Evolving Swarm-Robotic Behavioral Allocations

Ameel Valjee, Bilal Aslan, Geoff Nitschke
VLJAME001@myuct.ac.za,ASLBIL001@myuct.ac.za,gnitschke@cs.uct.za
Department of Computer Science, University of Cape Town
Cape Town, South Africa

ABSTRACT

This study investigates comparative methods for two-step collective behavior evolution (evolving group behaviors from pre-evolved behaviors), to encourage the evolution of behavioral diversity in swarm-robotic applications. Specifically, we investigate behavioral diversity evolution given pre-evolved behaviors in collective behaviors that are effective across increasingly complex and difficult collective herding task environments. Results indicate that specific complements of pre-evolved (lower task-performance) collective herding behaviors was suitable for achieving high task performance across all environments and task difficulty levels. Results support the efficacy of the two-step approach for evolving behaviorally heterogeneous groups in collective behavior tasks that benefit from groups comprising various complementary behaviors.

ACM Reference Format:
Ameel Valjee, Bilal Aslan, Geoff Nitschke. 2025. Evolving Swarm-Robotic Behavioral Allocations. In *Genetic and Evolutionary Computation Conference (GECCO '25 Companion), July 14–18, 2025, Malaga, Spain.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3712255.3726748

1 INTRODUCTION

An open issue with evolving behaviors in swarm-robotic systems is ensuring that the evolutionary design method produces behavioral heterogeneity in collective behaviors sufficient for problem-solving across complex tasks [5]. Behavioral diversity in collective behaviors, emerging in response to changing task environments have demonstrated benefits across swarm-robotic applications [11, 12, 14, 15, 18–20, 26]. Various evolutionary diversity maintenance methods, including *Novelty Search* (NS) [16] and *Quality-Diversity* (QD) [23] methods have elicited behavioral diversity in evolving collective swarm-robotic behaviors. NS generates sets of behaviorally diverse controllers, where evolved controllers are optimized for behavioral diversity instead of task specific behaviors, meaning behavior repertoires evolved by NS task-agnostic [9]. QD [23] combines the benefits of NS with evolutionary directed search [3, 4]. *Multi-dimensional Archive of Phenotypic Elites* (MAP-Elites) [22] evolves repertoires of behaviorally diverse controllers given pre-defined behavioral characteristics. MAP-elites has been applied to evolve various behaviorally diverse, high performance collective behaviors for various swarm-robotic tasks [7, 10, 11, 20, 21].

Since NS and QD methods produce repertoires of diverse evolved behaviors, recent work has focused on automating behaviorally diverse swarm-robotic group design via selecting complements of behaviors from pre-evolved behavioral repertoires. Such methods attempt to boost collective behavior efficacy via optimizing behavior allocations in groups to quickly adapt swarm-robotic behaviors to changing task environments. For example, [13] proposed an automated collective behavior design approach where repertoires of *Artificial Neural Network* (ANN) modules (diverse low-level behaviors) were automatically generated using a QD evolutionary algorithm [23]. The authors introduced the *Nata* method (an instance of *AutoMoDe* [8]), to automatically generate probabilistic *Finite State Machines* (FSMs), where states were selected from the repertoire of ANN controllers and transition conditions, selected from rule sets specific to the swarm-robotic platform. The behavioral repertoire and sets of transition rules were automatically generated *a priori* and FSM controllers were assembled from behavioral modules and transition rules using the *irace* optimizer [17]. The authors demonstrated *Nata* for automating collective behavior design for aggregation and foraging tasks in physical swarm-robotic systems.

Others have demonstrated swarm-robotic controllers for multiple tasks, where such controllers dynamically switch between tasks to accomplish complex collective behaviors, via selecting between diverse behaviors in MAP-Elites evolved repertoires of behavior primitives [7]. MAP-Elites has been applied to evolve repertoires of genetic programming trees, which constitute behavioral primitives combined into a single swarm-robotic controller (complete behavior tree), suitable for solving a collective foraging task [21]. Mkhatshwa and Nitschke [20] applied the MAP-Elites based environment driven QD method [12] to co-evolve suitable degrees of behavior-morphology diversity to boost swarm-robotic task performance across increasingly complex collective gathering tasks. The authors demonstrated the benefits of behavioral diversity for maintaining a consistently high quality collective behaviors as collective behavior task difficulty increased. However, methods evolving ensembles of diverse behaviors, further evolved into collective behaviors have received little attention [11, 21]. We first demonstrate task performance benefits of behavior allocation evolution over direct behavior evolution and then a suitable method for evolving behaviorally heterogeneous groups across increasingly difficult collective herding tasks. This study evaluates two methods: *Allocate SSGA Heterogeneous* and *Allocate MAP-Elites Heterogeneous* (Sec. 2), where each uses behavior allocation evolution to evolve a composite of behaviors from repertoires of pre-evolved behaviors.

GECCO '25 Companion, July 14–18, 2025, Malaga, Spain

ACM ISBN 979-8-4007-1464-1/2025/07.
https://doi.org/10.1145/3712255.3726748

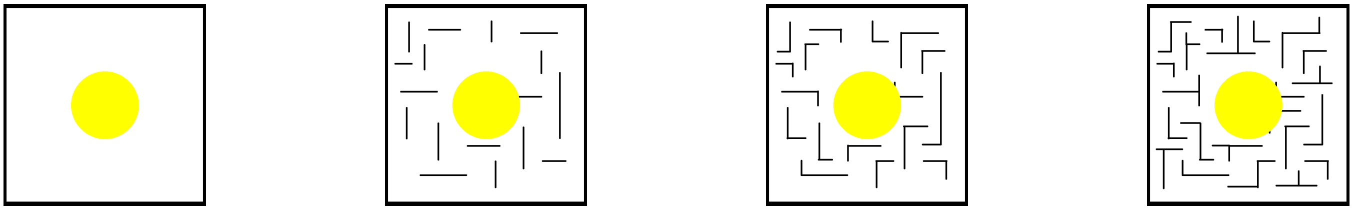

(a) No maze **(b) Easy maze** **(c) Medium maze** **(d) Difficult maze**

Figure 1: Task Environments. No maze and three maze environments of increasing complexity. Target zone at center.

2 METHODS AND EXPERIMENTS

2.1 Agents: Dogs and Sheep

Dogs use fully connected feed-forward *Artificial Neural Network* (ANN) controllers, with 9 sensory input, 10 hidden, and two motor output nodes (*tanh* activation functions). Dog controllers are evolved with SHOM and MHOM direct evolutions (Sec. 2.2.1, 2.2.2) for collective herding. Sheep use *Boids* [24] controllers to direct collective movement according to *avoidance*, *coherence* and *alignment* parameter values. Sheep had the same sensory-motor configuration as the dogs, but with different sensor ranges and fields of view. The dog and sheep controllers are fully described elsewhere [11], and so are not further described here. Tab. 1 presents the method and simulation parameters used for the dog and sheep agents.

2.2 Direct Behavior Evolution: Dog

Dog genotypes (110 evolvable connection weights, Tab. 1) are encoded as floating point weight vectors (normalized to: [-1.0, 1.0]), evolved by the SHOM or MHOM methods (Sec. 2.2.1, 2.2.2). Controller genotypes are evolved with the fitness objective of optimizing collective (herding) behavior. The ASHET and AMHET methods use archives of already evolved behaviors (previously evolved by the SHOM and MHOM methods), where the allocation of behaviors (per dog in the group) is evolved. Each of the behavior evolution methods (SHOM, MHOM), and behavior allocation evolution methods (ASHET, AMHET), are fully described in previous work [11]. For all methods (SHOM, MHOM, ASHET, AMHET), per generation, each controller in the population is systematically run and collective herding fitness assigned, for three task trials. Collective herding fitness is averaged over the three task trials. Each method is run for each environment: no maze, simple maze, medium maze and difficult maze for the two difficulty levels: easy and difficult, where sheep speed equals dog speed and is 1.25 times the dog speed, respectively (Tab. 1). A repertoire of diverse dog behaviors is evolved using MAP-Elites [22] with a d-dimensional archive, discretized into b bins. Each bin contains fittest behavior evolved thus far according to behavioral characteristic values. Behavioral characteristics for this task are: (1) *average distance between a dog and the nearest dog*, (2) *average distance between a dog and the nearest sheep*, and (3) *average distance between a dog and the target zone.*

2.2.1 SHOM: SSGA Homogeneous. SHOM uses the SSGA [25] method for dog collective herding evolution. The genotype population is randomly initialized. Each genotype corresponds to one dog ANN controller and the same controller is copied N times to compose a (homogeneous) dog group (Tab. 1). Per generation, genotypes are selected using tournament selection [6] (tournament size, k=3). Selected genotypes underwent two-point crossover and Gaussian mutation [6], with given probability (Tab. 1).

2.2.2 MHOM: MAP-Elites Homogeneous. MHOM uses MAP-Elites [22] for collective herding evolution. Each genotype in the population represents a dog controller, copied N times to derive a homogeneous dog group. The same selection and variation operators and operator parameter settings as used for SHOM (Sec. 2.2.1) are used by MHOM to evolve dog controllers (Tab. 1).

Behavior Evolution Parameters	
Generations / Task trials per generation	150 / 3
Dog genotype population	100
ANN nodes: Input / Hidden / Output	9 / 10 / 2
MAP-Elites archive: Dimensions / Bins	3 / 729
Crossover / Mutation probability	0.5 / 0.2
Simulation Parameters	
Experiment runs / Iterations per task trial	20 / 800
Initial agent positions / Group size	Random / 20
Task difficulty: Easy / Difficult	Sheep speed: x1.0 / x1.25
Environment / Size	Fig. 1 / 600px × 600px
Target zone radius / Position	100px / Center
Dog proximity sensor: Range / FOV	(0px, 100px] / [-90°, 90°]
Sheep proximity sensor: Range / FOV	(0px, 50px] / [-180°, 180°]
Sheep object avoidance: Wall/Dog/Sheep	15px / 50px / 5px
Sheep zone avoidance: Radius/Strength	50px / 0.25

Table 1: Evolution and simulation parameters.

2.3 Behavior Allocation Evolution: Group

Given that dog controllers are evolved by direct evolution (SHOM, MHOM), behavior allocation evolution methods (ASHET, AMHET), are applied to evolve complements of SHOM, MHOM evolved behaviors into various behavioral complements comprising dog groups.

2.3.1 ASHET: Allocate SSGA Heterogeneous. The ASHET method uses the population of dog controllers (behaviors) previously evolved by SHOM (Sec. 2.2.1), to evolve a suitable allocation of behaviors for a behaviorally heterogeneous dog group. The ASHET objective function is to optimize the behavioral complement of various evolved dog controllers (behaviors) to derive a high task performance (fitness, QD score) dog group. ASHET uses the final population of dog controllers evolved by SHOM, where each controller in the population represents the controller used by a dog in SHOM (Tab. 1). The final evolved SHOM populations (20 runs, Tab. 1) are projected into MAP-Elites archives based on the behavioral characteristics (Sec. 2.3) per genotype. These archives are aggregated into a reference-behavior archive comprising the 100 fittest controllers for all final generation populations. Using this behavior-reference archive, a new genotype population is initialized. Each genotype comprises 20 genes (representing a dog group, Tab. 1), where each gene indicates a specific dog controller. ASHET genotypes are evaluated as various permutations of dog controllers, where each controller is allocated from the reference-behavior archive. Per ASHET generation, parent genotypes are selected using tournament selection (*k=3*). Two-point crossover and uniform integer mutation [6] operators are applied (Tab. 1), per tournament selection operation to produce offspring genotypes, constituting the next generation of genotypes.

2.3.2 AMHET: Allocate MAP-Elites Heterogeneous. AMHET also evolved an allocation of pre-evolved controllers per dog, except allocated controllers are pre-evolved by MHOM (Sec. 2.2.2). AMHET follows the same allocation evolution method as ASHET (Sec. 2.3.1). The key difference is that the final evolved population of controllers is already contained in a MAP-Elites archive (Sec. 2.3). All other method parameters were the same as used for ASHET (Tab. 1).

3 EXPERIMENTS

Experiments[1] used four environments: *no maze*, *easy*, *medium* and *difficult* maze (Fig. 1). Dogs had the objective of herding sheep into a central target zone (Fig. 1). Sheep actively avoid entering the target zone, unless pursued by a dog. Once sheep enter the target zone, they are marked as *captured*. Dog fitness was the portion of sheep captured, averaged over three task trials per generation (Tab. 1). Experiments evaluated the behavior allocation evolution methods: AMHET and ASHET (Sec. 2) for collective herding behavior evolution in environments: *no maze*, *easy*, *medium*, and *difficult* maze. The task was varied according to the environment (no maze, simple, medium or difficult maze, Fig. 1), and task difficulty (*easy*: sheep moved at the same speed as dogs, or *difficult*: 1.25 times the speed of dogs). Each experiment applied AMHET or ASHET to evolve herding behavior for easy or difficult task settings in a given environment. Dog behavior metrics were maximum fitness and QD score. The QD score measured quality versus diversity as the sum of the highest fitness values per grid bin Q_i, as $\sum_{i=1}^{i=m} Q_i$ [23]. ASHET and AMHET used behavioral archives previously evolved by SHOM and MHOM (respectively). Each experiment ran for 150 generations (Tab. 1), with three task trials per generation. Dogs and sheep were initialized in random positions. Tab. 1 presents experiment and method parameters and respective values.

4 RESULTS AND DISCUSSION

Fig. 3 presents the average *maximum fitness* for groups evolved by ASHET and AMHET, across all environments, for the easy and difficult task settings. For comparison, Fig. 2 presents average maximum fitness results from direct behavior evolution using the SHOM and MHOM methods (Sec. 2.2.1, 2.2.2), to evolve behaviorally homogeneous groups across all environments and for easy and difficult task settings. Pairwise statistical tests (*Mann-Whitney U*, p<0.05) between average maximum fitness results of the direct evolution (SHOM, MHOM) and behavior allocation evolution (ASHET, AMHET), indicate both behavior allocation methods significantly out-perform both direct behavior evolution methods. The superior performance of ASHET and AMHET was in terms of both average maximum fitness and QD score (Fig. 2, 3).

This result demonstrates behavior performance and diversity benefits of behavior allocation evolution over direct behavior evolution. Comparing ASHET versus AMHET in terms of average *maximum fitness* and *QD score*, pairwise statistical tests (p<0.05) indicated the average fitness and QD score of ASHET evolved groups was significantly higher than AMHET evolved groups (for all environments and difficulty settings, Fig. 3). The higher average solution quality (fitness, QD score) of ASHET evolved groups is theorized to result from effective and efficient behavior evolution of the underlying SHOM method. Whereas, to the detriment of MHOM (underlying AMHET), MAP-Elites is less effective if given a relatively small evaluation budget: number of runs and population (behavioral map) size [2]. Overall, ASHET behavior allocation evolution successfully leveraged SHOM evolved behaviors to derive a suitable composition of group behaviors. The efficacy of ASHET as a two-step collective behavior evolution method (evolving behaviors with SHOM and then evolving specific behavior allocations within groups), is supported by lower average fitness and QD scores of SHOM direct evolution (underlying ASHET), across environments and task settings (Fig. 2). Thus, ASHET consistently, given relatively low fitness pre-evolved behaviors, evolved behavior allocations into complements of high fitness collective herding behaviors.

These results (Fig. 3) indicate a two-fold benefit of ASHET. First, the underlying SHOM method was suitable for evolving sufficiently diverse (addressing the collective herding task requirement for behavioral heterogeneity) and high fitness behaviors. Second, ASHET was better suited for allocating a suitable number of complementary behaviors, per environment and task setting, necessary to achieve the highest average fitness and QD scores overall. The benefits of ASHET behavior allocation evolution is supported by related work [11], though we further demonstrated behavior allocation evolution (ASHET) is most effective for producing group behavior compositions with consistently high task performance across various environment types and task difficulty levels. More broadly, these results support related work indicating the benefits of behavioral diversity across increasingly complex tasks [11, 12].

5 CONCLUSIONS AND FUTURE WORK

We investigated comparative methods for behavior allocation evolution for forming a collective herding behavior effective across increasing environment complexity: *no maze*, *simple*, *medium*, and *difficult* maze. Per environment, *easy* and *difficult* tasks were tested (sheep moved at the same speed and 1.25 faster than dogs). Results comparing behavior allocation evolution methods indicated that a complement of pre-evolved (lower task-performance) behaviors was suitable for achieving high task performance across environments. These results support the efficacy of the two-step approach

[1]Source code, data, videos: https://anonymous.4open.science/r/gecco25-sheepdogai

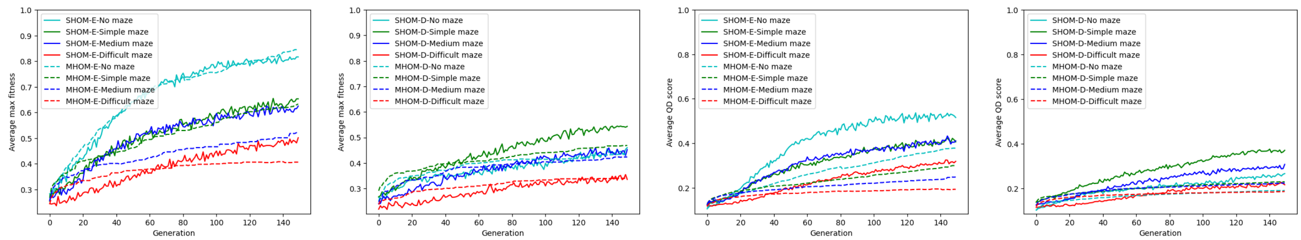

Figure 2: Direct Behavior Evolution (SHOM, MHOM): Average maximum fitness and QD scores for evolved dog groups in: no maze (cyan), simple (green), medium (blue) and difficult (red) maze, for easy (E) and difficult (D) task settings.

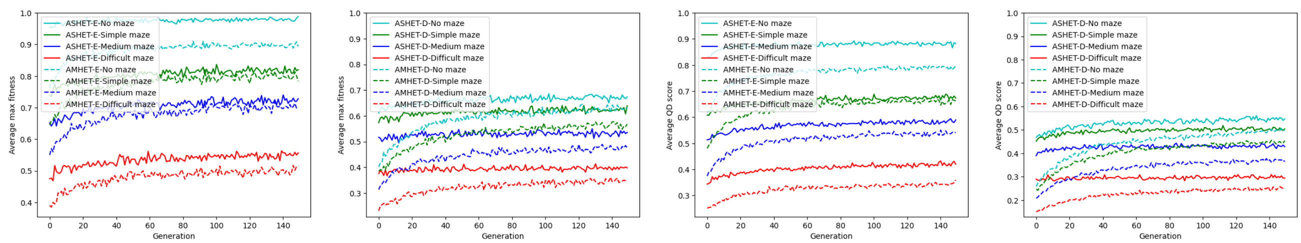

Figure 3: Behavior Allocation Evolution (ASHET, AMHET): Average maximum fitness and QD scores for evolved dog groups in: no maze (cyan), simple (green), medium (blue) and difficult (red) maze, for easy (E) and difficult (D) task settings.

for evolving behaviorally heterogeneous groups across increasingly complex collective behavior tasks. Future work aims to evaluate various evolutionary machine-learning methods [1] as behavior allocation approaches for the purpose of addressing the automated swarm-robotic system design problem [27].

REFERENCES

[1] T. Abdullahi, G. Nitschke, and N. Sweijd. 2022. Predicting Diarrhoea Outbreaks with Climate Change. *PLoS ONE* 17, 4 (2022), e0262008.

[2] C. Colas et al. 2020. Scaling MAP-Elites to Deep Neuroevolution. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Cancún, Mexico, 67–75. https://doi.org/10.1145/3377930.3390217

[3] S. Doncieux et al. 2015. Evolutionary Robotics: What, Why, and Where To. *Frontiers in Evolutionary Robotics* 2, 4 (2015), 1–18.

[4] S. Doncieux and J-B. Mouret. 2014. Beyond Black-box Optimization: A Review of Selective Pressures for Evolutionary Robotics. *Evolutionary Intelligence* 7, 2 (2014), 71–93.

[5] M. Dorigo, G. Theraulaz, and V. Trianni. 2020. Reflections on the Future of Swarm Robotics. *Science Robotics* 5, 40 (2020), eabe4385.

[6] A. Eiben and J. Smith. 2015. *Evolutionary Computing* (2nd ed.). Springer, Berlin, Germany.

[7] S. Engebråten et al. 2018. Evolving a Repertoire of Controllers for a Multi-function Swarm. In *Proceedings of the International Conference on the Applications of Evolutionary Computation*. Springer, Parma, Italy, 734–749.

[8] G. Francesca et al. 2014. AutoMoDe: A Novel Approach to the Automatic Design of Control Software for Robot Swarms. *Swarm Intelligence* 8, 2 (2014), 89–112.

[9] J. Gomes and A. Christensen. 2018. Task-Agnostic Evolution of Diverse Repertoires of Swarm Behaviours. In *Proceedings of the 11th international conference on Swarm intelligence (ANTS 2018)*. Rome, Italy, 225–238.

[10] J. Gomes and A. Christensen. 2018. Task-Agnostic Evolution of Diverse Repertoires of Swarm Behaviours. In *Swarm Intelligence*, M. Dorigo et al. (Eds.). Springer, Berlin, Germany, 225–238.

[11] S. Hallauer, G. Nitschke, and E. Hart. 2023. Evolving Behavior Allocations in Robot Swarms. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE Press, Mexico City, Mexico, 1526–1531.

[12] E. Hart et al. 2018. Evolution of a Functionally Diverse Swarm via a Novel Decentralised Quality-Diversity Algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Kyoto, Japan, 101–108.

[13] K. Hasselmann, A. Ligot, and M. Birattari. 2023. Automatic Modular Design of Robot Swarms based on Repertoires of Behaviors Generated via Novelty Search. *Swarm and Evolutionary Computation* 83, 9 (2023), 101395.

[14] L. Hong and S. Page. 2004. Groups of Diverse Problem Solvers can Outperform Groups of High-ability Problem Solvers. In *Proceedings of the National Academy of Sciences*, Vol. 101. 16385–16389.

[15] T. Kaiser and H. Hamann. 2022. From Simple Simulations to Real-World Experiments. *IEEE Transactions on Robotics* 38, 6 (2022), 3582–3601.

[16] J. Lehman and K. Stanley. 2011. Abandoning Objectives: Evolution through the Search for Novelty Alone. *Evolutionary Computation* 19, 1 (2011), 189–223.

[17] M. Lopez-Ibanez et al. 2016. The Irace Package: Iterated Racing for Automatic Algorithm Configuration. *Operations Research Perspectives* 3, 1 (2016), 43–58.

[18] G. Martin et al. 2008. Functional Group Diversity of Bee Pollinators Increases Crop Yield. *Proceedings of the Royal Society B: Biological Sciences* 275, 1648 (2008), 2283–2291.

[19] G. Martin et al. 2014. The Effects of Cultural Diversity in the Workplace. *Journal of Diversity Management* 9, 2 (2014), 89–92.

[20] S. Mkhatshwa and G. Nitschke. 2024. Body and Brain Quality-Diversity in Robot Swarms. *ACM Transactions on Evolutionary Learning and Optimization* (2024), DOI:10.1145/3664656.

[21] K. Montague et al. 2023. A Quality-Diversity Approach to Evolving a Repertoire of Diverse Behaviour-Trees in Robot Swarms. In *Proceedings of the International Conference on the Applications of Evolutionary Computation*. Springer, Brno, Czech Republic, 145–160.

[22] J.B. Mouret and J. Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).

[23] J. Pugh, L. Soros, and K. Stanley. 2016. Quality Diversity: A New Frontier for Evolutionary Computation. *Frontiers in Robotics and AI* 3 (2016), DOI:10.3389/frobt.2016.00040.

[24] C. Reynolds. 1987. Flocks, Herds and Schools: A Distributed Behavioral Model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, Anaheim, USA, 25–34.

[25] G. Syswerda. 1991. A study of Reproduction in Generational and Steady-state Genetic Algorithms. In *Foundations of Genetic Algorithms*, G. Rawlins (Ed.). Vol. 1. Elsevier, 94–101.

[26] E. Tuci. 2014. Evolutionary Swarm Robotics: Genetic Diversity, Task-Allocation and Task-Switching. In *Swarm Intelligence*, M. Dorigo et al. (Eds.). Springer, Auckland, New Zealand, 98–109.

[27] J. Watson and G. Nitschke. 2015. Evolving Robust Robot Team Morphologies for Collective Construction. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE Press, Cape Town, South Africa, 1039–1046.