RESPONSE TO CHANGES IN KEY STIMULI THROUGH THE CO-EVOLUTION OF SENSOR MORPHOLOGY AND CONTROL

GARY PARKER, CONNECTICUT COLLEGE, USA, parker@conncoll.edu
PRAMOD J NATHAN, TEXAS A&M UNIVERSITY, USA, pjnat@tamu.edu

ABSTRACT

Co-evolving a robot's sensor morphology and control program increases the potential that it can effectively complete its tasks and provides a means for adapting to changes in the environment. In previous work, we presented a learning system where the angle, range, and type of sensors on a hexapod robot, along with the control program, were evolved. Although three sensor stimuli were detectable by the system, it used only two due to the relative importance of these two stimuli in completing the task. In the research presented in this paper, we used the same system, but reduced the availability of a key stimuli; the most effective solution now required the use of all three stimuli. The learning system still performed well by pacing sensors appropriate for the third stimuli and creating a program that utilized these sensors to successfully solve the problem.

KEYWORDS: Co-evolution, Sensor Morphology, Control, Legged Robot, Learning, Hexapod, Adaptation, Stimuli, Evolutionary Robotics, Evolutionary Computation, Genetic Algorithm

1. INTRODUCTION

The link between the morphology, the controller and the environment is an important one [1,2]. For a given task, not all detectable stimuli in the environment are useful. A successful agent should pick out relevant information related to its task while being energy efficient. In co-evolving sensor morphology and control, this efficiency includes the type, number and sensitivity of sensors used, so that the agent can detect enough information to allow the controller to learn reasonable solutions to the task. Discarding capabilities and information unnecessary to the task decreases controller complexity making the agent more efficient.

Evolving sensor morphology has been applied to a collection of the same type of sensors [3,4], optimization of a compound eye [5], and the simultaneous design of the controller and sensors of a robot [3,10,7]. However each of these studies utilizes only one type of stimulus from the environment to perform its task. Using a single stimulus as a key to the environment reduces the agent's capabilities and ability adapt to environmental change, especially if that stimulus is no longer present in the environment or changes radically. In this paper we use three different stimuli from which the agent can choose.

Balakrishnan and Honavar [8] evolved the position and ranges of sensors in a highly limited and discrete simulated block environment. Bugajsaka and Schultz [9] showed strategies to find a general sensor morphology for any environment utilizing obstacle detecting sensors and an onboard sensor to detect distance to goal. Mautner and Belew [10] used co-evolution of the robot controller and the sensor morphology in an environment of constant complexity. Sugira et al. [11] used co-evolution of sensors and a neural network controller to show that the environment affected the sensitivity and resolution of the sensors. This study, contrasting with the above studies, tackles the ability of an agent to adapt to a significant change of stimuli in specific environment configurations of varied complexity. Our study introduces a significant amount of noise in the environment unlike the above studies. In addition, our agent does not have a sensor that detects its distance to goal; it must learn to find the goal using cues available in the environment.

This paper continues work [12] that examines the utilization of a genetic algorithm to automatically design a suitable sensor morphology and controller for a given task in environments of varying complexity. The problem used in these studies involves the navigation from a start point to a target point in environments with different obstacle configurations. In previous work [12], of the three stimuli available in the environment (obstacles, infrared (IR) light and ultraviolet (UV) light) the agent used only two stimuli (obstacles, UV light) to reach the target point. The IR and UV stimuli were placed in positions to facilitate the navigation task. The IR stimulus was placed to allow the robot to navigate around the obstacles; the UV stimulus was placed close to the target point. In this paper, we test the ability of the system to learn in environments where a heavily utilized stimulus (UV light) has changed to decrease its usefulness in the navigation task. We do this by increasing the height of the obstacles, thus not allowing the robot to detect the UV stimulus when not in its line of sight. With this important stimulus altered we verify the ability of the system to learn with this significant change in the environment.

An agent needs to determine which stimuli are needed to complete the task at hand, how to sense the stimuli and how to utilize the chosen stimuli to solve the task at hand. These three goals are achieved in this study in the following manner: (i) by pruning unnecessary sensors the agent increases its efficiency in terms of power consumption and controller complexity; (ii) evolving the angles and ranges of the sensors allows the agent to choose a more favorable sensor morphology for the given task; (iii) by co-evolving the controller the agent uses the chosen stimuli to perform the task with a certain level of robustness to noise in the environment.

2. THE ENVIRONMENT

The colony space is a 3m X 3m walled area in the Connecticut College Robotics Laboratory, represented by a 300 X 300 area in simulation with the top left corner marked as the (0, 0) coordinate [Fig 1]. The actual colony space can be equipped with robots and any reasonable number of obstacles and stimuli.
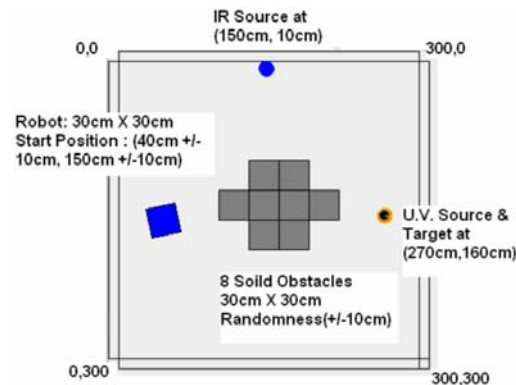

Fig. 1. Environment Description (Central Mountain Configuration)

Eight 30cm x 30cm x 10cm obstacles in four distinct configurations (8% obstacle density) and two omni directional light sources (infrared at position (150, 10) and ultraviolet (270, 160)) made up the environment configurations. For each generation of training, the obstacles were moved randomly by a factor equal to or less than +/- 10cm in the X and the Y direction to give the results a level of generality. The four obstacle configurations chosen allow a comparison with the experiments in our previous study [12]. In order to compare the adaptability of the system, four environmental setups (Central Mountain, Single Left Ridge, Single Right Ridge and Double Ridge) were used [Fig 4]. In our previous study [12], the obstacles did not block the light sources, as the light sensors could detect the light stimuli from above the solid obstacles. In this paper, the obstacle heights were increased to block any light source from reaching the light sensors unless

they were directly in its line of sight. This is a significant change because the UV light is on the target point and acts like a beacon to the target position. In our previous experiment the IR light source was ignored by the agent, though it was positioned to enable easier navigation in four of the obstacle configurations tested. By blocking the light from reaching the agent's sensors, unless in line of sight, we test the system's adaptability is when a significant stimulus like the UV light is appreciably diminished.

2.1 The Agent

The simulated robot closely models the ServoBot [13], a six legged robot with 2 degrees of freedom per leg. A gait for the ServoBot was evolved and used to generate 16 gaits that include 7 left turns, 7 right turns, a straight gait and a reverse gait. The resultant movement of the robot after a single step of each of these gaits was measured and stored. Given a start position and orientation, each gait corresponds to different end positions and orientations after a single step. As the locomotion of the ServoBot is non-deterministic due to inconsistencies in its build and environmental variances, the turn values were randomized by +/- 0.2 units in the simulation.

All sensors are attached to a sensor base, a 30cm x 30cm plate that is attached to the top of the robot. This sensor base, which completely covers the top of the ServoBot, acts as an easily reconfigurable platform for the sensors. It can carry modules of 4 tactile sensors, 4 infrared sensors, and 4 ultraviolet sensors, all interfaced with a sensor controller (Basic Stamp II). Using a Basic Stamp II along with the sensor modules restricts the number of sensors used to 12. The sensors are all binary; they either detect a stimulus or do not.

The light sensors have a range that is adjustable from 0 to 434cm in 14cm increments. The maximum range of these sensors is slightly more than the hypotenuse of the Robot Colony (424cm). The range of each tactile sensor (length of the feeler) is 25cm. An example of the placement, maximum range, and spread of the sensors on the sensor base is shown in Figure 2. The IR sensor and the UV sensor range and spread overlap in the diagram. The light sensors are placed at the sides of the sensor base at any angle relative to the robot's heading. The tactile sensors are placed in the corners of the sensor base at any angle relative to the robot's heading. The GA determined which sensors would be activated and their range (for light) and orientation. The spread of the sensors was not evolved since no mechanism was in place to adjust this aspect of the sensor. The evolved characteristics allowed the sensory information to be complex enough for the robot to be successful in the environment while making the simplifications necessary to allow ease of transfer to the actual robot.
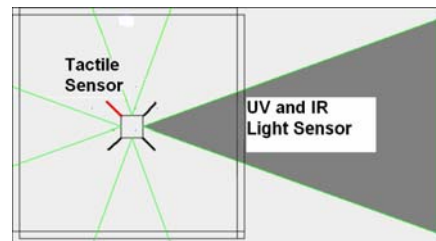

Fig. 2. Sensor Configuration (Range, Spread and Placement of Sensors on Sensor Base).

The controller of the ServoBot and its simulation is a reactive system that uses 13 rules of the form: If (Sensor detects a stimulus) then (Trigger gait number X), in which each sensor is associated with a specific rule and a single gait. The consequents of these 13 rules are selected by the GA from the 16 available gaits. The GA also selects the default gait, which is the gait used when no sensors are triggered (rule 13).

3. CO-EVOLUTION OF SENSOR MORPHOLOGY AND CONTROL

The evolutionary method, the chromosome and fitness function were kept the same as in the previous experiments [12] to test the adaptability of the system.

A 212 bits long chromosome, divided into 35 sets (1 set of 12 bits, 12 sets of 9 bits, 12 sets of 5 bits and 13 sets of 4 bits), was used. The first 12 bits (1 set of 12 bits) represented which of the 12 sensors to keep running during the length of the run. The next 108 bits (12 sets of 9 bits) represent the angles at which each of the 12 sensors are to be placed onboard the robot base as all the sensors can be rotated 360 degrees (512 values per sensor, if values higher than 360 are chosen the program uses the chosen number minus 360). 40 bits (12 sets of 5 bits) represent the range of all of the light sensors (32 values for each light sensor). The last 52 bits (13 sets of 4 bits) represent gaits which are consequents for each of the 13 rules. The position of the sensors on the base are fixed as shown in Figure 2 and do not change.

All 35 sets of bits underwent a single point crossover using stochastic (roulette wheel) selection of the parents. Rule selection had a 1% mutation rate. The other parameters had a mutation rate of 1% if any individual in the generation reached the goal otherwise it was 6%.

An agent successful in finding the target has a fitness based on the number of sensors it has off and the amount of time it took the agent to get to the target, plus a bonus. To achieve the maximum fitness, the robot has to have all its sensors off and reach the goal without any time being used, giving it a theoretical maximum fitness of 15600; this scenario is impossible. The fitness of an unsuccessful agent is dependent on how far away the robot is from the goal, giving a maximum achievable fitness of 20 * 313.8 where 313.8cm is the farthest possible distance in the Robot Colony from the target point. An unsuccessful agent has a maximum fitness of only 6276.

If (Agent_Reached_Goal)
   Fitness = 50 * Number_Sensors_Turned_Off +
            50 * (Total_Time – Time_to_Achieve_Goal) + Goal_Bonus
Else
   Fitness = 20 * Resultant_Distance_from_Goal

Fig. 3. The Fitness Function

The population consisted of 256 individuals, evolved for 512 generations in each environment. Each individual had 3 chances, starting at a random heading and positioned within +/- 10 of the start position of (40,150). The agent had 300 time intervals to complete the task. A time interval is the time it takes the Servobot to complete one step. A step starts with the legs in a ready to step position (right front, right back, and left middle legs forward, and the remaining legs back) and returns to this position after a full step cycle. Each of the 16 gaits completes a cycle in the same time. The time intervals elapse and the test run continues even if a collision occurs. The agent can get out of a collision as by working its way free of the obstacle, due to continual motion and its non-deterministic gait,. The run is stopped if the target position is reached or after 300 time intervals lapse. The entire experiment was repeated 5 times with random starting populations to check for consistency of the results. Individuals were selected for crossover stochastically with the most fit having the best chance of selection. In addition, the individual with the highest fitness was added to the next generation without change.

4. AGENT PATHS AND STRATEGIES

Figure 4 compares a sample of the evolved agents after 512 generations. Each of the evolved agents in the second experiment use the IR light as expected. It is interesting to consider the paths and strategies evolved for each environmental configuration.

1.1 Central Mountain 1.2

2 Tactile sensors at 22º and 333º 1 UV sensor at 101º; Length 378 cm 4 Gaits and Rules Used 42 Steps	2 Tactile sensors at 326º and 337º 1 UV sensor at Length 126 cm 2 IR sensors at Lengths 98 cm and 84 cm 6 Gaits and Rules Used 54 Steps

2.1 Single Left Ridge 2.2

2 Tactile sensors at 21º and 128º 1 UV sensor at 252º; Length 406cm 4 Gaits and Rules Used 51 Steps	1 Tactile sensor at 327º 1 UV sensor at 76º Length 182cm 2 IR sensors at 308º 264º Lengths 280 cm and 70 cm 5 Gaits and Rules Used 53 Steps

3.1 Single Right Ridge 3.2

2 Tactile sensors at 287º and 319º 2 UV sensors at 117º and 32º; Lengths 378cm and 204cm 5 Gaits and Rules Used 42 Steps	1 Tactile sensor at 331º 2 UV sensors at 312º and 111º Lengths 168 cm and 378 cm 1 IR sensor at Length 322 cm 5 Gaits and Rules Used 36 Steps

4.1 Double Ridge 4.2

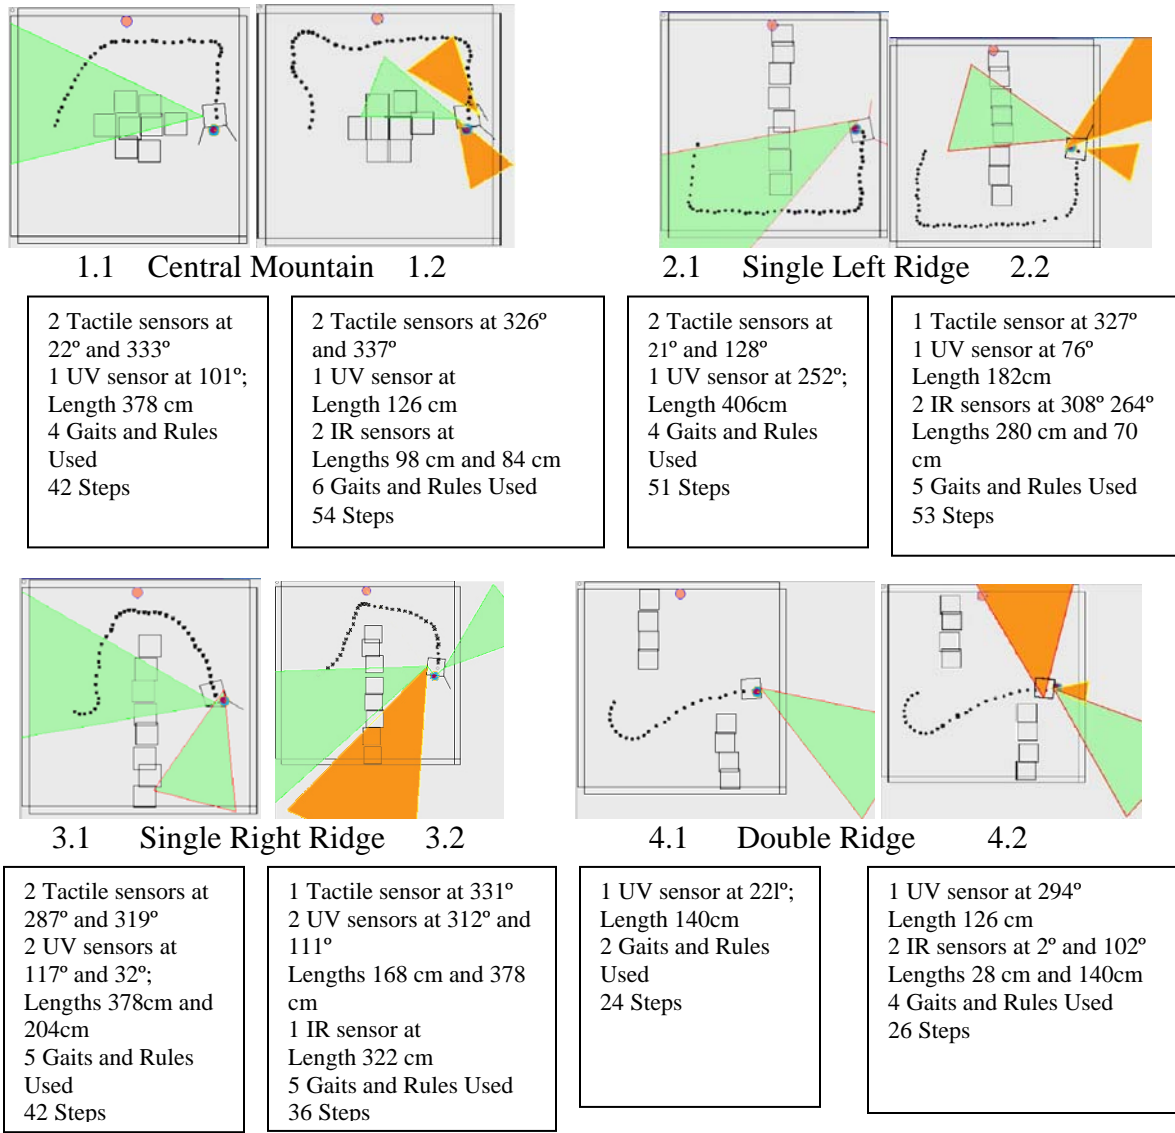|	
1 UV sensor at 22lº; Length 140cm 2 Gaits and Rules Used 24 Steps	1 UV sensor at 294º Length 126 cm 2 IR sensors at 2º and 102º Lengths 28 cm and 140cm 4 Gaits and Rules Used 26 Steps

Fig.4: Paths taken and sensors used by the agent to reach the target point in the 4 environmental configurations. This figure shows a side by side comparison of the agent's strategy in the previous experiment [12] (left) and our current experiment(right) that tests the adaptability of the system.

Central Mountain [Fig 4 Part 1.1, 1.2] shows 1.2 having a much less efficient path to goal than 1.1. This was expected as the task was more difficult due to the increased height of the obstacles. The 1.1 strategy used the UV light to determine a heading to the wall followed by wall following to the goal. The 1.2 strategy used the IR light. The first IR sensor was used to avoid the obstacles while moving toward the wall for wall following and the second IR sensor to correct the path once the robot was moving away from the IR light on its way to the goal; the UV sensor use was of minimal. In 1.2 the UV sensor was used to avoid the central mountain obstacle.

In the Single Left Ridge [Fig 4 Part 2.1, 2.2] the number of steps used by the agent was almost the same in both the experiments. Since the obstacle was high enough to block the IR and UV light from being detected, the UV light is completely blocked at the start position and the IR light is partially blocked. In the 2.2 solution, the robot used two IR sensors because of the

location of the solid obstacles. To avoid going the wrong way (toward the IR light) the robot used the two IR sensors to detect the IR light. The longer range one is used to check if it is moving the wrong way toward the IR light and the shorter is used to check if it is moving the wrong way and is close to the IR light. Moving away from the gap in the Single Left Ridge is a strong possibility as the start heading of the robot is randomized and it could start facing the wrong direction while being unable to detect either the UV or the IR light. Once the path of the agent is corrected to move towards the gap in the obstacles, the agent uses the tactile sensor to perform wall following until the UV light is detected which then causes it to turn towards it and the target point.

An agent evolved in the Single Left Ridge configuration compared to one in the Single Right Ridge configuration uses the IR light can be used to avoid the obstacle in both cases, but in the Single Right Ridge configuration the IR light marks the opening in the ridge. The agent can use a direct path to the IR light, which puts it in a good position to turn toward the goal (Fig 4 Part 3.2) when the wall is detected and then use the UV light to guide it toward the goal. In the case of the Single Left Ridge (Fig 4 Part 2.2), the lack of information from the UV light also required the utilization of the IR light, but due to its distance from the gap in the obstacles and the randomness in obstacle placement, results are not as consistent for the Single Left Ridge.

In the Double Ridge Configuration we can see that the paths are nearly identical as is the strategy for finding the target point. The steps taken and angle of the UV sensor are very similar. In Fig 4 Part 4.2 the IR sensor with the longer range is used to correct the path as the agent moves into position to utilize the UV sensor to reach the target point. The shorter range IR sensor is used to dramatically correct the heading of the robot if due to the randomness of the obstacle placement (+/- 10cm) the UV and IR stimuli are blocked from the two long range light sensors.

5. RESULTS

Figure 5 shows the fitness of the elite and the average fitness of the population for every generation. The consistency of the results is shown by the error bars which are calculated as the standard deviation over the 5 times that the experiment was run. These graphs show that in each of the four environments the co-evolution strategy not only successfully managed to complete the task but also had results comparable and similar to those of the previous experiment [12]. Central Mountain shows a nearly identical elite fitness after 75 generations. The average fitness though was a little higher in the previous experiment. In Single Left Ridge both the average and the elite fitness are lower in the current experiment as compared to the one in the previous experiment and show less consistency over 512 generations. This can be attributed the lack of stimuli in the environment as the UV light is blocked and the IR light can be blocked due to the random factor of the placement of the environment solid obstacles. In spite of this, the elites achieve 70% fitness due to the wall following strategy used until the UV light is found. The agents in Single Right Ridge performed better than those in the previous experiment due to the use of the IR sensor and the location of the IR light, whereas the agents in Double Ridge performed nearly identically.

6. CONCLUSIONS AND FUTURE WORK

The co-evolution strategy designed successful agents that achieved the goal efficiently in terms of number of steps and sensors used. This can be clearly seen, as in all the four environmental configurations the elites were at least 70% as fit as those in previous experiments and the average fitness of the agents showed an increase in each of the 4 environments. The strategies also increased sensor usage compared to our previous experiment to compensate for the increase in complexity of the environment (increased obstacle heights) by taking advantage of the available stimuli. The agent had to have more sensors to be effective.
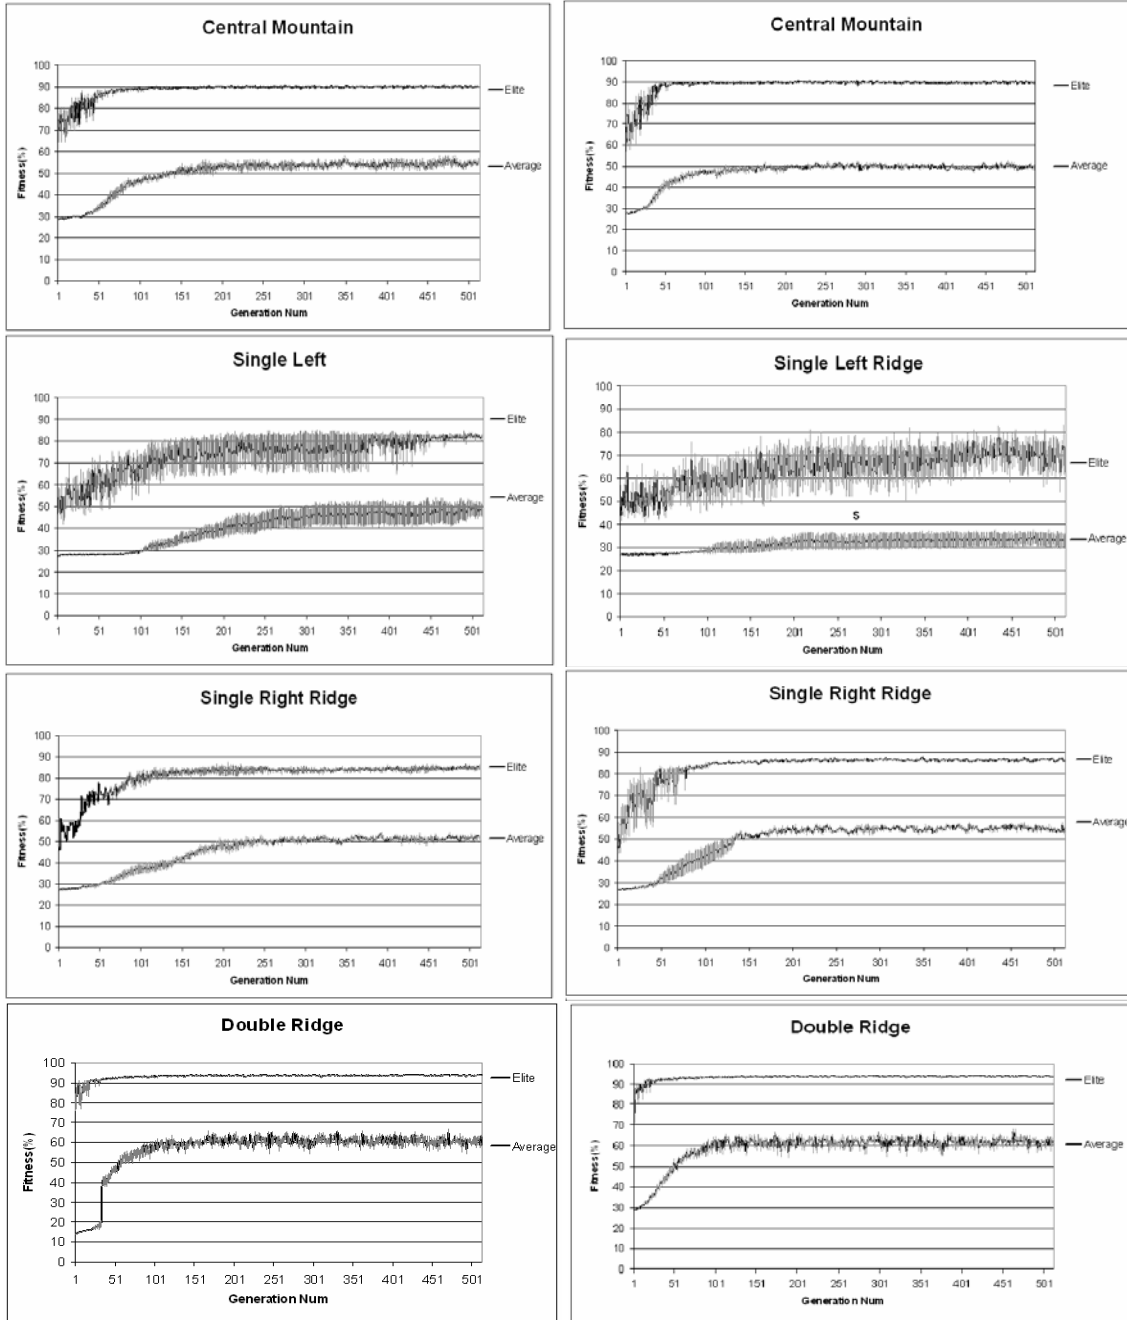
Fig 5: Comparisons of the fitness of the population in the two experiments (Fitness vs. Generation Number). The left hand column gives the results of the previous experiment [12] and the right hand column shows the results of the current experiment. Elite and Average Fitness growth averaged over 5 Runs. Error bars (grey) give the standard deviation over the 5 runs.

The paper showed that the method used to co-evolve sensor morphology on a legged robot worked well even with a significant change in stimulus availability in the environment. The strategy was adaptable and worked well as long as other stimuli (in this case IR light) were available and detectable by the agent. The results of Single Left Ridge and Single Right Ridge make the best case for this. In the Single Left Ridge, the agent was more inconsistent than the others as the IR light was not always detectable. Whereas when the IR light was available and in

an advantageous position the agent was more efficient. In the Single Right Ridge, the agent made good use of the consistently available IR light marking the best passage.

This strategy works well due to its ability to consider multiple stimuli and choose which one is the best for its current problem. Unlike in the previous experiment [12], every agent in this work used their IR sensing capabilities to complete the task. This makes the case for creating agents with access to more than essential sensing abilities since a change in a single stimulus required the usage of different stimuli that was previously deemed unnecessary. To have the ability to adapt to significant changes in the environment, this co-evolution strategy makes use of redundant data within the search space. The adaptability of the system to changes in environmental stimuli suggests that agents working in remote areas need redundant sensing abilities and would work well even with significant alteration of the environment.

In future work we will test the designs of the simulation on the actual robot In addition, we will consider a more complex controller. The simple If…then rules allowed us to easily compare controller complexity, but it limits the capabilities of the agent. We would like to test this new controller on more challenging environments.

5. REFERENCES

[1]   R. Pfeifer, "Building 'Fungus Eaters': Design Principles of Autonomous Agents," Proceedings of the Forth International Conference on Simulation of Adaptive Behavior, Cambridge, USA, 1996.

[2]   M. Atkins and P. Cohen, "Monitoring Strategies of Embedded Agents: Experiments and Analysis," Proceedings from Animals to Animats 4, Cambridge, USA, 1996

[3]   M. Fend, M. Yokoi, and R. Pfeifer, "Optimal Morphology of a Biologically-Inspired Whisker Array on an Obstacle-Avoiding Robot," Proceedings of the 7th European Conference on Artificial Life, ECAL Dortmund, Germany, 2003. .

[4]   K. Suguira, "Exploiting Interaction between Sensory Morphology and Learning" 2006

[5]   L. Linchentseiger, "Towards optimal sensor morphology for specific tasks: Evolution of an artificial compound eye for estimating time to contact" 2000

[6]   H.H. Lund "Co-evolving Controller Morphology with LEGO robots" In Hara and Pfeifer (eds.) Morpho-functional Machines, Springer-Verlag, Heidelberg.

[7]   A. Mark, D. Polani, and T. Uthmann. "A Framework for Sensor Evolution in a Population of Braitenberg Vehicle-like Agents," Artificial Life VI, pages 428--432, 1998

[8]   K. Balakrisha and V. Honavar, "On Sensor Evolution in Robotics," Proceedings of the First Annual Conference on Genetic Programming, Stanford University, USA, 1996.

[9]   M. Bugajska, and A. Schultz, "Coevolution of Form and Function in the Design of Micro Air Vehicles," Proceedings of the NASA/DoD Conference on   Evolvable Hardware, Alexandria, VA, USA, 2002.

[10]   C. Mautner and R. Belew, "Evolving Robot Morphology and Control," Proceedings of Artificial Life and Robotics, Oita, 1999.

[11]   K. Suigua, T. Shiose, H. Kawakami, and O. Katai. "Co-evolution of Sensors and Controllers," Transactions of the Institute of Electrical Engineers of Japan. Vol.124; No.10; 2004.

[12]   G.Parker, P. Nathan "Co-Evolution of Sensor Morphology and Control on a Simulated Legged Robot," Proceedings of the 7th IEEE International Symposium on Computational Intelligence in Robotics and Automation, 2007.

[13]   G. Parker, "Evolving Cyclic Control for a Hexapod Performing Area Coverage" IEEE International Symposium on Computational Intelligence in Robotics and Automation, 2001.