The Robot in the Swarm: An Investigation into Agent Embodiment within Virtual Robotic Swarms

Adam Lee Newton[1], Chrystopher L. Nehaniv[2], and Kerstin Dautenhahn[2]

[1] A.Newton@worc.ac.uk
[2] Adaptive Systems Research Group
Faculty of Engineering and Information Science,
University of Hertfordshire
{C.L.Nehaniv, K.Dautenhahn}@herts.ac.uk

Abstract. This paper explores the notion of 'degree of embodiment' within the context of autonomous agent research, specifically within swarms of virtual robotic agents. Swarms of virtual robots with systematically varied degrees of embodiment are designed and implemented, and a 3D world created for them. Experimental simulations are then carried out wherein groups of these robots perform swarm tasks, and levels of performance for each group are measured and analysed. Analysis of this data suggests that there is no simple linear or monotonic correlation between degree of agent embodiment and swarm performance (in this particular virtual environment), but rather that an 'ideal' degree of embodiment exists to create a superior swarm behaviour for a given task in a given environment.

1 Introduction and Background

The generally accepted non-scientific definition of embodiment is simply 'to have a physical body', and it is this concept that is extended and clarified by research in many of the Cognitive Sciences such as Philosophy, Artificial Intelligence and, perhaps more specifically, Artificial Life and Robotics. These fields share a common need to describe how an agent, whether biological, robotic, or perhaps one simulated in software, is 'embodied' or 'situated' within the surrounding environment, and how this nature of embodiment affects both agent and environment. A crucial definition of what is meant by the term *embodiment* within the context of agent research appears in (Quick et al. 1999), and it is this meaning that shall be used for the remainder of this paper:

A system X is embodied in an environment E if perturbatory channels exist between the two. That is, X is embodied in E if for every time t at which both X and E exist, some subset of E's possible states have the capacity to perturb X's state, and some subset of X's possible states have the capacity to perturb E's state.

W. Banzhaf et al. (Eds.): ECAL 2003, LNAI 2801, pp. 829–838, 2003.

Recent papers by Tom Quick of University College London and Kerstin Dautenhahn and colleagues at the University of Hertfordshire have further suggested that all bodies, (whether the body of an agent or an inanimate object) are somehow embodied and situated in their environment, but that embodiment is not a binary on-or-off attribute of an agent; rather that certain bodies may exhibit a higher 'degree of embodiment' than others (Quick et al. 1999, Nehaniv 2000, Dautenhahn et al. 2002). For instance, the greater the perturbatory 'bandwidth' connecting agent and environment, the higher the degree of embodiment. This idea is briefly summarised in Figure 1. For more information on this crucial concept, see (Quick et al. 1999, Quick and Dautenhahn 1999, Dautenhahn et al. 2002).

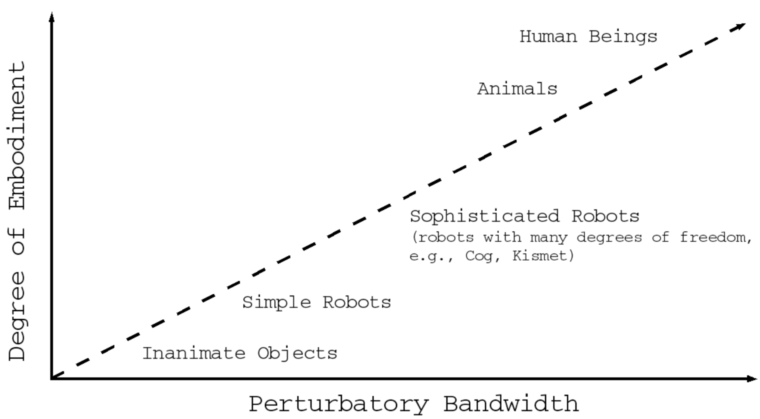

Fig. 1. A proposed correlation between degree of embodiment and perturbatory bandwidth (after Dautenhahn et al., 2002)

This theoretical correlation between the perturbatory channels and the level or degree to which an agent is embodied or situated in its environment might be experimentally investigated using virtual or physical robots. In Figure 1, the lower portion of the graph makes reference to some simplistic robotic agent having a lesser degree of embodiment than a more complex robotic agent. The investigations in this paper focus on this area, and examine whether a robot might be more or less embodied in an environment due to the perturbatory bandwidth between itself and its environment.

The number of perturbatory channels existing between a robotic agent and its environment is related to the number of sensory channels in the agent's physical construction - the number and usage of the various sensors and actuators it employs (wheels, motors, gripper arms, etc.). In the following experimental trials, groups of virtual robotic agents will be configured in a systematically differing manner, some having a higher number of sensory channels, some with a low number. Groups of these agents, placed into controlled experiments, will be given a task to perform. Performance scores, based on a simple metric, record the differences in behaviour from group to group, and this data is analysed and compared within and across the different robot groups. The central research issue this paper attempts to illuminate is how the nature of embodiment of the agents in a swarm changes the behaviour and performance of that swarm.

Though real-hardware Khepera robots were considered an option for this research, the practicalities proved prohibitive. The simulation package 'Webots' from Cyberbotics[1] was chosen due to its well-known ability to closely emulate Khepera hardware within software virtualisation. Webots provides the flexibility and rapid-prototyping facilities of other software simulation environments with the ability to adhere to Khepera reference design, allowing later cross-compilation of controller code to real-hardware robots. The GNU/Linux[2] OS platform was chosen for its reliability and scalability, and also because Linux is the most well supported OS in recent versions of Webots[3].

2 Methodology

When an agent forms part of a swarm, the 'surface' between itself and the environment, structurally connected via the aforementioned perturbatory channels, becomes even more complex. This surface now not only represents the connections between sub-states of the environment (i.e., those states that may effect change in the states of other elements connected via the perturbatory channels), but also (directly or indirectly), the states of the other agents in the swarm. In order to magnify the lower area of the graph featured in Fig. 1, and to identify how Quick et al.'s (1999) 'degrees of embodiment' theory might relate to agents that find themselves not only situated in an environment (virtualised, in this case), but also (unknowingly[4]) forming a part of a swarm, groups of systematically varied agents can be designed and developed to span different levels of sensorimotor complexity, directly relating to the amount of perturbatory bandwidth connecting each agent to its environment (and so, indirectly, to each other agent in its swarm).

The three groups of Khepera-like agents employed in these experiments vary systematically across the spectra of sensor complexity – e.g. sensory bandwidth – from simplistic to complex. Agents from group C enjoy only a low level of sensory complexity, group A enjoy a high amount, while group B is intermediate. The reference design (illustrated in Figure 1) for the Khepera robotic agent has been used as a template, and then modified to create groups A, B and C differing in

[1] http://www.cyberbotics.com

[2] The Sourcemage GNU/Linux distribution was chosen for its high performance.

[3] Version 3.2.x for these experiments.

[4] In the simplest of swarm simulations, no swarm agent explicitly 'knows' of the existence of any other agent. All communication between agents is indirect, through the environment via the process of stigmergy (see Grassé, 1959).

the number of sensory channels (infrared (IR) sensors)[5]. These groups were then populated with ten identical clones to create a uniform swarm of robotic agents. In order to explore the ramifications of design variations in both *number of* and *sensitivity of sensory channels*, each group contains three further variants - this time with the same number of 'channels' of perturbation (the same total number of sensors and actuators) but with the *sensor range* systematically altered; i.e., group A is now made up of A_i (long sensor range), A_{ii} ('normal' sensor range), A_{iii} (low sensor range). Figure 3 illustrates the swarm configurations in more detail.

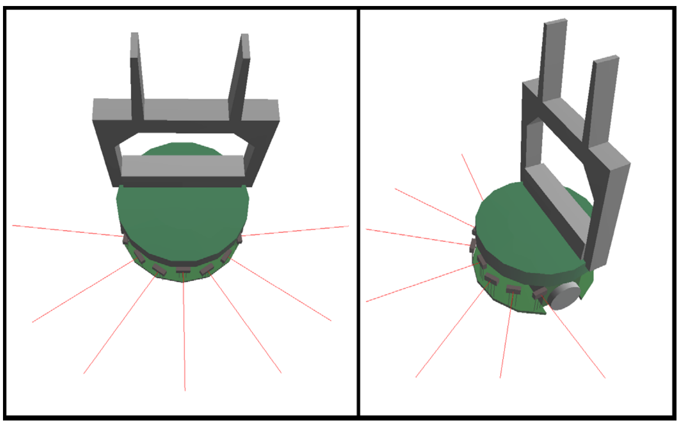

Fig. 2. A simple example of a virtual Khepera robot, modelled in Webots. (Left: top view, Right: side view)

A large arena allows the agents to roam in two dimensions, powered by their actuators. The arena is surrounded by a low wall and populated with a number of blocks, and both of these are objects that the Kheperas' IR sensors detect and respond to in a manner dictated by their controller executable. The environment shown in Figure 4 is typical of a swarm task scenario wherein a particular swarm group is challenged with roaming the arena, finding blocks and clustering them together through a simple, emergent, social-insect-like behaviour[6].

[5] The design approach taken with the 3D, VRML-based bodies and the C-based controller for each type of Khepera-like robotic agent ensured that all controllers could be later cross-compiled to real-world Khepera hardware, if available. All VRML Khepera modelling was based on the reference design supplied by Cyberbotics to ensure real-world validity at the design stage.

[6] This simplistic insect-inspired swarm task can be found in use by ants, termites, and several other species of insects. A swarm of agents roams around, picking up blocks where they find them, and putting them down near other blocks (in insects, sometimes aided by pheromones). Over time, ordered clusters of blocks emerge through a natural process of stigmergy. See Bonabeau et al. (1999) for details.

	Team A - 7 IR	Team B - 5 IR	Team C - 3 IR
Variant i	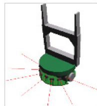 Range: 0.05 units	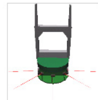 Range: 0.05 units	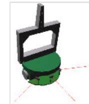 Range: 0.05 units
Variant ii	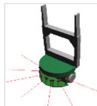 Range: 0.075 units	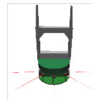 Range: 0.075 units	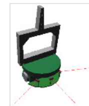 Range: 0.075 units
Variant iii	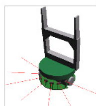 Range: 0.1 units	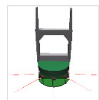 Range: 0.1 units	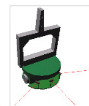 Range: 0.1 units

Fig. 3. Configuration matrix of the nine groups of robots. Each group contains ten agents to make nine swarms of ten robots each. Columns determine number of sensory channels (IR - infrared) for teams A, B, and C respectively, and rows show sensory ranges of variants i, ii, iii for each group.

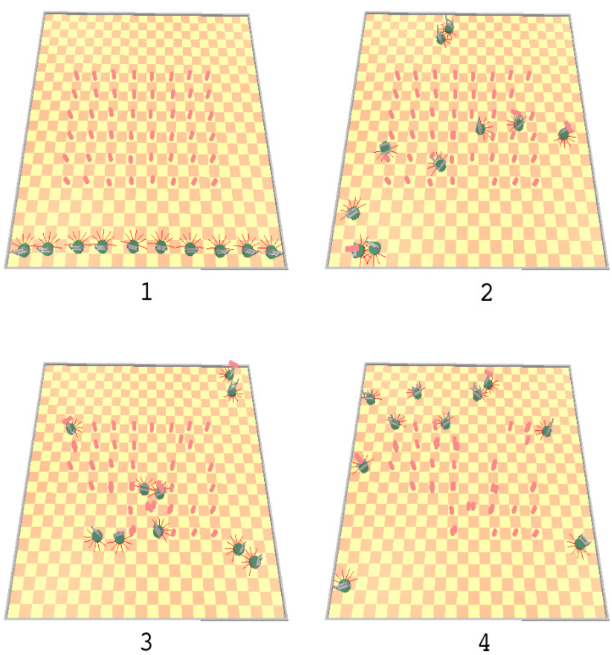

Fig. 4. An example time sequence showing a group of Kheperas at work clustering red blocks, at iteration time index 1: time $t = 0$, 2: $t = 5000$, 3: $t = 10000$, 4: $t = 15000$.

Briefly stated, the insect-like search-grab-move-drop cycle described by Bonabeau et al. (1999) breaks down into several simple behaviours. The agents must roam around, avoiding obstacles, they must be able to identify an object (in this case, the red blocks from Figure 4), pick it up, find another block, and finally put their block down next to the new one before moving away and starting the cycle all over again. Using Arkin's incremental development technique (described in (Arkin, 1998)), each of these behaviours was modelled as a simple computational routine, tested, refined, and put together to give the agents the behaviours they would need to carry out their task without being explicitly designed or programmed to do so. Figure 5 shows a simplified example of this set of behaviours in action. This code, when written in C (using the standard GNU GCC compiler) is then assigned as a controller to each agent. A cross compiler could also be used to bestow similar behaviours in physical hardware Kheperas. Each Khepera now has all the behavioural abilities it needs in order to carry out its task - it simply moves around using a random-walk[7], avoiding obstacles until it finds a block directly in front of it. The gripper arm grabs this block and the Khepera moves on until it finds another block. The block being carried is then placed adjacent to the new block, the Khepera moves away randomly to avoid picking up the same blocks, and the routine is repeated. Each successful loop of this cycle counts as one 'point' (arbitrarily chosen) in a *performance metric*, and each Khepera keeps count of its score during an experimental run of 100,000 Webots timesteps. The total from each Khepera is then collected and a swarm total score generated.

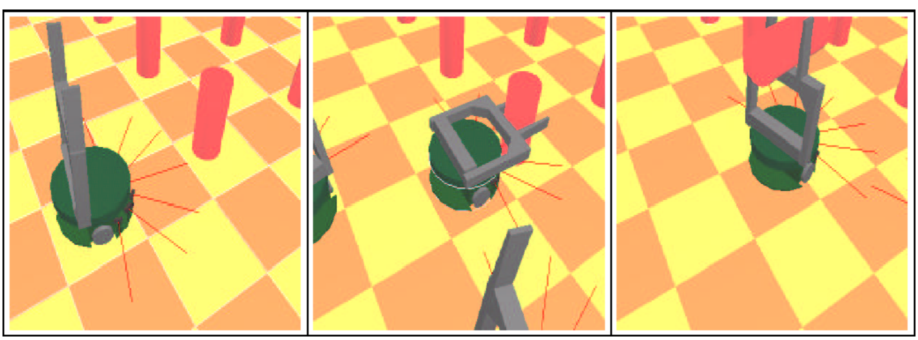

Fig. 5. A close-up on a Khepera carrying out the simple move-grab-move routine.

This process has been largely automated - each of the nine swarms (the three groups with differing number of sensors, each with their three sensor-range variants) carry out the same, standardised collective clustering task in

[7] A random-walk is an insect-like locomotion routine defined in this controller by: rotate-left(random amount), rotate-right(random amount), walk forward. This leads to a very simple 'wiggly', non-deterministic movement routine.

repeated experimental trials; simple Linux-side scripting is used to standardise the start state of the arena and the robots (in terms of their initial locations). Webots, the host software, then takes control and runs the simulation for a set number of controller iterations (common across all experiments). Further Linux-side scripting is used to gather and collate the data recorded by each agent's controller, regarding the number of successful block manipulations carried out within the time limit. These standardised results are then collated and studied.

3 Results

Each swarm repeated the simple collective clustering task fifteen times. Linux-side scripting collected the individual agent scores at the end of each run, and combined them to create a swarm performance total for that run, and an average across all fifteen runs. This data is summarised in Figure 6 showing average performance over fifteen runs for each of the nine embodiment configurations. These results show that no linear or monotonic relationship exists between sensory complexity and performance.

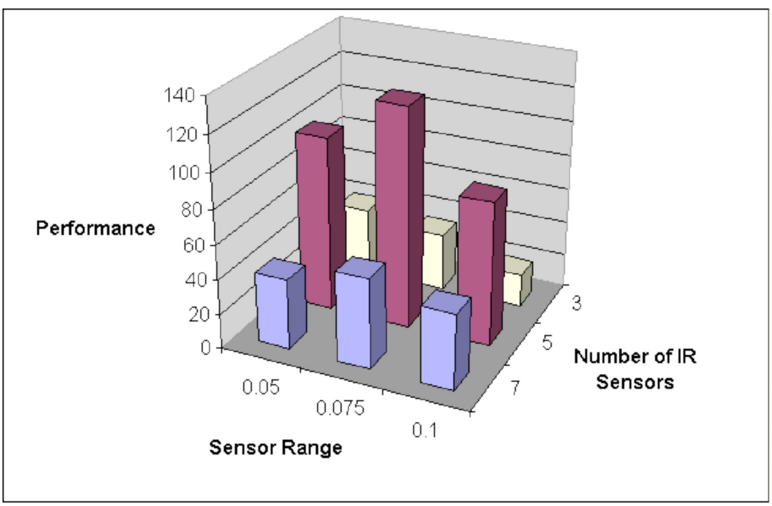

Fig. 6. Comparative Results. Average performance of the nine swarms with differing embodiments (15 runs each). The B team variants, i.e., those with an intermediate number of sensory channels, appear to be noticeably outperforming the other swarms.

Further statistical analysis was used to compare performance across sensor number and sensor range variants (as illustrated in the configuration matrix, Figure 3). Both factors, *sensor range* and *number of sensors*, showed a highly significant impact on performance with $p < .01$ for each of the two factors, although their interaction was not statistically significant (Table 1). Furthermore,

Table 1. Analysis of Variance (ANOVA) in performance for 15 experimental runs for each of the nine embodiment configurations.

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Number of Sensors	134115.6593	2	67057.8296	70.7480	2.53E-021	3.0681
Sensor Range	11114.4148	2	5557.2074	5.8630	3.68E-003	3.0681
Interaction	7718.9185	4	1929.7296	2.0359	9.33E-002	2.4436
Within	119428.0000	126	947.8413			
Total	272376.9926	134				

it is evident from Table 1 that varying the *number of sensory channels* accounted for about 16 times as much of the variance in performance in comparison to varying the sensor range.

4 Discussion

Interestingly, the results do not support an initial hypothesis that an increase in degree of embodiment will result in a superior task performance. Nevertheless, varying perturbatory bandwith (number of sensors in the experiments here) did show a highly significant impact on performance, as did varying the sensor ranges. There was no systematic linear or monotonic increase in performance with degree of embodiment (Figure 6). Rather, there appears to be an 'ideal' level or degree of embodiment that allows a given agent to most effectively traverse the terrain, avoid the arena walls and other agents, whilst successfully repeating the find-grab-drop behaviour needed to carry out the cluster task. It is possible that the variants with the high sensor ranges and larger numbers of sensors are 'too complex' for the task scenario - their movement routines are working against the functionality of the sensors. For example, the random element included in the impulse to each wheel motor at each controller iteration (each timestep) creates a random-walk, a non-deterministic (though overall forward-moving) locomotion behaviour. When sensory bandwidth is as high as found in, for example, group A_{iii}, the sensory input will become highly active as objects wiggle in and out of sensor contact as the Khepera moves. This creates a very good obstacle avoidance behaviour, but in this case means that nearly all sensor input is interpreted as indicative of an obstacle, and the blocks are avoided rather than recognised and gathered. Conversely, the variants with low sensor acuity and range will often mistake the wall or other agents for a block and waste valuable simulation cycles trying to pick it up (an agent cannot pick up a wall or another agent, so this behaviour is futile in terms of achieving the cluster task and simply wastes valuable time).

5 Conclusions

The research question, as stated previously, asked how the nature of embodiment of the agents in a swarm changes the behaviour and performance of that swarm. The experimental data drawn from this research shows that the relationship is definitely not a simple monotonic one, as the variation between scores from team to team was statistically significantly higher at an intermediate level of embodiment, rather than at high or low levels. Team B, with 5 IR sensors per Khepera, was able to out-perform both other teams, and the B_{ii} variant, with extended sensor range over both A_i, B_i, and C_i, was able to produce the overall highest average. This indicates that there is no definite correlation between either degree of embodiment, but that a certain type of embodiment (namely that found in team B_{ii}) happens to yield superior behaviour on average in each agent, and so better swarm performance overall. This is a significant result as it implies that a particular type of agent embodiment allows a swarm to operate most efficiently for a given task. The fact that the more sophisticated, more embodied agents failed to outperform the others suggests that there may be more to the nature of embodiment in terms of performance than sheer complexity (bandwidth and sensitivity) of configuration, and that further refinement of the experimental model described in this paper might shed further important light on the nature of embodiment itself.

So, in closing, the explorations that make up this paper have created questions that need to be investigated further; for example, is swarm B_{ii} the best at other tasks, too? Is the nature of embodiment (i.e., the manner in which sensors are used) more important than overall sensorimotor complexity? Further research, based on a framework similar to that described here, might well illustrate a more concrete connection between the nature of embodiment and the resultant behavioural differences expressed by an agent, and by the swarm of which that agent is a part.

Acknowledgements. The authors would like to thank the following people for their invaluable assistance with the creation and preparation of this document.

Thanks must go to Dr. René te Boekhorst (Adaptive Systems Research Group, University of Hertfordshire) for his assistance with statistical analysis of experimental data. Lastly, Adam Newton would like to thank his classmate and colleague Rehan Mushtaq for his tireless support and schooling in the finer points of C programming.

References

R. C. Arkin, *Behavior-Based Robotics.* MIT Press, Cambridge, Massachusetts, USA, 1998.

E. Bonabeau, M. Dorigo, G. Théraulaz, *Swarm intelligence: From natural to artificial systems*, Oxford University Press, 1999.

K. Dautenhahn, B. Ogden, T. Quick: "From Embodied to Socially Embedded Agents – Implications for Interaction-Aware Robots", *Cognitive Systems Research* 3(3): 397–428, 2002.

P. Grassé, "La Reconstruction du nid et les coordinations interindividuelles. La théorie de la stigmergie", Insectes Sociaux 6: 41–84, 1959.

C. L. Nehaniv, "Evolvability in Biologically Inspired-Robotics: Solutions for achieving open-ended evolution". In: G. T. Mckee and P. S. Schenker (Eds.), *Sensor Fusion and Decentralized Control in Robotic Systems III, Proc. of SPIE*, vol. 4196, pp. 13–26, 2000.

T. Quick, K. Dautenhahn, C. L. Nehaniv and G. Roberts, "On bots and bacteria: Ontology-independent embodiment", *Proc. 5th European Conference on Artificial Life ECAL'99*, Lecture Notes in Artificial Intelligence, Vol. 1674, Springer Verlag, pp. 339–343, 1999.

Quick, T., Dautenhahn, K. "Making Embodiment Measurable". *Proceedings of Fachtagung der Gesellschaft für Kognitionswissenschaft*, Bielefeld, Germany, 1999.