Body and Brain Quality-Diversity in Robot Swarms

SINDISO MKHATSHWA* and GEOFF NITSCHKE*, Department of Computer Science, University of Cape Town, South Africa

In biological societies, complex interactions between the behavior and morphology of evolving organisms and their environment have given rise to a wide range of complex and diverse social structures. Similarly, in artificial counterparts such as swarm-robotics systems, collective behaviors emerge via the interconnected dynamics of robot morphology (sensory-motor configuration), behavior (controller), and environment (task). Various studies have demonstrated morphological and behavioral diversity enables biological groups to exhibit adaptive, robust, and resilient collective behavior across changing environments. However, in artificial (swarm robotic) systems there is little research on the impact of changing environments on morphological and behavioral (body-brain) diversity in emergent collective behavior, and the benefits of such diversity. This study uses evolutionary collective robotics as an experimental platform to investigate the impact of increasing task environment complexity (collective behavior task difficulty) on the evolution and benefits of morphological and behavioral diversity in robotic swarms. Results indicate that body-brain evolution using coupled behavior and morphology diversity maintenance yields higher behavioral and morphological diversity, which is beneficial for collective behavior task performance across task environments. Results also indicate that such behavioral and morphological diversity maintenance coupled with body-brain evolution produces neuro-morpho complexity that does not increase concomitantly with task complexity.

CCS Concepts: • **Computing methodologies → Evolutionary robotics**; *Cooperation and coordination.*

Additional Key Words and Phrases: Swarm-Robotics, Controller-Morphology Adaptation, Quality-Diversity, Behavioral and Morphological Diversity Maintenance

1 INTRODUCTION

In nature, various complex interactions between evolving organism behavior, morphology, and environment, have resulted in the emergence of complex and diverse forms of social (collective) behavior [Duarte et al. 2011; Hart et al. 2002; O'Shea-Wheller et al. 2020]. Similarly, in artificial social systems such as swarm robotics, various forms of collective behavior arise from coupled dynamics between a robot's morphology (sensory-motor configuration), behavior (controller output) and environment (task) [Bredeche et al. 2018]. One perspective is that an embodied agent's (robot's) morphological and behavioral complexity must match its environment complexity [Pfeifer and Bongard 2006]. However, with varying examples in biology systems [McShea 1996], it remains an open question as to whether more complex task environments require agents with more complex behaviors and morphologies [Cheney et al. 2013; Nygaard et al. 2021a; Xu and Wang 2021].

Studies on artificial morphology-behavior (body-brain) evolution using simulated [Cheney et al. 2018; Kriegman et al. 2018] and physical [Nygaard et al. 2021a; Xu and Wang 2021] evolutionary robotics platforms have received

---

*Both authors contributed equally to this research.

Authors' address: Sindiso Mkhatshwa, MKHSIN035@myuct.ac.za; Geoff Nitschke, gnitschke@cs.uct.za, Department of Computer Science, and University of Cape Town, Computer Science Building, 18 University Ave N, Rondebosch, Cape Town, Western Cape, South Africa, 7700.

ACM 2688-3007/2024/5-ART
https://doi.org/10.1145/3664656

significant research attention [Doncieux et al. 2015]. However, work investigating the impact of body-brain adaptation in collective and swarm robotic systems is less prevalent [Buason et al. 2005; Furman et al. 2019; Hewland and Nitschke 2015]. Specifically, research investigating how emergent morphological diversity impacts evolving swarm robotic behavior has received little research attention [Hunt 2021]. This is due to the analytical difficulty of meaningfully distilling core relationships between a robot's evolving genotype (body-brain encoding) and phenotype (body-brain couplings) from the phenotypic interactions between robots and resulting emergent collective behaviors [Moore et al. 1997].

In both artificial and biological collective behavior systems, a single robot or organism genotype produces multiple morphology-behavior couplings (phenotypes) in response to varying environmental conditions [Kelly et al. 2012; Schlichting and Pigliucci 1998; West-Eberhard 1989; Wolf et al. 1999]. For example, local cues about resource distributions adjust agent foraging parameters in artificial swarms [Just and Moses 2017]. Similarly, in biological ant colonies, environmental heterogeneity generates variable foraging behaviors [Beverly et al. 2009; Gordon et al. 2011]. Such *phenotypic plasticity* resulting from evolving genotype and phenotypic interactions in collective behavior systems has been argued as a missing adaptive component preventing autonomous robots from working in unconstrained real-world environments [Hauser 2019].

Previous work [Kriegman et al. 2018; Pfeifer and Bongard 2006] also demonstrated that robot morphology significantly influences the range and types of behaviors it exhibits, where robot morphology plays a critical role in determining behavioral complexity. Adapting morphology in company with behavior can reveal novel motor-sensor relationships, which enhances the evolution of robust behaviors, particularly in challenging and unfamiliar environments [Birattari and et al. 2019; Bongard 2011; Kriegman et al. 2018; Nygaard et al. 2021b]. Given the demonstrated benefits of evolutionary robotics as an experimental tool for evaluating the evolutionary and environmental conditions under which specific behaviors emerge [Ferrante et al. 2015; Montanier et al. 2016; Nitschke et al. 2011; Steyven et al. 2017; Trueba et al. 2013; Trujillo et al. 2011; van Diggelen et al. 2022], this study uses evolutionary swarm robotics [Dias et al. 2021; Doncieux et al. 2015] as an experimental platform to investigate the impact of varying environment complexity (task difficulty) on the evolution of morphology-controller (body-brain) couplings in swarm robotic systems.

Within evolutionary swarm robotics, research incorporating evolvable coupled controllers and morphologies remains scarce and restricted in scope [Dias et al. 2021], with a noticeable absence of studies exploring the impact of environment (task) complexity on robot phenotypic plasticity and resultant emergent (collective) behavior. Conventionally, researchers define a specific environment and task (task environment), and then evolve robot behavior specifically tailored to operate within the given task environment. For example, many swarm robotics studies consider the task environment as an experimental parameter, where behavioral diversity is evolved using given morphologies that have been pre-designed to address specific collective behavior tasks [Brutschy et al. 2012; Ferrante et al. 2015; Nitschke et al. 2011; Steyven et al. 2017; van Diggelen et al. 2022]. However, notable exceptions include self-assembly swarm robotics systems comprising many individual functionally simple robots that physically attach to each other [Brambilla et al. 2013; Mathews et al. 2019; Mondada and et al. 2013]. This includes proof-of-concept demonstrations using hundreds of *Kilobots* [Carrillo-Zapata et al. 2019; Rubenstein et al. 2013; Slavkov et al. 2018], a *mergeable* nervous system [Mathews et al. 2017] and a *group mind* [Otte 2014], where swarm robotic behaviors emerge from self-organizing neural controllers interconnecting across hundreds of robots. Other examples include *multi-robot organisms* [Levi and Kernbach 2010] that adapt morphology via self-organizing into various problem-solving forms, such as evolving functional specialization in various interacting body-parts [Auerbach and Bongard 2009], and evolving desired group behaviors via morphological adaptation that switches sensors on and off, such that robots adapt to complementary sensory configurations [Furman et al. 2019; Hewland and Nitschke 2015; Watson and Nitschke 2015].

Furthermore, while previous evolutionary robotics work [Auerbach and Bongard 2014; Miras and Eiben 2019; Miras and Ferrante 2020; Spanellis et al. 2021] has studied the impact of the environment on body-brain co-evolution, there have been few studies that investigate environmental impact on body-brain evolution in swarm robotics [Furman et al. 2019; Nagar et al. 2019a,b]. In most evolutionary swarm robotics studies robot morphologies are fixed and only controllers evolve [Doncieux et al. 2015], though there are some studies demonstrating that diverse environments potentially produce diverse behaviors [Ferrante 2013; Ferrante et al. 2015]. Pertinent examples include the work of Ferrante et al. [Ferrante 2013] demonstrating that flat environments produced individual behaviors that did not elicit complex self-organized strategies, while sloped environments induced complex division of labour [Ferrante et al. 2015]. Otherwise, in evolutionary swarm robotics work with evolvable coupled controllers and morphologies, there is a lack of studies evaluating the impact of task environment complexity on emergent collective behavior [Dias et al. 2021].

Finally, in evolutionary robotics, *Quality Diversity* (QD) methods [Chatzilygeroudis et al. 2021; Cully and Demiris 2018; Pugh et al. 2016] have been applied to address the exploration-exploitation trade-off in the search for suitable body-brain couplings [Nordmoen et al. 2021; Nygaard et al. 2021a], and have been successfully applied in various single robot tasks including adaptive ambulation [Kaushik et al. 2020; Nordmoen et al. 2021; Nygaard et al. 2021a] and damage recovery [Allard et al. 2023; Cully et al. 2015; Mailer et al. 2013]. Such QD methods reformulate the optimization process to consider both behavior-morphology diversity and quality (behavior task performance), and thus aim to find new behavior-morphology couplings eliciting high-quality solutions. Given that QD methods have shown diversity of body-brain solutions, computed *a priori*, is critical for fast adaptation [Chatzilygeroudis et al. 2021; Cully et al. 2015; Kaushik et al. 2020; Mailer et al. 2013], we hypothesize that leveraging of behavioral and morphological diversity maintenance during evolutionary body-brain design within robotic swarms will elicit similarly high quality (collective) behaviors, robust across changing task environments.

1.1 Research Objectives and Contributions

Given previous work that demonstrates the benefits of QD methods for facilitating behavior diversity maintenance [Hallauer et al. 2023] and behavior-morphology diversity maintenance [Mkhatshwa and Nitschke 2023] in robot swarms, we formulate two research objectives to evaluate the impact of behavior and morphology diversity during swarm (collective behavior) evolution.

1. Evaluate the impact of behavior and behavior-morphology diversity maintenance on increasingly complex tasks, given robot controller and controller-morphology evolution.
2. Evaluate the impact of behavior and behavioral-morphological diversity maintenance on robot controller-morphology evolution (specifically, the degree of evolved neuro-morpho complexity) necessary to facilitate behaviors effective on increasingly complex tasks.

This study applies five comparative methods (Sections 2.1-2.5) to individual robot controller and controller-morphology adaptation, to evaluate the effectiveness of evolved swarm behaviors across increasingly complex (collective gathering) task environments (Section 3.1). Task complexity is defined by degrees of cooperation required (between robots) to achieve optimal task performance. The simplest environment requires no cooperation

Table 1. An overview of each of the evolutionary methods evaluated in this study.

Evolutionary method	Behavior map	Morphology map	Behavior adaptation	Morphology adaptation
mEDEA	✗	✗	✓	✗
mEDEA-M	✗	✗	✓	✓
EDQD	✓	✗	✓	✗
EDQD-M	✓	✗	✓	✓
Double-Map EDQD-M	✓	✓	✓	✓

between robots to push objects into a target-area, whereas the most complex environment requires high degrees of cooperation (multiple robots must cooperatively push objects into the target-area) to achieve optimal solutions (Section 3). Experiments aim to address our objectives via gauging the impact of behavior and morphology diversity maintenance for specific task environments to gain insights into how environment complexity impacts the evolution of behavior and morphology in the context of swarm (collective) behaviors.

This study's main contribution are results indicating the effectiveness of behavior-morphology diversity maintenance, given robot controller-morphology adaptation across increasing task environment complexity. Whereas related evolutionary robotics work has demonstrated benefits of behavioral [Babak et al. 2021; Miras and Eiben 2019; Miras and Ferrante 2020; Miras et al. 2018a,c,b] and morphological [Cheney et al. 2018; Nordmoen et al. 2021; Samuelsen and Glette 2014; Zardini et al. 2021], diversity maintenance when coupled with controller-morphology evolution [Nygaard et al. 2021a], such previous work focused on single robot tasks (such as gait adaptation) and the impact of variable environment complexity was usually not considered. We extend such previous research via demonstrating the impact and benefits of behavioral and morphological diversity maintenance in company with individual controller-morphology evolution for generating swarm robotic (collective gathering) behaviors effective across increasingly complex task environments.

We also address the limitations of related work most pertinent to this study [Mkhatshwa and Nitschke 2023], which also focused on swarm robotic controller-morphology evolution and behavioral-morphological diversity maintenance, with two key extensions and contributions. First, we evaluate more (five) robotic swarm evolution methods to demonstrate the efficacy and benefits of behavior and behavior-morphology diversity maintenance for evolving suitable robot controllers and controller-morphology couplings in swarms that must solve more (five) cooperative tasks. Second, we demonstrate the relationship between evolving behaviour and morphology diversity and robot neuro-morpho complexity in swarms evolved across increasingly complex task environments.

2 METHODS

In addition to flexibility, robustness, and scalability, swarm robotic systems offer various emergent behavior benefits [Hamann 2018], such as automated task discovery and problem-solving [Nitschke and Howard 2022]. The emergence of desired problem-solving behavior is highly desirable across a wide range of high-precision and hazardous tasks in dynamic application domains such as agriculture, mining, and disaster relief operations [Arnold et al. 2019], where the exact tasks and optimal solutions cannot be specified *a priori*. That is, as task complexity increases, defining the exact body-brain coupling of individual robots, where all possible combinations of collective behaviors resulting from all possible local interactions among robots and the environment becomes impractical, making it necessary to automate the design process [Furman et al. 2019].

We address this robotic swarm collective behavior design problem via applying evolutionary methods for automated design of robot body-brain couplings to produce desired (task-accomplishing) emergent collective (swarm) behaviors. Specifically, this study evaluates five controller-morphology adaptation methods for evolving individual robot behavior-morphology couplings and thus adapting emergent swarm behaviors. First, *minimal Environment-driven Evolutionary Algorithm* (mEDEA) [Bredeche and Montanier 2010] (Section 2.1), second, an extension to mEDEA for adapting robot behavior and morphology, termed: mEDEA-M (Section 2.2), third EDQD (*Environment Driven Quality Diversity*) method [Hart et al. 2018], fourth, EDQD-M (robot morphology adaptation), and finally, Double-Map EDQD-M (coupled robot behavior-morphology adaptation). Table 1 presents an overview of these comparative controller-morphology adaptation methods and associated use of behavioral and morphological diversity maintenance mechanisms by each method.

Also, as a point of comparison, we note marked similarity between these selected methods and related evolutionary methods such as the island model [Konfrst 2004], which was designed for the purpose of running multiple concurrent genotype (solution) populations and maintaining genetic and thus solution diversity over extended periods. The island model bears similarity to the *minimal Environment-driven Evolutionary Algorithm* (mEDEA) [Bredeche and Montanier 2010] used in this study. Specifically, both methods use N solution populations, operating in parallel, where encoded solution information is periodically exchanged between populations as means to adapt each population while maintaining diversity of solutions overall. The success of island models in these respects [Konfrst 2004; Skolicki and De Jong 2004], but their untested capabilities in the context of collective evolutionary robotics tasks, motivated the use of similar methods in previous evolutionary work [Bredeche et al. 2018; Bredeche and Montanier 2010] as well as in this study.

Another point of similarity is with swarm intelligence methods [Bonabeau et al. 1999] that attempt to balance exploration versus exploitation in the search for optimal solutions, given multiple, concurrent interactions (information exchanges) between N agents (the swarm). In swarm intelligence methods such as Particle Swarm Optimisation (PSO), the key proponents driving the exploration versus exploitation trade-off in the search process, are the personal best position (p-best) position of each particle and the global best position (g-best) of the swarm [Gad 2022]. Such PSO parameters have parallels with the interactions between the adapting (individual) robot behaviors versus the evolving (swarm) behavior driven by the evolutionary methods in this study (Sections 2.1−2.5). Swarm intelligence methods such as PSO have been successfully applied across various challenging tasks and the meta-heuristic nature of such methods makes them suitable for collective (swarm) robotic controller adaptation [Pugh and Martinoli 2016]. However, this study's goal was to better evaluate the benefits of evolutionary quality-diversity methods (EDQD, Section 2.3), and extensions accounting for body-brain adaptation (EDQD-M, Double-Map EDQD-M, Section 2.4, 2.5), where the efficacy of such quality-diversity methods for evolving effective body-brain couplings, and the impact of this on evolving swarm-robotic behavior, had not previously been evaluated.

All of this study's methods (Sections 2.1−2.5) are experimentally evaluated in the context of swarm robotic controller and morphology evolution across varying task environments (Section 3) using a custom evolutionary swarm robotics simulator[1]. We selected mEDEA as a benchmark method for swarm behavior evolution since it has been widely studied in swarm robotics [Galassi et al. 2016; Hart et al. 2018; Shan and Mostaghim 2021; Silva et al. 2015], and mEDEA-M is a logical extension that evaluates the impact of morphological adaption during swarm behavior evolution. Similarly, we selected to extend the EDQD method since EDQD has been demonstrated as successfully evolving behaviorally diverse robot swarms without requiring explicit mechanisms for genotypic

---

[1] The swarm-robotic simulator is available at: https://github.com/Body-Brain-QD-in-Robotic-Swarms/TELO-2023

(reproductive) isolation or division of labor [Hart et al. 2018], and has not been tested on increasingly complex collective behavior task environments.

2.1 mEDEA: minimal Environment-driven Evolutionary Algorithm (EA)

Using mEDEA for controller adaptation, a robot moves about its environment using its *Artificial Neural Network* (ANN) controller (Section 2.6) for x simulation iterations (*robot lifetime*, Table 3). When a robot moves, it broadcasts its genotype (ANN controller encoding) to all other robots in the environment within broadcast range (Table 3), while concurrently receiving genotypes from all other robots. Once the robot's lifetime (using its currently active genotype, encoding its current ANN controller) has ended, the robot randomly selects parents from a list of received genotypes and applies a variation operator. This takes the form of a *Gaussian* random mutation operator that is tuned through a σ parameter (Table 3). Following the application of variation operators, the selected genotype replaces the robot's currently active genotype (current ANN controller encoding). This study uses a later mEDEA variant [Perez et al. 2014] that incorporates an explicit objective function to direct robot (swarm) behavior to accomplish collective behavior tasks (Section 3.1). However the core mechanisms of mEDEA governing genotype exchange between robots were not changed, except the fitness value was broadcast with each robot's genotype and selection occurring with respect to swarm task performance (Section 3.2). The mEDEA method is otherwise the same as fully described in previous work [Bredeche and Montanier 2010].

2.2 mEDEA-M: minimal Environment-driven EA with Morphology Adaptation

The mEDEA-M method enables robot morphological (sensor) and behavioral (controller) adaptation. Specifically, in each generation, a random sensor type is chosen to undergo mutation. The mutation operator changes the sensor range at random until it exceeds a specified minimum or maximum *sensor-morpho* threshold (Table 2). When a given sensor's range drops below this threshold, the sensor becomes inactive, which is indicated by a zero input to the appropriate ANN sensory input node. Similarly, if the sensor range exceeds the *sensor-morpho* threshold due to sensor range mutation, an inactive sensor will reactivate, restoring the prior non-zero connection weight value for the specified ANN sensory input node. The ground-facing target-area detection sensor (Section 2.6) is not subject to morphological adaptation since all robots must be able to find the target-area. Also, the robot swarm is morphologically homogeneous, meaning that the same sensor adaptations (sensor range adaptation and sensors switched on and off) are applied concurrently to all robots in the swarm. Otherwise, mEDEA-M uses the same controller adaptation as mEDEA (Section 2.1).

2.3 EDQD: Embodied Distributed Quality Diversity

The EDQD method hybridizes the MAP-Elites [Mouret and Clune 2015] and mEDEA (Section 2.1) methods. Differing from mEDEA, as robots explore their environment they periodically broadcast their behavioral map (list of genotypes), instead of the genotype encoding the robot's current ANN controller (Section 2.1). The behavioral map is received and stored by all robots within broadcast range (Table 3). Such robot behavioral maps are termed *LocalMaps*, and contain a map of the genotypes (*genome* in figure 2, right) corresponding to specific robot behaviors (*phenome*, Figure 2, right) previously evaluated for each robot. At the end of each robot's *lifetime*, a genotype (genome) is randomly selected from the *SelectMap* (formed by merging the received maps with the robot's own *LocalMap*, Figure 2, right), and a mutation operator is applied to produce a new genotype which replaces the currently active genotype (robot ANN controller and thus behavior).

Applying EDQD to swarm evolution, robots store a 2D behavior map (*LocalMap*, Figure 1, left) defined by two behavioral dimensions of the collective gathering task (Section 3). Specifically, resource type collected, and

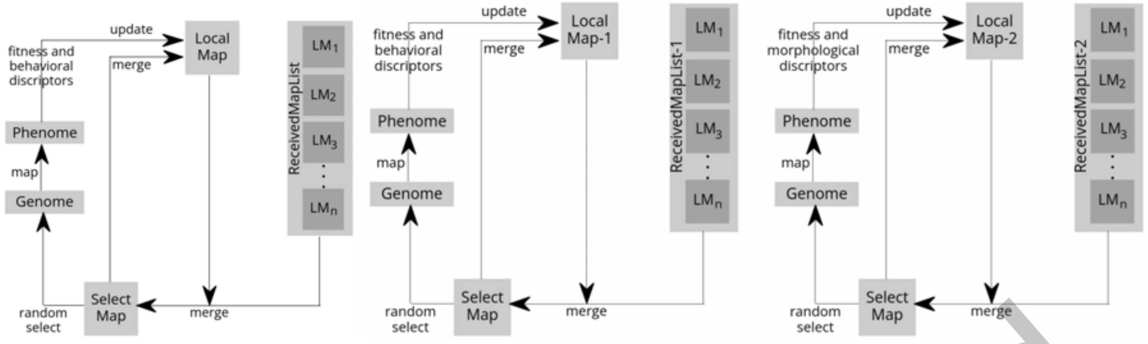

Fig. 1. Left: *EDQD*: Each generation, a random genotype (controller encoding) is selected from *SelectMap*, merging *ReceivedMapList* with the robot's *LocalMap*. Center-Right: *Double-Map EDQD-M*: Robots maintain two *LocalMaps*: *LocalMap-1*, *LocalMap-2* containing behavior and morphology feature descriptors, respectively.

maximum Euclidean distance traversed (explored) in the environment, by each robot (during its lifetime). The EDQD method and its extensions: *EDQD-M* (Section 2.4) and *Double-Map EDQD-M* (Section 2.5), thus leverage these behavioral dimensions to promote the evolution of behavioral diversity in terms of resource types collected and environment exploration.

2.4 EDQD-M: Embodied Distributed Quality Diversity with Morphology Adaptation

EDQD-M extends EDQD (Section 2.3) to enable morphological (sensor), in company with behavioral (controller), adaptation for each robot. Specifically, at the end of each generation, a random sensor type is selected to undergo mutation. The mutation operator changes the range of a randomly selected sensor until it reaches a given *sensor-morpho* threshold (Table 3). If the range falls below a minimum threshold then the given sensor becomes inactive, where sensor inactivity is realised by an input of zero to the corresponding ANN sensory input node. Similarly, if the mutation operator causes the sensor range to exceed the maximum threshold, then an inactive sensor will reactivate, reinstating the previous non-zero connection weight value for the given ANN sensory input node. As for mEDEA-M (Section 2.2), the bottom-facing target-area detection sensor is excluded from morphological adaptation, and the swarm is morphologically homogeneous meaning the same sensor adaptations are also concurrently applied to all robots. Otherwise, the EDQD-M controller adaptation process using the *LocalMap* is identical to EDQD (Section 2.3, Figure 1, left).

2.5 Double-Map EDQD-M: EDQD with Double-Map Morphology Adaptation

Double-Map EDQD-M extends EDQD via enabling co-adaptation of a robot's morphology and behavior. Double-Map EDQD-M employs two *LocalMaps*, the first of which is associated with a robot's controller (behavior) feature descriptors and the second with the robot's sensor (morphology) feature descriptors. This second map has two (morphological) dimensions: the *ratio of active sensor types* and the *average range of active sensors*. As in EDQD (Section 2.3) and EDQD-M (Section 2.4), a random genotype is selected from each *SelectMap* to undergo mutation at each generation of the evolutionary process. The selected genotypes from each *SelectMap* replace the robot's current active behavior (ANN controller) and morphology (sensory configuration), respectively (Figure 1, center-right). However, given that the selected morphology determines the corresponding controller, an additional procedure ensures that robot behavior and morphology are appropriately matched. Specifically, a chosen morphology with x active and y inactive sensors automatically re-configures the chosen controller so

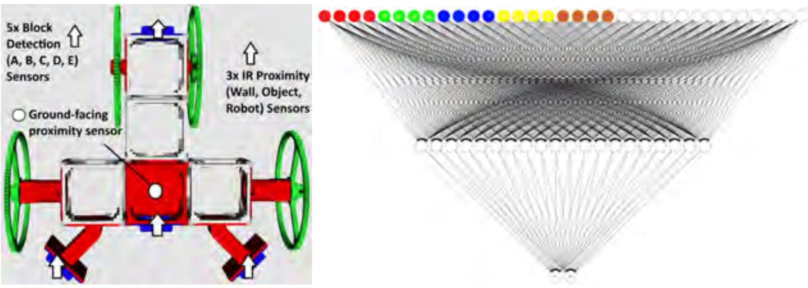

Fig. 2. **Left:** Each robot's chassis is equipped with four sensor sets (marked blue), one downward-facing target-area detection sensor (white circle), below the core-component (red-cube). Each of the four sensor sets comprises five block detection sensors and three IR proximity sensors (white arrows). **Right:** Each robot uses a feed-forward ANN controller that fully connects 33 sensory input nodes to 10 hidden layer nodes and two motor output nodes. Motor output nodes control each wheel's speed and thus robot direction and speed. The first 20 sensors comprise four of each sensor type, for detecting block types A (red), B (green), C (blue), D (yellow), and E (brown) blocks, with one of each block sensor type positioned at the robot's front, back and rear-left and rear-right positions (marked blue). The next 12 sensors (marked white) comprise three types of IR proximity sensors at four periphery sensor positions (marked blue). IR sensors are calibrated for detecting obstacles, walls, and robots. The final (far-right) sensor is a ground-facing target-area detection sensor to determine if the robot's current position is in the target-area (for delivering gathered blocks).

that x sensory input nodes are active and y are inactive.

As in EDQD-M (Section 2.4), ANN controller connection weights remain active, and robot sensory configuration (morphology) is adapted by switching specific sensors on and off, with zero values as controller inputs corresponding to switched-off sensors. Thus, Double-Map EDQD-M adapts both *LocalMaps* to promote morphological and behavioral diversity. In contrast to EDQD and EDQD-M (only accounting for behavioral diversity maintenance), these two maps (*LocalMap-1*, *LocalMap-2*, Figure 1) enable both robot behavior and morphology to be subject to diversity maintenance.

2.6 Robot Controllers

Robots explore the environment for their *lifetime* duration (Table 3), using ANN controllers (Figure 2, right), where ANN behavior is adapted by either mEDEA, mEDEA-M, EDQD, EDQD-M or Double-Map EDQD-M (Sections 2.1, 2.2, 2.3, 2.4, 2.5). Each robot (Figure 2, left) uses the same controller topology, a fully connected feed-forward ANN comprising 33 sensory input nodes (proximity, color, target-area detection), a 20 node hidden layer, and 2 motor output nodes (Table 3). To be consistent with previous work [Bredeche et al. 2012; Hart et al. 2018], all nodes in the ANN used Sigmoidal activation units. The two ANN outputs were the rotational and translational speed of each robot (normalised to: $[0, 1]$) at each simulation iteration (of the robot lifetime).

The sensory input nodes corresponded to three forward-facing proximity sensors, one backward-facing proximity sensor, and a bottom-facing target area detection sensor (constantly active). Proximity sensors were primed to detect the closest object in the environment, where the closer an object to the robot, the higher the sensor activation value (normalised to: $[0, 1]$). For each forward and backward-facing proximity sensor, there were seven object type (color) detection sensors that activated to discriminate between the colors of five resource types, walls, and other robots (Table 3). A robot's periphery comprised four sensor sets (each containing eight sensors) and one downward-facing target-area detection sensor, where these 33 sensors corresponded to the ANN

Table 2. Parameters for robot ANN controllers and swarm evolution methods (Sections 2.1-2.5).

Sensory input nodes	33
Hidden layer nodes	20
Motor output nodes	2
Node activation function	Sigmoidal
Sensory input-Motor output weight range	$[0.0, 1.0]$
Neuron weight range	$[-400, +400]$
Mutation operator	Gaussian (pre-tuned σ) [Hart et al. 2018]
Sigma range	$[0.001, 0.5]$
Update sigma step	0.35
Mutation probability	0.34
Sensor-morpho threshold	≤ 0: Sensor inactive; > 0: Sensor active
Map archive size	100
Number of dimensions per (behavior, morphology) map	2
Number of intervals per map-dimension	10

Table 3. Experiment and collective gathering task parameters.

Resource-types (size: x, y, z: meters)	A	0.08 x 0.08 x 0.08
	B	0.50 x 0.50 x 0.08
	C	0.8 x 0.8 x 0.08
	D	1.0 x 1.0 x 0.08
	E	1.2 x 1.2 x 0.08
Sensor types: Range	Infrared Proximity	$[0.0, 1.0]$
	Color	$[0.0, 1.0]$
	Target-area detector	Bottom facing
Run length (per experiment)	100 generations	
Robot lifetime	10000 (simulation iterations)	
Swarm size	100 robots	
Wait for assistance time (cooperative resource-pushing)	Remaining lifetime	
Initial robot & block position	Random (Outside target-area)	
Environment size | Target-area size (meters)	20 x 20 | 20 x 2	
Robot *LocalMap* broadcast range	Environment size	
Robot *LocalMap* broadcast frequency	1 (per lifetime)	

Table 4. Collective gathering task environments and associated complexity.

Task Environment	**Environment Type**	**Task Complexity**	**Resource Type Combination** (A, B, C, D, E)
1	Flat	Simple	42, 2, 2, 2, 2
2	Flat	Medium-Low	30, 5, 5, 5, 5
3	Flat	Medium	10, 10, 10, 10, 10
4	Flat	Medium-High	5, 5, 5, 5, 30
5	Flat	Difficult	2, 2, 2, 2, 42
Cooperation required to move resource type	*A*	1 robot	
	B	2 robots	
	C	3 robots	
	D	4 robots	
	E	5 robots	

input layer. The controller genotype adapted by each method (Sections 2.1-2.5), thus comprised 700 connection weights, that is, 33 input nodes fully connected to 20 hidden nodes (33x20 connections), fully connected to two output nodes (20x2 connections). We also note that while the robots operate in a 3D environment (Section 3.2), each robot's sensory construct is such that it experiences the environment in 2D. That is, sensors positioned on each robot's periphery (Figure 2) perceive block types and obstacles at a given angle and distance (in a 2D plane) relative to the robot's current position in the environment. Additionally, the target-area sensor detects if the robot's current position in the environment is within the 2D space that defines the *target-area* for gathered blocks.

3 EXPERIMENTS

This study uses an evolutionary swarm robotics simulator[2] to address our research objectives (Section 1), via running swarm robotic evolution experiments with various robot controller and controller-morphology evolution methods (Sections 2.1-2.5) and environments (Table 4), to observe how behavioral and morphological diversity impacts collective (swarm) behavior evolution. Five experiment sets (Table 5) were conducted to address our research objectives. Each experimental run entailed evaluating one of the swarm evolution methods in a given task environment, via running a swarm for one *lifetime* (10000 simulation iterations) and 100 generations, where one generation represented a swarm lifetime. Each run, robots and resources were re-initialized in new random positions and orientations. To complete the collective gathering task, robots had to search the environment for resources, and then cooperatively move found resources to the *target-area*.

3.1 Task Environments

Task difficulty was determined by the level of cooperation required for robots to collect all resources in the environment. Each resource type (*A*, *B*, *C*, *D*, *E*) differed in geometric size (Table 3) and so required varied degrees of cooperation to transport to the *target-area*. Table 4 presents the five task environments tested in this study, where task environment complexity is determined by the degree of cooperation required for robots to move a given resource type (*A*, *B*, *C*, *D*, *E*). *Simple* environments, for example, contained a high proportion of *type*

---

[2] The swarm-robotic simulator is available at: https://github.com/Body-Brain-QD-in-Robotic-Swarms/TELO-2023

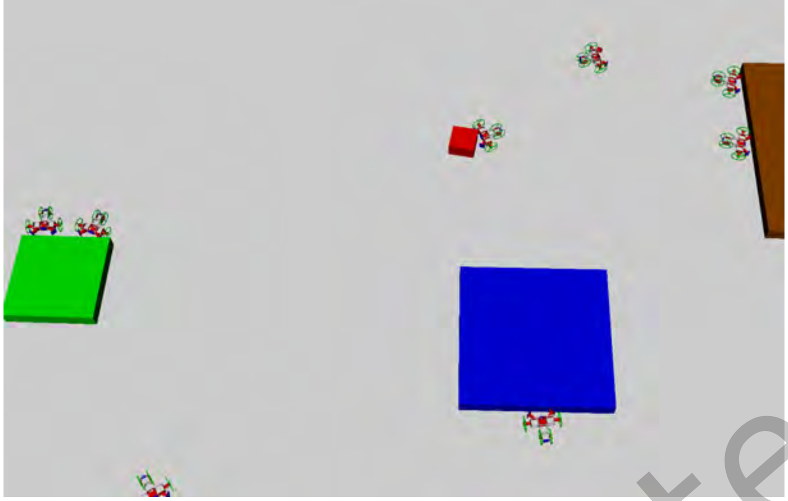

Fig. 3. Example collective gathering task environment. Task difficulty is tuned via distribution of resource types (Table 4): A (red), B (green), C (blue) and D (brown). Resource type E is not shown in this example.

A resources (collectable without the need for cooperation), whereas *difficult* environments contained a high proportion of *type E* resources (requiring five cooperating robots to transport). Figure 3 presents an example task environment illustrating robots attempting to cooperatively move various resource types. This experimental setup, using task difficulty equated with the degree of cooperation required for task accomplishment, is motivated by the success of previous evolutionary robotics experiments investigating emergent altruism in robot groups foraging for resources [Waibel et al. 2011].

Furthermore, it is important to note that this experimental setup using the degree of cooperation required between robots, is only one possible mechanism to *tune* task difficulty in the environment. We selected this definition given the simplicity in adjusting task difficulty via manipulating the composition of block types (Table 4), and the established use of this task difficulty definition in previous work [Hewland and Nitschke 2015; Nagar et al. 2019a]. However, ongoing work is integrating other experimental mechanisms for tuning task difficult including the slope and traction of the environment's surface, obstacles in the environment [Spanellis et al. 2021], and variable morphologies indicative of sensor and actuator damage that thus limits possible behaviours [Putter and Nitschke 2017]. The use of such task difficulty factors in related work has demonstrated that varying definitions of task difficult impact body-brain evolution in significantly different ways. Thus future work will consider the impact of a broad range of task difficulty definitions on body-brain evolution in robotic swarms.

3.2 Swarm Behavior Evolution and Evaluation

Table 5 presents an experiments overview, where experiments are defined by robot behavior or behavior-morphology evolution given application of mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M. The objective of all experiments is to investigate the impact of varying degrees of task complexity (Table 4) on evolved collective gathering behavior in robotic swarms, given behavioral diversity maintenance (EDQD, EDQD-M) versus behavior-morphology diversity maintenance (Double-Map EDQD-M) versus no diversity maintenance (mEDEA, mEDEA-M).

For each experiment, a 3D environment (*simple-difficult* task environments, Table 4) was initialized with a swarm of 100 robots and 50 resources in random positions and orientations. Robots and resources were set up outside of a *target-area* (where the collected resources were delivered to). For all swarm evolution methods (Sections 2.1-2.5), the percentage of resources pushed into the *target-area* in each swarm lifetime was used to compute average (over 30 runs) swarm task performance (behavior quality, equation 1), normalized to the range: [0.0, 1.0].

$$behavior\ quality = \frac{resourceValue}{numberPushingRobots} * \frac{distanceResourceMoved}{totalDistancePossible} \tag{1}$$

In equation 1, *resourceValue* is the minimum number of robots required to push a resource, *numberPushingRobots* is the number of robots pushing a given resource, *distanceResourceMoved* is the straight line distance that the resource is moved towards the gathering area, and *totalDistancePossible* is the straight line distance from one end of the environment to the center of the target-area.

Each swarm comprised 100 robots with identical morphologies (sensory-motor configurations) and ANN controllers. For mEDEA-M and EDQD-M, the swarm remained *morphologically homogeneous*, meaning any controller-morphology mutations were concurrently applied to all robots in the swarm, while the swarm evolved to be *morphologically heterogeneous* for Double-Map EDQD-M. However, the swarm was *behaviorally heterogeneous* for all methods, given that each robot was initialised with varying random ANN connection weight values and ANN connection weight mutations were independently applied to each robot during behavioral adaptation (Section 2).

3.3 Behavioral and Morphological Diversity Evaluation

Each run (100 generations, Table 2), we determined the behavioral diversity for EDQD, EDQD-M, and Double-Map EDQD-M as the number of distinct behaviors (occupied cells in *LocalMap*, sections 2.3-2.5). Average swarm behavioral diversity was then computed over 30 runs of EDQD, EDQD-M, and Double-Map EDQD-M. For Double-Map EDQD-M, we calculated a swarm's average morphological diversity (over 30 runs) as the number of distinct morphologies (occupied cells in *LocalMap2*, section 2.5) per run (100 generations, Table 2). The mEDEA and mEDEA-M methods do not use behavior map diversity maintenance; however, for the purposes of comparative analysis, per generation, a behavior map was created and updated (following the same procedure as the EDQD methods). Similarly, mEDEA-M and EDQD-M do not use morphology maps, but in order to compare adapted morphologies with Double-Map EDQD-M, per generation, a morphology map was created and updated (following the same procedure as Double-Map EDQD-M). Such behavior and morphology maps were not used by the evolutionary processes of these other methods and were only maintained for comparative results analysis (Section 4). For consistency with previous work [Pugh et al. 2015], we also calculate the average (over 30 runs per experiment, Table 5) QD score for behaviors evolved by mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M. We further calculate the average (over 30 runs per experiment) QD score for morphologies evolved by mEDEA-M, EDQD-M, and Double-Map EDQD-M.

3.4 Simulator

All experiment simulations used the *RoboGen* [Auerbach et al. 2018] open-source framework. *RoboGen* has a 3D physics simulation and artificial evolution engine and includes utilities for creating 3D-printable design files for body components and compiling neural network controllers to run on an *Arduino* micro-controller boards[3]. Since the original platform could only simulate one robot, we extended *RoboGen's* simulation engine to allow for

---

[3] https://www.arduino.cc/

Table 5. Experiments evaluate swarm evolution methods for adapting either robot behavior (Sections 2.1, 2.3) or behavior-morphology (Sections 2.2, 2.4, 2.5), across task complexity to address our objectives (section 1).

Experiment	**Method**	**Task Complexity**	**Research Objective**
1	mEDEA	*Simple - Difficult*	Investigate the impact of task complexity on behavioral diversity and swarm behavior quality in fixed-morphology swarms.
2	mEDEA-M	*Simple - Difficult*	Investigate the impact of task complexity on behavioral-morphological diversity and swarm behavior quality in adaptive-morphology swarms.
3	EDQD	*Simple - Difficult*	Investigate the impact of task complexity on behavioral diversity and swarm behavior quality in fixed-morphology swarms.
4	EDQD-M	*Simple - Difficult*	Investigate the impact of task complexity on behavioral-morphological diversity and swarm behavior quality in adaptive-morphology swarms.
5	Double-Map EDQD-M	*Simple - Difficult*	Investigate the impact of task complexity on behavioral-morphological diversity and swarm behavior quality in adaptive-morphology swarms.

swarm robotic simulation. This entailed numerous design considerations, such as performance, accuracy, and handling potential communication between robots. The extended version[4] used in this study is written in C++. All simulations were executed on the *Centre for High Performance Computing* (CHPC) Lengau cluster[5], using Intel 5th generation CPUs on 1368 compute nodes with 24 cores and 128 GiB memory. Local implementation and testing were conducted on a Linux Ubuntu 22.04.2 LTS computer with an 11th Gen Intel Core i7-1165G7 (@ 2.80GHz × 8) processor, Intel iRISxe graphics and 32 GB of RAM.

4 RESULTS AND DISCUSSION

This section first examines comparative task-performance (behavior quality) differences between respective swarm evolution methods (section 4.1), relates this to exhibited differences in evolved behavioral (section 4.2) and morphological (section 4.3) quality-diversity and neuro-morpho complexity (section 4.4). For all statistical significance tests presented in this section, results data were found to be non-parametric using the Kolmogorov–Smirnov normality test with Lilliefors correction [Ghasemi and Zahediasl 2012]. Mann–Whitney U statistical tests ($p<0.05$) [Flannery et al. 1986] were then applied in pair-wise comparisons with Effect Size [Cohen 1988] treatment[6].

4.1 Swarm-Robotic Task-Performance

Figure 4 presents average task performance (quality) progression over 100 evaluations (generations), for robotic swarms evolved by each method (Section 2, Table 1), in each task environment (Table 4). Average swarm quality was calculated at the end of each run (100 generations) and over 30 runs. The highest-performing swarms evolved in the *simple* task environment, followed by the *medium-low* environment, the *medium* environment, and the *medium-high* environment. The least effective swarms evolved in the *difficult* environment. To comprehend

---

[4] The simulator used in this study is available online: https://github.com/Body-Brain-QD-in-Robotic-Swarms/TELO-2023

[5] https://wiki.chpc.ac.za/chpc:lengau

[6] Complete statistical test details are in Supplementary Material (A.1: Statistical Comparisons): https://github.com/Body-Brain-QD-in-Robotic-Swarms/TELO-2023

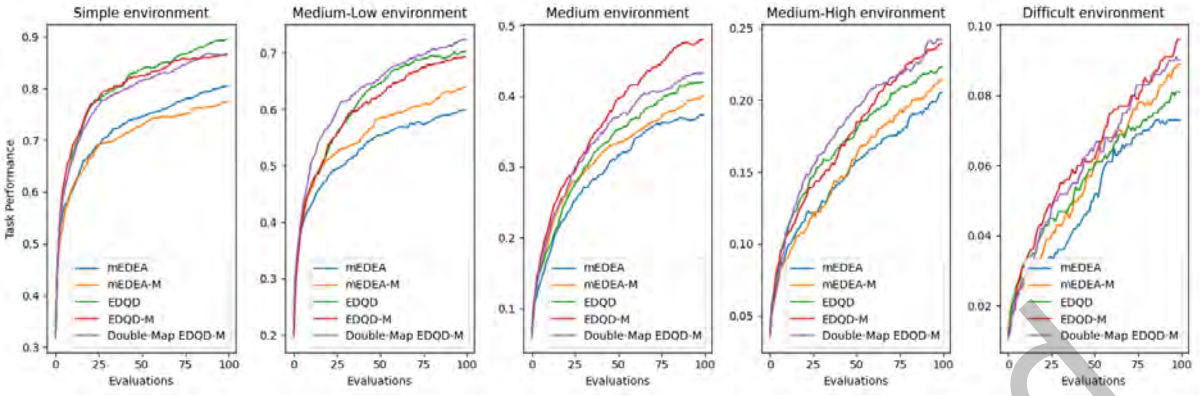

Fig. 4. Average swarm task performance (behavior quality) over 100 evaluations (generations), where task performance of 1.0 indicates all resources in environment gathered during the swarm's lifetime. Note that the *difficult* environment has been zoomed into the task-performance range: [0.0, 0.1] so as one can readily discern performance differences between the respective methods.

these average swarm quality differences, it is beneficial to examine the characteristics of the various environments.

In the *simple* and *medium-low* task environments, resources were primarily type A (Table 4), representing 84% and 60% of all resources, respectively. These could be collected individually without the need for cooperation. The remaining 16% and 40% of resources were evenly distributed among types B, C, D, and E, requiring the cooperation of two, three, four, and five robots to gather (Table 4). Thus, in the *simple* and *medium-low* environments, near-optimal task performance is achievable without extensive cooperation. Moreover, with enough robots (100), there was enough concurrent resource gathering to result in nearly optimal swarm task performance. In the *simple* environment, EDQD evolved the highest quality swarms, significantly outperforming mEDEA (Mann-Whitney U, $p<0.05$), while no statistical difference was observed between mEDEA and all the other methods (Mann-Whitney U, $p \geq 0.05$). All *Quality Diversity* (QD) based methods (EDQD, EDQD-M, Double-Map EDQD-M, section 2), significantly outperformed mEDEA-M, whereas there was no statistical difference between the QD methods. Conversely, in the *medium-low* environment, all QD methods significantly outperformed mEDEA (Mann-Whitney U, p<0.05), while there was no statistical difference between mEDEA and mEDEA-M. There was also no significant difference between the QD methods, indicating that methods focused on evolving behavioral diversity are better suited for evolving collective behavior, even without morphological adaptation (EDQD). This is particularly evident in the *simple* and *medium-low* environments, where significant resource portions can be collected concurrently without the need for high degrees of cooperation (Table 4).

In the *medium* environment, all resource types (A, B, C, D and E) were evenly distributed (Table 4). One-fifth of the resources could be collected individually, while all remaining resources required varying degrees of cooperation. Consistent with previous work [Furman et al. 2019], swarm quality was enhanced given the relative ease of acquiring a fifth of total resources without cooperation. In the *medium* environment, EDQD-M significantly outperformed all other methods (Mann-Whitney U, $p<0.05$). However, there was no statistical difference between mEDEA-M, EDQD, and Double-Map EDQD-M, highlighting the benefit of combining behavioral diversity maintenance and morphological adaptation (EDQD-M), in specific environment types (*medium* task environment in this case).

In the *medium-high* and *difficult* environments, 60% and 84% of resources were type *E*, necessitating the cooperation of five robots. The remaining 40% and 16% of resources (for the *medium-high* and *difficult* environments, respectively), were evenly distributed among resource types *B*, *C*, *D*, and *E*. These environment configurations resulted in lower average evolved swarm quality for all methods, compared to swarms evolved in the *simple*, *medium-low*, and *medium* environments (Figure 4). Also, cooperation used a *wait-for-help* period (Table 3), that entailed robots waiting for help whenever they encountered type *B-E* resources. Thus, in environments requiring high degrees of cooperation, for example, given many type *E* resources (Table 4), robots spent a significant portion of their lifetime awaiting assistance. This resulted in significantly reduced average behavior quality for swarms evolved (by all methods) in the *medium-high* and *difficult* environments. This observation is further supported by the lack of any statistically significant differences between all methods in the *medium-high* and *difficult* environments (Mann-Whitney U, $p \geq 0.05$).

Thus, the average behavior quality elicited by evolved swarms indicates that QD based methods using behavioral diversity maintenance (EDQD, EDQD-M, Double-Map EDQD-M), yield significantly higher average quality, compared to mEDEA and mEDEA-M. However, this significant difference only holds for the *simple*, *medium-low* and *medium* environments, with only EDQD-M yielding a significantly higher quality over other methods in the *medium* environment (Mann-Whitney U, $p<0.05$). Though figure 4 indicates that as task complexity increases (from *medium-low* to *difficult*), that a QD based method invariably yields the highest average quality overall. Specifically, swarms evolved by either EDQD-M (*medium* environment) or Double-Map EDQD-M methods (*medium-low*, *medium-high* and *difficult* environments). Related work supports such task performance benefits enabled by behavior-morphology diversity maintenance in controller-morphology adaptation in increasingly complex (single robot ambulation) tasks [Miras and Eiben 2019; Miras and Ferrante 2020], as does previous work with swarm-robotic cooperative gathering using behavior-morphology diversity maintenance in controller-morphology evolution [Mkhatshwa and Nitschke 2023]. In summary, these task performance results support the benefits of morphology adaptation coupled with behavior (EDQD-M) and behavior-morphology diversity maintenance (Double-Map EDQD-M). The impact of behavioral and morphological diversity maintenance is further discussed in sections 4.2 and 4.3.

4.2 Evolved Behavior Quality-Diversity (QD)

Swarm behavior evolved by mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M (section 2), was characterized by two behavioral descriptors. First, resource types collected, and second, the maximum distance covered in a robot's lifetime (Table 3). These descriptors enabled behavioral diversity in terms of the types of resources collected and the portion of the environment explored (section 3.3). The QD score [Pugh et al. 2015] was calculated (section 3.2, equation 1) as the overall *quality* (task performance) in filled cells within a QD behavior map[7] (section 2). A high average QD score indicates swarms with high average behavioral diversity (resource types gathered and distance covered) and quality (many resources of all types gathered).

Figure 5 presents the average QD score (over 30 runs) of swarms evolved by each method in each environment. In the *simple* environment, EDQD evolves swarms with the highest average QD score over 100 evaluations (generations), indicating swarm adaptation in the *simple* environment (demanding the least cooperation) benefits

---

[7] QD behavioral map visualisations of average swarm behavior are available in Supplementary Material (A.2: Evolved Behavior Quality Diversity): https://github.com/Body-Brain-QD-in-Robotic-Swarms/TELO-2023.

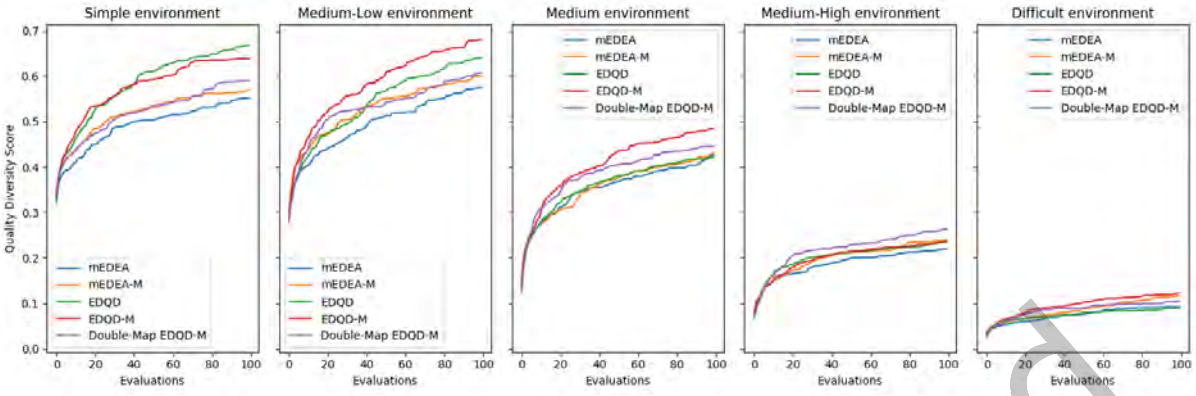

Fig. 5. Average behavior *Quality-Diversity* (QD) progression over the 100 evaluations (generations) of a run (average computed over 30 runs), for swarms evolved by each method in each task environment.

from behavioral diversity maintenance (EDQD) but does not require morphological adaptation. In the simple environment, 84% of the resources can be moved without requiring cooperation (Table 4), so EDQD, promoting behavioral diversity in company with behavioral optimisation, is sufficient for evolving diverse, high quality swarm behaviors.

However, as task difficulty increases, behavior diversity maintenance coupled with morphological adaptation becomes advantageous. For example, in the *medium-low*, *medium* and *difficult* environments, EDQD-M yields the overall highest average QD score, and in the *medium-high* environment, Double-Map EDQD-M evolved swarms yield the highest average QD score. Given increasing task complexity, this underscores the benefits of coupling morphological adaptation with behavioral diversity maintenance during swarm behavior optimisation (EDQD-M, Double-Map EDQD-M).

The efficacy of EDQD-M and Double-Map EDQD-M for evolving high-quality swarm behaviors is also supported by the behavior QD maps visualising environment exploration and gathering of diverse resource types for swarms evolved by EDQD-M and Double-Map EDQD-M in the *medium-low* to *difficult* environments. These demonstrated benefits of behavioral diversity maintenance and morphological adaptation in cooperative tasks of increasing complexity are supported by comparable previous work. Specifically, collective robotic task performance benefits from morphological (sensory configuration) adaptation across increasingly difficult cooperative transport tasks [Furman et al. 2019; Hewland and Nitschke 2015; Nagar et al. 2019a,b]. The benefits of behavioral diversity maintenance for boosting collective robotic task performance across increasingly difficult cooperative (RoboCup Keepaway soccer) tasks have similarly been demonstrated [Nitschke and Didi 2017]. As in this study (section 3.1), task complexity in such related work was equated to the degree of cooperation required to achieve optimal collective behavior task performance. For example, robots required for cooperative transport of each resource type in collective gathering [Furman et al. 2019; Hewland and Nitschke 2015; Nagar et al. 2019a,b], and number of *taker* robots that *keeper* robots must keep a passed ball away from in Keepaway [Nitschke and Didi 2017]).

Overall, morphological adaptation in company with behavior (EDQD-M) or behavior-morphology (Double-Map EDQD-M) diversity maintenance is most beneficial in terms of diversity of evolved behaviors and behavior quality across task environments. EDQD-M evolves swarms with the highest average QD score in the *medium-low*, *medium*, and *difficult* environments, and Double-Map EDQD-M evolves swarms yielding the highest QD score in the *medium-high* environment. Though, for all environments, the average QD scores of swarms evolved by EDQD-M and Double-Map EDQD-M were statistically comparable. In summary, these results partially address this study's first research question (section 1.1), via demonstrating the benefits of behavior diversity maintenance (EDQD-M and Double-Map EDQD-M) in evolving swarms eliciting high quality (collective gathering) behaviors effective across increasingly complex environments. To fully address this study's first research question, section 4.3 discusses the impact of morphological adaptation coupled with behavior and behavior-morphology diversity maintenance on evolved swarm behavior quality.

4.3 Evolved Morphology Quality-Diversity (QD)

Figure 6 presents the average (over 30 runs) morphology QD score of swarms evolved by mEDEA-M, EDQD-M, and Double-Map EDQD-M, per task environment. As with the QD behavior score (section 4.2), the QD morphology score was the overall swarm behavior quality (task-performance) across all filled cells but within the QD morphology map (used by morphology-adaptation methods: mEDEA-M, EDQD-M, and Double-Map EDQD-M). Figure 7 displays the average (over 30 runs) number of unique morphologies evolved by each method in each environment, calculated by tallying the number of QD morphology map (*LocalMap* for morphologies, section 2.5) cells filled per run for mEDEA-M, EDQD-M, and Double-Map EDQD-M. Note that mEDEA-M and EDQD-M do not retain a morphological map (*LocalMap*) during their swarm evolution process. Thus, for these methods, a morphological map was generated exclusively for analysis. For all QD morphology maps, morphological descriptors were designed to promote diverse usage of sensor ranges and active sensors in emergent unique morphologies. The QD morphology maps also indicate how effective evolved morphologies are for evolving swarm behaviors effective across various task environments. Figures 8-12 present QD morphology map visualisations, displaying the highest-performing swarm morphologies evolved by each method per environment. A high average morphology QD score indicates swarms with high morphological diversity coupled with high swarm behavior quality.

Figure 6 indicates the morphology QD score of swarms evolved by Double-Map EDQD as significantly higher (Mann–Whitney U, $p<0.05$), for all environments, than the morphology QD scores of mEDEA-M and EDQD-M evolved swarms. Figure 6 thus indicates the benefit of behavior-morphology diversity maintenance, coupled with morphological adaptation. That is, the diverse array of evolved morphologies with associated high behavioral quality (task performance) is reflected in terms of high morphology QD scores for Double-Map EDQD-M per environment. The diverse range of morphologies evolved per environment is evidenced by the significantly higher (Mann–Whitney U, $p<0.05$) number of morphologies evolved by Double-Map EDQD-M versus morphologies evolved by mEDEA-M and EDQD-M, per environment (Figure 7). The advantage of an evolved diverse range of morphologies is also evidenced by the high average behavior quality of swarms evolved by Double-Map EDQD-M per task environment (Figure 4). Specifically, the benefits of morphological adaptation coupled with both behavioral and morphological adaptation in Double-Map EDQD-M (section 2.5) versus other morphological adaptation methods (employing behavioral diversity maintenance), mEDEA-M and EDQD-M, are most evident from the morphological QD maps of swarms evolved by Double-Map EDQD-M. Figures 8-12 present the morphological QD maps of the highest quality swarms evolved by Double-Map EDQD-M in the *simple-difficult* environments, indicating Double-Map EDQD-M as consistently effective in evolving morphological diversity and generating high-quality behaviors across all environments.

Figures 8-12 visualise the portion of the morphological space (average sensor range, average sensor ratio), explored by the morphological adaptation methods: mEDEA-M, EDQD-M and Double-Map EDQD-M. For each environment, the effectiveness of exploration versus exploitation in Double-Map EDQD-M swarm behavior evolution is evident from the greater portion of the morphological space covered with associated higher quality behaviors (darker shading of more cells in Double-Map EDQD-M morphology QD maps in figures 8-12). Specifically, the best-performing morphologies evolved by Double-Map EDQD-M exhibited active sensor portions and average active sensor ranges ranging from 0.1 to 1.0 (Figures 8–12). That is, morphologies with quality exceeding 90%, were observed when sensors were active between 65% and 95% and operated at 85% to 100% of maximum range (Figures 8-12). Figures 8-12 thus indicate sensor activity and sensor ranges evolved by Double-Map EDQD-M were suitable for achieving high quality swarm (collective gathering) behaviors in each environment. However, comparable quality was observed for morphologies evolved by EDQD-M (employing behavioral diversity maintenance and morphological adaptation) in the *simple*, *medium-low*, *medium-high*, and *difficult* environments.

Figures 8-12 also highlight, for all methods and environments, the highest average quality morphology is that using all sensor types at maximum range. However, Double-Map EDQD-M evolved morphologies are characterized by many (Figure 7), diverse combinations of sensor ranges and associated active sensor ratios (Figures 8-12). Morphological value is evidenced by Double-Map EDQD-M exhibiting higher quality compared to counterpart morphologies in mEDEA-M and EDQD-M evolved swarms in terms of a significantly higher (Mann-Whitney U, $p<0.05$) average morphology QD scores for all environments (Figure 6). The benefit of these morphologies is also evident in the average quality of evolved swarms across all environments (Figure 4), where Double-Map EDQD-M yields the highest quality (except in the *simple* and *medium* environments). However, in the *simple* environment, morphological adaptation was demonstrated as unnecessary since most resources (84%), can be gathered concurrently without cooperation (Table 4). Whereas, in the *medium* environment, gathering most resources (80%) required some cooperation (Table 4). In this case, behavioral diversity maintenance coupled with morphological adaptation (EDQD-M) was sufficient, meaning swarms in the *medium* environment did not benefit from behavior-morphological diversity maintenance in addition to morphological adaptation (Double-Map EDQD-M).

Overall, these results highlight the QD benefits of behavioral diversity maintenance and behavior-morphology diversity maintenance when coupled with morphological adaptation (EDQD-M, and Double-Map EDQD-M) and, more generally, the benefits of morphological adaptation (mEDEA-M) as task complexity increases. These results are supported by related evolutionary robotics work that has similarly demonstrated the benefits of behavioral and morphological diversity maintenance in controller-morphology adaptation in robots operating in varying environments. For example, they demonstrated the benefits of behavioral diversity maintenance in company with controller-morphology co-evolution across increasingly complex tasks (robot ambulation and gait adaptation) [Babak et al. 2021; Miras and Eiben 2019; Miras and Ferrante 2020; Miras et al. 2018a,c,b].

Furthermore, morphology diversity maintenance [Cheney et al. 2018; Nordmoen et al. 2021; Samuelsen and Glette 2014; Zardini et al. 2021] has also been demonstrated as beneficial for boosting the quality of evolved controller-morphology couplings in robot ambulation [Cheney et al. 2018; Nordmoen et al. 2021] and soft-robotic goal tasks requiring morphological deformation [Zardini et al. 2021]. However, few [Nordmoen et al. 2021] have demonstrated the benefits of morphological diversity across increasingly complex task environments and as

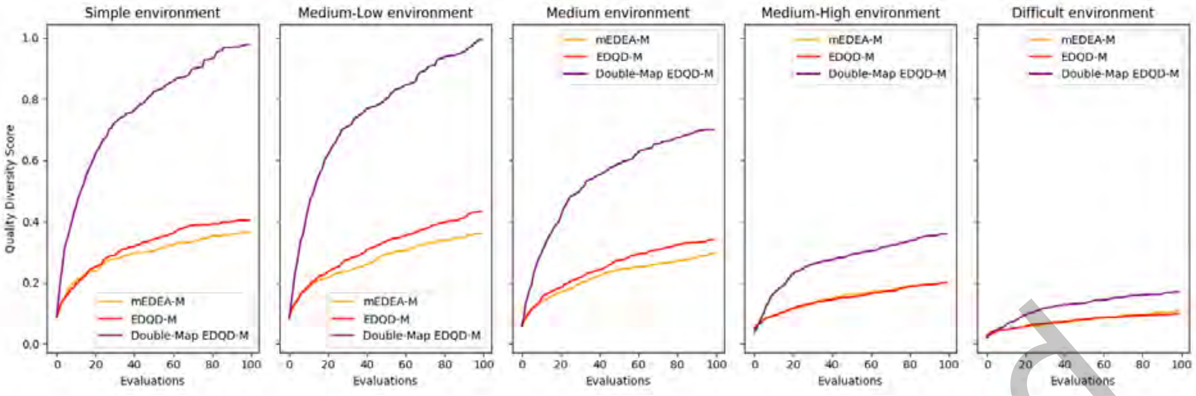

Fig. 6. Average morphology *Quality-Diversity* (QD) progression over the 100 evaluations (generations) of a run (average computed over 30 runs), for swarms evolved by each method in each task environment.

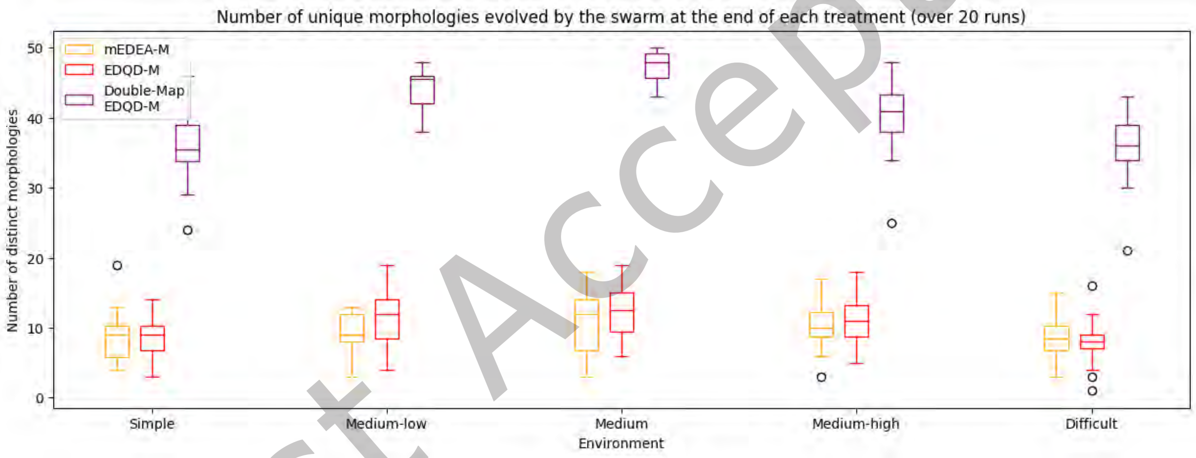

Fig. 7. Average number of distinct morphologies (average computed over 30 runs) of swarms evolved by mEDEA-M, EDQD-M, and Double-Map EDQD in each task environment.

a result showed that as task environment complexity increases, the necessity for morphological diversity also increases.

In summary, these results directly address this study's first research question (section 1.1), and demonstrate the benefits of behavior-morphology diversity maintenance (Double-Map EDQD-M) in evolving diverse, high quality controller-morphology couplings that enable emergent swarm robotic (collective gathering) behaviors that are effective across increasingly complex task environments. Furthermore, these results are relevant to future evolutionary swarm robotic system design methods. For example, to inform optimal sensory configuration designs for specific tasks within given environments. So, if minimal sensor complements operating at specific ranges are

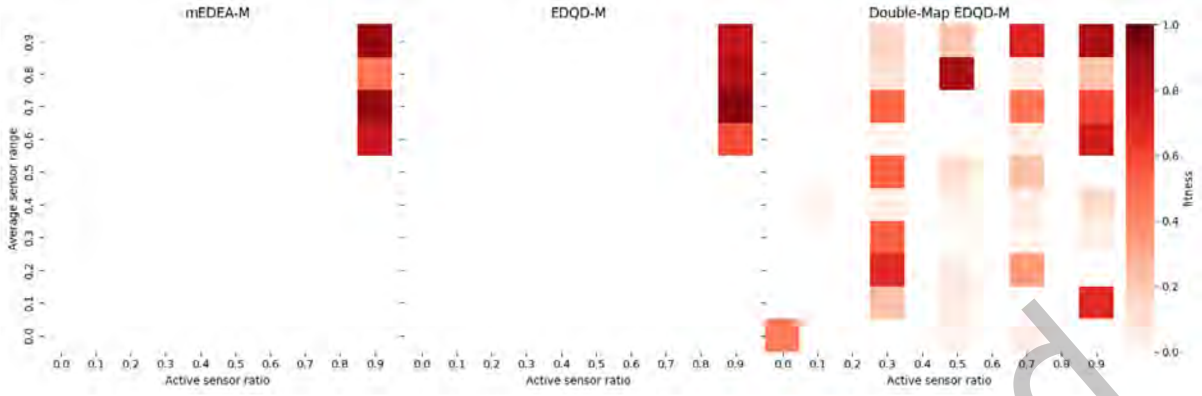

Fig. 8. Morphology QD of highest performing swarms evolved (per method) in the *simple* environment.

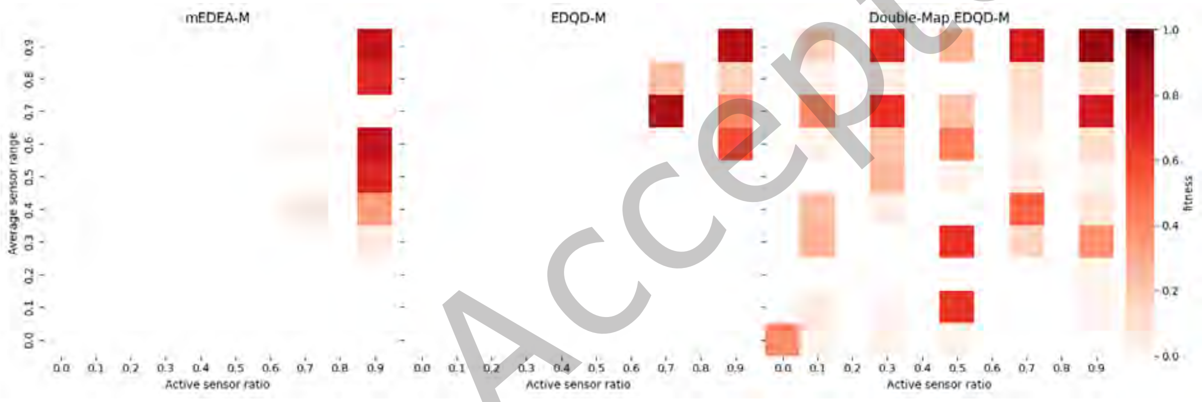

Fig. 9. Morphology QD of highest performing swarms evolved (per method) in the *medium-low* environment.

evolved as effective for given tasks, then designers can avoid the cost of equipping hundreds to thousands of robots in a swarm with maximal sensor complements operating at full capacity, thus reducing overall energy consumption and hardware costs. Similarly, related work has demonstrated devising evolutionary operators that integrate morphological diversity into evolutionary robot (morphology) design and optimisation can boost task performance [Weissl and Eiben 2023].

4.4 Evolved Neuro-Morpho Complexity

Given that sections 4.2 and 4.3 highlighted the swarm behavior-quality, morphology-quality and overall collective behavior benefits of controller-morphology adaptation methods that promote behavioral and morphological diversity, this section further examines the types of controllers evolved by the controller-morphology adaptation methods: mEDEA-M, EDQD-M and Double-Map EDQD-M. Specifically, evolved neuro-morpho complexity, referred to as such since robot neural controller complexity is closely coupled to morphological complexity, that

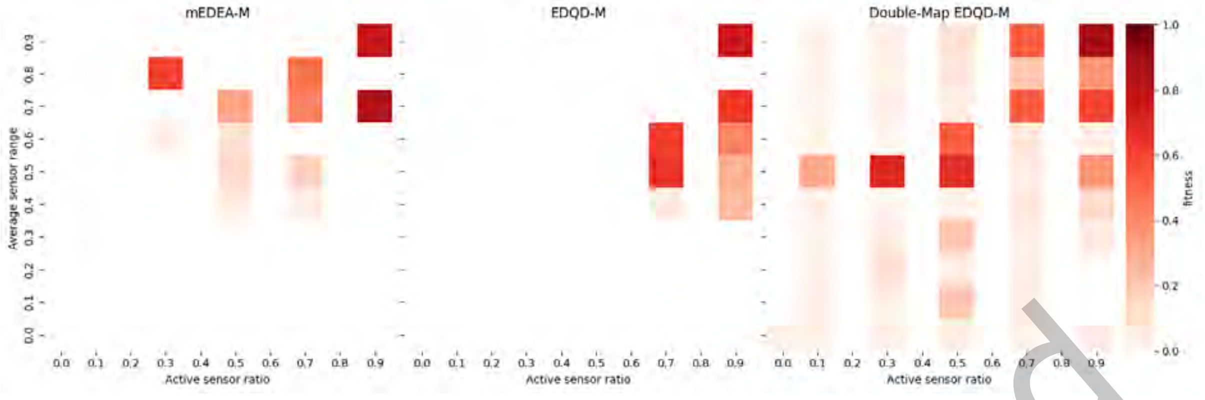

Fig. 10. Morphology QD of highest performing swarms evolved (per method) in the *medium* environment.

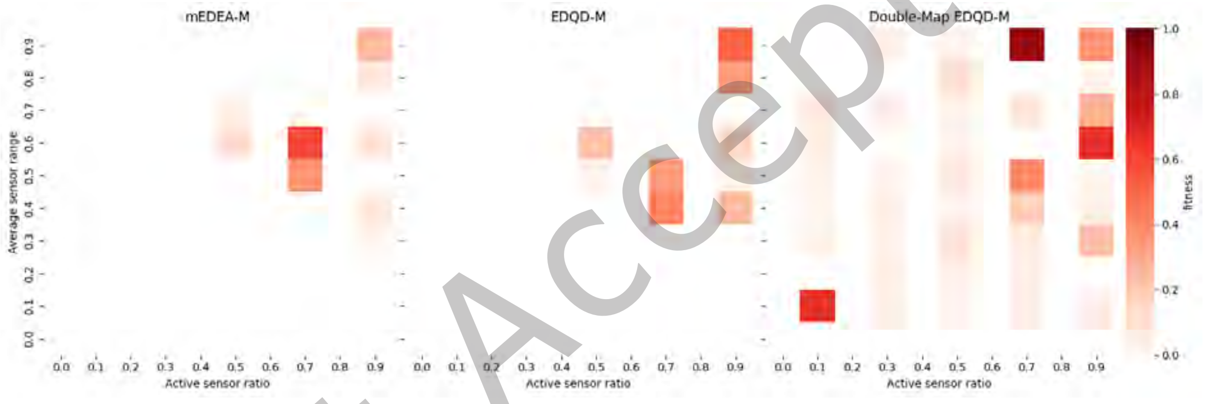

Fig. 11. Morphology QD of highest performing swarms evolved (per method) in the *medium-high* environment.

is, sensory inputs are activated and deactivated via switching controller inputs (and associated connections), on and off (section 2.6).

Figure 13 presents the progression of average (over 30 runs) neuro-morpho complexity for the swarm, evolved by each method in each environment, over 100 evaluations (generations). The neuro-morpho complexity metric (M_c, equation 2), extended the neural complexity metric from previous work [Nitschke and Didi 2017], to account for adaptive sensor configurations and coupled adaptive connectivity between sensory inputs and a robot's neural controller in swarms evolved by the morphological adaptation methods: mEDEA-M, EDQD-M, and Double-Map EDQD-M.

$$M_c = \frac{1}{2}(S_R + S_A) \tag{2}$$

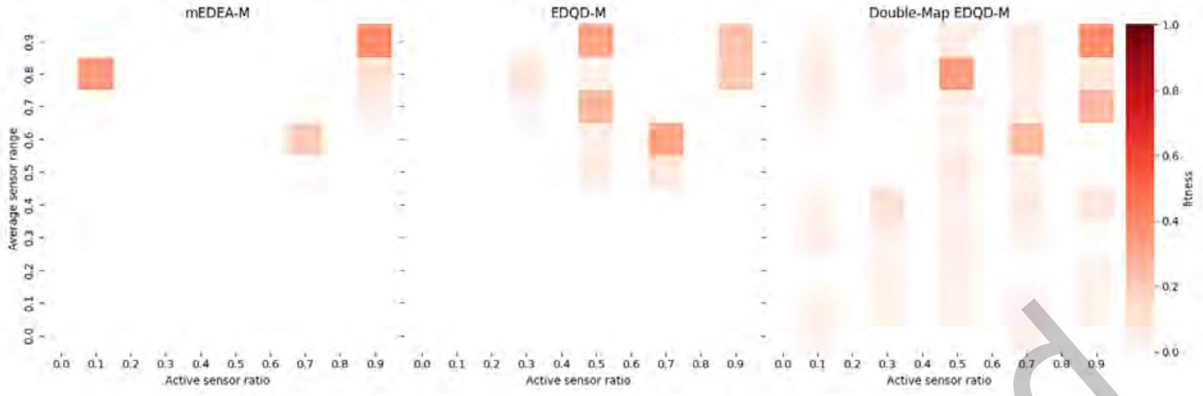

Fig. 12. Morphology QD of highest performing swarms evolved by each method in the *difficult* environment.

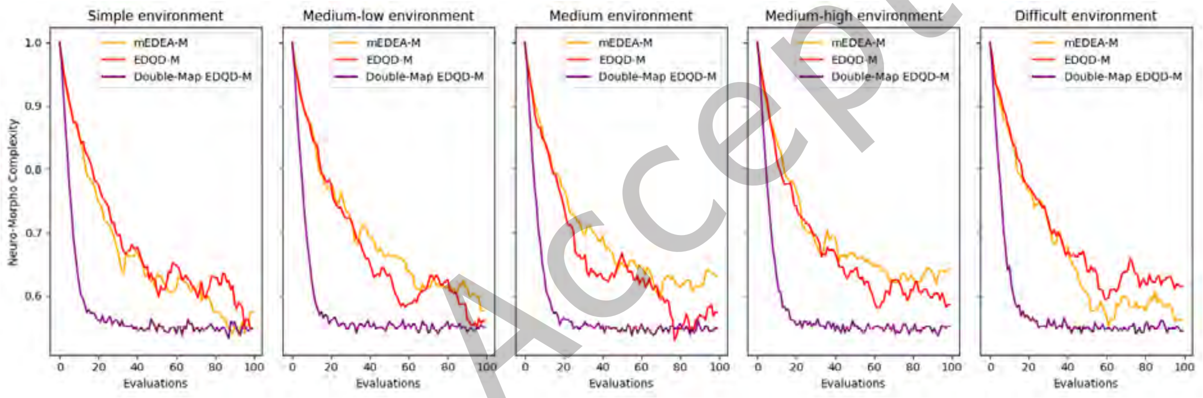

Fig. 13. Progression of average (over 30 runs) morphological complexity evolved by mEDEA-M, EDQD-M, and Double-Map EDQD-M, over 100 evaluations (generations).

Where, S_R represents the normalized average sensor range, and S_A represents the proportion of active sensors. Both S_R and S_A are constrained within the [0.0, 1.0] range. The multiplication by $\frac{1}{2}$ normalizes M_c to the [0.0, 1.0] range. A higher M_c value indicates a greater degree of morphological complexity. Specifically, equation 2 defines the degree of body-brain (neuro-morpho) complexity per robot, since a robot's neural complexity is tightly coupled to its morphology. That is, each active sensor corresponds to a connection weight w_i ($w_i > 0.0$), connecting the input neuron (for the given sensor) to the neural controller's hidden-layer (Figure 2). All robots begin with maximal neuro-morpho complexity ($M_c = 1.0$), meaning that initially all sensors are active and the neural controller's sensory input layer is fully connected to the hidden-layer. Whereas, inactive sensors correspond to inactive sensory-input to hidden-layer connection weights ($w_i = 0.0$), meaning more inactive sensors correspond to lower neuro-morpho complexity.

Figure 13 presents a distinction between the average neuro-morpho complexity of swarms evolved by each morphological adaptation method: mEDEA-M, EDQD-M, and Double-Map EDQD-M. Specifically, swarms evolved by Double-Map EDQD-M, in all environments, quickly converge (after approximately 20 generations), to a minimal neuro-morpho complexity of approximately 0.55), which is retained until run completion (100 generations). Whereas, the neuro-morpho complexity of swarms evolved by mEDEA-M and EDQD-M also initially decreases but then fluctuates between 0.65 and 0.55 after approximately 40 generations. These observed differences in evolved neuro-morpho complexity are supported by statistical comparisons (Mann–Whitney U, $p<0.05$), indicating significantly lower average neuro-morpho complexity of swarms evolved by Double-Map EDQD-M compared to swarms evolved by mEDEA-M and EDQD-M (*medium* and *medium-high* environments), while demonstrating an overall lower complexity for all environments (Figure 13).

These results address this study's second research objective, via indicating that behavioral-morphological diversity maintenance (Double-Map EDQD-M), enables the evolution of relatively low neural complexity (compared to methods employing only behavior diversity maintenance: mEDEA-M, EDQD-M). The interaction of these low-complexity neural controllers, within the context of a swarm, in turn, enables high quality emergent collective gathering behaviors, that are effective across increasingly complex task environments (Figure 4). The effectiveness of evolved low neuro-morpho complexity of robots is supported by the associated (significantly) higher average morphology QD score (Figure 6) and number of distinct morphologies (Figure 7) in swarms evolved by Double-Map EDQD-M (versus mEDEA-M and EDQD-M). The overall effectiveness of swarm (collective gathering) behaviors evolved by Double-Map EDQD-M is also evident in the consistently high average behavior quality across increasingly complex environments (Figure 4). These results are supported by previous work on the evolution of neural complexity in multi-agent (collective behavior) systems. For example, social interaction experiments using evolutionary robotics [Reséndiz-Benhumea et al. 2021] found that when lower neural complexity agents interacted, the complexity of their interactions resulted in emergent group behaviors comparable to those behaviors obtained by more neurally complex single agents. That is, agents with lower neural complexity could enhance their neural complexity through social interaction, thereby offsetting their reduced neural complexity. Similarly, in previous evolutionary robotics experiments [Nagar et al. 2019a], that investigated the impact of neural complexity in collective gathering tasks, results indicated lower (individual robot) neural complexity was sufficient for solving a range of collective behavior (cooperative) tasks with group task performance comparable to that of robots equipped with higher degrees of neural complexity (larger controllers connected to more sensors).

Similar results supporting this study's demonstrated benefits of minimal neuro-morpho complexity have also been reported for other cooperative tasks. For example, across increasingly complex RoboCup Keep-away tasks, where increased robot neural complexity offered negligible benefits to group task performance [Didi and Nitschke 2016; Nitschke and Didi 2017]. This and related previous work supports the notion that under specific environmental and evolutionary conditions, suitably simple neural controllers are sufficient for agents to take advantage of social interaction. Thus, during the evolution of sociality, an individual's neural structure need only become *complex enough*, since it is not an individual's neural complexity, but rather many interacting individuals that enables overall increased complexity (manifest in emergent collective behavior). These results are also pertinent to body-brain complexity evolution in biological social systems. For example, some social insects have evolved simpler bodies and brains than solitary insects [O'Donnell et al. 2015], though such social insect colonies exhibit relatively more complex behaviors [Farris 2016].

5 CONCLUSIONS

This study's objective was to ascertain the impact of behavioral and morphological diversity maintenance on evolving robot controller and controller-morphology (body-brain) couplings within a swarm-robotic system that must solve increasingly complex tasks. Task complexity was equated with the degree of cooperation necessitated between robots to optimally solve a collective gathering task. Results indicated that the Double-Map EDQD-M method, evolving swarm behavior and morphology with mechanisms for behavioral and morphological diversity maintenance, was most beneficial as task environment complexity increased. Specifically, results demonstrated that swarms evolved by Double-Map EDQD-M, operating across increasingly complex environments benefited from the evolution of a diverse range of morphologies to support behavior-morphology evolution. The benefits of morphological diversity were manifest as a significantly higher morphology QD score (compared to other morphological adaptation methods: mEDEA-M and EDQD-M), and a correspondingly high behavior quality (task performance) across all environments. Overall, behavior-morphology diversity maintenance in swarm-robotic systems co-adapting robot behavior and morphology enabled the evolution of morphological diversity needed for evolving sufficiently complex robot neuro-morpho complexity (coupled controller-morphology configurations). The interaction of many such robot controller-morphology couplings in a swarm then resulted in the synthesis of collective gathering behaviors effective across environments of increasing complexity.

These results are supported by related work [Kriegman et al. 2018], demonstrating an inter-dependency in the co-evolution of behavior, morphology, and the environment, where morphology determines, to a large extent, the types and complexity of behaviors that can be exhibited by robot controllers, and changes in the environment drive morphological adaptations, which in turn enable behavioral adaptations [Buresch et al. 2005]. As such, the discovery of diverse morphologies also leads to the evolution of diverse behaviors resulting from adaptation to specialized functional roles [Nitschke et al. 2006]. The evolution of such diverse behaviors and morphologies in turn makes the swarm more adaptable to changes in the environment [Hart et al. 2018]. These results also contribute to the development of future *AutoFac* systems [Nitschke and Howard 2022] to automate swarm robot design with controller-morphology complexity suitable for solving collective behavior tasks in unknown dynamic environments [Furman et al. 2019].

REFERENCES

M. Allard, S. Smith, K. Chatzilygeroudis, B. Lim, and A. Cully. 2023. Online Damage Recovery for Physical Robots with Hierarchical Quality-Diversity. *ACM Transactions on Evolutionary Learning and Optimization* 3, 2 (2023), 23 pages. https://doi.org/10.1145/3596912

R. Arnold, K. Carey, B. Abruzzo, and C. Korpela. 2019. What is A Robot Swarm: A Definition for Swarming Robotics. In *2019 IEEE 10th Annual Ubiquitous Computing, Electronics and Mobile Communication Conference*. 0074–0081. https://doi.org/10.1109/UEMCON47517.2019.8993024

J. Auerbach and J. Bongard. 2009. Evolution of Functional Specialization in a Morphologically Homogeneous Robot. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Montreal, Canada, 89–96.

J. Auerbach and J. Bongard. 2014. Environmental Influence on the Evolution of Morphological Complexity in Machines. *PLOS Computational Biology* 10, 1 (2014), e1003399.

J. Auerbach, A. Concordel, P. Kornatowski, and D. Floreano. 2018. Inquiry-based Learning with RoboGen: An Open-source Software and Hardware Platform for Robotics and Artificial Intelligence. *IEEE Transactions on Learning Technologies* 12, 3 (2018), 356–369.

K. Babak, K. Miras, and A. Eiben. 2021. The Effect of Selecting for Different Behavioral Traits on the Evolved Gaits of Modular Robots. In *The 2021 Conference on Artificial Life*. MIT Press, Prague, Czech Republic, https://doi.org/10.1162/isal–a–00404.

B. Beverly, H. McLendon, S. Nacu, S. Holmes, and D. Gordon. 2009. How Site Fidelity Leads to Individual Differences in the Foraging Activity of Harvester Ants. *Behavioral Ecology* 20, 1 (2009), 633–638.

M. Birattari and et al. 2019. Automatic Off-Line Design of Robot Swarms: A Manifesto. *Frontiers in Robotics and AI* 6 (2019). https://doi.org/10.3389/frobt.2019.00059

Eric Bonabeau, Marco Dorigo, Guy Theraulaz, and Guy Theraulaz. 1999. *Swarm intelligence: from natural to artificial systems*. Number 1. Oxford university press.

J. Bongard. 2011. Morphological Change in Machines Accelerates the Evolution of Robust Behavior. *Proceedings of the National Academy of Sciences* 108, 4 (2011), 1234–1239. https://doi.org/10.1073/pnas.1015390108

M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. 2013. Swarm Robotics: A Review from the Swarm Engineering Perspective. *Swarm Intelligence* 7, 1 (2013), 1–41.

N. Bredeche, E. Haasdijk, and A. Prieto. 2018. Embodied Evolution in Collective Robotics: A Review. *Frontiers in Robotics and AI* 5 (2018), 12.

N. Bredeche and J-M. Montanier. 2010. Environment-Driven Embodied Evolution in a Population of Autonomous Agents. *Parallel Problem Solving from Nature, PPSN XI* (2010), 290–299. https://doi.org/10.1007/978-3-642-15871-1_30

N. Bredeche, J-M Montanier, W. Liu, and A. Winfield. 2012. Environment-Driven Distributed Evolutionary Adaptation in a Population of Autonomous Robotic Agents. *Mathematical and Computer Modelling of Dynamical Systems* 18, 1 (2012), 101–129.

A. Brutschy, N-L. Tran, N. Baiboun, M. Frison, G. Pini, A. Roli, M. Dorigo, and M. Birattari. 2012. Costs and Benefits of Behavioral Specialization. *Robotics and Autonomous Systems* 60, 11 (2012), 1408–1420.

G. Buason, N. Bergfeldt, and T. Ziemke. 2005. Brains, Bodies, and Beyond: Competitive Co-evolution of Robot Controllers, Morphologies and Environments. *Genetic Programming and Evolvable Machines* 6, 1 (2005), 25–51.

T. Buresch, A. Eiben, G. Nitschke, and M. Schut. 2005. Effects of Evolutionary and Lifetime Learning on Minds and Bodies in an Artifical Society. In *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE Press, Edinburgh, UK, 1448–1454.

D. Carrillo-Zapata, J. Sharpe, A. Winfield, L. Giuggioli, and S. Hauert. 2019. Toward Controllable Morphogenesis in Large Robot Swarms. *IEEE Robotics and Automation Letters* 4, 4 (2019), 3386–3393.

K. Chatzilygeroudis, Vassiliades V. Cully, A., and J-B. Mouret. 2021. Quality-Diversity Optimization: A Novel Branch of Stochastic Optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer, 109–135.

N. Cheney, J. Bongard, V. SunSpiral, and H. Lipson. 2018. Scalable Co-optimization of Morphology and Control in Embodied Machines. *Journal of The Royal Society Interface* 15, 1 (2018), 20170937.

N. Cheney, R. MacCurdy, J. Clune, and H. Lipson. 2013. Unshackling Evolution: Evolving Soft Robots with Multiple Materials and a Powerful Generative Ecoding. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Amsterdam, The Netherlands, 11–23.

J. Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences*. Routledge, Milton Park, UK.

A. Cully, J. Clune, D. Tarapore, and J-B. Mouret. 2015. Robots that can Adapt like Animals. *Nature* 521, 7553 (2015), 503–507.

A. Cully and Y. Demiris. 2018. Improving Evolvability of Morphologies and Controllers of Developmental Soft-Bodied Robots with Novelty Search. *IEEE Transactions on Evolutionary Computation* 22, 2 (2018), 245–259.

P. Dias, M. Silva, G. Filho, P. Vargas, L. Cota, and G. Pessin. 2021. Swarm Robotics: A Perspective on the Latest Reviewed Concepts and Applications. *Sensors* 21, 6 (2021), 2062.

S. Didi and G. Nitschke. 2016. Hybridizing Novelty Search for Transfer Learning. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE Press, Athens, Greece, 2620–2628.

S. Doncieux, N. Bredeche, J-B. Mouret, and A. Eiben. 2015. Evolutionary Robotics: What, Why, and Where to. *Frontiers in Robotics and AI* 2, 4 (2015), 323–351.

A. Duarte, F. Weissing, I. Pen, and L. Keller. 2011. An Evolutionary Perspective on Self-Organized Division of Labor in Social Insects. *Annual Review of Ecology, Evolution, and Systematics* 42, 1 (2011), 91–110.

S. Farris. 2016. Insect Societies and the Social Brain. *Current Opinion in Insect Science* 15, 1 (2016), 1–8.

D. Turgut A. Wenseleers T. Ferrante, E. Guzman. 2013. GESwarm: Grammatical Evolution for the Automatic Synthesis of Collective Behaviors in Swarm Robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Amsterdam, The Netherlands, 17–24.

E. Ferrante, A. Turgut, E. Duéñez-Guzmán, M. Dorigo, and T. Wenseleers. 2015. Evolution of Self-Organized Task Specialization in Robot Swarms. *PLoS Computational Biology* 11, 8 (2015), e1004273.

B. Flannery, S. Teukolsky, and W. Vetterling. 1986. *Numerical Recipes*. Cambridge University Press, Cambridge, UK.

A. Furman, D. Nagar, and Nitschke. G. 2019. Automating Collective Robotic System Design. In *Proceedings of the 2019 Conference on Artificial Life*. IEEE Press, Xiamen, China, 1465–1472.

A. Gad. 2022. Particle Swarm Optimization Algorithm and Its Applications: A Systematic Review. *Archives of Computational Methods in Engineering* 29, 1 (2022), 2531–2561.

M. Galassi, N. Capodieci, G. Cabri, and L. Leonardi. 2016. Evolutionary Strategies for Novelty-based Online Neuroevolution in Swarm Robotics. In *2016 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, Budapest, Hungary, 002026–002032.

A. Ghasemi and S. Zahediasl. 2012. Normality Tests for Statistical Analysis: A Guide for Non-statisticians. *International Journal of Endocrinology and Metabolism* 10 (2012), 486–489.

D. Gordon, A. Guetz, M. Greene, and S. Holmes. 2011. Colony Variation in the Collective Regulation of Foraging by Harvester Ants. *Behavioral Ecology* 22, 1 (2011), 429–435.

S. Hallauer, G. Nitschke, and E. Hart. 2023. Evolving Herding Behaviour Diversity in Robot Swarms. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lisbon, Portugal, https://doi.org/10.1145/3583133.3590528.

H. Hamann. 2018. *Swarm Robotics: A Formal Approach*. https://doi.org/10.1007/978-3-319-74528-2

A. Hart, C. Anderson, and F. Ratnieks. 2002. Task Partitioning in Leafcutting Ants. *Acta ethologica* 5, 1 (2002), 1–11.

E. Hart, A. Steyven, and B. Paechter. 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.

H. Hauser. 2019. Resilient Machines through Adaptive Morphology. *Nature Machine Intelligence* 1, 1 (2019), 338–339.

J. Hewland and G. Nitschke. 2015. The Benefits of Adaptive Behavior and Morphology for Cooperation in Robot Teams. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE, Cape Town, South Africa, 1047–1054.

E. Hunt. 2021. Phenotypic Plasticity Provides a Bioinspiration Framework for Minimal Field Swarm Robotics. *Frontiers in Robotics and AI* 7, 23 (2021), doi:10.3389/frobt.2020.00023.

W. Just and M. Moses. 2017. Flexibility through Autonomous Decision-making in Robot Swarms. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE, Honolulu, USA, 1–8.

R. Kaushik, P. Desreumaux, and J-B. Mouret. 2020. Adaptive Prior Selection for Repertoire-Based Online Adaptation in Robotics. *Frontiers in Robotics and AI* 6, 1 (2020), 10.3389/frobt.2019.00151.

S. Kelly, T. Panhuis, and M. Stoehr. 2012. Phenotypic Plasticity: Molecular Mechanisms and Adaptive Significance. *Comprehensive Physiology* 2, 1 (2012), 1417–1439.

Z. Konfrst. 2004. Parallel Genetic Algorithms: Advances, Computing Trends, Applications and Perspectives. In *Proceedings of the 18th international Parallel and distributed Processing Symposium*. IEEE, Santa Fe, USA, 162–169.

S. Kriegman, N. Cheney, and J. Bongard. 2018. How Morphological Development can Guide Evolution. *Scientific Reports* 8, 1 (2018), 13934. https://doi.org/10.1038/s41598-018-31868-7

P. Levi and S. Kernbach. 2010. *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*. Springer, Berlin, Germany.

C. Mailer, G. Nitschke, and L. Raw. 2013. Evolving Gaits for Damage Control in a Hexapod Robot. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lille, France, 146–153.

N. Mathews, A. Christensen, R. O'Grady, F. Mondada, and M. Dorigo. 2017. Mergeable Nervous Systems for Robots. *Nature Communications* 8, 1 (2017), 439.

N. Mathews, A. Christensen, A. Stranieri, A. Scheidler, and M. Dorigo. 2019. Supervised Morphogenesis: Exploiting Morphological Flexibility of Self-assembling Multirobot Systems through Cooperation with Aerial Robots. *Robotics and Autonomous Systems* 112, 1 (2019), 154–167.

D. McShea. 1996. Perspective: Metazoan Complexity and Evolution: Is There a Trend? *Evolution* 50, 1 (1996), 477–492.

K. Miras and A. Eiben. 2019. Effects of Environmental Conditions on Evolved Robot Morphologies and Behavior. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Prague, Czech Republic, 125–132.

K. Miras and A. Ferrante, E. Eiben. 2020. Environmental Influences on Evolvable Robots. *PLoS ONE* 15, 5 (2020), e0233848.

K. Miras, A. Gansekoele, K. Glette, and A. Eiben. 2018a. Insights in Evolutionary Exploration of Robot Morphology Spaces. In *IEEE Symposium Series on Computational Intelligence*. IEEE, Bangalore, India, 867–874.

K. Miras, E. Haasdijk, K. Glette, and A. Eiben. 2018c. Effects of Selection Preferences on Evolved Robot Morphologies and Behaviors. In *The 2018 Conference on Artificial Life*. MIT Press, Tokyo, Japan, 224–231.

K. Miras, K Haasdijk, E. Glette, and A. Eiben. 2018b. Search Space Analysis of Evolvable Robot Morphologies. In *International Conference on the Applications of Evolutionary Computation*. Springer, Madrid, Spain, 703–718.

S. Mkhatshwa and G. Nitschke. 2023. The Impact of Morphological Diversity in Robot Swarms. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lisbon, Portugal, 422–431. https://doi.org/10.1145/3583131.3590347

F. Mondada and et al. 2013. Swarm-bot: A New Distributed Robotic Concept. *Autonomous Robots* 17, 1 (2013), 193–221.

J. Montanier, S. Carrignon, and N. Bredeche. 2016. Behavioral Specialization in Embodied Evolutionary Robotics: Why So Difficult? *Frontiers in Robotics and AI* 3 (2016). https://doi.org/10.3389/frobt.2016.00038

A. Moore, E. Brodie, and J. Wolf. 1997. Interacting Phenotypes and the Evolutionary Process: I. Direct and Indirect Genetic Effects of Social Interactions. *Evolution* 51, 1 (1997), 1352–1362.

J-B. Mouret and J. Clune. 2015. Illuminating Search Spaces by Mapping Elites. *CoRR* abs/1504.04909 (2015). arXiv:1504.04909 http://arxiv.org/abs/1504.04909

D. Nagar, A. Furman, and Nitschke. G. 2019a. The Cost of Big Brains in Groups. In *Proceedings of the 2019 Conference on Artificial Life*. MIT Press, Newcastle, United Kingdom, 404–411.

D. Nagar, A. Furman, and Nitschke. G. 2019b. The Cost of Complexity in Robot Bodies. In *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE Press, Wellington, New Zealand, 2713–2720.

G. Nitschke and S. Didi. 2017. Evolutionary Policy Transfer and Search Methods for Boosting Behavior Quality: Robocup Keep-away Case Study. *Frontiers in Robotics and AI* 4 (2017), 62.

G. Nitschke, A. Eiben, and M. Schut. 2006. Collective Specialization for Evolutionary Design of a Multi-robot System. In *Proceedings of the Second International Workshop on Swarm Robotics*. Springer, Rome, Italy, 189–205.

G. Nitschke, A. Eiben, and M. Schut. 2011. Evolving Behavioral Specialization in Robot Teams to Solve a Collective Construction Task. *Swarm and Evolutionary Computation* 2, 1 (2011), 25–38.

G. Nitschke and D. Howard. 2022. AutoFac: The Perpetual Robot Machine. *IEEE Transactions on Artificial Intelligence* 3, 1 (2022), 2–10. https://doi.org/10.1109/TAI.2021.3104789

J. Nordmoen, F. Veenstra, K. Ellefsen, and K. Glette. 2021. MAP-Elites enables Powerful Stepping Stones and Diversity for Modular Robotics. *Frontiers in Robotics and AI* 8 (2021), 639173.

T. Nygaard, C. Martin, D. Howard, J. Torresen, and K. Glette. 2021a. Environmental Adaptation of Robot Morphology and Control Through Real-World Evolution. *Evolutionary Computation* 29, 4 (2021), 441–461.

T. Nygaard, C. Martin, J. Torresen, K. Glette, and D. Howard. 2021b. Real-World Embodied AI through a Morphologically Adaptive Quadruped Robot. *Nature Machine Intelligence* 3, 5 (2021), 410–419. https://doi.org/10.1038/s42256-021-00320-3

M. Otte. 2014. An Emergent Group Mind across a Swarm of Robots: Collective Cognition and Distributed Sensing via a Shared Wireless Neural Network. *The International Journal of Robotics Research* 37, 9 (2014), 1017–1061.

S. O'Donnell, S. Bulova, S. DeLeon, P. Khodak, S. Miller, and E. Sulger. 2015. Distributed Cognition and Social Brains: Reductions in Mushroom Body Investment Accompanied the Origins of Sociality in Wasps (Hymenoptera: Vespidae). *Proceedings of the Royal Society B: Biological Sciences* 282:20150791, doi:10.1098/rspb.2015.0791 (2015), 1–6.

T. O'Shea-Wheller, E. Hunt, and T. Sasaki. 2020. Functional Heterogeneity in Superorganisms: Emerging Trends and Concepts. *Annals of the Entomological Society of America* 114, 5 (11 2020), 562–574.

I. Perez, A. Boumaza, and F. Charpillet. 2014. Comparison of Selection Methods in On-line Distributed Evolutionary Robotics. In *Proceedings of the Artificial Life Conference*. MIT Press, New York, USA, 282–289.

R. Pfeifer and J. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, Cambridge, USA.

J. Pugh and A. Martinoli. 2016. Distributed Adaptation in Multi-robot Search Using Particle Swarm Optimization. In *Proceedings of the 10th International Conference on Simulation of Adaptive Behavior*. Springer, Osaka, Japan, 393–402.

J. Pugh, P. Soros, L. Szerlip, and K. Stanley. 2015. Confronting the Challenge of Quality Diversity. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Madrid, Spain, 967–974.

K. Pugh, L. Soros, and K. Stanley. 2016. Quality Diversity: A New Frontier for Evolutionary Computation. *Frontiers in Robotics and AI* 3 (2016). https://doi.org/10.3389/frobt.2016.00040

R. Putter and G. Nitschke. 2017. Evolving Morphological Robustness for Collective Robotics. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE Press, Honolulu, USA, 1104–1111.

G. Reséndiz-Benhumea, E. Sangati, F. Sangati, S. Keshmiri, and T. Froese. 2021. Shrunken Social Brains? A Minimal Model of the Role of Social Interaction in Neural Complexity. *Front. Neurorobot.* 15:634085 (2021), doi:10.3389/fnbot.2021.634085.

M. Rubenstein, A. Cornejo, and R. Nagpal. 2013. Programmable Self-Assembly in a Thousand-Robot Swarm. *Science* 345, 1254295 (2013), 795–799.

E. Samuelsen and K. Glette. 2014. Some Distance Measures for Morphological Diversification in Generative Evolutionary Robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Vancouver, Canada, 721–728.

C. Schlichting and M. Pigliucci. 1998. *Phenotypic Evolution: A Reaction Norm Perspective*. Sinauer Associates Incorporated, Sunderland, USA.

Q. Shan and S. Mostaghim. 2021. Achieving Task Allocation in Swarm Intelligence with Bi-Objective Embodied Evolution. *Swarm Intelligence* 15, 3 (2021), 287–310. https://doi.org/10.1007/s11721-021-00198-2

F. Silva, P. Urbano, L. Correia, and A. Christensen. 2015. odNEAT: An Algorithm for Decentralised Online Evolution of Robotic Controllers. *Evolutionary Computation* 23, 3 (2015), 421–449. https://doi.org/10.1162/EVCO_a_00141

Z. Skolicki and K. De Jong. 2004. Improving Evolutionary Algorithms with Multi-representation Island Models. In *Proceedings of Parallel Problem Solving from Nature - PPSN VIII*. Springer, Birmingham, UK, 420–429.

I. Slavkov, D. Carrillo-Zapata, N. Carranza, X. Diego, F. Jansson, J. Kaandorp, S. Hauert, and J. Sharpe. 2018. Morphogenesis in Robot Swarms. *Science Robotics* 3, eaau9178 (2018), doi:10.1126/scirobotics.aau9178.

C. Spanellis, B. Stewart, and G. Nitschke. 2021. The Environment and Body-Brain Complexity. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lille, France, 138–145.

A. Steyven, E. Hart, and B. Paechter. 2017. An Investigation of Environmental Influence on the Benefits of Adaptation Mechanisms in Evolutionary Swarm Robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Berlin, Germany, 155–162.

P. Trueba, A. Prieto, F. Bellas, P. Caamaño, and R. Duro. 2013. Specialization Analysis of Embodied Evolution for Robotic Collective Tasks. *Robotics and Autonomous Systems* 61, 7 (2013), 682–693. https://doi.org/10.1016/j.robot.2012.08.005

L. Trujillo et al. 2011. Speciation in Behavioral Space for Evolutionary Robotics. *Journal of Intelligent and Robotic Systems* 64, 1 (2011), 323–351.

F. van Diggelen et al. 2022. Diversity in Swarm Robotics with Task-Independent Behavior characterization. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Boston, USA, 31–39.

M. Waibel, D. Floreano, and L. Keller. 2011. A Quantitative Test of Hamilton's Rule for the Evolution of Altruism. *PLOS Biology* 9, 5 (2011), e1000615.

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, Cape Town, South Africa, 1039–1046.

O. Weissl and A. Eiben. 2023. Morphological-Novelty in Modular Robot Evolution. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE, Mexico City, Mexico, 1066–1071.

M. West-Eberhard. 1989. Phenotypic Plasticity and the Origins of Diversity. *Annual Review of Ecology and Systematics* 20, 1 (1989), 249–278.

J. Wolf, E. Brodie, and A. Moore. 1999. Interacting Phenotypes and the Evolutionary Process. II. Selection Resulting from Social Interactions. *The American Naturalist* 153, 1 (1999), 254–266.

F. Xu and H. Wang. 2021. Soft Robotics: Morphology and Morphology-inspired Motion Strategy. *IEEE/CAA Journal of Automatica Sinica* 8, 9 (2021), 1500–1522.

E. Zardini, D. Zappetti, D. Zambrano, G. Iacca, and D. Floreano. 2021. Seeking Quality Diversity in Evolutionary Co-Design of Morphology and Control of Soft Tensegrity Modular Robots. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM. https://doi.org/10.1145/3449639.3459311