Proceedings Paper:

Wilson, James, Timmis, Jon orcid.org/0000-0003-1055-0471 and Tyrrell, Andy orcid.org/0000-0002-8533-2404 (2018) A Hormone-Inspired Arbitration System For Self Identifying Abilities Amongst A Heterogeneous Robot Swarm. In: IEEE International Conference on Evolvable Systems. , Bengaluru, India

A Hormone-Inspired Arbitration System For Self Identifying Abilities Amongst A Heterogeneous Robot Swarm

James Wilson, Jon Timmis, Andy Tyrrell
York Robotics Laboratory, Department of Electronic Engineering
University of York
York, YO10 5DD, UK
{james.s.wilson, jon.timmis, andy.tyrrell}@york.ac.uk

***Abstract*— Current exploration of adaptation in robot swarms requires the swarm or individuals within that swarm to have knowledge of their own capabilities. Across long term use a swarms understanding of its capabilities may become inaccurate due to wear or faults in the system. In addition to this, systems capable of self designing morphologies are becoming increasingly feasible. In these self designing examples it would be impossible to have accurate knowledge of capability before executing a task for the first time. We propose an arbitration system that requires no explicit knowledge of capability but instead uses hormone-inspired values to decide on an environmental preference. The robots in the swarm differ by wheel type and thus how quickly they are able to move across terrain. The goal of this system is to allow robots to identify their strengths within a swarm and allocate themselves to areas of an environment with a floor type that suits their ability. This work shows that the use of a hormone-inspired arbitration system can extrapolate robot capability and adapt the systems preference of terrain to suit said capability.**

I. INTRODUCTION

Adaptation is commonly used as a technique for improving the performance of swarm robot systems [2]. Commonly, adaptation is performed before a task begins, tuning robot parameters to get an optimal performance for a specific problem as is the case for more genetic algorithms. However, in many situations it is not possible to know the exact requirements of a task and thus it is not possible to perform this tuning before the task begins.

Systems have been proposed that attempt to transfer offline optimization to operate during a task. Some studies have investigated migrating the concept of genetic algorithms, which are used frequently in offline optimization, to allow swarms to adapt mid-task [6], [3]. This adaptation is achieved by giving each swarm member their own virtual genome. These genomes directly affect the behavior of individual robots and by sharing genetic information with one another virtual generations are produced. By choosing appropriate fitness parameters, these genetic systems can promote successful genomes to adapt swarms to a task and obtain a better performance. Another method is to enable members of a heterogeneous swarm to chose tasks based on their abilities [9]. In this system, swarm members bid for tasks they are capable of performing and then work from a ‘play book’ to complete them. Working in this manner allows swarms to form from robots of very different types, creating what the study refers to as a ‘pickup team’.

These online adaptation techniques prove successful within the context of the study’s goals even with robots of mixed ability. However, the previously mentioned studies all require the members of a swarm to have an understanding of at least their own abilities. Having such an understanding is not always possible. During long term deployment factors may change: tires may wear down or robots may experience motor or actuator failure. These factors may change the abilities of individual robots, having a negative effect on their interact with one another or the environment. Changes such as these will most likely cause reduction in performance over extended periods of time unless each robot in the swarm is capable of receiving an up to date diagnosis of their capabilities, using this to modify their behavior to the most fitting option.

Moreover, there may be situations in which there is no opportunity to inform a robot of its capabilities before a task starts. The Triangle of Life project [5] proposed a system in which robots are developed without humans in the loop, suggesting methods that would have robots in a swarm share both morphology and control systems through virtual genomes. In the iterative design presented in these systems, ‘infant’ stage robots have no context for their own abilities. The combination of parent robot morphology, control system and external mutation leave the new generation of robot’s abilities ambiguous.

In the aforementioned cases, greater performance could be achieved with a method of adaptation that does not require the swarm to have any initial understanding of their own capabilities. In this paper a method is presented that achieves this, instead of relying on an initial understanding of their abilities, robots implicitly gain information about themselves and other robots by monitoring the values of virtual hormones.

A system utilizing virtual hormones was proposed in our previous work, capable of arbitrating roles within a foraging swarm [16]. The previous system used hormone values to select either a low-power sleep state or a searching state for each robot with the goal to conserve the overall power consumption of the swarm. Here a new hormone-inspired system is presented that will deal with more complex foraging examples and arbitrate the states of multiple robot types within a swarm. Arbitration will entail the decision between environment types, with each robot in the swarm using hormone values to make their environmental choice. Using a hormone responses to dictate environmental preference is not unheard of, there are natural examples of animals exhibiting exactly this. For example, desert amphibians leave spawn in pools that are intermittently filled and then dried depending on weather. Based on this environmental stimuli (i.e. water availability) hormone levels in the spawn change, accelerating or inhibiting metamorphosis based on the need to stay or leave the pool they are currently in [4].

The system proposed in this work will not only take into account environmental stimuli, but will also use various transmitted hormone values from other swarm members. These values will then be used to gauge the capabilities of each robot individually during the task. With the information gained from the hormone values, the new system will allocate behavior states to each robot based on how suited they are for the task.

II. HORMONE-INSPIRED SYSTEMS

Virtual hormones and hormone-inspired systems have previously been used to directly control the motor functions of a single robot [15]. Hormone-inspired controllers have been successfully implemented to adapt swarm morphology, giving context to environments via stimuli and then constructing appropriate formations [8], [10], [11]. These studies show that hormone-inspired systems can be engineered to provide an effective, computationally inexpensive method for robot control.

Hormone values are constructed from a decay, reducing the level of the hormone value over time, and a stimuli, a condition which when met increases the level of the hormone value. Stimuli can take the form of an interaction with the environment or another robot. Examples of these interactions might include discovering a point of interest, colliding with another robot or the presence of another robots broadcasted hormone value. Some systems might also use inhibitors, triggered by interactions in the same way as stimuli, but instead decreasing the hormone level. Hormone values constructed in this manner are in accordance with the properties highlighted as intrinsic to hormone messages in [14].

The combination of decay and stimuli allow hormone values to fluctuate based on interactions, keeping a live record of the factors related to each stimuli. By using a variety of hormone values, each triggered by different stimuli, the relationships between differently affected hormone values can be examined to extrapolate information about environmental aspects. This information can then be used to educate or create preference within a swarm.

III. HORMONE-INSPIRED BEHAVIOR ARBITRATION SYSTEM (HIBAS)

The hormone-inspired controller we presented in [16] arbitrated states for a homogeneous swarm, allowing each robot to choose between sleeping at a nest site or searching the environment for food. By making this choice, the number of robots foraging was scaled by the hormone system to prevent large swarms from cluttering the environment. By reducing clutter, collisions between robots were less frequent and thus the swarm collected food in a more energy efficient manner. To build upon this work, the new system (HIBAS) removes the sleep state, instead the swarm is presented with the option to explore different environments. In addition to this the swarm is modified to contain different robot types, some more capable in one environment than others. With no prior knowledge of their capabilities individual members of the swarm are able to identify their strengths and form a preference for environment by using the hormone set shown in equations 1, 2 and 3.

Symbol	Meaning
H_x	Hormone for suiting to environment x.
H_{fx}	Food discovery hormone for environment x.
H_c	Hormone for environment/robot collisions.
λ_1	Acts as a decay affecting all environmental preference hormones (H_x), taking a value between 0 and 1.
λ_2	Acts as a decay affecting all environmental preference hormones (H_{fx}), taking a value between 0 and 1.
λ_3	Acts as a decay affecting all environmental preference hormones (H_c), taking a value between 0 and 1.
γ_1	Weighting of H_{fx} that act as stimulus to H_x.
γ_2	Weighting of E_{stim} that act as stimulus to H_x.
γ_3	Weighting of F_x that acts as stimulus to H_{fx}
γ_4	Weighting of C that acts as stimulus to H_c
E_{stim}	An integer variable that counts how many robots in the same state are transmitting suiting hormones that are larger than the detecting robot's.
F_x	Boolean variable that becomes true for a single time step while picking up a food item.
C	Boolean variable that becomes true if something encounters the robots obstacle avoidance sensors.
t	Current time step in experiment.

TABLE I

KEY FOR THE SYMBOLS USED IN THE HORMONE EQUATIONS.

$$H_x(t) = \lambda_1 H_x(t-1) + \gamma_1 H_{f_x} + \gamma_2 E_{stim} \tag{1}$$

$$H_{fx}(t) = \lambda_2 H_{fx}(t-1) + \gamma_3 F_x \tag{2}$$

$$H_c(t) = \lambda_3 H_c(t-1) + \gamma_4 C \tag{3}$$

The subscript 'x' in these equations is used to denote instances where duplicate functions and variables will have to be made. In order for the system to operate, robots require one of these equations for each environmental option they are presented with. By numbering these environments and creating hormone values that relate to each environment, copies of H_x would become H_1, H_2 and H_3 relating to environments 1, 2 and 3 respectively. Other symbols used in the hormone equations are defined in table I.

H_x shown in equation 1 is the primary hormone for controlling environment preference. In a two environment example each robot in the swarm will have an H_1 and an H_2 value. When arriving at a nest site the robot will chose between going to environment 1 or environment 2 based on which of the two hormone values is greater. During a task every H_x value is broadcast from every robot, allowing other members of the swarm to compare hormone values. H_x values are the only values broadcast with H_{fx} and H_c being used only internally by each robot.

In a two environment foraging example (Illustrated in Figure 1), considering a robot with a preference for environment 1, while exploring that environment E_{Stim} would keep track of how many robots are transmitting an H_1 value higher than the robot's own. E_{Stim} then increases the value of every H_x value other than the hormone giving preference to the robots current environment. In this example H_2 would be affected by E_{Stim} while H_1 would not be. This system allows robots to constantly compare their performance in their current environment and, if their performance is relatively poor given the context, start building a preference for another environment. Providing stimuli to hormones unrelated to the environment the robot is currently exploring is crucial. Without this the decay present in each hormone would slowly bring all non-stimulated hormones in the system to 0, preventing any preference for an environment from forming outside of the initial environment choice.

H_{fx} is a stimuli hormone, given two environments, H_{fx} is present in the system as H_{f1} and H_{f2}, feed into H_1 and H_2 respectively. The purpose of H_{fx} is to create a stimulus for H_x which operates across a greater length of time than that of the intial stimulus trigger. This is accomplished by taking the initial impulse of the stimulus received when a robot picks up a food item (F_x) and providing a decaying the value over time.

Stretching the stimuli over an additional length of time is important due to the repelling nature of the E_{stim} variable. If the H_{fx} hormone value was to immediately increase upon picking up food, it would immediately be at its greatest value. This would mean that almost any robot that had not just picked up a food item would be encouraged to change environment preference. With the slow increase provided by H_{fx} the system is able to compare performance between robots no matter their stage in the task. This increase also better mimics the stimuli found in biology. Typically a fast acting neurological signal will trigger the production of a hormone in one organ which in turn will change the production of hormones in other organs, modifying the behavior of the whole organism.

The hormone H_c is an element of the system kept from previous work [16]. Its purpose is to monitor the frequency of collisions in the environment, returning robots to the nest should they encounter an area too cluttered with objects or other robots. Frequency of collisions is monitored with a slowly decaying hormone, stimulated by a boolean value, C, triggered whenever the robots proximity sensors detect an entity. H_c is compared to every H_x value the robot currently stores. Should H_c exceed any of these values, the robot returns to the nest.

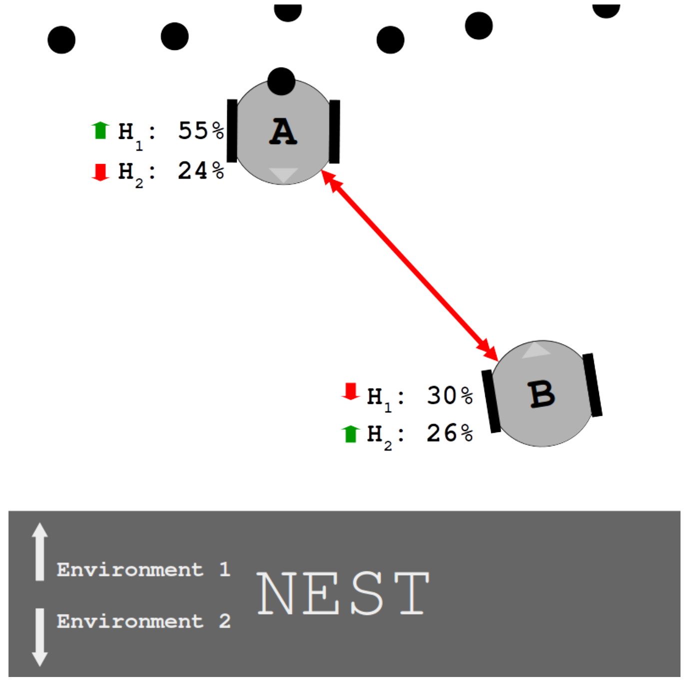

Fig. 1. Two robots, A and B, operating in environment 1. Robot A has successfully discovered a food item and as a results its H_1 value is increasing. Robot B has entered the environment and is under performing relative to Robot A as indicated by its lower H_1 percentage. These robots transmit H_1 values and as Robot B receives an H_1 greater than its own H_2 is stimulated, as Robot A receives an H_1 value lower than its own no such stimulation occurs and H_2 is left to decay. The net result is an encouragement for Robot B to change preference to environment 2 and for Robot A to continue operating in the same environment.

IV. CREATING A HETEROGENEOUS SWARM

In order to test a system in simulation with robots capable of self classifying, robots of different capabilities had to be identified in hardware that could provide a realistic reference point for parameters in the simulated experiments. To keep the system simple, an existing swarm formed from the psi swarm robot [7] developed by the York Robotics Laboratory was altered to allow the attachment of different wheel types to their drive train. A disparity in robot capability was then created by designing different wheel types that could be 3D printed and easily attached, creating a heterogeneous swarm from groups of robots with the same fundamental construction. Once each of the designed wheel types were printed, they were tested in trails of 10. The average speeds when moving in each environment were recorded, allowing loss of traction and instability in these terrains to affect these speed values (shown in table II).

A. WHEEL TYPE: WOOD ENVIRONMENT

The first wheel type (shown in figure 2) was a simple design, the only constraints being that the wheels would have to: allow a robot to travel quickly on a wooden surface and the robot should at least be capable of entering and exiting the grass environment.

The first factor to consider in designing the wheel was the diameter. By increasing the wheel diameter from the 31mm of the standard robot to 60mm, the robot would gain an additional 14mm clearance between the bottom of the robot and the ground (initially 6mm, now 20mm) and have a largely increased wheel circumference. This additional circumference allowed the robots to travel faster and, with the added hight, they were able to effectively transition from a wooden floor to the deep grass whereas previously this was impossible.

The next consideration was the wheel width, given that the wheels were being designed for a smooth surface, there was no real benefit to having wide wheels, as such a thickness of only 3mm was chosen. These thin wheels reduced the amount of force required to turn them by minimizing weight, allowing them to accelerate more quickly.

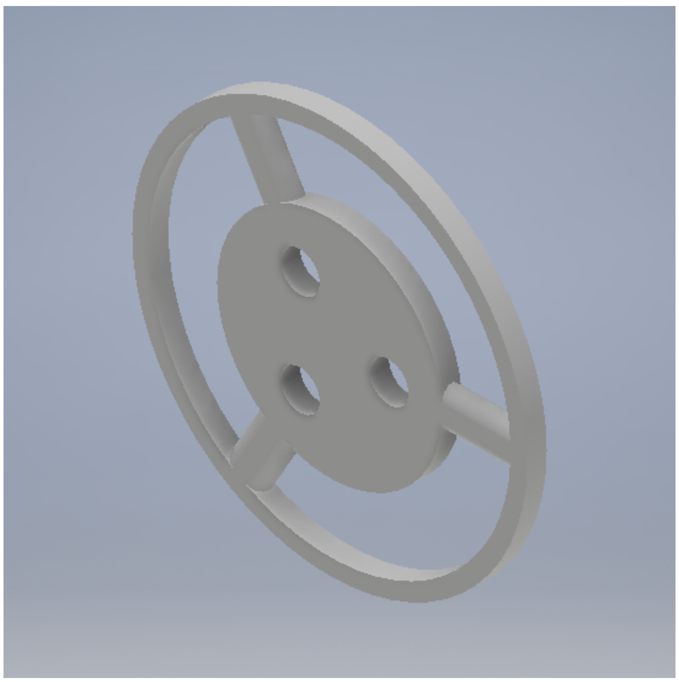

Fig. 2. Wheel designed to specialize in the environment with a wooden floor. Designed in Autodesk Inventor Professional 2018.

Wheel Type	Environment Floor Type Wood Speed (cm/s)	Grass Speed (cm/s)
Wood Wheel	30.9	18.8
Grass Wheel	24.6	21.1

TABLE II

TABLE LISTING THE SPEED RECORDED IN HARDWARE FOR EACH OF THE DESIGNED ROBOTS.

B. WHEEL TYPE: GRASS ENVIRONMENT

The second wheel (shown in figure 3) was a more complex design as it had to perform well in the grass environment. To achieve this, a spoke-like design was produced. These spokes gave the robot additional traction in soft grounded environments, digging into the surface and catching on imperfections in the ground to propel the robot forwards. At 14mm this wheel was also much wider than the wheel designed for the wooden environment. This additional width made the robot much more stable when traveling through the rough grass environment and, with more area in contact with the ground, made it less likely for wheels to fall into divots in the environment, causing momentary wheel slip.

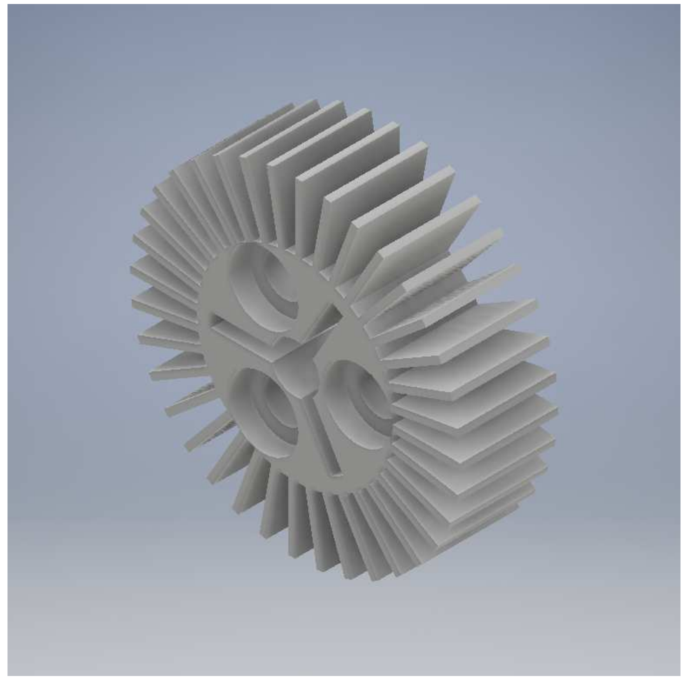

Fig. 3. Wheel designed to specialize in the environment with a grass floor. Designed in Autodesk Inventor Professional 2018.

C. WHEEL TYPE: CONSEPTUAL

This wheel was not designed in hardware but instead was created as a theoretical competitor to the first two wheels. This wheel was designed to travel at a constant speed in any environment. This speed was lower than the slowest speed in either the grass or wooden floored environments but would still travel at the same speed in areas of very difficult terrain where the other two robot types would be much slower. This was an attempt to simulate a robot with either large tracks, very wide wheels or even a robot with hovering capabilities.

V. EXPERIMENTS

To test the proposed HIBAS two experiments were designed. The goals of these experiments were to identify the categorizing capabilities of the HIBAS and the performance increases such a system may give. All of the following experiments were conducted with the parameters listed in Table III.

In each of the experiments the swarms are given a foraging challenged. Their task is to discover, pick up and return food items to a nest area. Foraging was chosen as the task in these experiments so as to build upon our previous work [16].

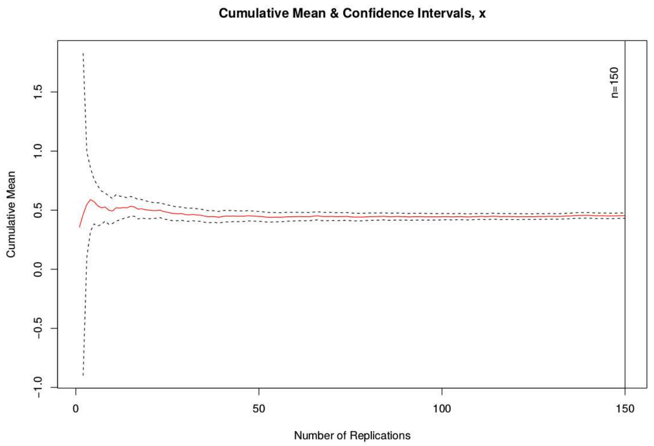

Fig. 4. Graph displaying cumulative mean and confidence intervals for experiment and time step requiring the largest number of replicates

Success in these experiments will be measured by the percentage of correct categorizations made by the systems and by the rate at which food is collected.

Parameter Symbol	λ_1	λ_2	λ_3	γ_1	γ_2	γ_3	γ_4
Parameter Value	0.999	0.999	0.995	0.005	0.1	50	0.2

TABLE III

TABLE LISTING THE PARAMETERS USED IN ALL SOFTWARE EXPERIMENTS.

A. SIMULATION

All experiments were conducted in the ARGoS simulator [12] a multi robot simulator used to simulate large robot swarms. As previously mentioned the robots used in these tests were assumed move at the speeds shown in table II based on the simulated wheel type and terrain. Additionally, it was assumed that each of the robots was equipped with a food sensor, allowing them to identify food items within a 2m radius.

Each experiment was set up to run for 1000 seconds, each simulation time step was 0.1 seconds with samples recorded for every 10 seconds of simulated time.

The number of replicates required for consistent results were determined by performing cumulative mean tests as specified in [13]. By using the cumulative mean of a data set, along with a calculated confidence interval, an estimate can be produced for a range in which the true mean lies. By taking cumulative mean tests across multiple time steps it was indicated that 150 trials would be the minimum number of replicates required for the results of the experiments to be an accurate representation of the simulation responses (graph shown in Figure 4).

B. COMPARISON SYSTEM

To provide baseline data in these experiments a random arbitration system was produced. This system performs the exact same tasks as the HIBAS with the exception of the two following changes:

1) Random Arbitration: Rather than using a series of hormones to decide which environment should be explored by each robot, each robot picks randomly giving each environment an equal weighting.

2) Collision Hormone Threshold: With no H_x values to compare H_c against the system instead uses a flat rate of 10 as the threshold value. If this value is exceeded before a robot finds a food item, the robot returns to the nest and picks another environment to explore at random. This threshold value indicates that the robot has been colliding with a robot or obstacle for a notable amount of time and will ensure robots do not get stuck in a single environment.

C. EXPERIMENT 1

The swarm in the experiment is made up of 7 robots with the wheels specializing in wooden floors and 7 robots with the wheels specializing in grass flooring. The environment for this experiment (shown in figure 5) measures 8mx20m, split into three parts. The two larger areas are both 8mx9m containing 50 food items, each of these areas has a different floor type. The third section is a strip down the middle acting as a nest site measuring 8mx2m.

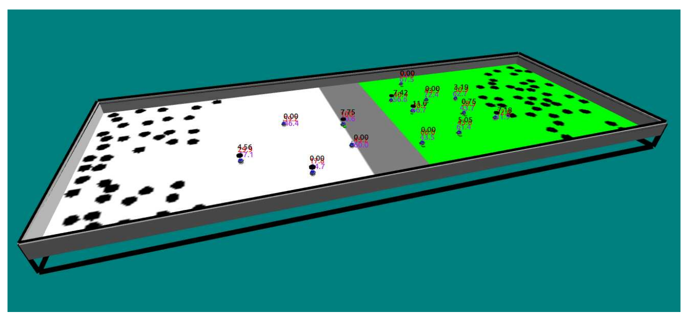

Fig. 5. Screenshot of simulation environment used in experiment 1. Black dots represent food items, white ground represents wooden flooring, green represents grass and the grey area in the center represents the nest site.

D. EXPERIMENT 2

This experiment was designed to test the robustness of the categorization technique. Using the exact same setup as the first experiment, this test had only one difference: when the simulation reached 100 seconds the environment floor types switch. This sudden change should test the swarms ability to categorize once already acclimatized to the environment. This change could represent a landslide or other catastrophe, clearing one side of a task environment but making the other more difficult to travel in. The percentage of correct categorizations from this experiment should be expected to suddenly drop at the 100 second point. A successful system will then steadily increase back to the same or greater categorization percentage than that of the switch point as the system re-adapts.

E. EXPERIMENT 3

The third experiment is a more challenging test of the system. The arena is much larger (see Figure 6), introducing a third floor type for robots to explore with the same dimensions (8mx9m) as the environments in the first experiments. The new environment also included an additional 50 food items, making for a total of 150 in the whole arena. The nest area is also expanded in this arena, measuring 8mx8m to give an equal perimeter to each of the three environments. This experiment also sees the addition of a third robot type bringing the swarm composition to: 5 robots with grass specializing wheels, 5 robots with wood specializing wheels and 5 robots with wheels specializing in the difficult terrain (red flooring). In the new environment, the robots that specialized in the grass and wooden floored environments moved at 11.1cm/s and 8.8cm/s respectively, both 10cm/s slower than in the grass environment to account for additional difficulty.

The purpose of this new environment, while having measurements from hardware, is to give an example of a hazardous area in which two thirds of the swarm are not able to viably operate in. With such a large disparity in ability, successful categorization of robots will benefit the overall performance of the swarm as items are foraged faster and more efficiently.

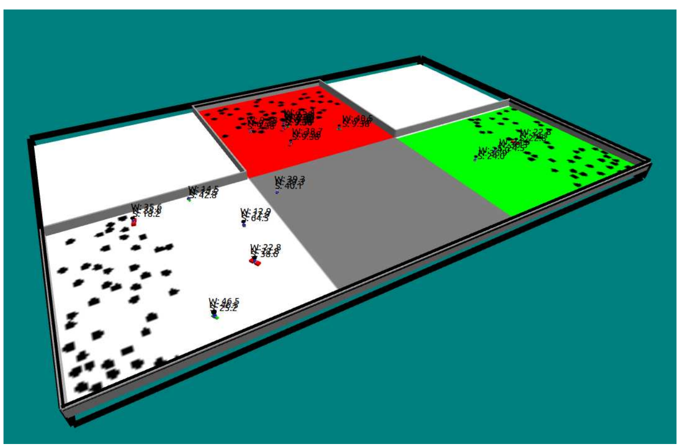

Fig. 6. Screenshot of simulation environment used in experiment 3. Black dots represent food items, white ground represents wooden flooring, green represents grass, red represents very rough terrain and the grey area represents the nest site.

VI. RESULTS

A. Experiment 1

The results from the first experiment are shown in Figure 7. From this graph it can be seen that in experiment 1 the HIBAS outperforms random arbitration in terms of ability to categorize. For the first few samples it appears as though the hormone system performs identically to the random system. To confirm this Wilcoxon tests were performed, comparing the categorization percentage datasets recorded for both systems at each time step. These tests showed that the 600th time step was the first sample with a significant difference between each systems results, giving a p value of less than 0.05. This marked the point at which the HIBAS and random system diverge. This initial starting period is to be expected; it will take time for the HIBAS to begin adapting to the new environment. From the 600th time step onwards it can be seen that the random arbitration remains with a correct categorization percentage of roughly 45% while the the hormone-inspired arbitration increases gradually, peaking at just over 75%. Showing that the HIBAS gives a large improvement to categorization over random allocation.

After this peak, the HIBAS starts to decrease in its ability to categorize environment, falling gradually to just below 50% and then fluctuating near the performance of the random system. This can be explained by the reduction in food items in the environment. As the source of primary stimuli reduces, the hormone system has no reward for item discovery and is therefore unable to accurately categorize. Once this source of stimuli is fully depleted, the system will behave essentially the same as a system arbitrating at random. This drop in performance is of no concern as, when the task nears completion, there is little to no need for the system to categorize successfully.

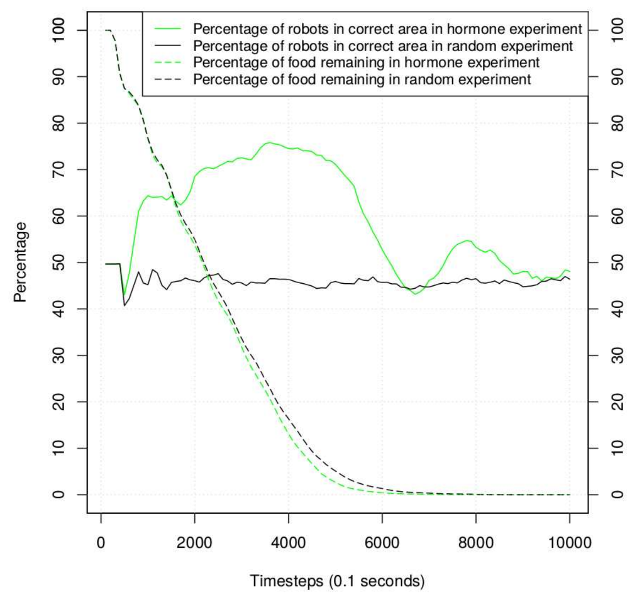

Fig. 7. Graph showing the mean results across 150 trials in experiment 1. The performance of both the random and hormone-inspired system is shown in terms of correct categorization and food items foraged.

B. Experiment 2

In the second experiment, the average performance between random system and the HIBAS is less disparate than the first (shown in Figure 8). By swapping the arena floor types just as the system starts to acclimatize, the hormone values must be re-evaluated by the swarm. This re-evaluation can be seen between the 1000th time step, where the switch occurs, to just after the 5000th. During this time period, robots reallocate themselves as their hormone values decay and it becomes apparent that their performance is lacking in their current environment. After the 5000th time step, in a similar manner to the first experiment, there is not enough food left in the environment for the robots to appropriately categorize themselves and as a result, the percentage of correct categorization begins to tend towards that of random allocation.

These first two experiments shared some commonality in that the rate of food collection was marginal between system types. Performing Wilcoxon tests on the food collection data at every time step showed that, for almost all of the time steps past the 150th, there was a significant difference in the food collection data. However, performing an effect magnitude test using the A-test [1] of the 100 time steps sampled for each experiment, only 22 datasets from the first experiment and 0 datasets from the second had a significantly large difference (an A-test score of over 0.75) from the random systems datasets. This lack of difference in food collection rate is due to the minimal difference in speed between the two robot types. This actually speaks to the benefit of the HIBAS as it was capable of assisting robots in their choice of environment even with a small difference in robot ability.

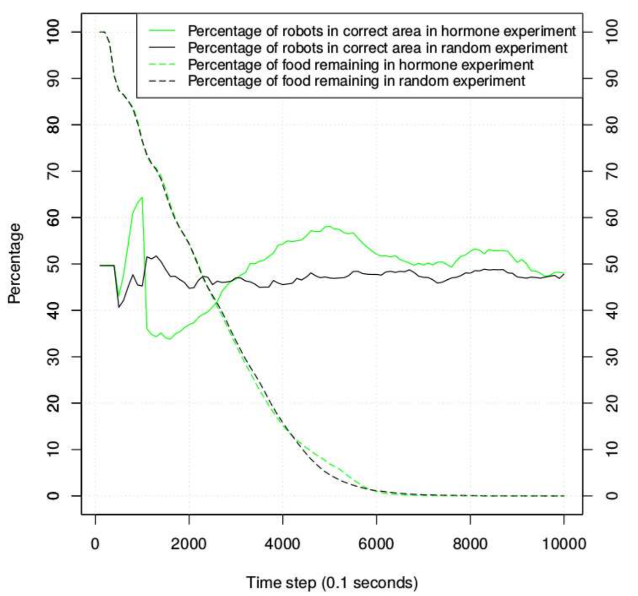

Fig. 8. Graph showing the mean results across 150 trials in experiment 2. The performance of both the random and hormone-inspired system is shown in terms of correct categorization and food items foraged.

C. Experiment 3

The third experiment shows what the system is capable of achieving in a more complex system and how the HIBAS can give a large increase to performance when there are larger disparities in robot ability.

Even given the added complexity in this final experiment (results shown in Figure 9), the HIBAS behaves similarly to the previous two experiments. The categorization percentage takes an initial dip and then begins to diverge from the percentage of the random categorization. However, in this experiment the percentage of robots correctly categorized fluctuates at around 40% which is much lower than the previous tests. With an additional robot type and environment choice this is still a good result as the HIBAS still outperforms the random system by upwards of 10% once the system has adapted.

The third experiment highlights two other key features of the hormone-system. First, due to the increased number of food items in the three environment experiment, not all of the food is foraged. As a result of this, the categorization percentage does not taper off by the end of the experiment. Second, due to the increased difference in robot capability, for the first time in these experiments there is a clear difference in food collection. This is confirmed by running A-tests for each of the 100 time steps sampled showing a significantly large difference between the food collection in the two systems across all results from the 1500th time step onwards.

The results for the third experiment show that, given a large enough difference in capabilities, the HIBAS can provide a large improvement to foraging collection through correct categorization of robot ability to environment.

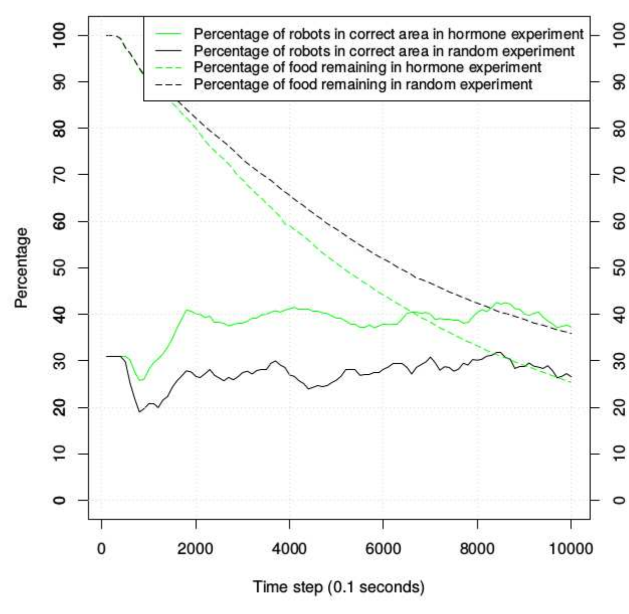

Fig. 9. Graph showing the mean results across 150 trials in experiment 3. The performance of both the random and hormone-inspired system is shown in terms of correct categorization and food items foraged.

VII. CONCLUSION AND FURTHER WORK

This paper has shown that by using a hormone-inspired behavior arbitration system a heterogeneous swarm of robots can categorize their abilities based on their performance in a selection of environments. In all experiments it has been shown that given stimulus availability the HIBAS was able to outperform the random system in terms of percentage of correct categorization. It is also clear from the presented results that, given a simple choice between two environments, the hormone-system is capable of categorizing successfully with even a small difference in robot traits.

By observing the collection of food in the environment with three terrain types and the additional robot, it is clear that:

1) The HABAS can increase likelihood of correct categorization when presented with more complex choices.
2) The categorization provided by the hormone-inspired system can be beneficial to the performance of a task, so long as there is a large enough change in robot capability.

The first step that will be taken to further this work is to test the full system in hardware with the currently designed wheels. Following this, a return to simulation should explore how the HIBAS performs with a swarm operating over much greater periods of time. During these longer experiments the swarm should have its traits modified through either faults or a separate adaptation techniques. Such an investigation will be substantially challenged the HIBAS and its ability to recover from change more rigorously tested.

REFERENCES

[1] Kieran Alden, Mark Read, Jon Timmis, Paul S Andrews, Henrique Veiga-Fernandes, and Mark Coles. Spartan: a comprehensive tool for understanding uncertainty in simulations of biological systems. *PLoS computational biology*, 9(2):e1002916, 2013.

[2] Erkin Bahceci, Onur Soysal, and Erol Sahin. A review: Pattern formation and adaptation in multi-robot systems. *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-03-43*, 2003.

[3] Nicolas Bredeche, Jean-Marc Montanier, Wenguo Liu, and Alan FT Winfield. Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents. *Mathematical and Computer Modelling of Dynamical Systems*, 18(1):101–129, 2012.

[4] Robert J Denver. Evolution of the corticotropin-releasing hormone signaling system and its role in stress-induced phenotypic plasticity. *Annals of the New York Academy of Sciences*, 897(1):46–53, 1999.

[5] AE Eiben, Nicolas Bredeche, M Hoogendoorn, J Stradner, J Timmis, A Tyrrell, A Winfield, et al. The triangle of life: Evolving robots in real-time and real-space. *Advances in artificial life, ECAL*, 2013:1056, 2013.

[6] Evert Haasdijk, Nicolas Bredeche, and AE Eiben. Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics. *PloS one*, 9(6):e98466, 2014.

[7] James Hilder, Alexander Horsfield, Alan G Millard, and Jon Timmis. The psi swarm: A low-cost robotics platform and its use in an education setting. In *Conference Towards Autonomous Robotic Systems*, pages 158–164. Springer, 2016.

[8] Y. Jin, H. Guo, and Y. Meng. Robustness analysis and failure recovery of a bio-inspired self-organizing multi-robot system. In *Self-Adaptive and Self-Organizing Systems, 2009. SASO'09. Third IEEE International Conference on*, pages 154–164. IEEE, 2009.

[9] Edward Gil Jones, Brett Browning, M Bernardine Dias, Brenna Argall, Manuela Veloso, and Anthony Stentz. Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks. In *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 570–575. IEEE, 2006.

[10] S. Kernbach, E. Meister, F. Schlachter, K. Jebens, M. Szymanski, J. Liedke, D. Laneri, L. Winkler, T. Schmickl, R. Thenius, et al. Symbiotic robot organisms: Replicator and symbrion projects. In *Proceedings of the 8th workshop on performance metrics for intelligent systems*, pages 62–69. ACM, 2008.

[11] T. Kuyucu, I. Tanev, and K. Shimohara. Hormone-inspired behaviour switching for the control of collective robotic organisms. *Robotics*, 2(3):165–184, 2013.

[12] Carlo Pinciroli, Vito Trianni, Rehan O'Grady, Giovanni Pini, Arne Brutschy, Manuele Brambilla, Nithin Mathews, Eliseo Ferrante, Gianni Di Caro, Frederick Ducatelle, Mauro Birattari, Luca Maria Gambardella, and Marco Dorigo. ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4):271–295, 2012.

[13] S. Robinson. *Simulation: the practice of model development and use*. Wiley Chichester, 2004.

[14] Wei-Min Shen, Yimin Lu, and Peter Will. Hormone-based control for self-reconfigurable robots. In *Proceedings of the fourth international conference on Autonomous agents*, pages 1–8. ACM, 2000.

[15] J. Stradner, H. Hamann, T. Schmickl, and K. Crailsheim. Analysis and implementation of an artificial homeostatic hormone system: A first case study in robotic hardware. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2009.

[16] James Wilson, Jon Timmis, and Andy Tyrrell. A hormone arbitration system for energy efficient foraging in robot swarms. In *Towards Autonomous Robotic Systems, In Press*. Springer International Publishing, 2018.