

Body and Brain Quality-Diversity in Robot Swarms



Presented by:
Sindiso Mkhathshwa

Submitted to the Department of Computer Science at the University of Cape Town in fulfilment of the academic requirements for a Master of Science by dissertation degree in Computer Science

June 30, 2023

Declaration

1. I understand what plagiarism is.
2. This dissertation titled, 'Body and Brain Quality-Diversity in Robot Swarms' is my own work.
3. The APA7 convention for citation and referencing has been followed in this dissertation. Every contribution and quotation from the work of other individuals has been appropriately attributed, cited, and referenced.

Signature:.....

S. Mkhathshwa

Date:.....29/06/2023

Acknowledgments

I would like to express my heartfelt gratitude to my supervisor, Prof. Geoff Nitschke, for his unwavering guidance and support throughout the duration of my research. His expertise, valuable insights, and encouragement have been instrumental in shaping the direction and success of this dissertation.

Abstract

Various studies have shown that diverse groups perform better, solve problems more adeptly, and are more resilient. However, in evolutionary robotics, evolving group diversity is a difficult task that frequently calls for geographic isolation, a division of labour mechanism, and a careful choice of parameters. According to recent research, decentralized Quality Diversity (QD) algorithms can generate behavioral diversity across a swarm without requiring geographical isolation or a division of labour mechanism. Despite the fact that these findings represent an essential first step in the quest to find a mechanism to evolve behavioral diversity across a swarm in physical robot tasks, little research has been done on evolving behaviour-morphology diversity across a robot swarm given cooperative tasks. To address this issue, we investigate the application of a decentralized QD algorithm (EDQD) to generate group diversity given an increasingly challenging collective behavior task in order to determine the circumstances in which it succeeds and fails. We further develop Double-Map EDQD-M, an algorithm that combines morphology characterization and behavior characterization (body-brain diversity maintenance). Results indicate that body-brain diversity maintenance yielded significantly higher behavioral and morphological diversity in evolved swarms overall, which was beneficial in the most complex task environment.

Contents

Abbreviations	xv
1 Introduction	1
1.1 Motivation	4
1.2 Research Questions	6
1.3 Contributions	7
1.4 Overview	7
2 Literature Review	9
2.1 Swarm Robotics	9
2.1.1 What is Swarm Robotics	9
2.1.2 Self-organization and Emergence	10
2.1.3 Basic Collective Behaviors	11
2.1.4 Current Applications	13
2.1.4.1 Aerial	13
2.1.4.2 Aquatic	14
2.1.4.3 Terrestrial	15
2.1.5 Future Applications	16
2.2 Adaptive Behaviour Methods	17

2.2.1	Behavior-based methods	17
2.2.2	Reinforcement Learning	18
2.2.3	Evolutionary Robotics	20
2.2.3.1	Evolving control	21
2.2.3.2	Body-brain co-evolution	23
2.2.4	Embodied Evolutionary Robotics	24
2.2.5	Evolutionary Embodied Collective Robotics	25
2.2.5.1	Evolving control	25
2.2.6	Novelty Search	26
2.2.6.1	Evolving control	27
2.2.6.2	Body-brain co-evolution	27
2.2.7	Quality Diversity Algorithms	28
2.2.7.1	Evolving control	29
2.2.7.2	Body-brain co-evolution	30
2.3	Environmental Influences on Evolving Behavior and Morphology	31
2.4	Discussion	33
3	Methodology	35
3.1	Robot Controllers	35
3.2	mEDEA	36
3.3	mEDEA-M	36
3.4	EDQD	37
3.5	EDQD-M	38

3.6	Double-Map EDQD-M	38
3.7	Summary	39
4	Experiments	40
4.1	Collective Behavior Task	40
4.2	Task Complexity	41
4.3	Experiments Setup	41
4.3.1	Collective Behavior Evaluation	42
4.3.2	Behavior Quality Evaluation	42
4.3.3	Diversity Evaluation	43
4.3.4	mEDEA: Robotic Swarm Behavior Evolution	43
4.3.5	mEDEA-M: Robotic Swarm Behavior-Morphology Co-evolution	43
4.3.6	EDQD: Robotic Swarm Behavior Bvolution with Behavioral Diversity Maintenance	44
4.3.7	EDQD-M: Robotic Swarm Behavior-Morphology Co-evolution with Behavioral Di- versity Maintenance	44
4.3.8	Double-Map EDQD-M: Robotic Swarm Behavior-Morphology Evolution with Behavior-Morphology Diversity Maintenance	44
4.4	Parameter Tuning	45
4.5	Simulator	45
4.5.1	Robotic Agents	45
4.6	Summary	47
5	Results	48
5.1	Quality Diversity (QD) of Evolved Behaviors	48

5.1.1	Average Behavioral Diversity	49
5.1.1.1	Behavioral diversity comparison in the <i>simple</i> environment	49
5.1.1.2	Behavioral diversity comparison in the <i>medium</i> environment	49
5.1.1.3	Behavioral diversity comparison in the <i>difficult</i> environment	50
5.1.1.4	Behavioral diversity comparison: <i>simple</i> vs. <i>medium</i> vs. <i>difficult</i>	50
5.1.2	Average Distance Explored	51
5.1.2.1	Distance explored comparison in the <i>simple</i> environment	52
5.1.2.2	Distance explored comparison in the <i>medium</i> environment	52
5.1.2.3	Distance explored comparison in the <i>difficult</i> environment	52
5.1.2.4	Distance explored comparison: <i>simple</i> vs. <i>medium</i> vs. <i>difficult</i>	52
5.1.3	Average Resources Collected	54
5.1.3.1	Type <i>A</i> resources	55
5.1.3.2	Type <i>B</i> resources	55
5.1.3.3	Type <i>C</i> resources	56
5.1.3.4	Type <i>D</i> resources	57
5.1.3.5	Type <i>E</i> resources	58
5.1.4	Behavior Map-archives of Best-Performing Swarms	60
5.2	Quality Diversity (QD) of Evolved Morphologies	62
5.2.1	Average Morphological Diversity	63
5.2.1.1	Morphological diversity comparison in the <i>simple</i> environment	63
5.2.1.2	Morphological diversity comparison in the <i>medium</i> environment	63
5.2.1.3	Morphological diversity comparison in the <i>difficult</i> environment	63

5.2.1.4	Morphological diversity comparison: <i>simple</i> vs. <i>medium</i> vs. <i>difficult</i> . . .	63
5.2.2	Average Range of Active Sensors	64
5.2.2.1	Sensor range comparison in the <i>simple</i> environment	65
5.2.2.2	Sensor range comparison in the <i>medium</i> environment	65
5.2.2.3	Sensor range comparison in the <i>difficult</i> environment	66
5.2.2.4	Sensor range comparison: <i>simple</i> vs. <i>medium</i> vs. <i>difficult</i>	66
5.2.3	Average Proportion of Active Sensors	66
5.2.3.1	Proportion of active sensors comparison in the <i>simple</i> environment	68
5.2.3.2	Proportion of active sensors comparison in the <i>medium</i> environment	68
5.2.3.3	Proportion of active sensors comparison in the <i>difficult</i> environment	68
5.2.3.4	Proportion of active sensors comparison: <i>simple</i> vs. <i>medium</i> vs. <i>difficult</i> . . .	68
5.2.4	Morphological Complexity	69
5.2.4.1	Morphological complexity comparison in the <i>simple</i> environment	70
5.2.4.2	Morphological complexity comparison in the <i>medium</i> environment	70
5.2.4.3	Morphological complexity comparison in the <i>difficult</i> environment	70
5.2.4.4	Morphological complexity comparison: <i>simple</i> vs. <i>medium</i> vs. <i>difficult</i> . . .	70
5.2.5	Morphology Map-archive of Best-Performing Swarms	71
5.3	Swarm Task Performance	72
5.3.1	Task Performance Comparison in the <i>Simple</i> Environment	73
5.3.2	Task Performance Comparison in the <i>Medium</i> Environment	73
5.3.3	Task Performance Comparison in the <i>Difficult</i> Environment	73
5.3.4	Task Performance Comparison: <i>Simple</i> vs. <i>Medium</i> vs. <i>Difficult</i> Environment . .	73

5.4	Quality Diversity Score of Behavior Archives	74
5.4.1	Quality Diversity Score Comparison in the <i>Simple</i> Environment	75
5.4.2	Quality Diversity Score Comparison in the <i>Medium</i> Environment	75
5.4.3	Quality Diversity Score Comparison in the <i>Difficult</i> Environment	76
5.4.4	Quality Diversity Score Comparison: <i>Simple</i> vs. <i>Medium</i> vs. <i>Difficult</i>	76
5.5	Quality Diversity Score of Morphology Archives	77
5.6	Summary	79
6	Discussion	80
6.1	Impact of Environmental Conditions	81
6.1.1	Team Task Performance	81
6.1.2	Behavioral Diversity	82
6.1.3	Morphological Diversity	83
6.2	Impact of Behavioral Diversity Maintenance on Collective Behavior Evolution	85
6.3	Impact of Behavioral Diversity Maintenance on Co-evolving Collective Behavior and Morphology	87
6.3.1	Behavioral Diversity	87
6.3.2	Morphological Diversity	88
6.3.3	Neural-Morpho Complexity	88
6.4	Impact of Body-Brain Diversity Maintenance on Co-evolving Swarm Behavior and Morphology	89
6.4.1	Behavioral Diversity	89
6.4.2	Morphological Diversity	90
6.4.3	Neural-Morpho Complexity	91

6.5 Summary	93
7 Conclusions	95
7.1 Future Work	96

List of Figures

1.1	Illustration of the differences between Objective-driven and Quality Diversity (QD)-driven evolutionary search.	3
2.1	Self-organization in social animals.	10
2.2	Aggregation behavior in a swarm.	12
2.3	Collective navigation behaviors.	12
2.4	Collective decision-making in a swarm.	13
2.5	Simulation of SwarmCity in Unity.	14
2.6	Environmental monitoring with a swarm of bouys.	15
2.7	Behavior-based control of an aerial swarm.	18
2.8	Various components of a Reinforcement Learning (RL) agent.	19
2.9	The steps of an evolutionary algorithm.	21
2.10	Swarm controller evolution with 100 robots.	22
2.11	In related work, morphology adaptation adapted the sensory configuration parameters shown above	23
2.12	Visualisation of the necessary steps of an Embodied Evolutionary Robotics (EER) implementation.	25
2.13	Multi-Dimensional Archive of Phenotypic Elites (MAP-Elites)	28
2.14	MAP-Elites searches a high-dimensional space for the highest-performing solutions within a single run.	29

3.1	Simulated robotic agent and an illustration of the EDQD method.	37
3.2	Illustration of the Double-Map EDQD-M method.	38
4.1	Overview of RoboGen framework.	46
4.2	Robotic agent components.	46
4.3	Front and back views of a complete robotic agent.	46
4.4	Task environment with robotic agents cooperating to move resources.	47
5.1	Average behavioral diversity of swarms evolved by each approach in each environment . .	50
5.2	Average distance explored by swarms evolved by each approach in each environment . . .	51
5.3	Resources collected by swarms evolved by each approach in each task environment.	54
5.4	Map-archive of robot behaviors discovered by the best-performing swarm of each approach in the <i>simple</i> environment over a single run (at the final generation), colored based on fitness.	60
5.5	Map-archive of robot behaviors discovered by the best-performing swarm of each approach in the <i>medium</i> environment over a single run (at the final generation), colored based on fitness.	61
5.6	Map-archive of robot behaviors discovered by the best-performing swarm of each approach in the <i>difficult</i> environment over a single run (at the final generation), colored based on fitness.	61
5.7	Average morphological diversity of swarms evolved by the minimal Environment- driven Evolutionary Algorithm-with Morphology adaptation (mEDEA-M), Embodied Dis- tributed Quality Diversity-with Morphology adaptation (EDQD-M), and Double-Map EDQD-M approaches in each environment	62
5.8	Average range of active sensors (mean of the population) at the end of each treatment .	64
5.9	Average (over 20 runs) range of active sensors calculated from mEDEA-M, EDQD-M, and Double-Map EDQD-M at the end of each generation (for each run) for swarms evolved in the <i>simple</i> , <i>medium</i> and <i>difficult</i> environments.	65
5.10	Average proportion of active sensors (mean of the population) at the end of each treatment	67

5.11	Average (over 20 runs) proportion of active sensors calculated from mEDEA-M, EDQD-M, and Double-Map EDQD-M at the end of each generation (for each run) for swarms evolved in the <i>simple</i> , <i>medium</i> , and <i>difficult</i> environments.	67
5.12	Average morphological complexity (mean of the population) at the end of each treatment	69
5.13	Map-archive of robot morphologies evolved by the best-performing swarm of each approach in each environment over a single run (at the final generation), colored based on fitness. .	71
5.14	Average task performance of swarms evolved by the minimal Environment-driven Evolutionary Algorithm (mEDEA), mEDEA-M, Embodied Distributed Quality Diversity (EDQD), EDQD-M, and Double-Map EDQD-M approaches in each environment.	72
5.15	QD score calculated from mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M final behavior-maps (at the end of each run) for swarms evolved in the <i>simple</i> , <i>medium</i> , and <i>difficult</i> environments.	74
5.16	Average (over 20 runs) QD score calculated from mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M behavior-maps at the end of each generation (for each run) for swarms evolved in the <i>simple</i> , <i>medium</i> , and <i>difficult</i> environments.	74
5.17	QD score calculated from mEDEA-M, EDQD-M, and Double-Map EDQD-M final morphology-maps (at the end of each run) for swarms evolved in the <i>simple</i> , <i>medium</i> , and <i>difficult</i> environments.	77
5.18	Average (over 20 runs) QD score calculated from mEDEA-M, EDQD-M, and Double-Map EDQD-M morphology-maps at the end of each generation (for each run) for swarms evolved in the <i>simple</i> , <i>medium</i> , and <i>difficult</i> environments.	77

List of Tables

3.1	An overview of each of the evolutionary methods evaluated in this thesis.	36
4.1	Parameters for robot Artificial Neural Network (ANN) controllers adapted by MAP-Elites component of EDQD, EDQD-M, Double-Map EDQD-M swarm behavior-morphology adaptation methods.	40
4.2	Experiment parameters for mEDEA, mEDEA-M, EDQD, EDQD-M and Double-Map EDQD-M methods (applied to adapt swarm behavior-morphology) and collective gathering task (evaluating swarm adaptation methods) parameters.	41
4.3	An overview of each of the experiments conducted in this thesis, and the specific objective each experiment is trying to satisfy.	42
5.1	Statistical comparisons of unique behaviors in each environment.	49
5.2	Statistical comparisons of unique behaviors in each approach across the three environments.	50
5.3	Statistical comparisons of distance explored in each environment.	51
5.4	Statistical comparisons of distance explored in each approach across the three environments.	53
5.5	Statistical comparisons of type-A resources collected in each environment.	55
5.6	Statistical comparisons of type-B resources collected in each environment.	56
5.7	Statistical comparisons of type-C resources collected in each environment.	57
5.8	Statistical comparisons of type-D resources collected in each environment.	58
5.9	Statistical comparisons of type-E resources collected in each environment.	59
5.10	Statistical comparisons of unique morphologies in each environment.	62

5.11	Statistical comparisons of average morphological diversity in each approach across the three environments.	64
5.12	Statistical comparisons of average range of active sensors in each environment.	65
5.13	Statistical comparisons of average range in each approach across the three environments.	66
5.14	Statistical comparisons of proportion of active sensors in each environment.	67
5.15	Statistical comparisons of proportion of active sensors in each approach across the three environments.	68
5.16	Statistical comparisons of morphological complexity in each environment.	69
5.17	Statistical comparisons of morphological complexity in each approach across the three environments.	70
5.18	Statistical comparisons of task performance of robot teams evolved by the mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M approaches in each task environment.	72
5.19	Statistical comparisons of task performance in each approach across the three environments.	73
5.20	Statistical comparisons of quality diversity score of robot teams evolved by the mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M approaches in each task environment.	75
5.21	Statistical comparisons of QD score in each approach across the three environments.	76
5.22	Statistical comparisons of QD score in each environment.	78
6.1	The number of individuals required to collect all the resources of each resource type in each environment.	84
6.2	Median neural-morpho complexity (over 20 runs) in each environment. Bold values indicate the lowest neural-morpho evolved in each environment.	89
6.3	Median neural-morpho complexity (over 20 runs) in each environment. Bold values indicate the lowest median complexity in each environment.	92

Abbreviations

ANN	Artificial Neural Network
BC	Behavior Characterization
CMA-ES	Covariance Matrix Adaptation Evolution Strategy
dEE	distributed Embodied Evolution
EA	Evolutionary Algorithms
EDQD	Embodied Distributed Quality Diversity
EDQD-M	Embodied Distributed Quality Diversity-with Morphology adaptation
EE	Embodied Evolution
eEE	encapsulated Embodied Evolution
EER	Embodied Evolutionary Robotics
EMOTIONAL	EMbodied Open-ended evoluTIONary ALgorithm
ER	Evolutionary Robotics
MAP-Elites	Multi-Dimensional Archive of Phenotypic Elites
MARL	Multi-Agent Reinforcement Learning
mEDEA	minimal Environment-driven Evolutionary Algorithm
mEDEA-M	minimal Environment-driven Evolutionary Algorithm-with Morphology adaptation
MGSO	Modified Glowworm Swarm Optimization
NEAT	NeuroEvolution of Augmenting Topologies
NS	Novelty Search
NSLC	Novelty Search with Local Competition
PFSM	Probabalistic Finite State Machines
PGTA	Probabilistic Gene Transfer Algorithm
QD	Quality Diversity

RL Reinforcement Learning

Chapter 1

Introduction

A swarm robotics system (Hamann, 2018) is composed of numerous autonomous robotic agents that are relatively simple in nature. These agents have the ability to interact with one another as well as their surrounding environment. However, their access to environmental information is limited or constrained in some way. This means that the individual agents must rely on local interactions and their immediate surroundings to gather and exchange information, rather than having access to a centralized control system or comprehensive environmental data. These local interactions, both among the robotic agents themselves and between the agents and their environment, give rise to the emergence of complex collective behavior (Kernbach, 2012).

The emergent property of a swarm robotics system offers numerous potential benefits, including flexibility, robustness, and scalability. Flexibility refers to the system's ability to adapt to a broad range of environments and tasks. Robustness implies that the system can adjust and remain functional even in the face of individual agent loss or significant shifts in environmental conditions, avoiding catastrophic failures. Lastly, scalability denotes the system's capability to perform effectively regardless of changes in population size, meaning that the addition or removal of agents does not lead to substantial disruptions in task performance.

These desirable traits have made swarm robotics systems highly attractive across a wide range of application domains. They find utility in tasks that require cost-effective designs (Navarro & Matía, 2012), such as agricultural foraging and mining activities, as well as in high-precision and hazardous tasks (Arnold et al., 2019) like target searching and post-disaster relief. However, the literature on real-world swarm applications remains limited due to the inherent risks associated with their deployment (Schranz et al., 2020). Consequently, a majority of current studies are conducted in physics simulations and controlled laboratory environments. While these settings provide valuable insights into the capabilities and potential of swarm robotics systems, the translation of findings into real-world applications is still relatively scarce due to the associated risks and complexities.

As task complexity increases, capturing all the possible combinations of outcomes resulting from local interactions among agents and between agents and the environment becomes impractical, making it necessary to automate the design process (Furman et al., 2019). **Evolutionary Robotics (ER)** (Bongard, 2013; Doncieux et al., 2015; Nolfi et al., 2016) has proven to be highly successful in addressing this challenge. ER has made significant contributions to the fields of evolutionary biology and engineering.

In evolutionary biology, [ER](#) provides valuable scientific tools for modeling and simulating biological evolution (Long, [2012](#)). It enables researchers to study and understand the mechanisms and dynamics of evolution in a controlled and reproducible manner. This approach allows for the exploration of various evolutionary processes and the examination of their outcomes.

In the field of engineering, [ER](#) offers an alternative approach to designing robotic systems. Unlike traditional design methods that focus on individual components, [ER](#) takes a holistic perspective, considering both the morphology (physical structure) and control aspects of the robot system. By balancing the complexity across these different parts, [ER](#) facilitates the creation of simpler and more effective robotic systems (Doncieux et al., [2015](#)).

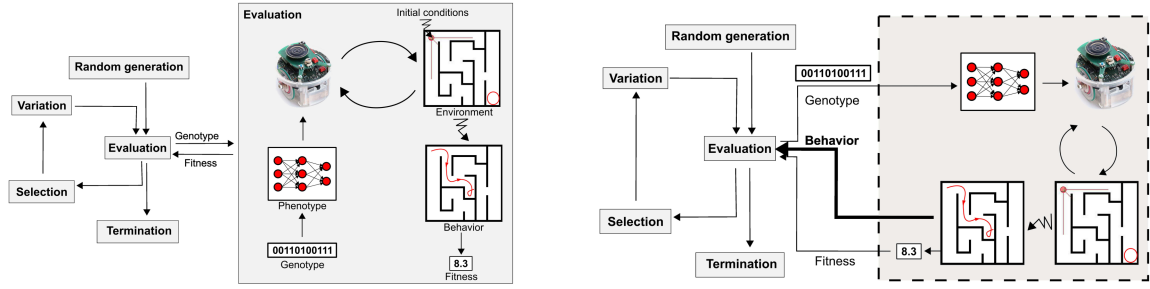
One notable advantage of using evolution in the design process is its capacity to produce novel and unexpected solutions that may not arise through conventional design approaches (Hornby et al., [2011](#)). [Evolutionary Algorithms \(EA\)](#) explore a vast design space and have the potential to discover innovative solutions that go beyond the limitations of human-designed systems. This capability is particularly valuable when dealing with complex problems that are challenging for traditional design methods.

Furthermore, unlike traditional design methods, [ER](#) offers learning and adaptation capabilities (Doncieux et al., [2015](#)). This means that [ER](#) can be utilized to design robotic systems that can learn and adapt to unforeseen or changing environments. This ability reduces the need for the designer to make assumptions about potential sources of failure and integrate them into the design beforehand.

However, it is important for the designer to possess prior knowledge of the problem domain in order to define an appropriate fitness function for the given task. The choice of fitness function plays a crucial role in the design process. An explorative fitness function drives the search towards unexplored regions of the solution space, while an exploitative fitness function directs the search towards high-performing areas of the space. Striking the right balance between exploration and exploitation is critical. Excessive exploration can result in slow convergence, while excessive exploitation can lead to premature convergence, where the search becomes trapped in local optima.

In most cases, approaches that prioritize behavioral diversity through explorative search have demonstrated superior performance compared to objective-driven approaches that solely focus on optimization (Kistemaker & Whiteson, [2011](#); Lehman & Stanley, [2010](#), [2011a](#); Lehman & Stanley, [2008](#)). A new class of algorithms known as [QD](#) algorithms (Pugh et al., [2016b](#)) addresses this trade-off by reformulating the optimization process to consider both behavioral diversity and quality. These algorithms aim to find a balance between exploring new behaviors and achieving high-quality solutions, offering a more comprehensive approach to evolutionary optimization in robotics (see figure [1.1](#)).

Traditionally, the focus of optimization in [ER](#) has been on the controller, neglecting the potential benefits of morphology optimization. However, recent research (Bongard, [2011](#)) has emphasized that adapting a robot's morphology while it is actively engaged in behavior can reveal novel motor-sensor relationships. This adaptive process has been shown to enhance the evolution of robust behaviors, particularly in challenging and unfamiliar environments. These findings have sparked further investigations into the advantages of co-evolving the body and brain of robots (Birattari et al., [2019](#); Kriegman et al., [2018](#); Lipson et al., [2016](#); Nygaard et al., [2021b](#)). Notably, previous work (Kriegman et al., [2018](#)) has demonstrated that the morphology of a robot significantly influences the range of behaviors that its



(a) Objective-driven evolutionary search typically focuses solely on the fitness value of individuals during the selection process. (b) QD-driven evolutionary search takes both the fitness value and some aspects of the individual behavior into account during the selection process.

Figure 1.1: Illustration of the differences between Objective driven and QD driven evolutionary search (Doncieux et al., 2015).

controllers can exhibit. In other words, the morphology of a robot plays a critical role in determining the types and complexity of behaviors it can perform. These insights have motivated further exploration of the interplay between morphology and controller, providing valuable insights into how optimizing both aspects can enable robots to exhibit specific behaviors and effectively adapt to diverse environmental conditions.

By combining the adaptive and learning capabilities of ER with the consideration of body-brain diversity, it is possible to design swarm robotic systems that not only adapt to changing environments but also possess a wide range of behaviors and morphologies, enhancing their overall problem-solving capacity. The objective of this thesis is to investigate the benefits of such approaches within the context of cooperative robot teams and examine how environmental conditions impact their efficiency and diversity. By examining the relationship between the environment, morphology, and controller, this research aims to contribute to our understanding of the interdependencies among these factors and their implications for the performance and adaptability of swarm robotic systems.

1.1 Motivation

Social insect societies represent significant examples of self-organization and decentralized control in the biological realm (O’Shea-Wheller et al., 2020). Within these societies, complex interactions between the behavior and morphology of evolving organisms, as well as their surrounding environment, have given rise to a wide range of complex and diverse social structures (Duarte et al., 2011; Hart et al., 2002). Likewise, in artificial social systems like swarm robotics, different types of collective behavior emerge through the interconnected dynamics among a robot’s morphology (its sensory-motor configuration), behavior (the output of its controller), and environment (the task at hand) (Bredeche et al., 2018). One viewpoint suggests that an agent’s morphological and behavioral complexity should align with the complexity of its environment (Pfeifer & Bongard, 2006). However, considering the diverse instances observed in biology (McShea, 1996), it remains an open question whether higher levels of task complexity necessitate agents with more complex behaviors and morphologies (Cheney et al., 2013; Nygaard et al., 2021a; Xu & Wang, 2021).

Research focused on the evolution of artificial morphology-behavior (body-brain) couplings using simulated (Cheney et al., 2017; Kriegman et al., 2018) and physical (Nygaard et al., 2021a; Xu & Wang, 2021) ER platforms has garnered considerable attention (Doncieux et al., 2015). However, the exploration of how body-brain adaptation influences collective (swarm) robotic systems has been relatively limited (Buason et al., 2005; Furman et al., 2019; Hewland & Nitschke, 2015), primarily due to the challenge of effectively discerning the connections between genotype (body-brain encodings) and phenotype plasticity (body-brain couplings) in emerging collective behaviors (Moore et al., 1997). In this context, *phenotypic plasticity* pertains to a single genotype generating multiple morphology-behavior couplings in response to environmental conditions (Kelly et al., 2012; Schlichting & Pigliucci, 1998; West-Eberhard, 1989; Wolf et al., 1999).

As an illustration, in artificial swarms (Just & Moses, 2017), local cues regarding resource distributions modify the foraging parameters of agents. Likewise, in ant colonies, environmental heterogeneity leads to diverse foraging behaviors (Beverly et al., 2009; Gordon et al., 2011). Phenotypic plasticity has been identified as a crucial adaptive element that is currently lacking in enabling autonomous robots to operate effectively in unconstrained real-world environments (Hauser, 2019). It has been explored through the lens of the necessary genotypic-phenotypic interactions and environmental conditions for the emergence of behavioral diversity (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) in evolutionary swarm robotics (Dias et al., 2021; Doncieux et al., 2015).

Despite prior research (Bongard, 2011; Kriegman et al., 2018) in evolutionary robotics highlighting the significance of phenotypic plasticity in driving morphological adaptation to enhance the resilience of adapted behaviors, there is a scarcity of studies examining the influence of emergent morphological diversity on evolving swarm-robotic behavior (Hunt, 2021). Moreover, although previous research (Auerbach & Bongard, 2014; Miras & Eiben, 2019; Miras et al., 2020; Spanellis et al., 2021) in evolutionary robotics has examined the influence of the environment on body-brain co-evolution, there is a limited body of work (Furman et al., 2019; Nagar et al., 2019a, 2019b) investigating the environmental impact on body-brain evolution specifically in the context of swarm robotics. The existing studies on phenotypic plasticity in evolutionary swarm robotics can be divided into two categories. The first

category focuses on fixed robot morphologies, where only the controllers are evolved, while the second category involves the co-evolution of both the controllers and the morphologies of each robot (Doncieux et al., 2015). Within the studies using fixed morphologies, there are few instances that demonstrate the potential for diverse environments to yield diverse behaviors (Ferrante et al., 2013; Ferrante et al., 2015).

On the other hand, research on evolutionary swarm robotics incorporating evolvable coupled controllers and morphologies remains scarce and confined in its scope (Dias et al., 2021), with a noticeable absence of studies exploring the effects of environment complexity (task complexity) on phenotypic plasticity. Conventionally, researchers define a specific environment and task, and then evolve robots specifically tailored to operate within that particular combination. For instance, in the majority of swarm-robotics studies where the task environment is considered an experimental parameter, behavioral diversity is evolved using pre-existing morphologies that are predefined to address diverse collective behavior tasks (Brutschy et al., 2012; Ferrante et al., 2015; Nitschke et al., 2011; Steyven et al., 2017; van Diggelen & et al., 2022).

Another motivation behind this study is the recognition that morphological diversity plays a significant role in enhancing the problem-solving capabilities of swarm behaviors as observed in social insect colonies (Fjerdingstad & Crozier, 2006). For instance, in social insect colonies, different workers with varying body shapes and sizes are more adept at solving specific tasks, leading to an overall enhancement in the robustness of collective behavior (Jandt & Dornhaus, 2014). Building on this observation, our hypothesis posits that incorporating morphological diversity within robotic swarms will yield similar benefits, contributing to improved problem-solving efficacy and the overall performance of the swarm.

1.2 Research Questions

In order to address the challenges mentioned in section 1.1, our research builds upon previous work (Hart et al., 2018) that utilizes decentralized QD methods to evolve functional diversity, specifically behavioral specialization, within simulated swarms, without the need for geographical isolation (Montanier et al., 2016) or division of labor mechanisms (Haasdijk et al., 2014).

Specifically, our research explores the application of three QD methods to evolve collective behavior in increasingly complex cooperative tasks. We first evaluate behavior evolution with behavioral diversity maintenance (section 3.4), secondly, behavior and morphology evolution with behavioral diversity maintenance (section 3.5), and lastly, behavior and morphology evolution with behavior-morphology diversity maintenance (section 3.6). The aim is to identify the most suitable swarm controller design methods for specific task environments and gain insights into how environmental factors impact the evolution of behavioral and morphological diversity, as well as the resulting benefits (task performance).

By exploring the potential of the three QD methods to evolve diverse body-brain couplings and cooperative behaviors in robotic swarms, this research aims to contribute to a better understanding of how to effectively employ QD maintenance approaches in complex and cooperative swarm tasks. The progression of this thesis is guided by the following research questions.

1. **Primary research question:** Do environmental conditions impact the efficiency (that is, task performance) and diversity (that is, behavioral and morphological diversity) of body-brain evolved cooperative robot teams?
2. **Secondary research questions:**
 - 2.1. Is behavioral diversity maintenance beneficial for evolving collective behavior across increasingly complex task environments?
 - 2.2. Is behavioral diversity maintenance beneficial for co-evolving collective behavior and morphology across increasingly complex task environments?
 - 2.3. Is behavioral-morphological diversity maintenance beneficial for co-evolving collective behavior and morphology across increasingly complex task environments?

The QD team task performance (section 4.3.1) and diversity (section 4.3.3) will be evaluated with respect to mEDEA (section 3.2) and mEDEA-M (section 3.3) as benchmarks over an increasingly complex collective gathering task (section 4.1). In this task, robots will be tasked with moving different block types randomly scattered in the environment to a designated gathering zone. Team task performance will be measured as the portion of blocks gathered in a given run. Task complexity (section 4.2) will denote the degree of cooperation required to optimally solve the task.

1.3 Contributions

The primary contribution of this study is empirical evidence indicating the effectiveness of body-brain diversity maintenance search approaches in complex tasks defined by large feature spaces in robotic swarms. While previous related work (Nordmoen et al., 2021; Pugh et al., 2016a; Pugh et al., 2016b; Zardini et al., 2021) has shown the benefits of these approaches for single agents, this thesis extends those findings to collective behavior in robotic swarms.

- This work conducts empirical experiments (Chapter 4, table 4.3) using a collective gathering task of varying complexity (section 4.1). Through these experiments, our work investigates the influence of environmental complexity on the emergence of neural-morpho complexity (section 5.2.4), as well as behavioral (section 5.1) and morphological (section 5.2) diversity in robotic swarms. By examining how different levels of environmental complexity impact these aspects, this thesis aims to shed light on the relationship between environmental conditions and the observed complexity and diversity within the evolved behavior-morphology couplings in robotic swarms.
- The design and development of an extension for the RoboGen framework (Auerbach et al., 2018), allowing for the embodied evolution of robotic swarms. The existing RoboGen framework already encompasses tools for generating 3D-printable design files for robot body parts and compiling neural-network controllers that can be executed on an Arduino micro-controller board. Previous research (Jelisavcic et al., 2017) has successfully validated the functionality of this modular robot system using physical hardware. Hence, this thesis serves as evidence for the framework’s capability to be expanded and adapted.

1.4 Overview

The rest of this thesis is divided into several Chapters as follows, Literature Review (Chapter 2), Methodology (Chapter 3), Experiments (Chapter 4), Results (Chapter 5), Discussion (Chapter 6), and Conclusions (Chapter 7). The following is an overview of the Chapters.

Chapter 2 is divided into three main sections. The first section provides a background on simulated and real swarm robotics, offering a comprehensive understanding of the field. The second and third sections review relevant literature that pertains to the main questions addressed in this thesis. These sections critically examine existing studies and research contributions, highlighting their significance and relevance to the research objectives.

Chapter 3 outlines the methodology employed to address the research questions presented in this thesis. It describes the various evolutionary methods utilized in conducting the experiments. The chapter provides a clear and detailed explanation of each of the five evolutionary methods evaluated in this thesis.

In Chapter 4, the experiments conducted to investigate the research questions are presented. The Chapter provides a detailed explanation of the experimental setup, including the collective behavior task, how team task performance and behavior diversity and quality were measured, the swarm robotics simulator

used, and the environmental conditions evaluated.

Chapter 5 presents the main findings of this thesis. Specifically, the Chapter presents the results of evaluating the five evolutionary methods in terms of the efficiency and diversity of the cooperative robot teams under different environmental conditions.

Chapter 6 critically discusses the findings presented in Chapter 5 and examines their significance in the context of the research questions. It explores the implications of the results and compares them with previous studies discussed in the literature review.

Lastly, chapter 7 draws conclusions based on the discussion of the results and their implications. It summarizes the key findings, evaluates the research objectives, and answers the research questions formulated in this thesis. This Chapter also discusses the broader implications of the study and potential avenues for future research vis-a-vis cooperative robot teams and their interaction with environmental conditions.

Chapter 2

Literature Review

Section 1 covers a background on simulated and real swarm robotics and the benefits and applications of swarm robotics. More of an emphasis will be placed on simulated swarm robotic systems since this will be the focus of this thesis. Whereas, section 2 will focus on methods to adapt robot (agent) behaviour (controllers) in simulated and physical swarm and collective robotic systems with more of a focus on simulated rather than physical systems. Lastly, section 3 reviews previous work that investigates environmental influences on evolving morphology and behavior.

2.1 Swarm Robotics

This section presents a high-level overview of the swarm robotics field. The aim is to highlight the many potential benefits offered by robot swarms. Section 2.1 gives a concise definition of swarm robotics. Section 2.2 gives a brief introduction to the concepts of self-organization and emergence. Section 2.3 presents basic collective robot swarm behaviors. Lastly, section 2.4 presents current and future applications of robot swarms.

2.1.1 What is Swarm Robotics

Swarm robotics draws inspiration from the self-organized behaviors demonstrated by social animal groups in nature (see figure 2.1), such as an ant colony (Jackson & Ratnieks, 2006), a swarm of bees (Narumi et al., 2018), a flock of birds (Wallraff, 2010) and a school of fish (Eberhart et al., 2001). Şahin (2005) defines swarm robotics as “the study of how a large number of relatively simple physically embodied agents can be designed such that a desired collective behavior emerges from the local interactions among agents and between the agents and the environment.” This definition is adopted by most predominant work in the literature (Brambilla et al., 2013; Hamann, 2018).

However, Arnold et al. (2019) criticize the above definition, citing that it is ambiguous with regards to the swarm size. That is, what is the appropriate size of a swarm? Indeed this is a valid question to ask. From the definition, it is clear that it has to be many robots, but what exactly is the threshold? In



(a) An ant colony (Wild, 2007).



(b) A swarm of bees (Swenty, 2012).



(c) A flock of birds (Hodan, n.d.).



(d) A school of fish (Kratochvil, n.d.).

Figure 2.1: Self-organization in social animals.

this regard, Beni (2004) posits that the swarm size should not be so large that it has to be treated as a statistical problem, and not so small that it has to be treated as a few-body problem.

Arnold et al. (2019) also argue that the swarm agents do not necessarily have to be simple. Citing that modern swarms may incorporate sophisticated, highly capable agents as part of a heterogeneous team. In fact, we have already seen such heterogeneous swarms in various applications (Dorigo et al., 2013; Ducatelle et al., 2010; Kengyel et al., 2015; Prorok et al., 2016). So, while this is a valid claim, the authors miss the point that the agents are relatively *simple* when compared to the collective behavior that emerges from local interactions among agents and between the agents and the environment. That is, individual agents are inefficient relative to the considered task (Hamann, 2018).

2.1.2 Self-organization and Emergence

A robot swarm has self-organization capabilities. This means that the system can be stimulated by external influences to set in motion specific internal system mechanisms that result in the origin of distinguished structures (Kernbach, 2012). The said structures manifest at two distinct levels; micro and macro. Self-organization is made possible by multiple interactions between system elements (that

is, agents), feedback (both positive and negative), and the balance between exploration and exploitation (Hamann, 2018).

Self-organization in swarms gives rise to emergent behavior. This refers to the ability of a swarm to demonstrate complex collective behavior at the swarm level from only local interactions among agents and between the agents and the environment. The grand challenge in swarm robotics is designing control mechanisms for the agents such that the desired behavior emerges at the swarm level from local interactions among agents and between the agents and the environment. The emergence of a robot swarm offers numerous benefits, including robustness, scalability, and flexibility. Interestingly, these characteristics align with those found in social animal groups in the natural world.

Robustness: This indicates that the system possesses adaptability, meaning it can withstand the loss of individuals or significant changes in the environment without experiencing catastrophic failure. Robustness is achieved through decentralized control and distributed communication and sensing. Decentralized control means that each agent follows its own control algorithm, so if one agent fails, the rest of the system remains unaffected. Distributed communication relies on local communication, as does distributed sensing. This is feasible due to the agents' limited sensing and communication capabilities.

Scalability: System scalability refers to its ability to maintain good performance despite changes in population size. Adding or removing agents does not lead to significant fluctuations in task performance. Scalability is achieved through decentralized control, distributed communication, and sensing. By avoiding a centralized communication or sensing mechanism, the system mitigates the risk of bottlenecks in sensing or communication when the population size increases. Consequently, this ensures that the system remains unaffected by changes in size, making it size-proof.

Flexibility: This implies that the system has the capacity to adapt to various environments and tasks. It achieves flexibility by employing simple agent behaviors and incorporating redundancy.


2.1.3 Basic Collective Behaviors

Basic collective behavior tasks can be broadly categorized into four main groups: spatially organizing behaviors, collective navigation behaviors, collective decision-making behaviors, and miscellaneous collective behaviors.

Spatially organizing behaviors involve the movement and arrangement of robots or objects within the environment. This category includes various behaviors such as aggregation (see figure 2.2), pattern formation, chain formation, self-assembly and morphogenesis, as well as object clustering and assembling.

Collective navigation behaviors involve the coordination and organization of a swarm of robots to move together within the environment (see figure 2.3). This category includes behaviors such as collective exploration, coordinated motion, and collective transport.


Collective decision-making behaviors aim to facilitate consensus among a swarm of robots when making



(a) Starting position.

(b) Final position.

Figure 2.2: Aggregation behavior in a swarm (Misir & Gökrem, 2021).



(a) Dynamic chain formation (Ducatelle et al., 2011).

(b) Self-organized flocking (Turgut et al., 2008).

Figure 2.3: Collective navigation behaviors.

decisions on a particular issue. This category includes two main behaviors: consensus achievement (see figure 2.4) and task allocation.

The final category encompasses collective behaviors that do not fit into the previous three categories. This category includes various behaviors such as human-swarm interaction, group size regulation, collective fault detection, self-healing, and self-reproduction.

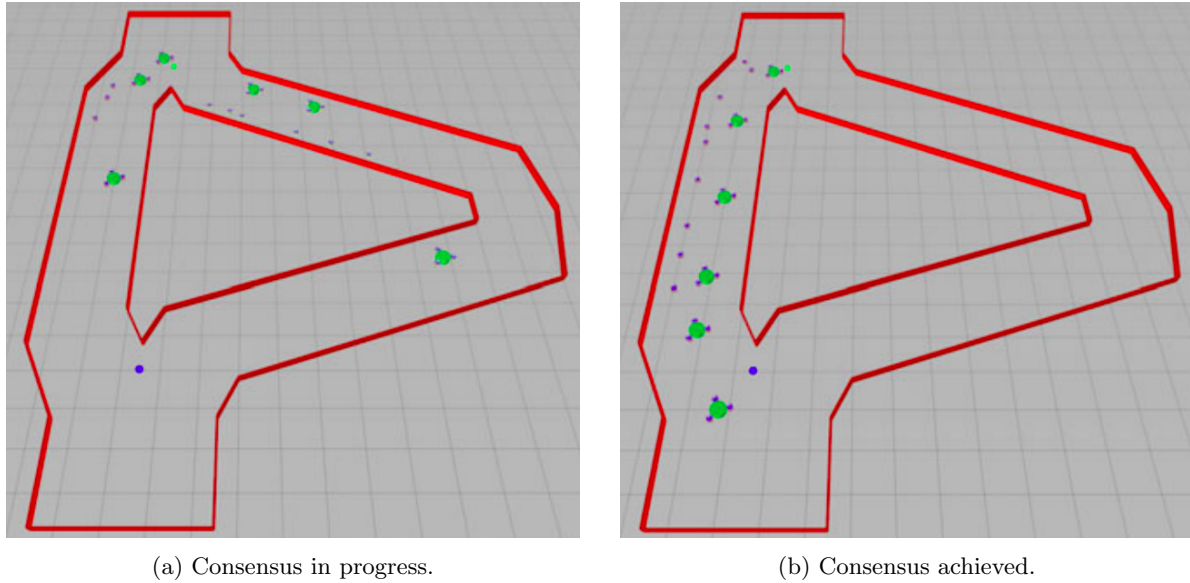


Figure 2.4: Collective decision-making in a swarm (Montes de Oca et al., 2011).

2.1.4 Current Applications

By skillfully combining various basic collective behaviors, robot swarms can be employed in a wide range of application domains, including those where optimal solutions are currently lacking (Brambilla et al., 2013). These applications can be broadly categorized based on the environment in which the swarms operate, namely, aerial, aquatic, and terrestrial. Each of these environments offers unique challenges and opportunities for swarm robotics applications.

2.1.4.1 Aerial

In the work by Camci et al. (2018), a particle swarm optimization-sliding mode control algorithm is employed to optimize type-2 fuzzy neural networks (T2-FNNs). These networks are specifically designed for controlling a swarm of quad-copters engaged in farm inspection and mapping tasks. The experimental results demonstrate that these controllers can significantly improve performance compared to conventional proportional-derivative controllers, with a reduction in trajectory tracking integral squared error of up to 26%. This highlights the effectiveness of the proposed approach in enhancing the control capabilities of the quad-copter swarm for agricultural applications.

In the study conducted by Garcia-Aunon et al. (2019), a behavior-based (section 2.2.1) surveillance algorithm is developed for aerial swarms responsible for monitoring traffic in a simulated city environment known as SwarmCity (refer to figure 2.5). The algorithm incorporates six distinct behaviors, each defined by twenty-three parameters that are fine-tuned using a genetic algorithm. The optimized algorithm's performance is subsequently evaluated within the SwarmCity simulation. The results show that the algorithm achieves an efficiency rate of 25%. However, it is worth noting that the efficiency decreases significantly when the number of cars in the simulation is increased. This observation suggests that the algorithm's performance may be impacted by scalability issues when dealing with higher traffic volumes. Further investigation and potential enhancements may be required to address this limitation and improve



Figure 2.5: Simulation of SwarmCity in Unity (Garcia-Aunon et al., 2019).

its efficiency in scenarios with larger numbers of vehicles.

2.1.4.2 Aquatic

In the study conducted by Duarte et al. (2016), controllers for a robot swarm are evolved (section 2.2.3) to perform multiple tasks, including homing, dispersion, clustering, and monitoring. The evolved controllers are subsequently evaluated using a real robotics system comprised of aquatic surface robots. The authors also demonstrate the practical application of the swarm by combining the evolved swarm behaviors to accomplish a comprehensive environmental monitoring task. By carefully combining the different evolved swarm behaviors, the swarm of robots is able to effectively carry out the environmental monitoring task. This showcases the versatility and potential of the evolved controllers in enabling robot swarms to perform complex tasks in real-world scenarios.

In the study conducted by Zoss et al. (2018), behavior-based (section 2.2.1) cooperative algorithms are developed to enable a real swarm of 50 buoys to perform various collective behaviors, including flocking, navigation, and area coverage (refer to figure 2.6). These behaviors are crucial for the swarm to effectively operate and coordinate its actions in a coordinated manner. The developed algorithms not only allow the swarm to exhibit core swarm properties such as scalability, flexibility, and robustness but also demonstrate their applicability in monitoring coastal and inland water environments. By combining the flocking behavior, which promotes cohesion and alignment within the swarm, with navigation and area coverage behaviors, the swarm is capable of efficiently monitoring and exploring these environments. The success of the developed algorithms showcases the significance of collective behaviors and their practical application in real-world scenarios, particularly in environmental monitoring.

Finally, in their work, Gupta and Bayal (2020) employ a [Modified Glowworm Swarm Optimization \(MGSO\)](#) algorithm to optimize controllers for a swarm of robots assigned with the task of locating the source of oil spills in marine environments. The [MGSO](#) algorithm is utilized to enhance the performance of the swarm and improve its efficiency in carrying out this specific task. The results obtained from their study indicate that reducing the number of iterations and swarm size leads to higher performance in terms of the swarm's ability to locate the source of oil spills. This finding suggests that optimizing the parameters of the [MGSO](#) algorithm, such as the number of iterations and swarm size, can significantly



Figure 2.6: Top panel: a swarm of bouys collectively operating. Bottom left panel: a swarm of bouys during transportation. Bottom right: a swarm of bouys just before deployment (Zoss et al., 2018).

impact the performance and effectiveness of the swarm in carrying out the desired objective. By leveraging the [MGSO](#) algorithm and optimizing the controllers for the swarm of robots, Gupta and Bayal (2020) demonstrate the potential of swarm intelligence techniques in addressing complex environmental challenges, such as oil spill detection and localization in marine environments. Their findings highlight the importance of parameter optimization and provide insights into improving the performance of swarm-based systems in real-world applications.

2.1.4.3 Terrestrial

In their research, Ball et al. (2015) focus on the development of small, lightweight, and cost-effective behavior-based (section 2.2.1) cooperative robots. These robots are specifically designed to address the issue of resistant weeds in large-scale crop fields while minimizing the environmental impact. By employing mechanical mechanisms, the swarm of robots is capable of effectively dealing with weeds that are resistant to conventional methods of control. The potential applications of this swarm of robots extend beyond weed control. They can also be utilized for tasks such as mowing and slashing vegetation, planting seeds, irrigation management, and even harvesting crops. The versatility of these robots allows for their integration into various agricultural operations, offering potential benefits in terms of efficiency, precision, and environmental sustainability.

Gebhardt et al. (2018) employ policy search (section 2.2.2) techniques to train a swarm of robots in assembly tasks. The assembly process is divided into two components: creating a high-level plan for the assembly and acquiring a low-level policy for object movement. The findings demonstrate the system’s ability to successfully complete assembly tasks involving objects with different shapes, even when encountering objects that have not been seen before.

In their study, Farrugia and Fabri (2018) propose behavior-based (section 2.2.1) cooperative transportation algorithms, including pushing, caging, and grasping, for a swarm of LEGO robots. The objective is to enable the swarm to work together to transport large objects to a specific location. The authors evaluate the efficiency of these algorithms by conducting experiments with the LEGO robot swarm. The results demonstrate that the caging and grasping algorithms are capable of generating accurate collective object transportation behavior across various scenarios, indicating their reliability and effectiveness in facilitating cooperative transport tasks.

Lastly, Hiraga et al. (2020) focus on evolving controllers (section 2.2.3) for a swarm of robots with the objective of collectively distinguishing between food and poison objects. Additionally, the robots are required to cooperate in order to transport food objects to a designated nest. Accordingly, the controllers are optimized to effectively handle the tasks of object recognition, discrimination, and cooperative transportation. The study demonstrates the successful evolution of controllers that enable the swarm to collectively distinguish between food and poison objects and work together to move food objects to the desired nest location.

2.1.5 Future Applications

The existing literature on real-world applications of robot swarms is relatively limited, with most applications being conducted in controlled laboratory settings. Schranz et al. (2020) highlight several barriers that hinder the widespread adoption of swarm robotic systems in real-world scenarios. Firstly, there is the challenge of accurately predicting and quantifying the effectiveness of swarm behavior that emerges from local interactions among the robots. This difficulty arises due to the complex and nonlinear nature of swarm dynamics. Secondly, the current communication architectures used in swarm robotics often fail to meet the requirements for efficient swarm communication. Innovations in communication technologies are needed to address this limitation. Lastly, while simulations are commonly used for testing swarm algorithms, relying solely on simulations can be risky as it may lead to a significant gap between simulated and real-world performance, known as the "reality gap" (Bredeche et al., 2018).

Despite these challenges, there are envisioned future applications of swarm robotics in various domains. Dorigo et al. (2020) propose potential applications in precision medicine, precision agriculture, infrastructure inspection and maintenance, civil construction, and space missions. These domains could benefit from the collective capabilities of robot swarms in terms of efficiency, scalability, adaptability, and robustness. However, further research and development are necessary to overcome the current barriers and realize the full potential of swarm robotic systems in real-world applications.

2.2 Adaptive Behaviour Methods

Swarm robotics systems are characterized by decentralization, meaning there is no central control and each individual agent is responsible for its own decision-making. This aspect presents a significant challenge in system design due to the complex interactions among agents and between agents and the environment. The main objective in swarm robotics is to design controls for individual agents in a way that desired collective behaviors emerge at the swarm level through local interactions. Two broad design approaches have been adopted to tackle this challenge: behavior-based methods and automatic methods.

2.2.1 Behavior-based methods

Behavior-based methods require the designer to specify control of each individual robotic agent such that when all the agents interact amongst themselves and with the environment the desired collective behavior emerges at the swarm level. Hand-crafting control for each agent requires the designer to have enough *a priori* knowledge of the task domain as this process typically requires one to describe the desired behavior in the form of a mathematical model. Nevertheless, models for commonly studied collective behaviours such as aggregation (Misir & Gökrem, 2021), and flocking (Fine & Shell, 2013), are readily available.

In this category, two commonly utilized approaches are Probabilistic Finite State Machines (PFSM) and potential-field based approaches. Potential-field based approaches employ the principles of physics, providing a well-grounded framework for analysis. However, PFSM is the more widely adopted approach. For instance, Garcia-Aunon and Cruz (2018) implemented PFSM to control an aerial robotic swarm engaged in a search task within a rectangular environment. The authors defined six distinct behaviors that govern the agents' movements (refer to figure 2.7). They also introduced a virtual pheromone model that facilitates transitions between behavior-states. Building upon this algorithm, Garcia-Aunon et al. (2019) extended it for surveillance purposes in a traffic monitoring task. They modified the search behavior to ensure agents persistently search within a confined area, while keeping the remaining behaviors unchanged.

Another example is the work of Zoss et al. (2018), where a PFSM is employed to control a swarm of buoys responsible for collectively monitoring large marine environments with arbitrary shapes. The authors establish a set of three collective behaviors: flocking, collective navigation, and area coverage, to govern the swarm's movement. Notably, the area coverage control algorithm is based on a potential-field approach specifically designed for area coverage (Howard et al., 2002). Through field experiments, the study demonstrates that these three collective behaviors effectively enable the deployment, control, and positioning of a large number of aquatic surface vehicles for efficient monitoring of diverse marine environments.

Behavior-based methods inherently involve an iterative process where the designer fine-tunes the model parameters of individual behaviors until the desired collective behavior emerges. Manual parameter tuning can be time-consuming, especially for complex models. For example, Garcia-Aunon and Cruz (2018) and Garcia-Aunon et al. (2019) utilized a genetic algorithm to automate the tuning of twenty-three parameters in a behavior-based algorithm consisting of six behaviors. Moreover, creating mathematical

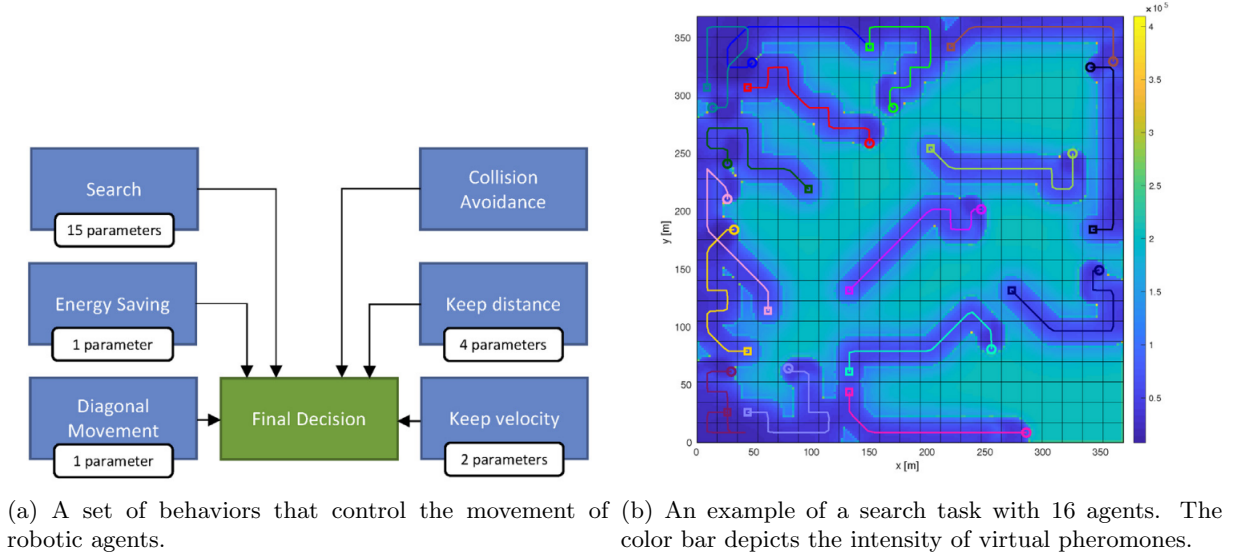


Figure 2.7: Behavior-based control of an aerial swarm (Garcia-Aunon & Cruz, 2018).

models to describe individual behaviors in unknown, complex, or dynamic problem domains becomes increasingly challenging. To address this limitation, automatic design methods have been adopted, namely, **RL** (Sutton & Barto, 2018) and **ER** (Doncieux et al., 2015). These approaches aim to overcome the need for manual parameter tuning and facilitate the design process in complex problem domains.

2.2.2 Reinforcement Learning

RL involves the process of learning how to map environmental states to actions in order to optimize a numerical reward. **RL** consists of two main components: trial-and-error search and delayed reward. The trial-and-error search component stems from the fact that the **RL** learner has no prior knowledge of which actions to take and must explore and learn through a process of trial and error. By interacting with the environment, the learner discovers which actions lead to favorable outcomes and gradually improves its decision-making abilities. Delayed reward refers to the concept that the consequences of the learner's actions extend beyond immediate rewards. The actions taken by the learner not only impact immediate rewards but also carry over to the next state, influencing subsequent rewards. This delayed feedback allows the learner to consider long-term consequences and make decisions that maximize cumulative rewards over time.

A typical **RL** agent consists of four essential elements: a *policy*, a *reward function*, a *value function*, and sometimes a *model of the environment*. The *policy* is responsible for mapping observed environmental states to the corresponding actions to be taken in those states. This mapping can take various forms, such as a simple function, a lookup table, or even computationally intensive processes like search algorithms. The *reward function* assigns a numerical value to each state-action combination, representing the desirability of that state. It captures the immediate feedback and essential aspects of the task that the learning agent aims to maximize. The *value function* provides insight into the long-term desirability or utility of being in a particular state. Additionally, the **RL** agent may incorporate a *model of the environment*, which simulates the behavior of the actual environment. Given a state-action pair, the model predicts the resultant next state and the associated reward. This model is often used for planning

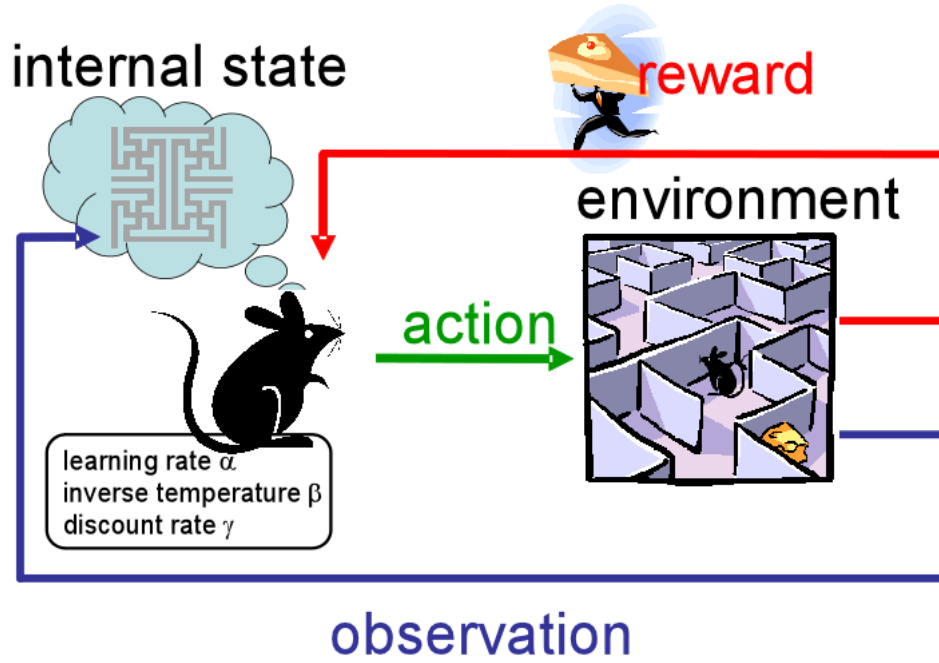


Figure 2.8: Various components of a [RL](#) agent.

purposes, enabling the agent to anticipate and evaluate potential future scenarios and decide on a course of action accordingly.

[RL](#) has been widely used in mainstream robotics to optimize the control of robotic agents. It has proven to be particularly effective in single agent domains where it is not necessary to model or predict the behavior of other actors in the environment. [RL](#) methods have a strong theoretical foundation with elaborate convergence proofs.

However, when it comes to multi-agent situations, traditional [RL](#) algorithms face challenges. The inherently non-stationary nature of the environment hampers Q-learning (Watkins & Dayan, 1992), and the variance worsens as the number of agents increases in Policy Gradient methods (Silver et al., 2014). Agent interactions can be categorized into three types, namely, cooperative, competitive, or a hybrid of both.

There are two main approaches to learning in cooperative multi-agent scenarios: *centralized Multi-Agent Reinforcement Learning (MARL)* and independent learners approach. The centralized [MARL](#) approach treats the learning problem as a single-agent [RL](#) problem, where the observations of individual agents are concatenated, and the action space becomes combinatorial. However, in practice, this approach often struggles to solve even relatively simple cooperative [MARL](#) problems (Buşoniu et al., 2010; Zhang et al., 2021).

The independent learners approach trains each agent independently to optimize for the benefit of the team. In this scenario, each agent faces a non-stationary learning problem, as the environment dynamics change when teammates adapt their behavior based on their own learning experiences. This creates a challenging and dynamic learning environment for each agent (Laurent et al., 2011).

The challenges posed by multi-agent environments have been the subject of extensive research (Busoniu et al., 2008; Buşoniu et al., 2010; Lanctot et al., 2017; Zhang et al., 2021). For example, Lowe et al. (2017) propose an extension of actor-critic approaches that addresses the non-stationarity problem by incorporating information from other agents. However, this method faces scalability issues, as the input space grows linearly with the number of agents (N). This leads to high state and observation dimensionality, and changes in the number of agents or observed neighbors result in a varying information set size.

Taking inspiration from this observation, Hüttenrauch et al. (2019) propose a novel state representation for deep reinforcement learning in multi-agent systems. This representation utilizes mean embeddings of distributions to treat state information from neighboring agents as samples. As a result, the representation is permutation invariant, meaning it is insensitive to the order of agents or the swarm size, thereby overcoming the scalability limitations faced by other methods.

In the independent learners approach, where agents can only observe a portion of the environment, there is a possibility of receiving erroneous reward signals that stem from the actions of unobserved teammates. To mitigate this limitation, Sunehag et al. (2017) propose a novel architecture called the value decomposition network. This architecture learns to decompose the team value function into agent-wise value functions.

Specifically, the proposed representation encodes information received from different agents in a multi-channel image, where each channel corresponds to a specific feature based on the local view of the agents. This representation helps alleviate permutation invariance issues. However, it is important to note that the information obtained from neighboring agents is of a spatial nature, which means that the dimensionality of the representation increases linearly with each feature. Additionally, representing the information as pixels in an image format can introduce accuracy loss due to quantization errors.

While the reviewed MARL research findings are valuable, it is important to acknowledge their limitations. These findings have primarily been validated using simple tasks and a small number of robots, which may not fully capture the complexity and challenges of real-world scenarios. To evaluate the effectiveness of MARL mechanisms in practical swarm applications, further research and validation are needed to test these mechanisms in more complex task environments and with larger swarms to assess their scalability and robustness.

2.2.3 Evolutionary Robotics

ER utilizes EAs to design and optimize different aspects of a robot, including its morphology and control. EAs draw inspiration from biological evolution to find solutions to complex problems. A generic EA typically involves three phases, namely, initialization, evaluation, and reproduction (refer to figure 2.9). In the *initialization* phase, a population of genomes is created, where each genome represents a potential robot aspect, such as a controller or morphology. The genomes are randomly generated to explore a wide range of possibilities. During the *evaluation* phase, the quality of each genome is assessed using a fitness function. This function measures how well a particular robot aspect performs in achieving its objectives or solving the given task. In the *reproduction* phase, the best-performing genomes are

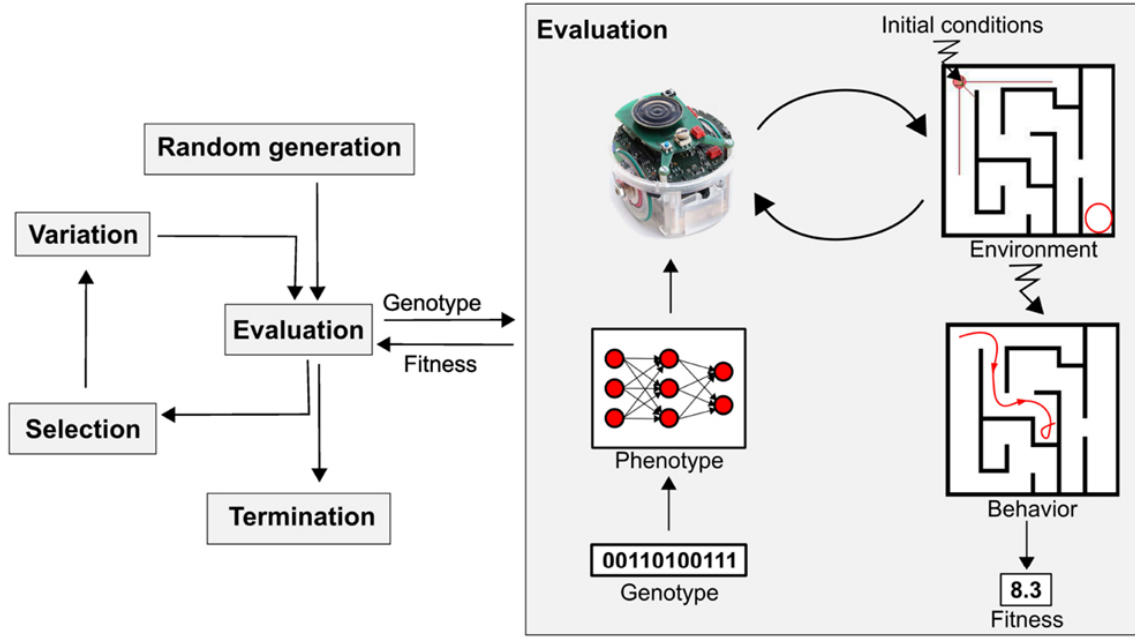


Figure 2.9: The steps of an evolutionary algorithm (Doncieux et al., 2015).

selected for further improvement. Through stochastic variations, such as mutation (introducing small random changes) and crossover (combining genetic information from different genomes), new solutions are generated. These variations allow for exploration of different combinations of traits. Low-quality solutions are typically removed from the population to maintain a constant population size. This *selection* and *reproduction* process continues for a fixed number of generations or until the desired performance is achieved. Finally, the solution with the highest quality, typically the best-performing robot aspect, is selected for manufacturing and deployment in the real-world task environment.

EAs can optimize both the parameters and architecture of a robotic control policy, often implemented as an ANN. This approach reduces the assumptions that need to be made about the optimization problem, as it allows for a more flexible search in the solution space. However, unlike RL methods, EAs do not provide guarantees about finding an optimal solution or specify when such a solution will be found for a given optimization problem (Bongard, 2013). The effectiveness of EAs depends on various factors, including the quality of the fitness function, the representation of the solution space, and the selection and variation operators used.

2.2.3.1 Evolving control

Despite the lack of theoretical convergence proofs, ER techniques have shown successful applications in designing control for swarm robotic systems. Researchers have used various ER methods to evolve controllers for different collective behavior tasks.

For example, Duarte et al. (2016) employed **NeuroEvolution of Augmenting Topologies (NEAT)** to evolve ANN controllers for tasks such as homing, dispersion, clustering, and monitoring in robot swarms. The NEAT algorithm allows for the evolution of both the weights and structure of ANNs, providing flexibility in designing complex controllers.

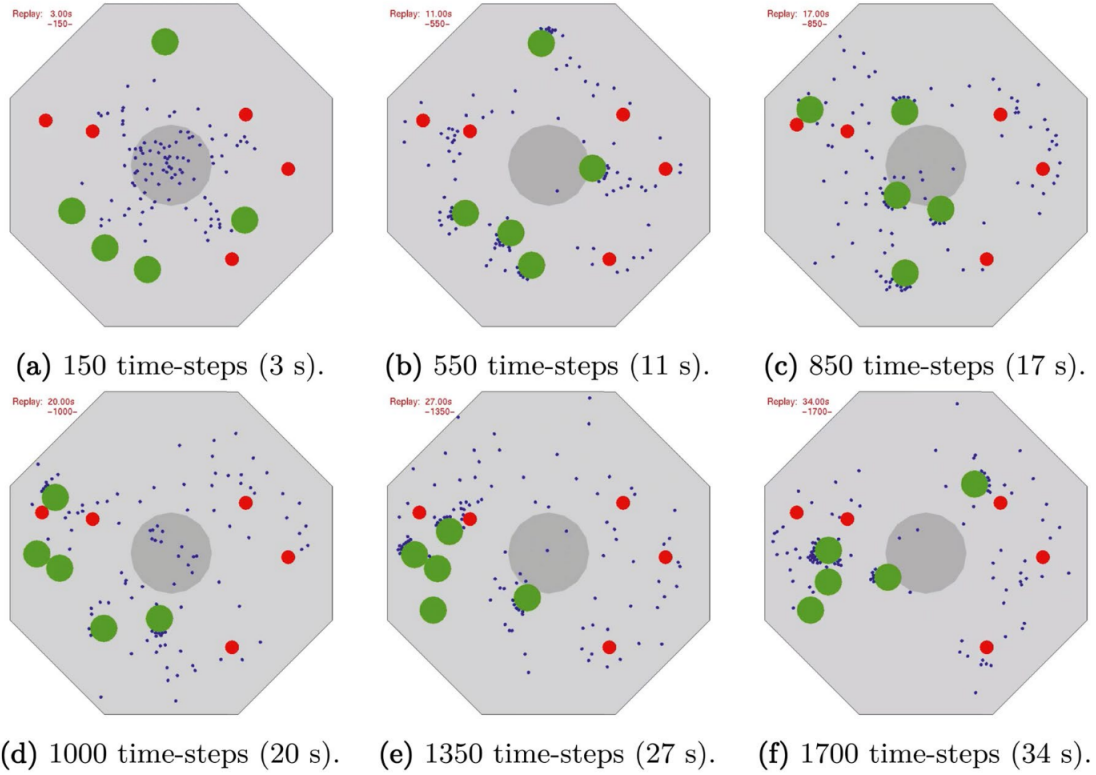


Figure 2.10: Swarm controller evolution with 100 robots. The green circles are food objects while the red circles are poison objects (Hiraga et al., 2020).

In another study, Hiraga et al. (2018) utilized a (μ, λ) evolution strategy (Eiben, Smith, et al., 2003), a type of EA, to optimize the weights of ANN controllers for a path-formation task in swarm robotics. This approach allows for the exploration of the solution space by generating a population of candidate solutions and selecting the best-performing individuals for reproduction and variation.

Furthermore, Hiraga et al. (2020) employed the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) (Hansen & Ostermeier, 1996), another EA, to evolve controllers for robotic swarms in a foraging task that requires collective cognition. CMA-ES adapts the covariance matrix of the search distribution to improve the search efficiency and exploration of the solution space.

Practical validation has shown that EAs can be effectively used to design control for robot swarms, despite the associated costs in terms of time and computational resources. However, the benefits of using EAs outweigh these drawbacks. One of the key advantages is that ER techniques have learning and adaptation capabilities (Doncieux et al., 2015). This means that the designer is not required to have prior knowledge of all possible sources of failure or combinations of outcomes resulting from interactions between agents during the system design phase. ER allows the system to learn and adapt to its environment, improving its performance over time.

Additionally, ER takes a holistic approach to robot design, considering the entire system as a whole. It balances complexity across different components, such as the controller and morphology, to achieve optimal performance. This holistic perspective leads to simpler and better-performing robotic systems (Doncieux et al., 2015). By leveraging the power of EAs, ER can explore a wide range of design possibilities and find solutions that are well-suited for the given task and environment. The next section




Figure 2.11: In related work (Hewland & Nitschke, 2015), morphology adaptation adapted the sensory configuration parameters shown above.

will further delve into this idea.

2.2.3.2 Body-brain co-evolution

While the optimization of robot controllers is commonly pursued in evolutionary swarm robotics, and ER in general, it is important to note that the optimization of robot morphology can also be highly beneficial. The morphology plays a significant role in determining the range and complexity of behaviors that can be exhibited by the robot's controllers (Kriegman et al., 2018).

The concept of body-brain co-evolution, where both the morphology and controller evolve in tandem, has several positive implications. One such benefit is increased resilience, as highlighted by Zhang et al. (2017). By allowing the morphology to adapt and evolve alongside the controller, robots can better withstand and recover from perturbations or environmental changes. This adaptability enhances the robot's ability to navigate and operate in unpredictable or challenging environments.

Furthermore, body-brain co-evolution can lead to higher performance. Furman et al. (2019) have shown that optimizing both the morphology and controller can result in robots that exhibit superior capabilities compared to those optimized with only the controller (see figure 2.11). The interplay between the robot's morphology and its control system can lead to synergistic effects, enabling the robot to perform tasks more efficiently and effectively.

In some cases, adapting the robot's morphology becomes essential to ensure suitable behaviors in previously unknown environments. Kriegman et al. (2019) emphasize that by evolving both the body and brain, robots can overcome limitations and explore environments that were previously considered uninhabitable. This adaptability allows the robot to leverage its physical structure to interact with and

navigate through complex or unfamiliar surroundings.

Notably, there are instances where morphology adaptation is the only viable option to enable desired behaviors in unknown environments. Picardi et al. (2021) highlight scenarios where preconceived notions about the environment may not hold true, and the robot’s morphology needs to be capable of adapting on-the-fly to achieve suitable behaviors.

2.2.4 Embodied Evolutionary Robotics

Embodied Evolution (EE) is a sub-field of **ER** that focuses on continuous adaptation through the use of evolution. In **EE**, evolution takes place within a population of robotic agents that are physically embodied and situated in their task environment. Unlike offline evolution where robots are optimized in a separate simulation or offline environment, **EE** enables the robots to actively evolve and adapt their behaviors in a distributed, asynchronous, and autonomous manner.

Autonomy plays a crucial role in **EE**. Firstly, once deployed, the robotic agents operate and carry out their mission objectives without any human supervision. They are self-sufficient in performing their tasks and do not require continuous human intervention. Secondly, the agents possess the capability to continuously evaluate their own behavior and make adjustments to adapt to the given task and changing environment. This autonomous learning allows the agents to improve their performance and adapt to new circumstances without human oversight or supervision.

Ficici et al. (1999) pioneered the **EE** methodology with three primary objectives, namely, mitigating the reality gap problem in transferring robot controllers from simulators to real robots, enhancing the robustness of the evolution process in robotics, and improving its scalability. By integrating evolution within physically embodied robots situated in their task environment, **EE** aimed to bridge the gap between simulation and reality, ensuring that evolved controllers perform well in real-world conditions.

However, **EE** also has certain limitations, with the most significant being the requirement for a large population size. To maintain genotypic diversity and avoid premature convergence, **EE** necessitates a substantial number of robots in the environment. This large population size enables a wider exploration of the solution space and increases the chances of finding high-quality solutions. Nevertheless, managing and coordinating a large number of robots can pose logistical and practical challenges, such as increased computational resources, communication overhead, and potential physical constraints in the environment.

To address the challenge of large population size in **EE**, an alternative approach called **encapsulated Embodied Evolution (eEE)** has been developed. In **eEE**, each agent in the swarm carries its own population of genotypes and conducts its own evolution process, either partially or entirely independently from other agents. This decentralized approach allows for more efficient evolution and reduces the need for a large population size.

One notable method that follows the **eEE** approach is the **EMbodied Open-ended evoluTIONary ALgorithm (EMOTIONAL)** (Nogueira et al., 2020). **EMOTIONAL** is designed to evolve behaviors for a single autonomous agent without requiring explicit objectives, evaluation metrics, or cooperative

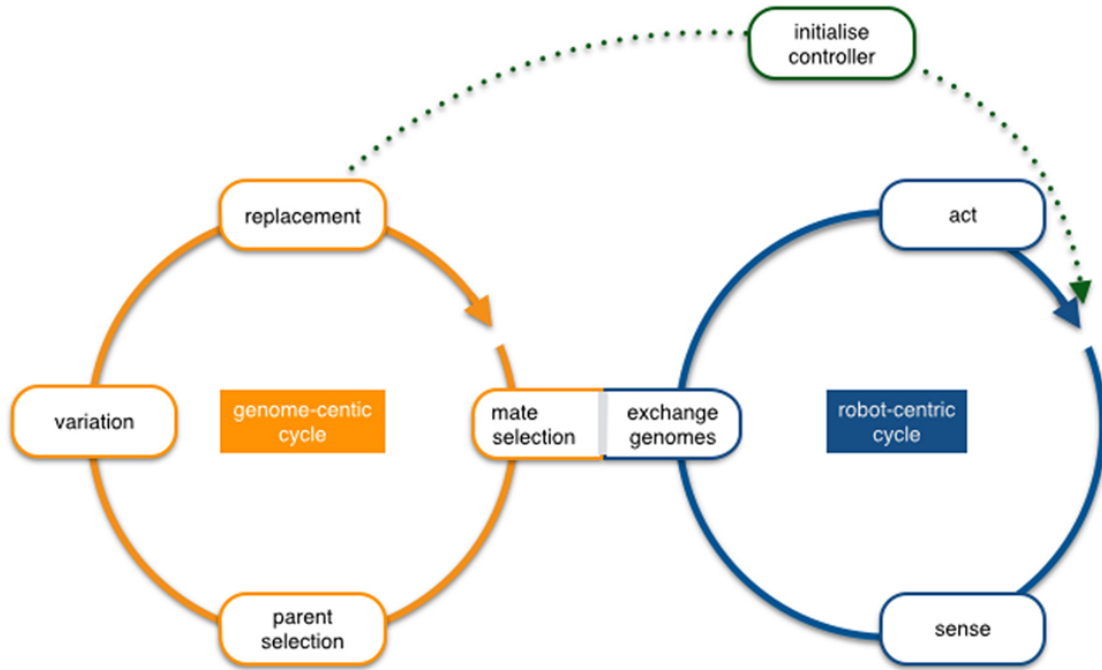


Figure 2.12: Visualisation of the necessary steps of an [EER](#) implementation (Bredeche et al., 2018).

dynamics. Instead, the algorithm promotes the emergence of open-ended evolution, where agents can continuously adapt and generate novel and diverse behaviors without specific predefined goals.

2.2.5 Evolutionary Embodied Collective Robotics

In [distributed Embodied Evolution \(dEE\)](#), individual agents in the swarm carry their own unique genome, and the exchange of genetic information occurs through interactions between the robotic agents. This exchange of genetic material, known as mating, plays a crucial role in [dEE](#). Mating is influenced by heuristics predetermined by the designer and the evolved behavior of the agents. Evolution in [dEE](#) is characterized by decentralization, online adaptation, and parallel processing, with the additional step of mating facilitating the genetic exchange between agents (refer to figure 2.12).

2.2.5.1 Evolving control

The original [dEE](#) algorithm, known as the [Probabilistic Gene Transfer Algorithm \(PGTA\)](#) (Ficici et al., 1999), was specifically designed for evolving a population of real robots. During its development, practical considerations, including inter-agent communication, were carefully taken into account. One notable feature of [PGTA](#) is its minimal reliance on inter-agent communication, and it eliminates the need for a separate reproduction phase by integrating reproduction with task behavior. In their study, Ficici et al. (1999) assess the effectiveness of [PGTA](#) in a phototaxis task involving a population of eight robots. The results obtained demonstrate that solutions evolved through [PGTA](#) outperform manually designed solutions.

The **mEDEA** algorithm is a widely studied approach for adapting robotic-swarm controllers within the **dEE** framework. It has been a subject of various research efforts, including extensions and combinations with other algorithms, leading to the development of novel approaches in various studies (Galassi et al., 2016; Hart et al., 2018; Shan & Mostaghim, 2021; Silva et al., 2015).

In the work of Galassi et al. (2016), **mEDEA** is combined with **Novelty Search (NS)** (section 2.2.6), resulting in a distributed and online variant of **NS**. This hybrid approach effectively addresses deception and outperforms **mEDEA** in problem domains that involve deception and complexity. Similarly, Hart et al. (2018) hybridize **mEDEA** with **MAP-Elites**, creating a distributed and online variant of **MAP-Elites**. This combination successfully promotes behavioral diversity within a robot swarm, as demonstrated in a token foraging task.

Another integration is presented by Shan and Mostaghim (2021), who combine **mEDEA** with the **Non-dominated Sorting Genetic Algorithm II (NSGA-II)** (Deb et al., 2002). This hybridization enhances the multi-objective optimization capabilities of **mEDEA** by leveraging the strengths of **NSGA-II**. Lastly, Silva et al. (2015) combine **mEDEA** with **NEAT**, providing a combined approach that benefits from the evolutionary capabilities of **mEDEA** and the adaptive neural network structures of **NEAT**.

2.2.6 Novelty Search

In **NS**, the focus is on rewarding the learning method for discovering individuals that exhibit behavior different from previously found individuals. This is in contrast to objective search, where the learning method is rewarded for finding solutions that are closer to a predefined objective in the search space.

NS uses a novelty metric to measure the uniqueness of each behavior. The novelty metric quantifies the distance between a newly found individual and the existing population of individuals in terms of their demonstrated behaviors. It compares the behavior of the new individual to an archive of previously encountered individuals with highly novel behaviors.

The novelty metric serves to determine how areas in the behavior space are rewarded. It is a domain-independent metric that computes the sparseness of visited points in the behavior space. Areas with dense clusters of visited points receive lower rewards because they are considered less novel compared to areas with sparse clusters of visited points.

When a newly found individual exhibits sufficient novelty, it is added to the permanent archive. The archive, along with the current population, provides a comprehensive representation of the areas visited by the learning method and its current location in the search space. The termination condition is typically based on whether the discovered individuals meet the desired goal criteria, which needs to be evaluated by the practitioner.

2.2.6.1 Evolving control

Lehman and Stanley (2011a) conducted experiments using NS to evolve ANN controllers for a single agent in maze navigation and biped locomotion tasks. They demonstrated that NS outperforms objective-based search in complex and deceptive problem domains. Building upon this, Pugh et al. (2016b) used NS in a more challenging deceptive maze navigation task, achieving superior performance compared to objective-based search.

The success of NS in single-agent tasks inspired further research in other areas, such as swarm robotics. Gomes et al. (2013) employed NS to evolve controllers for a homogeneous swarm of robots, evaluating its effectiveness in an aggregation task and resource sharing. They found that NS reliably handles deception, discovers solutions with lower neural network complexity, and combining NS with an objective function improves performance.

Didi and Nitschke (2016) applied NS to evolve controllers for a multi-robot task, specifically RoboCup keep-away soccer. They compared NS with hybrid variants that combine novelty and fitness rewards, as well as other approaches. The novelty-fitness variant outperformed the others in increasingly complex and deceptive tasks. Similarly, Gomes et al. (2017) used NS-driven cooperative coevolution in three simulated multi-robot tasks namely, a predator-prey pursuit task (Nitschke et al., 2012), a simulated multi-rover task (Nitschke et al., 2010) and a herding task (Potter et al., 2001), demonstrating its superiority over fitness-driven coevolution.

Galassi et al. (2016) implemented a novelty-based evolution approach in an online distributed context, evaluating its performance in a foraging task. Their results showed improved performance compared to the mEDEA algorithm, particularly in more complex and deceptive scenarios. Brown et al. (2018) took a different perspective by using NS to search for all possible collective emergent behaviors in a robot swarm with specific capabilities. They discovered various behaviors like wall following, dispersal, cyclic pursuit, and aggregation using a binary-sensor capability model.

Overall, these studies highlight the advantages of NS in handling deception, discovering diverse and effective solutions, and outperforming objective-based search in complex and deceptive problem domains, both in single-agent, collective, and swarm scenarios.

2.2.6.2 Body-brain co-evolution

Joachimczak et al. (2015, 2016) conducted research using NS to co-evolve the body and brain of soft-bodied robots. They found that NS is well-suited for this problem domain. While objective-based search can discover complex morphologies capable of various movements and reshaping, the large search space and complex genotype-phenotype relationship make the problem highly deceptive. Objective-based search often gets stuck in local optima, whereas NS, with its bias towards novelty, enhances evolvability in this context.

However, NS has its limitations. Without an explicit objective, there is no inherent bias towards optimizing solutions once they are found. Recognizing this, Mouret (2011) proposed a concurrent

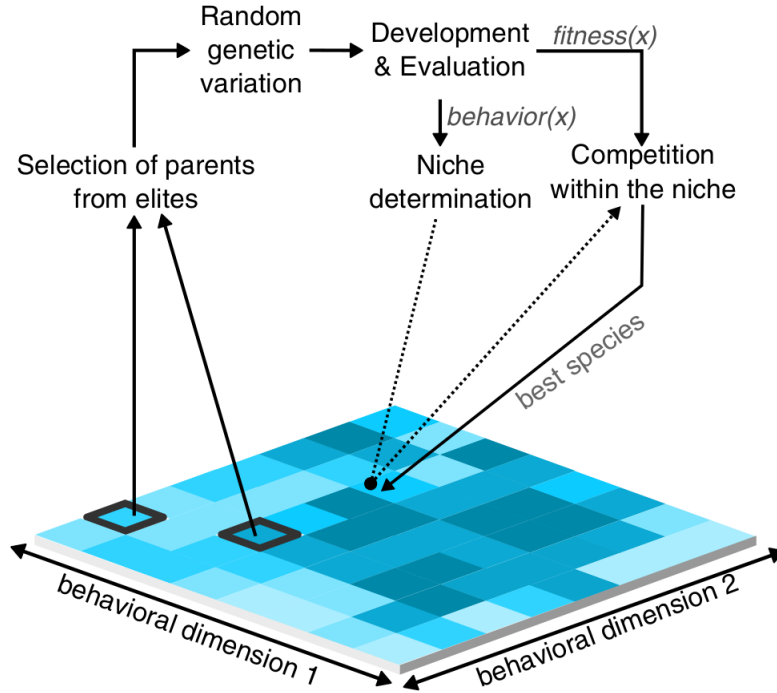


Figure 2.13: [MAP-Elites](#): an example of a [QD](#) algorithm (Mouret, 2020).

approach that rewards both novelty and fitness in a multi-objective formulation of the deceptive maze problem. This variant proved to be more effective in fine-tuning results and generating more precise solutions compared to the original [NS](#) algorithm. This approach allows for a more nuanced exploration of both novelty and fitness in parallel. The next section will discuss a class of algorithms that simultaneously search for novelty and quality in a more refined manner.

2.2.7 Quality Diversity Algorithms

[QD](#) is a class of algorithms that aims to generate a diverse set of high-quality behaviors or solutions in a single run (Pugh et al., 2016b). The algorithm produces an archive or map of behaviors by discretizing a low-dimensional representation known as a [Behavior Characterization \(BC\)](#) (see figure 2.13). Each solution found is mapped to a cell in the map based on its features according to the [BC](#). [QD](#) takes inspiration from natural evolution, where organisms are rewarded for being different rather than just being better.

The [BC](#) is typically a vector of real numbers that describes the actions performed by an individual during evaluation. It represents the space of all possible behaviors. However, in challenging problem domains, certain choices of the [BC](#) may hinder the algorithm from finding high-performing solutions. The [BC](#) plays a crucial role in driving diversity in [QD](#) by quantifying and assessing the novelty of solutions.

In some cases, discovering novel solutions also leads to improved performance, indicating a **highly aligned** [BC](#). However, there are situations where the desired behavior, as characterized by the [BC](#), does not necessarily result in higher performance or quality, indicating an **unaligned** [BC](#). While an algorithm solely driven by an unaligned [BC](#) can still converge to the desired solution due to the quality component

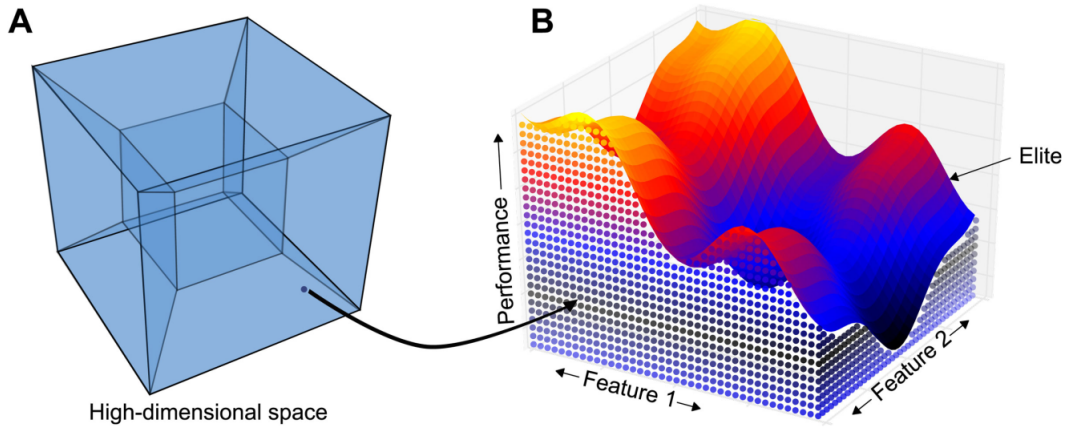


Figure 2.14: [MAP-Elites](#) searches a high-dimensional space for the highest-performing solutions within a single run (Mouret & Clune, 2015) .

of [QD](#), this convergence process can be quite slow (Pugh et al., 2015).

2.2.7.1 Evolving control

[Novelty Search with Local Competition \(NSLC\)](#) is considered the original [QD](#) algorithm (Pugh et al., 2016b), while [MAP-Elites](#) is the most extensively studied [QD](#) approach. [NSLC](#) has shown success in various studies, including the evolution of diverse morphologies or creatures with varying numbers of legs within a single run (Lehman & Stanley, 2011b) and the evolution of hundreds of walking strategies for a hexapod robot (Cully & Mouret, 2016).

On the other hand, [MAP-Elites](#) is widely used in different application domains and has been extended to various contexts, such as online evolution in robot swarms (Hart et al., 2018) and body-brain co-evolution in soft robots (Zardini et al., 2021). In [MAP-Elites](#), the designer specifies N dimensions of variation of interest, defining a low-dimensional space. The algorithm then searches a high-dimensional space to find the highest-performing solutions at each point defined by the N dimensions, creating an archive of diverse and high-performing individuals. The algorithm terminates after a fixed number of iterations or when the desired performance is achieved, and at that point, the archive of diverse and high-performing individuals is returned (refer to figure 2.14).

The multi-[BC](#) [MAP-Elites](#) variant proposed by Pugh et al. (2016b) introduces a second map to the algorithm, maintaining two distinct maps instead of one. These maps are updated independently at the end of each generation. The first map is defined by a highly aligned [BC](#), while the second map is defined by an unaligned [BC](#). This approach aims to overcome deception in difficult problem domains and improve performance for simple task domains. The effectiveness of this approach is evaluated using a deceptive maze navigation task for a wheeled robot controller, and the reported results demonstrate successful evolution of diverse solutions.

Another variant of [MAP-Elites](#) is developed by Hart et al. (2018) and is implemented in a distributed,

online context to evolve robot swarm controllers for a token foraging task. This variant outperforms [mEDEA](#) and has the ability to generate behavioral diversity without the need for carefully selecting parameters (Trueba et al., 2013), introducing a market mechanism (Haasdijk et al., 2014), or employing geographical separation (Montanier et al., 2016). The results indicate that this variant performs significantly better in evolving swarm controllers and showcases the advantages of applying [MAP-Elites](#) in a distributed, online setting.

The design of resilient robots is indeed an interesting application of [MAP-Elites](#). In this approach, a large repertoire of behaviors is evolved in an offline manner under different operating conditions prior to robot deployment. The robot is equipped with this repertoire of behaviors, allowing it to adapt and select the most suitable behavior from the repertoire in response to failures or drastic changes in the environment. For example, Kaushik et al. (2020) utilized [MAP-Elites](#) to evolve repertoires for various situations, such as a hexapod robot with a missing leg or operating on different floors. During online adaptation, the robot can query this repertoire and select the behavior that best matches the current operating conditions. This approach enhances the robot’s resilience by enabling it to dynamically adapt its behavior based on the specific situation it encounters.

2.2.7.2 Body-brain co-evolution

The study by Zardini et al. (2021) introduces the [Double-Map MAP-Elites \(DM-ME\)](#) algorithm to address the co-design of morphology and control in [Tensegrity Soft Modular Robots \(TSMRs\)](#). In this approach, two maps are utilized, each associated with a different feature descriptor—one for entity-related features and the other for controller-related features. Unlike the multi-BC [MAP-Elites](#) approach by Pugh et al. (2016b) that focuses on diversifying controller properties, [DM-ME](#) aims to diversify both morphology and controller properties simultaneously.

The effectiveness of the [DM-ME](#) approach is evaluated through two robotic tasks: goal reaching and squeezing through an aperture. The results of the study indicate that [DM-ME](#) performs exceptionally well in exploring the search space for both tasks. The algorithm successfully discovers a diverse collection of potential designs within a single run, demonstrating its capability to illuminate the design space and provide a range of options for [TSMRs](#) morphologies and control strategies.

By leveraging the power of [MAP-Elites](#) and incorporating multiple maps with different feature descriptors, the [DM-ME](#) algorithm enables the exploration of a broader solution space in the co-design of morphology and control in [TSMRs](#). This approach can facilitate the discovery of novel and effective designs that exhibit diverse morphologies and control properties, leading to enhanced performance and adaptability of [TSMRs](#) in various tasks and environments.

2.3 Environmental Influences on Evolving Behavior and Morphology

The body-brain synergy has been identified as a crucial factor in the development of advanced intelligence, as indicated by several studies (Pfeifer & Bongard, 2006; Weigmann, 2012). In nature, the evolution of the body, brain, and behavior of living creatures is strongly influenced by the environment (Farina, 2012; Sapolsky, 2017; Southwood, 1977). For example, species with distinct genetic ancestors may undergo parallel evolutionary adaptations when exposed to similar ecological conditions (Stern, 2013). This implies that the same concept applies to artificial evolutionary systems given that natural evolution is the hallmark of evolution (O'Shea-Wheller et al., 2020).

Within the realm of swarm robotics, studies investigating the influence of the environment on the evolution of behavior and morphology can be categorized into two main groups. The first category encompasses studies that consider the simultaneous evolution of both robot bodies and brains. These studies explore the interplay between morphology and control in response to environmental factors. However, research in this area is relatively scarce and often has limitations in terms of scope and depth of analysis. The second category focuses on scenarios where the robot morphologies are kept fixed, and only the evolution of the robot brains or controllers is considered. We consider the former.

Hewland and Nitschke (2015) conducted a study where they co-evolved behavior and morphology in robot teams. Their objective was to investigate the relationship between the complexity of the collective gathering task and the resulting morphology of the robot teams. They found that the most complex gathering task, which required a high level of cooperation among the robots, led to the selection of morphologically simple teams.

Interestingly, when they examined a simpler version of the same task that required a medium degree of cooperation, they observed the evolution of more complex controllers and morphologies in the robot teams. This suggests that the complexity of the task influenced the evolution of both behavior and morphology.

It is important to note that in this particular study, there was no cost associated with increasing morphological complexity. This means that the selection process favored the development of more complex morphologies without any trade-offs or disadvantages. The absence of such costs may have contributed to the observed evolution of more complex controllers and morphologies in the simpler task scenario.

Contrary to the absence of cost associated with increasing morphological complexity in the study by Hewland and Nitschke (2015), Nagar et al. (2019a) conducted a study that introduced a neural complexity cost in the evolution of robot groups. They aimed to examine the effect of this cost on the resulting neural complexity in three different task environments of varying complexity, where the measure for task complexity is the same as in (Hewland & Nitschke, 2015).

This study demonstrates that when robot groups are evolved with a neural complexity cost, it leads to the emergence of comparable neural simplicity across all environments. Interestingly, the simplest

neuro-morpho couplings are observed in the simple environment, followed by the medium and difficult environments.

The overall findings of Nagar et al. (2019a) highlight the advantages of incorporating a neural complexity cost during the evolution process. This cost constraint encourages the evolution of simpler controllers, and yet the best-performing simple controllers are capable of producing collective behaviors that are comparable to those achieved by more complex controllers evolved without any complexity cost. This suggests that the neural complexity cost can lead to the development of efficient and effective collective behaviors while promoting simpler neural architectures.

2.4 Discussion

Behavior-based methods have been extensively researched in designing controllers for swarm robotics systems. These methods provide complete control over individual agent behaviors and offer a well-established framework for design and analysis. They rely on mathematical models or physics functions to describe agent behaviors, allowing for analysis and fine-tuning based on mathematical and physics laws. However, these methods have primarily been applied to simpler swarm robotic tasks like aggregation, collective navigation, and self-organized flocking. The iterative process of manually crafting control for each individual in the swarm becomes impractical in more complex, unknown, or dynamic task domains.

RL and ER have been employed in swarm robotics to address the challenges of behavior design in complex problem domains. RL has a well-established theoretical framework with convergence proofs in single agent problems. However, when applied to swarm robotics, the extensive search space makes it difficult for multi-agent learning to find optimal solutions within a reasonable time frame (Bernstein et al., 2002; Panait & Luke, 2005).

On the other hand, ER has shown remarkable success in swarm robotics through practical validation, despite lacking an elaborate theoretical framework. With ER, designers are not required to pre-engineer control aspects, as EAs optimize both the parameters and architecture of a robot’s control policy, typically an ANN. This approach enables the discovery of sophisticated ANN designs that may not have been conceived by human designers. As a result, ER facilitates the design of swarm robot controllers for more complex problem domains, including cooperative tasks, where neither behavior-based methods nor RL alone may suffice. Moreover, ER provides valuable scientific tools for modeling and simulating biological evolution (Long, 2012), allowing researchers to study and understand the mechanisms and dynamics of evolution in a controlled and reproducible manner. Consequently, ER has been employed as an experimental platform to answer various biological hypothesis, such as the *social brain hypothesis* (Dunbar, 2009).

Research in the field of ER has extensively focused on studying the evolution of artificial morphology and behavior, both in simulated (Cheney et al., 2017; Kriegman et al., 2018) and physical (Nygaard et al., 2021a; Xu & Wang, 2021) robotics platforms. However, there is comparatively less research dedicated to exploring the effects of body-brain adaptation in collective or swarm robotic systems (Buason et al., 2005; Furman et al., 2019; Hewland & Nitschke, 2015). This scarcity of studies can be attributed to the inherent challenges in understanding and analyzing the intricate relationships between genotype (representing body-brain encodings) and phenotype plasticity (representing the adaptability of body-brain couplings) that give rise to emergent collective behaviors (Moore et al., 1997). Phenotypic plasticity, in this context, refers to the ability of a single genotype to exhibit multiple morphology-behavior couplings in response to varying environmental conditions (Kelly et al., 2012; Schlichting & Pigliucci, 1998; West-Eberhard, 1989; Wolf et al., 1999).

While previous research (Bongard, 2011; Kriegman et al., 2018) in ER has highlighted the significance of morphological adaptation through phenotypic plasticity in enhancing the robustness of adapted behaviors, there has been limited investigation into the influence of emergent morphological diversity on evolving swarm-robotic behavior (Hunt, 2021). However, there are a few noteworthy examples in this area. One such example is the utilization of self-assembly swarm robotics systems, where numerous

functionally simple robots physically connect to each other (Brambilla et al., 2013; Mathews et al., 2019; Mondada et al., 2004), as demonstrated in proof-of-concept experiments involving hundreds of *Kilobots* (Carrillo-Zapata et al., 2019; Rubenstein et al., 2014; Slavkov et al., 2018). These experiments have explored concepts such as mergeable nervous systems (Mathews et al., 2017) and group minds (Otte, 2018), wherein swarm-robotic behaviors emerge from the interconnections among self-organizing neural controllers across a large number of robots. Additionally, there have been studies involving multi-robot organisms (Levi & Kernbach, 2010) that exhibit morphological adaptation through self-organization into various problem-solving forms. For instance, these adaptations often involve the evolution of functional specialization in different interacting body parts (Auerbach & Bongard, 2009). In the context of smaller swarm sizes, collective robotics research has focused on evolving desired group behaviors through morphological adaptation that dynamically activates or deactivates sensors, allowing robots to adapt to complementary sensory configurations (Furman et al., 2019; Hewland & Nitschke, 2015; Watson & Nitschke, 2015).

Furthermore, while previous research (Auerbach & Bongard, 2014; Miras & Eiben, 2019; Miras et al., 2020; Spanellis et al., 2021) in ER has examined the influence of the environment on body-brain co-evolution, there has been limited exploration of how the environment impacts body-brain evolution specifically in the context of swarm robotics (Furman et al., 2019; Nagar et al., 2019a, 2019b). Studies focusing on phenotypic plasticity in evolutionary swarm robotics can be broadly categorized into two groups. The first category involves studies where robot morphologies remain fixed while only the controllers evolve and second, where both the controllers and morphologies of individual robots are coupled and co-evolved (Doncieux et al., 2015). When considering fixed morphologies, there are only a few studies that have demonstrated the potential for diverse environments to give rise to diverse behaviors (Ferrante et al., 2013; Ferrante et al., 2015).

However, research on evolutionary swarm robotics that incorporates evolvable coupled controllers and morphologies is scarce and limited in scope (Dias et al., 2021), and lacks comprehensive exploration of the impact of environmental complexity on phenotypic plasticity. Typically, experiments in this field involve setting a specific environment and task and evolving robots to perform in that particular combination. For instance, many swarm robotics studies consider the task environment as an experimental parameter and evolve behavioral diversity using predetermined morphologies to solve various collective behavior tasks (Brutschy et al., 2012; Ferrante et al., 2015; Nitschke et al., 2011; Steyven et al., 2017; van Diggelen & et al., 2022).

Finally, in both biology (Armstrong, 1983; Carroll, 2001; Isler & van Schaik, 2009) and robotic design (Cheney et al., 2013; Nygaard et al., 2021a; Xu & Wang, 2021), an unanswered question pertains to the ideal type of body-brain coupling that enables survival and task-solving in diverse environments. While natural evolution offers some examples, the automation of body-brain design in robotics requires exploration. Limited research has been conducted on employing behavior-diversity maintenance methods, specifically QD approaches, to automate the design of body-brain couplings in cooperative tasks within robotic swarms. Although previous work (Cully & Mouret, 2016; Hart et al., 2018; Lehman & Stanley, 2011b; Zardini et al., 2021) has demonstrated the value of these methods, their effectiveness in cooperative body-brain design remains unknown. Consequently, this thesis proposes a method that seeks to evolve QD through the body-brain co-evolution of robot swarms engaged in cooperative tasks.

Chapter 3

Methodology

This thesis employs a bespoke evolutionary embodied collective robotics simulator to evaluate swarm-robotic controller and morphology evolution. Specifically, our work applies five different methods, namely [mEDEA](#), an extension to [mEDEA](#) (termed [mEDEA-M](#)), [EDQD](#), and two extensions to [EDQD](#) (termed [EDQD-M](#), and Double-Map [EDQD-M](#)), to evolve *behavior*, *behavior and morphology*, *behavior with behavioral diversity maintenance*, *behavior and morphology with behavioral diversity maintenance*, and *behavior and morphology with behavioral-morphological diversity maintenance*, respectively, across a robot swarm given increasingly complex cooperative tasks. Both [EDQD](#) and [mEDEA](#) have been extended to allow for body-brain co-evolution. We chose to extend [EDQD](#) because it has already been shown to evolve a behaviorally diverse robot swarm without the need for explicit mechanisms for genotypic isolation (Montanier et al., [2016](#)) or division of labor (Haasdijk et al., [2014](#)). We chose [mEDEA](#) as a benchmark since it is a widely studied method in embodied evolution (Galassi et al., [2016](#); Hart et al., [2018](#); Shan & Mostaghim, [2021](#); Silva et al., [2015](#)).

3.1 Robot Controllers

To build on earlier research (Hart et al., [2018](#)), robots explore their environment for the duration (table [4.2](#)) of their *lifetimes* in accordance with an [ANN](#) controller behavior, where the [ANN](#) behavior is adapted by either [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), or Double-Map [EDQD-M](#) (sections [3.2](#), [3.3](#), [3.4](#), [3.5](#), and [3.6](#), respectively). Each robot in the swarm employed the same controller topology, a fully connected feed-forward [ANN](#) with 33 sensory input nodes (proximity, color, target-area detection), a 20 node hidden layer, and two motor output nodes (table [4.2](#)).

For each simulation (robot lifetime) iteration, the two [ANN](#) outputs corresponded to the rotational and translational speeds of each robot. 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). In order to detect the closest object in the environment, proximity sensors were set up so that the closer an object is to the robot, the higher the sensor activation value (normalised to the range: $[0, 1]$). However, there were seven object type (color) detection sensors for each forward and backward-facing proximity sensor that would activate to discriminate between the colors of five resource types, walls, and other robots (table [4.2](#)). Hence, the robot’s peripheral consisted of four sensor sets

Table 3.1: An overview of each of the evolutionary methods evaluated in this thesis.

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

(each with eight sensors) and one downward-facing target-area detection sensor (figure 3.1, left), with these 33 sensors corresponding to the ANN input layer. As a result, the controller genotype adapted by mEDEA, mEDEA-M, EDQD, EDQD-M, or Double-Map EDQD-M had 700 connection weights. More specifically, there were 33 input nodes fully connected to 20 hidden nodes (33x20 connections), which were then fully connected to two output nodes (20x2 connections).

3.2 mEDEA

In mEDEA, an agent moves around its environment according to its control algorithm for a fixed period of time (a lifetime). When it moves, it broadcasts its genome to neighboring agents within its communication range while simultaneously receiving genomes from nearby agents. Once the lifetime of its currently active genome has ended, the agent randomly selects parents from the list of received genomes. Following the use of variation operators, the selected genetic material is copied to form progeny. The offspring then replaces the currently active genome. Later variants incorporate an explicit objective function to perform a specific task however, the key concepts do not change except now the fitness value is broadcasted with the genetic material and selection with regard to task performance is active. The latter variant is employed in this work.

3.3 mEDEA-M

The mEDEA-M extension of mEDEA enables morphological (sensor) and behavioral (controller) adaptation for each robot. In particular, a random sensor type is chosen to undergo mutation at the end of each generation. The mutation operator limits the range of a sensor chosen at random until it exceeds a specified *sensor-morpho* threshold (table 4.2). When the range drops below this threshold, the supplied sensor goes into inactive mode, which is indicated by an input of zero to the appropriate ANN sensory input node. Similar to this, if the sensor range exceeds the *sensor-morpho* threshold due to the mutation operator, an inactive sensor will reactivate, restoring the prior non-zero connection weight value for the specified ANN sensory input node. The target-area detection sensor (section 3.1) that faces the ground is not affected by morphological adaptation since robots still need to be able to find the target area and finish their work. The robot swarm is morphologically homogeneous, which means that the same sensor adaptations (sensors turned on and off) are simultaneously applied to all of the robots in the swarm. Other than that, mEDEA-M controller adaptation is the same as mEDEA (section 3.2).




Figure 3.1: **Left:** *Simulated robotic agent*: Each robot’s peripheral consisted of four sensor sets (three forward facing and one backward facing, all blue in color), each with eight sensors and one downward-facing target-area detection sensor (below the core-component/the red-cube, but not visible in this aerial picture). **Right:** *EDQD*: Each generation, a random genotype (controller encoding) is selected from the *SelectMap* – which is created by merging the *ReceivedMapList* with the robot’s *LocalMap*.

3.4 EDQD

The *EDQD* method hybridizes the *MAP-Elites* and *mEDEA* methods. Differing from *mEDEA*, as robots explore their environment they periodically broadcast (table 4.2), their behavioral map (list of genotypes), instead of genomes, which is received and stored by all robots within broadcast range. Such robot behavioral maps are termed *LocalMaps*, and contain a map of the genotypes (*genome*, figure 3.1, right) corresponding to specific robot behaviors (*phenome*, figure 3.1, right) previously evaluated for each robot. At the end of each robot’s current *lifetime* (table 4.2), a genotype is randomly selected from the *SelectMap* (which is formed by merging the received maps with the robot’s own *LocalMap*, figure 3.1, right), and a mutation operator (table 4.2) is applied to produce a new genotype which replaces the currently active genotype (robot behavior).

Applying *EDQD* to our swarm, robots store a 2D behavior map (*LocalMap*, figure 3.1, right) defined by two behavioral dimensions of the collective gathering task (section 4). Specifically, resource type collected (table 4.2), and maximum Euclidean distance traversed (explored) in the environment, by each robot (during its lifetime). The *EDQD* method and its extensions: *EDQD-M* and *Double-Map EDQD-M*, thus purposefully leveraged these behavioral dimensions to promote the evolution of behavioral diversity in terms of resource types collected and environment exploration.

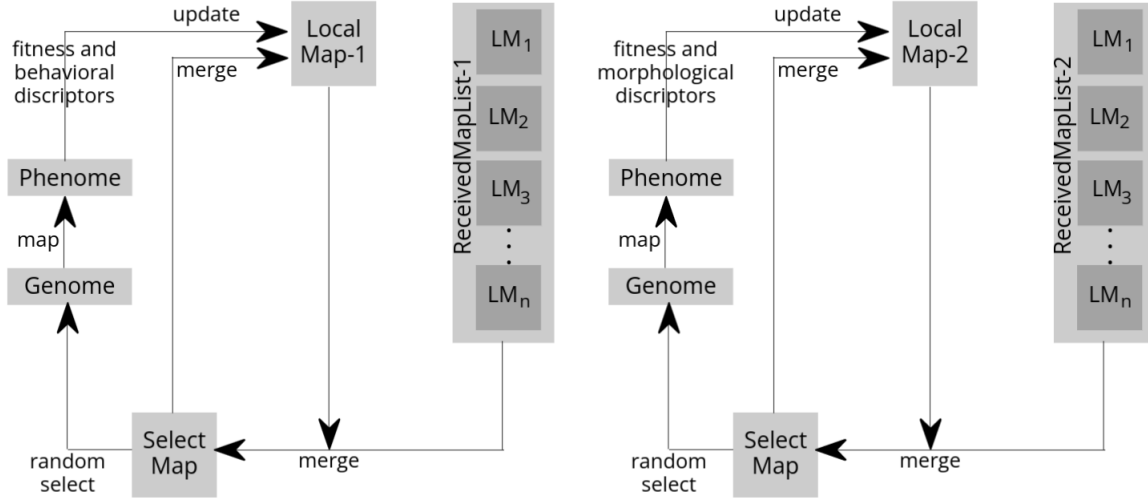


Figure 3.2: *Double-Map EDQD-M*: Robots maintain two *LocalMaps*. *LocalMap-1* is associated with behavior feature descriptors, *LocalMap-2* is associated with morphology feature descriptors.

3.5 EDQD-M

EDQD-M extends *EDQD* by allowing morphological (sensor) and behavioral (controller) adaptability for each robot. As with *mEDEA-M*, a random sensor type is chosen at the end of each generation to undergo mutation. The mutation operator narrows the range of a randomly chosen sensor until it reaches a pre-specified *sensor-morpho* threshold (table 4.2). Once the range falls below this threshold, the given sensor becomes inactive, which is indicated by a zero input to the related *ANN* sensory input node. Similarly, if the mutation operator increases the sensor range above the *sensor-morpho* threshold, an inactive sensor will reactivate, restoring the prior non-zero connection weight value for the associated *ANN* sensory input node. The robot swarm is morphologically homogeneous, which means that the same sensor adaptations (sensors turned on and off) are applied to all robots in the swarm at the same time. Otherwise, the *EDQD-M* controller adaption procedure using the *LocalMap* is the same as the *EDQD* method (section 3.4).

3.6 Double-Map EDQD-M

Double-Map *EDQD-M* is an extension of *EDQD* that allows for the co-adaptation of a robot's morphology and behavior. Double-Map *EDQD-M* employs two *LocalMaps*, the first of which is associated with controller (behavior) feature descriptors and the second with 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* and *EDQD-M*, a random genotype is independently selected from each *SelectMap* to undergo mutation at the end of each generation of the evolutionary process. The chosen genotypes from each *SelectMap* replace the robot's current active behavior (controller) and morphology (sensory configuration), respectively (figure 3.2). The selected morphology, however, determines the corresponding controller, ensuring that robot behavior and morphology are appropriately matched. As a result, a chosen morphology with x active and y inactive sensors automatically reconfigures the chosen *ANN* controller so that x sensory input nodes are active and y are inactive.

As in [EDQD-M](#) (section 3.5), [ANN](#) connection weights remain active, and robot sensory configuration (morphology) is adapted by switching specific sensors on and off, with zero values being input to [ANN](#) inputs corresponding to switched off sensors. As a result, Double-Map [EDQD-M](#) adapts both of each robot's *LocalMaps* to promote morphological and behavioral diversity. In contrast to [EDQD](#) and [EDQD-M](#), two maps are concurrently maintained and adapted in this case, where both robot behavior and morphology are subject to diversity maintenance. [EDQD](#) and [EDQD-M](#), on the other hand, only accounted for behavioral diversity maintenance, with sensor adaptation in [EDQD-M](#) not accounting for morphological diversity maintenance.

3.7 Summary

In this Chapter we described each of the evolutionary methods evaluated in order to answer the research questions of this work. Specifically, we described the implementation of [mEDEA](#), [mEDEA-M](#) (an extension of [mEDEA](#)), [EDQD](#), and two variants of [EDQD](#), namely, [EDQD-M](#) and Double-Map [EDQD-M](#).

Chapter 4

Experiments

All experiments were carried out using a collective gathering task (section 4.1) simulation implemented on RoboGen (section 4.5). Experiments (table 4.3) assess the capability of [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) (sections 3.2, 3.3, 3.4, 3.5, and 3.6, respectively) to evolve *efficient* (section 4.3.1) and *diverse* (section 4.3.1) collective behavior for solving progressively challenging (section 4.2) collective gathering tasks.

4.1 Collective Behavior Task

The collective gathering task used to evaluate [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) involved running the swarm for one *lifetime* (10000 simulation iterations) and 100 generations, with each generation representing a lifetime (table 4.2). At the start of 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, find them, and then cooperatively move them to the *target-area*.

Table 4.1: Parameters for robot [ANN](#) controllers adapted by [MAP-Elites](#) component of [EDQD](#), [EDQD-M](#), Double-Map [EDQD-M](#) swarm behavior-morphology adaptation methods.

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 (tuned σ)
Sigma range	[0.001, 0.5]
Update sigma step	0.35
Mutation probability	0.34
Sensor-morpho threshold	0 (≤ 0 : Sensor inactive; > 0 : Sensor active)
Map archive size	100
Number of dimensions per map	2
Number of intervals per map-dimension	10

Table 4.2: Experiment parameters for [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#) and Double-Map [EDQD-M](#) methods (applied to adapt swarm behavior-morphology) and collective gathering task (evaluating swarm adaptation methods) 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
Task environments (Resource types: A, B, C, D, E)	1: Simple	30, 5, 5, 5, 5
	2: Medium	10, 10, 10, 10, 10
	3: Difficult	5, 5, 5, 5, 30
Cooperation needed to move resource type	A	1 robot
	B	2 robots
	C	3 robots
	D	4 robots
	E	5 robots
Run length (per experiment)	100 generations	
Robot lifetime (simulation iterations)	10 000	
Swarm size	100 robots	
Wait for assistance (cooperation) time	Remaining lifetime	
Initial robot & block position (outside target-area)	Random	
Environment size Target-area size (meters)	20 x 20 20 x 2	
Robot <i>LocalMap</i> broadcast range	Environment size	
Robot <i>LocalMap</i> broadcast frequency	1 (per lifetime)	

4.2 Task Complexity

The degree of task difficulty is determined by the level of team cooperation required to cooperatively collect all resources in the environment (table 4.2). Each resource type (A, B, C, D, E) differed in geometric size and so required varied degrees of cooperation for robots to transport to the *target-area* (table 4.2). A *simple* environment, for example, implies a high proportion of *type-A* resources (which can be collected without the need for cooperation) whereas a *difficult* environment implies a high proportion of *type-E* resources (which require five cooperating robots to be transported) thus necessitating more team cooperation. The rationale for varying task difficulty is to address the [primary research question](#), as well as for completeness when answering the secondary research questions ([research question 2.1](#), [2.2](#), and [2.3](#)).

4.3 Experiments Setup

Five different sets of experiments were conducted in order to answer the questions of this thesis (table 4.3). All experiments were designed to address the primary question of this work, to investigate whether environmental conditions (that is, varying degrees of task complexity) influence the efficiency (that is, task performance) and diversity (that is, behavioral and morphological diversity) of body-brain evolved cooperative robot teams across increasingly challenging collective gathering tasks. Experiments number 3, 4, and 5 (table 4.3) seek to further address the secondary questions of this work.

Table 4.3: An overview of each of the experiments conducted in this thesis, and the specific objective each experiment is trying to satisfy.

Experiment number	Evolutionary method	Task difficulty		Research objective
1	mEDEA	<i>Simple, Difficult</i>	<i>Medium,</i>	Investigates whether environmental conditions (that is, varying degrees of task complexity) influence task performance and behavioral diversity of body-brain evolved cooperative robot teams.
2	mEDEA-M	<i>Simple, Difficult</i>	<i>Medium,</i>	Investigates whether environmental conditions (that is, varying degrees of task complexity) influence task performance and body-brain diversity of body-brain evolved cooperative robot teams.
3	EDQD	<i>Simple, Difficult</i>	<i>Medium,</i>	Investigates the impact of behavioral diversity maintenance on evolving swarm behavior (research question 2.1.).
4	EDQD-M	<i>Simple, Difficult</i>	<i>Medium,</i>	Investigates the impact of behavioral diversity maintenance on co-evolving behavior and morphology (research question 2.2.).
5	Double-Map EDQD-M	<i>Simple, Difficult</i>	<i>Medium,</i>	Investigates the impact of behavioral-morphological diversity maintenance on co-evolving behavior and morphology (research question 2.3.).

4.3.1 Collective Behavior Evaluation

For each experiment, a 3D environment (*simple*, *medium* and *difficult* task environments, table 4.2) 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 mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M the percentage of resources that were pushed into the *target-area* throughout the course of all swarm lives (that is, per run: 100 generations) was used to compute the average (over 20 runs) swarm (team) task performance. After that, the average swarm task performance was normalized to: [0.0, 1.0].

4.3.2 Behavior Quality Evaluation

Behavior quality (fitness) for each evaluated controller (the value that determines whether a found solution is high performing enough to be stored in each grid cell) was calculated according to equation 4.1.

$$fitness = \frac{resourceValue}{numberOfPushingRobots} * \frac{distanceResourceMoved}{totalDistancePossible} \quad (4.1)$$

Where the *resourceValue* is equal to the minimum number of robots required to push a resource, *numberOfPushingRobots* is the number of robots attached (pushing) to a resource, *distanceResourceMoved* is the straight-line distance that the resource is moved towards the gathering

zone, and *totalDistancePossible* is the straight-line distance from one end of the environment to the center of the gathering zone.

4.3.3 Diversity Evaluation

Also, at the conclusion of each run (100 generations), we determine the behavioral diversity for [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) (specifically, the behavior map is used for this analysis) as the quantity of unique behaviors (occupied cells in the swarm’s *LocalMap*). [mEDEA](#) and [mEDEA-M](#) do not require the maintenance of a behavior map however, for ease of analysis one is created, and updated (in the same fashion as the [QD](#) methods) at the end of each generation. Note that this map does not form part of the evolutionary process, it is maintained solely for analysis purposes only.

In line with prior research (Pugh et al., 2015), we additionally determine the average [QD](#) score for behaviors discovered by [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) over the course of 20 runs.

4.3.4 mEDEA: Robotic Swarm Behavior Evolution

The objective of this experiment (experiment number 1, table 4.3) is to investigate whether environmental conditions (that is, varying degrees of task complexity) influence task performance and behavioral diversity of body-brain evolved cooperative robot teams. The [mEDEA](#) algorithm (section 3.2) is chosen to establish a benchmark because of its popularity.

4.3.5 mEDEA-M: Robotic Swarm Behavior-Morphology Co-evolution

This experiment (experiment number 2, table 4.3) uses the [mEDEA-M](#) algorithm (section 3.3) to evolve both collective behavior and morphology in robotic swarms. The objective is to investigate whether environmental conditions (that is, varying degrees of task complexity) influence the efficiency (that is, task performance) and diversity (that is, behavioral and morphological diversity) of body-brain evolved cooperative robot teams across increasingly challenging collective gathering tasks.

Given that [mEDEA-M](#) does not maintain a morphology map, one is created at the conclusion of each generation in order to compare [mEDEA-M](#) to the [EDQD-M](#) and Double-Map [EDQD-M](#) methods in terms of quality of evolved morphologies. This morphology map is not used during the evolutionary process.

4.3.6 EDQD: Robotic Swarm Behavior Evolution with Behavioral Diversity Maintenance

This experiment (experiment number 3, table 4.3) uses the EDQD algorithm to evolve collective behavior with behavioral diversity maintenance. The objective is to investigate the impact of collective behavior evolution with behavioral diversity maintenance across cooperative robot teams given an increasingly challenging collective behavior task. Evolution takes place in three environments of increasing task difficulty for a collective gathering task.

4.3.7 EDQD-M: Robotic Swarm Behavior-Morphology Co-evolution with Behavioral Diversity Maintenance

This experiment (experiment number 4, table 4.3) utilizes the EDQD-M algorithm (section 3.5) to evolve both robot behavior (controller task performance) and morphology with behavioral diversity maintenance. The objective is to elucidate the impact of behavior-morphology evolution with behavioral diversity maintenance in cooperative robot swarms given progressively challenging collective gathering tasks.

Given that EDQD-M does not maintain a morphology-related *LocalMap*, one is created and updated at the end of each generation in order to compare EDQD-M to the mEDEA-M and Double-Map EDQD-M methods in terms of quality of evolved morphologies. As in mEDEA-M, this morphology map is maintained solely for analysis purposes.

4.3.8 Double-Map EDQD-M: Robotic Swarm Behavior-Morphology Evolution with Behavior-Morphology Diversity Maintenance

This experiment (experiment number 5, table 4.3) employs the novel Double-Map EDQD-M algorithm (section 3.6) to adapt robotic swarm behavior (controller task performance) and morphology (sensory configuration) with behavior-morphology diversity maintenance. The objective is to investigate the impact of behavioral-morphological diversity maintenance in co-evolving behavior and morphology in cooperative robot teams given increasingly challenging collective gathering tasks.

At the conclusion of each run (100 generations), we determine a swarm’s average morphological diversity (across 20 runs) as the number of distinctive morphologies evolved (occupied cells in the swarm’s second *LocalMap*).

4.4 Parameter Tuning

Parameter tuning experiments were used to derive the values in the parameter values table (table 4.2). Five different task environments were evaluated, each with a different complexity level. The different task environments were coined *super-simple*, *simple*, *medium*, *difficult*, and *super-difficult*. The resource type distribution of the *simple*, *medium*, and *difficult* environments is given in table 4.2. Otherwise, in the *super-simple* environment, 80% of the resources were type A resources, with an equal distribution of type B, C, D, and E resources. In the *super-difficult* environment, 80% of the resources were type E resources, with an equal distribution of type A, B, C, and D resources. The number of simulation evaluations (generations) was determined through trial and error over the range [100, 200]. The number of iterations per genome lifetime was also determined experimentally over the range [8000, 12000]. Lastly, three swarm sizes were evaluated; 100, 150, and 200 robots.

4.5 Simulator

All experiment simulations were conducted on an open source framework known as RoboGen (Auerbach et al., 2018). RoboGen was created with the goal of evolving robots that can be easily produced using 3D printing and a modest collection of low-cost, off-the-shelf electronic components. It has a physics simulation engine as well as an evolution engine (figure 4.1). It also includes utilities for creating 3D-printable design files for [body components](#) and compiling neural-network controllers to run on an Arduino microcontroller board. Earlier work (Jelisavcic et al., 2017) has proved the functionality of this modular robot system in real hardware. This was the main motivation for implementing our experiments on this framework. However, the original platform can only simulate one robot at any given instance. As a result, we had to extend the simulation engine component of RoboGen to allow for the simulation of a robot swarm. This involved a number of design considerations such as, performance, accuracy, and how to handle communication between robotic agents. The extended version¹ is written in c++.

All simulation experiments were implemented on a homogeneous cluster, comprising Intel 5th generation CPUs. The cluster has 1368 compute nodes with 24 cores and 128 GiB memory. Local testing was conducted on a Linux Ubuntu 22.04.2 LTS computer with an 11th Gen Intel Core i7-1165G7 (@ 2.80GHz \times 8) processor, Intel iRISxe graphics and 32 GB memory (RAM).

4.5.1 Robotic Agents

Each robotic agent used in the experiments of this thesis consisted of a core component, four fixed bricks, five infrared sensors, two active wheels, and two passive wheels. The core component (figure 4.2a) is a cube that houses the inertial measurement unit (IMU) as well as the microcontroller. The side length of the cube is 4.65 cm. The fixed brick is the same as the core component, except it does not contain the latter’s electronics. The infrared sensor (figure 4.2b) is a range sensor that measures distance in a 10-millimeter-wide beam perpendicular to the sensor across a 10-millimeter to 1-meter range. The

¹The simulator used in this study is available online: <https://github.com/einstein07/AUTOFAC>

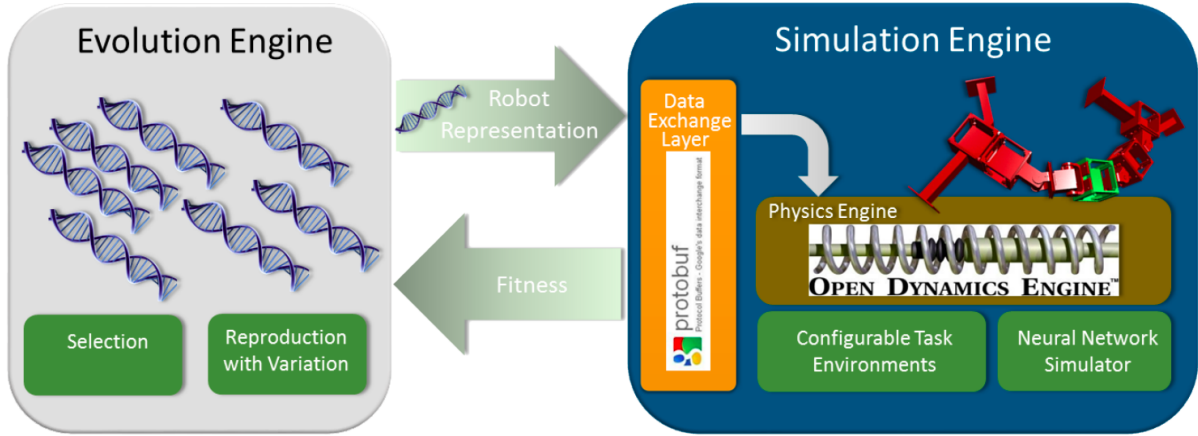


Figure 4.1: Overview of RoboGen framework (Auerbach et al., 2018).




Figure 4.2: Robotic agent components.

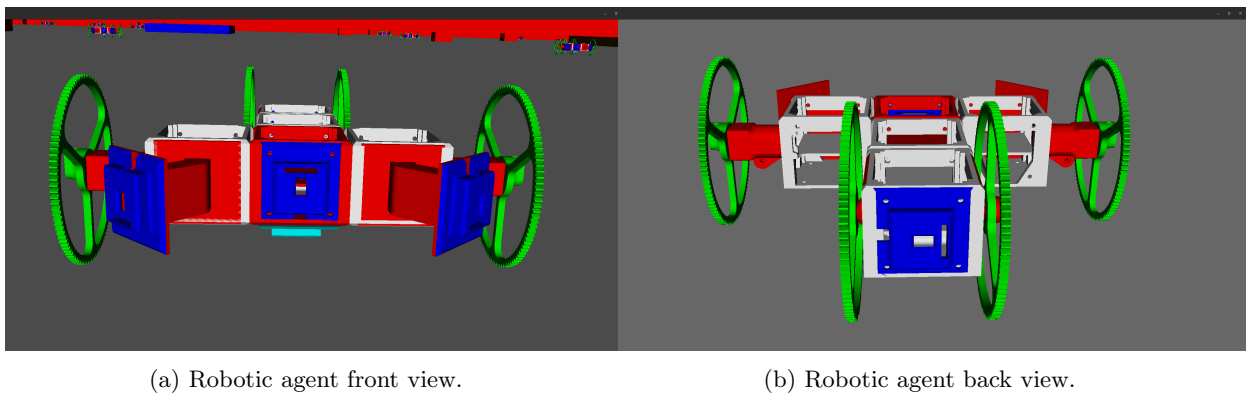


Figure 4.3: Front and back views of a complete robotic agent.

active wheel (figure 4.2c) differs from the passive wheel (figure 4.2d) in that it contains a rotation motor, otherwise the two wheel types are of the same diameter, that is, 0.04 meters. Figures 4.3a and 4.3b show




Figure 4.4: Task environment with robotic agents cooperating to move resources.

front and back views of a robotic agent in its task environment. An example of how the task environment looks like is also shown in figure 4.4.

4.6 Summary

In this Chapter, we described each of the experiments carried out as part of this thesis. More specifically, we defined each experiment by the evolutionary method type (chapter 3), task complexity (table 4.2), and research objective that the experiment is attempting to meet (research questions). We went on to discuss how all experiments are set up and carried out, including initial robot and object positions, robot lifespan, the number of generations, how task performance, diversity, and quality diversity are measured, and how averages are determined.

Chapter 5

Results

In this chapter, we present the results of the experiments (table 4.3) conducted in order to evaluate the research questions of this thesis. We specifically present the results of evaluating mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M (experiments 1, 2, 3, 4, and 5, respectively) in terms of quality diversity of behaviors evolved (section 5.1), task performance (section 5.3), and quality diversity score (section 5.4). We further present the results of evaluating mEDEA-M, EDQD-M, and Double-Map EDQD-M in terms of diversity of evolved morphologies, and the relative quality of those morphologies (section 5.2).

5.1 Quality Diversity (QD) of Evolved Behaviors

The behavior map-archive is defined by two dimensions, namely, the resource type collected and maximum Euclidean distance traversed (explored) in the environment by each robot during its lifetime. This means that in order for diverse behaviors to be evolved (and hence fill up the map-archive), each approach needs to encourage the collection of diverse resource types as well as a varied exploration of the task environment.

The remainder of this section is divided into the following sections; Average Behavioral Diversity (section 5.1.1), Average Distance Explored (section 5.1.2), Average Resources Per-type Collected (section 5.1.3), and Behavior Map-archive of Best Performing Swarms (section 5.1.4). Section 5.1.1 presents comparative method average behavioral diversity. Sections 5.1.2 and 5.1.3 present comparative method results in each dimension of the map-archive in order to gain a deeper understanding of how behavioral diversity evolved. Finally, section 5.1.4 compares method behavior map-archives of the best performing swarms.

5.1.1 Average Behavioral Diversity

Figure 5.1 presents the average (over 20 runs) number of distinct behaviors evolved by each approach in each environment. This was determined by counting the number of *LocalMap* cells that were occupied at the end of each run (for mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M).

Table 5.1: Statistical comparisons of unique behaviors in each environment.

<i>Simple environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.533600	0.007991	0.007725	0.006012
mEDEA-M		0.095910	0.013550	0.045020
EDQD			0.197000	0.613100
EDQD-M				0.454900
<i>Medium environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.841400	0.972000	0.005123	0.004662
mEDEA-M		0.834400	0.040910	0.007893
EDQD			0.018000	0.006758
EDQD-M				0.9544
<i>Difficult environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.189200	0.185500	0.012660	0.0001872
mEDEA-M		0.7756	0.058670	0.002204
EDQD			0.323800	0.009830
EDQD-M				0.022410

5.1.1.1 Behavioral diversity comparison in the *simple* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.1) indicate that EDQD, EDQD-M, and Double-Map EDQD-M outperform mEDEA ($p < 0.05$), while there is no statistical significance ($p \geq 0.05$) between the QD approaches. EDQD-M and Double-Map EDQD-M also outperform mEDEA-M ($p < 0.05$), with no statistical significance between mEDEA-M and EDQD ($p \geq 0.05$).

5.1.1.2 Behavioral diversity comparison in the *medium* environment

EDQD-M and Double-Map EDQD-M outperform mEDEA, mEDEA-M, and EDQD ($p < 0.05$, table 5.1), with no statistical significance between EDQD-M and Double-Map EDQD-M ($p \geq 0.05$). There is no statistical significance between mEDEA, mEDEA-M and EDQD ($p \geq 0.05$).




Figure 5.1: Average behavioral diversity of swarms evolved by each approach in each environment

5.1.1.3 Behavioral diversity comparison in the *difficult* environment

Double-Map **EDQD-M** outperforms the other methods in the *difficult* environment ($p < 0.05$, table 5.1). **EDQD-M** outperforms **mEDEA** ($p < 0.05$), with no statistical significance between **mEDEA-M**, **EDQD** and **EDQD-M** ($p \geq 0.05$). Lastly, there is no statistical significance between **mEDEA**, **mEDEA-M** and **EDQD** ($p \geq 0.05$).

5.1.1.4 Behavioral diversity comparison: *simple* vs. *medium* vs. *difficult*

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.2) indicate that all methods evolve significantly more behavioral diversity in the *simple* and *medium* environments than the *difficult* environment ($p < 0.05$). For **mEDEA**, **mEDEA-M**, and Double-Map **EDQD-M**, the behavioral diversity that emerges in the *simple* environment is comparable to that that emerges in the *medium* environment ($p \geq 0.05$). However, **EDQD** and **EDQD-M** evolve significantly more behavioral diversity in the *simple* environment than the *medium* environment ($p < 0.05$).

Table 5.2: Statistical comparisons of unique behaviors in each approach across the three environments.

			EDQD		
			Medium	Difficult	
mEDEA					
	Medium	Difficult			
Simple	0.053000	9.159e − 05			
Medium		0.000660			
mEDEA-M					
	Medium	Difficult			
Simple	0.080110	0.000159			
Medium		0.000426			

			EDQD		
			Medium	Difficult	
Simple	0.001880	0.000116			
Medium		0.000365			
EDQD-M					
	Medium	Difficult			
Simple	0.045080	0.000109			
Medium		0.000301			
Double-Map EDQD-M					
	Medium	Difficult			
Simple	0.119800	0.001290			
Medium		0.002659			




Figure 5.2: Average distance explored by swarms evolved by each approach in each environment

5.1.2 Average Distance Explored

Figure 5.2 presents the average (over 20 runs) distance explored by swarms evolved by each approach in each environment. Explored swarm distance was calculated as the sum of the distance explored by each individual of the swarm for each run (for mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M).

Table 5.3: Statistical comparisons of distance explored in each environment.

<i>Simple</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.869500	2.67e – 05	1.907e – 06	9.537e – 06
mEDEA-M		0.001690	1.907e – 05	2.67e – 05
EDQD			0.02958	0.002712
EDQD-M				0.840800
<i>Medium</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.048440	0.039990	9.537e – 06	5.722e – 06
mEDEA-M		0.8983	0.001986	0.001017
EDQD			9.537e – 06	5.722e – 06
EDQD-M				0.8124
<i>Difficult</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.000261	0.058260	1.907e – 06	1.907e – 06
mEDEA-M		0.017180	0.089690	0.010690
EDQD			0.000261	0.0001335
EDQD-M				0.728500

5.1.2.1 Distance explored comparison in the *simple* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.3) indicate that the QD approaches outperform mEDEA and mEDEA-M ($p < 0.05$, table 5.3), while there is no statistical significance between mEDEA and mEDEA-M ($p \geq 0.05$). EDQD and Double-Map EDQD-M outperform EDQD ($p < 0.05$), while there is no statistical significance between EDQD-M and Double-Map EDQD-M ($p \geq 0.05$).

5.1.2.2 Distance explored comparison in the *medium* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.3) indicate that all the other methods outperform mEDEA ($p < 0.05$). EDQD-M and Double-Map EDQD-M outperform mEDEA-M and EDQD, while there is no statistical significance between mEDEA-M and EDQD ($p \geq 0.05$). There is no statistical significance between EDQD-M and Double-Map EDQD-M ($p \geq 0.05$).

5.1.2.3 Distance explored comparison in the *difficult* environment

Double-Map EDQD-M outperforms mEDEA, mEDEA-M, and EDQD ($p < 0.05$, table 5.3), while there is no statistical significance between Double-Map EDQD-M and EDQD-M ($p \geq 0.05$). EDQD-M outperforms mEDEA and EDQD ($p < 0.05$), while there is no statistical significance between mEDEA-M and EDQD-M ($p \geq 0.05$). mEDEA-M outperforms EDQD and mEDEA, while there is no statistical significance between EDQD and mEDEA ($p \geq 0.05$).

5.1.2.4 Distance explored comparison: *simple* vs. *medium* vs. *difficult*

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.4) indicate that the QD evolved swarms explore significantly greater ($p < 0.05$) portions of the *simple* environment, followed by the *medium*, and then the *difficult* environment. mEDEA and mEDEA-M evolved swarms explore significantly greater portions of the *simple* and *medium* environments compared to the *difficult* environment. There is, however, no significant difference between the distance explored in the *simple* and *medium* environments ($p \geq 0.05$).

Table 5.4: Statistical comparisons of distance explored in each approach across the three environments.

mEDEA		
	Medium	Difficult
Simple	0.123100	5.722e – 06
Medium		0.001209
mEDEA-M		
	Medium	Difficult
Simple	0.621500	0.005581
Medium		0.009436

EDQD		
	Medium	Difficult
Simple	3.815e – 06	1.907e – 06
Medium		9.537e – 06
EDQD-M		
	Medium	Difficult
Simple	0.048440	9.537e – 06
Medium		0.001690
Double-Map EDQD-M		
	Medium	Difficult
Simple	0.001017	5.722e – 06
Medium		1.907e – 06




(a) *Simple* task environment.(b) *Medium* task environment.(c) *Difficult* task environment.

Figure 5.3: Resources collected by swarms evolved by each approach in each task environment.

5.1.3 Average Resources Collected

To further elucidate how behavioral diversity evolved in the swarms, figure 5.3 shows the various resources collected by swarms evolved by each approach in each task environment. The average number of resources collected for each resource type was calculated as the sum of resources (for each resource-type) pushed

into the target-area at the end of each run over 20 runs.

5.1.3.1 Type A resources

Double-Map [EDQD-M](#) outperforms all the other methods in the *simple* environment ($p < 0.05$, table 5.5), while there is no statistical significance between [mEDEA-M](#), [EDQD](#), and [EDQD-M](#) ($p \geq 0.05$). [mEDEA-M](#), [EDQD](#), and [EDQD-M](#) outperform [mEDEA](#) ($p < 0.05$).

In the *medium* environment, Double-Map [EDQD-M](#) outperforms [mEDEA](#), [mEDEA-M](#), and [EDQD](#) ($p < 0.05$, table 5.5), while there is no statistical significance between Double-Map [EDQD-M](#) and [EDQD-M](#) ($p \geq 0.05$). [EDQD-M](#) outperforms [mEDEA](#) and [EDQD](#), while there is no statistical significance between [EDQD-M](#) and [mEDEA-M](#) ($p \geq 0.05$). There is no statistical significance between [mEDEA](#), [mEDEA-M](#), and [EDQD](#) ($p \geq 0.05$).

In the *difficult* environment, Double-Map [EDQD-M](#) outperforms [mEDEA](#), [EDQD](#), and [EDQD-M](#) ($p < 0.05$, table 5.5), while there is no statistical significance between Double-Map [EDQD-M](#) and [mEDEA-M](#) ($p \geq 0.05$). Otherwise there is no statistical significance between the other methods ($p \geq 0.05$).

Table 5.5: Statistical comparisons of type-A resources collected in each environment.

<i>Simple</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.032410	0.009330	0.016470	0.000136
mEDEA-M		0.423800	0.540500	0.000525
EDQD			0.871500	0.000877
EDQD-M				0.000662
<i>Medium</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.139100	0.469400	0.002100	0.001928
mEDEA-M		0.515000	0.154700	0.023180
EDQD			0.003906	0.002561
EDQD-M				0.701500
<i>Difficult</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.307200	0.490500	0.459600	0.006808
mEDEA-M		0.397700	0.661100	0.142800
EDQD			0.667500	0.008309
EDQD-M				0.031660

5.1.3.2 Type B resources

In the *simple* environment, [EDQD](#) and [EDQD-M](#) outperform [mEDEA](#) ($p < 0.05$, table 5.6), while there is no statistical significance between [mEDEA](#), [mEDEA-M](#), and Double-Map [EDQD-M](#) ($p \geq 0.05$). [EDQD](#)

outperforms **mEDEA-M** ($p < 0.05$), while there is no statistical significance between **mEDEA-M**, **EDQD-M**, and Double-Map **EDQD-M** ($p \geq 0.05$). There is no statistical significance between the **QD** approaches ($p \geq 0.05$).

In the *medium* environment, **EDQD-M** outperforms **mEDEA**, **mEDEA-M**, and **EDQD** ($p < 0.05$, table 5.6), while there is no statistical significance between **EDQD-M** and Double-Map **EDQD-M**. The **QD** approaches outperform **mEDEA**, while there is no statistical significance between **mEDEA** and **mEDEA-M** ($p \geq 0.05$). There is no statistical significance between all the methods in the *difficult* environment ($p \geq 0.05$).

Table 5.6: Statistical comparisons of type-B resources collected in each environment.

<i>Simple</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.500600	0.011280	0.026490	0.212200
mEDEA-M		0.023290	0.095950	0.575700
EDQD			0.646300	0.103300
EDQD-M				0.221800
<i>Medium</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.611300	0.000944	0.001388	0.006443
mEDEA-M		0.084490	0.014950	0.087060
EDQD			0.049220	0.454300
EDQD-M				0.294700
<i>Difficult</i> environment				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.177500	0.853200	0.144300	0.259400
mEDEA-M		0.365700	0.874800	0.772300
EDQD			0.296900	0.561700
EDQD-M				0.743000

5.1.3.3 Type C resources

In the *simple* environment, **EDQD** outperforms **mEDEA-M** ($p < 0.05$, table 5.7), while there is no statistical significance between **mEDEA**, **EDQD**, **EDQD-M**, and Double-Map **EDQD-M** ($p \geq 0.05$). There is no statistical significance between all the methods in the *medium* and *difficult* environments ($p \geq 0.05$).

Table 5.7: Statistical comparisons of type-C resources collected in each environment.

<i>Simple</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.292700	0.244300	0.504800		0.968400
mEDEA-M		0.042490	0.237000		0.521800
EDQD			0.949300		0.321300
EDQD-M					0.416600
<i>Medium</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.854500	0.631100	0.081380		0.979200
mEDEA-M		1	0.174600		0.848800
EDQD			0.194900		0.734800
EDQD-M					0.142600
<i>Difficult</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.543900	0.555800	0.156800		0.759600
mEDEA-M		0.818200	0.585300		0.407000
EDQD			0.296900		0.464400
EDQD-M					0.081060

5.1.3.4 Type *D* resources

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.8) indicate that there is no statistical significance between all the approaches in the *simple* and *difficult* environments ($p \geq 0.05$).

In the *medium* environment, [EDQD](#) outperforms Double-Map [EDQD-M](#) ($p < 0.05$, table 5.8), while there is no statistical significance between [mEDEA](#), [mEDEA-M](#), [EDQD](#), and [EDQD-M](#) ($p \geq 0.05$).

Table 5.8: Statistical comparisons of type-D resources collected in each environment.

<i>Simple</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.726700	0.238600	0.750200		0.366900
mEDEA-M		0.394900	0.815700		0.478000
EDQD			0.610200		0.073590
EDQD-M					0.165100
<i>Medium</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.772200	0.183600	0.346700		0.322100
mEDEA-M		0.407500	0.688800		0.389100
EDQD			0.612900		0.032770
EDQD-M					0.056990
<i>Difficult</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	1	0.587700	0.831900		0.789700
mEDEA-M		0.565300	0.793100		0.835300
EDQD			0.675600		0.482100
EDQD-M					0.887000

5.1.3.5 Type E resources

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.9) indicate that in the *simple* environment **EDQD** outperforms **mEDEA** ($p < 0.05$), while there is no statistical significance between **mEDEA**, **mEDEA-M**, **EDQD-M**, and Double-Map **EDQD-M** ($p \geq 0.05$). **EDQD-M** outperforms **mEDEA-M** ($p < 0.05$), while there is no statistical significance between **mEDEA-M**, **EDQD**, and Double-Map **EDQD-M** ($p \geq 0.05$). There is no statistical significance between the **QD** approaches ($p \geq 0.05$).

There is no statistical significance between all the methods in the *medium* environment ($p \geq 0.05$).

In the *difficult* environment, Double-Map **EDQD-M** and **EDQD** outperform **mEDEA-M** ($p < 0.05$, table 5.9), while there is no statistical significance between **mEDEA**, **mEDEA-M**, and **EDQD-M**. There is no statistical significance between the **QD** approaches ($p \geq 0.05$).

Table 5.9: Statistical comparisons of type-E resources collected in each environment.

<i>Simple</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.850100	0.033010	0.058150		0.851400
mEDEA-M		0.055620	0.020960		0.667500
EDQD			0.915500		0.077070
EDQD-M					0.089960
<i>Medium</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.873700	0.440200	0.071530		0.930300
mEDEA-M		0.592600	0.275100		0.923300
EDQD			0.475500		0.746000
EDQD-M					0.159000
<i>Difficult</i> environment					
	mEDEA-M	EDQD	EDQD-M	Double-Map	EDQD-M
mEDEA	0.161400	0.218400	0.526100		0.401100
mEDEA-M		0.018120	0.161500		0.032030
EDQD			0.462100		1
EDQD-M					0.714300

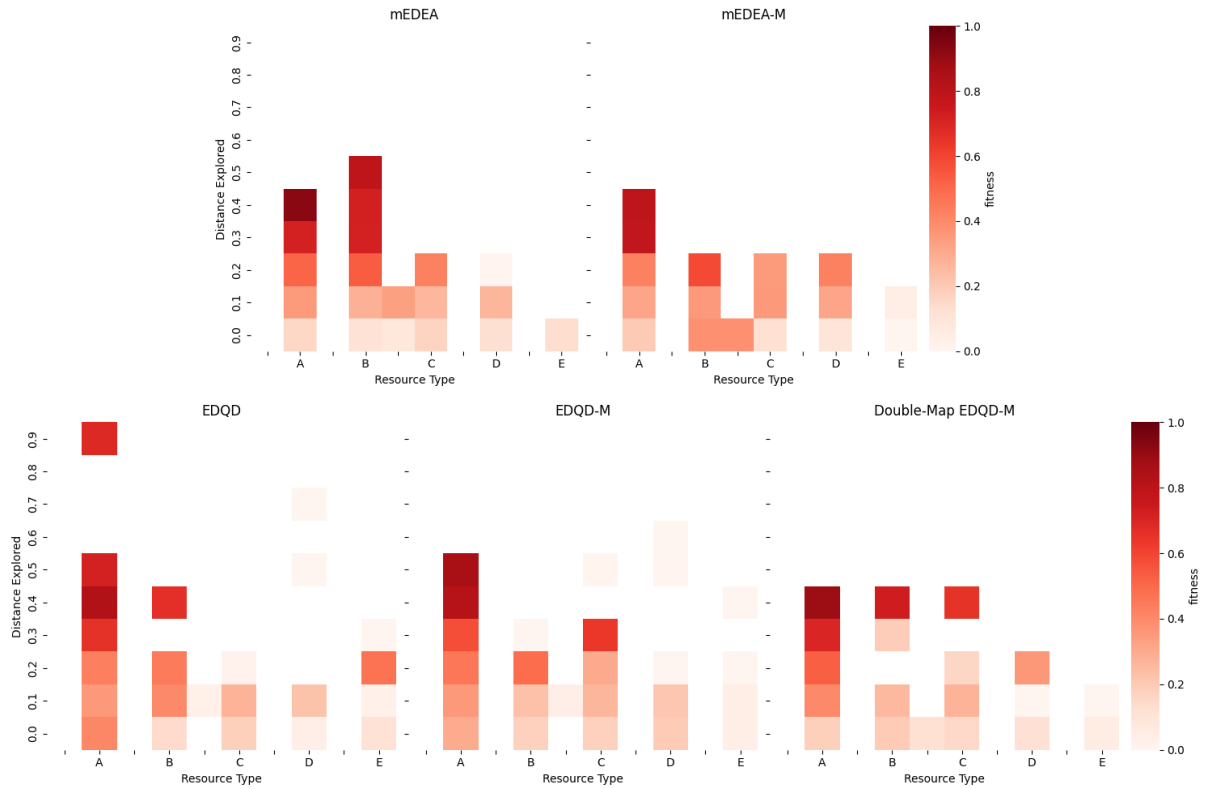


Figure 5.4: Map-archive of robot behaviors discovered by the best-performing swarm of each approach in the *simple* environment over a single run (at the final generation), colored based on fitness.

5.1.4 Behavior Map-archives of Best-Performing Swarms

At the end of each run (table 4.2), QD behavioral maps (figures 5.4, 5.5, 5.6) evolved by the highest performing swarm provided an indication of swarm behavior and the quality of such behaviors for gathering each resource type in each task environment. The Euclidean distance traversed and resource type gathered (transported) to the target-area were used to measure swarm behavior (and thus diversity). There is a task performance value that corresponds to the solution discovered in each map cell (section 4.3.2, equation 4.1). Hence, such QD maps show how efficient evolved gathering behaviors are for various task difficulties (*simple*, *medium*, *difficult*).

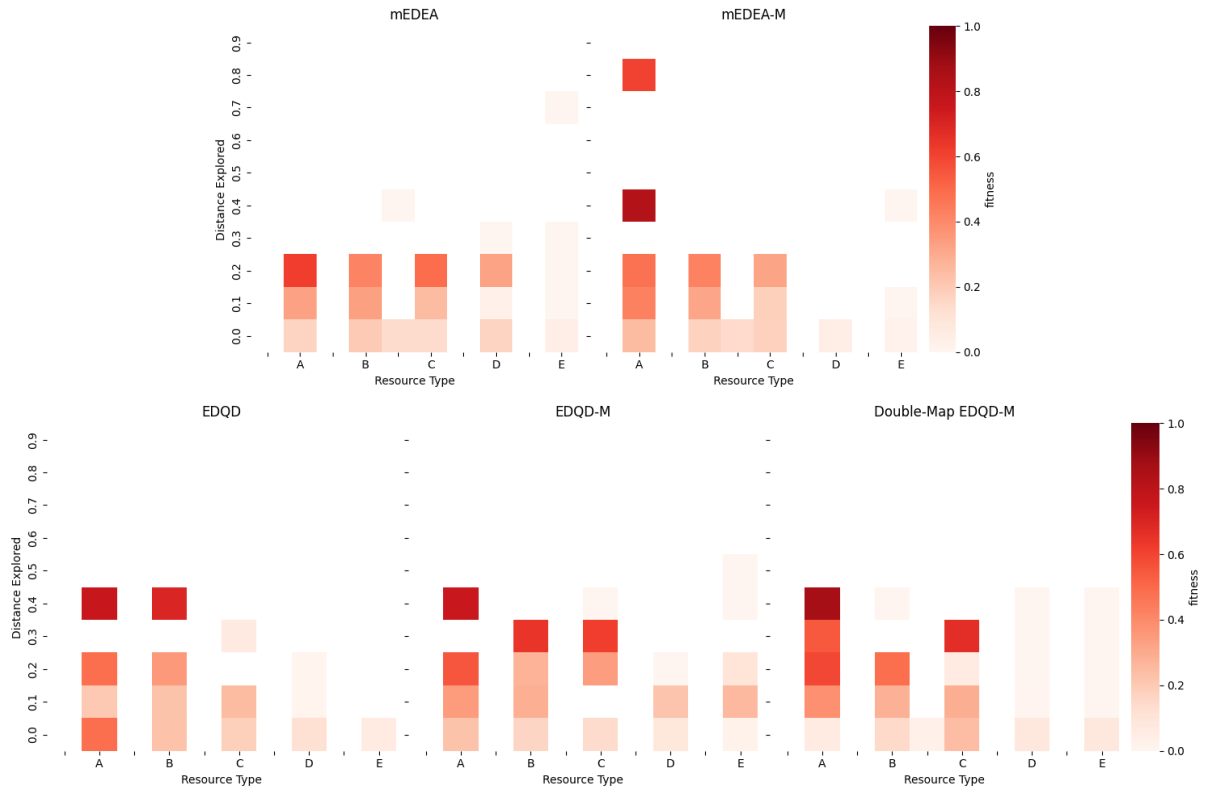


Figure 5.5: Map-archive of robot behaviors discovered by the best-performing swarm of each approach in the *medium* environment over a single run (at the final generation), colored based on fitness.

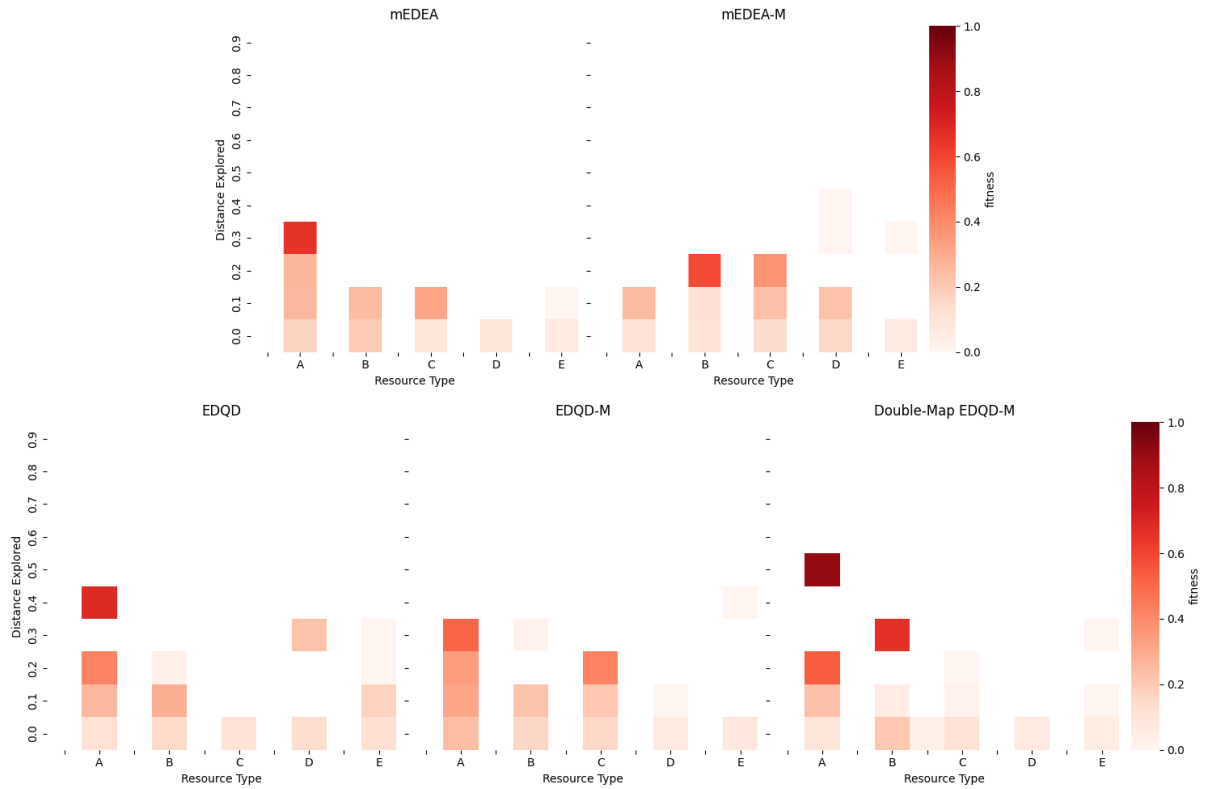


Figure 5.6: Map-archive of robot behaviors discovered by the best-performing swarm of each approach in the *difficult* environment over a single run (at the final generation), colored based on fitness.




Figure 5.7: Average morphological diversity of swarms evolved by the [mEDEA-M](#), [EDQD-M](#), and Double-Map [EDQD-M](#) approaches in each environment

5.2 Quality Diversity (QD) of Evolved Morphologies

The morphological map-archive is defined by two dimensions, namely, the proportion of active sensors and the average range of active sensors. This means that in order for unique morphologies to emerge (and therefore fill up the map-archive), each approach must use the whole spectrum of the sensor range (range: $[0;1]$) as well as a variable number of sensors (that is, switch on and off some of its sensors).

The remainder of this section is divided as follows: Average Morphological Diversity (section 5.2.1), Average Range of Active Sensors (section 5.2.2), Average Proportion of Active Sensors (section 5.2.3), and Morphology Map-archive of Best Performing Swarms (section 5.2.5). Section 5.2.1 presents comparative method average morphological diversity. Sections 5.2.2 and 5.2.3 consider the results of each method in each dimension of the map-archive in order to gain a deeper understanding of how morphological diversity evolved. Lastly, section 5.2.5 compares method behavior map-archives of the best performing swarms.

Table 5.10: Statistical comparisons of unique morphologies in each environment.

<i>Simple environment</i>			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.17	$9.33\text{e} - 05$	
EDQD-M		$9.356\text{e} - 05$	
<i>Medium environment</i>			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.628800	$9.316\text{e} - 05$	
EDQD-M		$9.251\text{e} - 05$	
<i>Difficult environment</i>			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.903500	$9.33\text{e} - 05$	
EDQD-M		$9.343\text{e} - 05$	

5.2.1 Average Morphological Diversity

Figure 5.7 presents the average (over 20 runs) number of unique morphologies evolved by each approach in each environment. This was determined by counting the number of *LocalMap* cells that were occupied at the end of each run (for mEDEA-M, EDQD-M, and Double-Map EDQD-M).

5.2.1.1 Morphological diversity comparison in the *simple* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.10) indicate that the average morphological diversity of Double-Map EDQD-M evolved swarms was significantly higher than mEDEA-M and EDQD-M evolved swarms ($p < 0.05$), while there is no statistical significance between mEDEA-M and EDQD-M ($p \geq 0.05$).

5.2.1.2 Morphological diversity comparison in the *medium* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.10) indicate that Double-Map EDQD-M outperforms mEDEA-M and EDQD-M ($p < 0.05$), while there is no statistical significance between mEDEA-M and EDQD-M ($p \geq 0.05$).

5.2.1.3 Morphological diversity comparison in the *difficult* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.10) indicate that Double-Map EDQD-M outperforms mEDEA-M and EDQD-M ($p < 0.05$), while there is no statistical significance between mEDEA-M and EDQD-M ($p \geq 0.05$).

5.2.1.4 Morphological diversity comparison: *simple* vs. *medium* vs. *difficult*

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.11) indicate that there is no significant difference in the diversity that evolves in the three environments for mEDEA-M and EDQD-M ($p \geq 0.05$). However, Double-Map EDQD-M evolves significantly ($p < 0.05$) more diverse morphologies in the *medium* environment, followed by the *simple* environment, and then the *difficult* environment ($p < 0.05$).




Figure 5.8: Average range of active sensors (mean of the population) at the end of each treatment

Table 5.11: Statistical comparisons of average morphological diversity in each approach across the three environments.

mEDEA-M		
	Medium	Difficult
Simple	0.067220	0.192600
Medium		0.497400
EDQD-M		
	Medium	Difficult
Simple	0.331200	0.795200
Medium		0.175900
Double-Map EDQD-M		
	Medium	Difficult
Simple	0.00276	0.006328
Medium		0.000109

5.2.2 Average Range of Active Sensors

Figure 5.8 presents the average (over 20 runs) range of active sensors evolved by each approach in each environment. The average range of active sensors was calculated at the end of each run (final generation) over 20 runs (for mEDEA-M, EDQD-M, and Double-Map EDQD-M).

Figure 5.9 shows mEDEA-M, EDQD-M, and Double-Map EDQD-M learning curves. Despite the fact that the final average range of the three approaches is comparable (sections 5.2.2.1, 5.2.2.2, and 5.2.2.3), the learning curves show that Double-Map EDQD-M tends to decrease quickly and then settle at about $\approx 50\%$ of the total range. However, mEDEA-M and EDQD-M both decline gradually before stabilizing within the same average range.




Figure 5.9: Average (over 20 runs) range of active sensors calculated from [mEDEA-M](#), [EDQD-M](#), and Double-Map [EDQD-M](#) at the end of each generation (for each run) for swarms evolved in the *simple*, *medium* and *difficult* environments.

Table 5.12: Statistical comparisons of average range of active sensors in each environment.

<i>Simple environment</i>		
	EDQD-M	Double-Map EDQD-M
mEDEA-M	0.8124	0.8408
EDQD-M		0.5459
<i>Medium environment</i>		
	EDQD-M	Double-Map EDQD-M
mEDEA-M	0.08255	0.08255
EDQD-M		0.2774
<i>Difficult environment</i>		
	EDQD-M	Double-Map EDQD-M
mEDEA-M	0.6742	0.4524
EDQD-M		0.8124

5.2.2.1 Sensor range comparison in the *simple* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.12) indicate that there is no statistical significance between [mEDEA-M](#), [EDQD-M](#), and Double-Map [EDQD-M](#) ($p \geq 0.05$).

5.2.2.2 Sensor range comparison in the *medium* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.12) indicate that there is no statistical significance between [mEDEA-M](#), [EDQD-M](#), and Double-Map [EDQD-M](#) ($p \geq 0.05$).

5.2.2.3 Sensor range comparison in the *difficult* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.12) indicate that there is no statistical significance between **mEDEA-M**, **EDQD-M**, and Double-Map **EDQD-M** ($p \geq 0.05$).

5.2.2.4 Sensor range comparison: *simple* vs. *medium* vs. *difficult*

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.13) indicate that there is no statistical difference between the *simple*, *medium*, and *difficult* environments for **mEDEA-M**, **EDQD-M**, and Double-Map **EDQD-M** ($p \geq 0.05$).

Table 5.13: Statistical comparisons of average range in each approach across the three environments.

mEDEA-M		
	Medium	Difficult
Simple	0.132700	0.701200
Medium		0.257500
EDQD-M		
	Medium	Difficult
Simple	0.230500	0.756200
Medium		0.294300
Double-Map EDQD-M		
	Medium	Difficult
Simple	0.701200	0.812400
Medium		0.430400

5.2.3 Average Proportion of Active Sensors

Figure 5.10 presents the average (over 20 runs) proportion of active sensors evolved by each approach in each environment. The average proportion of active sensors was calculated at the end of each run (final generation) over 20 runs (for **mEDEA-M**, **EDQD-M**, and Double-Map **EDQD-M**).

Figure 5.11 depicts the learning curves for **mEDEA-M**, **EDQD-M**, and Double-Map **EDQD-M** in relation to the proportion of active sensors. These learning curves show the same tendency as the evolution of the average range of active sensors (section 5.2.2), but slightly more apparent. That is, Double-Map **EDQD-M** tends to decrease quickly and then settle at about $\approx 60\%$ of the total number of sensors however, **mEDEA-M** and **EDQD-M** both decline gradually before stabilizing within the same average proportion.




Figure 5.10: Average proportion of active sensors (mean of the population) at the end of each treatment




Figure 5.11: Average (over 20 runs) proportion of active sensors calculated from **mEDEA-M**, **EDQD-M**, and Double-Map **EDQD-M** at the end of each generation (for each run) for swarms evolved in the *simple*, *medium*, and *difficult* environments.

Table 5.14: Statistical comparisons of proportion of active sensors in each environment.

<i>Simple environment</i>		
	EDQD-M	Double-Map EDQD-M
mEDEA-M	0.6861	0.3488
EDQD-M		1
<i>Medium environment</i>		
	EDQD-M	Double-Map EDQD-M
mEDEA-M	0.7548	0.06372
EDQD-M		0.05826
<i>Difficult environment</i>		
	EDQD-M	Double-Map EDQD-M
mEDEA-M	0.465400	0.01718
EDQD-M		0.33

5.2.3.1 Proportion of active sensors comparison in the *simple* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.14) indicate that there is no statistical significance between mEDEA-M, EDQD-M, and Double-Map EDQD-M ($p \geq 0.05$).

5.2.3.2 Proportion of active sensors comparison in the *medium* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.14) indicate that there is no statistical significance between mEDEA-M, EDQD-M, and Double-Map EDQD-M ($p \geq 0.05$).

5.2.3.3 Proportion of active sensors comparison in the *difficult* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.14) indicate that mEDEA-M outperforms Double-Map EDQD-M ($p < 0.05$), while there is no statistical significance between Double-Map EDQD-M and EDQD-M ($p \geq 0.05$). There is no statistical significance between mEDEA-M and EDQD-M ($p \geq 0.05$).

5.2.3.4 Proportion of active sensors comparison: *simple* vs. *medium* vs. *difficult*

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.15) indicate that there is no statistical difference between the *simple*, *medium*, and *difficult* environments for mEDEA-M, EDQD-M, and Double-Map EDQD-M ($p \geq 0.05$).

Table 5.15: Statistical comparisons of proportion of active sensors in each approach across the three environments.

mEDEA-M		
	Medium	Difficult
Simple	0.420100	0.078020
Medium		0.340500
EDQD-M		
	Medium	Difficult
Simple	0.101100	0.468900
Medium		0.474900
Double-Map EDQD-M		
	Medium	Difficult
Simple	0.595800	0.654100
Medium		0.311800




Figure 5.12: Average morphological complexity (mean of the population) at the end of each treatment

5.2.4 Morphological Complexity

Figure 5.12 presents the average (over 20 runs) morphological complexity evolved by each approach in each environment. The average morphological complexity was calculated at the end of each run (final generation) over 20 runs (for mEDEA-M, EDQD-M, and Double-Map EDQD-M).

Morphological complexity is calculated from sensor configurations (that is, average sensor range and proportion of active sensors) according to equation 5.1:

$$M_c = \frac{1}{2}(S_R + S_A) \quad (5.1)$$

where M_c is the morphological complexity, S_R is the normalized average sensor range, and S_A is the proportion of active sensors. Since S_R and S_A fall in the range $[0;1]$, we multiply by $\frac{1}{2}$ to normalize M_c to the range $[0;1]$. A high value of M_c indicates a high degree of morphological complexity.

Table 5.16: Statistical comparisons of morphological complexity in each environment.

<i>Simple environment</i>			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.8695		0.5706
EDQD-M			0.8124
<i>Medium environment</i>			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.6215		0.02958
EDQD-M			0.3884
<i>Difficult environment</i>			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.5459		0.01208
EDQD-M			0.5217

5.2.4.1 Morphological complexity comparison in the *simple* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.16) indicate that there is no statistical significance between mEDEA-M, EDQD-M, and Double-Map EDQD-M ($p \geq 0.05$).

5.2.4.2 Morphological complexity comparison in the *medium* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.16) indicate that Double-Map EDQD-M evolves the least complex morphologies compared to mEDEA-M ($p < 0.05$), while there is no statistical significance between Double-Map EDQD-M and EDQD-M ($p \geq 0.05$).

5.2.4.3 Morphological complexity comparison in the *difficult* environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.16) indicate that Double-Map EDQD-M evolves the least complex morphologies compared to mEDEA-M ($p < 0.05$), while there is no statistical significance between Double-Map EDQD-M and EDQD-M ($p \geq 0.05$).

5.2.4.4 Morphological complexity comparison: *simple* vs. *medium* vs. *difficult*

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.17) indicate that mEDEA-M evolves significantly more complex morphologies in the *difficult* environment compared to the *simple* environment ($p < 0.05$), while there is no statistical significance between the *simple* and *medium* environments, and between the *medium* and *difficult* environments. There is no statistical significance in the morphological complexity that evolves in the *simple*, *medium*, and *difficult* environments for EDQD-M and Double-Map EDQD-M ($p \geq 0.05$).

Table 5.17: Statistical comparisons of morphological complexity in each approach across the three environments.

mEDEA-M		
	Medium	Difficult
Simple	0.2162	0.03623
Medium		0.5713
EDQD-M		
	Medium	Difficult
Simple	0.8983	0.9854
Medium		0.9563
Double-Map EDQD-M		
	Medium	Difficult
Simple	0.8983	0.6215
Medium		0.4524

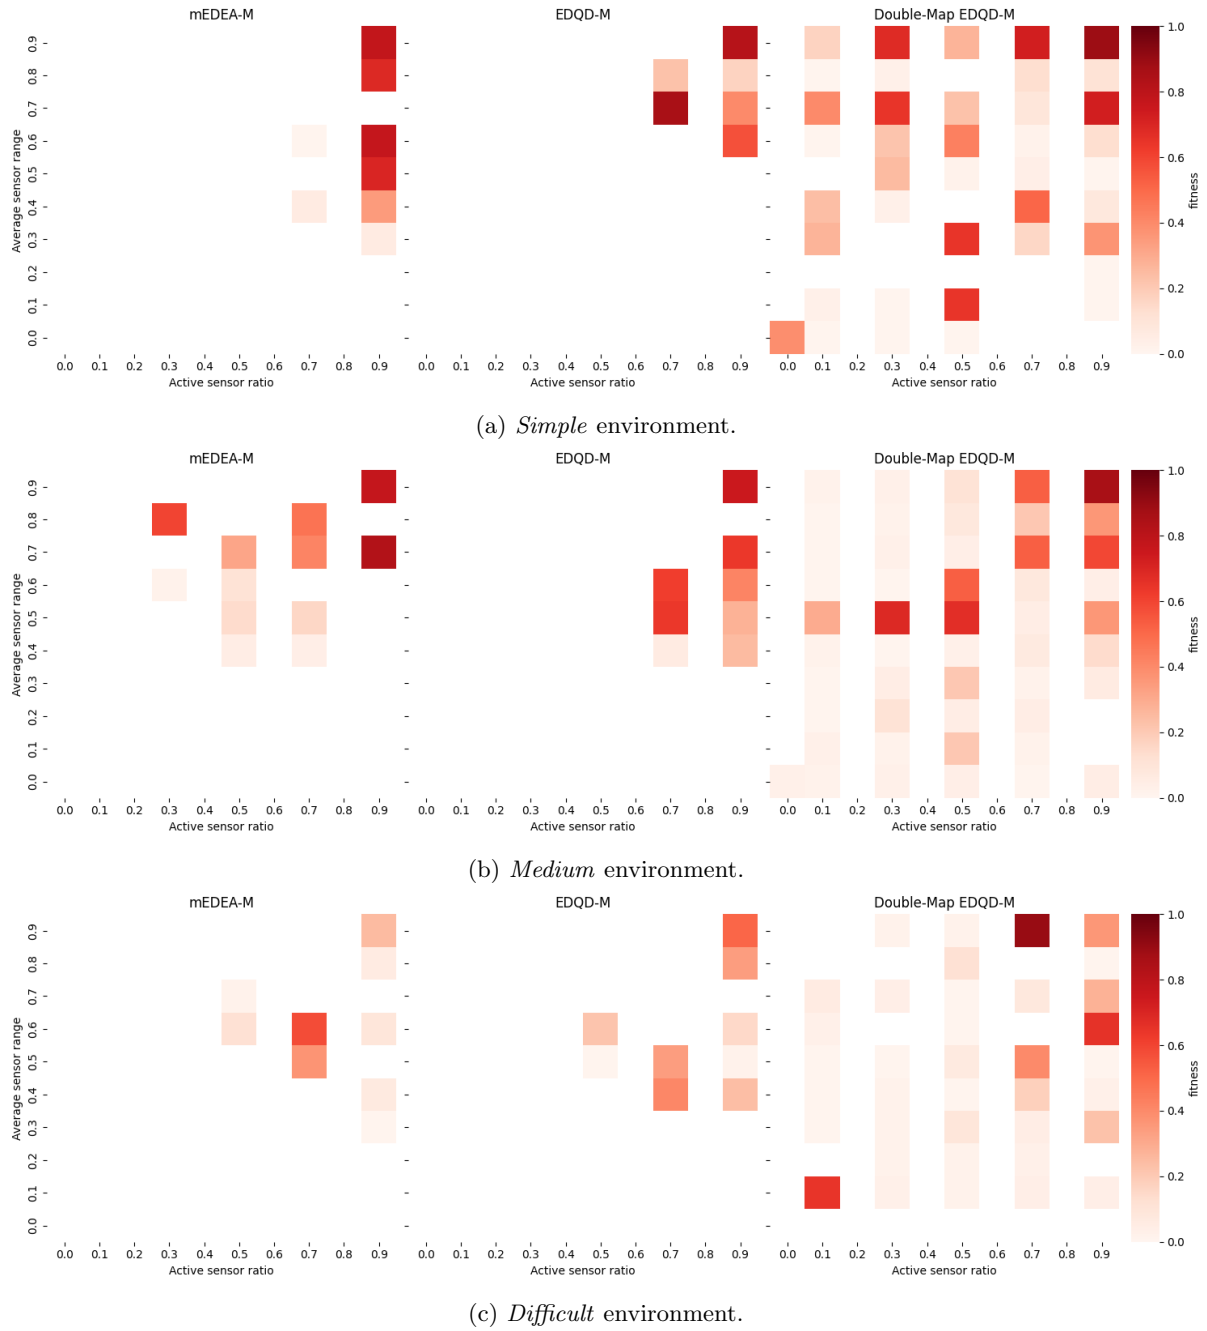


Figure 5.13: Map-archive of robot morphologies evolved by the best-performing swarm of each approach in each environment over a single run (at the final generation), colored based on fitness.

5.2.5 Morphology Map-archive of Best-Performing Swarms

At the end of each run (table 4.2), morphology map-archives (figures 5.4, 5.5, 5.6) evolved by the highest performing swarm of each approach provided an indication of the highest quality morphologies (sensory configurations) evolved in each cell of the map-archive. Morphological descriptors were defined as the average sensor range and proportion of active sensors. There is a task performance value that corresponds to the solution discovered in each map cell (section 4.3.2, equation 4.1). Hence, such maps show how efficient evolved morphologies are for various task difficulties (*simple*, *medium*, *difficult*).




Figure 5.14: Average task performance of swarms evolved by the [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) approaches in each environment.

5.3 Swarm Task Performance

Average team fitness is calculated as the portion of resources (range: $[0.0, 1.0]$) gathered during all swarm life-times (100 generations per run), averaged over 20 runs (figure 5.14).

Table 5.18: Statistical comparisons of task performance of robot teams evolved by the [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) approaches in each task environment.

<i>Simple environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.075620	0.000667	0.002973	0.000444
mEDEA-M		0.017400	0.047350	0.001569
EDQD			0.532400	0.358100
EDQD-M				0.156400
<i>Medium environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.186900	0.036140	0.000391	0.022540
mEDEA-M		0.419700	0.007357	0.197300
EDQD			0.008313	0.678300
EDQD-M				0.024750
<i>Difficult environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.871900	0.266100	0.099330	0.078670
mEDEA-M		0.913200	0.330900	0.204000
EDQD			0.265000	0.169600
EDQD-M				0.856100

5.3.1 Task Performance Comparison in the *Simple* Environment

Statistical tests indicate that **EDQD** and Double-Map **EDQD-M** outperform **mEDEA** and **mEDEA-M** (Mann-Whitney U, $p < 0.05$, table 5.18). **EDQD-M** outperforms **mEDEA** ($p < 0.05$), while there is no statistical significance between **mEDEA-M** and **EDQD-M** ($p \geq 0.05$). There is no statistical significance between **EDQD**, **EDQD-M** and Double-Map **EDQD-M** ($p \geq 0.05$).

5.3.2 Task Performance Comparison in the *Medium* Environment

EDQD-M outperforms all the other methods (Mann-Whitney U, $p < 0.05$, table 5.18). **EDQD** and Double-Map **EDQD-M** outperform **mEDEA** ($p < 0.05$), while there is no statistical significance between **mEDEA** and **mEDEA-M** ($p \geq 0.05$). There is no statistical significance between **mEDEA-M**, **EDQD** and Double-Map **EDQD-M** ($p \geq 0.05$).

5.3.3 Task Performance Comparison in the *Difficult* Environment

Statistical tests (Mann-Whitney U-tests, table 5.18) indicate that there is no statistical significance between all the methods in the *difficult* environment ($p \geq 0.05$).

5.3.4 Task Performance Comparison: *Simple* vs. *Medium* vs. *Difficult* Environment

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.19) indicate that robotic swarms evolved by all methods in the *simple* environment performed significantly better, followed by swarms evolved in the *medium* environment, and then the *difficult* environment (Mann-Whitney U, $p < 0.05$, table 5.19).

Table 5.19: Statistical comparisons of task performance in each approach across the three environments.

			EDQD	
			Medium	Difficult
mEDEA				
	Medium	Difficult		
Simple	9.556e – 05	9.556e – 05	Simple	9.542e – 05
Medium		0.000300	Medium	0.000129
mEDEA-M				
	Medium	Difficult		
Simple	9.529e – 05	9.542e – 05	Simple	9.516e – 05
Medium		0.000142	Medium	9.475e – 05
			EDQD-M	
			Medium	Difficult
Simple			Simple	9.475e – 05
Medium			Medium	9.436e – 05
			Double-Map EDQD-M	
			Medium	Difficult
Simple			Simple	9.422e – 05
Medium			Medium	9.449e05
				9.475e – 05




Figure 5.15: QD score calculated from mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M final behavior-maps (at the end of each run) for swarms evolved in the *simple*, *medium*, and *difficult* environments.




Figure 5.16: Average (over 20 runs) QD score calculated from mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M behavior-maps at the end of each generation (for each run) for swarms evolved in the *simple*, *medium*, and *difficult* environments.

5.4 Quality Diversity Score of Behavior Archives

As in previous work (Pugh et al., 2015), the QD score is calculated as the overall quality (section 4.3.2, equation 4.1) across all filled grid-cells within the QD behavior map-archive. A high average QD score thus implies swarms with a high level of behavioral diversity, as well as a high level of quality. Figure 5.15 depicts the average quality diversity score of mEDEA, mEDEA-M, EDQD, EDQD-M, and Double-Map EDQD-M at the end of each run, where an average QD score is obtained for the 20 runs of each approach.

Figure 5.16 shows the average variation of the QD score over the course of 20 runs of each treatment. As task difficulty increases, we see the benefits of Double-Map EDQD-M. That is, Double-Map EDQD-M produces the lowest average QD score after 100 evaluations (generations) in the *simple* environment. Double-Map EDQD-M has the second lowest average QD score in the *medium* environment. In the *difficult* environment, however, Double-Map EDQD-M produces the highest average QD score (figure 5.16).

This evidence supports the benefits of increased behavioral (and morphological) diversity afforded by Double-Map [EDQD-M](#), and the [QD](#) maps for evolved behaviors (section 5.1) and morphologies (section 5.2) provide additional evidence for the method’s suitability for evolving swarms in environments with increased levels of difficulty.

Table 5.20: Statistical comparisons of quality diversity score of robot teams evolved by the [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) approaches in each task environment.

<i>Simple environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.5958	0.03277	0.003153	0.3683
mEDEA-M		0.3488	0.08969	0.9854
EDQD			0.2162	0.2774
EDQD-M				0.02958
<i>Medium environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.9563	0.5217	0.02395	0.4524
mEDEA-M		0.7841	0.1231	0.5706
EDQD			0.002712	0.4091
EDQD-M				0.03999
<i>Difficult environment</i>				
	mEDEA-M	EDQD	EDQD-M	Double-Map EDQD-M
mEDEA	0.498	0.7562	0.6477	0.01069
mEDEA-M		0.8695	0.6215	0.2611
EDQD			0.8695	0.1327
EDQD-M				0.1893

5.4.1 Quality Diversity Score Comparison in the *Simple* Environment

Statistical tests indicate that [EDQD](#) and [EDQD-M](#) outperform [mEDEA](#) (Mann-Whitney U, $p < 0.05$, table 5.20), while there is no statistical significance between [mEDEA-M](#), [EDQD](#), and [EDQD-M](#) ($p \geq 0.05$). [EDQD-M](#) outperforms Double-Map [EDQD-M](#) ($p < 0.05$), while there is no statistical difference between [mEDEA](#), [mEDEA-M](#), and Double-Map [EDQD-M](#) ($p \geq 0.05$).

5.4.2 Quality Diversity Score Comparison in the *Medium* Environment

[EDQD-M](#) outperforms [mEDEA](#), [EDQD](#), and Double-Map [EDQD-M](#) (Mann-Whitney U, $p < 0.05$, table 5.20), while there is no statistical difference between [mEDEA-M](#) and [EDQD-M](#) ($p \geq 0.05$). There is no statistical significance between [mEDEA](#), [mEDEA-M](#), [EDQD](#), and Double-Map [EDQD-M](#) ($p \geq 0.05$).

5.4.3 Quality Diversity Score Comparison in the *Difficult* Environment

Statistical tests indicate that Double-Map [EDQD-M](#) outperforms [mEDEA](#) (Mann-Whitney U, $p < 0.05$, table 5.20), whereas there is no statistical difference between [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) ($p \geq 0.05$). There is no statistical significance between the other approaches ($p \geq 0.05$).

5.4.4 Quality Diversity Score Comparison: *Simple* vs. *Medium* vs. *Difficult*

Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.21) indicate that robotic swarms evolved by all methods in the *simple* environment performed significantly better, followed by swarms evolved in the *medium* environment, and then the *difficult* environment ($p < 0.05$).

Table 5.21: Statistical comparisons of [QD](#) score in each approach across the three environments.

			EDQD	
			Medium	Difficult
mEDEA				
	Medium	Difficult		
Simple	0.000168	1.907e – 06	Simple	1.907e – 06
Medium		3.624e – 05	Medium	3.815e – 06
mEDEA-M				
	Medium	Difficult		
Simple	1.907e – 06	1.907e – 06	Simple	2.67e – 05
Medium		1.907e – 06	Medium	1.907e – 06
			EDQD-M	
			Medium	Difficult
Simple			Simple	1.907e – 06
Medium			Medium	1.907e – 06
			Double-Map EDQD-M	
			Medium	Difficult
Simple			Simple	9.537e – 06
Medium			Medium	1.907e – 06




Figure 5.17: QD score calculated from mEDEA-M, EDQD-M, and Double-Map EDQD-M final morphology-maps (at the end of each run) for swarms evolved in the *simple*, *medium*, and *difficult* environments.




Figure 5.18: Average (over 20 runs) QD score calculated from mEDEA-M, EDQD-M, and Double-Map EDQD-M morphology-maps at the end of each generation (for each run) for swarms evolved in the *simple*, *medium*, and *difficult* environments.

5.5 Quality Diversity Score of Morphology Archives

For mEDEA-M, EDQD-M, and Double-Map EDQD-M we measured the QD score of morphology map archives as the overall quality across all filled grid-cells at the end of each run. A high average QD score thus implies swarms with a high level of morphological diversity and quality. Figure 5.17 shows the average QD across the task environments, over 20 runs. Statistical tests (Mann-Whitney U-tests (Mann & Whitney, 1947), table 5.22) indicate that Double-Map EDQD-M significantly outperformed the other two approaches in all three environments ($p < 0.05$), while there is no statistical significance between mEDEA-M and EDQD-M ($p \geq 0.05$).

Figure 5.18 shows the average variation of the QD score over the course of 20 runs of each treatment. In all three environments, Double-Map EDQD-M tends to rise sharply until the 30th generation, after which it continues rising but not as rapidly as in the first 30 generations. In comparison, the rise of mEDEA-M and EDQD-M is gradual throughout the 100 generations in all the task environments. As task complexity increases, the QD score evolved by all three approaches decreases however, Double-Map EDQD-M remains superior to other the other two approaches across all three task environments.

Table 5.22: Statistical comparisons of QD score in each environment.

<i>Simple</i> environment			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.08969	1.907e − 06	
EDQD-M		1.907e − 06	
<i>Medium</i> environment			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.2455	1.907e − 06	
EDQD-M		1.907e − 06	
<i>Difficult</i> environment			
	EDQD-M	Double-Map	EDQD-M
mEDEA-M	0.7652	9.537e − 06	
EDQD-M		0.0003893	

5.6 Summary

In this Chapter, we presented the results of evaluating [mEDEA](#), [mEDEA-M](#), [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) in relation to the diversity and quality of evolved behaviors, swarm task performance, and quality diversity score. We observed visually distinct results and conducted Mann-Whitney U tests (with a significance level of $p = 0.05$) to determine their statistical significance.

Regarding the diversity of evolved behaviors, the tests showed that in the *simple* environment, both [mEDEA](#) and [mEDEA-M](#) performed poorly and were comparable. [EDQD](#) performed better than [mEDEA](#) but was comparable to [mEDEA-M](#). [EDQD-M](#) and Double-Map [EDQD-M](#) were the highest performing approaches, outperforming [mEDEA](#) and [mEDEA-M](#) while being comparable to [EDQD](#). In the *medium* environment, [EDQD-M](#) and Double-Map [EDQD-M](#) were comparable and outperformed the other approaches. [mEDEA](#), [mEDEA-M](#) and [EDQD](#) were all comparable. In the *difficult* environment, Double-Map [EDQD-M](#) outperformed all other methods. [EDQD-M](#) outperformed [mEDEA](#), while being comparable to [mEDEA-M](#) and [EDQD](#). Otherwise, [mEDEA](#), [mEDEA-M](#), and [EDQD](#) were all comparable.

Upon further analysis of the diversity of evolved morphologies, it was found that in the *simple*, *medium*, and *difficult* environments, Double-Map [EDQD-M](#) outperformed the other two approaches. In all three environments, [mEDEA-M](#) and [EDQD-M](#) showed comparable performance.

Further statistical tests were conducted to evaluate the swarm task performance of different approaches. In the *simple* environment, [mEDEA](#) and [mEDEA-M](#) showed comparable and relatively poor performance. [EDQD-M](#) performed better than [mEDEA](#), while being comparable to [mEDEA-M](#). All the [QD](#) approaches demonstrated similar performance, but [EDQD](#) and Double-Map [EDQD-M](#) outperformed [mEDEA](#) and [mEDEA-M](#). In the *medium* environment, [mEDEA](#) and [mEDEA-M](#) exhibited comparable and relatively poor performance. [mEDEA](#) was outperformed by [EDQD](#) and Double-Map [EDQD-M](#). [mEDEA-M](#), [EDQD](#), and Double-Map [EDQD-M](#) showed comparable performance. In the *difficult* environment, all the methods displayed comparable performance.

The statistical analyses of the [QD](#) score of behavior archives revealed the following results. In the *simple* environment, [mEDEA](#), [mEDEA-M](#), and Double-Map [EDQD-M](#) showed comparable and relatively poor performance. [EDQD](#) and [EDQD-M](#) outperformed [mEDEA](#), while being comparable to [mEDEA-M](#). [EDQD](#) and Double-Map [EDQD-M](#) exhibited comparable performance, while [EDQD-M](#) outperformed Double-Map [EDQD-M](#). In the *medium* environment, [mEDEA](#), [mEDEA-M](#), [EDQD](#), and Double-Map [EDQD-M](#) demonstrated comparable performance. [EDQD-M](#) outperformed [mEDEA](#), [EDQD](#), and Double-Map [EDQD-M](#), while being comparable to [mEDEA-M](#). In the *difficult* environment, [mEDEA](#), [mEDEA-M](#), [EDQD](#), and [EDQD-M](#) showed comparable performance. Double-Map [EDQD-M](#) outperformed [mEDEA](#), while being comparable to all the other approaches.

Chapter 6

Discussion

This chapter discusses our Chapter 5 findings and offers a critical assessment of the research questions raised in this thesis. Overall, the results of this work demonstrate that environmental conditions influence the efficiency and diversity of body-brain evolved cooperative robot teams. In particular, our findings show that task performance and diversity of evolved robot teams decrease as task environment complexity is increased, where task complexity is the level of cooperation (amongst individual robots) required for optimal performance. Furthermore, the results indicate that explicit behavioral-morphological diversity maintenance is well suited for evolving collective behavior across increasingly complex task environments given that it consistently facilitated the emergence of significantly more diverse behaviors and morphologies than all the other methods in all the task environments while maintaining comparable team performance.

The remainder of the Chapter discusses the findings of this work in connection to the said research questions. Section 6.1 compares the efficiency and diversity of robot teams evolved in the *simple*, *medium*, and *difficult* environments. Section 6.2 discusses the impact of behavioral diversity maintenance on evolving swarm behavior. Section 6.3 reviews the influence of behavioral diversity maintenance on co-evolving swarm behavior and morphology. Section 6.4 details the impact of body-brain diversity maintenance on co-evolving swarm behavior and morphology. Section 6.5 provides a summary of the entire Chapter.

6.1 Impact of Environmental Conditions

Experiments (Chapter 4, table 4.3) intended to investigate whether environmental conditions (that is, varying degrees of task complexity) influence the efficiency (that is, task performance) and diversity (that is, behavioral and morphological diversity) of body-brain evolved cooperative robot teams. Our findings imply that environmental conditions influence task performance (section 6.1.1) as well as the emergence of behavioral (section 6.1.2) and morphological (6.1.3) diversity in cooperative robot teams.

6.1.1 Team Task Performance

For all methods (Chapter 3, table 3.1), the best-performing robot teams evolved in the *simple* task environment, followed by the *medium* environment, while the worst-performing teams evolved in the *difficult* environment (section 5.3.4). The rest of this section discusses the properties of each of the environments in order to understand how team performance evolved.

In the *simple* task environment, 60% of the resources were type *A* which could be collected individually without cooperation. The remaining 40% were evenly distributed between types *B*, *C*, *D*, and *E* (requiring two, three, four, and five robots to cooperatively push, respectively, table 4.2). It was expected that the highest team performance would be achieved in this environment because optimal fitness could be attained without a high level of cooperation. Additionally, with a large number of robots (100 in the group), there was sufficient concurrency in the gathering behavior, leading to near-optimal task performance.

Comparatively, in the *medium* task environment, all resource types (*A*, *B*, *C*, *D*, and *E*) were equally distributed. One-fifth of the resources could be collected by individual robots, while the rest required varying levels of cooperative effort (table 4.2). Consistent with earlier findings (Furman et al., 2019), team fitness was improved by the relative simplicity of acquiring a fifth of the resources individually, that is, without cooperation.

The lowest-performing robot teams evolved in the *difficult* task environment, where 60% of the resources were type *E*, requiring five robots to cooperatively push. The remaining 40% were evenly distributed among types *A*, *B*, *C*, and *D*. Cooperation involved a *wait-for-help* period that lasted for the robot's lifetime (table 4.2). Due to the large portion of type *E* resources, many robots spent most of their lifetime *idle* or waiting for help. It was less likely for three or more robots to converge upon the same resource within the given *wait-for-help* period. This was corroborated by behavior map archives of the best-performing swarms (section 5.1.4, figures 5.4, 5.5, and 5.6), where behaviors associated with types *D* and *E* resources were of a lower quality (section 4.3.2) across all three environments for all treatments.

Overall, the task environment complexity (the degree of cooperation required) had a significant impact on the performance of the robot teams, with *simpler* environments leading to better outcomes.

6.1.2 Behavioral Diversity

Behavior was described in terms of two behavioral descriptors, namely, (1) the type of resource collected and (2) the maximum distance traversed in the environment during a robot’s lifetime (that is, a single generation). These descriptors aimed to encourage diversity in both the types of resources collected and the portions of the environment explored. These traits were calculated at the end of a lifetime (section 3.4).

Except for EDQD and EDQD-M, all other methods (Chapter 3, table 3.1) evolved comparable behavioral diversity in the *simple* and *medium* task environments (section 5.1.1.4, table 5.2). We observe that these environments required lower to moderate degrees of cooperation, allowing for extensive exploration of the environment (thus maximizing diversity in terms of the portions of the environment explored, section 5.1.2, figure 5.2) and diverse resource collection (section 5.1.3, figures 5.3a and 5.3b). This is corroborated by behavior map archives of the best-performing swarms (section 5.1.4, figures 5.4 and 5.5), where all methods evolved high-quality behaviors in the *simple* and *medium* environments that explored environment portions in the range [0.05, 0.95] and collected resource types *A* and *B* with a quality in the range [0.1, 0.95]. For resource types *C*, *D*, and *E*, the corresponding quality was in the range [0.05, 0.65]. Such maps show that swarms are capable of evolving diverse behavior (exploring up to 95% of the environment and gathering each resource type), with resource types *A* and *B* (requiring no and two robot cooperation, table 4.2) exhibiting the best-quality behavior (section 4.3.2).

In comparison to the *simple* and *medium* environments, the *difficult* environment necessitated a higher level of cooperation, resulting in reduced exploration of the environment (due to a *wait-for-assistance* period, table 4.2). This decrease in exploration was corroborated by a significantly smaller exploration of the environment (section 5.1.2.4, table 5.4) and fewer and less varied resource collection (section 5.1.3, figure 5.3c) compared to the *simple* and *medium* environments. This was further evidenced by behavior map archives of the best-performing swarms (section 5.1.4, figure 5.6), where all methods evolved behaviors that explored less than 50% of the environment (with the exception of Double-Map EDQD-M, which explored up to 55% of the environment for resource type *A*, with a corresponding quality of ≈ 0.95). The explored portions of the environment ranged from [0.05, 0.45] for all other methods, and the quality of resource types *A* and *B* collected fell within the range of [0.10, 0.3]. On the other hand, resource types *C*, *D*, and *E* had lower quality scores within the range [0.0, 0.2]. These behavior maps demonstrate that swarms can evolve diverse behaviors, although not to the same extent as in the *simple* and *medium* environments (section 5.1, figure 5.1). Additionally, resource types *A* and *B*, which required no or minimal cooperation (two robot cooperation), exhibited the highest quality behaviors (section 4.3.2).

These findings demonstrate that robot swarms (as described in this thesis), are capable of evolving diverse behaviors, exploring a large portion of the environment, and collecting diverse resource types. However, the degree of cooperation required in the task environment influenced the extent of exploration and the quality of behaviors exhibited by the swarms.

6.1.3 Morphological Diversity

We found that for [mEDEA-M](#) and [EDQD-M](#), there was no significant difference in the morphological diversity that evolved across the three environments (section 5.2.1, table 5.11). However, Double-Map [EDQD-M](#) evolved significantly greater morphological diversity in the *medium* environment than the other two environments. The *simple* environment also significantly outperformed the *difficult* environment. We do not discuss environmental influences on the evolution of morphological diversity for robot swarms evolved with [mEDEA-M](#) and [EDQD-M](#), since these variants yielded comparable morphological diversity across all task environments.

We observe that Double-Map [EDQD-M](#) evolved significantly greater morphological diversity in the *medium* environment (section 5.2.1, figure 5.7). This increased morphological diversity was accompanied by the highest maximum number of behaviors (section 5.1.1, figure 5.1) compared to the *simple* and *difficult* environments. A similar pattern was observed in the *simple* environment, where significantly higher morphological diversity coincided with significantly higher behavioral diversity compared to the *difficult* environment. These results support the notion that robot morphology determines the types and complexity of behaviors that can be exhibited by robot controllers (Kriegman et al., 2018).

The significant morphological diversity observed in the *difficult* environment for Double-Map [EDQD-M](#) swarms, coupled with significantly greater behavioral diversity, sets it apart from the other methods. While the other methods evolved comparable behavioral diversity (section 5.1.1, figure 5.1, table 5.1), Double-Map [EDQD-M](#) showcased a unique relationship between morphological and behavioral diversity. These results suggest that the evolution of morphological diversity contributes to the emergence of behavioral diversity. By allowing for a wider range of morphological configurations, robot swarms can explore a greater repertoire of behaviors, leading to increased behavioral diversity. The particular significance of morphological diversity in the *difficult* environment underscores its role in enabling the swarms to tackle the challenges posed by the high degree of cooperation required.

The morphology map archives of the best-performing swarms provided evidence for the benefits of significant morphological (and behavioral) diversity. These archives showed that swarms with higher morphological diversity exhibited higher behavioral quality across all task environments (figures 5.13a, 5.13b, 5.13c). Furthermore, the swarms evolved using Double-Map [EDQD-M](#) demonstrated the best overall team fitness compared to the other methods in the *simple* and *difficult* environments, although the difference was not statistically significant. These findings align with previous research (Zardini et al., 2021) that has highlighted the effectiveness of body-brain diversity maintenance mechanisms in facilitating the exploration of search spaces and the evolution of high-performing solutions.

This work focused on different task environments that varied in the level of cooperation (measure of complexity) needed to collect all the resources (table 4.2). The *medium* environment required a moderate degree of cooperation and had an equal distribution of resource types, resulting in the highest resource variety. A hypothesis suggests that the evolution of morphological diversity is influenced by the diversity of resources in the environment. To test this hypothesis, we examined a scenario where each robot only reacted to one type of resource. Table 6.1 presents the number of agents needed to collect all the resources of each type in the three task environments.

Table 6.1: The number of individuals required to collect all the resources of each resource type in each environment.

	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>
<i>Simple</i>	30	10	15	20	25
<i>Medium</i>	10	20	30	40	50
<i>Difficult</i>	5	10	15	20	120

Upon examining the hypothetical scenario where each robot responds to only one resource type, it becomes apparent that the task would still be accomplished in the *simple* environment since the total number of required robots matches the population of the swarm (100 robots). However, in the *medium* environment, a total of 150 robots would be necessary, which is 1.5 times more than the swarm population. This indicates that robots in this environment need to respond to multiple resource types. The variation in the number of robots required per resource type in the *medium* environment is minimal, ranging from 7% to 33%, with an average difference of 7%. Hence, the *medium* environment necessitates a higher level of diversity. On the other hand, the *difficult* environment demands a lower level of diversity, as 70.5% of the population is needed to collect all the type *E* resources, while only 11% is required for type *D* resources (the second highest).

The observed results, combined with the hypothetical scenario, provide support for the hypothesis proposing that the diversity of resources in the environment influences the evolution of morphological diversity. The findings indicate that in environments with a higher variety of resources, such as the *medium* environment, a greater level of morphological diversity is required to effectively collect all the resources. Conversely, in environments with a lower diversity of resources, like the *difficult* environment, a lesser amount of morphological diversity is necessary. Therefore, it can be concluded that the availability and distribution of resources play a role in shaping the evolution of morphological diversity in robot swarms.

When examining the relationship between neural-morpho complexity (section 5.2.4) and morphological diversity in each environment, an intriguing observation emerges. The *medium* environment exhibits the highest morphological diversity (section 5.2.1, figure 5.7) in concert with the lowest median complexity (section 5.2.4, figure 5.12). This finding suggests that diversity leads to the evolution of simpler morphologies within the population. In the *simple* environment, the second highest morphological diversity corresponds to the second lowest median complexity. In contrast, the *difficult* environment shows the least morphological diversity in conjunction with the highest median complexity. These results indicate that a lack of diversity results in the emergence of complex morphologies.

Consequently, these findings lead to the hypothesis that there exists a relationship between morphological complexity and the emergence of morphological diversity in robot swarms, similar to what is observed in nature (Carroll, 2001). However, in the field of ecology, this relationship has primarily been studied at the phylogenetic level, where it has been observed that increased morphological complexity leads to expansions into previously unoccupied ecological niches and an accompanying increase in species diversity (Carroll, 2001). Indeed, further research is warranted to gain a deeper understanding of the relationship between morphological complexity and morphological diversity at the group level in robot swarms. While this work provides valuable insights into the influence of environmental factors on morphological diversity, there is still more to explore.

6.2 Impact of Behavioral Diversity Maintenance on Collective Behavior Evolution

Our second objective was to investigate whether behavioral diversity maintenance is advantageous for evolving collective behavior across increasingly complex task environments. Overall results indicated that the [EDQD](#) method, evolving collective behavior with behavioral diversity maintenance, was demonstrated as beneficial as environment complexity increased. These benefits were demonstrated as significantly higher task performance (section 5.3, figure 5.14) and behavioral diversity (section 5.1, figure 5.1) compared to the [mEDEA](#) method (benchmark), evolving collective behavior without diversity maintenance mechanisms.

Our observation focuses on the behavioral diversity maintenance mechanism employed by the [EDQD](#) method, specifically the [BC](#). The [BC](#) actively promotes varied exploration of the environment and the collection of diverse resources. As a result, the discovery of novel solutions also led to higher team task performance. This is supported by significantly higher team task performance compared to the [mEDEA](#) method in the *simple* and *medium* environments, while performance remains comparable in the *difficult* environment (section 5.3, figure 5.14). The [BC](#) mechanism is considered to be aligned with the task at hand, as it encourages behaviors that enhance overall performance in the given environments.

In the *simple* environment, the effectiveness of diversity maintenance, characterized by an aligned [BC](#), was corroborated by significantly greater exploration of the environment (section 5.1.2, figure 5.2). This increased exploration proved beneficial in the collection of resource types *A*, *B*, and *E* (section 5.1.3, figure 5.3a). The suitability of diversity maintenance in evolving effective behaviors in the *simple* environment was further demonstrated by analyzing the behavior map archives of the best-performing swarms (section 5.1.4, figure 5.4). The [EDQD](#) method evolved behaviors explored up to 95% of the environment and successfully collected all resource types within a quality range of [0.01, 0.8]. Comparatively, [mEDEA](#) evolved behaviors explored up to 55% of the environment, with a corresponding quality range of [0.05, 0.95]. These findings highlight how diversity maintenance, driven by an aligned [BC](#), enhances exploration and resource collection capabilities in the *simple* environment.

Comparatively, in the *medium* environment, both [EDQD](#) and [mEDEA](#) evolved comparable behavioral diversity (section 5.1, table 5.1). However, [EDQD](#) evolved swarms significantly outperformed those evolved with [mEDEA](#) on the collective gathering task overall (section 5.3, figure 5.14), demonstrating the effectiveness of diversity maintenance in the *medium* environment. Furthermore, the effectiveness of diversity maintenance was further evidenced by behavior map archives of the best-performing swarms (section 5.1.4, figure 5.5). In the case of resource types *A* and *B*, the behaviors evolved by [EDQD](#) explored up to 45% of the environment, with corresponding quality values within the range [0.2, 0.8]. In comparison, behaviors evolved by [mEDEA](#) for the same resource types explored only up to 20% of the environment, with quality values within the range [0.1, 0.5]. These behavior maps clearly demonstrate the benefits of explicitly encouraging environment exploration during behavior search, as opposed to relying solely on the objective function to discover such behaviors.

In the *difficult* environment the effectiveness of driving behavior search with diversity maintenance, guided by an aligned [BC](#), was demonstrated by the evolution of greater behavioral diversity (section 5.1, figure 5.1) and improved task performance (section 5.3, figure 5.14), although the differences were not

statistically significant when compared to [mEDEA](#) (which evolves collective behavior without a diversity maintenance mechanism). However, the benefits of diversity maintenance were evident when analyzing the behavior map archives of the best-performing swarms (section 5.1.4, figure 5.6). The behaviors evolved by [EDQD](#) were capable of exploring up to 45% of the environment with a maximum quality of ≈ 0.6 . In comparison, the behaviors evolved by [mEDEA](#) explored up to 30% of the environment with a similar quality of ≈ 0.6 . Although not significantly different in terms of performance, the behavior maps highlight the potential of diversity maintenance in enabling the evolution of behaviors that can effectively explore a substantial portion of the environment while achieving comparable quality.

The demonstrated benefits of the [EDQD](#) algorithm in evolving effective collective behaviors across multiple robots align with previous research (Hart et al., 2018), which showed that a decentralized [QD](#) approach can achieve such results. Additionally, our findings support previous work (Pugh et al., 2016b) that demonstrated the effectiveness of an aligned [BC](#) in driving search for [QD](#) as task environment complexity increases. While Pugh et al. (2016b) focused on single agents in a maze exploration context, our work extends these findings to the collective behavior domain. Moreover, the aligned [BC](#) used in our diversity maintenance methods explicitly encouraged diverse exploration of the environment and collection of varied resources, contrasting with the maze endpoint exploration encouraged by Pugh et al. (2016b).

6.3 Impact of Behavioral Diversity Maintenance on Co-evolving Collective Behavior and Morphology

Our third objective was to investigate whether behavioral diversity maintenance is beneficial for co-evolving collective behavior and morphology across increasingly complex task environments. Overall results indicated that the **EDQD-M** method, co-evolving collective behavior and morphology with behavioral diversity maintenance, was demonstrated as beneficial as task complexity increased. These benefits were demonstrated as significantly higher task performance (section 5.3, figure 5.14) and behavioral diversity (section 5.1, figure 5.1) compared to the **mEDEA-M** method (benchmark), co-evolving collective behavior and morphology without diversity maintenance mechanisms.

6.3.1 Behavioral Diversity

The **EDQD-M** method utilized a behavioral diversity maintenance mechanism, guided by an aligned **BC**, to drive behavior search. On the other hand, the **mEDEA-M** method exclusively utilized an objective function to guide behavior search. Specifically, in **EDQD-M**, the **BC** was designed to explicitly encourage diversity in terms of environmental exploration, which was evidenced by significantly greater average distance explored overall (section 5.1.2, figure 5.2). The **BC** further promoted diversity in terms of the resource types collected. The benefits of fostering diversity in this aspect were evidenced by the ability of **EDQD-M** evolved swarms to consistently outperform **mEDEA-M** in terms of the average number of resource types gathered overall (section 5.1.3, figures 5.3a, 5.3b, and 5.3c).

The benefits of behavioral diversity maintenance were further evidenced by examining behavior map archives of the best-performing swarms (section 5.1.4, figures 5.4, 5.5, and 5.6). The **EDQD-M** method consistently evolved high-quality behaviors for resource types *A*, *B*, and *C* across all three environments. Comparatively, the **mEDEA-M** method typically evolved a high-quality behavior for one resource type, while the behaviors corresponding to other resource types had lower quality. For example, in the *medium* environment, the highest-quality behavior evolved by **mEDEA-M** corresponded to resource type *A* (quality ≈ 0.9), while the highest-quality for resource types *B* and *C* were ≈ 0.4 and ≈ 0.3 , respectively. In contrast, the highest-quality behavior evolved by **EDQD-M** corresponded to resource type *A* (quality ≈ 0.8), while the highest quality for resource types *B* and *C* were ≈ 0.6 and ≈ 0.6 , respectively, which are all comparable.

Consequently, diversity maintenance led to the emergence of significantly higher behavioral diversity in the *simple* and *medium* environments compared to **mEDEA-M** (section 5.1, figure 5.1). This increased diversity also resulted in a significantly higher level of team task performance in the medium environment (section 5.3, figure 5.14). Specifically, **EDQD-M** significantly outperformed all other methods in the *medium* environment. We observe that the *medium* environment had an equal distribution of all resource types, and required a moderate level of cooperation. These findings highlight the effectiveness of behavioral diversity maintenance in evolving effective behaviors, in conjunction with morphologies, in collective behavior tasks that necessitate a moderate level of cooperation.

6.3.2 Morphological Diversity

In both [EDQD-M](#) and [mEDEA-M](#), the state of the evolving sensors at the end of each generation, determined by the sensor-morpho threshold in table 4.2, determined the activation or deactivation of the associated [ANN](#) sensory input nodes (sections 3.5 and 3.3). For example, if there were x active sensors and y inactive sensors, the chosen [ANN](#) controller would be reconfigured so that x sensory input nodes were active and y were inactive. Therefore, the evolved morphology determined the corresponding controller.

However, behavioral diversity maintenance, defined by a task-aligned [BC](#), did not have an impact on the evolution of morphological diversity. This was evident from the emergence of similar levels of morphological diversity across all task environments (section 5.2, figures 5.7). Furthermore, the learning curves of both approaches were identical (sections 5.2.2 and 5.2.3, figures 5.9 and 5.11).

Examining the morphology map archives of the best-performing swarms, it was observed that [mEDEA-M](#) utilized a wider range of active sensors and sensor ranges compared to [EDQD-M](#) (section 5.2.5, figures 5.13a, 5.13b, and 5.13c). This suggests that environment-driven approaches are better suited to exploiting morphological diversity.

6.3.3 Neural-Morpho Complexity

Neural-morpho complexity, measured according to equation 5.1, results indicated no statistical significance between [mEDEA-M](#) and [EDQD-M](#) (section 5.2.4, figure 5.12). This is because both approaches had an environment-driven morphology evolution process. The diversity maintenance mechanism in [EDQD-M](#) was solely employed to drive behavior search.

However, when we considered the median neural-morpho complexity for each approach, interesting observations were made (table 6.2). Specifically, [EDQD-M](#) evolved its simplest morphologies in the *difficult* environment, while the most complex morphologies were evolved in the *medium* environment. It is worth noting that these complex morphologies were associated with significantly higher team task performance compared to all other methods (section 5.3, figure 5.14). Comparatively, [mEDEA-M](#) evolved its simplest morphologies in the *simple* environment, while the most complex morphologies were evolved in the *medium* environment.

Based on these observations, we hypothesize that the influence of the diversity maintenance mechanism in [EDQD-M](#), which drives behavior search, is responsible for the observed differences between [mEDEA-M](#) and [EDQD-M](#). This hypothesis aligns with the notion proposed by Spanellis et al. (2021) that diversity maintenance approaches evolve robot complexity necessary to achieve high task performance in specific environments. It is a reasonable assumption considering the coupling of morphology and behavior in this work.

Table 6.2: Median neural-morpho complexity (over 20 runs) in each environment. Bold values indicate the lowest neural-morpho evolved in each environment.

Method	<i>Simple</i>	<i>Medium</i>	<i>Difficult</i>
mEDEA-M	0.549	0.630	0.621
EDQD-M	0.596	0.599	0.564

6.4 Impact of Body-Brain Diversity Maintenance on Co-evolving Collective Behavior and Morphology

Our final objective was to investigate whether behavioral-morphological diversity maintenance is beneficial for co-evolving collective behavior and morphology across increasingly complex task environments. Results indicated that the Double-Map [EDQD-M](#) method, co-evolving collective behavior and morphology with behavioral-morphological diversity maintenance, was demonstrated as beneficial as task complexity increased. These benefits were demonstrated as significantly higher task performance (section 5.3, table 5.18), behavioral diversity (section 5.1, table 5.1), and morphological diversity (section 5.2, table 5.10) than the [mEDEA-M](#) method (benchmark), co-evolving collective behavior and morphology without diversity maintenance mechanisms.

6.4.1 Behavioral Diversity

The Double-Map [EDQD-M](#) method utilized two [QD](#) maps to maintain diversity in controller and morphology related properties. The [BC](#) used in the controller-related [QD](#) map was aligned with the collective gathering task, promoting varied exploration of the environment and collection of diverse resources. As a result, the discovery of novel solutions contributed to higher task performance. This can be observed through the emergence of significantly higher behavioral diversity in the *simple*, *medium*, and *difficult* environments (section 5.3, table 5.18) compared to [mEDEA-M](#). In the *simple* environment, this increased diversity was also accompanied by significantly higher team task performance (section 5.3, table 5.18).

The behavior map archives of the best-performing swarms in each approach provided evidence of the effectiveness of evolved gathering behaviors across different task complexities (*simple*, *medium*, *difficult*, section 5.1.4). In the *simple* environment, the benefits of body-brain diversity were evidenced by the discovery of high-quality behaviors (figure 5.4), where the maximum portion of the environment explored for resource types *A*, *B*, and *C* was ≈ 0.45 , with corresponding behavior quality in the range: [0.65, 0.95]. Comparatively, [mEDEA-M](#) evolved behaviors yielded a quality range of [0.03, 0.75] for an explored area in the range: [0.25, 0.45] for the same resource types. Moreover, Double-Map [EDQD-M](#) evolved behaviors exhibited an extensive exploration of the environment (explored up to 45% of the environment) to find resource types *B* and *C*, while also yielding high behavior quality. In contrast, [mEDEA-M](#) only explored 25% of the environment for these resources. The ability of Double-Map [EDQD-M](#) to evolve behaviors that efficiently search the environment to find all resource types was further corroborated by significantly greater overall environment exploration (section 5.1.2, table 5.3).

The benefits of body-brain diversity were further evidenced in the *medium* environment (figure 5.4), where the maximum portion of the environment explored for resource type *A* was 0.45, with a corresponding behavior quality of ≈ 0.95 . Comparatively, mEDEA-M evolved behaviors explored a larger portion of the environment (0.85) for resource type *A* but yielded a lower quality of ≈ 0.65 . This indicates that mEDEA-M behaviors wasted more time exploring the environment without achieving higher quality. On the other hand, for resource type *C*, Double-Map EDQD-M evolved behaviors explored a maximum environment portion of 0.35 and yielded a quality of ≈ 0.65 . In comparison, mEDEA-M evolved behaviors explored a smaller environment portion (0.25) for resource type *C* and had a lower quality of ≈ 0.2 . These findings demonstrate that when Double-Map EDQD-M evolved behaviors explored a larger portion of the environment, they also exhibited higher quality. These findings are consistent with prior research (Nordmoen et al., 2021), which has shown that body-brain diversity maintenance is effective in achieving a balance between exploration and exploitation in the behavior-morphology search space.

In the *difficult* environment, the effectiveness of body-brain diversity maintenance was particularly evident in Double-Map EDQD-M (figure 5.6). The behaviors evolved by Double-Map EDQD-M exhibited a high-quality diversity, with an explored environment portion range of [0.35, 0.55] for resource types *A* and *B*, and a corresponding quality range of [0.65, 0.95]. In contrast, mEDEA-M evolved behaviors for resource types *A* and *B* had a quality range of [0.10, 0.50] for an explored area of [0.05, 0.25]. Additionally, figure 5.1 illustrates the average behavioral diversity of Double-Map EDQD-M evolved swarms, which showed significantly higher diversity compared to swarms evolved by mEDEA-M (mEDEA, EDQD, and EDQD-M) according to statistical analysis (Mann-Whitney U, $p < 0.05$). Although the team task performance of Double-Map EDQD-M was higher than other approaches, the difference was not statistically significant. These findings support previous research demonstrating that sufficient diversity is required to find solutions in complex environments, and search algorithms that generate and maintain diversity are more likely to find a solution even when considering different environments with varying complexities (Nordmoen et al., 2021).

6.4.2 Morphological Diversity

The morphological descriptors utilized by the morphology-related QD map in the Double-Map EDQD-M method encouraged diverse utilization of the maximum sensor range and the number of active sensors. The benefits of morphological diversity maintenance were demonstrated as significantly higher morphological diversity compared to mEDEA-M (and EDQD-M) in the *simple*, *medium*, and *difficult* task environments (section 5.2, table 5.10). Moreover, the evolved morphologies proved to be effective, as they achieved significantly higher task performance in the *simple* environment, and comparable task performance in the *medium* and *difficult* environments when compared to mEDEA-M (section 5.3, table 5.18).

The effectiveness of the Double-Map EDQD-M method in evolving morphological diversity and generating high-quality behavior was consistently demonstrated across all task environments. 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 5.13a, 5.13b, 5.13c). The highest-quality solutions, with success rates exceeding 90%, were achieved when sensors were active between 65% and 95% and operated at 85% to 100% of their maximum range (figures 5.13a, 5.13b, 5.13c, right). Comparable task performance was observed in the *simple* and *difficult* environments for the morphologies evolved by EDQD-M, indicating

that the sensor activity and ranges achieved through the Double-Map [EDQD-M](#) approach facilitated effective swarm behaviors.

Upon examining the morphology map archives of [mEDEA-M](#) and Double-Map [EDQD-M](#), it becomes evident that in the *simple* and *medium* task environments, the highest quality morphologies consistently outperform their counterparts in [mEDEA-M](#) (and [EDQD-M](#)) evolved swarms. These superior morphologies in Double-Map [EDQD-M](#) exhibit higher quality scores compared to the highest quality morphologies in [mEDEA-M](#) (and [EDQD-M](#)) evolved swarms (as shown in figures [5.13a](#) and [5.13b](#)). Similarly, in the *difficult* task environment, the highest quality morphology evolved by Double-Map [EDQD-M](#) surpasses the highest quality morphology evolved by [mEDEA-M](#), although the difference is only slight ($\approx 90\%$ against 60% of maximum sensor range, as depicted in figure [5.13c](#)).

Additionally, in the *difficult* task environment, the benefits of morphological diversity in Double-Map [EDQD-M](#) are reinforced by the presence of higher behavioral quality diversity compared to [mEDEA-M](#). Specifically, the highest quality behaviors evolved by Double-Map [EDQD-M](#) (related to gathering resource type *A*) outperformed the highest quality behaviors evolved by [mEDEA-M](#) (with a maximum behavior quality of approximately ≈ 0.9 and ≈ 0.6 , respectively, as depicted in figure [5.6](#)). This finding highlights the effectiveness of morphological diversity maintenance in Double-Map [EDQD-M](#) in generating behaviors of superior quality compared to [mEDEA-M](#) in challenging task environments.

In the *medium* task environment, Double-Map [EDQD-M](#) exhibited similar, albeit less noticeable, benefits compared to [mEDEA-M](#). Specifically, the maximum behavior quality achieved by Double-Map [EDQD-M](#) for gathering resource type *A* was ≈ 0.9 , while [mEDEA-M](#) achieved a maximum behavior quality of ≈ 0.9 (as shown in figure [5.6](#)). In the *simple* environment, the maximum behavior quality of Double-Map [EDQD-M](#) was comparable to that of [mEDEA-M](#) for gathering resource type *A*. These findings, supported by table [5.18](#), further validate previous research (Pugh et al., [2016a](#)) using single agents (while this study focuses on collective behavior) that demonstrates how multi-BC approaches can facilitate the explicit search for both unaligned (morphological diversity in this study) and aligned diversity, while potentially losing no significant performance.

6.4.3 Neural-Morpho Complexity

Since Double-Map [EDQD-M](#) maintained two map archives, each of which was updated independently at the end of each generation, there had to be a mechanism in place to ensure that robot behavior and morphology were appropriately matched. Consequently, the evolved morphology determined the corresponding controller. As an example, a chosen morphology with x active and y inactive sensors would automatically reconfigure the chosen [ANN](#) controller so that x sensory input nodes were active and y were inactive. As a result, the complexity of the morphology directly influenced the complexity of the neural network in this work (where neural complexity is calculated as the number of active input nodes, which corresponds to the proportion of active sensors, section [5.2.3](#)).

Table [6.3](#) serves as additional evidence for the benefits of body-brain diversity maintenance. Specifically, in the *simple* environment, Double-Map [EDQD-M](#) evolved the most complex morphologies, although the difference was not statistically significant, and these morphologies were accompanied by significantly

Table 6.3: Median neural-morpho complexity (over 20 runs) in each environment. Bold values indicate the lowest median complexity in each environment.

Method	<i>Simple</i>	<i>Medium</i>	<i>Difficult</i>
mEDEA-M	0.549	0.630	0.621
Double-Map EDQD-M	0.551	0.549	0.559

higher task performance compared to [mEDEA-M](#). However, as task complexity increased, body-brain diversity maintenance consistently led to the evolution of significantly simpler morphologies compared to [mEDEA-M](#), while still maintaining effective task performance (although not significantly higher than [mEDEA-M](#)). These results provide additional evidence that diversity maintenance promotes the evolution of simple controllers and high-quality behaviors in robot teams (Gomes et al., [2013](#); Lehman & Stanley, [2011a](#); Nitschke & Didi, [2017](#)).

The results presented in Table [6.3](#) reinforce the notion that diversity maintenance approaches are capable of evolving the necessary complexity in robots to achieve high task performance in specific environments (Spanellis et al., [2021](#)). Therefore, these results provide further evidence that increased task complexity does not necessarily equate to the evolution of increased morphological complexity in robot teams (Hewland & Nitschke, [2015](#); Nitschke & Didi, [2017](#)). One hypothesis suggests that this is due to the presence of morphological diversity (see section [6.1.3](#)).

6.5 Summary

This Chapter discussed our findings from Chapter 5 and provided a critical evaluation of the research questions posed in this thesis. The focus was on comparing the task performance, behavioral diversity, and morphological diversity that emerged in the *simple*, *medium*, and *difficult* environments. Our results clearly showed that the complexity of the task environment had a significant impact on task performance as well as the emergence of behavioral and morphological diversity in cooperative robot teams. As the task complexity increased from *simple* to *difficult*, we observed a notable decrease in team fitness. A similar trend was observed in terms of the emergence of behavioral diversity. To gain a deeper understanding of how environmental complexity influences the evolution of collective behavior and morphology, we conducted a comprehensive analysis of the properties associated with each of the three environments considered in this thesis.

In section 6.2, we focused on comparing the impact of behavioral diversity maintenance versus objective-based evolutionary search on collective behavior evolution. Our analysis revealed that the EDQD approach outperformed mEDEA in terms of producing significantly diverse robot swarms with significantly higher team task fitness. By examining behavior map archives, we found that EDQD evolved high-quality behaviors that extensively explored the environment and gathered diverse types of resources. On the other hand, mEDEA struggled to evolve such combinations of behavioral features due to its inability to ensure equal distribution of effort across different behavior dimensions. The ability of EDQD to achieve a balanced effort distribution and explore a wide range of behaviors was made possible by its behavioral diversity maintenance mechanism. This allowed for the discovery of high-quality behaviors through sufficient exploration of the search space.

In section 6.3, we focused on the coevolution of collective behavior and morphology and compared the impact of behavioral diversity maintenance and objective-based evolutionary search. The analysis of results showed that EDQD-M outperformed mEDEA-M in several aspects. Firstly, EDQD-M facilitated the evolution of significantly behaviorally diverse robot swarms with higher team task fitness. Additionally, EDQD-M maintained lower median neural-morpho complexity as task complexity increased. The evolved swarms from EDQD-M explored larger portions of the environment and collected a greater number of resources per resource type compared to mEDEA-M. The analysis of behavior map archives further supported the effectiveness of EDQD-M in evolving high-quality behaviors associated with diverse resource types, whereas mEDEA-M struggled to achieve similar quality for more than one resource type. The ability of EDQD-M to balance exploration and exploitation in the behavior space, facilitated by behavioral diversity maintenance, also led to the evolution of simpler neural-morpho complexity in complex environments.

In section 6.4, we investigated the impact of body-brain diversity maintenance compared to objective-based evolutionary search on the coevolution of collective behavior and morphology. The results showed that Double-Map EDQD-M outperformed mEDEA-M in several aspects. Firstly, Double-Map EDQD-M enabled the evolution of significantly higher behavioral and morphological diversity in robot swarms with higher team task fitness. As task complexity increased, Double-Map EDQD-M maintained simpler behavior-morphology couplings. The analysis of evolved swarm behaviors revealed that Double-Map EDQD-M consistently explored larger portions of the environment compared to mEDEA-M. This propensity for exploration and the collection of diverse resource types was facilitated by body-brain

diversity maintenance, leading to the discovery of high-quality behaviors. In contrast, [mEDEA-M](#) and other approaches failed to explore the feature space to the same extent and consequently could not uncover high-quality behaviors, particularly in complex environments. The findings highlight the effectiveness of body-brain diversity maintenance in promoting behavioral and morphological diversity and enhancing overall swarm performance.

This thesis makes a significant contribution to the state of the art by exploring the interplay between environmental complexity, behavioral diversity maintenance, and body-brain diversity maintenance in the evolution of collective behavior and morphology in robot swarms. The key findings highlight the importance of these factors in achieving robust and adaptive behavior in cooperative robot teams. The study demonstrates that incorporating diversity maintenance mechanisms during the evolutionary process, as exemplified by the [EDQD](#), [EDQD-M](#), and Double-Map [EDQD-M](#) approaches, leads to the evolution of significantly diverse robot swarms with higher task fitness when compared to objective-based approaches. These approaches effectively balance exploration and exploitation, enabling the discovery of high-quality behaviors associated with diverse resource types. Moreover, the research sheds light on the impact of environmental complexity on task performance and the emergence of behavioral and morphological diversity. By providing insights into these dynamics, the thesis offers valuable knowledge for designing more effective evolutionary algorithms and strategies for collective behavior in robotics.

Chapter 7

Conclusions

The goal of this thesis was to examine the influence of task complexity on the task performance, behavioral diversity, and morphological diversity of body-brain evolved cooperative robot teams. We sought to achieve this by comparing robot teams evolved with diversity maintenance mechanisms against those evolved without such mechanisms (that is, objective-based search). By conducting this comparison, the research aimed to provide insights into the role of diversity maintenance in shaping the characteristics and capabilities of evolved robot teams in different task complexity scenarios.

To evaluate the effectiveness of the different neuro-evolution approaches, it was necessary to employ various levels of task complexity. The collective gathering task was chosen as the primary test-bed due to its suitability for assessing cooperation levels among individual robots for optimal performance. This task has been widely used as a benchmark in collective robotics research (Bonabeau et al., 1999; Furman et al., 2019; Hewland & Nitschke, 2015; Kernbach, 2012) and is closely related to real-world applications such as surveillance, environment monitoring (Vásárhelyi et al., 2018), agricultural foraging, and mining tasks (Floreano & Wood, 2015). By selecting this task, the thesis aimed to provide insights that are relevant to both current and future swarm robotics applications, thereby contributing to the advancement of the field.

The experimental results from this thesis provided valuable insights into the relationship between task complexity and the performance and diversity of body-brain evolved cooperative robot teams. Specifically, the findings indicated that increasing task complexity negatively affected team task performance and behavioral and morphological diversity. Furthermore, the results demonstrated that explicit diversity maintenance mechanisms in evolutionary search were crucial for evolving effective collective behavior in progressively complex tasks.

In the *simple* environment, behavioral diversity maintenance alone resulted in significantly higher team task performance when compared to objective-based search approaches. However, in the *medium* environment, a combination of behavioral diversity maintenance and morphology adaptation was necessary for achieving similar benefits. In the most complex environment, body-brain diversity maintenance yielded the best results in terms of behavioral and morphological diversity and team task performance.

These findings align with previous research that similarly demonstrated, using single agents (whereas these results are pertinent to collective behavior), that multi-BC approaches allow explicitly searching for both unaligned (morphological diversity in this work) and aligned diversity while potentially losing no significant performance (Pugh et al., 2016a; Pugh et al., 2016b). These results also provide additional evidence to the notion that sufficient diversity is required to find solutions in complex environments, and search algorithms that generate and maintain diversity are more likely to find a solution even when considering different environments with varying complexities (Nordmoen et al., 2021). Finally, experiment results also indicated that diversity maintenance search evolved high-quality behaviors encoded by simple behavior-morphology couplings as environment complexity increased. Thus, these results provide additional evidence that diversity maintenance promotes the evolution of simple controllers and high-quality behaviors in robot teams (Gomes et al., 2013; Lehman & Stanley, 2011a; Nitschke & Didi, 2017).

The main contribution of this study is empirical evidence indicating the effectiveness of body-brain diversity maintenance search approaches in complex tasks defined by large feature spaces in robotic swarms. While previous related work (Nordmoen et al., 2021; Pugh et al., 2016a; Pugh et al., 2016b; Zardini et al., 2021) has shown the benefits of these approaches for single agents, this thesis extends those findings to collective behavior in robotic swarms.

7.1 Future Work

A logical progression of this research would involve exploring alternative indicators of environmental complexity, such as different levels of environmental steepness, to gain deeper insights into the influence of body-brain diversity that arises in response to changing environments within robotic swarms. In particular, there is a need for further investigation to comprehend the effects and interactions of different types of behavior-morphology couplings in concert with various environment types within the framework of body-brain diversity maintenance-driven evolutionary search.

Secondly, there is currently a lack of research on the relationship between morphological diversity and complexity in robot swarms. It is crucial to investigate and comprehend the relationship between morphological diversity and morphological complexity that arise as a result of changing environments within robotic swarms. This research would contribute to a better understanding of how changing environments impact the emergence of diverse morphologies and their associated complexities within the swarm.

Bibliography

- Armstrong, E. (1983). Relative brain size and metabolism in mammals. *Science*, 220(4603), 1302–1304.
- Arnold, R., Carey, K., Abruzzo, B., & Korpela, C. (2019). What is a robot swarm: A definition for swarming robotics. *2019 IEEE 10th Annual Ubiquitous Computing, Electronics and Mobile Communication Conference (UEMCON)*, 0074–0081. <https://doi.org/10.1109/UEMCON47517.2019.8993024>
- Auerbach, J., & Bongard, J. (2009). Evolution of functional specialization in a morphologically homogeneous robot. *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, 89–96.
- Auerbach, J., & Bongard, J. (2014). Environmental influence on the evolution of morphological complexity in machines. *PLoS computational biology*, 10(1), e1003399.
- Auerbach, J., Concorde, A., Kornatowski, P., & Floreano, D. (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), 356–369.
- Ball, D., Ross, P., English, A., Patten, T., Upcroft, B., Fitch, R., Sukkarieh, S., Wyeth, G., & Corke, P. (2015). Robotics for sustainable broad-acre agriculture. In *Field and service robotics: Results of the 9th international conference* (pp. 439–453). Springer International Publishing. https://doi.org/10.1007/978-3-319-07488-7_30
- Beni, G. (2004). From swarm intelligence to swarm robotics. *Lecture Notes in Computer Science*, 3342, 1–9. https://doi.org/10.1007/978-3-540-30552-1_1
- Bernstein, D., Givan, R., Immerman, N., & Zilberstein, S. (2002). The complexity of decentralized control of markov decision processes. *Mathematics of operations research*, 27(4), 819–840.
- Beverly, B., McLendon, H., Nacu, S., Holmes, S., & Gordon, D. (2009). How site fidelity leads to individual differences in the foraging activity of harvester ants. *Behavioral Ecology*, 20(1), 633–638.
- Birattari, M., Ligt, A., Bozhinoski, D., Brambilla, M., Francesca, G., Garattoni, L., Garzón Ramos, D., Hasselmann, K., Kegeleirs, M., Kuckling, J., Pagnozzi, F., Roli, A., Salman, M., & Stützle, T. (2019). Automatic off-line design of robot swarms: A manifesto. *Frontiers in Robotics and AI*, 6. <https://doi.org/10.3389/frobt.2019.00059>
- Bonabeau, E., Dorigo, M., Theraulaz, G., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. Oxford university press.
- Bongard, J. (2011). Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences*, 108(4), 1234–1239. <https://doi.org/10.1073/pnas.1015390108>
- Bongard, J. (2013). Evolutionary robotics. *Commun. ACM*, 56(8), 74–83. <https://doi.org/10.1145/2493883>

- Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41. <https://doi.org/10.1007/s11721-012-0075-2>
- Bredeche, N., Haasdijk, E., & Prieto, A. (2018). Embodied evolution in collective robotics: A review. *Frontiers in Robotics and AI*, 5. <https://doi.org/10.3389/frobt.2018.00012>
- Brown, D. S., Turner, R., Hennigh, O., & Loscalzo, S. (2018). Discovery and exploration of novel swarm behaviors given limited robot capabilities. In R. Groß, A. Kolling, S. Berman, E. Frazzoli, A. Martinoli, F. Matsuno, & M. Gauci (Eds.), *Distributed autonomous robotic systems: The 13th international symposium* (pp. 447–460). Springer International Publishing. https://doi.org/10.1007/978-3-319-73008-0_31
- Brutschy, A., Tran, N.-L., Baiboun, N., Frison, M., Pini, G., Roli, A., Dorigo, M., & Birattari, M. (2012). Costs and benefits of behavioral specialization. *Robotics and autonomous systems*, 60(11), 1408–1420.
- Buason, G., Bergfeldt, N., & Ziemke, T. (2005). Brains, bodies, and beyond: Competitive co-evolution of robot controllers, morphologies and environments. *Genetic Programming and Evolvable Machines*, 6(1), 25–51.
- Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2), 156–172.
- Bușoniu, L., Babuška, R., & De Schutter, B. (2010). Multi-agent reinforcement learning: An overview. *Innovations in multi-agent systems and applications-1*, 183–221.
- Camci, E., Kripalani, D., Ma, L., Kayacan, E., & Khanesar, M. (2018). An aerial robot for rice farm quality inspection with type-2 fuzzy neural networks tuned by particle swarm optimization-sliding mode control hybrid algorithm. *Swarm and Evolutionary Computation*, 41, 1–8. <https://doi.org/https://doi.org/10.1016/j.swevo.2017.10.003>
- Carrillo-Zapata, D., Sharpe, J., Winfield, A., Giuggioli, L., & Hauert, S. (2019). Toward controllable morphogenesis in large robot swarms. *IEEE Robotics and Automation Letters*, 4(4), 3386–3393.
- Carroll, S. (2001). Chance and necessity: The evolution of morphological complexity and diversity. *Nature*, 409(6823), 1102–1109.
- Cheney, N., MacCurdy, R., Clune, J., & Lipson, H. (2013). Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding. *Proceedings of the Genetic and Evolutionary Computation Conference*, 11–23.
- Cheney, N., Bongard, J., SunSpiral, V., & Lipson, H. (2017). Scalable co-optimization of morphology and control in embodied machines. *Journal of The Royal Society Interface*, 15. <https://doi.org/10.1098/rsif.2017.0937>
- Cully, A., & Mouret, J.-B. (2016). Evolving a Behavioral Repertoire for a Walking Robot. *Evolutionary Computation*, 24(1), 59–88. <https://doi.org/10.1162/EVCO.a.00143>
- Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2), 182–197.
- Dias, P., Silva, M. C., Rocha Filho, G., Vargas, P., Cota, L., & Pessin, G. (2021). Swarm robotics: A perspective on the latest reviewed concepts and applications. *Sensors*, 21(6), 2062.
- Didi, S., & Nitschke, G. (2016). Hybridizing novelty search for transfer learning. *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, 1–8. <https://doi.org/10.1109/SSCI.2016.7850180>
- Doncieux, S., Bredeche, N., Mouret, J.-B., & Eiben, A. (2015). Evolutionary robotics: What, why, and where to. *Frontiers in Robotics and AI*, 2, 4. <https://doi.org/10.3389/frobt.2015.00004>
- Dorigo, M., Floreano, D., Gambardella, L., Mondada, F., Nolfi, S., Baaboura, T., Birattari, M., Bonani, M., Brambilla, M., Brutschy, A., Burnier, D., Campo, A., Christensen, A., Decugniere, A.,

- Di Caro, G., Ducatelle, F., Ferrante, E., Forster, A., Gonzales, J., ... Vaussard, F. (2013). Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics and Automation Magazine*, 20(4), 60–71. <https://doi.org/10.1109/MRA.2013.2252996>
- Dorigo, M., Theraulaz, G., & Trianni, V. (2020). Reflections on the future of swarm robotics. *Science Robotics*, 5(49), eabe4385. <https://doi.org/10.1126/scirobotics.abe4385>
- Duarte, A., Weissing, F., Pen, I., & Keller, L. (2011). An evolutionary perspective on self-organized division of labor in social insects. *Annual Review of Ecology, Evolution, and Systematics*, 42(1), 91–110.
- Duarte, M., Costa, V., Gomes, J., Rodrigues, T., Silva, F., Oliveira, S., & Christensen, A. (2016). Evolution of collective behaviors for a real swarm of aquatic surface robots. *PLOS ONE*, 11(3), 1–25. <https://doi.org/10.1371/journal.pone.0151834>
- Ducatelle, F., Di Caro, G., & Gambardella, L. (2010). Cooperative self-organization in a heterogeneous swarm robotic system. *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation*, 87–94. <https://doi.org/10.1145/1830483.1830501>
- Ducatelle, F., Di Caro, G., Pinciroli, C., Mondada, F., & Gambardella, L. (2011). Communication assisted navigation in robotic swarms: Self-organization and cooperation. *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4981–4988. <https://doi.org/10.1109/IROS.2011.6094454>
- Dunbar, R. (2009). The social brain hypothesis and its implications for social evolution. *Annals of human biology*, 36(5), 562–572.
- Eberhart, R., Shi, Y., & Kennedy, J. (2001). *On our nonexistence as entities: The social organism*. Elsevier Science. <https://books.google.co.za/books?id=Z1z3byh1JAsC>
- Eiben, A., Smith, J., et al. (2003). *Introduction to evolutionary computing* (Vol. 53). Springer.
- Farina, M. (2012). Louise barrett, beyond the brain: How body and environment shape animal and human minds.
- Farrugia, J., & Fabri, S. (2018). Swarm robotics for object transportation. *2018 UKACC 12th International Conference on Control (CONTROL)*, 353–358. <https://doi.org/10.1109/CONTROL.2018.8516829>
- Ferrante, E., Duéñez-Guzmán, E., Turgut, A., & Wenseleers, T. (2013). Geswarm: Grammatical evolution for the automatic synthesis of collective behaviors in swarm robotics. *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, 17–24.
- Ferrante, E., Turgut, A., Duéñez-Guzmán, E., Dorigo, M., & Wenseleers, T. (2015). Evolution of self-organized task specialization in robot swarms. *PLoS computational biology*, 11(8), e1004273.
- Ficici, S., Watson, R., & Pollack, J. (1999). Embodied evolution: A response to challenges in evolutionary robotics. *Proceedings of the eighth European workshop on learning robots*, 14–22.
- Fine, B., & Shell, D. (2013). Unifying microscopic flocking motion models for virtual, robotic, and biological flock members. *Autonomous Robots*, 35(2), 195–219. <https://doi.org/10.1007/s10514-013-9338-z>
- Fjerdingstad, E., & Crozier, R. (2006). The evolution of worker caste diversity in social insects. [PMID: 16673347]. *The American Naturalist*, 167(3), 390–400. <https://doi.org/10.1086/499545>
- Floreano, D., & Wood, R. (2015). Science, technology and the future of small autonomous drones. *nature*, 521(7553), 460–466.
- Furman, A., Nagar, D., & Nitschke, G. (2019). Automating collective robotic system design. *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*, 1465–1472. <https://doi.org/10.1109/SSCI44817.2019.9002925>

- Galassi, M., Capodieci, N., Cabri, G., & Leonardi, L. (2016). Evolutionary strategies for novelty-based online neuroevolution in swarm robotics. *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 002026–002032.
- Garcia-Aunon, P., & Cruz, A. (2018). Control optimization of an aerial robotic swarm in a search task and its adaptation to different scenarios. *Journal of Computational Science*, 29, 107–118. <https://doi.org/https://doi.org/10.1016/j.jocs.2018.10.004>
- Garcia-Aunon, P., Roldán, J., & Barrientos, A. (2019). Monitoring traffic in future cities with aerial swarms: Developing and optimizing a behavior-based surveillance algorithm. *Cognitive Systems Research*, 54, 273–286. <https://doi.org/https://doi.org/10.1016/j.cogsys.2018.10.031>
- Gebhardt, G., Daun, K., Schnaubelt, M., & Neumann, G. (2018). Learning robust policies for object manipulation with robot swarms. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 7688–7695. <https://doi.org/10.1109/ICRA.2018.8463215>
- Gomes, J., Mariano, P., & Christensen, A. (2017). Novelty-Driven Cooperative Coevolution. *Evolutionary Computation*, 25(2), 275–307. <https://doi.org/10.1162/EVCO.a.00173>
- Gomes, J., Urbano, P., & Christensen, A. (2013). Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*, 7(2), 115–144. <https://doi.org/10.1007/s11721-013-0081-z>
- Gordon, D., Guetz, A., Greene, M., & Holmes, S. (2011). Colony variation in the collective regulation of foraging by harvester ants. *Behavioral Ecology*, 22(1), 429–435.
- Gupta, R., & Bayal, R. (2020). Source detection of oil spill using modified glowworm swarm optimization. *2020 5th International Conference on Computing, Communication and Security (ICCCS)*, 1–6. <https://doi.org/10.1109/ICCCS49678.2020.9276960>
- Haasdijk, E., Bredeche, N., & Eiben, A. (2014). Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics. *PLoS ONE*, 9(6). <https://doi.org/10.1371/journal.pone.0098466>
- Hamann, H. (2018). *Swarm robotics: A formal approach*. <https://doi.org/10.1007/978-3-319-74528-2>
- Hansen, N., & Ostermeier, A. (1996). Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation. *Proceedings of IEEE International Conference on Evolutionary Computation*, 312–317. <https://doi.org/10.1109/ICEC.1996.542381>
- Hart, A., Anderson, C., & Ratnieks, F. (2002). Task partitioning in leafcutting ants. *Acta ethologica*, 5, 1–11.
- Hart, E., Steyven, A., & Paechter, B. (2018). Evolution of a functionally diverse swarm via a novel decentralised quality-diversity algorithm. *Proceedings of the Genetic and Evolutionary Computation Conference*. <https://doi.org/10.1145/3205455.3205481>
- Hauser, H. (2019). Resilient machines through adaptive morphology. *Nature Machine Intelligence*, 1(1), 338–339.
- Hewland, J., & Nitschke, G. (2015). The benefits of adaptive behavior and morphology for cooperation. *2015 IEEE Symposium Series on Computational Intelligence*, 1047–1054. <https://doi.org/10.1109/SSCI.2015.151>
- Hiraga, M., Wei, Y., & Ohkura, K. (2020). Evolving collective cognition for object identification in foraging robotic swarms. *Artificial Life and Robotics*, 26. <https://doi.org/10.1007/s10015-020-00628-0>
- Hiraga, M., Yasuda, T., & Ohkura, K. (2018). Evolutionary acquisition of autonomous specialization in a path-formation task of a robotic swarm. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 22, 621–628. <https://doi.org/10.20965/jaciii.2018.p0621>
- Hodan, G. (n.d.). Flying pigeons [This work is licensed under the Creative Commons CC0 1.0 Universal Public Domain Dedication. To view a copy of this license, visit <https://creativecommons.org/>

- [publicdomain/zero/1.0/](https://www.publicdomainpictures.net/en/view-image.php?image=20195&picture=flying-pigeons)]. <https://www.publicdomainpictures.net/en/view-image.php?image=20195&picture=flying-pigeons>
- Hornby, G., Lohn, J., & Linden, D. (2011). Computer-automated evolution of an x-band antenna for nasa's space technology 5 mission. *Evolutionary computation*, 19(1), 1–23.
- Howard, A., Matarić, M., & Sukhatme, G. (2002). Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem. In H. Asama, T. Arai, T. Fukuda, & T. Hasegawa (Eds.), *Distributed autonomous robotic systems 5* (pp. 299–308). Springer Japan.
- Hunt, E. (2021). Phenotypic plasticity provides a bioinspiration framework for minimal field swarm robotics. *Frontiers in Robotics and AI*, 7(23), doi:10.3389/frobt.2020.00023.
- Hüttenrauch, M., Adrian, S., Neumann, G., et al. (2019). Deep reinforcement learning for swarm systems. *Journal of Machine Learning Research*, 20(54), 1–31.
- Isler, K., & van Schaik, C. (2009). The expensive brain: A framework for explaining evolutionary changes in brain size. *Journal of human evolution*, 57(4), 392–400.
- Jackson, D., & Ratnieks, F. (2006). Communication in ants. *Current Biology*, 16(15), R570–R574. <https://doi.org/https://doi.org/10.1016/j.cub.2006.07.015>
- Jandt, J., & Dornhaus, A. (2014). Bumblebee response thresholds and body size: Does worker diversity increase colony performance? *Animal Behaviour*, 87, 97–106. <https://doi.org/https://doi.org/10.1016/j.anbehav.2013.10.017>
- Jelisavcic, M., de Carlo, M., Hupkes, E., Eustratiadis, P., Orlowski, J., Haasdijk, E., Auerbach, J., & Eiben, A. (2017). Real-World Evolution of Robot Morphologies: A Proof of Concept. *Artificial Life*, 23(2), 206–235. <https://doi.org/10.1162/ARTL.a.00231>
- Joachimczak, M., Suzuki, R., & Arita, T. (2015). Improving evolvability of morphologies and controllers of developmental soft-bodied robots with novelty search. *Frontiers in Robotics and AI*, 2. <https://doi.org/10.3389/frobt.2015.00033>
- Joachimczak, M., Suzuki, R., & Arita, T. (2016). Artificial Metamorphosis: Evolutionary Design of Transforming, Soft-Bodied Robots. *Artificial Life*, 22(3), 271–298. <https://doi.org/10.1162/ARTL.a.00207>
- Just, W., & Moses, M. (2017). Flexibility through autonomous decision-making in robot swarms. *Proceedings of the IEEE Symposium Series on Computational Intelligence*, 1–8.
- Kaushik, R., Desreumaux, P., & Mouret, J.-B. (2020). Adaptive prior selection for repertoire-based online adaptation in robotics. *Frontiers in Robotics and AI*, 6. <https://doi.org/10.3389/frobt.2019.00151>
- Kelly, S., Panhuis, T., & Stoehr, M. (2012). Phenotypic plasticity: Molecular mechanisms and adaptive significance. *Comprehensive Physiology*, 2(1), 1417–1439.
- Kengyel, D., Hamann, H., Zahadat, P., Radspieler, G., Wotawa, F., & Schmickl, T. (2015). Potential of heterogeneity in collective behaviors: A case study on heterogeneous swarms. https://doi.org/10.1007/978-3-319-25524-8_13
- Kernbach, S. (2012). *Handbook of collective robotics: Fundamentals and challenges*. <https://doi.org/10.4032/9789814364119>
- Kistemaker, S., & Whiteson, S. (2011). Critical factors in the performance of novelty search. *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, 965–972. <https://doi.org/10.1145/2001576.2001708>
- Kratochvil, P. (n.d.). Fish underwater [This work is licensed under the Creative Commons CC0 1.0 Universal Public Domain Dedication. To view a copy of this license, visit <https://creativecommons.org/publicdomain/zero/1.0/>]. <https://www.publicdomainpictures.net/en/view-image.php?image=12451&picture=fish-underwater>

- Kriegman, S., Cheney, N., & Bongard, J. (2018). How morphological development can guide evolution. *Scientific Reports*, 8(1), 13934. <https://doi.org/10.1038/s41598-018-31868-7>
- Kriegman, S., Walker, S., Shah, D., Levin, M., Kramer-Bottiglio, R., & Bongard, J. (2019). Automated shapeshifting for function recovery in damaged robots. *Robotics: Science and Systems XV*. <https://doi.org/10.15607/rss.2019.xv.028>
- Lanctot, M., Zambaldi, V., Gruslys, A., Lazaridou, A., Tuyls, K., Pérolat, J., Silver, D., & Graepel, T. (2017). A unified game-theoretic approach to multiagent reinforcement learning. *Advances in neural information processing systems*, 30.
- Laurent, G., Matignon, L., Fort-Piat, L., et al. (2011). The world of independent learners is not markovian. *International Journal of Knowledge-based and Intelligent Engineering Systems*, 15(1), 55–64.
- Lehman, J., & Stanley, K. (2010). Revising the evolutionary computation abstraction: Minimal criteria novelty search. *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation*, 103–110. <https://doi.org/10.1145/1830483.1830503>
- Lehman, J., & Stanley, K. (2011a). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2), 189–223. https://doi.org/10.1162/EVCO_a.00025
- Lehman, J., & Stanley, K. (2008). Exploiting open-endedness to solve problems through the search for novelty. *Proceedings of the Eleventh International Conference on Artificial Life (Alife XI)*.
- Lehman, J., & Stanley, K. (2011b). Evolving a diversity of virtual creatures through novelty search and local competition. *Proceedings of the 13th annual conference on Genetic and evolutionary computation - GECCO 11*. <https://doi.org/10.1145/2001576.2001606>
- Levi, P., & Kernbach, S. (2010). *Symbiotic multi-robot organisms: Reliability, adaptability, evolution* (Vol. 7). Springer Science & Business Media.
- Lipson, H., Sunspiral, V., Bongard, J., & Cheney, N. (2016). *On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures* (Vol. ALIFE 2016, the Fifteenth International Conference on the Synthesis and Simulation of Living Systems). <https://doi.org/10.1162/978-0-262-33936-0-ch042>
- Long, J. (2012). *Darwin's devices: What evolving robots can teach us about the history of life and the future of technology*. Basic Books (AZ).
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Pieter Abbeel, O., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30.
- Mann, H., & Whitney, D. (1947). On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, 50–60.
- Mathews, N., Christensen, A., O'Grady, R., Mondada, F., & Dorigo, M. (2017). Mergeable nervous systems for robots. *Nature communications*, 8(1), 439.
- Mathews, N., Christensen, A., Stranieri, A., Scheidler, A., & Dorigo, M. (2019). Supervised morphogenesis: Exploiting morphological flexibility of self-assembling multirobot systems through cooperation with aerial robots. *Robotics and autonomous systems*, 112, 154–167.
- McShea, D. (1996). Perspective: Metazoan complexity and evolution: Is there a trend? *Evolution*, 50(1), 477–492.
- Miras, K., & Eiben, A. (2019). Effects of environmental conditions on evolved robot morphologies and behavior. *Proceedings of the Genetic and Evolutionary Computation Conference*, 125–132.
- Miras, K., Ferrante, E., & Eiben, A. (2020). Environmental influences on evolvable robots. *PloS one*, 15(5), e0233848.
- Misir, O., & Gökrem, L. (2021). Dynamic interactive self organizing aggregation method in swarm robots. *Biosystems*, 207, 104451. <https://doi.org/https://doi.org/10.1016/j.biosystems.2021.104451>

- Mondada, F., Pettinaro, G., Guignard, A., Kwee, I., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L., & Dorigo, M. (2004). Swarm-bot: A new distributed robotic concept. *Autonomous robots*, 17, 193–221.
- Montanier, J.-M., Carrignon, S., & Bredeche, N. (2016). Behavioral specialization in embodied evolutionary robotics: Why so difficult? *Frontiers in Robotics and AI*, 3. <https://doi.org/10.3389/frobt.2016.00038>
- Montes de Oca, M., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., & Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5, 305–327. <https://doi.org/10.1007/s11721-011-0062-z>
- Moore, A., Brodie, E., & Wolf, J. (1997). Interacting phenotypes and the evolutionary process: I. direct and indirect genetic effects of social interactions. *Evolution*, 51(1), 1352–1362.
- Mouret, J.-B. (2011). Novelty-based multiobjectivization. In *New horizons in evolutionary robotics* (pp. 139–154). Springer.
- Mouret, J.-B. (2020). Evolving the behavior of machines: From micro to macroevolution. *iScience*, 23(11), 101731. <https://doi.org/https://doi.org/10.1016/j.isci.2020.101731>
- Mouret, J.-B., & Clune, J. (2015). Illuminating search spaces by mapping elites. *CoRR*, abs/1504.04909. <http://arxiv.org/abs/1504.04909>
- Nagar, D., Furman, A., & Nitschke, G. (2019a). The cost of big brains in groups. *Proceedings of the 2019 Conference on Artificial Life*, 404–411.
- Nagar, D., Furman, A., & Nitschke, G. (2019b). The cost of complexity in robot bodies. *Proceedings of the IEEE Congress on Evolutionary Computation*, 2713–2720.
- Narumi, T., Uemichi, K., Honda, H., & Osaki, K. (2018). Self-organization at the first stage of honeycomb construction: Analysis of an attachment-excavation model. *PLOS ONE*, 13(10), 1–15. <https://doi.org/10.1371/journal.pone.0205353>
- Navarro, I., & Matía, F. (2012). An introduction to swarm robotics. *ISRN Robotics*, 2013, 608164. <https://doi.org/10.5402/2013/608164>
- Nitschke, G., Eiben, A., & Schut, M. (2011). Evolving behavioral specialization in robot teams to solve a collective construction task. *Swarm and Evolutionary Computation*, 2(1), 25–38.
- Nitschke, G., & Didi, S. (2017). Evolutionary policy transfer and search methods for boosting behavior quality: Robocup keep-away case study. *Frontiers in Robotics and AI*, 4, 62.
- Nitschke, G., Eiben, A., & Schut, M. (2012). Evolving team behaviors with specialization. *Genetic Programming and Evolvable Machines*, 13(4), 493–536.
- Nitschke, G., Schut, M., & Eiben, A. (2010). Collective neuro-evolution for evolving specialized sensor resolutions in a multi-rover task. *Evolutionary Intelligence*, 3(1), 13–29.
- Nogueira, Y., de Brito, C., Vidal, C., & Cavalcante-Neto, J. (2020). Towards intrinsic autonomy through evolutionary computation. *Artificial Intelligence Review*, 53(6), 4449–4473. <https://doi.org/10.1007/s10462-019-09798-1>
- Nolfi, S., Bongard, J., Husbands, P., & Floreano, D. (2016). Evolutionary robotics. In B. Siciliano & O. Khatib (Eds.), *Springer handbook of robotics* (pp. 2035–2068). Springer International Publishing. https://doi.org/10.1007/978-3-319-32552-1_76
- Nordmoen, J., Veenstra, F., Ellefsen, K., & Glette, K. (2021). Map-elites enables powerful stepping stones and diversity for modular robotics. *Frontiers in Robotics and AI*, 8, 639173.
- Nygaard, T., Martin, C., Howard, D., Torresen, J., & Glette, K. (2021a). Environmental adaptation of robot morphology and control through real-world evolution. *Evolutionary Computation*, 29(4), 441–461.

- Nygaard, T., Martin, C., Torresen, J., Glette, K., & Howard, D. (2021b). Real-world embodied ai through a morphologically adaptive quadruped robot. *Nature Machine Intelligence*, 3(5), 410–419. <https://doi.org/10.1038/s42256-021-00320-3>
- O’Shea-Wheller, T., Hunt, E., & Sasaki, T. (2020). Functional Heterogeneity in Superorganisms: Emerging Trends and Concepts. *Annals of the Entomological Society of America*, 114(5), 562–574. <https://doi.org/10.1093/aesa/saaa039>
- Otte, M. (2018). 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), 1017–1061.
- Panait, L., & Luke, S. (2005). Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3), 387–434. <https://doi.org/10.1007/s10458-005-2631-2>
- Pfeifer, R., & Bongard, J. (2006). *How the body shapes the way we think: A new view of intelligence*. MIT press.
- Picardi, G., Hauser, H., Laschi, C., & Calisti, M. (2021). Morphologically induced stability on an underwater legged robot with a deformable body. *The International Journal of Robotics Research*, 40(1), 435–448. <https://doi.org/10.1177/0278364919840426>
- Potter, M., Meeden, L., Schultz, A., et al. (2001). Heterogeneity in the coevolved behaviors of mobile robots: The emergence of specialists. *International joint conference on artificial intelligence*, 17(1), 1337–1343.
- Prorok, A., Hsieh, A., & Kumar, V. (2016). Fast redistribution of a swarm of heterogeneous robots. *Proceedings of the 9th EAI International Conference on Bio-Inspired Information and Communications Technologies (Formerly BIONETICS)*, 249–255. <https://doi.org/10.4108/eai-3-12-2015.2262349>
- Pugh, J., Soros, L., Szerlip, P., & Stanley, K. (2015). Confronting the challenge of quality diversity. *Proceedings of the 2015 on Genetic and Evolutionary Computation Conference - GECCO 15*. <https://doi.org/10.1145/2739480.2754664>
- Pugh, J., Soros, L., & Stanley, K. (2016a). Searching for quality diversity when diversity is unaligned with quality. *Parallel Problem Solving from Nature-PPSN XIV: 14th International Conference, Edinburgh, UK, September 17-21, 2016, Proceedings 14*, 880–889.
- Pugh, J. K., Soros, L. B., & Stanley, K. (2016b). Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3. <https://doi.org/10.3389/frobt.2016.00040>
- Rubenstein, M., Cornejo, A., & Nagpal, R. (2014). Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198), 795–799.
- Şahin, E. (2005). Swarm robotics: From sources of inspiration to domains of application. In E. Şahin & W. M. Spears (Eds.), *Swarm robotics* (pp. 10–20). Springer Berlin Heidelberg.
- Sapolsky, R. (2017). *Behave: The biology of humans at our best and worst*. Penguin.
- Schlichting, C., & Pigliucci, M. (1998). *Phenotypic evolution: A reaction norm perspective*. Sinauer Associates Incorporated.
- Schranz, M., Umlauf, M., Sende, M., & Elmenreich, W. (2020). Swarm robotic behaviors and current applications. *Frontiers in Robotics and AI*, 7. <https://doi.org/10.3389/frobt.2020.00036>
- Shan, Q., & Mostaghim, S. (2021). Achieving task allocation in swarm intelligence with bi-objective embodied evolution. *Swarm Intelligence*, 15(3), 287–310. <https://doi.org/10.1007/s11721-021-00198-2>
- Silva, F., Urbano, P., Correia, L., & Christensen, A. (2015). Odneat: An algorithm for decentralised online evolution of robotic controllers. *Evolutionary Computation*, 23(3), 421–449. https://doi.org/10.1162/EVCO_a.00141

- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., & Riedmiller, M. (2014). Deterministic policy gradient algorithms. *International conference on machine learning*, 387–395.
- Slavkov, I., Carrillo-Zapata, D., Carranza, N., Diego, X., Jansson, F., Kaandorp, J., Hauert, S., & Sharpe, J. (2018). Morphogenesis in robot swarms. *Science Robotics*, 3(25), eaau9178.
- Southwood, T. (1977). Habitat, the templet for ecological strategies? *Journal of animal ecology*, 46(2), 337–365.
- Spanellis, C., Stewart, B., & Nitschke, G. (2021). The environment and body-brain complexity. *Proceedings of the Genetic and Evolutionary Computation Conference*, 138–145.
- Stern, D. (2013). The genetic causes of convergent evolution. *Nature Reviews Genetics*, 14(11), 751–764.
- Steyven, A., Hart, E., & Paechter, B. (2017). An investigation of environmental influence on the benefits of adaptation mechanisms in evolutionary swarm robotics. *Proceedings of the Genetic and Evolutionary Computation Conference*, 155–162.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J., Tuyls, K., et al. (2017). Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296*.
- Sutton, R., & Barto, A. (2018). *Reinforcement learning: An introduction*. MIT press.
- Swenty, S. (2012). Honey bees swarm [This work is licensed under the Creative Commons Attribution 2.0 Generic (CC BY 2.0). To view a copy of this license, visit <https://creativecommons.org/licenses/by/2.0/deed.en>]. [https://commons.wikimedia.org/wiki/File:Honey_Bees_Swarm_\(8159050858\).jpg](https://commons.wikimedia.org/wiki/File:Honey_Bees_Swarm_(8159050858).jpg)
- Trueba, P., Prieto, A., Bellas, F., Caamaño, P., & Duro, R. (2013). Specialization analysis of embodied evolution for robotic collective tasks [Collective and Social Autonomous Robots]. *Robotics and Autonomous Systems*, 61(7), 682–693. <https://doi.org/https://doi.org/10.1016/j.robot.2012.08.005>
- Trujillo, L., & et al. (2011). Speciation in behavioral space for evolutionary robotics. *Journal of Intelligent and Robotic Systems*, 64(1), 323–351.
- Turgut, A., Çelikkanat, H., Gökçe, F., & Sahin, E. (2008). Self-organized flocking in mobile robot swarms. *Swarm Intelligence*, 2, 97–120. <https://doi.org/10.1007/s11721-008-0016-2>
- van Diggelen, F., & et al. (2022). Diversity in swarm robotics with task-independent behavior characterization. *Proceedings of the Genetic and Evolutionary Computation Conference*, 31–39.
- Vásárhelyi, G., Virágh, C., Somorjai, G., Nepusz, T., Eiben, A., & Vicsek, T. (2018). Optimized flocking of autonomous drones in confined environments. *Science Robotics*, 3(20), eaat3536.
- Wallraff, H. (2010). *Avian navigation: Pigeon homing as a paradigm* (Vol. 208). <https://doi.org/10.1007/b137573>
- Watkins, C., & Dayan, P. (1992). Q-learning. *Machine learning*, 8, 279–292.
- Watson, J., & Nitschke, G. (2015). Evolving robust robot team morphologies for collective construction. *2015 IEEE Symposium Series on Computational Intelligence*, 1039–1046.
- Weigmann, K. (2012). Does intelligence require a body? the growing discipline of embodied cognition suggests that to understand the world, we must experience the world. *EMBO reports*, 13(12), 1066–1069.
- West-Eberhard, M. (1989). Phenotypic plasticity and the origins of diversity. *Annual Review of Ecology and Systematics*, 20(1), 249–278.
- Wild, A. (2007). Army ants (*labidus praedator*) along pipeline road, gamboa, panama. [This work is licensed under the Creative Commons CC0 1.0 Universal Public Domain Dedication. To view a copy of this license, visit <https://creativecommons.org/publicdomain/zero/1.0/>]. https://commons.wikimedia.org/wiki/File:Army_ants_swarming.jpg

- Wolf, J., Brodie, E., & Moore, A. (1999). Interacting phenotypes and the evolutionary process. ii. selection resulting from social interactions. *The American Naturalist*, 153(1), 254–266.
- Xu, F., & Wang, H. (2021). Soft robotics: Morphology and morphology-inspired motion strategy. *IEEE/CAA Journal of Automatica Sinica*, 8(9), 1500–1522.
- Zardini, E., Zappetti, D., Zambrano, D., Iacca, G., & Floreano, D. (2021). Seeking quality diversity in evolutionary co-design of morphology and control of soft tensegrity modular robots. *Proceedings of the Genetic and Evolutionary Computation Conference*. <https://doi.org/10.1145/3449639.3459311>
- Zhang, K., Yang, Z., & Başar, T. (2021). Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of reinforcement learning and control*, 321–384.
- Zhang, T., Zhang, W., & Gupta, M. (2017). Resilient robots: Concept, review, and future directions. *Robotics*, 6(4), 22. <https://doi.org/10.3390/robotics6040022>
- Zoss, B., Mateo, D., Kuan, Y., Tokić, G., Chamanbaz, M., Goh, L., Valleggra, F., Bouffanais, R., & Yue, D. (2018). Distributed system of autonomous buoys for scalable deployment and monitoring of large waterbodies. *Autonomous Robots*, 42. <https://doi.org/10.1007/s10514-018-9702-0>