2015 IEEE Symposium Series on Computational Intelligence

Evolving Robust Robot Team Morphologies for Collective Construction

James Watson
Department of Computer Science
University of Cape Town
Cape Town, South Africa
Email: jwatsonza@gmail.com

Geoff Nitschke
Department of Computer Science
University of Cape Town
Cape Town, South Africa
Email: gnitschke@cs.uct.ac.za

***Abstract*—This research falls within evolutionary robotics and the larger taxonomy of cooperative multi-robot systems. A study of comparative methods to adapt the behaviors and morphologies of simulated robot teams that must solve a collective construction task is presented. Multiple versions of an indirect encoding (developmental) method for the artificial evolution of team behaviors and morphologies were tested. Results indicated the developmental method was able to evolve effective robot team morphologies in a collective construction task, where evolved teams yielded a task performance comparable to optimal team morphologies manually designed specifically for the collective construction task. Results also indicated that the developmental method was appropriate for evolving controllers that were able to generalize to a range of different team morphologies that also solved the collective construction task with a high degree of task performance.**

I. Introduction

An open problem in *cooperative multi-robot systems* [1] is determining *a priori* the most appropriate sensory-motor configurations (morphologies) for individual robots such that robots best collectively solve cooperative tasks.

Many approaches for co-evolving robot behaviors and morphologies have successfully derived *behavior-morphology* couplings [2] specifically suited to accomplishing a range of tasks [3], [4], [5], [6], [7], [8], [9] using various direct encoding approaches in a range of artificial evolution methods. Indirect encoding (developmental) approaches have also been demonstrated as effective in many behavior-morphology adaptation studies using single (simulated) robots that must accomplish relatively simple tasks (that is, those not requiring cooperation) [10], [11], [12], [13], [14], [15].

However, research that applies developmental methods to evolve behavior and morphology in robot teams that must accomplish collective behavior tasks remains relatively scarce [16] and usually focuses on controller evolution for fixed morphology teams [17], [18], [19]. Work on the co-evolution of behavior and morphology has typically focused on single robot tasks due to the added complexity of co-evolving behavior-morphology couplings for multiple robots that must cooperatively interact. This is especially the case for behaviorally and morphologically *heterogenous* teams.

This study extends previous work [20] testing controller evolution in fixed morphology teams, where a team's morphology was pre-determined using *morphological parameter tuning* experiments that tested a diverse yet functional range of robot morphologies. The focus of this study is on behavior-morphology evolution in *homogenous* teams to reduce the computational complexity required to evolve behavior-morphology couplings that effectively solve collective behavior tasks. Thus, the morphology of each robot is adapted in company with its controller over the course of an artificial evolution process. At the end of the evolutionary process each robot in the team is given the same (fittest) behavior-morphology coupling. That is, the evolved behavior and accompanying morphology that best solves the given collective behavior task. As such, this study excludes the definition of adaptive behavior and morphology used by self assembling multi-robot systems [21], [22].

This study contributes to an open objective in collective and swarm robotics, which is to devise computational methods that automate the behavior-morphology design of robot teams. The future applications such as methods is that they automate the design of robot teams to be designed (artificially evolved), tested and verified in simulation for a given collective behavior task (for example, cooperative search of remote environments such as distant planets [23], search and rescue [24], and construction and repair [25]) before being built (for example with rapid prototyping and three-dimensional printing technology [26], [5]) and deployed to solve counter-part real-world collective behavior tasks.

This study tests two research goals given behavior-morphology adaptation in homogenous teams that must accomplish a *collective construction* task.

The first goal is to demonstrate that for the given collective construction task, co-adapting behavior and morphology produces teams that out-perform teams using controller evolution (behavioral adaptation) only. In this case, the latter teams used a fixed morphology determined by previous work [20]. To address this objective, various developmental approaches for evolving team behaviors and morphologies are comparatively tested against the fixed morphology team in the collective construction task.

The second goal is to demonstrate that a developmental encoding of controllers and morphologies in homogenous teams is appropriate for evolving (collective) behaviors that are robust to variations of the morphologies with which they were co-adapted. The motivation for this objective was that

978-1-4799-7560-0/15 $31.00 © 2015 IEEE
DOI 10.1109/SSCI.2015.150

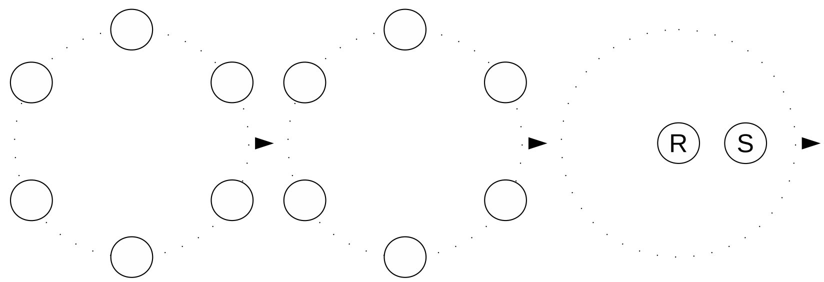

Fig. 1. ANN Topology as it relates to robot morphology: Sensory input layer (left), hidden layer (center) and motor output layer (right). Output nodes *R* and *S* determine a robot's rotation and speed, respectively. Arrows indicate the direction the agent is facing.

in future applications that use rapid prototyping or three-dimensional printing technologies to build physical versions [27] of robot teams designed in simulation, not all the required resources, materials and hardware components will necessarily be available. Hence, controllers evolved for a given team morphology would have to be sufficiently robust so as to function in varied but similar team morphologies. Also, robots may be damaged as they interact with their environment [28], [29], in which case evolved behaviors must continue to operate in modified robot bodies (morphologies).

To address this second objective, the controller (*substrate* [17]) is taken from the fittest behavior-morphology coupling evolved for the collective construction task, and this controller is copied to each robot in a team with a varied morphology. The controller coupled with this varied topology is then tested on the same collective construction task to ascertain its behavioral robustness to morphological change.

The *Collective construction* task [30] was selected since it is a variation of the well studied *collective gathering* task [31] and has pertinence to future multi-robot applications. The task was for robots to search for randomly distributed resources (blocks) in the environment and then push them such that they connected to other blocks. The goal was for all blocks to be connected during the robots' *lifetime*. One block type required cooperation (two or more robots) to move, while another block type could be moved by individual robots.

Robots were unable to explicitly detect or identify each other, and as such all cooperative interactions were *stigmergic* [32], taking place via multiple robots concurrently moving towards blocks and pushing them together into a built structure.

This task is solvable by homogenous teams given that no behavioral or morphological specialization is required for optimal solutions [33]. This task is an abstraction of real world multi-robot collective construction tasks where functional structures such as human habitats must be built from prefabricated modules [34], [35], [25]. The collective construction task used in this case study is a simplified version of a more complex construction task that requires robots to collectively build structures via connecting resources in specific ways such that a target structure is built [36].

II. Methods

A. Evolving Controllers and Morphologies

All robots in a team used the same *Artificial Neural Network* (ANN) controller, where each controller had *N* sensory input nodes (determined in various ways depending on the adaptive morphology method used), that mapped sensory inputs, via a hidden layer, to two motor outputs (figure 1) using *Hypercube-based Neuro-Evolution of Augmented Topologies* (HyperNEAT) [37]. HyperNEAT (figure 2) was selected as it is a generative encoding method that produces regular and modular ANNs with increased learning capacities [38], and has been demonstrated as exploiting regularity and modularity in multi-agent tasks to evolve solutions that could not otherwise be evolved [39]. In the collective construction task, HyperNEAT is potentially beneficial as structures to be built are modular (comprised of a set of blocks), and regular (the same sequence of blocks can be repeated). Another reason for using HyperNEAT is its successful application to evolving team (collective) behaviors for various multi-agent tasks including *RoboCup Soccer* [40] and *Pursuit-Evasion* [39].

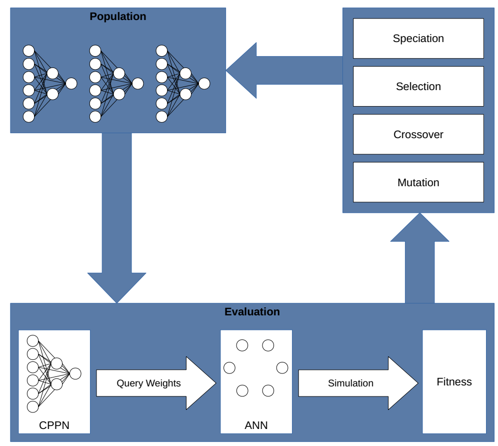

Fig. 2. A graphical overview of the HyperNEAT process.

Robot controllers were not directly evolved, but generated using an evolved CPPN (*Compositional Pattern Producing Network*) [41]. HyperNEAT was used to adapt the ANN connection weights and inter-layer connectivity, as well as robot morphology (which was simplified to only the number of sensors given results of previous work [20]). Thus, HyperNEAT evolved the connection weights and the connectivity between a fixed sensory input layer, hidden layer and motor output layer, and a sensor count, *N*. Teams were *behaviorally homogenous* in that the current fittest ANN controller was copied *P* times for *P* robots in a team. Teams were *morphologically homogenous*

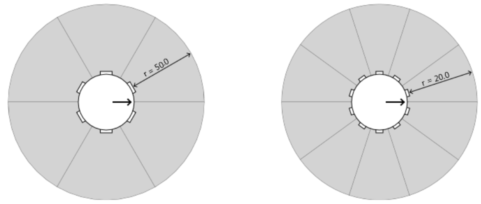

Fig. 3. Example robot sensory configurations (Left: 8 sensors, Right: 10 sensors). The sensory slices of N sensors comprise a 360 degree sensory field of view.

since each robot used the same sensory-motor configuration.

Figure 1 presents an example ANN configuration for number of sensors (N) = 6. The ANN uses a three dimensional coordinate system for processing x, y, z positions in the CPPN in order to generate weight values, connectivity and bias. The CPPN indirect encoding of HyperNEAT allows evolved controllers to exploit the geometry of the task and the environment. In the collective construction task, such geometric feature include the relative positions of other robots, blocks, and the direction robots and blocks are facing. HyperNEAT exploits the configuration of nodes in the ANN controller. Thus, controllers are evolved to exploit the position, number, and direction each sensory input faces as well as motor output nodes on a robot's body. In this study, a robot's motor outputs and the positions and direction of each sensor remains fixed, and only the number of sensors is adapted.

The input layer of the ANN controller is represented as a circle of N evenly distributed nodes. Each node is a sensor, where the sensory *Field of View* (FOV) of all sensors forms a complete 360 degree FOV (figure 3). The rotation output node is in the center to preserve the angle between sensory input nodes. The *speed* motor output node (S in figure 3) is offset in the direction the robot is facing to signify forward movement at a given speed. The intermediate hidden layer reflects the configuration of the input layer, in order to preserve the geometry of the sensory input layer, that is the direction of each sensor's FOV (figure 1). The ANN is initialized with full connectivity between adjacent layers. However, partial connectivity is evolvable via the CPPN generating a zero weight value for a given connection. During the artificial evolution process, the CPPN is developed via having nodes and connections added and removed, as well as connection weight values mutated [37]. *Neuro-Evolution* (NE) parameters used are given in table II and were determined in previous work [20].

Robot morphology (number of sensors, N) was encoded in the same genotype as the robot's ANN controller in order that behavior and morphology could be co-adapted with variations of HyperNEAT. To constrain the morphological search space and reduce computation time, robots were permitted to have [3, 16] sensors of a given type (section II-B1). The following variations of HyperNEAT were tested to evolve behavior-morphology couplings in teams.

1) Approach 1: Fixed Morphology: Team morphology, determined by previous research [20], was fixed for the adaptive process and HyperNEAT was used to evolve ANN controller connections and connection weights to adapt team behavior.

2) Approach 2: Behavior-Morphology Coupling: A robot's morphology was encoded on the same genotype as its ANN controller and team behavior and morphology was evolved with HyperNEAT. For adapting morphology, only mutation (table II) was applied to a *morphology gene* (appended to the end of the genes encoding ANN controller connectivity and weight values). The parameters for mutating the morphology gene in a robot's genotype is specified in table II.

3) Approach 3: Morphological Boosting: This approach evolved controllers for a fixed morphology, however, an extra step was included at each generation that allowed for morphological adaptation during the evolutionary process. The fittest genotype from each generation was re-evaluated on M different morphologies (table II). Re-evaluation was equivalent to one generation, that is, X team lifetimes (evaluations per genotype in table I). After re-evaluation of the fittest genotype on the M different morphologies, a fitness was assigned to each of these M morphologies and weighted random selection was used to select one of these morphologies in all genotypes of the next generation. That is, morphologies with higher fitness have a higher degree of probability of being selected for use in the next generation. This approach took advantage of HyperNEAT's capability to encode modular and regular connectivity patterns rather than a direct encoding of connections between layers of nodes. This indirect encoding allowed one evolved CPPN to generate ANN controllers that functioned together with a range of morphologies. Thus, this approach allowed HyperNEAT to ascertain if the given morphology was sub-optimal, to adapt morphologies across generations, and to evolve controllers that were robust to functioning in multiple morphologies.

4) Approach 4: Behavior-Morphology Coupling and Boosting: This approach was a combination of behavior-morphology coupling and morphological boosting. That is, a robot's behavior and morphology was encoded in the same genotype but also included the morphological re-evaluation step at each generation (that is, approach 3). Hence, after re-evaluation of the fittest genotype, a new morphology was probabilistically selected in proportion to the fitness of the re-evaluated morphologies, then the fittest genotype was adapted to include this new morphology. This was done via updating the number of sensors encoded by the genotype to that used by the new morphology.

B. Controller Sensors and Actuators

1) Block Detection Sensors: Robots had N *block detection* sensors each with a range of r set as portion of the environment's length (table I). N was the subject of the experimental comparisons (section III), and thus either fixed (section II-A1) or evolved as part of team morphology (sections II-A2, II-A3, II-A4). Sensor range r was preset to a value determined by

TABLE I. EXPERIMENT PARAMETERS

Generations	250
Sensors per robot (Fixed / Adaptive Morphology)	[3, 16] / 3
Sensor ranges	50
Evaluations per genotype	3
Experiment runs	30
Environment length, width	100
MaxDistance	1
Team size	30
Team Lifetime (Task scenario length)	120
Type A blocks (1 robot to push)	10
Type B blocks (2 robots to push)	10

previous research [20]. A robot's 360 degree sensory FOV was split into *N* sensor quadrants (figure 3). Block detection sensors were constantly active for the duration of a robot's lifetime. Sensor q returned either 0 (no blocks detected) or 1 (one or more blocks detected) in sensor quadrant q.

These sensors were an abstract representation within the task being modeled. For example, in the physical counterpart of the collective construction task, sensors could be a combination of directional *Radio Frequency Identification* (RFID) sensors, where different blocks types are identified with specific radio frequencies output by embedded RFID chips, to enable their location and identification by robots [42]. In such a task, RFID tagging would be viable as the blocks represent prefabricated components of a structure. For the purposes of keeping robot sensory configurations minimal in this simulation, robots were only able to detect blocks, and collision detection behavior was pre-specified. As such, robots were circular and given minimal friction, so that unless robots were moving in precisely opposite directions they would push past each other with minimal changes to their trajectories. Such collisions were modeled in the simulator[1], and despite robot collisions, the team was on average able to accomplish its task.

2) Movement Actuators: Two wheel motors controlled a robot's heading (*rotation*) and *speed* (*R* and *S* in figure 1). Values for these wheel motors were normalized within the range [-1.0, 1.0], where *R* = 0.0 corresponded to no change in heading, *R* = -1.0 to maximum speed clockwise rotation, and *R* = 1.0 corresponded to maximum anti-clockwise rotation. Similarly, *S* = 0.0 corresponded to no movement and *S* = 1.0 to movement in the robot's current heading at maximum speed. A robot's maximum speed was the maximum distance it could traverse in one simulation iteration ($MaxDistance$ in table I).

III. EXPERIMENTS

Experiments tested *n* robots in a bounded two dimensional continuous environment (100 x 100 units) containing a random distribution of type *A* and *B* blocks (table I). Robots were initialized with random orientations within an area at the environment's center. Blocks were randomly placed throughout the entire environment, so that although the robots' initial positions were fixed, the relative difference in positions between robots and blocks was randomized. In previous research, a *construction schema* dictated the sequence of block types that must be connected together in order for a structure to be built [43]. However, to first demonstrate that the proposed approaches worked with a simple collective construction task, blocks could be connected in any sequence.

[1]The multi-robot simulator we developed for the experiments in this study (and source code) can be found at: https://github.com/james-za/necc

TABLE II. NEURO-EVOLUTION PARAMETERS

Mutation rate	Add neuron	0.25
	Add connection	0.8
	Remove connection	0.02
	Weight	0.1
	Sensor count	0.25
Mutation type		Gaussian
Total Morphologies		14
Sensor count (All morphologies)		[3, 16]
Population size		100
Survival rate		0.3
Crossover proportion		0.4
Elitism proportion		0.1
CPPN topology		Feed-forward
CPPN inputs		Position, delta, angle

A. Collective Construction Task:

This task required a simulated robot team to gather blocks and cooperatively build a structure from gathered blocks in a bounded continuous environment. Task complexity was equated with the degree of cooperation (number of robots required) to collectively transport blocks and connect them together with other blocks in order to build a structure. The final structure resulted from connecting all blocks in the environment. Team task performance (fitness) was the number of blocks connected during a team's lifetime (equation 1).

The fitness function (equation 1) used in team evaluation was a weighted sum that included, the number of times a robot successfully found blocks (a in equation 1), the number of times *type A* blocks were pushed by *one robot* and connected with a built structure, and the number of times *type B* blocks were pushed by *two robots* and connected with a built structure (b in equation 1).

Parameter tuning experiments found that setting the weights (reward values r_a and r_b in equation 1) both to 1.0 resulted in functional controller evolution. Fitness was normalized to the range $[0.0, 1.0]$ using the number of blocks and robots required to move a given block (s_i).

$$f = \frac{r_a a + r_b b}{r_b n + r_a \sum_{i=1}^{n} s_i} \tag{1}$$

Task difficulty is regulated via requiring varying degrees of cooperation to make specific block connections. Cooperation was when at least two robots simultaneously pushed a block to touch another block (blocks were automatically connected in this manner). One robot only was required to push type *A* blocks and two robots were needed to push type *B* blocks. Hence, the more robots required to push a given block type, the more difficult the task. Task difficulty could further be increased via increasing the portion of blocks (of all blocks in the environment) that must be cooperatively moved.

TABLE III. ONE-WAY ANOVA TO TEST THE IMPACT OF ADAPTIVE MORPHOLOGY METHOD ON GENERATIONS TAKEN TO ACHIEVE MAXIMUM COLLECTIVE CONSTRUCTION TASK PERFORMANCE.

One-way analysis of means (not assuming equal variances)			
F	Num. Df	Denom. Df	p-value
1.4732	3.000	61.314	0.2306

TABLE IV. ONE-WAY ANOVA TEST TO GAUGE THE ROBUSTNESS OF THE FITTEST CONTROLLERS RE-EVALUATED IN A RANGE OF VARIED TEAM MORPHOLOGIES.

One-way analysis of means (not assuming equal variances)			
F	Num. Df	Denom. Df	p-value
10.056	3.000	61.024	1.779×10^{-05}

B. Experiment Design and Objectives

Experiments compared various adaptive morphology methods and controller evolution in a fixed team morphology, measuring the impact of each on the collective behavior of a homogenous robot team evolved for the collective construction task (section III-A). In order to address this study's research objectives (section I), experiments were designed to ascertain the most appropriate approach for evolving team behavior and morphology such that team task performance was maximized. Previous research [20] guided the choice of a fixed morphology to use as a baseline. This was compared to three adaptive morphology approaches (sections II-A2, II-A3, II-A4).

Each of four experiments tested one approach for a team size of 30 robots, where each experiment applied HyperNEAT (for fixed morphology teams with controller evolution) or a HyperNEAT variation (adaptive morphology teams with controller evolution) for 150 generations. Previous experiments found 150 generations to be sufficient to observe convergence to an optimal collective behavior for a range of team morphologies [20]. One generation comprised three *team lifetimes* (simulation task scenarios), where one team lifetime was 240 simulation iterations. This represented a task scenario that tested different robot starting orientations and block locations in the environment. For a given experiment, the number of generations taken to achieve maximum (optimal) team task performance (fitness) was recorded. An average of generation counts was calculated over 30 runs for each experiment.

Experiment and NE parameters are given in tables I and II, respectively. In table II, the CPPN inputs which affected the weight or bias of a given node were the x, y, z position of connecting nodes, the difference between their positions (*delta*), and the angle between them. These parameter values were determined experimentally and minor value changes produced similar results. Except those parameters given in table II, other NE parameters were set to values previously used for HyperNEAT [39].

IV. RESULTS AND DISCUSSION

Figure 4 presents box plots of the average number of generations (*efficiency*) with variances and outliers taken to achieve maximum fitness, by teams using each of the four approaches (section II-A). Efficiency refers to the number of generations taken for the evolutionary process to reach optimal task performance, where 150 generations was the maximum taken by any approach. In figure 4 fitness is normalized to the range: [0.0, 1.0].

Figure 4 indicates that while some approaches converge on maximum fitness slightly earlier (*Approach 2* for example), all approaches tested achieved optimal task performance within 150 generations. No statistically significant difference was found in a comparison ($F = 1.4732$, $p = 0.2306$, table III) between the three adaptive morphology approaches (approaches 2, 3, and 4 in figure 4) and the fixed morphology approach (approach 1 in figure 4) using a one-way analysis of variance (ANOVA). The results of these statistical tests are summarized in table III. This result confirms that the adaptive morphology approaches tested have a comparable average *efficiency* for the collective construction task.

Similarly, a comparison between the average team fitness (equation 1) of the fixed and adaptive morphology approaches yielded no statistical difference[2]. Also, this result is inline with previous research [20] indicating that careful tuning of a team's morphology is effective for evolving teams with a high task performance, given that the experimenter has some *a priori* knowledge of the task.

Given these results, a set of *morphological re-evaluation* experiments were performed. Hence, for each approach (figure 4), the controller of fittest team evolved after 30 runs was taken and placed in a new team morphology. That is, the fittest controller (substrate) evolved for each approach (experiment) was re-evaluated on 14 morphologies (corresponding to robots with between [3, 16] sensors, table II). This re-evaluation did not run neuro-evolution and as such did not adapt team controllers or morphologies any further. Instead each team was executed for 50 epochs (*team lifetimes*) and collective construction task performance evaluated. This was done in order to rigorously test the fittest controllers in their new morphologies. These re-evaluation experiments used the same collective construction task and simulation setup as previous experiments (section III).

As per the research goals of this study (section I), the intent of these morphological re-evaluation experiments was to test how robust the fittest evolved controllers were to variations in a team's morphology. The idea being that testing these controllers on a range of other morphologies emulates loss of sensors due to damage, or resource constraints when manufacturing physical robots.

Figure 5 presents the average team task performance for each approach calculated over the 14 morphologies tested and 30 runs. Average task performance is normalized to the range: [0.0, 1.0] in order to equate to a portion of maximum fitness achievable (equation 1).

Results indicate that approaches 2, 3 and 4 (sections II-A2, II-A3 and II-A4) perform significantly better, compared to approach 1 (section II-A1) when re-evaluated across all morphologies. The statistical significance of these results was confirmed with a one-way ANOVA test conducted for all approaches. The ANOVA test results indicating a statistically significant difference between approach 1 and approaches 2,

[2]An average fitness graph was not presented here due to space constraints, but is available as an online appendix (along with an analysis of the evolved CPPNs): http://people.cs.uct.ac.za/~gnitschke/SSCI2015/necc-ssci-extra.html

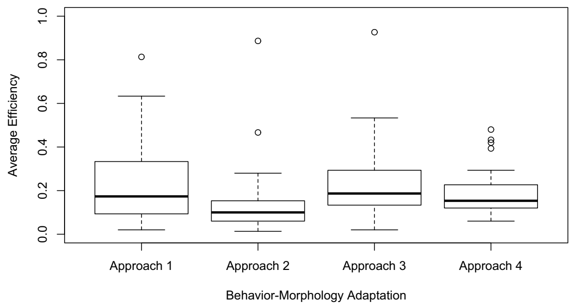

Fig. 4. Average *efficiency* (generations taken to attain optimal fitness) for each approach (average calculated over 30 runs). The efficiency scale is normalized to be a portion of the maximum number of generations (150) taken for any approach to evolve an optimal collective behavior. *Approach 1:* Fixed Morphology. *Approach 2:* Behavior-Morphology Coupling. *Approach 3:* Morphological Re-evaluation. *Approach 4:* Behavior-Morphology Coupling and Re-evaluation.

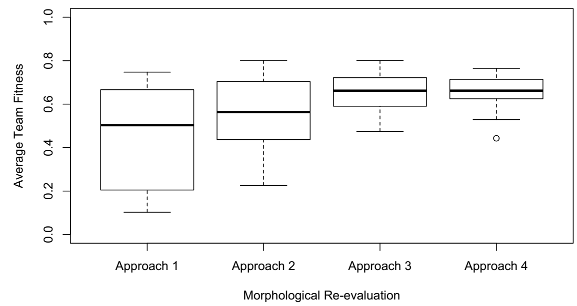

Fig. 5. Average fitness for *morphology re-evaluation* experiments (average calculated over 30 runs). The fittest team evolved for each approach (figure 4) was taken and re-evaluated in all 14 morphologies (table II). *Approach 1:* Fixed Morphology. *Approach 2:* Behavior-Morphology Coupling. *Approach 3:* Morphological Re-evaluation. *Approach 4:* Behavior-Morphology Coupling and Re-evaluation.

3 and 4 is presented in table IV ($F = 10.056$, $p = 1.779 \times 10^{-05}$).

This result indicates that the adaptive morphology approaches are more suitable for evolving controllers that are robust to morphological variation in homogenous robot teams. This is theorized to be a result of the indirect encoding property of HyperNEAT that naturally encodes *connectivity patterns* [44], coupled with the mechanisms used to adapt team morphology in the adaptive morphology approaches.

That is, HyperNEAT is able to evolve CPPNs that represent large-scale ANN controllers with their own symmetries and regularities, which exploit the sensory-motor geometry [37] of robots, which in turn significantly impacts evolved team behavior. Thus, we hypothesize that adapting controllers in company with the number of sensors allows for the encoding of controllers that do not exploit specific configurations of blocks and robots in the task environment or sensory-motor mappings in controllers. This allows HyperNEAT to evolve connectivity patterns that are broadly applicable to a range of teams (that is, the fittest evolved controllers function with varied team morphologies). This is supported by related research that similarly demonstrates the robustness of HyperNEAT evolved controllers for a range of quadruped robot morphologies [45].

Also, this robustness of controllers evolved under the adaptive morphology approaches was likely enabled by the use of homogenous teams. That is, all robots had the same *multi-agent policy geometry* [17] (relationship between robot starting positions of robots and their behaviors) for any given robot position and orientation in the environment.

To better elucidate this result, the average efficiency of each approach for all re-evaluated team morphologies was plotted. Figure 6 presents the average team task performance for each approach, when the fittest controller evolved by the given approach was re-evaluated on all 14 morphologies. Task performance was averaged over 30 runs and normalized to be a portion of maximum fitness attainable (equation 1).

Given that fixed morphology teams (approach 1) were evolved using only three sensors, approach 1 yielded a comparable average team task performance (fitness) for teams using four sensors, however, team fitness decreases proportionately as the sensor count increases.

Figure 6 also indicates that for teams where robots used three or four sensors, approach 2 yielded a lower team fitness and a comparable team fitness for five and six sensors. For seven and 16 sensors, average team fitness was higher than the fixed morphology approach but lower than adaptive morphology approaches 3 and 4. This indicates that approach 2 was not beneficial in the given collective construction task when robots had too few or too many sensors. Prior to the morphological re-evaluation experiments, the fittest team evolved by approach 2 used robots with on average (calculated over 30 runs) 11 sensors. Some reliance of approach 2 on the team morphology with which it was evolved can be observed in figure 6. That is, approach 2 achieved the highest team fitness in the morphological re-evaluation experiments for sensor counts observed to be in the range: [9, 12], close to the original evolved morphology.

Hence, the fittest controllers evolved by approaches 1 and 2 were not robust to variations in team morphology. Approach 1 delivered an average team fitness of approximately 40% for all 14 team morphologies with a high variance when robots used fewer than 10 sensors. Approach 2 delivered an average team fitness of approximately 50% for all morphologies, and average team fitness was significantly lower, compared to approaches 3 and 4, when robots were re-evaluated with robots using seven or fewer sensors.

However, approach 3 (including an extra step for adapting team morphology) and approach 4 (combining approaches 3 and 4) were able to a deliver consistently high average team fitness, with relatively little variance, for all 14 morphologies. This result is supported by the one-way ANOVA test indicating a statistically significant difference between approach 1 and approaches 2, 3, and 4 in these morphological re-evaluation experiments (table IV). This indicates that approaches 3 and 4

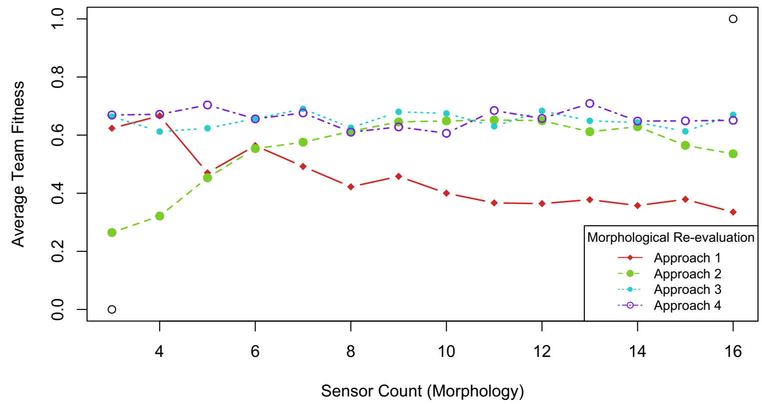

Fig. 6. Re-evaluated average team task performance (averaged over 30 runs) for each approach, when the fittest controller evolved by a given approach was re-evaluated on the complete range of 14 morphologies (number of sensors per robot in the range: [3, 16]). Note that task performance is normalized so as to equate to a portion of maximum fitness achievable (equation 1).

for adapting behavior and morphology were especially suited for evolving controllers that were robust to a range of variation in team morphology. The morphologies of the fittest teams evolved by approaches 3 and 4, prior to the morphological re-evaluation experiments, used robots with on average (calculated over 30 runs) three sensors. Observing figure 6, these fittest controllers performed with consistent team fitness when tested in robots with between three and 16 sensors, attaining an average team fitness of approximately 70%.

This result supports the hypothesis that the mechanisms used to adapt team morphologies in approaches 3 and 4 worked well in company with the adaptive process of HyperNEAT, facilitating the evolution of connectivity patterns that are sufficiently generalized, in that they do not exploit specific sensory-motor mappings derived from specific task environment instances. This enabled these fittest controllers to function with a relatively high team fitness in a range of team morphologies. However, the exact evolutionary and environmental mechanisms underlying observed controller robustness to morphological variation as evolved by the adaptive morphology approaches remains the subject of ongoing research.

V. Conclusions and Future Work

This research presented a study on the effectiveness of various adaptive morphology methods applied to evolving collective behaviors in a robot team. The team had to accomplish a collective construction task where cooperation was required for optimal solutions. This study demonstrated the efficacy of indirect encoding approaches for evolving connectivity patterns (controllers) that are able to function in homogenous teams with varied morphologies.

Results indicated that such developmental approaches were robust to morphological change, where as controllers evolved with a fixed team morphology were not. This has important implications for future evolutionary and collective robotics applications, where rapid prototyping technologies are used to manufacture and deploy robot teams evolved in simulation, but where robotic hardware may be limited or robots are damaged causing morphological change.

Future work will focus on testing these adaptive behavior-morphology and other developmental approaches on morphologically and behaviorally heterogenous teams for other collective behavior tasks with varying task complexity.

References

[1] A. Farinelli, R. Farinelli, L. Iocchi, and D. Nardi, "Multi-robot systems: A classification focused on coordination," *IEEE Transactions on Systems Man and Cybernetics B*, vol. 34, pp. 2015–2028, 2004.

[2] R. Pfeifer and J. Bongard, *How the body shapes the way we think*. Cambridge, USA: MIT Press, 2006.

[3] K. Sims, "Evolving 3d morphology and behavior by competition," in *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, USA: MIT Press, 2004, pp. 28–39.

[4] H. Lund, J. Hallam, and W. Lee, "Evolving robot morphology," in *Proceedings of the IEEE International Conference on Evolutionary Computation*. Indianapolis, USA: IEEE Press, 1997, pp. 197–202.

[5] H. Lipson and J. Pollack, "Automatic design and manufacture of robotic life forms," *Nature*, vol. 406, no. 1, pp. 974–978, 2000.

[6] G. Hornby and J. Pollack, "Creating high-level components with a generative representation for body-brain evolution," *Artificial Life*, vol. 8(3), pp. 1–10, 2002.

[7] H. Lund, "Co-evolving control and morphology with lego robots," in *Morpho-functional Machines: The New Species (Designing Embodied Intelligence)*, H. Fumio and R. Pfeifer, Eds. Berlin, Germany: Springer, 2003, pp. 59–79.

[8] A. Schultz and M. Bugajska, "Co-evolution of form and function in the design of autonomous agents: Micro air vehicles project," in *Proceedings of the Workshop on Evolution of Sensors in Nature, Hardware, and Simulation*. AAAI Press, 2000, pp. 154–166.

[9] J. Auerbach and J. Bongard, "Environmental influence on the evolution of morphological complexity in machines," *PLoS Computational Biology*, vol. 10(1), p. e1003399. doi:10.1371/journal.pcbi.1003399, 2014.

[10] P. Eggenberger, "Evolving morphologies of simulated 3d organisms based on differential gene expression," in *Proceedings of the Fourth European Conference on Artificial Life*. MIT Press: Brighton, United Kingdom, 1997, pp. 205–213.

[11] C. Mautner and R. Belew, "Evolving robot morphology and control," *Artificial Life and Robotics*, vol. 4(3), pp. 130–136, 2000.

[12] G. Hornby and J. Pollack, "Evolving l-systems to generate virtual creatures," *Computers and Graphics*, vol. 25(6), pp. 1041–1048, 2001.

[13] ——, "Body-brain co-evolution using l-systems as a generative encoding," in *Proceedings of the Genetic and Evolutionary Computation Conference*. San Francisco, USA: IEEE Press, 2001, pp. 868–875.

[14] J. Auerbach and J. Bongard, "Dynamic resolution in the co-evolution of morphology and control," in *Proccedings of the 12th International Conference on the Synthesis and Simulation of Living Systems*. Odense, Denmark: MIT Press, 2010, pp. 451–459.

[15] N. Cheney, R. MacCurdy, J. Clune, and H. Lipson, "Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding," in *Proceedings of the Genetic and Evolutionary Computation Conference*. Amsterdam, Netherlands: ACM Press, 2013, pp. 167–174.

[16] Y. Asai and T. Arita, "Coevolution of morphology and behavior of robots in a multiagent environment," in *Proceedings of the SICE 30th Intelligent System Symposium*. Tokyo, Japan: The Society of Instrument and Control Engineers, 2003, pp. 61–66.

[17] D. D'Ambrosio, J. Lehman, S. Risi, and K. Stanley, "Evolving policy geometry for scalable multi-agent learning," in *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems*. Richland, USA: ACM Press, 2010, pp. 731–738.

[18] ——, "Task switching in multi-agent learning through indirect encoding," in *Proceedings of the International Conference on Intelligent Robots and Systems*. Piscataway, USA: IEEE Press, 2013, pp. 2802–2809.

[19] D. D'Ambrosio, J. Lehman, S. Goodell, S. Risi, and K. Stanley, "Multi-robot behavior synchronization through direct neural network communication," in *Proceedings of the 5th International Conference on Intelligent Robotics and Applications*. Montreal, Canada: Springer, 2012, pp. 603–614.

[20] J. Watson and G. Nitschke, "Deriving minimal sensory configurations for evolved cooperative robot teams," in *Proceedings of the IEEE Congress on Evolutionary Computation*. Sendai, Japan: IEEE Press, 2015, pp. 3065–3071.

[21] R. O'Grady, A. Christensen, and M. Dorigo, "Swarmorph: Morphogenesis with self-assembling robots," in *Morphogenetic Engineering, Understanding Complex Systems*, R. Doursat, Ed. Berlin, Germany: Springer-Verlag, 2012, pp. 27–60.

[22] P. Moubarak and P. Ben-Tzvi, "Modular and reconfigurable mobile robotics," *Robotics and Autonomous Systems*, vol. 60, no. 12, pp. 1648–1663, 2012.

[23] R. Brooks and A. Flynn, "Fast, cheap and out of control: A robot invasion of the solar system." *Journal of the British Interplanetary Society*, vol. 1, no. 1, p. 478485, 1989.

[24] H. Kitano, S. Tadokoro, I. Noda, H. Matsubara, T. Takahashi, A. Shinjoh, and S. Shimada, "Robocup rescue: Search and rescue in large-scale disasters as a domain for autonomous agents research," in *Proceedsings of the Conference on Systems, Man and Cybernetics*. ACM Press.

[25] J. Werfel and R. Nagpal, "Three-dimensional construction with mobile robots and modular blocks," *The International Journal of Robotics Research*, vol. 27(3-4), pp. 463–479, 2008.

[26] H. Lipson, "Homemade: The future of functional rapid prototyping," *IEEE Spectrum*, vol. May, pp. 24–31, 2005.

[27] A. Eiben and J. Smith, "From evolutionary computation to the evolution of things," *Nature*, vol. 521, pp. 476–482, 2015.

[28] V. Zykov, E. Mytilinaios, B. Adams, and H. Lipson, "Self-reproducing machines," *Nature*, vol. 435, no. 7038, pp. 163–164, 2005.

[29] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 1, pp. 503–507, 2015.

[30] J. Werfel, "Building blocks for multi-robot construction," in *Distributed Autonomous Robotic Systems 6*. Tokyo, Japan: Springer, 2007, pp. 285–294.

[31] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems.* Oxford, England: Oxford University Press, 1998.

[32] R. Beckers, O. Holland, and J. Deneubourg, "From local actions to global tasks: Stigmergy and collective robotics," in *Proceedings of the International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, USA: MIT Press, 1994, pp. 181–189.

[33] G. Nitschke, M. Schut, and A. Eiben, "Emergent specialization in biologically inspired collective behavior systems," in *Intelligent Complex Adaptive Systems*. New York, USA: IGI, 2008, pp. 100–140.

[34] R. Brooks, P. Maes, M. Mataric, and G. More, "Lunar base construction robots," in *Proceedings of Intelligent Robots and Systems*. Tokyo, Japan: IEEE Press, 1990, pp. 389–392.

[35] W. Shen, P. Will, and B. Khoshnevis, "Self-assembly in space via self-reconfigurable robots," in *IEEE International Conference on Robotics and Automation*. Taipei, Taiwan: IEEE Press, 2003, pp. 2516–2522.

[36] J. Werfel, K. Petersen, and R. Nagpal, "Designing collective behavior in a termite-inspired robot construction team," *Science*, vol. 343(6172), pp. 754–758, 2014.

[37] K. Stanley, D'Ambrosio, and J. Gauci, "Hypercube-based indirect encoding for evolving large-scale neural networks," *Artificial Life*, vol. 15, no. 1, pp. 185–212, 2009.

[38] P. Tonelli and J. Mouret, "On the relationships between generative encodings, regularity, and learning abilities when evolving plastic artificial neural networks," *PLoS ONEs*, vol. 8(11), p. e79138. doi:10.1371/journal.pone.0079138, 2013.

[39] D. D'Ambrosio and K. Stanley, "Generative encoding for multiagent learning," in *Proceedings of the Genetic and Evolutionary Computation Conference*. Atlanta, USA: ACM Press, 2008, pp. 819–826.

[40] P. Verbancsics and K. Stanley, "Evolving static representations for task transfer," *Journal of Machine Learning Research*, vol. 11(1), pp. 1737–1769, 2010.

[41] K. Stanley, "Compositional pattern producing networks: A novel abstraction of development," *Genetic Programming and Evolvable Machines: Special Issue on Developmental Systems*, vol. 8, no. 2, pp. 131–162, 2007.

[42] V. Ziparo, A. Kleiner, B. Nebel, and D. Nardi, "Rfid-based exploration for large robot teams," in *Proceedings of the International Conference on Robotics and Automation*. Rome, Italy: IEEE Press, 2007, pp. 4606–4613.

[43] G. Nitschke, M. Schut, and A. Eiben, "Evolving behavioral specialization in robot teams to solve a collective construction task," *Swarm and Evolutionary Computation*, vol. 2, no. 1, pp. 25–38, 2012.

[44] J. Gauci and K. Stanley, "Autonomous evolution of topographic regularities in artificial neural networks," *Neural Computation journal*, vol. 22(7), pp. 1860–1898, 2010.

[45] S. Risi and K. Stanley, "Confronting the challenge of learning a flexible neural controller for a diversity of morphologies," in *Proceedings of the Genetic and Evolutionary Computation Conference*. Amsterdam, The Netherlands: ACM Press, 2013, pp. 255–261.