

Deriving Minimal Sensory Configurations for Evolved Cooperative Robot Teams

James Watson

Department of Computer Science
University of Cape Town
Cape Town, South Africa
jwatson@cs.uct.ac.za

Geoff Nitschke

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

Abstract—This paper presents a study on the impact of different robot sensory configurations (morphologies) in simulated robot teams that must accomplish a collective (cooperative) behavior task. The study’s objective was to investigate if effective collective behaviors could be efficiently evolved given minimal morphological complexity of individual robots in an homogenous team. A range of sensory configurations are tested in company with evolved controllers for a collective construction task. Results indicate that a minimal sensory configuration yields the highest task performance, and increasing the complexity of the sensory configuration does not yield an increased task performance.

I. INTRODUCTION

An open problem in *cooperative multi-robot systems* [1] is determining, *a priori* the most appropriate sensory-motor configuration (morphology) for individual robots, such that when an automated controller design (behavioral adaptation) method is applied, a robot team derives a collective behavior to solve a given cooperative task. This study falls within *Evolutionary Robotics* (ER) [2] research and the larger taxonomy of cooperative multi-robot systems [1], where *Neuro-Evolution* (NE) [3] is used to evolve controllers for a morphologically homogenous team. A popular approach in ER is to employ *Cooperative Co-evolutionary Algorithms* (CCAs) [4], [5], [6] to co-adapt robot behaviors and morphologies. Such approaches have been successful for finding robot morphologies and controllers specifically suited to accomplishing various tasks [7], [8], [9], [10], [11]. With notable exceptions such as Asai and Arita [12], Buason *et al.* [13], and O’Grady *et al.* [14] (including other work from the research field of self-assembling multi-robot systems [15]), the co-evolution of behavior and morphology has focused on single robot tasks due to the added complexity of co-evolving behavior-morphology couplings for multiple robots that must cooperatively or competitively interact.

However, an alternate approach to using CCAs in ER multi-robot systems is to systematically test a range of robot morphologies in company with controller evolution, in order to ascertain the best team morphology and behavior for a given task and environment. For morphologically homogenous teams such an approach does not entail intractable search spaces or exponentially increasing computational complexity associated with increasing team sizes and task complexity. Rather, the experimenter must design a set of *morphological parameter tuning* experiments that test a sufficiently diverse yet functional range of robot morphologies. In this case some *a priori* task knowledge is assumed, and the morphologies selected by

these parameter tuning experiments should be sufficient to accomplish a broad range of instances of the given task.

Pertinent CCA research [7], [8], [9], [10], [11] also assumes some *a priori* knowledge of the task for which robot controllers and morphologies are co-evolved and tasks are limited to those solvable by single robots in order to constrain computation time. In cases where the experimenter is seeking optimal performance for collective behavior tasks (only solvable by robot teams), the computation time needed to attain effective co-evolved robot *body-brain* couplings quickly becomes intractable. Hence, we propose that if the experimenter is seeking optimal team performance for a specific task, then manually tuning morphological parameters with a bias towards minimal configurations, and then evolving controllers, is the most cost effective approach.

Thus, this study’s objective is to ascertain if effective collective (multi-robot) behaviors can be efficiently evolved in simulated robot teams where a team has a minimal sensory configuration. To address this objective, we hypothesize that for many collective behavior tasks, the most cost effective approach is to first perform morphological parameter tuning experiments to ascertain an appropriate team morphology, and then to keep this morphology fixed for controller evolution experiments. This study assumes that one has *a priori* task knowledge and wants to derive a minimal morphology for a homogenous team such that the team effectively accomplishes many instances of its given task.

As initial support of this hypothesis, this study uses an ER simulation task with potential real world collective behavior applications. In this case study, the task is autonomous multi-robot collective construction, where robots are required to cooperatively move and connect building blocks. In such collective behavior tasks the behavior-morphology search space of a CCA rapidly increases as the team size and task complexity increases. Also, searching for the best controller-morphology combinations takes an inordinate amount of time if robots must have different (specialized) behaviors for the team to optimally accomplish its task.

In this study, a *collective construction* task [16] was selected since it is a variation of the well studied *collective gathering* task [17] 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 to move, while another block type could be moved by individual robots. 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 [18], [19], [20]. In such cases, modules have different sizes and weights and thus light weight modules could be moved by individual robots, while heavier modules would require cooperative transportation. In this case study, blocks could be connected to any other in any order. However, this 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 [21].

This research is a preliminary step in producing NE methods that autonomously adapt individual robot behaviors and morphologies such that problem solving collective behaviors are produced for high level user specified goals [22]. For example, where desired structures are specified by a user and robots adapt their individual behaviors in order to collectively build [21] or self-assemble the desired structure [23].

Thus this study falls within the larger scope of an ongoing research endeavor that aims to produce computational methods that automate the design (behavioral and morphological) of robotic swarms that must optimally solve complex collective behavior tasks with potential physical applications such as cooperative search, transportation, construction and repair. However, as an initial step to addressing this larger context, this paper's research goal was constrained to deducing if manual morphological parameter tuning with controller evolution is a cost effective alternative to CCA for the given collective behavior case study.

The motivation for this approach was the desire to find minimal sensory configurations for robot teams that allow collective behavior tasks to be efficiently solved. This falls in line with objectives of designing swarm robotic systems able to accomplish collective behavior tasks with minimal cost, weight (number of sensors) and power requirements, and where effective problem solving collective behaviors can be evolved in minimal time [24].

A. Collective Construction Task:

This task requires a simulated robot team to gather blocks and cooperatively build a structure from gathered blocks in a bounded continuous environment (figure 1). The complexity of this task is 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 (resultant from connecting all blocks in the environment). In this research there are two block types, *A* and *B* that require one and two robots to transport, respectively. The blocks must be connected into a structure according to a *construction schema*, that dictates the sequence for how the block types must be connected. However, for the testing purposes of this preliminary research the construction schema allows blocks to be connected together in any sequence. Task performance (team fitness) is the number of blocks connected (as a built structure) during a team's lifetime (table I).




Fig. 1. Collective Construction Task Example: Robots are the circles and blocks as the small (*Type A*) and large (*Type B*) rectangles. The green structure is that which has been built thus far via robots pushing blocks together.

II. METHODS

A. Robot Controller

All robots in a team used the same *Artificial Neural Network* (ANN) controller, where each controller had N sensory input nodes (determined by a given experiment set), that mapped sensory inputs, via a hidden layer, to two motor outputs (figure 3) using HyperNEAT [25]. HyperNEAT was selected since it is a generative encoding method that produces regular and modular ANNs with increased learning capacities [26], HyperNEAT has also been demonstrated as being capable of exploiting regularity and modularity in multi-agent tasks in order to evolve solutions that could not otherwise be evolved [27]. In the collective construction task, HyperNEAT is potentially beneficial, given that 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 HyperNEAT's selection is its successful application to evolving team behaviors for various multi-agent tasks including *RoboCup Soccer* [28] and *Pursuit-Evasion* [27].

Robot controllers were not directly evolved, but generated using an evolved CPPN (*Compositional Pattern Producing Network*) [29]. Whilst the number of sensory input and hidden nodes and sensor ranges were determined by the experimenter (manual morphological parameter tuning), HyperNEAT was used to adapt the ANN connection weights and inter-layer connectivity. Thus, HyperNEAT evolved the connection weights and the connectivity between a fixed sensory input layer, hidden layer and motor output layer. 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 since each robot used the same sensory-motor configuration.

Figure 3 illustrates an example ANN configuration for $N = 6$. The ANN uses a three dimensional coordinate system for processing x, y, z positions in the CPPN in order to generate weight and bias values and connectivity. 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




Fig. 2. Example robot sensory configurations (Left: 8 sensors, Right: 10 sensors). The sensory slices of N sensors comprise a 360 degree sensory field of view.




Fig. 3. 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.

blocks are facing. Also, HyperNEAT exploits the configuration of nodes in the ANN controller, thus the sensory input and motor output nodes must be in an appropriate configuration reflecting their position as part of a robot's morphology. Nodes for processing sensory inputs correspond to the direction each sensor faces.

Thus, the input layer of the ANN controller is 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 2). The rotation output node is in the center to preserve the angle between sensory input nodes. The speed output node 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 3). The ANN is initialized with full connectivity between adjacent layers, however, partial connectivity can be evolved via the CPPN generating a zero weight. During the artificial evolution process, the CPPN is developed via having nodes and connections added and removed, as well as connection weight values mutated [25]. NE parameters used in this study are given in table II.

1) *Block Detection Sensors*: Each robot has N *block detection* sensors each with a range of r (portion of the environment's length), where N and r are the subject of the experimental comparisons of this study, and thus determined by the experimenter (section III). A robot's 360 degree sensory FOV was split into N sensor quadrants (figure 2). 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 are an abstract representation within the task being modeled. For example, in the physical counter-part 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 [30]. In such a task, RFID tagging would be viable as the blocks represent prefabricated components of a structure. However, for the purposes of keeping robot sensory configurations minimal in this simulation, robots are only able to detect blocks, and collision detection functionality is pre-specified. That is, 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 trajectory. Such collisions were modeled in the simulator¹, and despite robot collisions, the robot team was on average able to accomplish its task.

Given that robots are unable to explicitly detect or identify each other, all cooperative interactions were *stigmergic* [31], taking place via multiple robots concurrently moving towards blocks and pushing them together into a built structure.

2) *Movement Actuators*: Two motor outputs (wheels) control a robot's heading and speed (R and S in figure 3) of movement. Values for these wheel motors are normalized within the range $[-1.0, 1.0]$, where $R = 0.0$ corresponds to no change in heading, $R = -1.0$ to maximum speed clockwise rotation, and $R = 1.0$ corresponds to maximum speed anti-clockwise rotation. Values for the wheel motors are normalized within the range $[0.0, 1.0]$, where $S = 0.0$ corresponds to no movement, and $S = 1.0$ to movement in the robot's current

¹The multi-robot simulator and source code used for experiments in this study can be found at: <https://github.com/james-za/necc>

TABLE I. EXPERIMENT PARAMETERS

Generations	250
Sensors per robot	3, 6, 10
Sensor ranges	20, 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

TABLE II. NEURO-EVOLUTION PARAMETERS

Mutation rate	Add neuron	0.25
	Add connection	0.8
	Remove connection	0.02
	Weight	0.1
Population size		100
Survival rate		0.3
Crossover proportion		0.4
Elitism proportion		0.1
CPPN topology	Feed-forward	
CPPN inputs	Position, delta, angle	

heading at maximum speed. A robot’s maximum speed is the maximum distance it can traverse in one simulation iteration (*MaxDistance* in table I).

III. EXPERIMENTS

Experiments test 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 are initialized with random orientations within an area at the center of the environment. 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 were randomized. Figure 1 illustrates an example environment, containing 30 robots and 10 *type A* and 10 *type B* blocks. In previous research, how different block types could be connected was dictated by a *construction schema* defining the sequence of block types that must be connected together in order for a structure to be built [32]. However, since this is a preliminary study to demonstrate a simple collective construction task, blocks could be connected in any sequence.

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

A. Experiment Design

Experiments measure the impact of varying robot morphology (number of sensors and sensor range) upon evolved collective behavior of a homogenous robot team given a collective

construction task in a simulation environment. The collective construction task required robots to search the environment for blocks and then push the blocks so as to connect them with other blocks in the environment. The task goal was for all blocks to be connected to form a structure. Team task performance equaled the number of blocks connected together during a team’s lifetime (equation 1). An environment was defined as a distribution of block types, robots, and a construction schema.

The research objective was to efficiently ascertain an appropriate sensory morphology, such that a homogenous team maximizes its task performance. Exploratory experiments of 12 different robot morphologies were run to guide experiment design. These preliminary tests indicated that experiments testing a set of six robot morphologies would be sufficient to satisfy the research objective. A set of three sensor counts [3, 6, 10], and two different sensor ranges [20, 50] for each robot, were tested.

Each of the six experiments tested one combination of the number and range of sensors for a fixed team size of 30 robots. Each experiment applied HyperNEAT to evolve team behavior for 250 generations. Initially, 500 generations were tested, but 250 was found to be sufficient to observe the convergence of team behavior for all morphologies tested. A generation comprised three *team lifetimes* (simulation task scenarios). One team lifetime was 120 simulation iterations, representing a task scenario that tested different robot starting orientations and block locations in the environment. For a given morphology, the task performance of a team was an average calculated over 30 simulation runs, where the maximum team task performance was selected from each simulation run.

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. Minor value changes produced similar results for all morphologies. Except those parameters given in table II, other NE parameters were set to values previously used for HyperNEAT [27].

The fitness function (equation 1) used in team (controller) 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 i (s_i).

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




Fig. 4. Average maximum fitness yielded for each robot team morphology (experiment), after 250 generations of controller evolution. Averages are calculated over 30 runs for each experiment.

TABLE III. TWO-WAY ANOVA TO TEST THE IMPACT OF TEAM MORPHOLOGY (NUMBER OF SENSORS AND SENSOR RANGE) ON COLLECTIVE CONSTRUCTION TASK PERFORMANCE.

	Df	Sum Sq	Mean Sq	F-value	p-value
Range	1	0.0005	0.00050	0.261	0.610
Sensors	2	0.1210	0.06049	31.737	1.78×10^{-12}
Range:sensors	2	0.0041	0.00204	1.070	0.345
Residuals	174	0.3316	0.00191		

IV. RESULTS & DISCUSSION

For a given experiment, the highest team task performance (fitness) achieved during an evolutionary run was recorded. An average of these maximum fitness values was then calculated over the 30 runs performed for each experiment. Figure 4 presents the average fitness achieved, by teams using each of the six team morphologies (section III-A) in terms of box plots. Figure 4 also illustrates fitness means, variances and outliers.

Figure 4 indicates that teams using only three sensors yield a significantly higher task performance, compared to the other morphologies tested. This difference was found to be statistically significant using a two-way analysis of variance (ANOVA) [33] (factors were *sensor range* and *sensor count*). ANOVA confirmed that the number of sensors has a significant impact on maximum fitness ($F = 31.737$, $p = 1.78 \times 10^{-12}$). However, the impact of sensor range was not found to be significant ($F = 0.261$, $p = 0.610$), and no interaction effect was found ($F = 1.070$, $p = 0.345$). The results of these statistical tests are summarized in table III.

This result implies that for homogenous teams that must accomplish a collective behavior task, then simpler morphologies (in this case, sensory configurations), leads to significantly higher team task performance. This is theorized to be the result of the demonstrated benefits of HyperNEAT applied to controller evolution, such as its capability to exploit modularity and geometric regularities in the task environment [27], [25], coupled with the lower dimensionality of the search space for ANNs with only three sensory nodes. Thus, in this case, the modular and regular nature of the task, requiring blocks to

be connected together in a repeated fashion and the lower dimension search space facilitated a rapid convergence upon an effective collective behavior solution by HyperNEAT.

However, attaining these results depended upon morphological design choices by the experimenter. A robot’s sensory configuration, that is, the input neurons of the ANN, were set relative to a robot’s rotation, rather than being absolute in the environment. This design decision had two main motivations. First, it allowed evolved behaviors to be more robust in that it assumed the robot has no information about its orientation with respect to global markers in the environment. Second, a robot’s heading is not aligned with a specific sensor, rather it was exactly between two sensory FOVs (figure 2). This allowed a robot to more accurately move towards blocks with a feedback-control loop of minor adjustments to the robot’s rotation. This resulted in quick direct movement towards detected blocks, thus reducing the time needed to push detected blocks, which in turn increased average team task performance.

Also, the geometric relationship between the sensory inputs and the wheel motors, controlling a robot’s rotation and speed (figure 3), benefited from relatively few sensors. That is, less sensory information being processed by the ANN, together with HyperNEAT working in a low dimensional search space (determined by a low number of sensors in this case) also facilitated the evolution of effective collective behavior.

Figure 5 presents the progression of maximum team fitness during the course of controller evolution, averaged over 30 runs. Whilst figure 4 indicated that teams using three sensors achieved the highest overall fitness, figure 5 indicates that team fitness, for this lower number of sensors, after a steep increase during the first 50 generations, remains relatively constant throughout the rest of the evolutionary run. This implies that if controller evolution can only be run for relatively few generations, then selecting an appropriate sensory configuration yields significantly beneficial task performance advantages.

Experimental results achieved on average a maximum of approximately 60 percent of optimal team task performance




Fig. 5. Progression of best team task performance over the course of controller evolution, averaged over 30 runs. Red and blue curves at the top correspond to teams using three sensors, and yield a higher average task performance (with statistical significance), compared to the other sensory configurations tested.

(figure 4). This was a result of limiting the team lifetime duration (table I) to lower computational expense, and the minimal sensory configuration which excluded sensors to detect fellow robots, so as cooperation would emerge via stigmergic interactions [31]. However, results of additional experiments² that ran with team lifetimes of 600 seconds yielded an average maximum task performance of approximately 90 percent of optimal team task performance.

These results are supported by related research [34] that similarly demonstrates a minimal number of sensors and an appropriate sensory configuration leads to desired task performance, and increasing the complexity of the sensory configuration does not necessarily result in increased task performance. Also, the results are relevant to the research fields of ER [2] and swarm robotics [24], where critical objectives are often the minimization of sensor cost, weight and power requirements of robots.

In related research that co-evolves behavior and morphology, especially that dealing with robot teams and evolving teams to solve collective behavior tasks, there is a significant computational and time expense involved to evolve robots that only solve a specific task in a specific environment [7], [8], [9], [10], [11]. In such cases, enough prior knowledge of the task and environment is often assumed, meaning that manual parameter tuning together with controller evolution would often prove to be more cost effective. Here, cost effectiveness refers to CCA implementation time and the computation time taken to evolve teams with effective collective behaviors.

Furthermore, methods that co-evolve behavior and morphology must often place significant constraints on the types of behaviors and morphologies that can be evolved in order to reduce the size of the search space and the number of behavior-morphology combinations that must be tested and evaluated. In such a case, morphological design by the experimenter, and then evolving controllers for these morphologies, would often

be just as effective as the behavior-morphology co-evolution (CCA) approaches.

Hence, the main contribution of this research is two fold. First, it demonstrated that for this collective construction task and a morphologically homogenous robot team, a minimal number of sensors facilitates the efficient evolution of effective collective behaviors, and increasing the complexity of the team’s sensory configuration does not significantly increase its task performance. Second, the study indicated that for collective behavior tasks where experimenters have some *a priori* knowledge from previous experiments [32] to guide morphological design, then a limited set of exploratory experiments that test a range of morphologies is sufficient (section III-A), and in such cases the time taken for a CCA to evolve comparably effective solutions will likely not be cost effective by comparison.

V. CONCLUSIONS AND FUTURE WORK

This research presented a study on the impact of different robot sensory configurations (morphologies) on evolving collective behaviors in a robot team that had to accomplish a collective construction task. The collective construction task required cooperation in order for the team to optimally accomplish it. This study followed the notion that we have a specific collective behavior task that must be solved by a multi-robot team, some *a priori* task information, and we want to efficiently derive a controller-morphology configuration for the team such that it effectively solves its given task. The team was morphological homogenous where the morphology was determined *a priori* by the experimenter and an accompanying ANN controller was evolved for each robot. The research objective was to investigate what degree of morphological complexity would be amenable to the efficient evolution of effective collective behaviors. This objective addressed future ER goals of designing robot swarms that have minimal cost, power and weight requirements, where problem solving collective behaviors must be efficiently evolved. Results indicated that minimal sensory configurations yielded the highest

²The results of these additional experiments can be found online at: http://people.cs.uct.ac.za/~jwatson/rundata_t600.zip

task performance, and that effective collective behaviors can be more efficiently evolved for such minimal morphologies. In this collective construction task, increasing morphological complexity did not result in increased task performance over minimal morphology-behavior couplings.

Current work is focused on increasing support for this study's hypothesis. That is, that for some specific collective behavior tasks experimenters will have some prior task knowledge and hence experience for what constitutes an effective robot team morphology. In such cases manual morphological parameter tuning in company with controller evolution will be more cost effective than co-evolving a robot team's behavior and morphology. Future work will also further test this hypothesis via including a broad range of task instances with varying degrees of task complexity that require morphologically and behaviorally heterogeneous teams.

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] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, USA: MIT Press, 2000.
- [3] D. Floreano, P. Dürri, and C. Mattiussi, "Neuroevolution: from architectures to learning," *Evolutionary Intelligence*, vol. 1, no. 1, pp. 47–62, 2008.
- [4] F. Gomez and R. Miikkilainen, "Solving non-markovian control tasks with neuroevolution," in *Proceedings of the International Joint Conference on Artificial Intelligence*. Stockholm, Sweden: Morgan Kaufmann, 1999, pp. 1356–1361.
- [5] M. Potter and K. De Jong, "Cooperative coevolution: An architecture for evolving coadapted subcomponents," *Evolutionary Computation*, vol. 8, no. 1, pp. 1–29, 2000.
- [6] R. Wiegand, W. Liles, and K. De Jong, "An empirical analysis of collaboration methods in cooperative coevolutionary algorithms," in *Proceedings of the Genetic and Evolutionary Computation Conference*. San Francisco, USA: AAAI Press, 2001, pp. 1235–1242.
- [7] 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.
- [8] H. Lipson and J. Pollack, "Automatic design and manufacture of robotic life forms," *Nature*, vol. 406, no. 1, pp. 974–978, 2000.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] G. Buason, N. Bergfeldt, and T. Ziemke, "Brains, bodies, and beyond: Competitive co-evolution of robot controllers, morphologies and environments," *Genetic Programming and Evolvable Machines*, vol. 6(1), pp. 25–51, 2005.
- [14] 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.
- [15] P. Moubarak and P. Ben-Tzvi, "Modular and reconfigurable mobile robotics," *Robotics and Autonomous Systems*, vol. 60, no. 12, pp. 1648–1663, 2012.
- [16] J. Werfel, "Building blocks for multi-robot construction," in *Distributed Autonomous Robotic Systems 6*. Tokyo, Japan: Springer, 2007, pp. 285–294.
- [17] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. Oxford, England: Oxford University Press, 1998.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] C. Jones and M. Mataric, "From local to global behavior in intelligent self-assembly," in *Proceedings of IEEE International Conference on Robotics and Automation*. Taipei, Taiwan: IEEE Press, 2003, pp. 721–727.
- [23] M. Rubenstein, A. Cornejo, and R. Nagpal, "Programmable self-assembly in a thousand-robot swarm," *Science*, vol. 345(6198), pp. 795–799, 2014.
- [24] G. Beni, "From swarm intelligence to swarm robotics," in *Proceedings of the First International Workshop on Swarm Robotics*. Santa Monica, USA: Springer, 2004, pp. 1–9.
- [25] K. Stanley, D'Amrosio, and J. Gauci, "Hypercube-based indirect encoding for evolving large-scale neural networks," *Artificial Life*, vol. 15, no. 1, pp. 185–212, 2009.
- [26] 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.
- [27] D. D'Amrosio 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.
- [28] P. Verbancsics and K. O. Stanley, "Evolving static representations for task transfer," vol. 11, pp. 1737–1769.
- [29] 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.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] B. Flannery, S. Teukolsky, and W. Vetterling, *Numerical Recipes*. Cambridge, UK: Cambridge University Press, 1986.
- [34] 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*. Chicago, USA: AAAI Press, 2000, pp. 154–166.