Optimizations of Morphology, Behaviour and Evolution in Multi-Agent Systems for Predator-Prey Pursuit Problem

Milen Svetoslavov Georgiev

Department of Information and Computer Science

Graduate School of Science and Engineering

Doshisha University

July 2020

Table of Contents

Abstract

Multi-agent systems (MAS) are widely applied for problem solving, software engineering, and the simulation of (human, robotic, etc.) societies. Compared to a monolith system, owing to their complex, non-linear nature, MAS can often provide efficient solutions to complex problems. However, the main challenge of applying MAS is that due to the enormous semantic gap between the properties of (i) entities (agents) and (ii) the system as a whole, it is difficult to obtain an optimal solution to the problem analytically.

The objective of our research is to investigate two relatively orthogonal ways of improving the overall performance of MAS: (i) minimizing the time needed by MAS to solve a given problem by evolutionary optimizing (coevolving) both the morphology and behaviour of agents, and (ii) minimizing the runtime needed by evolutionary framework – genetic programming – to successfully accomplish such a coevolution. As an evolutionary framework we adopted the in-house XML-based genetic programming (XGP), which offers a flexible, human-readable, and cross-application compatible XML representation of the genotype of evolved agents.

The application domain is the well-known, but difficult to solve predator-prey pursuit problem (PPPP) comprising a team of predator agents, that needs to capture a prey by surrounding it in a simulated two-dimensional world. We considered two instances of PPPP featuring predator agents with different abilities. The first instance, inspired by the opportunity to challenge the relevance of the "average" (rather than the individual) abilities of agents, comprises relatively complex, reactive predator agents with continuous sensory (morphological) and moving abilities. The second instance – comprising absolutely simple reactive predator agents (that do not even compute in their decision making) with rather primitive, discrete sensory and moving capabilities

(a single line-of-sight sensor and two thrusters in a differential drive configuration, respectively) – is motivated by our intention to model the recently emerged nano- and micro robots and their potential applications in biomedicine.

The experimental results obtained from the evolution of the team of complex predator agents in PPPP indicate that, indeed, MAS with lower values of the average (mean) sensory- and moving capabilities of agents could have a superior performance compared to that of MAS with higher average values. This finding is consistent with the concept of the "end of average" arguing that the combination of individual qualities of entities in complex system matters more than the average value of these qualities. From another standpoint, the results could be seen as a verification of the survival value of the diversity of qualities of entities in complex systems, such as MAS.

The results obtained from the evolution of the team of simple predator agents suggest that an asymmetric morphology (i.e., an angular offset of the line-of-sight sensor) coevolved with an intricate "driving" behaviour of predator agents, results in a most efficient behaviour of the agents. The results confirm that even the complex problems, such as PPPP, could be solved by the team of extremely simple predator agents if their morphology and behaviour are developed by means of simulated evolution. On the other hand, the experimental results also indicate that efficiency of this evolution depends on the size of the evolved population of agents. Counterintuitively, the smaller sizes of populations – due to lower genotypic redundancy and favourable cache-related effects during the simulation – result in faster overall runtime of evolution.

Presented research could be viewed as a step towards the development of MAS that are capable of solving complex problems efficiently, having a lower evolutionary overhead.

Acknowledgment

I want to express my gratitude to Doshisha for providing the facilities and equipment needed for conducting this research, and to the Japanese Ministry of Culture, Sports, Science and Education (MEXT), for the provided financial support. Also I would like to thank Professor Ivan Tanev and Professor Katsunori Shimohara, for their continuous support.

List of Figures

List of Tables

Chapter 1: Introduction

1.1 Background

Throughout the years, multi-agent systems (MAS) have become important in solving complex problems where monolith (single entity) systems are unable to produce an acceptable solution in terms of speed or resource requirements. Base components of MAS are a world (environment), entities (agents), relations between the entities, a way with which they perceive the world, a set of operations and functions of the agents and changes in the world as a result of their execution. The main applications of MAS are problem solving, simulation, collective robotics, software engineering and construction of synthetic worlds [1].

1.2 Objective of research

The objective of our research is to investigate two relatively orthogonal ways of improving the overall performance of MAS: (i) minimizing the time needed by MAS to solve a given problem by evolutionary optimizing (coevolving) both the morphology and behaviour of agents, and (ii) minimizing the runtime needed by evolutionary framework – genetic programming – to successfully accomplish such a coevolution. As an evolutionary framework we adopted the in-house XML-based genetic programming (XGP), which offers a flexible, human-readable, and cross-application compatible XML representation of the genotype of evolved agents.

For this reason, we have investigated the performance of heterogeneous multi-agent systems of agents in comparison to morphologically identical homogeneous systems, pertaining the same average physical and sensory abilities for the system as a whole. Furthermore, we have suggested improvements to the system, such as simplifying the morphology, lowering the evolutionary overhead and minimizing the solution search space.

1.3 Motivation for the research

Out motivation for this research comes from a need to develop cheaper, better performance MAS, for use in everyday real-life situations. During our research, we've found numerous ways to improve the efficacy of multiple aspects of MAS in PPPP.

1.4 The predator-prey pursuit problem

We chose to use the PPPP, because it's a well-known and extensively studied problem, which features hunting agents (Predators) - ideally modelled by a MAS, and a Prey agent.

1.5 Methodology

We've developed a testing system, having two sub-components – evolutionary system and evaluation system. Using the adopted method, we've developed additionally two MAS, featuring complex and simple agents. Each of the bringing improvements to the traditional solutions of PPPP.

1.5.1 Homogeneous complex MAS

We've developed an in-house heterogeneous MAS, that we will use as a base for comparison, of the suggested improvements.

1.5.2 Heterogeneous complex MAS

In the first part of this research, we've shown improvements based on the real-life evidence, that the average capabilities of a system do not always define the system as a whole and a better disparity in individual capabilities of its consisting entities can bring better results in solving the problems at hand [2].

1.5.3 Simple MAS

In the second part of the research, we’ve simplified the agents as much as possible, in an attempt to lower the resources, need for their creation and operation to a minimum. After seeing their success, we’ve continues improving the agents by suggesting changes to their morphology and the evolution of their phenotype.

1.5.4 Evolutionary system

We developed an evolutionary computing framework – e.g., genetic programming, that could be used to evolve such a behaviour of agents that result in best efficiency (performance) of the multi-agent system as a whole. The use of genetically evolved solutions will make our work more realistic than commonly considered previous work [3] [4]. Further, we will evolve this multi-agent system for various combinations of the individual abilities of the agents (and for various results of the average of these abilities) and investigate the obtained optimal performance of the whole system.

1.5.5 Evaluation system

In consistency with the MAS model, each type of agents received their own evaluation system – “World”, which helped determine their performance and behavioural qualities.

Chapter 2: Introduction to genetic programming and genetic algorithms

2.1 Basic concepts

Finding a solution to a problem that requires exploration of a big search space may be challenging. Genetic and evolutionary programming can help us overcome this challenge, by developing strategies to do that exploration more efficiently. In genetic programming, the generated solutions take form of computer programs which are then evaluated to determine their fitness – their ability to solve the problem at hand. Evolutionary programming is similar to genetic programming with the difference that, in this case, the structure of the program is fixed and only its parameters are being evolved.

2.2 Selection algorithm

The selection algorithm's purpose is to choose which individuals of the population at every generation will survive to the next generation. Commonly this requires ranking the performance of the generated individuals using an arbitrary measure such as their fitness to the environment (ability to solve a certain problem or quality of the solution following given rules).

2.3 Breeding algorithms

Breeding algorithms are the strategies used to generate an offspring based on an existing one. There are multiple ways to do that generation. For sake of simplicity, we will review only the methods that we have used in our research.

2.3.1 Crossover

Crossover in genetic programming, also called "recombination", is a genetic operation that is used to combine information from the phenotype of two "parents" to generate new offspring. It is one of the ways to generate new population, based on an existing one.

2.3.2 Mutation

Mutation is a genetic operation used to maintain genetic diversity. It can be used in combination with crossover to introduce a greater genetic variety to the offspring. Since mutation can change the solution entirely, compared to the previous generation, it's effects should be controlled to avoid reducing the evolution to a primitive random search.

2.4 Genetic representation

Just as all living organisms' characteristics are described by their DNA and RNA, we need a way to describe the properties of our genetic entities. Genotype is usually represented of linear multigenic chromosomes of fixed length and the phenotype takes the form of different expression trees.

2.4.1 Traditional genetic trees

Genetic programming uses tree-like structures as in memory representation. Usually binary trees are used, as operations with them are very simple to implement. Every tree node is an operator function and every terminal (leaf) node is an operand. These structures, while useful for representation, are hard to work with and difficult to read by humans. This is why we developed our in-house solutions – XGP and XGA.

2.4.2 Introducing XGP and XGA

For our implementation of the evolutionary framework, we've decided to use XML schema, to evolve Genetic Programs (XGP) for the controller of the complex agents and to evolve the values described in the morphology of the simple agents, using Genetic Algorithms (XGA). We chose to use XML, because it provides the following advantages compared to the traditional genetic trees [5]:

- The genetic operations can be performed using off the shelf, and programming language neutral, XML DOM parsing tools, thus allowing fast prototyping.
- The generated code is human readable
- Platform and programming language independent approach for descripting genetic structures
- Generic support for the representation of the grammar of a strongly-typed genetic program

2.5 Summary

Using genetic and evolutionary programming allows an easy way to automatically generate, evolve, evaluate and improve the controlling algorithms for a MAS. Moreover, describing the programs using a human readable format, such as XML, allows further analysis and improvement of the generated solutions. This level of automation saves time and resources, in attempt to find the solutions to complex problems as PPPP.

Chapter 3: Modelling the predator-prey pursuit problem featuring complex predator agents

3.1 Investigating a heterogeneous approach to MAS

Due to the various constraints (e.g. consensus problem, inter-agent credit-assignment problem, computationally heavy evolution due to the large search space, etc.) pertinent to the development of heterogeneous multi-agent systems (MAS) as a distributed problem-solving approach [6], the research on them is underrepresented compared to the alternative homogeneous implementations [7]. The main motivation of our current work is that, to the best of our knowledge, the comparative analysis of the effects of uncertainty and noise in the environment of heterogeneous and homogeneous multi-agent systems is not studied extensively enough. An additional motivation of our research is, consonant with the concept of the "end of average" that appreciates the difference from the (often – mediocre, and sometimes – even non-existing, statistically calculated) average in human societies [2], to investigate the importance (if any) of the diversity of individual capabilities of heterogeneous agents featuring the same average as the (identical) agents in analogous homogeneous systems.

As a model of the typical human being, its performance and efficiency in society, the most common image is characterized by a simple value – the value of the average of abilities of its respective members [2]. The concept is borrowed from electrical engineering where the average usually manifests the useful signal while any fluctuations from it are a result of random noise. For human societies, however, the average (of a given ability of the members of society) is not necessarily seen as a useful signal, but rather as a synthetic, and often – meaningless, value, that is not actually exhibited by the vast majority of the members of society. Similarly, the fluctuations from the average value (of a given ability) are far from noise, but rather – specific variations that characterize the identity and personality of these members and a trait that, in many cases, may contribute to solving challenging new, previously unknown problems.

What we've studied is the importance (if any) of the average for the efficiency of multi-agent systems as a model (yet, to very a limited extent) of human societies, in uncertain environments. In addition, we have investigated whether the diversity (even at the expense of reduced average) of abilities rather than their average plays an important role in building better-performing multi-agent systems.

The main application areas of MAS are problem solving, simulation, collective robotics, software engineering and modelling of synthetic worlds [1]. For this purpose, we developed heterogeneous MAS that models some of the important aspect of human societies, such as cooperation, collaboration, communication, and division of labour. We are also implemented an evolutionary computing framework – e.g., genetic programming (GP), that could be used to evolve such a behaviour of agents that results in best effectiveness (performance) of the multi-agent system as a whole. The use of genetically evolved solutions will make our work more realistic than commonly considered previous work [3] [4]. Further, we will evolve this multi-agent system for various combinations of the individual abilities of the agents (and for various results of the average of these abilities) and investigate the obtained optimal performance of the whole system.

Within the considered context, we will be employing the predator-prey pursuit problem (PPPP) to investigate the effect of the average and diversity on the MAS in evolution of the behaviour and performance in an unforeseen, randomly generated environment.

The efficiency of the multi-agent systems will be measured using a few factors – the number of different test scenarios that leads to a positive outcome (capturing the prey), overall fitness of the solution, found through genetic programming, and speed of evolution (the time needed to find an optimal solution for the specified number of test cases). Additionally, we will investigate the robustness of the evolved team of agents to newly presented, previously unknown initial situations.

3.2 Proposed approach

In this section, we will describe, in detail, the implementation of the complex predator and prey agents in the PPPP (predator-prey-pursuit problem), as well as the implementation of the world. Our initial implementation of the predator-prey-pursuit problem features a team of superior in terms of perception agents - predators, attempting to capture a single - more mobile agent – prey, as seen in Figure 1.

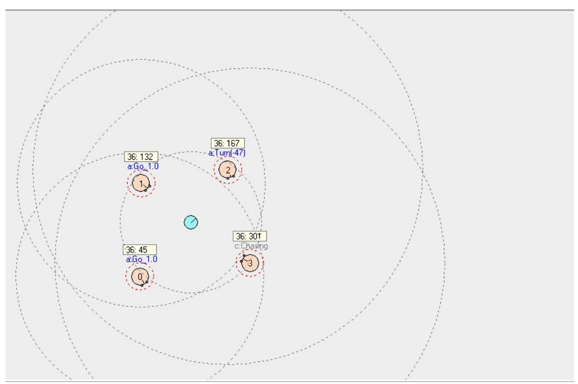

Figure 1:Predators surrounding a prey agent

3.3 Evolutionary framework

We will be using a previously developed custom implementation of a strongly typed genetic programming framework [8] for homogeneous MAS. We are striving to achieve heterogeneity by a morphological difference in spite of genetical similarity, this enables us to use the same framework to evolve both the heterogeneous and homogeneous systems of agents for our test cases, without additional changes.

The main parameters chosen for the settings of the evolutionary framework are shown in Table 1, below. The evolution continues until 50 generations have been evaluated, the evolution stagnated for 10 consecutive generations or an appropriate solution to all 10 initial situations was found. Every generation includes 400 different chromosomes, initially random generated for the first generation.

Table 1: Parameters of the GP framework

Parameter	Value
Population size	400 chromosomes
Selection	Binary tournament
Selection ratio	10%
Elite	Best 4 chromosomes
Crossover	Both single- and two-point
Mutation	Single-point
Mutation ratio	1-30%
Fitness cases	10 initial situations
Duration of the fitness trial	600 cycles per initial situation (300 seconds with 500ms sampling interval)
Fitness value	Sum of the average distance to the prey, average energy consumption and elapsed time for the trial. In addition, punishment is applied for large controllers. Agents are explicitly rewarded for capturing the prey
Termination criteria	(Fitness value<300 AND 10 successful situations) or (# Generations>50) or (Stagnation of fitness for 10 consecutive generations)

Since, initially, the agents won't be able to solve all 10 situations, to improve the computational efficiency of GA, the first test starts with 2 initial situations and the number of tested situations is increased by 2, every time the agents manage to solve n-1 situations, where n is the number of currently tested situations. If the GA is unable to produce a solution for the tested situations, we increase the number of situations until stagnation criteria or trial end conditions are met.

3.3.1 Genetic representation

The controlling programs for the predator agents are a set of IF-THEN stimuli-response rules. An example of an agent controller in XML format is shown in Figure 2.

```
<?xml version="1.0"?>
<GP xmlns:xs="http://www.w3.org/2001/XMLSchema-instance"
xs:noNamespaceSchemaLocation="..\SharedResources\GPSchema.xsd">
  <STM ind="3" age="1">
    <IF-THEN ind="4" age="1" histone="1">
      <COND-THEN ind="5" age="1">
        <COND_TBool ind="6" age="1">
          <OPER_TBool>not</OPER_TBool>
          <VAR_TBool>PeerVisible</VAR_TBool>
        </COND_TBool>
      </COND-THEN>
      <THEN ind="11" age="1">
        <STML ind="12" age="1">
          <STM ind="13" age="1">
            <COM-TURN ind="14" age="1">
              <COM-TURN-OPERAND ind="15" age="1">
                <SIGN>+</SIGN>
                <ANGLE-OP ind="18" age="1">
                  <CONST_TVisAngle_Pos>18</CONST_TVisAngle_Pos>
                </ANGLE-OP>
              </COM-TURN-OPERAND>
            </COM-TURN>
          </STM>
        </STML>
      </THEN>
    </IF-THEN>
  </STM>
</GP>
```

Figure 2: Sample agent controller code in XML format

This set of rules is represented as Document Object Model (DOM) parse tree structures, featuring, in addition, using plain-text XML encoding [5]. The DOM/XML representation allows us to perform the genetic operations using the API of an off-the-shelf, programming-language-neutral, XML DOM parser. The set of functions and terminals of the adopted GP are identical to the ones used in our previous work [8]. They are shown in Table 2 and Table 3, respectively.

Table 2: Function set of GP

Designation	Meaning
IF-THEN	stimuli-response IF-THEN rule
LE, GE, WI, EQ, NE, +, -	≤, ≥, Within, =, ≠, +, -

Table 3: Terminal set of GP

Category	Designation	Explanation
Sensory abilities	Prey_d; Peer_d	Distance to the prey and to the closest agent, *mm.*
	Prey_a; Peer_a	Bearing of the prey and of the closest agent, *degrees*
	PreyVisible; PeerVisible	True if prey (predator) agent is "visible", false otherwise
State variable	Speed	Speed of the agent, *mm/s*
Ephemeral constants	Integer	
Moving abilities	Turn(α)	Turns relatively to α degrees (α>0: clockwise)
	Stop, Go_1.0	Sets speed to 0, or to maximum, respectively
	Go_0.25, Go_0.5, Go_0.75	Sets speed to 25%, 50%, 75% of maximum

The execution of the example in Figure 2, of a behavioural stimuli-response IF-THEN rule would result in turning the predator 18 degrees to the right and setting if a peer (another predator agent) is not in sight.

3.3.2 Genetic operations

The breeding strategy is homogeneous in such a way that the performance of a single chromosome, cloned to all four agents is evaluated. The gene pool consists of 400 chromosomes.

We introduced a binary tournament selection as we consider it computationally efficient and simple to implement. We also adopted elitism in that the 10% of the best performing chromosomes of the current generation are copied unconditionally and are inserted in the mating pool for the next generation. A strongly typed crossover operation is defined in a way that only the nodes of the same data type (featuring an identical DOM/XML tag) from both parents can be swapped. Sub-tree mutation is also allowed, in a strongly typed way – a synthetically correct subtree can replace a random node in the genetic program.

3.4 Evaluation subsystem

We used a previously developed implementation of a strongly-typed GP framework [8] for homogeneous MAS. We intend to achieve heterogeneity in the behaviour of agents by means of exploiting their morphological – rather than their genotypic – differences. The possibility to exploit such a polymorphism to obtain a behavioural heterogeneity of genetically identical (homogeneous) predators allows us to employ the same evolutionary framework to evolve the team of predators in both – homogeneous and heterogeneous systems. We view this as an important argument in favour of the fairness of the presented comparative analysis of both systems.

3.4.1 Simulating the world

The simulated environment is a two-dimensional infinite toroidal world with size of 1600x1040mm (scaled). The perception range, decision making and resulting new state (location, orientation and speed) of the agents are updated with sampling interval of 500ms.

3.4.2 Fitness calculation

In order to evolve a general enough solution to the problem, the behaviour of the team of predators is evaluated on 10 initial situations. This allows us to avoid overfitting of the evolved agents to any particular situation, and to create a more robust system. In each of the evaluated situations, the prey is located in the centre of the world and oriented in a random direction. The agents are then randomly placed on the field in such a way, as to have a diverse set of situations, to avoid overfitting for a certain way of disposition.

The overall fitness for the particular chromosome is calculated as an average of the fitness values scored in each of the test situations for that run. The fitness for an initial situation is the sum of the average distance to the prey, average energy consumption and elapsed time for the trial. To avoid generation of very large controllers, a punishment is added to the fitness equal to 0.1xC, where C is the complexity of the agents' representation in tree nodes, estimated as the number of nodes in its parse-tree. Lower fitness values represent better performing team of predator agents. The criteria to end evolution is fitness under 300 and 10 successful situations, stagnation for 10 consecutive generations or 50 total generations tested.

3.5 Agent characteristics

3.5.1 Prey agent

The prey is a single agent with fixed behaviour using a handcrafted escaping strategy [8]. It responds to two situations – running away at 180 degrees angle from the nearest adversary, when a predator agent is visible and random wandering when there is not imminent threat. The maximum moving speed of the prey is higher than the maximum speed of the predators. The movement of the prey is continuous; it can turn left or right at any angle from its current direction. When chased, the prey is able to run at full speed, until its adversary is no longer in perception range. Features of the prey agent can be seen in Table 4.

Table 4: Features of the prey

Feature	Value
Number of prey agents	1
Diameter, mm	40
Maximum speed, mm/s	24
Type of sensor	Omnidirectional vision
Range of visibility of the sensor, mm	200

3.5.2 Predator agents

The team of predators consists of four agents with inferior moving abilities, compared to the prey. We do not consider the case in which the agents are superior in terms of speed, as capturing the prey in that condition seems to be trivial and a single agent will be able to accomplish it. In addition, if the prey agent is completely superior, it will be very hard or even impossible for our team of predators to capture it. Therefore, to give the chasing agents a chance to complete the given task - they are equipped with a sensor that provides them with a greater vision range than the prey. As the behaviour of the predator agents is not fixed like the prey agent, this will allow for emerging collective strategies during the evolution of the behaviour. As shown in Table 3, the predators also have the ability to adjust their moving speed to 0, 0.25, 0.5, 0.75 and 1.0 of their maximum speed, if they need to. The main features of the predator agents are shown in Table 5, for homogeneous multi-agent system and in

Table 6, for heterogeneous multi-agent system.

Table 5: Predator features in homogeneous MAS

Feature	Value
Number of predator agents	4
Diameter, mm	50
Max speed of predator agents, mm/s	16
Range of visibility of the sensor, mm	450
Type of sensor	360 degree (Omnidirectional)

Table 6: Features of the predators in heterogeneous MAS

Feature	Value
Number of predator agents	4
Diameter, mm	50
Max speed of predator agents, mm/s	16
Type of sensor	Omnidirectional vision
Range of visibility of the sensor of the two agents in Group 1, mm	*300-400*
Range of visibility of the sensor of the two agents in Group 2, mm	*500-600*

The heterogeneous system features two groups of two, morphologically identical agents (2 by 2 identical). Each of them will have a value of sensor range, such as to keep the average of all four agents of 450 – equal to the range of sensors of predators in homogeneous MAS.

We have chosen an arbitrary value of 450 for the average view range and 16 for the average speed. The homogeneous system will be serving as our control group and take on the average values, while the homogeneous system will be tested with different values for the range of the view sensor.

3.6 Experimental results

Our experiments involve 40 independent runs of GP for each one of the test cases of PPPP. Each test case involves different configuration of agents, and for heterogeneous system this implies a different combination of ranges of sensors of predators belonging to the two groups of agents. We considered a change of the range in intervals of 50 (e.g. 400-500, 350-550, 300-600) while keeping the value of the average of the range's constant (i.e., 450). In this section we will present the features of the behaviours of the team of predators, obtained from the evolution of an average homogeneous system, compared to the solutions of each of the heterogeneous configurations, evolved over 50 generations of GP. Additionally, we will review the generality of the evolved predator agents, as well as their robustness to a changing environment. We will discuss the problems that arise from the proposed approach to create a heterogeneous system based on disparity in morphology rather than changes in genotype of the predator agents and how they affect the general performance in a noisy or uncertain environment.

3.6.1 Evolution of the Homogeneous System

On average, the homogeneous system is able to solve all 10 initial situations only in 3 out of 40 (i.e., 7.5%) of runs. The evolution shows consistent development of controllers of predator agents, with fitness level averaging around 447 and with worst solution having a fitness of 465. However, best solution of the series ends the trial with fitness of 235 and 10 successful initial situations in twenty-first generation, as illustrated in Figure 3 and Figure 4.

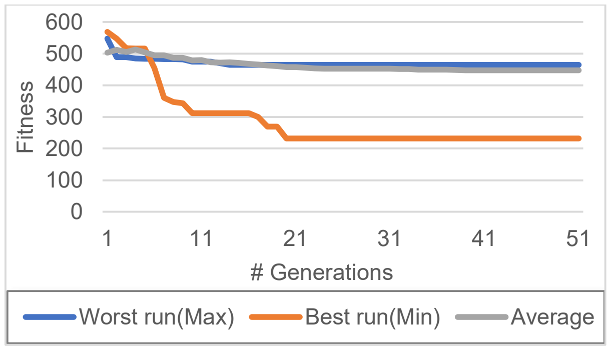

Figure 3: Convergence of fitness in the homogeneous MAS

The result of this experiment demonstrates that, while the best run is able to solve the problem with a reasonable effectiveness, the efficiency of evolution is rather poor, as the majority of the independent runs of GP could not reach the desired results. Moreover, as shown in Figure 5, even some of the runs could not resolve more than one initial situation.

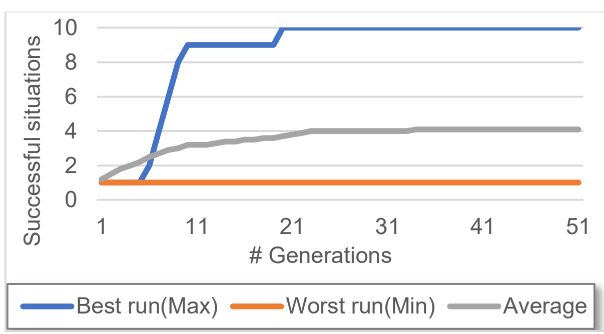

Figure 4: Dynamics of number of successfully solved initial situations by the homogeneous MAS.

To get a precise understanding of what is happening, we will be analysing the total number of successful situations for all 40 runs – shown in Figure 5. The distribution of success shows that more than 50% of the agents finish their evolution with 1 or 2 solved initial situations. While only 7.5% manage to solve all 10 initial situations

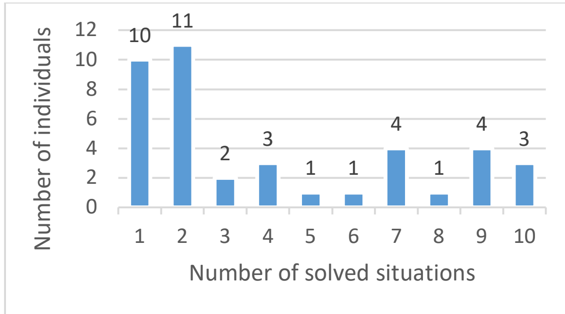

Figure 5: Distribution of successful individuals for each tested initial situation

We view this inconsistency as an indication that the homogeneous system – for the considered combinations of perception- and moving abilities of the entities – features a rather difficult, rugged fitness landscape, and the evolution often struggles to discover the areas of the optimal fitness in it. In particular, as the analysis of the evolved behaviours of the predators, the shorter range of sensors often hinders the formation of the behaviour pattern (surrounding), required to capture the prey and this pattern is seldom discovered by the successfully evolved team of predators.

3.6.2 Improving the performance of the multi-agent system

We introduced the changes into the heterogeneous system with the expectation that they will encourage the agents to evolve a more complex behaviour and more complex (yet implicit) interactions between the predator agents in order to solve the more difficult task.

3.6.3 Evolution of the Heterogeneous System

In our quest to discover such a heterogeneous system that would result in a better efficiency of evolution compared to that of the homogeneous system, we conducted experiments with evolution of three different configurations of predator agents as shown in Table 7. Notice that the average of the range of visibility of sensors of the agents in these three configurations of the heterogeneous system is constant. Moreover, it is equal to the range of visibility of sensors of predators in the considered homogeneous MAS.

Table 7: The three experimental configurations of the ranges of visibility of sensors of the agents in heterogeneous MAS

Experimental Configuration (Test Case)	Range of visibility of sensors, mm		
	Group 1 (two agents)	Group 2 (two agents)	Average of all four agents
A	400	500	450 for all three test cases
B	350	550	
C	300	600	

The results are very diverse – some configurations show better evolution than the average (homogeneous) system, while others cannot compare at all. Some of the configurations manage to evolve better solutions than the average, with fitness values converging around 425 (compared to 447 for the average systems) and the solved initial situations converging around 5 (compared to 4 in average MAS). Other configurations of the heterogeneous system show poor results in terms of successfully solved situations and fitness, as illustrated in Figure 6 and Figure 7. These results demonstrate that the improvement of the overall performance of the heterogeneous MAS, in regard to the average value of a homogeneous MAS, vary depending on the difference between the average values of their perception abilities and the desired optimal values of the implementation of the agents (e.g. in general, based on financial, available resources or some other constraints). In the considered context, the most prominent results are exhibited by multi-agent systems with sensor variations between 10 and 20 percent of the range of visibility of the predators in the average (homogeneous) system. More significant disparities in perceptions of the heterogeneous predators seem to be detrimental both for the efficiency of evolution and effectiveness of the evolved behaviour of agents in the considered instance of MAS.

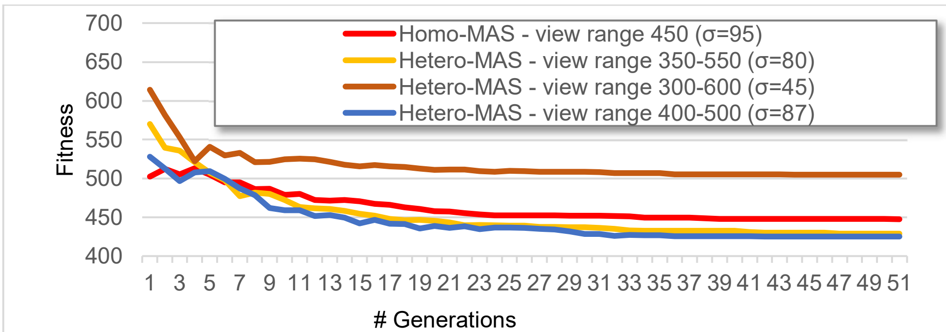

Figure 6: Convergence of the average fitness for different configurations of evolved MAS. On generation #50 the P-value is $1,91x10^{-11}$<<0.05

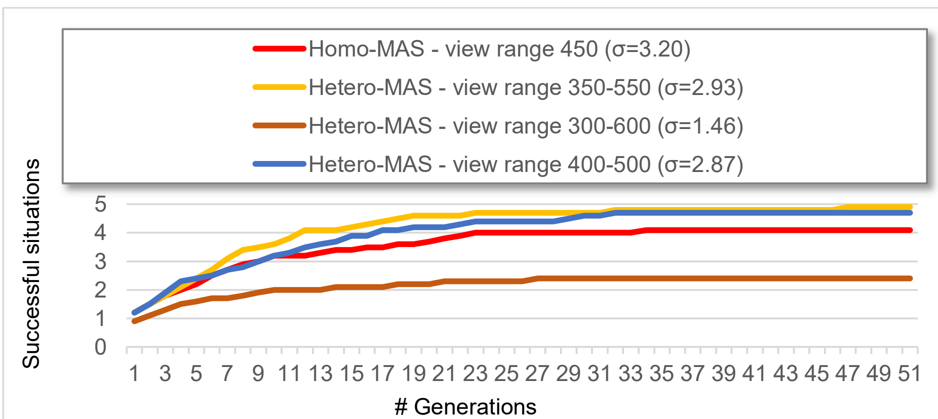

Figure 7: Dynamics of number of successfully solved initial situations. On generation #50 the P-value is $1,56x10^{-83}$<<0.05

From all 4 test cases, the one of the heterogeneous MAS where the visibility range of predators is 350 and 550 (configuration B, shown in Table 7), respectively, demonstrates both (i) most consistent evolution and (ii) highest number of solved initial situations. Figure 8 and Figure 9 show comparison between the best evolved behaviour of predators of this heterogeneous configuration with the best behaviour of the homogeneous system featuring configuration B. The average values, for this test case converge at 429 fitness and 5 successfully solved initial situations.

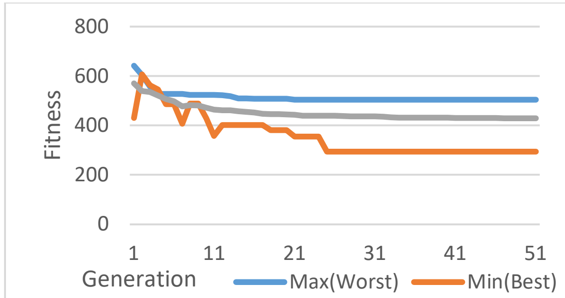

Figure 8: Variation of average, best and worst fitness in a heterogeneous system of predator agents

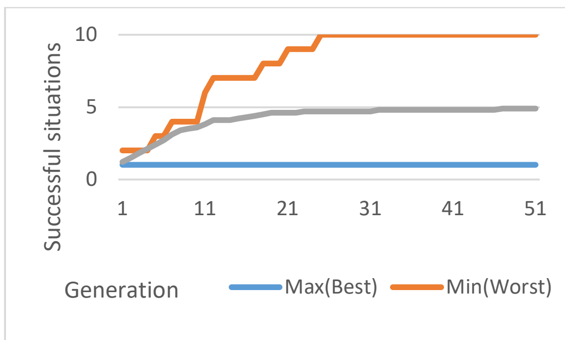

Figure 9: Variation of average, best and worst number of successful situations in a heterogeneous system of predator agents

While in the system of homogeneous agents, the average fitness value was almost overlapping with the worst agent in the series - Figure 3, this time we have a clear separation of the both. This would lead us to think, that more of the chromosomes in every generation are able to find solutions to more of the initial situations. Let us see the distribution of solutions for the system of heterogeneous agents – *Figure 10*.

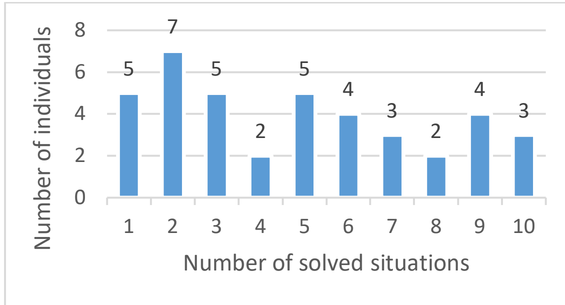

Figure 10: Distribution of successful individuals for each tested initial situation

We can clearly see that the number of individuals that solve only one or two situations dropped in half compared to the homogeneous MAS - Figure 5. Instead, the distribution shifted towards the centre as the system of heterogeneous agents now solves more initial situations in the range of 3 to 8. The number of individuals that solve 9 or 10 situations, however, remains unchanged in comparison to our previous tests. This would lead us to believe that the heterogeneous system has, at least, as good performance as the homogeneous system. However, this comes at the cost of increased evolution time, as the best individual from the heterogeneous system is able to reach 10 successful initial situations at generation 26, compared to generation 9 in the homogeneous system. Another notable thing is a visible deterioration in regards of fitness, of the best individual. The heterogeneous system managed to reach a fitness of 294 compared to such of 232 for the homogeneous system – 21% worse.

We can see further improvements, if we compare the different homogeneous systems variations – namely configuration A and B. The heterogeneous system with predator agents featuring range of visibility 350 and 550 (configuration B) manages to solve all 10 initial situations at generation 24 with a fitness value of 264, while the system with a range of visibility of predators 400 and 500, respectively (configuration A) solves 10 initial situations at generation 15 with a fitness of 223. On the other hand, the homogeneous system with an average range of visibility of predators (i.e., 450) solves all 10 initial situations at generation 20 and fitness of 232. The increased computational effort of evolution is somehow expected, given the significant inflation of the search space of the heterogeneous system. This inflation is caused by the fact that the agents are now separated in two groups of morphologically similar entities and their position in the world has an effect on the outcome of the attempt to solve the task, while in the homogeneous system all agents are the same and their position does not matter. Moreover, despite the increased search space, the results remain comparable - as shown in Table 8, the heterogeneous MAS is more successful in solving more than one initial situation even though the end results are the same – three out of four systems manage to evolve around 3 individuals that solve all 10 initial situations.

Table 8: Average number of successfully solved initial situations.

Successful Initial Situations	Configuration of MAS			Standard deviation σ
	Homogeneous	Heterogeneous		
		Range of visibility 350 and 550 (Test Case B)	Range of visibility 400 and 500 (Test Case A)	
1	10	5	5	2.88
2	11	7	7	2.30
3	2	5	7	2.51
4	3	2	1	1
5	1	5	5	2.30
6	1	4	2	1.52
7	4	3	5	1
8	1	2	2	0.57
9	4	4	4	0
10	3	3	2	0.57

In addition, we would like to note that the best behaviour of predators is evolved in the heterogeneous MAS with range of visibility of predators 400 and 500, respectively, which corresponds to about 10% disparity compared to the average value of 450, used by agents in homogeneous system. Figure 11 illustrates the dynamics of the fitness value and the number of successfully solved initial situations. It is interesting that the evolution actually manages to solve all 10 situations at generation 7. However, because the fitness value at that point does not meet the termination criterion of 300, the evolution proceeds further until, at generation 15 both the fitness value (223) and the number of successful situations (10) satisfy these criteria.

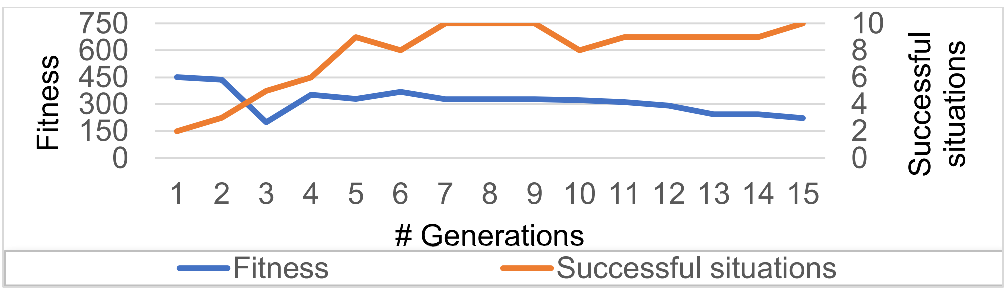

Figure 11: Evolution of the best behaviour in heterogeneous MAS with range of visibility of predators 400 and 500, respectively

3.6.4 Heterogeneous MAS featuring an unequal size of groups of predators

From what we have observed so far, the heterogeneous MAS shows promising results in surpassing the capabilities of the average homogeneous system, though the increase in performance is not significant. Considering the small number of agents – just four – the advantage of using agents with greater sensor abilities may not be sufficient to compensate the drawbacks of having more myopic agents in the same team of predators. Indeed, the two superior agents would not be sufficient to capture the prey in PPPP, as at least three predators (i.e., a “critical mass”) would be needed to surround the prey from all sides of the world. In an attempt to investigate whether the issue of critical mass of predators is relevant to the considered case of PPPP, we introduce divide the predators into two groups with unequal number of members as follows: one group of three agents with increased (above the average) sensory capabilities and one group of one inferior agent with lower than average range of visibility, and vice versa. Table 9 and Table 10 show the two variants of such grouping. We will conduct additional experiments with these two variants of configurations of MAS, and will refer to them as unbalanced configuration #1 and #2 from now on.

Table 9: First variant of the configuration of heterogeneous agents featuring an unequal size of groups

Group #	Number of agents	Range of visibility, mm	Average range, mm
1	3	400	450
2	1	600	

Table 10: Second variant of the configuration of heterogeneous agents featuring an unequal size of groups

Group #	Number of agents	Range of visibility, mm	Average range, mm
1	3	500	450
2	1	300	

The trend of solving more of the cases with 3 to 9 initial situations remain, even for the MAS, where 3 of the agents have lower than average sensory abilities, as 20% of the runs manage to solve 9 initial situations, compared to 17,5% in the average MAS. Table 11 shows the success rate of the newly tested configurations of MAS.

We would like to note that in the unbalanced configuration #2, the number of successfully evolved individuals that solve 10 initial situations, significantly increases, however, at the cost of worse fitness values. Most of the evolved solutions for the improved systems were able to solve 10 initial situations and complete the evolution with fitness greater (worse) than 300 – one of the termination criteria, after which they regress to being able to solve less of the initial situations. In addition, for unbalanced configuration #2, there was an individual that stand out from the other solutions, which completed its evolution with great results by having fitness of 159 and 10 successfully solved initial situations, with small regression in the first few generations. Figure 12 shows the evolution of that individual.

Table 11: Success rate (in %) for each initial situation count for the unbalanced configuration compared to average in percentage of total runs.

Successful Initial Situations	Configuration of MAS		
		Heterogeneous	
	Homo-Heterogeneous	Unbalanced Configuration #1	Unbalanced Configuration #2
1	100	100	100
2	75	90	85
3	47,5	67,5	60
4	42,5	62,5	42,5
5	35	50	35
6	32,5	45	27,5
7	30	32,5	25
8	20	25	25
9	17,5	12,5	20
10	7,5	5	12,5

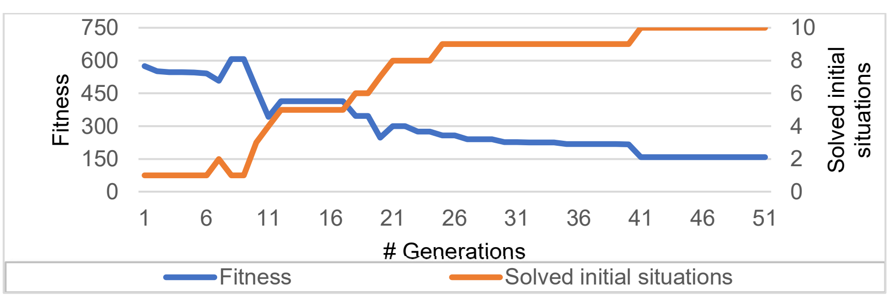

Figure 12: Evolution of the best individual in unbalanced configuration #2.

Even though unbalanced configuration #1 results in lower probability of success, manifested by the lower number of evolutionary runs that solve all 10 initial situations, it produces a notable individual on its own. One of the runs managed to satisfy the termination criteria, with fitness of 260 in only 16 generations, compared to 20 for the best homogeneous (average) system. The evolution of that run is shown in Figure 13.

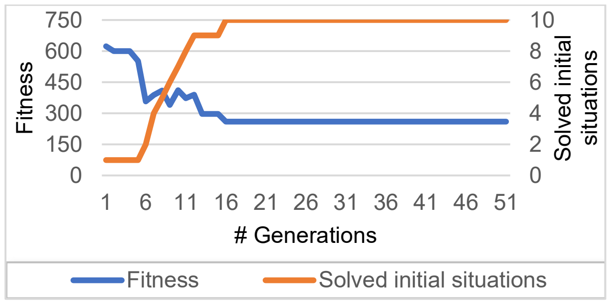

Figure 13: Evolution of the sample individual in unbalanced configuration #1.

3.6.5 Generality of the evolved behaviour of the predators

We investigated the generality of the best evolved behaviours of predator agents in different configurations of MAS. The generality, in our experiments is estimated by the number of successfully resolved situations out of 1,000 initial situations, containing the 10 situations employed for the evolution of predators, plus 990 newly introduced situations. The experimental results are shown in Table 12.

Table 12: Generality of evolved behaviour of predators to newly introduced 990 situations

Configuration of MAS	Range of Visibility of Predators	# Successful Situations (out of 1000)	Success rate
Homogeneous	4 x 450 mm	736	100% (base)
Heterogeneous A	2x400 mm and 2x500 mm	660	90%
Heterogeneous B	2x350 mm and 2x550 mm	533	72%
Unbalanced #1	3x400 mm and 1x600 mm	549	75%
Unbalanced #2	3x500 mm and 1x300 mm	461	63%

As a base for comparison, we used the number of situations, successfully solved by the homogeneous system (736 situations). Because for the same initial situation, the number of combinations of four heterogeneous agents divided into two groups of two identical agents is (4!) ÷ (2!) × (2!) = 6, theoretically the total number of possible initial situations in heterogeneous systems A and B is 6 times (4 times for unbalanced situations #1 and #2) higher than that of homogeneous MAS. Consequently, in order to provide the heterogeneous agents with an equal opportunity to learn (how to capture the prey) as the agents in homogeneous systems, we should have evolved them on 6 times higher number of initial situations, i.e., 60 initial situations.

Because we evolved all the systems under the same setup of the evolutionary framework, we, to some extent, expected the inferior generality of the heterogeneous systems. Nevertheless, the heterogeneous systems A featuring a lower disparity of range of visibility of predators (400 mm and 500 mm, respectively) solves 90% of initial situations that are solved by homogeneous system. Also, it is interesting to note that the unbalanced system #1, with the critical mass of 3 myopic, below average (range of visibility 400 mm) heterogeneous agents is more general than that featuring three longsighted, above the average predators (range of visibility 500 mm).

To show how the increased search space affects the team of heterogeneous agents, we will evolve a new set of controllers for the team of agents in configuration B. To compensate for the increased complexity, we will be evolving a new set of chromosomes, this time on 60 initial situations in a period of 200 generations. How the new configuration of the genetic programming framework relates to the previous one, is displayed in *Table 13*. We will compare the newly evolved solution to the current, using 1000 initial situations.

Table 13: Changes to the GP framework to address the new evolutionary requirements

Parameter	Old Value	New Value
Fitness cases	10 initial situations	60 initial situations
Termination criteria	(Fitness value<300 AND 10 successful situations) or (# Generations>50) or (Stagnation of fitness for 10 consecutive generations)	(Fitness value<300 AND 10 successful situations) or (# Generations>200) or (Stagnation of fitness for 10 consecutive generations)

Once again, to be consistent, we generated 40 runs with the new configuration. The increase in difficulty proved great, as all of the produced solutions stagnated without finding a solution to all 60 initial situations. The best of them, which we will use in the comparison, needed 157 generations to reach its peak performance, while the worst could solve only one situation in 35 generations. Figure 14 shows the development of the best controller from all 40 runs of the GP framework.

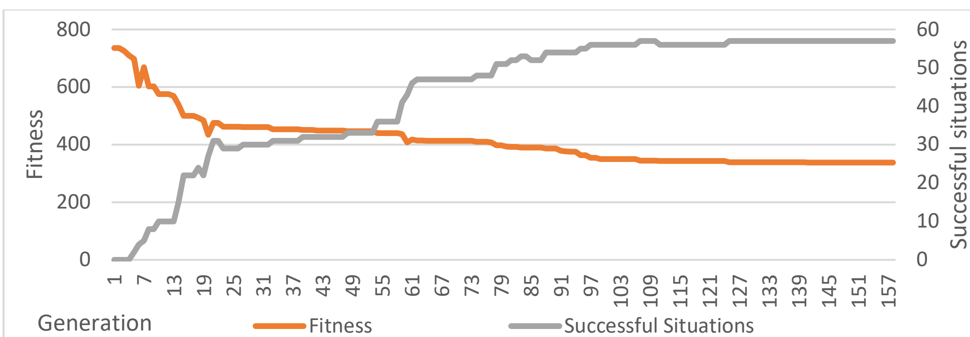

Figure 14: Evolution of the best run of the heterogeneous system evolved on 60 initial situations

The best chromosome of the newly evolved solution for the heterogeneous system managed to solve 57 out of 60 initial situations during the evolution and achieve a fitness of 338. In the test for robustness, we compared that individual to the best individual of the homogeneous system, using 1000 initial situations. As seen in *Table*

14, the results show that the newly evolved agent managed to solve only 692 out of 1000 situations, compared to the 736 out of 1000 for the heterogeneous system. While these results are far better than the results that the initial solution shows – 436 out of 1000, and the deterioration of performance drops from 30% down to less than 10%, the heterogeneous systems still remains inferior to its heterogeneous counterpart.

Table 14: Comparison of robustness test on different types of MAS

System type	# of successful situations in 1000 initial situations
Homogeneous	736
Heterogeneous (configuration B)	533
Heterogeneous trained in 60 initial situations (configuration B)	692

3.6.6 Robustness to noise

While introducing noise to the environment or hardware errors to the agents is the most obvious way to test for robustness of the evolved controller solutions, we will introduce a simpler way, next.

To investigate the robustness of the team of predators, evolved in noiseless environment, would degrade when subjected to perception noise, we introduced a uniform perception noise of up to 5% to both the distance (perceptions Prey_d and Peer_d, shown in Table 3) and bearing (Prey_a and Peer_a) of the perceived entities in MAS. Figure 15 illustrates the variations of the number of solved initial 1,000 (including 10 used for evolution, and 990 newly added) situations for different levels of perception noise. For most of the considered configurations of MAS, the noise results in anomalous increase of the number of successfully solved situations. Because the actions of predators (e.g., “Turn 22 degrees to the left”, “Go with 50% of max speed”, “Turn 10 degrees to the right”, etc.) are a result of execution of alternating stimuli-response rules (corresponding to the instantly perceived, dynamic environment), and therefore, the behaviour of agents – seen as a sequence of actions – is rather discrete (jerky) [8], a possible explanation of this anomaly is in the favourable effect of the noise-induced *dithering* (smoothing) on such a behaviour. We are planning a more in-depth investigation of why and how dithering facilitates a better behaviour of predators in unforeseen situations. Moreover, we intend to investigate the conditions (if any), at which the generality of the multi-agent systems could be improved by adding a certain amount of perception noise.

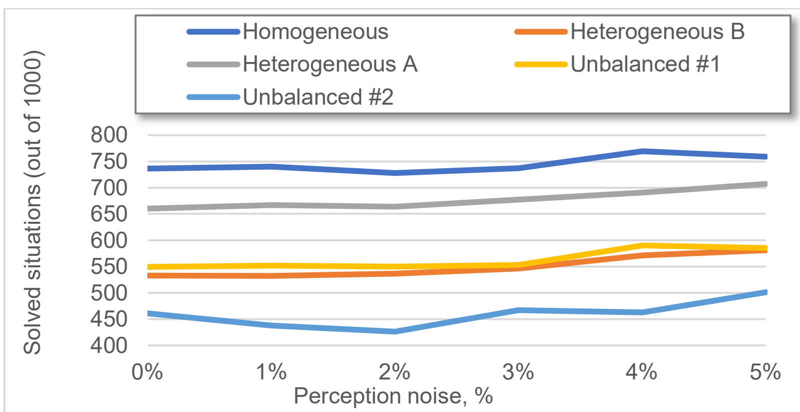

Figure 15: Robustness to perception noise in 1000 initial situations. Legend of lines from top to bottom: Homogeneous, Heterogeneous A, Unbalanced #1m Heterogeneous B, Unbalanced #2

Due to the trend that number of successfully solved situations increases with noise, we decided to make one additional test of the two best chromosomes – homogeneous and heterogeneous A. We have tested with 25% noise. The results show that the homogeneous system suffered a regression to only 596 solved situations, while the heterogeneous system managed to solve 843 out of 1000. It shows 14% increase compared to the base of 736.

3.7 Summary

In this chapter we analysed the performance of homo- and heterogeneous multi-agent systems modelling the predator-prey pursuit problem. All considered systems featured identical average values of the respective perception abilities of predator agents. The experimental results indicate that both (i) the speed of evolution of the successful capturing behaviour of predator agents and (ii) the effectiveness (i.e., its fitness value) of the best-evolved behaviour, of the heterogeneous MAS are improved, using different methods and techniques, in such a way that it performs better than its homogeneous counterpart. We have demonstrated that the heterogeneous system featuring a deviation of the perception abilities of predator agents, of about 10% from the average, could be evolved faster and could result in a better performing team of agents. We also showed that by implementing a team of agents that is big enough to potentially solve the problem alone (i.e., critical mass), the evolution of even better performing team of agents could be achieved even faster. The homogeneous system, however, is more general, in that it is able to successfully resolve the higher number of unforeseen initial situation. The robustness to introduced noise, however, depends on the level of the noise. With high levels of noise, the heterogeneous system shows results that are more consistent and more efficient. One of the reasons for the inferior robustness of heterogeneous systems to uncertainty is that the space of possible combinations of initial situations is significantly larger than that of the homogeneous system. Evolving both types of systems on the same number of initial situations might result in under-representation of the training cases for the heterogeneous system. However, increasing the number of initial situations used for the evolution of the latter would inevitably result in increase of the computational overhead of simulated evolution. In our future work we are planning to investigate the trade-off (if any) between the computational overhead of evolution and the robustness to uncertainty of heterogeneous systems. An eventual success in this direction would allow us to verify our hypothesis that – similarly to the human societies – the disparities in individual capabilities of agents are more important for the success of the team of agents than maintaining identical, "average" agents.

3.8 Discussion

We've shown that, by making a part of the agents simpler, in terms of morphology, the system as a whole can gain some improvement in certain areas, such as effectiveness in capturing the prey and evolution speed. This led us to wonder, what would happen if we simplify all the agents in the system. We revisited the predator-prey pursuit problem to use very simple predator agents.

Chapter 4: Modelling the predator-prey pursuit problem featuring simple predator agents

4.1 Proposed approach

The direction for this optimisation was inspired by simple reactive robots that were previously modelled as agents by Gauci et al [9]. The agents were able to self-organize in order to solve the simple robot aggregation problem. The same framework was also successfully applied for the more-complex object-clustering problem [10] in which the agents need to interact with an additionally introduced immobile object. The very possibility of a team of such agents to conduct an elaborate social (surrounding) behaviour in an environment featuring dynamic objects was recently demonstrated by Ozdemir et al [11] in solving the shepherding problem, where a team of simple agents (shepherds) need to guide multiple dynamic agents (sheep) toward an a priori defined goal.

Bhattacharya et al. [12] presented one of the first works on sensory constrains for robots featuring two wheels in a differential drive configuration (the simplest possible effectors), that are required to solve complex tasks such as navigation. The notion of sensory constraints was later developed into the concept of the minimum amount of sensory information that should be adequate for robots with two wheels as effectors to accomplishing a task of a given complexity. Yu et al. [13] proposed the simple "windshield" (field of view) sensors. The proposed sensor was further minimized to a single line-of-sight sensor that could be viewed as a special case of the "windshield" featuring a nearly zero angle of the visual field [14].

In our study, we are proposing the use of similar team of simple agents in the solution of a different task – the well-studied, yet difficult to solve predator-prey pursuit problem

(PPPP) [3] [4] [15] [16]. In the considered PPPP, eight identical, simple agents (predators) are required to capture the single dynamic agent (prey).

4.1.1 Changes compared to the complex MAS

To comply with this definition of simple robots, in our research we consider predator agents featuring a single beam (line-of-sight) sensor providing just two bits of information, and two wheels (arranged in a differential drive configuration), rotational velocities of which are controlled by two motors. Their purely reactive behaviour is realized by a simple decision-making that does not require any computing. Instead, it involves a direct mapping of just four perceived environmental states into corresponding pairs of rotational velocities of wheels' motors.

4.1.2 Motivation for the introduced changes

While previously are interested to reduce the costs of manufacturing and operation, of the MAS, we are looking towards creating such a team of agents, that can be used for novel tasks, where other – more complex systems are unusable. We are especially interested in the emerging small-scale robots – micro and nano-robots – that are promising candidates in future manufacturing and biomedicine [17] [18] [19] [20]. Other researchers have already found ways to create robots on a nano-scale, which are guided by an external force [21] [22]. Our work focuses on creating autonomous units which can operate, even on a miniature scale, without the need of outer force or a monitoring. However, several challenges are currently hindering the progress of the real-world applicability of these robots. Because of the physical constrains due to their small size, these robots could not be morphologically advanced – both the sensors and the actuators would have to be rather simple in order to fit the body of the agent. Further, their behaviour would be simple as well. It would not involve any computing; instead, it would feature a direct mapping of the (few) perceived environmental states into actuators commands, instead of featuring a complex decision-making mechanism in each one of them. Most likely the communication (if any), between the individual agents, would be impossible to be realized in a direct manner and would be fulfilled implicitly, using the environment. As an example of such robots, we consider robots equipped with a single line-of-sight sensor providing only two bits of information, and two thrusters (wheels, in two dimensions, or propellers, in three-dimensional environments) in a differential drive configuration, controlled by two motors. Such robots can be regarded as an ultimate case of Occam's razor principle, applied both the morphology and decision-making of mobile robots. The simplicity of such robots would imply a reduced size of the search space, and therefore more efficient heuristics [23] [24].

As a model of such agents, we've used PPPP [3] [8] [15] [16], to test and show, if any, the advantages that the agents have over the traditional complex MAS. The PPPP is widely used as a benchmark for the effectiveness of emergent complex, coordinated behaviour of agents in MAS. It could serve as a model of various potential real-world applications of both macro- [25] [3] [26] and micro-robots [26] [27] [28] [29] [30].

The proposed MAS features eight identical, simple agents (predators) - used to capture a single dynamic agent (prey). At first, our *objective* is to investigate whether the PPPP is solvable by the team of such simple predator agents. Further, we investigated the feasibility of applying genetic algorithms (GA) to evolve direct mapping of the four perceived environmental states into the respective velocities of the wheels of the predators that yield the social behaviour of the predators, resulting in the successful capturing of the prey.

An additional motivation of our research is the recognition that while many real-world scenarios could be, indeed, reduced to the previously researched wall-following, dispersal [31], clustering [9], and shepherding problems [11], there would be few scenarios – requiring a direct physical contact with an active prey – that could be modelled by the proposed instance of PPPP [28] [29] [30]. These scenarios might include pinpoint drug delivery, surrounding and destroying (cancer) cells or bacteria, gathering around cells to facilitate their repair or imaging, etc.

4.2 Agent characteristics

4.2.1 Prey agent

The prey is equipped with an omnidirectional sensor, with limited visibility range. To balance the advantage that the omnidirectional sensor gives to the prey, compared to the single line-of-sight sensor of the predators. The viewing distance of the prey is only 50 units, compared to the 200 units of the predators (Table 15). The maximum speed of the prey, however, is identical to that of the predators. We introduced such sensory and moving contrast to encourage the agents, to evolve as cooperative behaviour as they will be unable to capture the prey alone. Another viewpoint suggests that a successful solution to PPPP, defined in such a way, could demonstrate the virtue of the MAS as it could solve a problem that a single (predator) agent could not.

In contrast to the predator behaviours, we implemented a handcrafted behaviour for the prey. The prey attempts to escape from the closest predator (if any) by running at its maximum speed in the direction that is exactly opposite to the bearing of the predator. The prey it remains still if it does not detect any predator [8]. Table 15 shows the main features of the prey agent.

4.2.2 Predator agents

Each of the eight identical *predators* models a simple cylindrical robot with a sensor featuring a limited range, and two wheels, controlled by two motors in a differential drive configuration. The features of the agents are shown in Table 15.

Table 15: Features of the predator and prey agents

Feature	Value of the Feature	
	Predators	Prey
Number of agents	8	1
Diameter (and wheel axle track), units	16	16
Max linear velocity of wheels, units /s	10	10
Max speed of agents, units /s	10	10
Type of sensor	Single line-of-sight	Omni-directional
Range of visibility of the sensor, units	200	50
Orientation of sensor	Parallel to longitudinal axis	N/A

The predators' sensor provides two bits of information: each bit encodes if an entity (predator or prey) – is detected in the line of sight. Such sensors allow the predators to perceive only four discrete environmental states, as shown in Figure 16. The perceived environmental states do not provide the predators with any insight about the distance to the perceived entities, nor their total number.

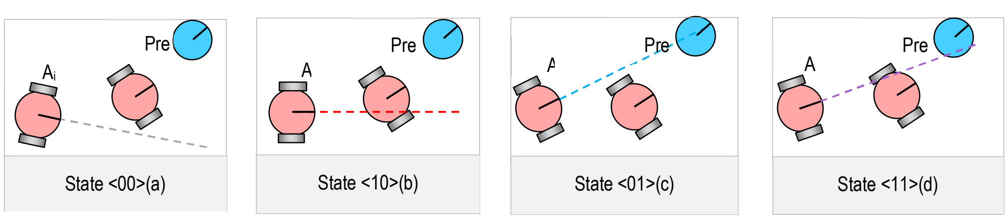

Figure 16: The four possible environmental states perceived by (any) predator agent Ai.

The reactive behaviour of the predator agents could be described as a direct mapping of each of the four perceived environmental states into a corresponding rotational speed of the wheel motors. For simplicity, hereafter, we will assume a mapping into the linear velocities of the wheels, expressed as a percentage – within the range [-100 % ; +100%] – of their respective maximum linear velocities. The decision- making of the predator agents could be formally expressed by the following octet *D*:

$$D = \{V_{00L}, V_{00R}, V_{01L}, V_{01R}, V_{10L}, V_{10R}, V_{11L}, V_{11R}\} \quad (1)$$

where V_{00L}, V_{00R}, V_{01L}, V_{01R}, V_{10L}, V_{10R}, V_{11L}, and V_{11R} are the linear velocities (as a percentage of the maximum linear velocity) of the left and right wheels of the predators for the perceived environmental states <00>, <01>, <10>, and <11>, respectively.

Our *objective* of evolving (via GA) the optimal direct mapping of the four perceived environmental states into their respective velocities of wheels could be rephrased as evolving such values of the velocities, shown in the octet in Equation (1), resulting in an efficient capturing behaviour of the team of predator agents.

4.3 Evolutionary subsystem

MAS, as complex systems, feature a significant semantic gap between the hierarchically lower-level properties of the agents, and the (emergent) higher-level properties of the system as a whole. Thus, we could not analytically infer the optimal velocity values of the wheels of the agents from the desired behaviour of the team of predator agents. Therefore, we applied the GA – a nature-inspired heuristic approach to gradually evolve good values of the parameters, similar to the evolution of species in nature. GA have proven to be efficient in finding the optimal solution(s) to combinatorial optimization problems featuring large search spaces [32] [33]. Thus, consonant with the concept of evolutionary robotics [34], we adopted the GA [5] to evolve good values of the eight velocities of the wheels of the predators that resulted in an efficient behaviour – presumably involving exploring the environment, surrounding, and capturing the prey – of the team of predators. The algorithmic steps of the GA are shown in Figure 17, and its main attributes are elaborated below.

Step 1: Creating the initial population of random chromosomes;
Step 2: Evaluating the population;
Step 3: WHILE not (Termination Criteria) DO Steps 4~7:
Step 4: Selecting the mating pool of the next generation;
Step 5: Crossing over random pairs of chromosomes of the mating pool;
Step 6: Mutating the newly created offspring;
Step 7: Evaluating the population;

Figure 17: Main steps of GA

4.3.1 Genetic Representation

The decision-making of the predator agents is encoded genetically as a "chromosome". The latter consist of an array of eight integer values ("alleles") of the evolved wheel velocities of the agents, as shown in Equation (1). Sample configuration of a predator chromosome is shown in Figure 18. These values are within the range [-100 % … +100 %], and are discretized into 40 values, with an equal interval of 5 % between them. This number of discrete values provides an acceptable trade-off between the resolution of the evolved velocities and the size of the search space (40^8) of the GA. The size of the population is 400 and the breeding strategy is homogeneous – each chromosome is evaluated after being cloned to all eight predator agents.

```
<?xml version="1.0"?>
<GP xmlns:xs="http://www.w3.org/2001/XMLSchema-instance"
xs:noNamespaceSchemaLocation="..\SharedResources\GPSchema.xsd">
  <STM ind="3">
    <M00 ind="4">
      <M00L>-4</M00L>
      <M00R>17</M00R>
    </M00>
    <M01 ind="9">
      <M01L>16</M01L>
      <M01R>15</M01R>
    </M01>
    <M10 ind="14">
      <M10L>-13</M10L>
      <M10R>-12</M10R>
    </M10>
    <M11 ind="19">
      <M11L>14</M11L>
      <M11R>9</M11R>
    </M11>
  </STM>
</GP>
```

Figure 18: Sample chromosome of a predator agent in simple MAS

4.3.2 Genetic Operations

We decided to use binary tournament selection strategy in the evolutionary framework. It is computationally efficient, and has been proven to provide a good trade-off between diversity of the population and the fitness convergence rate [33]. We also adopted elitism in that the four best-performing chromosomes survive unconditionally and are inserted into the mating pool of the next generation. Further, we implemented – with equal probability – both one and two-point crossovers. The two-point crossover results in an exchange of the values of both velocities (of the left and right wheels, respectively) associated with a given environmental state (e.g., both V_{01L} and V_{01R}). This reflects our assumption that the velocities of both wheels determine the moving behaviour of the agents (for a given environmental state); therefore, they should be treated as a whole – as an evolutionary building block. The one-point crossover is applied to develop such building blocks (*exploration* of the search space), while the two-point crossover is intended to preserve them (*exploitation*).

4.4 Evaluation subsystem

We've used the previously developed simulation and introduced changes to model the new environment more closely.

4.4.1 Simulating the world

We modelled the world as a two-dimensional infinite plane with a visualized part of 1600 × 1600 units. We update the perceptions, decision-making, and the resulting new state (e.g., location, orientation, and speed) of agents with a sampling interval of 0.1 s. The duration of trials is 120s, modelled in 1200 time-steps. We approximate the new state of predators in the following two steps, as illustrated in Figure 19. First, from the current orientation, the yaw rate, and the duration of the sampling interval we calculate the new yaw (orientation) angle (as an azimuth to the north) of the agents. The yaw rate is obtained from the difference between the linear velocities of the left and right wheels, and the length of the axis between the wheels. Then, we calculate the new position (i.e., the two-dimensional Cartesian coordinates) as a projection (in time, equal to the duration of the sampling interval) of the vector of the linear velocity of predators. The vector is aligned with the newly calculated orientation, and its magnitude is equal to the mean of the linear velocities of the two wheels.

```
// Global definitions:
type
   TEntitiy_in_MAS = record
     Yaw           : float;   // radians
     X    : float;   // units
     Y    : float:   // units
   end;
const
   Num_of_Predators      =  8;
   Pred_Max_Speed        = 10;     // units/s
   Sampling_Interval     =  0.1;   // seconds
   Pred_Radius           =  8;     // units
var
   Predator: array [0.. Num_of_Predators-1] of TEntitiy_in_MAS;
. . .
  // The routine Move_Predator estimates the new state of predators
Procedure Move_Predator (ID: integer; V_L,V_R: float);
 // ID: the ID of the predator being currently updated, within the range [0..7]
 // V_L and V_R: linear velocities of the left and right wheels, respectively.
 // Calculated from the evolved genotype (as percentages of the max velocities), the values of max velocities
 // (10 units/s), and currently perceived (one of the four: <00>,<01>,<10> or <11>) environmental situations.
 // For the evolved sample genotype <10,15,20,25,30,35,40,45> and current situation <01> these values
 // are V_L=2.0 units/s, and V_R=2.5 units/s, respectively
begin
// Step #1: Calculating the new yaw angle of the predators #ID as azimuth (to the north) in radians:
   Predator[ID].Yaw := Predator[ID]. Yaw + (V_L –V_R) ⁄ ( Pred_Radius × 2) × Sampling_Interva

// Step #2: Calculating the new position (X,Y) of the predator #ID:
   Predator[ID].X := Predator[ID]. X + ((V_L + V_R)/2) × sin(Predator[ID].Yaw) × Sampling_Interval;
   Predator[ID].Y := Predator[ID]. Y + ((V_L + V_R)/2) × cos(Predator[ID].Yaw) × Sampling_Interval;
end;
```

Figure 19: The pseudocode of estimating the new state of the moving predators.

4.4.2 Fitness Calculation

To evolve predator behaviours that are general to several initial situations, we evaluated the objective (fitness) function (OF) of each of the evolved chromosomes on 10 different initial situations. In each of these situations, the prey is located in the centre of the world. The predators are scattered in a small cloud situated south of the prey (Figure 20).

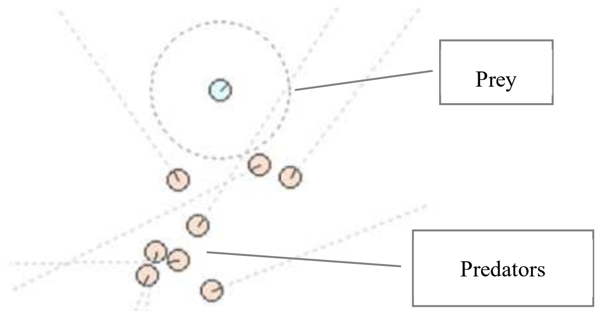

Figure 20: Sample initial situation.

The predators' starting location is very close to the prey during the initial situation and increases in distance with each run. The distance of the cluster, of agents, to the prey is calculated as the follows: ID of the current situation × 2 + (random of 50 units). This helps reduce the impact of the first few evolutionary runs, when the predators are learning how to move around the environment to find the prey. The overall fitness is the sum of the values, scored in each of the 10 initial situations. For a successful situation (the predators manage to capture the prey during the 120s trial), the fitness is the time needed to capture the prey (selection favouring the lowest values). For an unsuccessful situation, the OF is calculated as the sum of (i) the closest distance, registered during the trial, between the prey and any predator, and (ii) a penalty of 10,000. The former component provides evolution with a cue about the comparative quality of the different unsuccessful behaviours. We verified empirically that this heuristic quantifies the "near-misses" well, and correlates with the chances of the predators – pending small evolutionary tweaks to their genome – to successfully capture the prey in the future. The second component penalizes heavily the lack of success of the predators in any given initial situation. The main parameters of the GA are elaborated in Table 16.

Table 16: Parameters of GA

Parameter	Value
Population size	400 chromosomes
Selection	Binary tournament
Selection ratio	10%
Elites	Best 4 chromosomes
Crossover	Both single and two-point
Mutation	Random single-point (with even distribution)
Mutation ratio	5%
Fitness cases	10 initial situations
Duration of the trial	120 s per situation
Fitness calculation	Sum of fitness values of each situation: a) Successful situation: time needed to capture the prey b) Unsuccessful situation: 10,000 + shortest distance between the prey and any predator
Termination criteria	(overall fitness < 600) or (number of generations > 200) or (fitness stagnation for 32 generations)
Genotype	Eight integer values of the velocities of wheels (V_{00L}, V_{00R}, V_{01L}, V_{01R}, V_{10L}, V_{10R}, V_{11L}, and V_{11R})

Our PPPP is an instance of a minimization problem, as a lower overall fitness value corresponds to a better performing team of predator agents. The evolution terminates on overall fitness values lower than 600, which implies a successful capture of the prey in all 10 initial situations in an average time shorter than 60s (half of the trial duration).

4.5 Initial results and challenges

4.5.1 Evolving the team of straightforward predator agents

We implemented 32 independent runs of the GA in an attempt to evolve a suitable mapping of the perceived environmental states into corresponding velocities of wheels of predators with canonical morphology [23]. The sensor of these agents is aligned with their longitude axis. Figure 21 illustrates, the mean value of the OF slowly converges to approximately 60,000, indicating that, on average, only 6 (of 10) initial situations could be successfully resolved (Figure 21). The best result, achieved by the team of predators, is only 6 successful situations. These results suggest that the PPPP is, in general, intractable for the current morphology of the predator agents.

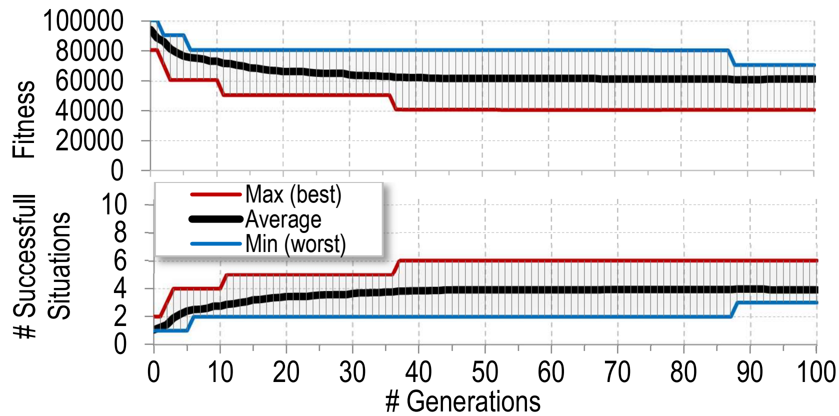

Figure 21: Convergence of the values of best objective function (top left) and the number of successful situations (bottom left) of 32 independent runs of GA. The bold curves correspond to the mean, while the envelope shows the minimum and maximum values in each generation. A snapshot of a sample initial situation is shown on the right.

4.5.2 Enhancing the morphology of predators

To improve the generality of the evolved predator behaviours, we focus on modifying their morphological features. The last of the features listed in Table 15 – the orientation of the sensors – implies a straightforward implementation of the agents. This, indeed, is the common configuration of the previously studied simple agents [9] [10] [11] [14] [31]. We are interested in whether an a priori fixed *asymmetry* – an angular offset – would facilitate the evolution of more general behaviours of the team of predators. We speculate that a sensory offset would allow the predators to realize an anticipatory (equiangular, proportional) pursuit of the prey, aiming at the anticipated point of contact with the moving prey, rather than the currently perceived position of the prey. Notice that the proposed asymmetric morphology does not compromise the intended simplicity of the predator agents.

In our experimental setup, we fixed the offset of all predators to 10°, 20°, 30°, and 40° counter clockwise and conducted 32 evolutionary runs of the GA for each of these 4 configurations. The results are shown in Figure 22a, b, c, and d, respectively, and summarized in Table 17. As Figure 22a and Table 17 illustrate, offsetting the sensors by only 10° significantly improves the generality of the evolved predator behaviours. They can resolve all 10 situations in 30 (93.75%) of the 32 evolutionary runs. The probability of success – the statistical estimation of the efficiency of evolution, defined for the PPPP as the probability to resolve all 10 initial situations, reaches 90% by generation #60 (Table 17). The terminal value of the OF in the worst evolutionary run is 10,987, corresponding to only one unresolved initial situation.

Offset	Terminal value of objective function				Successful Runs		# Generations needed to reach probability of success 90%
	Best	Worst	Mean	Standard deviation	Number	in % of 32 runs	
No offset	40,928	70,729	61,064	8,516	0	0	NA
10°	504	10,987	1,310	2,531	30	93.75	60
20°	**468**	818	588	57.2	**32**	**100**	**9**
30°	495	**713**	**574**	**38.5**	**32**	**100**	12
40°	475	40,903	1,840	7,128	31	96.875	15

Table 17: Efficiency of evolution of the team of predator agents

More efficient evolution and behaviours that are more general were obtained for the sensory offsets of 20° and 30°. As Figure 22b and Table 17 depict for 20°, the predators successfully resolved all 10 initial situations in all 32 evolutionary runs. The probability of success reaches 90% relatively quickly – by generation #9 (Table 17). Both the efficiency of evolution and the generality of the predator behaviours are similar for agents with a sensory offset of 30°, while these two characteristics deteriorate with the further increase of the offset to 40° (Figure 22c, d, and Table 17).

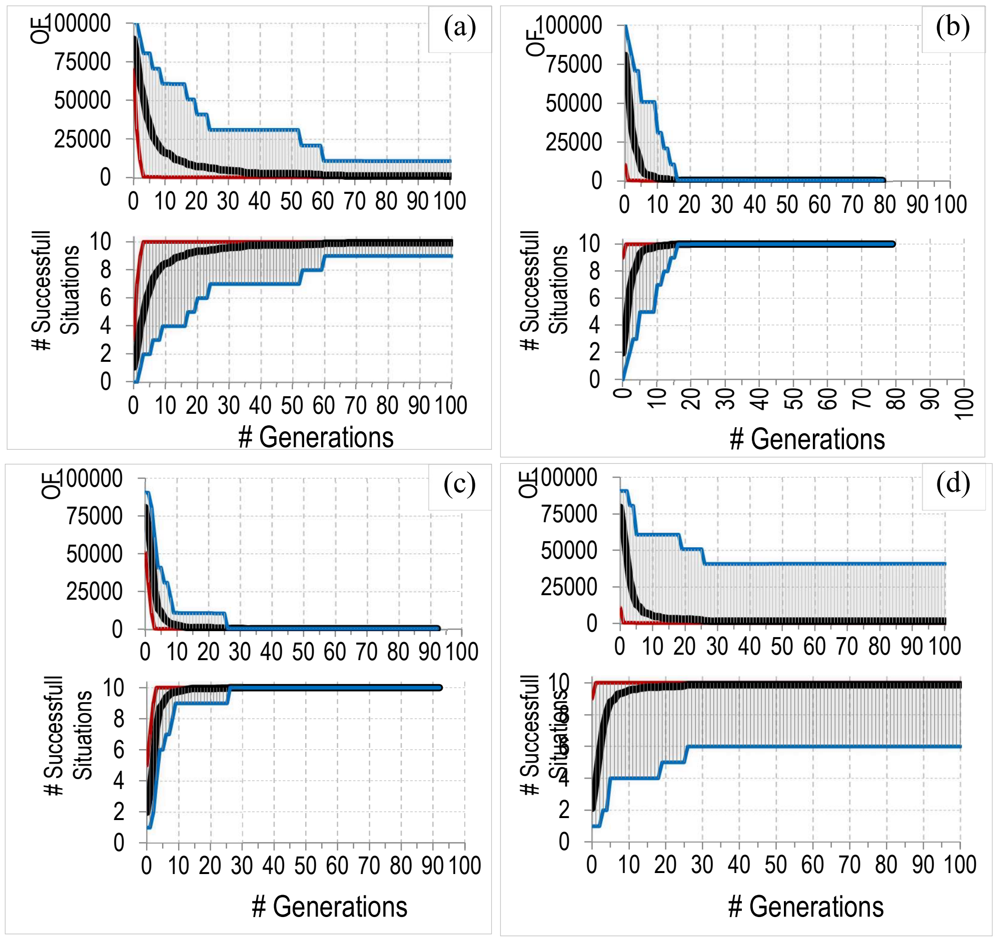

Figure 22: Convergence of the values of best objective function (top) and the number of successful situations (bottom) of 32 runs of GA evolving predators with sensor offset of (a) 10°, (b) 20°, (c) 30°, and (d) 40°, respectively. The bold curves correspond to the mean, while the envelope illustrates the minimum and maximum values in each generation.

4.6 Robustness of the Evolved Behaviour to Noise

We examined the effect of a random perceptual noise on all evolved behaviours of the most general predators – those with sensory offsets of 20° and 30°. We introduced two types of noise – a false positive (FP) and false negative (FN), respectively. The FP results in either of the two bits of perception information to be occasionally (with a given probability) read as “1” regardless of whether an entity is detected by the predators. The FN results in readings of “0” even if an entity is the line of sign. The best results with the increase in the amount of noise from 0 to 16 % (Figure 23a and b) were achieved by a predator with a sensor offset of 20°, as shown in Table 18. The OF value of such predators in a noiseless environment is 552 – close to the average (588) and far from the best evolved (468). Interestingly, the same behaviour, being evolved for the sensor offset of 20°, exhibits an impressive robustness to errors in the angular positioning of the sensor, as well. As shown in Figure 23c, the predators can resolve 9 (of 10) initial situations when the sensor offset of all the agents is set to any value between 10° and 40°.

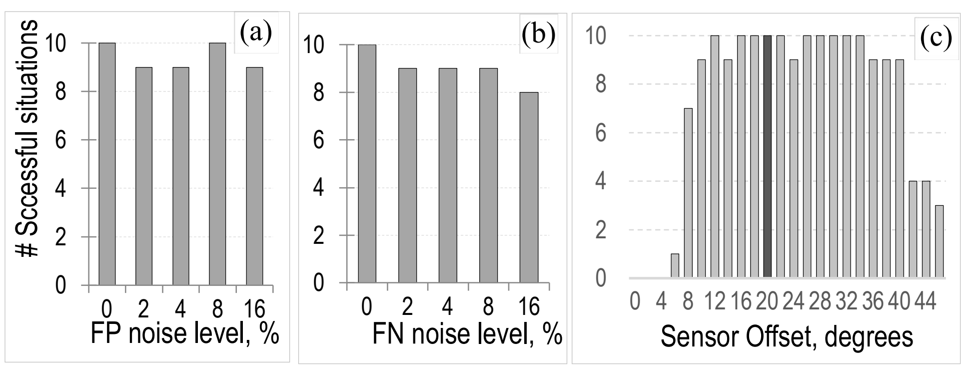

Figure 23: Robustness of a sample best evolved behaviour of predators with sensor offset of 20° to random false positive (FP) noise (a), false negative (FN) noise (b), and to error in angular positioning of the sensor (c).

Table 18: Evolved velocities of wheels of predators that result in a behaviour that is most robust to noise. The sensor offset is 20°.

V_{00L}	V_{00R}	V_{01L}	V_{01R}	V_{10L}	V_{10R}	V_{11L}	V_{11R}
25%	100%	100%	100%	-25%	-20%	100%	100%

The team of predators exhibits three *emergent behaviours*, as illustrated in Figure 24 (a movie is available at http://isd-si.doshisha.ac.jp/itanev/SA/): (i) exploring the environment by dispersing themselves into a wide area in the world (t = 0 and t = 20 s), (ii) shepherding (t = 30s and t = 40 s), and (iii) capturing the prey (t = 50 s and t = 55 s), respectively. The agents commence the trial (t = 0) with no entity in sight. Controlled by $<V_{00L} = 25\%, V_{00R} = 100\%>$ (Table 18) they turn to the left until either a predator (most likely) or a prey is detected. Detecting a predator activates the setup of the wheels $<V_{10L} = -25\%, V_{10R} = -20\%>$, resulting in both turning slowly and moving (dispersing) away from the perceived predator. Such a dispersion widens the area of the cloud of predators and enhances their ability *to explore* the environment and to detect the prey (t = 0 s and t = 20 s). When the predators detect the prey, they activate the setup $<V_{01L} = 100\%, V_{01R} = 100\%>$, resulting in a chase of the prey in the forward direction with maximum speed (Figure 24, t = 20 s and t = 30 s). As a result of the optical parallax, during the chase, the prey might become temporarily invisible, as shown in Figure 16a and b. When this occurs, the predator activates the setup $<V_{00L} = 25\%, V_{00R} = 100\%>$, which yields a counter clockwise rotation towards the invisible prey. The predator exhibits an *embodied cognition* that the parallax is a result, in part, of its own forward motion; therefore, the new location of the prey is – due to the counter clockwise offset of the sensor, – most likely on the left of its own orientation. Therefore, the *virtue* of the sensor offset is in the more deterministic direction of the prey disappearance, which facilitates a faster rediscovery of the latter by the predator (Figure 16b and c). The predator could quickly rediscover the prey by turning slightly (and quickly) by only a few degrees to the left (α). Conversely, in an eventual straightforward implementation of the sensor, the predator would need to turn α degrees if the direction of turning by chance coincides with the new location of the disappeared prey, or (360-α) degrees otherwise. Because, from the predator viewpoint, the moving of the disappearing prey is non-deterministic, on average, the predator would have to turn 180° – i.e., much wider (and slower) than turning only a few degrees (α), as with an offset sensor. Moreover, such a chase, due to the sensor offset, yields a counter clockwise, circular trajectory of both the chasing predator(s) and the prey (Figure 16c), thereby resulting in *shepherding* the prey back into the (already widely dispersed) other predators. Surrounded from all sides of the world by both current and newly encountered chasing predators, the prey is finally being captured (Figure 24, t = 40 s, t = 50 s and t = 55 s).

In hindsight, we could also argue that the initial dispersion illustrates the emergent strategy of the predators, i.e., for a capture, only three of them (the "critical mass") would be sufficient. By moving away from each other, most of the predators move further away from the prey as well (Figure 24, t = 0 and t = 20 s), thereby compromising their chances to capture the prey. However, such an altruistic behaviour results in a faster discovery – and a faster capture of the prey by some (e.g., just three) predators, presumably, for the benefit of the whole team.

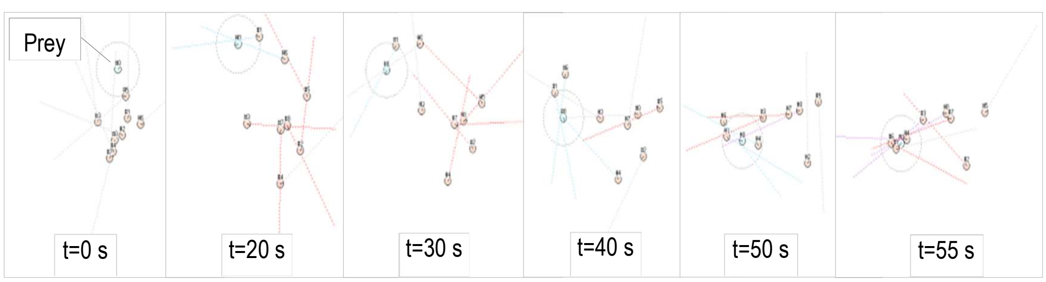

Figure 24: Phases of a sample best evolved behaviour of the predators with sensor offset of 20°.

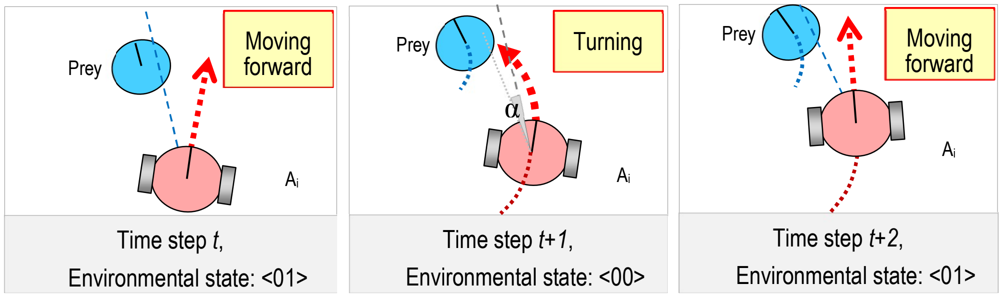

Figure 25: Reliable tracking of the prey by chasing predator A_i.

4.7 Summary

We examined a PPPP featuring very simple, non-computing predator agents, equipped with a single line-of-sight sensor and a simple control of velocities of their two wheels. We applied a GA to evolve the successful behaviour of the team of predator agents. To enhance the generality of the evolved behaviour, we proposed an asymmetric morphology of the predators. Offsetting their sensors angularly to 20° and 30° yielded the most efficient and consistent evolution of successful behaviours of agents.

4.8 Discussion

4.8.1 Heterogeneous vs Homogeneous systems

A different approach to finding a solution would be to implement a multi-agent system with several different types of predator agents, in which, each of them has a specifically assigned role that contributes towards capturing the prey. In our previous work that compares the performance of heterogeneous and homogeneous MAS [35] [36], we have analysed in-depth the different problems that heterogeneity introduces. The main reason we did not implement a heterogeneous MAS is that the efficiency of evolution of the heterogeneous system might be hindered by the inflated search space. Additionally, due to the a priori defined specialization of the agents in heterogeneous systems, we cannot ensure that generality of the heterogeneous team; e.g., consisting of two types of agents—drivers and capturers—would not be compromised, because we could not guarantee that the a priori defined drivers would be in the most favourable position relative to the prey in each of the initial situations. Indeed, in real-world situations, placing a particular predator (driver) in a particular position is challenging and not always possible. Alternatively, we decided to implement a priori unspecialized, versatile agents and to give them the ability to execute any role (depending on their perceived environment) that is needed to capture the prey. In our case, the agent that is closest to the prey would assume the role of a "driver". This behavioural heterogeneity emerges dynamically from the interaction between the homogeneous genotype (all of the agents share the same velocity mappings of the rotation velocities of wheels) and the environment.

4.8.2 Finding the optimal configuration

In this section, we've shown how a MAS featuring simple predator agents compares to its counterpart system featuring more complex in terms of morphology agents. Initially the simple MAS was able to solve the problem only partially – being able to catch the prey only 6 out of 10 times. To overcome this problem, an offset to the viewing sensor was introduced. We've determined that this simple change provides the MAS with the means to successfully solve the PPPP. However, we cannot say if the values for the sensory offset (20° and 30°) that we've found to work the best in our tests are optimal for the system. To answer this question, we've decided to coevolve the sensory offset as well as the motor mappings of the predators, in attempt to find an optimal value.

Chapter 5: Coevolution of the morphology and behaviour of simple predator agents

In this section we will present the experimental results of evolving the optimal values of the velocities of the motors and the angular offset of the sensor that yield an optimal successful behaviour of the predator agents. We will evaluate the proposed approach in terms of efficiency and consistence of evolution, generality of evolved behaviour, and robustness to noise.

5.1 Coevolution objective

Our objective of coevolving (via GA) the behaviour and asymmetric sensory morphology of the agents could be rephrased as coevolving (i) such values of the velocities, shown in the octet in Equation (1), together with (ii) the angular offset of the sensor, resulting in an efficient capturing behaviour of the team of predator agents. We shall elaborate on such a coevolution in the next section. MAS, as a complex system, feature a significant semantic gap between the simple, hierarchically lower-level properties of the agents, and the more elaborate, higher-level behaviour of the whole system. Consequently, we would be unable to formally infer the values of the octet of velocities of the wheels of agents from the desired behaviour of the team of such agents. Similarly, we are unaware of the value of the angular offset of the sensor – resulting in an efficient capturing behaviour of the agents. Moreover, the values of velocities of both wheels and the value of the angular offset of the sensor would, most likely, be dependent on each other.

5.2 Experimental Results

5.2.1 Coevolving the Asymmetric Morphology and the Behaviour of Predator Agents

As Figure 26, Figure 27 and Figure 28 illustrate, just by adding the offset, the results in number of successful initial situations and overall fitness significantly improves compared to the evolution of the team of straightforward predator agents featuring no angular offset for the sensors. On average, the predators were able to resolve all 10 initial situations by 10th generation of the GA. From all 32 independent runs of GA, there is one distinguished solution (from now on we will refer to it as the fastest evolved solution SFE) which successfully solves 8 initial situations (out of 10) in the first generation. The chromosome of this solution encodes for offset of the sensor of 20°. This confirms the findings in our previous research [23] [37], that a team of predators with 20° sensor offset yields favourable results during evolution. As we will discuss later, this is also true in case of additional – unforeseen, situations and presence of perception noise. However, from all 32 solutions, this is not the one that has achieved the best overall fitness value. The best agent behaviour (manifested by the achieved lowest of fitness value) was obtained by the solution SBF featuring a sensory offset of 16°. Compared to the fastest evolving solution SFE, the solution SBF evolved a bit slower and solved all 10 situations by 6th generation, achieving the terminal fitness of 369 (compared to 417 of solution SFE).

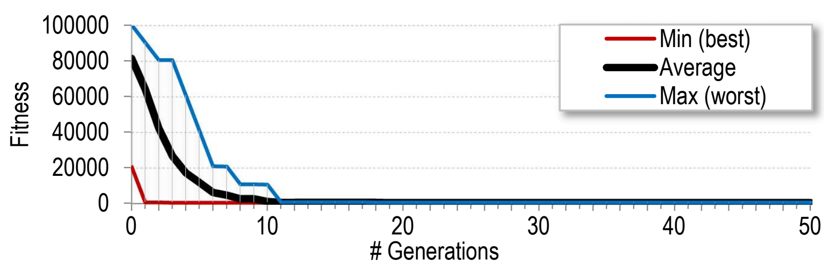

Figure 26: Convergence of the best fitness of 32 independent runs of GA

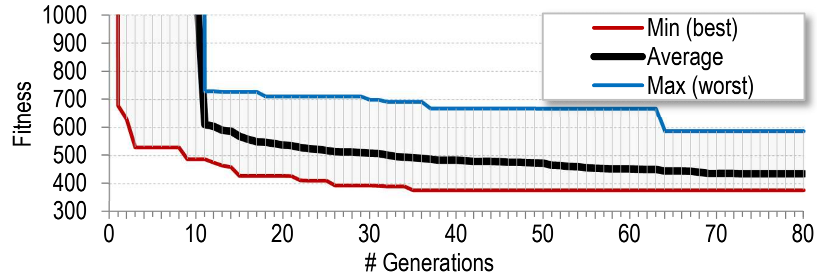

Figure 27: A more detailed illustration of the convergence of the best fitness of 32 independent runs of GA

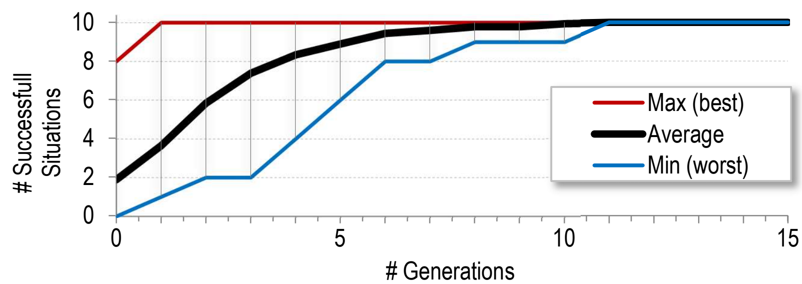

Figure 28: Convergence of the number of successful situations of 32 independent runs of GA.

Figure 29 illustrates the angular offset of the best solutions obtained from each of the 32 independent runs of the GA. As seen in Figure 29, the fitness of 80% (i.e., 26 of 32) of solutions is in the range between 369 and 448, i.e., the team of agents could capture the prey (on average over all 10 situations) between 36.9 s and 44.8 s into the allocated 120 s of the trial. The fitness of the worst solution is 622, meaning that the team of predator agents captures the prey, on average, at 62.2s, i.e. around the middle of the 120 s trial. Moreover, as Figure 29 illustrates, for a particular value of the sensor offset there are multiple solutions with different fitness values, meaning that there are variations in the behaviour of the morphologically identical predators, and that the sensory asymmetry is only a precondition for an efficient capturing behaviour of the predators. Analogically, very similar fitness values could be achieved by predators featuring different sensor offset, suggesting that the combination of both (i) the morphology and (ii) the behaviour, rather than a particular instance of each of them, is important for the success of the behaviour of predator agents.

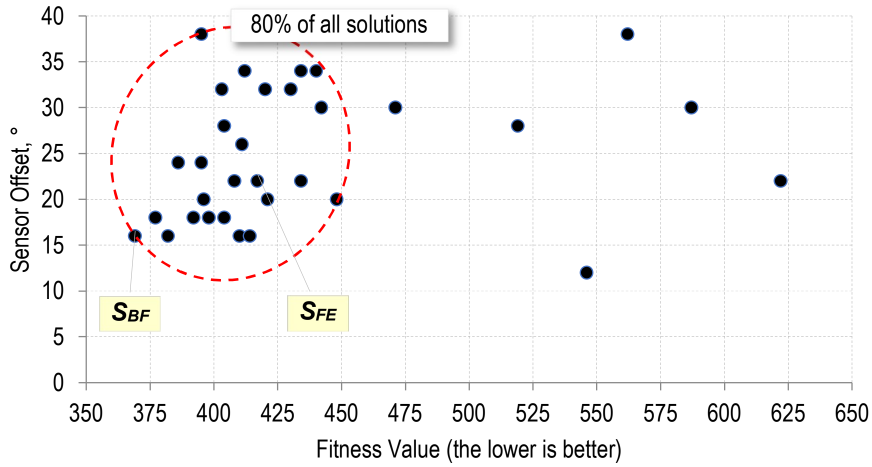

Figure 29: The sensor offset and the fitness value of all 32 solutions obtained form 32 independent runs of the GA The fastest evolved- and the best overall solutions are denoted as solutions SFE and SBF, respectively.

The breakdown of the number of the successful situations and the sensor offset of all 32 solutions are illustrated in Figure 30. As depicted in Figure 30 (right), the sensor offset of 90% (i.e., 29 of 32) of solutions is within the range (15° … 35°). There is no evolved solution that features a sensor offset lower than 10°, which confirms experimentally our initial hypothesis about the beneficial effect of the asymmetric morphology of predators on the efficiency of their behaviour. The statistical characteristics of all 32 solutions are shown in Table 19.

Table 19: Statistical characteristics of the 32 solutions obtained form 32 independent runs of the GA

Parameter	Value
Mean of the best fitness values	436
Standard deviation of the best fitness value	63
Mean of the sensor offset, °	24.7
Standard deviation of the sensor offset, °	7.2

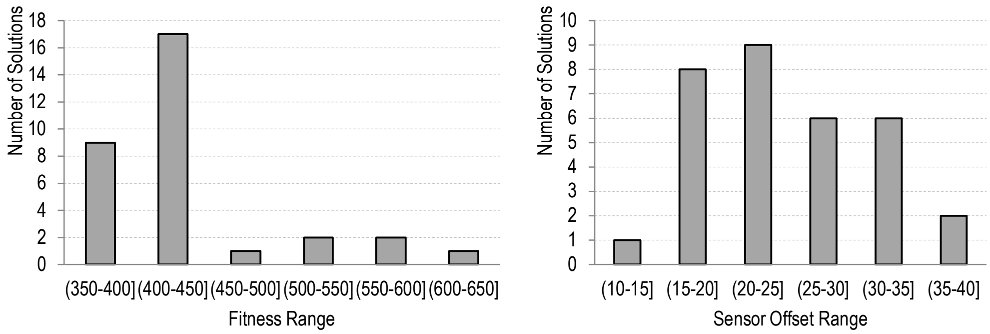

Figure 30: The breakdown of the number of the successful situations (left) and the sensor offset (right) of all 32 solutions obtained form 32 independent runs of the GA.

5.2.2 Generality of the Evolved Solutions

To assess the generality of the evolved behaviour of the predator agents, we will examine how their performance (i.e., the number of successfully resolved initial situations) degrades with the increase of the speed of the escaping prey. We tested all 32 solutions, obtained via the GA (for the speed of the prey equal to 10 units/s), for speeds of the prey, unforeseen during the evolution, of 12, 14, 16, 18, and 20 units/s, respectively. The number of initial situations, successfully solved by each of the 32 solutions for each of the considered speed of the prey is shown in Figure 31. The mean (over the whole range of speeds of the prey) of the successfully solved situations by each of these solutions, and its breakdown are depicted in Figure 32. As these figures illustrate, one of these solutions – denoted as S_{MG} – is most general in that is features no degradation in the number of successful situations with the increase of the speed of the prey. Moreover, its fitness value remains under 500 (i.e., the agents capture the prey earlier than 50 s into the 120 s trial) for all considered speeds of the prey. As shown in Table 20, the sensor offset of S_{MG} is 24°.

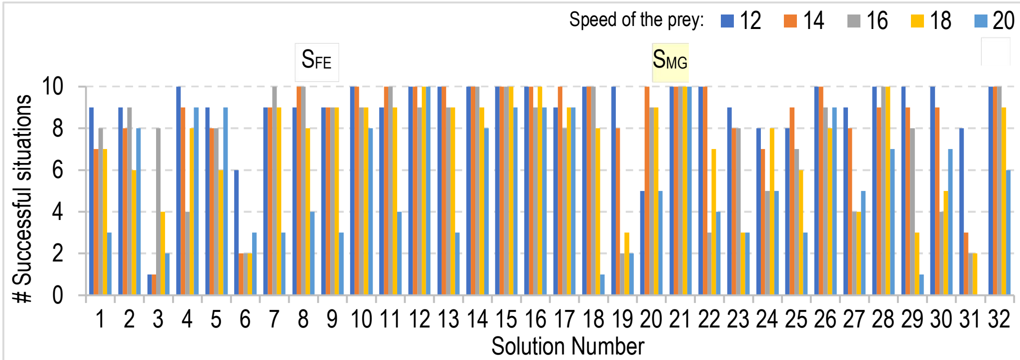

Figure 31: The number of successfully solved situations by the evolved 32 solutions for the speed of prey being increased from 10 to 12, 14, 16, 18 and 20 units/s, respectively.

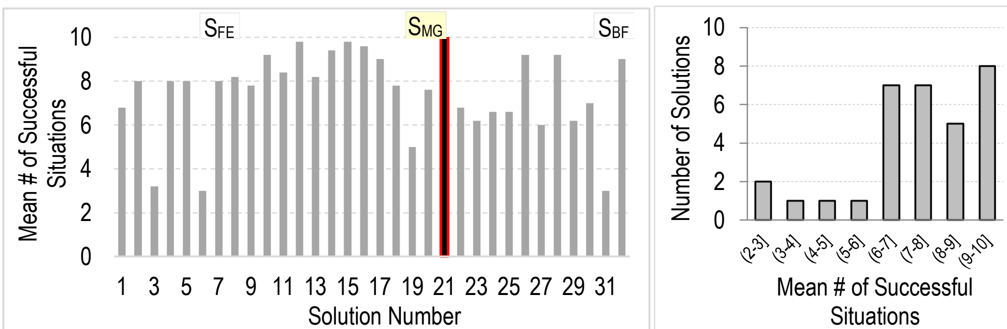

Figure 32: Generality of the evolved 32 solutions to the changes in the speed of prey from 10 to 12, 14, 16, 18, and 20 units/s: the mean number of successfully solved situations (left) and its breakdown (right).

5.2.3 Robustness to perception noise.

We evaluated the robustness of the 32 evolved solutions, evolved in a noiseless environment, to a random perception noise. We introduced two types of noise – a false positive (FP) and a false negative, respectively. The former results in either of the two bits of perception information to be occasionally (with a given probability) read as "1" regardless of whether an entity is detected in the line of sight of the predators or not. False negative noise (FN) results in readings of "0" even if an entity is seen. We focused on these types of noise as we assume that the perception subsystem of predators, yet being rather simple, would require an appropriate thresholding of the sensory signal. A combination of unfavourable factors, such as incorrectly established threshold, variable noise levels in the environment or in sensors would result in the considered two types of perception noise. Figure 33 and Figure 34 show the dynamics of the number of successfully solved situations by all 32 solutions for different amount of FP and FN perception noise, respectively.

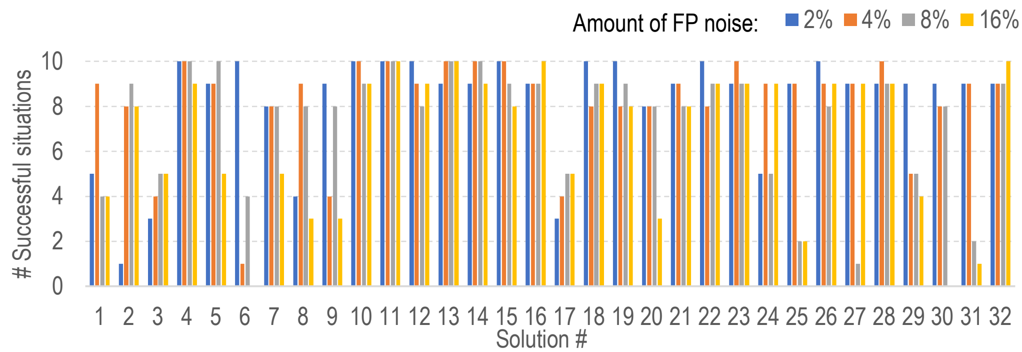

Figure 33: Robustness to FP noise of each of the 32 evolved solutions.

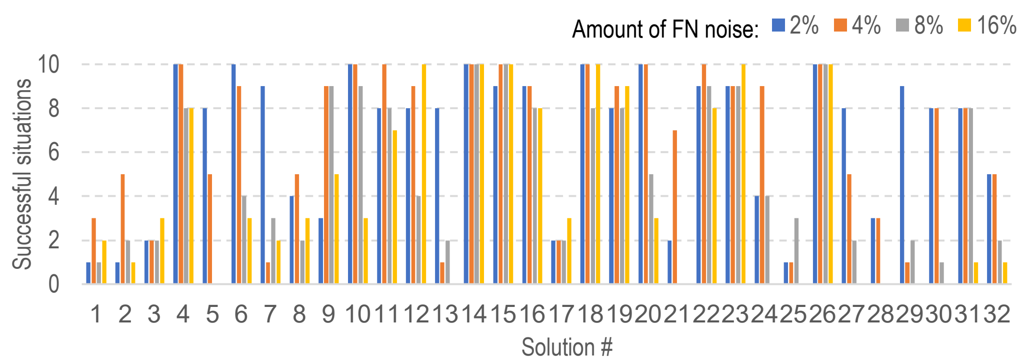

Figure 34: Robustness to FN noise of each of the 32 evolved solutions.

As Figure 33 and Figure 34 illustrate, neither the fastest evolved solution S_{FE} nor the solution with the best fitness S_{BF}, which we previously discussed, feature a good robustness to perception noise. On average they solve 6.25 initial situations each, with the introduction of either FP or FN noise. Both solutions yield similar results with the difference between them being that S_{BF} is more robust to FP noise while S_{FE} is better is FN one. Instead, the solutions S_{MRFP} and S_{MRFN} (featuring a genotype as shown in Table

20) emerge as most robust to FP noise and FN noise, respectfully. Solution S_{MRFP} manages to solve the tests with FP noise perfectly, while maintaining satisfactory performance in the tests with FN noise, being able to solve on average 8.25 initial situations, depending on the level of FN noise. On the contrary, the agents controlled by S_{MRFN} solve the situations with FP noise perfectly, while being able to an average of 9.5 initial situations in the situations with FN noise, resulting in the best overall performance. The sensor offset of S_{MRFP} and S_{MRFN} is 18° and 20°, respectively (Table 20).

Table 20: Genotype of evolved solutions: the fastest evolved (SFE), with the best fitness (SBF), most general (SMG), most robust to FP (SMRFP) and FN (SMRFN) noise

Solution	Fitness	V_{00L}, %	V_{00R}, %	V_{01L}, %	V_{01R}, %	V_{10L}, %	V_{10R}, %	V_{11L}, %	V_{11R}, %	Sensor Offset α,°
S_{FE} (#9)	417	30	95	100	90	-80	-75	50	-95	22
S_{BF} (#32)	369	-95	80	90	85	-90	-90	100	90	16
S_{MG} (#21)	382	-95	80	95	90	-90	-90	60	-10	24
S_{MRFP} (#11)	404	-70	70	90	85	-100	-100	65	70	18
S_{MRFN} (#14)	421	30	100	100	95	-75	-70	100	100	20

5.3 Discussion

5.3.1 The Advantage of Asymmetric Morphology

The experimental results indicate that the introduction of the angular offset of the sensor of the agents improves the effectiveness of the team by enabling the emergence of a general and robust (to environmental noise) capturing behaviour. It also helps increase the efficiency of the evolution of successful capturing behaviour of agents. Our previous work [23], as well as the result provided above suggest that the behaviour, evolved with a sensor offset of 20° (in solution S_{MRFN}) is most robust to noise and is close enough in terms of fitness to the best performing team of agents in noiseless environments. The fitness of S_{MRFN} is 421 compared to 369 of S_{BF}. While S_{MG} shows best results in the generality test, with perfect score in all initial situations, it falls short in the noise robustness test. This leads us to believe that S_{MRFN} is an example of a good combination of coevolved behaviour and asymmetric morphology of the predator agents. On average, S_{MRFN} manages to solve 9.57 and 9.65 situation in the generality- and robustness tests cases. The angular offset of 20° of S_{MRFN} provides a good trade-off between the tangential- and radial (i.e., towards the prey) components of the speed vector of the chasing predators.

The beneficial effect of the sensor offset is in that it helps the chasing predator to implicitly determine the position of the prey if the latter disappears. Having a counter clockwise displacement means that most of the time the disappeared prey, due to the parallax induced by the movement of the predator, would be to the left, and consequently – a slight turn to the left would allow relocating it again. Therefore, one of the virtues of the sensor offset is in the *more deterministic direction* of the disappearance of the prey – almost certainly to the left – which, in turn facilitates a faster rediscovery, and consequently – a more reliable tracking of the latter by the predator. Moreover, as shown in Figure 35, the chase by the predator featuring an asymmetric morphology would result in a characteristic *circular trajectory* of both the predator and the prey. This behaviour is more efficient, compared to that of the predators with canonical straightforward morphology. In the latter case the predator could not make any reliable prediction about the most likely direction of the disappearance of the chased prey. With the rather challenging, but realistic assumption that initially the prey is not being surrounded by the predators (as illustrated in Figure

20) such emergent circular trajectories would facilitate the surrounding as the prey would be shepherded (driven) by a single predator (driver) towards the pack of the remaining predators.

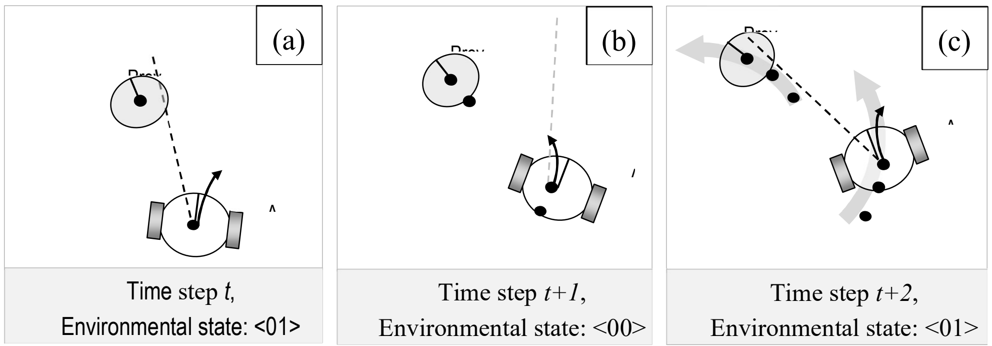

Figure 35: Chasing the prey by a sample predator agent Ai

Significantly reducing the sensor offset from 20° would stretch the chasing trajectory of the predator. Moreover, in such case, as the chased prey becomes closer to the longitudinal axis of the predator, in order not to compromise the certainty that the prey has disappeared to the left, it would need to turn slightly to the right (instead of going straight, as shown in Figure 35, for environmental state <01>)—by reducing the speed of the right wheel—during the chase. This, in turn, would reduce the overall chasing speed of the predator. These two factors—stretching the chasing trajectory and reducing the chasing speed of the predator—would result in increasing the time needed to drive the prey towards the dispersed predators. Conversely, increasing the sensor offset would result in a more compact chasing trajectory that might not stretch enough to reach back to the pack of remaining predators.

5.3.2 Emergent behavioural strategies

The current research, as well as our previous work on agents with asymmetric morphology [23], suggest that the solution most robust to noise and with greatest success rate in the generality tests is the behaviour obtained from the evolutionary run

17. In this section, we will use that specific behaviour to review the behavioural strategies of the team of predator agents emerging from evolution of the velocity mappings by the GP framework.

The values of the evolved velocities of motors and the sensor offset of behaviour #17 are shown in Table 18. The team of predator agents exhibits the following four behavioural strategies, executed in four consecutive phases of the trial: (i) *circling* around until they find a peer or the prey (controlled by velocities V_{00}), (ii) *exploring* the environment by distancing themselves from each other (controlled by velocities V_{10}), (iii) *surrounding* by shepherding (driving) the prey (by some of the predators - drivers) in a circular trajectory (V_{01}), and (iv) *capturing* the prey (V_{11}).The team of predator agents exhibits the following three behavioural strategies, executed in three consecutive phases of the trial: (i) circling around until they find a peer or the prey (controlled by velocities V_{00}), and then *exploring* the environment by distancing themselves from each other (controlled by velocities V_{10}), (ii) *surrounding* by shepherding (driving) the prey (by some of the predators - drivers) in an circular trajectory (V_{01}), and (iii) *capturing* the prey (V_{11}). Figure 36 illustrates the different phases the agents go through in the process of catching the prey. A video of how the team of predators deals with all 10 initial situations can be found at http://isd-si.doshisha.ac.jp/m.georgiev/2018-12-08-SA20deg.mp4

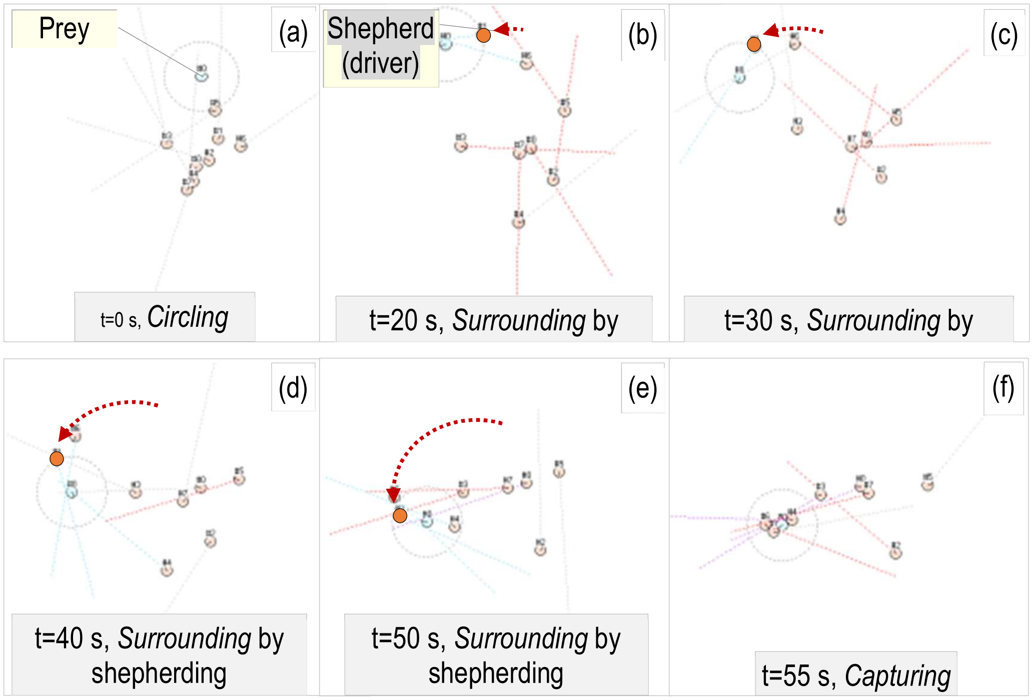

Figure 36: Emergent behavioural strategies of a sample evolved team of predator agents with sensor offset of 20°. Environmental state perceived by predator: grey = <00>, red = <10>, blue = <01>, purple = <11>.

As shown in Figure 36a, in the beginning all agents have no vision of either the prey or any of the peers. Following the mapping of V_{00L} = 25% and V_{00R} = 100% (as shown in Table 18), they start *circling* around—scanning the environment in an attempt to find another entity. Detecting a peer activates the set of velocities V_{10L} = −25% and V_{10R} = −20%, which forces the predators to enter the second stage: to move away from the perceived agent, which facilitates a better dispersion and a coverage of a wider area. This enhances the ability of the predators to *explore* the environment and to discover the prey. The third stage—*surrounding*—begins when any of the predators discovers the prey. The mapping V_{01L} = 100%, and V_{01R} = 100% results in moving forward at the highest speed, which helps in keeping the prey almost always in the same relative position to the agent—i.e., on the left side, as shown in Figure 25, 9b–e. Once the prey disappears from view—as shown in the centre panel of Figure 25—the predator exhibits an embodied cognition that the disappearance is a result, in part, of its own forward motion; therefore, the new location of the prey is—due to the counter clockwise offset of the sensor—most likely on the left of its own orientation. Then the evolved V_{00L} = 25% and V_{00R} = 100% are activated (Figure 25 right) resulting in a circular motion to the left, until the agent rediscovers the disappeared prey.

Moreover, as Figure 36b–d show, a single predator—driver—due to its sensor offset, shepherds (drives) the prey in a circular, counter clockwise trajectory into the (already dispersed) other predators. The fourth (final) behavioural phase concludes the chase by *capturing* the prey that is surrounded from all sides by both the driver(s) and the newly encountered predators, as illustrated in Figure 36e,f. When approaching from opposite sides, the predators are able to see both the prey and a peer, which activates the mapping V_{11L} = 100% and V_{11R} = 100%. Since they have a slight angular offset, it is possible for only three predators to catch the prey, as illustrated in Figure 36e, f. One of the predators chases the prey from behind and guides it to its front left side, while the other approaches it from the opposite direction.

At the same time, as shown in Figure 36d–f, two of the agents keep distancing themselves from the group of other predators. The agents seem to exhibit an emergent knowledge [38] that not all eight agents are needed to capture the prey. For the group of agents to be successful, the most important mission is to capture the prey, rather than which particular agent does it. As the performance of the predators is calculated, based on the success of the group, instead of that of the particular individual agent, such behaviour helps the team (as a whole) by expanding the search field and finding the prey faster, especially when it is further away from the predators. If, instead, the agents were trying to find the prey and capture it by themselves via "greedy chase", they would inevitable fail because: (i) the prey is fast enough to run away from a single predator, and (ii) the predators would have been unable to engage in any organized behaviour that allows surrounding, and ultimately – capturing the prey.

The manifestation of shepherding behaviour in the third behavioural phase is probably the most significant difference between the evolved behaviour of the canonical straightforward predator agents and that of the agents with asymmetric morphology. This behaviour, being a vital part of the successful capturing, could not be observed in the behaviour of the canonical predator agents. A video showing how the canonical agents struggle to find and capture the prey is available at http://isd-si.doshisha.ac.jp/m.georgiev/2018-12-03-SA-Straightforward.mp4

5.3.3 Alternative methods

We could have adopted another – deterministic – approach, such as, for example, a complete enumeration of the possible combinations of the eight velocities of wheels and the sensor offset. If each of these 8 velocities is discretized into, say, 40 possible integer values ranging from -100% to +100%, and the sensor offset – just into 20 values – then the size of the resulting search space would be equal to 40^8 or about 1.3×10^{14}. This would render the eventual "brute force" approach, based on complete enumeration of possible combinations of values of velocities computationally intractable.

As an alternative to the brute force search, we could apply reinforced learning (RL) in order to define the good mapping of the four perceived environmental states into the four pairs of velocities of wheels. However, MAS are complex, non-linear systems, and there is a significant gap between the properties of the entities and the (emergent) properties of the system. RL would obtain a "reward" from the system (i.e. the efficiency of the team of predators), and will try to modify the properties (the four pairs of velocities of wheels) of the entities. Due to the complexity and non-linearity of MAS, this is not a straightforward task. This is also related to the intra-agent credit- (or blame-) assignment problem, as we could not tell which part of the agents is responsible (and therefore – should be modified) for the bad overall behaviour of the system.

Evolutionary computing solves these challenges in an elegant way – by obtaining the fitness value from the system, as a whole (i.e., the efficiency of predators in capturing the prey) and then modifying the properties of entities (pairs of velocities of wheels of predators) via genetic operations - crossover and mutations.

Yet another challenge in RL is the delayed reward problem – the success (if any) of the system (team of predators) would occur several hundred time-steps into the trial, but might be related to the earlier behaviour phases of the team of predators – such as the dispersing (exploration of the environment) at the very beginning of the trial. Regarding the delayed reward problem, the evolution, as a holistic approach, does not care about how to achieve the success, but rather – about the overall (final) outcome of the trial.

In our work we apply GA – a nature-inspired heuristic approach that gradually evolves the values of a set of parameters in a way similar to the evolution of species in nature. GA has proved to be efficient in finding optimal solution(s) to combinatorial optimization problems featuring large search spaces [32] [33] [39]. Thus, consonant with the concept of evolutionary robotics [34], we adopted GA to evolve the values of the eight velocities of the wheels and the offset of the sensor that result in an efficient behaviour – presumably involving exploring the environment, surrounding-, and capturing the prey – of the team of predators.

5.4 Additional tests

5.4.1 Generality of the Evolved Behaviour

After seeing the success which the sensory offset brought to solving the initial 10 test situations, we investigated the generality of the best-evolved teams of predators on an extended set of 100 initial situations, including the same 10 initial situations used in the evolution, and 90 additional situations unforeseen during the evolution. For the 10 initial situations (from situation #1 to situation #10) used during the evolution, the predators are dispersed south of the prey (as illustrated in Figure 20) such that the average distance between the prey and the group of predators slightly increases with each situation. For the additional situations, (from situation #11 to situation #100) the average distance between the predators and the prey is kept the same as that of the situation with the most distant predators used during the evolution—situation #10. In each of the additional situations #11…#100 the average distance of the predators to the prey is the same as that of situation #10; however, both the position and the orientation of the predators are random. We would like to note that the alternative approach of exploring the generality of the agents by evolving them directly on 100 initial situations would feature a greater computational overhead as the total number of fitness trials would be almost 10-fold higher.

We conducted the experiments on the set of these 100 initial situations with the 32 evolved best-of-run teams of predators featuring a sensor offset of 20°. Our choice of this particular morphological configuration is based on its superiority in terms of both (i) the quality (fitness value) of the evolved best-of-run teams of agents and (ii) the consistency (probability of success) in evolving these teams, as illustrated in Table 17. Figure 38 shows the experimental results of the number of successfully solved initial situation by each of the 32 evolved best-of-run teams of predators. The results shown in Figure 37 demonstrate that the sensory offset of 20° results in behaviours of predator agents that are quite general (rather than over-fitted to particular initial situations). Indeed, only three solutions resolved less than 60 of all 100 initial situations (obtained from evolutionary runs #11, #22, and #25), while, the best solution (obtained from evolutionary run #17) resolves 97 initial situations. The evolved mapping of the velocities of wheels of the most general solution #17 is shown in Table 18.

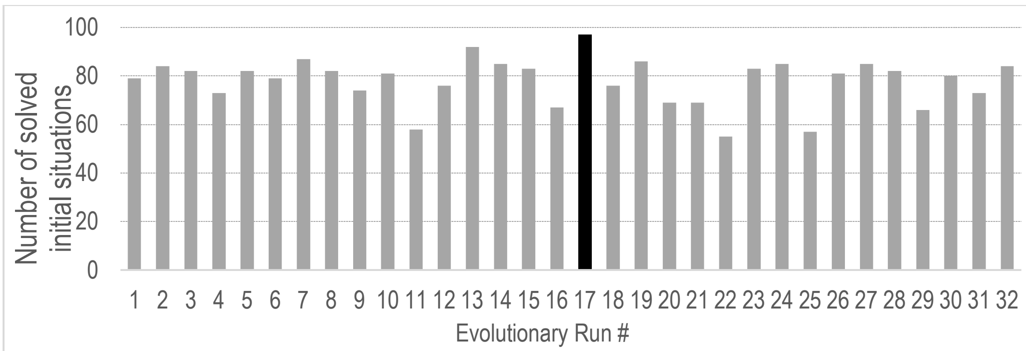

Figure 37: Generality of the 32 evolved best-of-run behaviours of the team of predator agents.

5.4.2 Robustness to Sensory Noise

After verifying the generality of the 32 evolved best-of-run behaviours of the team of predator agents, we decided to investigate how each of these behaviours degrades when a random noise is introduced into the perception readings. Revisiting the same set of 100 initial situations (including the same 10 initial situations used in the evolution, and 90 additional unforeseen situations), we introduced two types of noise—false positive (FP) and false negative (FN), respectively. Under the influence of the FP noise, the value of either (randomly chosen with probability of 50%, individually for every predator agent, on each time step) of the two bits of sensory information is read as one, regardless of the actual reading of the sensor. On the contrary, in the presence of FN noise, the reading is registered as zero, even if the corresponding entity (a predator or the prey) is in the line-of-sight. We conducted experiments for levels of either FP or FN noise, starting from 2% and increasing by multiplier of 2, up to 16% (i.e., 2%, 4%, 8% and 16%).

The experimental results obtained, as shown in Figure 38, are somehow unexpected. We anticipated that that the number of solved initial situations would decrease in a monotonic way with the increase of the level of FP noise. However, most of the 32 best-of-run behaviours are quite robust to FP noise in that the number of solved initial situations varies slightly compared to the number of solved situations in a noiseless environment. Moreover, often the number of solved situations anomalously *increases* with the increase of noise levels. Notable behaviours that exhibit a slight increase in the number of solved situations with the increase of noise are, for example, those obtained from evolutionary runs #7 and #15. For behaviour #15, the number of successful situations (91) for 4% noise is significantly higher than that obtained in noiseless environments (91 vs 83, respectively, or a 10% increase).

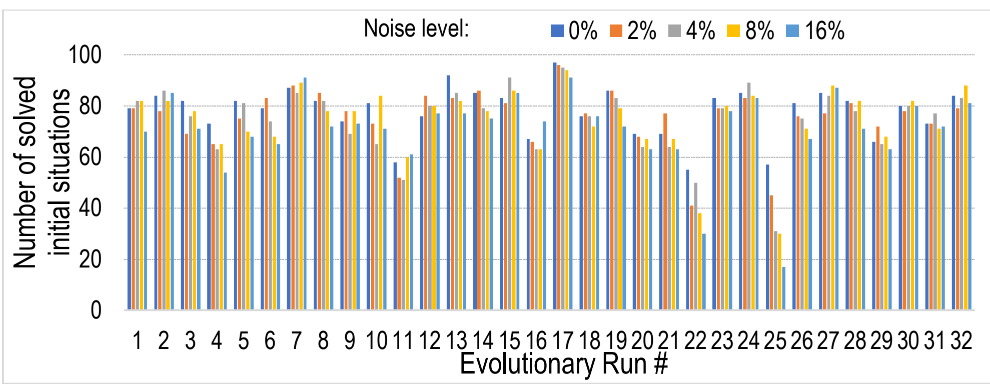

Figure 38: Number of successfully solved initial situations for various levels of false positive (FP) perception noise.

As illustrated in Figure 39, the detrimental effect of FN noise is as we expected – for most of the 32 best-of-run behaviours the number of solved initial situations steadily decreases with the increase of noise levels. Even so, there are some behaviours that exhibit an anomalous increase of the number of solved situations with the increase of noise levels. The notable behaviours are #7, #16, #19 and #28. Each of these behaviours result in an increase in the number of solved initial situations for 8%, 2%, 16% and 4% noise, respectively, as seen from Figure 38.

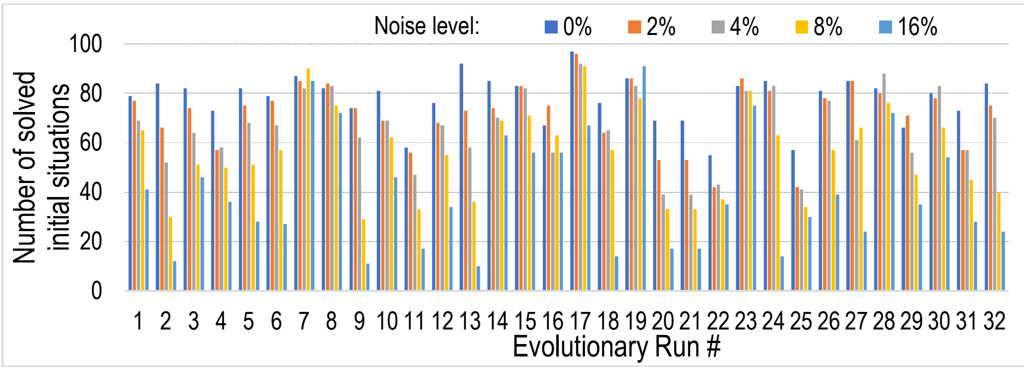

Figure 39: Number of successfully solved initial situations for various levels of FN perception noise.

These results demonstrate that both the type and the magnitude of perception noise have an influence on the robustness of the evolved behaviours of predator agents. The effect of FP noise on the behaviour of the team of predator agents is sometimes detrimental, sometimes favourable, or often insignificant. On the other hand, the FN noise with few exceptions (e.g., behaviours #7, #17, #19, #23 and #28) is detrimental for performance of the agents.

It is interesting to note that the evolved behaviour #17 (shown in Table 18) is rather versatile in that it (i) exhibits good performance on the 10 initial situations used during the evolution, (ii) is general to most unforeseen initial situations (solving 97 of 100 situations), and (iii) is robust to both FP and FN noise. Furthermore, behaviours like #7, #23 and #28, exhibit robustness to both types of noises and even increased performance for certain levels of noise.

5.5 Conclusions

We considered a society of very simple robots (modelled as agents in MAS) that feature an "extreme simplicity" of both sensors and control. The adopted agents have a single line-of-sight sensor, two wheels in a differential drive configuration as effectors, and a controller that does not require a memory and does not involve any computing, but rather a direct mapping of the currently perceived environmental state into a pair of velocities of the two wheels. Also, we applied genetic algorithms to evolve a mapping that will result in effective behaviour of the team of predator agents towards the goal of capturing the prey in the predator-prey pursuit problem (PPPP). The preliminary experimental results indicated that the simple agents featuring the canonical straightforward sensory morphology could hardly evolve the ability to solve the PPPP.

To enhance the performance of the evolved system of predator agents, we propose an asymmetric morphology featuring an angular offset of the sensor, relative to the longitudinal axis of the agents. The experimental results demonstrated that this modification brings—without compromising the simplicity of agents—a considerable improvement of both (i) the efficiency of evolution and (ii) the effectiveness of the evolved capturing behaviour of predator agents. Also, we verified that some of the evolved best-of-run behaviours of predators featuring a sensor offset of 20° are both (i) general in that they are able to successfully solve most of the additionally introduced, unforeseen initial situations, and (ii) robust to perception noise in that they show a limited degradation of the number of successfully resolved initial situations.

The results described in our work could be seen as a step towards the verification that complex behaviour needed for solving challenging tasks could emerge from the coordination of very simple robots featuring an asymmetric sensory morphology. The advantages of such robots, in addition to the simple design, include better robustness, higher throughput of production and lower production costs, reduced energy consumption, and the potential to be implemented at very small (nano- and micro-) scales.

Chapter 6: Improving the evolution speed of the system featuring simple predator agents

6.1 Evolving agents with different population sizes

The first evolution with low population size happened by a serendipity. In our previous work on simple agents we were using 32 individuals for testing. Since the software for testing and evolution is the same, when we needed to evolve additional chromosomes for more tests, we forgot to change the population size from 32 to 400. After seeing that the results with large and small populations are comparable, we decided to study how population size affects the evolution runs and the evolved chromosomes. The population size in each evolutionary run will be one of 16/32/100/400, as our goal is to test how the population size affects the performance of the produced solution. In Table 21, we can see how the time needed for the evolution changes as the population size changes. With lower population size we usually need less time to generate 32 chromosomes that solve the 10 initial situations – this is our criteria to stop the run. Table 21 shows the distribution of all generated solutions, grouped in three groups, using their success status. “Non-successful” means the run ended without producing a chromosome that was able to solve all 10 initial situations. “Successful with stagnations” means that the run was able to solve all 10 initial situations, but it was not able to reach the goal of fitness better than or equal to 600. The third group shows the number of chromosomes that was able to reach the ending goal of having fitness better than or equal to 600. The “Total” column shows the overall number of solutions that were generated during each run. We must note that our target during the evolutionary runs was to generate 32 chromosomes that solve all 10 initial situations and we don’t care how many of the extra chromosomes fail during that time.

We see the ability to generate more chromosomes than needed, in less time, as an advantage of the runs with lower population count.

Table 21: Statistics about the evolution with different population sizes

Population Size	Evolution Time (Hours)	Non-Successful	Successful with Stagnation	Successful reaching 600 fitness	Total
400	1732	**0**	**4**	28	32
100	1318	18	6	26	50
32	707	40	14	18	72
16	**385**	60	22	**10**	**92**

The resulting increase of stagnated chromosomes should not be a reason for concern. If we look at Table 22, we see that with lower population size, the fitness of the best individual improves. In our case we have the best fitness with the lowest population size of 16. Although, later we will show that the best fitness does not mean most general and robust chromosome.

Table 22: Performance results from the evolution with different population sizes

Population Size	Terminal Fitness				Successful with Stagnation	Successful without Stagnation
	Best*	Worst	Average	Stand. Dev		
400	466	**693**	**578**	**44.2**	4	28
100	514	805	598	66.4	6	26
32	512	781	629	59.5	14	18
16	**445**	870	657	97.8	22	10

* Best fitness during the evolution but not with highest generality or robustness

6.1.1 Generality of the Evolved Behavior

We investigated the generality of the best-evolved teams of predators on an extended set of 100 initial situations, including the same 10 initial situations used in the evolution, and 90 additional situations unforeseen during the evolution. For the 10 initial situations (from situation #1 to situation #10) used during the evolution, the predators are dispersed south of the prey (as illustrated in Figure 2) such that the average distance between the prey and the group of predators slightly increases with each situation. For the additional situations, (from situation #11 to situation #100) the average distance between the predators and the prey is kept the same as that of the situation with the most distant predators used during the evolution—situation #10; however, both the position and the orientation of the predators are random. We would like to note that the alternative approach of exploring the generality of the agents by evolving them directly on 100 initial situations would feature a greater computational overhead as the total number of fitness trials would be almost 10-fold higher.

We conducted the generality test using the set of 100 initial situations on all 128(4*32) successful chromosomes from the 4 groups of evolutionary runs using different population sizes. In

Table 23, we can see a comparison of all the chromosomes with best fitness during evolution and best chromosome from the generality tests. As mentioned before, we can see that, having the best fitness after the evolutionary run is over, does not mean the evolved individual will have a good generality. Later we will show that this is also valid for the robustness to noise that each chromosome exhibits.

Table 23: Most general chromosome (right) compared to the one with best fitness during evolution (left) for every group of different population sizes

Population size	400		100		32		16	
Fitness	**466**	541	**514**	**514**	**512**	593	**445**	589
Generality*	56	**94**	**94**	**94**	65	**97**	76	**95**

* Number of solved initial situations out of 100

A notable point is that the results from all the best individuals are comparable. The generality that they show is around 95%. This shows our initial hypothesis that the lower population size does not have a great impact on the quality of the generated solutions.

The evolved mapping of the velocities of wheels of the most general solution is shown in Table 24 – the one evolved with 32 population size.

Table 24: Velocity mappings of the most general and robust chromosome (top) and the one with best fitness during evolution (bottom)

V_{00L}	V_{00R}	V_{01L}	V_{01R}	V_{10L}	V_{10R}	V_{11L}	V_{11R}
25%	100%	100%	100%	− 25%	− 20%	100%	100%
-60%	75%	95%	90%	-100%	-100%	5%	-80%

6.1.2 Robustness to Sensory Noise

After verifying the generality of the 32 evolved best-of-run behaviors of the team of predator agents for each of the four evolutionary cases (total 128 individual chromosome), we decided to investigate how each of them behaves when a random noise is introduced into the perception readings. Using the same set of 100 initial situations, previously used in the generality test, we repeated the tests, after introducing two types of noise—false positive (FP) and false negative (FN), respectively. Under the influence of the FP noise, the value of either (randomly chosen with probability of 50%, individually for every predator agent, on each time step) of the two bits of sensory information is read as one, regardless of the actual reading of the sensor. On the contrary, in the presence of FN noise, the reading is registered as zero, even if the corresponding entity (a predator or the prey) is in the line-of-sight. We conducted experiments for levels of either FP or FN noise, starting from 2% and increasing by multiplier of 2, up to 16% (i.e., 2%, 4%, 8% and 16%).

For the sake of simplicity, we will show the results only for the top individuals of each evolutionary run. In Table 25, we have a comparison of the robustness test for the best performing chromosomes of each of the four evolutionary cases. We chose the individual chromosomes by calculating the average solved situations in all cases (with and without noise) and selecting the top one. The table shows that even with a great change in population size (ex. From 400 to 16), there isn't a great impact on the performance of the team of agents. Moreover, sometimes, the performance can increase – such is the result in the cases when population lowers from 400 or 100 to 32. In this case we can see that the number of solved situations increases for ¾ of the tests having false positive noise.

Table 25: Noise test results for the best chromosomes of all 4 evolution cases.

Population	False Positive Noise					False Negative Noise			
	0%	2%	4%	8%	16%	2%	4%	8%	16%
400	94	91	90	94	82	90	83	77	37
100	94	92	92	90	91	91	93	91	90
32	97	100	99	93	97	95	96	92	82
16	95	95	93	92	91	92	92	86	86

6.2 Discussion

6.2.1 The best versus the average

In the previous section, we've shown a comparison between the best (on average) individuals for each of the 4 evolution cases. This might lead to a question – why choose a single chromosome from each evolutionary setup instead of using the average individual comprised of the average values, of the motor mappings of all generated chromosomes. In Table 26 we can see the results of doing so:

Table 26: Generality and noise test results for 4 synthetic chromosomes made with the average motor mappings from the best selected chromosomes in each run

Population	False Positive Noise					False Negative Noise			
	0%	2%	4%	8%	16%	2%	4%	8%	16%
400	73	66	71	69	65	70	54	43	29
100	59	66	71	65	61	67	60	58	34
32	54	58	59	54	54	57	49	47	25

16	52	52	57	56	54	52	48	50	25

These results suggest that the eight motor mappings in each chromosome evolve finely tuned interactions which are lost when the mappings of different genomes are averaged.

6.2.2 Lowering the population size even more

While it may be possible to lower the population size even more, we decided to stop at size 16. After this point, the population may be too low to allow proper use of genetic algorithms, such that the elitism and cross-over may be affected. In the end it will be reduced to random mutation.

6.3 Conclusions

In our research, we considered an implementation of the predator-prey pursuit problem (PPPP), involving a team of very simple robots that feature an extremely simple design for both control and sensor abilities. The robots, considered as agents in a multi-agent system, have a single line-of-sight sensor, two wheels in a differential drive configuration and a controller that does not require any computational effort or memory. Instead, there is a direct mapping of the currently perceived environmental state into a pair of velocities for the two wheels. We used the findings in our previous research to improve the performance of the robots by introducing a counterclockwise offset of 20 degrees to the sensor, relative to the agents' longitudinal axis. This was needed because the agents featuring a straightforward sensor were struggling to find solutions to the 10 initial situations used for the evolution of their behavior. We have also increased the range of the vision that the sensor provides to 400 units, for this experiment, as we determined that this increase yields better results. We applied genetic algorithms to evolve a mapping that will result in effective behavior of the team of predator agents towards the goal of capturing the prey in the predator-prey pursuit problem (PPPP). Inspired by serendipity, we hypothesized that lowering the population size will not have any significant impact on the quality of the solutions produced by the genetic algorithm. Furthermore, we found out that we were able to find mappings with a comparable performance, for less time, than the runs with higher population size. All of this leads to reduction of the hardware and computational resources needed, which would mean lower production and exploitation costs, and higher production throughput.

Chapter 7: Summary, Conclusion and Future work

7.1 Summary

Multi-agent systems (MAS) have proven to be an important tool for problem solving and have become widely applied. Due to their complex, non-linear nature, MAS can often provide efficient solutions to problems where monolithic systems are unable to produce an acceptable answer in terms of speed or resource requirements. However, their strength also makes implementation of MAS difficult to implement as an optimal solution to the problem is hard to obtain analytically

In our research we have investigated two relatively orthogonal ways of improving the overall performance of MAS: (i) minimizing the time needed by MAS to solve a given problem by evolutionary optimizing (coevolving) both the morphology and behaviour of agents, and (ii) minimizing the runtime needed by evolutionary framework – genetic programming – to successfully accomplish such a coevolution. As an evolutionary framework we adopted the in-house XML-based genetic programming (XGP), which offers a flexible, human-readable, and cross-application compatible XML representation of the genotype of evolved agents.

In attempt to improve the cost-effectiveness of developing a MAS, we tested how two novel approaches compare to a traditional MAS – (i) a heterogeneous MAS, featuring greater disparity of agents and (ii) a MAS comprised of very simple agents.

In our research we've used in-house implementations of the well-studied but difficult to solve predator-prey pursuit problem (PPPP), since it is an ideal target for solving by MAS.

Our initial results have shown that both (i) the speed of evolution of the successful capturing behaviour of predator agents and (ii) the effectiveness (i.e., its fitness value) of the best-evolved behaviour, of the heterogeneous MAS are improved, using different methods and techniques, in such a way that it performs better than its homogeneous counterpart. We have shown that a heterogeneous system featuring agents with a 10% deviation from the average capabilities of the system will offer faster evolution and better performing agents. However, the homogeneous system was still more general, in a way that it could solve more additional, previously unforeseen situations. The homogeneous system is more consistent and efficient in a presence of sensory noise, making it more robust. This disparity is caused by the increased search space in the heterogeneous system. To overcome this problem, we have increased the number of training initial situations. After the introduction of bigger training set, the heterogeneous system was able to catch up in robustness, in some of the test cases. However, this increase in evolution time is opposing our goal of a faster evolution.

Next, inspired by previous research on simple robots, we have suggested an implementation of non-computing predator agents, equipped with a single line-of-sight sensor and a simple control of velocities of their two wheels. Our goal was to have a system where the predator agents can catch (make contact with) the prey – a feature that other similar researches, did not have. We applied a GA to evolve the successful behaviour of the team of predator agents. Immediately, we've seen a problem, since the MAS was unable to solve many of the initial test situations. The suggested morphology was good enough in a task where simply clustering of objects was required, but not in our specific scenario. To enhance the generality of the evolved behaviour, we proposed an asymmetric morphology of the predators. The performance of the system was greatly increased. Inspired by this success, we've introduced the sensory offset into the genetic algorithm, in order to evolve the optimal value that will show best performance.

Offsetting their sensors angularly between 20° and 30° yielded the most efficient and consistent evolution of successful behaviours of agents.

During our work, we usually used a greater generation pool for evolution than we did for testing of the evolved solutions. Due to serendipity, we have discovered that lowering the population size will not have any significant impact on the quality of the solutions produced by the genetic algorithm. Furthermore, we found out that we were able to find mappings with a comparable performance, for less time, than the runs with higher population size. All of this leads to reduction of the hardware and computational resources needed, which would mean lower production and exploitation costs, and higher production throughput

7.2 Conclusion

One of the most desired features of autonomous robotic systems is their ability to accomplish complex tasks with a minimum amount of resources, lowering their implementation and exploitation costs. Often, however, the limited physical abilities should be compensated by more precise and complex control. An optimal trade-off between the simplicity of their morphology and control would result in robots featuring better robustness, higher throughput of production and lower production costs, reduced energy consumption, and the potential to be implemented in solution to novel problems. In our work we've proposed several approaches to minimizing the resources needed to build and evolve efficient MAS. These methods include the use of heterogeneous systems, simplifying the morphology to feature a reactive model and minimizing evolutionary time by reducing the computational overhead. These techniques were able to produce solutions comparable or even better (in some cases), to the one evolved using standard complex MAS.

7.3 Future Work

In our future work, we are planning to investigate the anomalous increase of the number of successful situations with the increase of false positive noise. While similar phenomena are well known in engineering (e.g., stochastic resonance, dithering) there are no documented results on the beneficial effects of noise on the performance of MAS. Also, we are planning to develop an even more realistic, three-dimensional model of the environment of PPPP, including simulation of Brownian motion, which may affect the behaviour of nano scale robots.

Bibliography

[1] J. Ferber, "An introduction to Distributed Artificial Intelligence," Harlow, 1999.

[2] T. Rose, "The End of Average," 2016.

[3] T. Haynes and S. Sen, "Evolving behavioral strategies in predators and prey.," in *International Joint Conference on AI*, Montreal, QC, Canada, 1995.

[4] T. Haynes, R. Wainwright, S. Sandip and D. Schoenefeld, "Strongly Typed Genetic Programming in Evolving Cooperation Strategies," in *Sixth International Conference on Genetic Algorithms*, 1995.

[5] I. Tanev and K. Shimohara, "XML-based genetic programming framework: design philosophy, implementation, and applications," in *Artificial Lifeand Robotics*, 2010.

[6] M. Fischer, N. Lynch and M. Paterson, "Impossibility of Distributed Consensus with One Faulty Process," in *JACM*, 1985.

[7] B. Wang and Y.Sun, "Consensus analysis of Heterogeneous Multi-Agent Systems with Time-Varying Delay," in *Entropy*, 2015.

[8] I. Tanev, M. Brzozowski and K. Shimohara, "Evolution Generality and Robustness of Emerged Surrounding Behavior in Continuous Predators-Prey Pursuit Problem," in *Genetic Programming and Evolvable Machines*, 2005.

[9] M. Gauci, J. Chen, W. Li, T. Dodd and R. Groß, "Self-organized aggregation without computation," in *IJRR*, 2014.

[10] M. Gauci, J. Chen, W. Li, T. Dodd and R. Groß, "Clustering objects with robots that do not compute," in *International Conference on Autonomous Agents and Multi-Agent Systems*, Paris, France, 2014.

[11] A. Özdemir, M. Gauci and R. Groß, "Shepherding with Robots That Do Not Compute," in *14th European Conference on Artificial Life*, Lyon, France, 2017.

[12] S. Bhattacharya, R. Murrieta-Cid and S. Hutchinson, "Optimal Paths for Landmark-Based Navigation By Differential-Drive Vehicles With Field-Of-View Constraints," in *IEEE Transactions on Robotics*, 2007.

[13] J. Yu, S. LaValle and D. Liberzon, "Rendezvous Without Coordinates," in *IEEE Transactions on Automatic Control*, 2012.

[14] M. Gauci, "Swarm Robotic Systems with Minimal Information Processing," Sheffield, UK, 2014.

[15] M. Benda, V. Jagannathan and R. Dodhiawala, "On Optimal Cooperation of Knowledge Sources," Bellvue, WA, USA, 1986.

[16] S. Luke and L. Spector, "Evolving Teamwork and Coordination with Genetic Programming," in *First Annual Conference on Genetic Programming*, Stanford, CA, USA, 1996.

[17] P. Diamandis, "Where We Are Today and Why Their Future Has Amazing Potential," in *Singularity Hub*, 2016.

[18] S. M. Martel, P. G. Madden, L. Sosnowski and I. W. Hunter, "NanoWalker: A Fully Autonomous Highly Integrated Miniature Robot for Nanoscale Measurements," in *SPIE 3825 Microsystems Metrology and Inspection*, 1999.

[19] L. Mertz, "Tiny Conveyance: Micro- and Nanorobots Prepare to Advance Medicine," in *IEEE Pulse*, 2018.

[20] A. Bhat, "Nanobots: The Future of Medicine," *International Journal of Management and Engineering Sciences,* vol. 5, no. 1, pp. 44-49, 2014.

[21] D. H. Kim, U. K. Cheang, L. Kohidai, D. Y. Byun and M. J. Kim, "Artificial magnotactic motion control of Tetrahymena pyriformis using ferromagnetic nanoparticles–A tool to fabricate new class of microbiorobots," *IEEEAppl. Phys. Lett.,* vol. 97, p. 173702, 2010.

[22] P. S. S. Kim, A. Becker, Y. Ou, A. A. Julius and M. J. Kim, "Imparting magnetic dipole heterogeneity to internalized iron oxide nanoparticles for microorganism swarm control," *Nanopart.Res.,* vol. 17, no. 3, p. 144, 2015.

[23] I. Tanev, M. Georgiev and K. Shimohara, "Evolving a Team of Asymmetric Predator Agents That Do Not Compute in Predator-Prey Pursuit Problem," in *AIMSA*, Varna, Bulgaria, 2018.

[24] M. Georgiev, I. Tanev and K. Shimohara, "Coevolving behavior and morphology of simple agents that model small-scale robots," in *GECCO (Companion)*, Kyoto, Japan, 2018.

[25] M. Rubenstein, A. Cabrera, J. Werfel, G. Habibi, J. McLurkin and R. Nagpal, "Collective transport of complex objects by simple robots: Theory and experiments," in *International Conference on Autonomous Agents and Multiagent Systems*, St. Paul, MN, USA, 2013.

[26] J. Lien, S. Rodriguez, J. Malric and N. Amato, "Shepherding Behaviors with Multiple Shepherds," in *IEEE ICRA*, Barcelona, SpainIEE, 2005.

[27] A. Requicha, "Nanorobots, NEMS, and Nanoassembly," in *IEEE*, 2013.

[28] R. Niu, D. Botin, J. Weber, A. Reinmüller and T. Palberg, "Assembly and Speed in Ion-Exchange-Based Modular Phoretic Microswimmers," in *Langmuir*, 2017.

[29] M. Ibele, T. Mallouk and A. Sen, "Schooling Behavior of Light-Powered Autonomous Micromotors in Water," in *Angewandte Chemie Int. Ed.*, 2009.

[30] F. Martinez-Pedrero, H. Massana-Cid and P. Tierno, "Assembly and Transport of Microscopic Cargos via Reconfigurable Photoactivated Magnetic Microdockers," 2017.

[31] D. Brown, R. Turner, O. Hennigh and S. Loscalzo, "Discovery and exploration of novel swarm behaviors given limited robot capabilities.," in *13th International Symposium on Distributed Autonomous Robotic Systems*, London, UK, 2016.

[32] J. Holland, "Adaptation in Natural and Artificial Systems," Ann Arbor, MI, USA, 1975.

[33] E. Goldberg, "Genetic Algorithms in Search, Optimization and Machine Learning," Boston,MA, USA, 1989.

[34] S. Nolfi and D. Floreano, "Evolutionary Robotics: The Biology, Intelligence, and Technology of Selforganizing Machines," Cambridge, MA, USA, 2000.

[35] M. Georgiev, I. Tanev and K. Shimohara, "Performance Analysis and Comparison on Heterogeneous and Homogeneous Multi-Agent Societies in Correlation to Their Average Capabilities," in *SICE*, Nara, Japan, 2018.

[36] M. Georgiev, I. Tanev and K. Shimohara, "Exploration of the effect of uncertainty in homogeneous and heterogeneous multi-agent societies with regard to their average characteristics," in *GECCO (Companion)*, Kyoto, Japan, 2018.

[37] M. Georgiev, I. Tanev, K. Shimohara and T. Ray, "Evolution, Robustness and Generality of a Team of Simple Agents with Asymmetric Morphology in Predator-Prey Pursuit Problem," *Information,* vol. 10, no. 2, 2019.

[38] P. J. Angeline and K. Kinnear, "Genetic Programming and Emergent Intelligence," in *Advances in Genetic Programming*, Cambridge, MA, USA, 1994.

[39] M. Mitchell, "An Introduction to Genetic Algorithms," Cambridge, MA, USA, 1998.

Publications

Journal articles

- M. Georgiev, I. Tanev, K. Shimohara and T. Ray, "Evolution, Robustness and Generality of a Team of Simple Agents with Asymmetric Morphology in Predator-Prey Pursuit Problem," in Information, vol. 10, no. 2, 2019
- M. Georgiev, I. Tanev, K. Shimohara "Coevolution of the Asymmetric Morphology and the Behaviour of Simple Predator Agents in Predator-Prey Pursuit Problem" in Computational Intelligence and Neuroscience, 2019

Conference papers

- M. Georgiev, I. Tanev and K. Shimohara, "Coevolving behavior and morphology of simple agents that model small-scale robots," in *GECCO (Companion)*, Kyoto, 2018.
- M. Georgiev, I. Tanev and K. Shimohara, "Exploration of the effect of uncertainty in homogeneous and heterogeneous multi-agent societies with regard to their average characteristics," in *GECCO (Companion)*, Kyoto, 2018.
- M. Georgiev, I. Tanev and K. Shimohara, "Performance Analysis and Comparison on Heterogeneous and Homogeneous Multi-Agent Societies in Correlation to Their Average Capabilities," in *SICE*, Nara, Japan, 2018.
- I. Tanev, M. Georgiev and K. Shimohara, "Evolving a Team of Asymmetric Predator Agents That Do Not Compute in Predator-Prey Pursuit Problem," in *AIMSA*, Varna, Bulgaria, 2018.