Behavior Allocations in Robotic Collective Herding Behavior Evolution

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

***Abstract*—Behavioral heterogeneity yields problem solving benefits in biological collective behavior systems such as insect colonies and human societies and in artificial collective behavior systems such as distributed computer networks and swarm-robotics systems. In this study, we investigate comparative methods for two-step collective behavior evolution designed 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 a minimal complement 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.**

***Index Terms*—Swarm-robotics, Behavioral Diversity**

I. 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 [1]. Behavioral diversity in collective behaviors, emerging in response to changing task environments have demonstrated benefits across swarm-robotic applications [2]–[9]. Various evolutionary diversity maintenance methods, including *novelty search* [10] and *quality-diversity* [11] methods have elicited behavioral diversity in evolving collective swarm-robotic behaviors. Novelty search generates sets of behaviorally diverse controllers, where evolved controllers are optimized for behavioral diversity instead of task specific behaviors, meaning behavior repertoires evolved by novelty search task-agnostic [12]. Quality diversity [11] combines the benefits of novelty search with evolutionary directed search [13], [14]. For example, *Multi-dimensional Archive of Phenotypic Elites* (MAP-Elites) [15] evolves repertoires of behaviorally diverse controllers according to user defined behavioral characteristics. MAP-elites has been applied to evolve various behaviorally diverse, high performance collective behaviors for various swarm-robotic tasks [8], [9], [16]–[18].

Since novelty search and quality-diversity methods produce repertoires containing various 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, [19] proposed an automated collective behavior design approach where repertoires of *Artificial Neural Network* (ANN) modules (diverse low-level behaviors) were automatically generated using a quality-diversity evolutionary algorithm [11]. The authors introduced the *Nata* method (an instance of *AutoMoDe* [20]), 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 [21]. 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 [17]. 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 [18]. Mkhatshwa and Nitschke [9] applied the MAP-Elites based environment driven quality-diversity method [6] 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.

Other approaches to evolving diverse behaviors in swarm-robotic systems includes *multi-objective optimization* [22] and *surprise minimization* [7]. In the former case, evolutionary multi-objective optimization was applied to evolve a diverse set of behaviors via encouraging the evolution of behaviors that encapsulated trade-offs associated with optimizing for multiple task objectives. The authors demonstrated that their evolutionary multi-objective optimization method evolved diverse behaviors suitable for achieving high task performance for a flocking and cooperative swarm-robotics task [22]. In the latter case, the emergence of diverse swarm-robotic collective behaviors was driven by a fitness function that minimized surprise (maximized prediction accuracy). Each robot in the swarm used an actor-predictor pair of ANN controllers. Direct selection pressure from minimizing surprise was applied to the predictor ANN, while the actor ANN received indirect selection pressure from its predictor pairing. Behaviorally diverse collective behaviors emerged as a by-product of surprise minimization, with aggregation task performance comparable to that evolved by novelty search [7].

However, methods evolving ensembles of diverse behaviors that are further evolved into composite collective behaviors have received relatively little attention [8], [18]. This study addresses two research objectives. First, to demonstrate the collective herding task-performance benefits of behavior allocation evolution over direct behavior evolution. Second, to ascertain a suitable method for evolving behaviorally heterogeneous groups across increasingly difficult collective herding (sheep-dog) tasks. To address this second objective, we evaluate two methods: *Allocate SSGA Heterogeneous* (Section II-C3) and *Allocate MAP-Elites Heterogeneous* (Section II-C4). Each method uses behavior allocation evolution meaning collective behaviors in groups of dog robots are evolved as composites from repertoires of pre-evolved behaviors (section II).

II. Methods

This section describes the dog and sheep agent controllers (Sections II-A, II-B), evolutionary methods to evolve dog collective behaviors (Section II-C) and the collective herding task environment and objective function (Section II-D).

A. Evolved Agents: Dogs

Dogs use fully connected feed-forward *Artificial Neural Network* (ANN) controllers (*tanh* activation functions) adapted with the SHOM and MHOM method (Sections II-C1, II-C2) for the collective herding task. Dogs had an array of proximity sensors positioned about a circular periphery to detect the nearest object type (dog, sheep and wall) within a given range and *Field Of View* (FOV). Dog controllers comprised nine sensory input nodes, 10 hidden nodes and two motor outputs, corresponding to 110 evolvable connection weights (Table I). Each sensory input used distance and angle readings from three proximity sensors (one per object type), and distance and angle readings from a *target zone* (Section II-D) sensor. Distance values were normalized in the range [0.0, 1.0], where 0.0 denotes undetected and 1.0 denotes an object is as close as possible to the dog. Angle values were also normalized to [-1.0, 1.0], where -1.0 corresponds to -180 degrees and 1.0 corresponds to +180 degrees. The motor output values indicated the dog's translation and rotation from its current position. The translation value was normalized to [-1.0, 1.0], where -1.0 was the maximum translation speed backwards and +1.0 was the maximum forward translation speed. The dog's rotation value was normalized to [-1.0, 1.0], such that -1.0 denoted the maximum rotation speed to the left and +1.0 denoted the maximum rotation speed to the right.

B. Heuristic Agents: Sheep

Sheep used a pre-defined heuristic controller causing them to wander as a herd. Sheep had the same sensory input and motor output configuration as the dogs (Section II-A), but with different sensor ranges and FOV values (Table I). The sheep used a *Boids* [23] controller to direct collective movement according to *avoidance*, *coherence* and *alignment* parameter values in Boids rules (Table I). Avoidance rules used proximity thresholds for each object type, ordered by priority, so as sheep first avoid dogs and then avoid the target zone. Coherence controlled the speed with which sheep moved towards one other, and alignment controlled the degree to which sheep followed the average direction of neighboring sheep.

C. Dog Behavior Evolution Methods

The dog ANN controllers were encoded as genotypes of floating point ANN weight vectors (normalized to the range [-1.0, 1.0]), that were then evolved by the SHOM or MHOM methods (Sections II-C1, II-C2). Controller genotypes were evolved with the fitness objective of optimizing collective (herding) behavior (Section II-D1). The ASHET and AMHET methods used archives of already evolved behaviors (previously evolved by the SHOM and MHOM methods), where the allocation of behaviors (per dog in the group) was evolved. Each of the behavior evolution methods (SHOM, MHOM), and behavior allocation evolution methods (ASHET, AMHET), are fully described in previous work [8]. For all methods (SHOM, MHOM, ASHET, AMHET), per generation, each controller in the population was systematically run and collective herding fitness assigned (Section II-D1), for three task trials, where collective herding fitness was averaged over the three task trials. Each method was run for each of the simulation task environments: no maze, simple maze, medium maze and difficult maze for the two difficulty levels: easy and difficult, where sheep speed equaled and was 1.25 times the dog speed, respectively (Table I).

1) SHOM: SSGA Homogeneous: SHOM used the SSGA [24] method for dog controller adaptation (collective herding evolution). The genotype population was randomly initialized, where each genotype corresponded to one dog ANN controller and the same controller was copied N times to compose a (homogeneous) dog group (Table I). Per generation, genotypes were selected using tournament selection [25] (tournament size, *k=3*). Selected genotypes underwent two-point crossover and Gaussian mutation [25], with given probability (Table I).

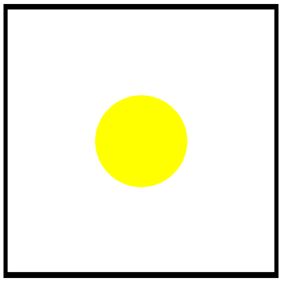

(a) No maze

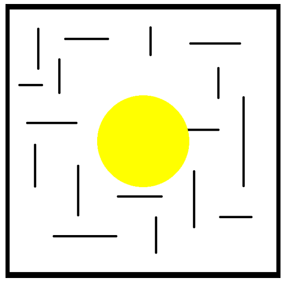

(b) Easy maze

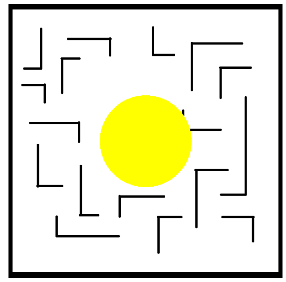

(c) Medium maze

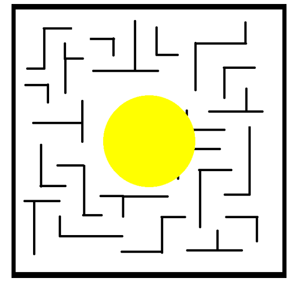

(d) Difficult maze

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

2) MHOM: MAP-Elites Homogeneous: MHOM used MAP-Elites [15] for collective herding evolution. MHOM evolved a repertoire of diverse dog behaviors, where behaviors were defined according to a descriptor containing d user-defined characteristics, defining a d-dimensional archive, discretized into b bins. Each behavior was mapped to a bin according to its descriptor. Each bin retained a single behavior with the best fitness found thus far for the given bin (behavioral descriptor value). We defined three behavioral characteristics: (1) *average distance between each dog and the nearest dog*, (2) *average distance between each dog and the nearest sheep*, (3) *average distance between each dog and the target zone*. Each behavioral characteristics was normalized to the range [0.0, 1.0], where 0.0 was the minimum average distance and 1.0 was the maximum average distance observed. The genotype population was randomly initialized, where each genotype represented a dog controller and a given controller was copied N times to derive a homogeneous dog group. The same selection and variation operators and parameter settings as used for SHOM (Section II-C1) were used by MHOM to evolve dog controllers.

Behavior Evolution Parameters	
Generations per experiment run	150
Trial evaluations per dog group	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	
Runs per experiment	20
Time steps per trial evaluation	800
Initial agent positions	Random
Dog, Sheep group size	20
Sheep speed: Easy / Difficult	x1.0 / x1.25
Task environment	See Figure 1
Arena size (width $\times$ height)	600px $\times$ 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 I: **Evolution and simulation parameters.**

3) ASHET: Allocate SSGA Heterogeneous: The ASHET method uses the population of dog controllers (behaviors) previously evolved by SHOM (Section II-C1), to evolve a suitable allocation of behaviors for a (heterogeneous) dog group. That is, the ASHET objective function is to optimize the behavioral complement of various evolved dog controllers (behaviors) in order derive a high task performance 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 (homogeneous) dog team in SHOM (Table I). Since the final population of SHOM evolved controllers are likely to be behaviorally similar given the elitist SSGA behavior [26], underlying SHOM, we pre-processed the final generation SHOM population in preparation for ASHET to evolve a behaviorally heterogeneous group.

Specifically, the final evolved SHOM populations (from 20 runs, Table I) are projected into MAP-Elites archives based on behavioral characteristics (Section II-C2) for each genotype. These archives are aggregated into a reference archive comprising the (100, Table I) fittest controllers for all final generation populations. Using this reference archive, a new genotype population is initialized. Each genotype comprises 20 genes (representing a dog group, Table I), where each gene indicates a specific dog controller. Thus ASHET genotypes are evaluated as various permutations of dog controllers, where each controller is allocated from the reference archive based on its index. Per ASHET generation, parent genotypes were selected using tournament selection (*k=3*). Two-point crossover and uniform integer mutation [25] operators were applied (with given probability, Table I), per tournament selection operation to produce offspring genotypes, constituting the next generation of genotypes.

4) AMHET: Allocate MAP-Elites Heterogeneous: AMHET also evolves an allocation of already evolved controllers per dog in a group, except controllers being allocated are pre-evolved by MHOM. AMHET follows the same method for allocation evolution as ASHET (Section II-C3). The key difference is that the final evolved population of controllers is already contained in a MAP-Elites archive (Section II-C2). All other parameters are the same as for ASHET (Table I).

D. Task Environment

A 2D bounded environment[1] was configured with four environments of increasing task difficulty: *no maze*, *easy*, *medium* and *difficult* maze (Figure 1). The 20 dogs had the objective of herding a flock of 20 sheep into a central target zone (Table I). Sheep actively avoid entering the target zone, unless pursued by a dog (Section II-B). Once sheep enter the target zone, they are marked as *captured*. This herding task was designed so as agents with complementary behaviors can achieve optimal task performance, and is also a surrogate for various collective robotics tasks such as search and rescue and toxic waste disposal [8]. The maze environments enable us to elucidate the capability of behavior allocation evolution to evolve group behavior compositions that achieve high task performance across increasingly complex task environments.

$$F = \sum_{i=1}^{n} \left(\frac{c_i}{t_i} \right) \div n \tag{1}$$

1) Dog Fitness Evaluation: Dog controllers were evaluated according to the number of sheep captured, c, out of all sheep, t, over a task trial (800 simulation iterations, Table I). A score of *zero* corresponds to no sheep captured and *one* to all sheep captured. Since sheep positions are randomly initialized, dog fitness is averaged across n evaluation trials (Equation 1).

III. Experiments

Experiments[2] evaluated the behavior allocation evolution methods: AMHET and ASHET (Section II) for collective (herding) behavior evolution across environments: *no maze*, *easy*, *medium*, and *difficult* maze (Table II). The herding task was varied according to environment complexity (no maze, simple, medium or difficult maze, Figure 1), and second according task difficulty. That is, sheep moved at the same speed as the dogs (*easy*) or 1.25 times the speed of the dogs (*difficult*). Each experiment applied a given behavior evolution method (AMHET or ASHET) to evolve herding behavior for easy or difficult task settings. Dog behavior metrics were maximum fitness (sheep captured), QD score (behavior quality versus diversity), and unique behaviors (minimal behavior set). The QD score measures quality versus diversity and is the sum of the highest fitness values per grid bin Q_i, as $\sum_{i=1}^{i=m} Q_i$ [28]. ASHET and AMHET used behavioral archives previously evolved by SHOM and MHOM (respectively). Each behavior evolution (SHOM, MHOM) or behavior allocation evolution (ASHET, AMHET) experiment evolved dog group behavior (SHOM, MHOM) or an allocation of previously evolved behaviors in a dog group (ASHET, AMHET). Each experiment ran a given method, environment, and task difficulty for 150 generations (Table I). Per generation, dog and sheep groups were run in three task trials, where each trial comprised 800 simulation iterations and initialized the dogs and sheep in different initial positions in a given environment with a given task difficulty. Table I presents experiment and method parameter values.

[1]Simulated with the *Roborobo!* multi-agent simulator [27].

[2]Source code: https://anonymous.4open.science/r/cec25-sheepdogai

IV. Results and Discussion

This discussion compares two methods for allocation evolution, ASHET and AMHET (Section II), for evolving suitable behavioral compositions such that an effective collective herding behavior emerges across increasingly complex environments: *no maze*, *simple*, *medium*, and *difficult* maze for easy and difficult task settings (Section III) results. The efficacy of the ASHET versus AMHET methods is evaluated according to averages (over 20 runs) of the metrics: *maximum fitness*, *QD score* and *unique behaviors*, in the evolved collective herding behavior of the dog group. We tested for statistically significant difference between comparative results sets using *Mann-Whitney U* [29] ($p<0.05$) in pair-wise comparisons. *Levene's test* [30] was also applied to ensure assumed equal variances.

First, in order to establish the suitability of behaviorally heterogeneous groups for achieving a high task performance in the collective herding task (Section II-D), we executed a series of benchmark experiments. The aim was to illustrate that, across environments and task difficulty levels, behaviorally homogeneous groups were unable to achieve an average maximum fitness comparable to behaviorally heterogeneous groups. For each environment and task difficulty setting, we extracted the behavior (evolved controller) comprising the largest group of unique behaviors from the fittest ASHET and AMHET evolved groups, and replicated the extracted behavior 20 times to derive a behaviorally homogeneous *benchmark* group. This benchmark group was then run in all environments, for the easy and difficult task settings, where average maximum fitness was computed over 20 runs (Figure 2). Pair-wise statistical tests between the average maximum fitness results of ASHET and AMHET evolved groups versus the benchmark group, indicate, for all environments and both easy and difficult task settings, that both ASHET and AMHET evolved groups significantly outperform ($p<0.05$) the benchmark group. This supports our hypothesis that behaviorally heterogeneous teams are most suitable for achieving high task performance across all given collective herding environments and task difficulty settings.

To support our first research objective (Section I), we next compare results of our behavior allocation evolution (AMHET, ASHET) versus direct evolution methods (SHOM, MHOM), in order to demonstrate the effectiveness evolving complements of various behaviors per group for given environments and task difficulty settings. Figure 4 presents the average *maximum fitness* for groups evolved by ASHET and AMHET, across

Method	Environment	Task Difficulty (Sheep versus Dog Speed)	
		Easy	Difficult
Allocation Evolution	No Maze	ASHET, AMHET	ASHET, AMHET
	Simple Maze	ASHET, AMHET	ASHET, AMHET
	Medium Maze	ASHET, AMHET	ASHET, AMHET
	Difficult Maze	ASHET, AMHET	ASHET, AMHET

TABLE II: **Experiment Setup.** Allocation evolution variants: ASHET, AMHET, are used for evolving collective herding in no maze, simple, medium and difficult maze task environments for varying task difficulty (Sheep speed).

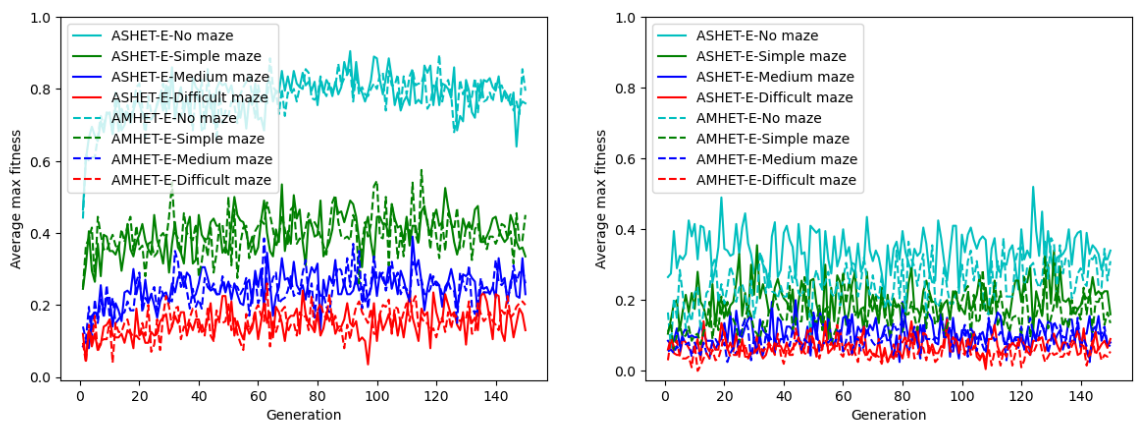

Fig. 2: **Benchmark Behavior Evolution.** To illustrate the benefit of heterogeneous groups, the behavior (evolved controller) comprising the largest group of unique behaviors was extracted from ASHET and AMHET evolved groups across environments and difficulty settings, replicated 20 times and run as a homogeneous group in the same environments for easy and difficulty task settings. Average maximum fitness (over 20 runs) results for: *no maze* (cyan), *simple* (green), *medium* (blue), and *difficult* (red) mazes for *easy* (Sheep and dog speed the same, left) and *difficult* (Sheep speed 1.25 times dog speed, right) task settings.

all environments, for the easy and difficult task settings. For comparison, Figure 3 presents average maximum fitness results from direct behavior evolution using the SHOM and MHOM methods (Sections II-C1, II-C2), to evolve behaviorally homogeneous groups across all environments and for easy and difficult task settings. Pairwise statistical tests between average maximum fitness results of the direct evolution (SHOM, MHOM) and behavior allocation evolution (ASHET, AMHET), indicate for all environments, easy and difficult task settings, that 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 (Figures 3, 4).

This supports our first objective (Section I), demonstrating task performance benefits of behavior allocation evolution over direct behavior evolution. Since we established the efficacy of behavior allocation evolution and the suitability of behaviorally heterogeneity generated by ASHET and AMHET, we next focus on the suitability of ASHET and AMHET across all environments and task difficulty settings. To address our second objective (Section I), we compare the efficacy of ASHET versus AMHET in terms of average *maximum fitness*, *QD score*, and *unique behaviors* (Figures 4, 5). Across all environments, for the easy and difficult task difficulty settings, pairwise statistical tests were applied between ASHET versus AMHET average maximum fitness, QD score and number of unique behaviors. For all environments and task settings, the average fitness and QD score of ASHET evolved groups was significantly higher than AMHET evolved groups. However, the average unique behaviors within AMHET evolved groups was significantly higher than in ASHET evolved groups, for all environments and task difficulty settings (Figure 5).

The higher average behavior quality of ASHET (versus AMHET) is supported by significantly higher QD scores across all environments and task difficulty settings, highlighting that ASHET was better suited for a more expansive behavior space search, enabling the evolution of more diverse and higher quality behaviors (Figure 4). The overall higher average solution quality (fitness, QD score) of ASHET evolved groups is theorized to result from effective and efficient behavior evolution of the SHOM (underlying ASHET) method, with the given population size and number

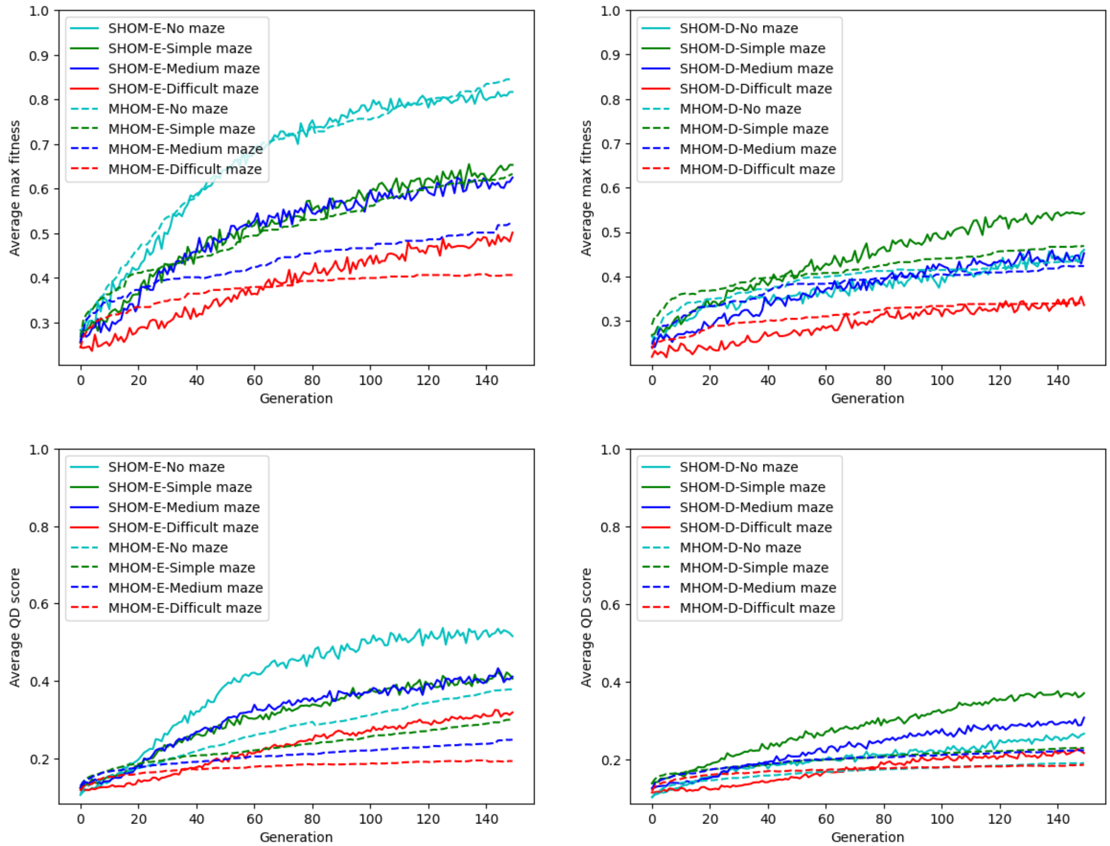

Fig. 3: **Direct Behavior Evolution (SHOM, MHOM):** Average maximum fitness and QD scores for evolved dog groups in: no maze (cyan), simple (green), medium (blue) and difficult (red) maze, for easy (E) and difficult (D) task settings.

of generations. Whereas, to the detriment of MHOM (underlying AMHET), MAP-Elites has been demonstrated as most effective given relatively large evaluation budgets: number of runs and population (behavioral map) sizes [31]. Thus, given this limited evaluation budget, ASHET behavior allocation evolution effectively leveraged SHOM evolved behaviors to evolve suitable compositions of group behaviors. The efficacy of ASHET's underlying direct evolution method (SHOM), is also supported by the lower number of unique behaviors in ASHET versus AMHET evolved groups across environments and task difficulties. That is, ASHET evolved a more diverse, higher quality and minimal complement of behaviors for all environments and task settings (Figure 5).

The higher average fitness, QD score and lower number of unique behaviors (Figures 4, 5), indicates a two-fold benefit of ASHET. First, that SHOM direct behavior evolution (underlying ASHET) was suitable for evolving sufficiently diverse (addressing the collective herding task requirement for behavioral heterogeneity), high fitness behaviors. Second, ASHET was better suited for allocating minimal behaviors necessary to achieve the highest average fitness overall. The efficacy of ASHET as a two-step collective behavior evolution method (evolving dog behaviors with SHOM and then evolving specific behavior allocations within dog groups), is also supported by lower average maximum fitness and QD scores of SHOM direct evolution, for all environments and task settings (Figure 3). Thus, per environment and task setting, ASHET effectively, given pre-evolved (low fitness) behaviors, evolved behavior allocations into complements of high fitness performance collective herding behaviors.

These results highlighting the benefits of ASHET behavior allocation evolution are also supported by related work [8]. However, this study further demonstrates behavior allocation evolution (ASHET) is most effective for producing minimal behavior compositions in groups yielding consistently high task performance across various environment types and task difficulty levels. More broadly, these results support related work on evolving collective behaviors, with results indicating the benefits of behavioral diversity across increasingly complex collective behavior task environments [6], [8].

V. Conclusions and Future Work

This study investigated comparative methods for behavior allocation evolution for deriving a composition of behaviors suitable for forming a collective herding behavior that is effective across increasing environment complexity and task difficulty. Four environment types of increasing complexity

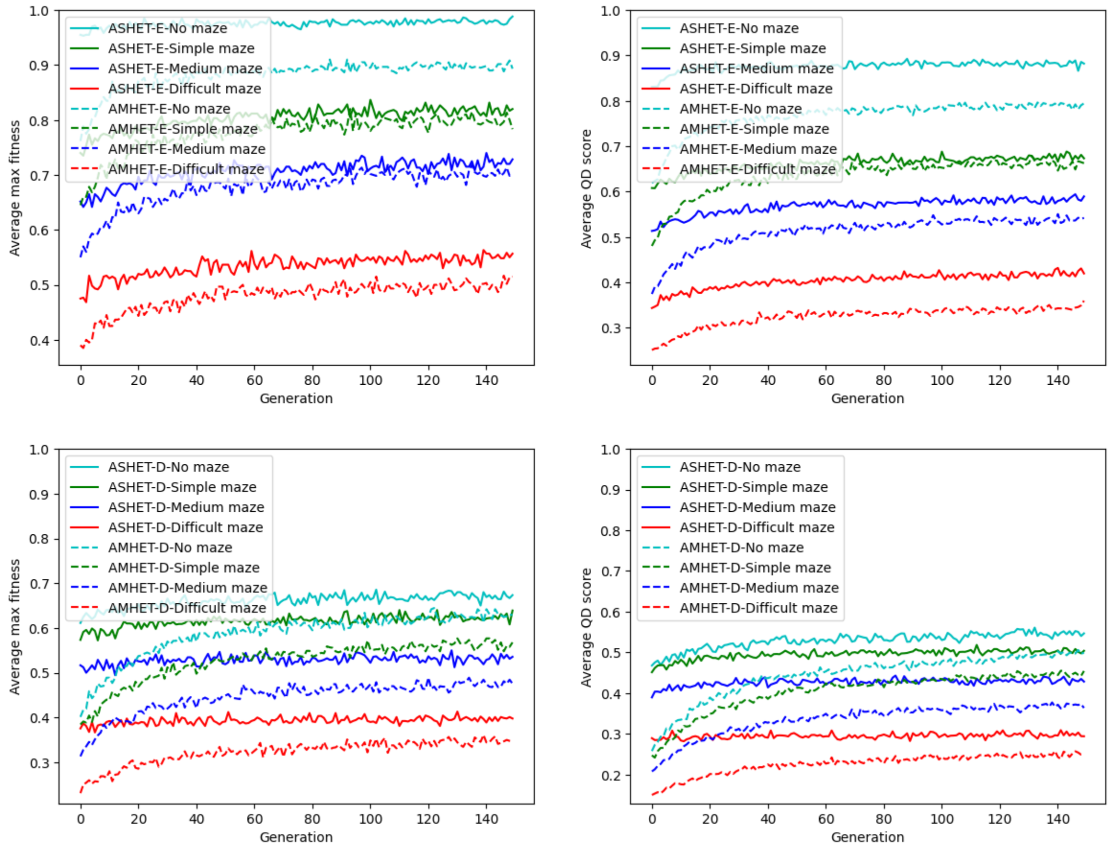

Fig. 4: **Behavior Allocation Evolution (ASHET, AMHET):** Average maximum fitness and QD scores for evolved dog groups in: no maze (cyan), simple (green), medium (blue) and difficult (red) maze, for easy (E) and difficult (D) task settings.

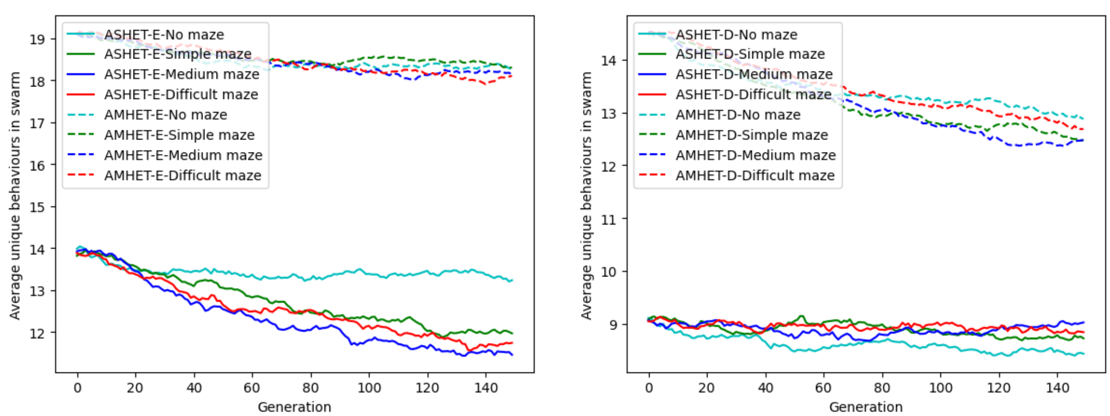

Fig. 5: **Behavior Allocation Evolution (ASHET, AMHET):** Average Unique behaviors comprising evolved dog groups in: no maze (cyan), simple (green), medium (blue) and difficult (red) maze, for easy (E) and difficult (D) task settings.

were tested: no maze, simple maze, medium maze, and difficult maze. For each environment type, two levels of task complexity were tested: easy (dogs and sheep moved at the same speed) and difficult (sheep moved 1.25 faster than dogs). We first established that behavioral heterogeneity is beneficial across all environment types and task difficulty settings for this collective herding task. Results comparing two behavior allocation evolution methods indicated that a minimal complement of pre-evolved (lower task-performance) behaviors was suitable for achieving high task performance across all environments and task difficulty levels. These results support the efficacy of the two-step approach for evolving behaviorally heterogeneous groups across increasingly complex collective behavior task environments. Future work will various evolutionary machine-learning methods [32] as behavior allocation approaches for addressing the automated swarm-robotic behavior design problem [33], [34].

References

[1] M. Dorigo, G. Theraulaz, and V. Trianni, "Reflections on the Future of Swarm Robotics," *Science Robotics*, vol. 5, no. 40, p. eabe4385, 2020.

[2] G. Martin *et al.*, "Functional Group Diversity of Bee Pollinators Increases Crop Yield," *Proceedings of the Royal Society B: Biological Sciences*, vol. 275, no. 1648, pp. 2283–2291, 2008.

[3] L. Hong and S. Page, "Groups of Diverse Problem Solvers can Outperform Groups of High-ability Problem Solvers," in *Proceedings of the National Academy of Sciences*, vol. 101, pp. 16385–16389, 2004.

[4] G. Martin *et al.*, "The Effects of Cultural Diversity in the Workplace," *Journal of Diversity Management*, vol. 9, no. 2, pp. 89–92, 2014.

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

[6] E. Hart, A. Steyven, and B. Paechter, "Evolution of a Functionally Diverse Swarm via a Novel Decentralised Quality-Diversity Algorithm," in *Proceedings of the Genetic and Evolutionary Computation Conference*, (Kyoto, Japan), pp. 101–108, ACM, 2018.

[7] T. Kaiser and H. Hamann, "From Simple Simulations to Real-World Experiments," *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3582–3601, 2022.

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

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

[10] J. Lehman and K. Stanley, "Abandoning Objectives: Evolution through the Search for Novelty Alone," *Evolutionary Computation*, vol. 19, no. 1, pp. 189–223, 2011.

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

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

[13] S. Doncieux and J.-B. Mouret, "Beyond Black-box Optimization: A Review of Selective Pressures for Evolutionary Robotics," *Evolutionary Intelligence*, vol. 7, no. 2, pp. 71–93, 2014.

[14] S. Doncieux *et al.*, "Evolutionary Robotics: What, Why, and Where To," *Frontiers in Evolutionary Robotics*, vol. 2, no. 4, pp. 1–18, 2015.

[15] J.-B. Mouret and J. Clune, "Illuminating Search Spaces by Mapping Elites," *arXiv preprint arXiv:1504.04909*, 2015.

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

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

[18] K. Montague *et al.*, "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*, (Brno, Czech Republic), pp. 145–160, Springer, 2023.

[19] K. Hasselmann, A. Ligot, and M. Birattari, "Automatic Modular Design of Robot Swarms based on Repertoires of Behaviors Generated via Novelty Search," *Swarm and Evolutionary Computation*, vol. 83, no. 9, p. 101395, 2023.

[20] G. Francesca *et al.*, "AutoMoDe: A Novel Approach to the Automatic Design of Control Software for Robot Swarms," *Swarm Intelligence*, vol. 8, no. 2, pp. 89–112, 2014.

[21] M. Lopez-Ibanez *et al.*, "The Irace Package: Iterated Racing for Automatic Algorithm Configuration," *Operations Research Perspectives*, vol. 3, no. 1, pp. 43–58, 2016.

[22] V. Trianni and M. Lopez-Ibanez, "Advantages of Task-specific Multi-objective Optimisation in Evolutionary Robotics," *PLOS ONE*, vol. 10, no. 1, p. e0136406, 2015.

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

[24] G. Syswerda, "A study of Reproduction in Generational and Steady-state Genetic Algorithms," in *Foundations of Genetic Algorithms* (G. Rawlins, ed.), vol. 1, pp. 94–101, Elsevier, 1991.

[25] A. Eiben and J. Smith, *Evolutionary Computing*. Berlin, Germany: Springer, 2nd ed., 2015.

[26] D. Whitley and J. Kauth, "GENITOR: A Different Genetic Algorithm," in *Proceedings of the 4th Rocky Mountain Conference on Artificial Intelligence*, Springer, Boulder, USA, 1988.

[27] N. Bredeche *et al.*, "Roborobo! A Fast Robot Simulator for Swarm and Collective Robotics," *arXiv:1304.2888*, 2013.

[28] J. Pugh *et al.*, "Confronting the Challenge of Quality Diversity," in *Proceedings of the Genetic and Evolutionary Computation Conference*, (Madrid, Spain), pp. 967–974, ACM, 2015.

[29] B. Flannery, S. Teukolsky, and W. Vetterling, *Numerical Recipes*. Cambridge, England: Cambridge University Press, 1986.

[30] H. Levene, "Robust Tests for Equality of Variances," in *Contributions to Probability and Statistics*, pp. 279–292, Stanford, USA: Stanford University Press, 1961.

[31] C. Colas *et al.*, "Scaling MAP-Elites to Deep Neuroevolution," in *Proceedings of the Genetic and Evolutionary Computation Conference*, (Cancún, Mexico), pp. 67–75, ACM, 2020.

[32] T. Abdullahi, G. Nitschke, and N. Sweijd, "Predicting Diarrhoea Outbreaks with Climate Change," *PLoS ONE*, vol. 17, no. 4, p. e0262008, 2022.

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

[34] G. Nitschke and D. Howard, "Autofac: The perpetual robot machine," *IEEE Transactions on Artificial Intelligence*, vol. 3, no. 1, pp. 2–10, 2022.