also likely to increase with the more complex worlds.

Acknowledgment

I would like to thank Paul Moggridge for the kind lending of his hardware, without which this wouldn't have been written and the experiments not conducted.

References

[1] Bonabeau1, E., Dorigo, M., Theraulaz, G. "Inspiration for optimization from social insect behaviour" *Nature* Vol.406, pp.39-42, 6 July 2000.

[2] Couzin, I.D., Franks, N.R. "Self-organized lane formation and optimized traffic flow in army ants" *The Royal Society: London B* 18 September 2002.

[3] Lynch, A. "The Evolution of Cooperative Behavior" *Association for Biology Laboratory Education (ABLE)* Vol.15, pp.319-333, 1994.

[4] Braitenberg, V. "Vehicles: Experiments in Synthetic Psychology" *The MIT Press* 1 January 1984.

[5] Koenig, N. "Robot simulation made easy" *Gazebo Simulator 6.0.0* 2002. URL: http://gazebosim.org/

[6] Bronstein, J.L., Alarcón, R., Geber, M. "The evolution of plantinsect mutualisms" *New Phytologist* Vol.176, pp.412-428, 2006.

[7] Calvez, B. Hutzler, G. "Automatic Tuning of Agent-Based Models Using Genetic Algorithms" *Universite dEvry-Val dEssonne/CNRS*.

[8] Canonical Ltd. "Wily Werewolf" *Ubuntu Operating System* 2015. URL: http://www.ubuntu.com/

[9] Klusch, M., Omicini, A., Ossowski, S., Laamanen, H. "Cooperative Information Agents" *Springer* 27 August 2003.

[10] Dreżewksi, R., Siwik, L. "Agent-Based Co-Operative Co-Evolutionary Algorithm for Multi-Objective Optimization" *Springer* LN AI 5097, pp.288-397, 2008.

[11] Abbattista, F., Abbattista, N., Caponetti, L. "An Evolutionary and Cooperative Agents Model for Optimization" *IEEE* Vol.2, pp.668-671, 29 November 1995.

[12] Julaton III, C. "Evolving Cooperative Strategies in Multi-Agent Systems Using a Coevolutionary Algorithm" *Massachusetts Institute of Technology*

[13] Sánchez-Velazco, J., Bullinaria, J.A. "Sexual Selection with Competitive/Cooperative Operators for Genetic Algorithms" *University of Birmingham*