On different communication modalities between foraging robot swarms

Thesis presented by Roman MILETITCH
in fulfilment of the requirements of the
PhD Degree in Engineering Sciences and Technology
("Docteur en Sciences de l'ingénieur et technologie")
Année académique 2022-2023

Supervisor: Prof. Marco DORIGO
Co-supervisor: Dr. Vito Trianni
IRIDIA - Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle

On different communication modalities between foraging robot swarms

Roman Miletitch

IRIDIA, Université libre de Bruxelles, Belgium.

2023

Composition of the jury

Prof. Hugues BERSINI
Full Professor of Computer Science, IRIDIA, École polytechnique de Bruxelles, Université libre de Bruxelles, Belgium.

Prof. Mauro BIRATTARI
Research Director of the F.R.S.-FNRS at IRIDIA, École polytechnique de Bruxelles, Université libre de Bruxelles, Belgium.

Prof. Marco DORIGO (supervisor)
Research Director of the F.R.S.-FNRS and co-director of IRIDIA, École polytechnique de Bruxelles, Université libre de Bruxelles, Belgium.

Prof. Eliseo FERRANTE
Assistant Professor at Vrije Universiteit, Amsterdam, the Netherlands.

Prof. Yara KHALUF
Assistant Professor at Wageningen University and Research, Wageningen, the Netherlands.

Prof. Vito TRIANNI (co-supervisor)
Senior Researcher at ISTC, Italian National Research Council, Rome, Italy.

Abstract

Coordination in a group relies heavily on the type and quality of interactions and communication among individuals. In swarm robotics, communication can make the difference between a heap of isolated robots working independently of each other, and a connected swarm displaying self-organisation. Communication between robots in a swarm can be *indirect*, for instance through stigmergy whereby robots exploit the sign of previous actions to coordinate, or *direct*, by means of messages exchanged among robots for the purpose of influencing each other's behaviour. In the latter case, messages can consist either of simple signals, or more structured information, possibly encoding some concept representing features of the environment or of the desired coordination outcome. More complex communication can support more complex self-organising behaviors, deeply impacting on how the swarm tackles the task at hand. In this work, we consider different ways of exploiting communication in the context of a foraging task, in which robots search an open environment for resources to be exploited. Foraging requires abilities such as navigation, exploration and collective decision making. Coordination within a foraging context can lead to higher efficiency in exploiting resources, both in the short or in the long run, by avoiding over-exploitation. Throughout my thesis, foraging is used as a means to study the coupling between different communication processes and the undertaking of a meaningful task by the robots. Specifically, we study three different uses of communication during foraging.

Firstly, we focus on simple aggregation of information, and study three parameter-free information processing mechanisms. These result in varying behavior, from the selection of a single resource by the whole swarm to the robots splitting among the resources present. This study is supported by an extensive analysis of navigation and congestion, helping to explain how swarm density can affect the perceived quality of a resource.

Next, we consider an exchange of more complex signals, inspired by the honeybee value-sensitive decision making abilities. This results in a fine-grained load-balancing between resources, suitable for an adaptive exploitation of sources at the collective level, without requiring individuals to compare the profitability of different sources or a central planner with global knowledge of the environmental conditions.

Last, we tackle the case of robots *talking* about the resources, that is, assigning names to resources following dynamics typical of language evolution. Such a

process initially leads to a temporary segregation of vocabulary, closely tied to the swarm's topology. However, over time, the swarm converged to a comprehensive and accurate description of its surroundings, encompassing all relevant resources. The emergent naming conventions facilitated effective coordination and decision-making within the swarm, highlighting the potential of language dynamics in enhancing collective behavior in complex environments.

Declaration of Authorship

This thesis presents an original work that has never been submitted to the Université libre de Bruxelles or any other institution for the award of a doctoral degree. Some parts of this thesis are based on several peer-reviewed articles that the author, together with his supervisor, co-supervisor and other collaborators, has published in the scientific literature.

Parts of the introduction (Chapter 1), state of the art (Chapters 2) and conclusion (Chapter 7) are based on:

- N. Cambier, R. Miletitch, V. Fr éemont, M. Dorigo, E. Ferrante, and V. Trianni. Language evolution in swarm robotics: A perspective. *Frontiers in Robotics and AI*, 7:12, 2020

The analysis of different information aggregation protocols for social odometry (Chapter 4) are based on:

- R. Miletitch, V. Trianni, A. Campo, and M. Dorigo. Information aggregation mechanisms in social odometry. In *Proceedings of the 20th European Conference on Artificial Life (ECAL 2013)*, pages 102–109. MIT Press, Cambridge, MA, 2013b

- R. Miletitch, M. Dorigo, and V. Trianni. Social dynamics for an exploration and exploitation task in swarm robotics. *Rapport d'avancement des recherches*, 2013a.

The analysis of the algorithm inspired by the nest site selection behaviour (NSS) of honeybees (Chapter 5) is based on:

- R. Miletitch, M. Dorigo, and V. Trianni. Balancing exploitation of renewable resources by a robot swarm. *Swarm Intelligence*, 12(4):307–326, 2018

Lastly, the analysis of the combination of NSS and language games -Minimal Language Game and Category Game- (Chapter 6) is based on:

- R. Miletitch, A. Reina, M. Dorigo, and V. Trianni. Emergent naming conventions in a foraging robot swarm. *Swarm Intelligence*, 16(3):211–232, 2022

Acknowledgements

I want first to thank everyone at IRIDIA, past and present. I felt very lucky to have spent time in that laboratory and to be able to share that space with brilliant people who made for very interesting conversations and a warm environment.

A huge thanks to Vito Trianni and Marco Dorigo for their patience in what has been quite a lengthy journey. In particular, thank you Vito for the insight you shared with me. I still remember vividly one of our first conversations back at IRIDIA that really helped me understand better the kind of research I wanted to do.

Thanks Gaëtan for existing, couldn't have asked for a better office mate. This made tolerable some long days; I'm so glad I can call you a friend.

Thanks to all my friends that supported me over the years, Daniel, Antoun, Mathilde, Raphaël, Charles... the list would be too long, and I'd probably forget quite a few, so we will leave it at that.

Last and clearly not least, thanks to my sister, who has helped me grow in life in more ways than one. I'm lucky that you share with me the way you see the world.

Contents

Abstract **5**

Declaration of Authorship **7**

Acknowledgements **9**

1 Introduction **13**
1.1 Foraging 15
1.2 Simple to structured communication 17
1.3 Main Contributions 18
1.3.1 Additional contributions 20
1.4 Structure of the thesis 21

2 State of the art **25**
2.1 Collective Navigation 25
2.2 Foraging and exploitation of resources 27
2.3 Collective decision making 28
2.4 Communication systems 30
2.5 Language games 32
2.6 Coupling between task and language 34

3 Tools **37**
3.1 The ARGoS Simulator 37
3.2 Robots 38
3.3 Navigation 40
3.3.1 Goal Vector 40
3.3.2 Targeted Object Approach 41
3.3.3 Avoidance of robots and objects 41
3.3.4 Aggregation of force vectors 42

4 Information aggregation mechanisms in social odometry **43**
4.1 Social odometry 44
4.1.1 Information Sharing 44
4.1.2 Information Processing 45
4.2 Navigation Task 47

4.2.1 Individual Behaviour . . . 47
4.2.2 Experiments . . . 48
4.2.3 Results . . . 50
4.3 Exploitation Task . . . 58
4.3.1 Individual Behaviour . . . 58
4.3.2 Experiments . . . 60
4.3.3 Results . . . 62
4.4 Conclusions . . . 69

5 Balancing exploitation of renewable sources 71
5.1 Experimental setup . . . 72
5.1.1 Individual and Collective Behaviour . . . 72
5.2 Baseline exploration efficiency . . . 77
5.3 Baseline efficiency in source exploitation . . . 78
5.4 Exploration vs. exploitation of a single source . . . 81
5.5 Balancing source exploitation . . . 83
5.6 Discussion and conclusions . . . 87

6 Emergent naming conventions in a foraging robot swarm 89
6.1 Language games in a foraging robot swarm . . . 90
6.2 Experimental setup . . . 91
6.2.1 Individual and collective behaviour . . . 92
6.3 Correctness and completeness of vocabulary . . . 95
6.4 A study of the swarm's spatial characteristics . . . 99
6.4.1 Impact of spatial word creation . . . 100
6.4.2 Communication topology and interactions within the swarm 101
6.5 Emergence of spatial categories . . . 104
6.5.1 Experimental SetupImplementation of the Category Game 104
6.5.2 Results . . . 106
6.6 Conclusion . . . 107

7 Conclusions 111
7.1 A perspective for language evolution in swarm robotics . . . 112
7.2 Future work . . . 114

Chapter 1

Introduction

In multi-robot systems, collaborative efforts towards a common objective often yield enhanced efficiency compared to individual agents operating independently. The group performance arises not merely from the individual robot behaviors but from the intricate interactions and coordination that occur among them, facilitating the accomplishment of the shared goal. Swarm robotics, as a research domain, delves into the design of multi-robot systems with an emphasis on coordination and communication between relatively simple robots, fostering the emergence of complex collective behaviors (Dorigo et al., 2021).

A crucial aspect of designing robot swarms is the principle of self-organization, which stems from the numerous local interactions occurring among robots and between robots and their environment (Şahin, 2004; Brambilla et al., 2013). These local interactions are frequently designed with inspiration from natural systems, drawing upon the intricacies of animal societies as guiding principles. Specifically, social insects, such as ants, bees, and termites, have proven to be a valid source of inspiration due to their ability to perform complex tasks through decentralized, self-organized mechanisms. By emulating the cooperative strategies and communication methods observed in these insect societies, swarm robotics harnesses their inherent robustness, adaptability, and scalability. This biomimetic approach facilitates the development of efficient and resilient multi-robot systems capable of addressing a myriad of real-world challenges, including search and rescue missions, environmental monitoring, and precision agriculture (Dorigo et al., 2020).

In this context, communication plays a crucial role in enabling individual robots to effectively collaborate and exhibit sophisticated collective behaviors. By implementing well-designed communication protocols, robots can exchange vital information, synchronize their actions, and cooperatively adapt to dynamically changing environments. Swarm robotics research predominantly identifies three interaction modalities: indirect communication, direct interactions, and direct communication (Trianni and Dorigo, 2006). Indirect communication, commonly observed in insect societies, is often referred to as stigmergy. It involves communication through environmental modifications, such as ants depositing pheromones during foraging activities (Deneubourg et al., 1990; Garnier et al., 2007). Direct interactions entail physical contact-based influence between individuals, eliciting a response in the receiving party (e.g., pulling/pushing forces during collective transport, see Kube and Bonabeau, 2000). Lastly, direct communication entails the real-time exchange of information between individuals without physical contact and is the most prevalent interaction modality in swarm robotics.

The choice of communication modality in swarm robotics depends on the task and environment the swarm operates in. Direct communication has been employed as a foundational mechanism in the implementation of swarm behavior, helping with the tackling of various tasks. For instance, self-organized aggregation (Soysal and Sahin, 2005; Dorigo et al., 2004; Cambier et al., 2021) enables robots to autonomously form clusters, which can enhance the system's fault tolerance and simplify tasks that require group effort; to achieve aggregation, robots communicate to signal their presence and favour the formation of a large cluster. Morphogenesis (O'Grady et al., 2009; Rubenstein et al., 2014; Slavkov et al., 2018) enables robots to generate complex shapes and structures, supporting the creation of adaptive and reconfigurable systems; to achieve morphogenesis, robots communicate by exchanging messages about when and how to expand the forming structure. Foraging (Ducatelle et al., 2011c; Talamali et al., 2020) focuses on the search, collection, and transportation of resources; to achieve foraging, robots communicate to share the location and quality of resources. Flocking (Baldassarre et al., 2006; Çelikkanat et al., 2009; Ferrante et al., 2014) enables robots to coordinate their movements as a cohesive unit, so to maintain group integrity while traveling; to achieve flocking, robots communicate their preferred motion direction and velocity. Furthermore, direct communication has exhibited its versatility by replicating other interaction modalities in swarm robotics. One notable example is the emulation of pheromone-based communication through the use of robot chains (Ducatelle et al., 2011b; Nouyan et al., 2009; Campo et al., 2010; Ferrante et al., 2013).

Notwithstanding the recognised relevance of communication in swarm robotics, the interplay between the behaviour dynamics and the communication protocols—i.e., the rules and conventions that determine the transmission of information—has not been extensively studied. As a matter of fact, often the communication protocols are *a priori* defined and the swarm behaviour is then designed to optimise a given performance metric. Mostly, the effects of the communication system on the swarm performance are studied with respect to limitations in range (how far each robot can exchange messages) and bandwidth (how much information each robot can exchange). The communication protocols themselves are rarely the object of study. The overall goal of this thesis is to shed light on the mutual influence among communication protocols and collective behaviours. We chose *foraging* as the reference collective behaviour, as it represents relevant tasks for many application domains (see also Section 1.1).

The contribution of this thesis is the exploration of how communication methods affect and are influenced by foraging tasks. The different studies that compose this thesis contribute to reveal the impact of different communication protocols on swarm efficiency and decision-making, and how these adapt in dynamic foraging scenarios. Additionally, the role of embodiment and situatedness is investigated in relation to language dynamics, offering new insights into the intricate relationship between task execution and communication in robotic swarms. In particular, on the one hand, this thesis demonstrates how different communication protocols are sufficient to radically impact the complex dynamics of a foraging behaviour, whereby robots navigate an open environment to locate and utilize resources. More specifically, the study focuses on the effects of the communication protocols on the swarm topology and the collective performance. On the other hand, the thesis explores how complex forms of communication can emerge from the task execution dynamics. Specifically, the communication system is made adaptive by the exploitation of language games, which lead to the emerge of words that are relevant to the foraging task. The thesis studies how the emergent communication system can be considered a tool to describe how the swarm experiences the environment. In other words, the thesis provides a demonstration of how the collective knowledge emerging from robot interactions can be exploited to accurately represent the working environment, hence pointing to a novel way of exploiting communication for more adaptive and flexible collective behaviours.

In the remainder of this introductory chapter, we first introduce more details about foraging as a suitable playground for studying the effects of communication in swarm robotics, focusing on aspects such as navigation, exploration, and decision-making (see Section 1.1). Then, we touch upon the varying degrees of communication complexity, ranging from basic to intricate forms, and discuss their relationship with task performance, especially information aggregation mechanisms (see Section 1.2). Subsequently, we outline the principal contributions of this thesis in Section 1.3. Lastly, we provide a detailed description of the thesis content and organization in Section 1.4.

1.1 Foraging

In this thesis, foraging represents a prototypical example to investigate the interplay between various communication processes and the execution of a meaningful task by the robots, ultimately examining how this relationship contributes to accurate decision-making during task execution. This investigation will focus on two main objectives: (i) to evaluate the efficiency of a swarm during foraging and its ability to make decisions based on environmental conditions while utilizing various communication protocols, and (ii) to analyze how the swarm's engagement in a task—particularly in terms of emergent network topology—influences language game dynamics. By addressing these aims, we strive to deepen our understanding of the intricate relationship between swarm behavior and communication, ultimately contributing to new paradigms for the design of swarm robotics systems.

Foraging refers to the process of searching for *resources* within an unknown environment. Resources in foraging are typically generated or made accessible in specific regions, known as *sources*. Foraging theory is often examined in terms of optimizing the payoff that results from a foraging decision. In many behavioral ecology models, this payoff is represented as the amount of energy an organism acquires per unit time, minus the cost of foraging. However, in swarm robotics, the cost is frequently disregarded, with emphasis placed on the efficiency of resource retrieval from sources distributed throughout an arena. Foraging tasks can involve abstract sources (e.g., zones to reach) or tangible resources (i.e., items to collect), with various spatial distributions (e.g., uniformly throughout the environment or clustered). These experimental setups serve as analogs for tasks necessitating exploration and exploitation, such as mining, cleaning, or search and rescue operations.

While the primary objective of foraging is typically to maximize exploitation, other goals may emerge depending on the specific application. For instance, in hazardous environments, a swarm may prioritize exploiting a single source to prevent excessive and unsafe dispersal of individual agents. Conversely, if congestion is a concern, swarm agents may prefer to distribute among multiple sources to minimize interference and achieve a balanced exploitation of sources. In search and rescue scenarios, the focus shifts to enhancing exploration, enabling the swarm to comprehensively locate points of interest, such as victims in need of rescue.

At the swarm level, foraging necessitates capabilities including navigation, exploration, and collective decision-making to effectively exploit the available sources as a group. The most rudimentary approach to navigating and exploring within a confined area is the random walk. This method involves robots moving in random directions with varying step lengths, allowing for unbiased exploration without any prior knowledge of the environment. Though not highly efficient, this method guarantees that robots will eventually reach all parts of the environment, albeit potentially over an extended duration due to repeated visits to already explored areas. To improve upon purely random exploration, robots can employ memory and mapping techniques to avoid revisiting previously explored regions (Thrun, 2008) and to target specific areas of interest. Odometry is one of these techniques, aiming to estimate a robot's position and orientation based on its sensors and actuators, typically involving measurements of wheel rotations or accelerometers, which provide information on the robot's movement relative to a known starting point. However, the integration of such navigational information over time is inherently susceptible to errors. As robots rely on the accumulated sensor data to determine their position, any inaccuracies or noise in the measurements can cause a gradual drift in their estimated location. This drift, compounded over time, can lead to significant discrepancies between the robots' perceived position and their actual location within the environment, ultimately affecting their ability to efficiently navigate and explore the designated area. Social odometry (Gutiérrez et al., 2010) rectifies these inaccuracies by enabling robots to share their positional data and aggregate it with their own, ensuring enhanced precision in self-localization and navigation towards designated areas.

In scenarios featuring multiple sources within the environment, a collective decision regarding their exploitation is necessary for the swarm to operate cohesively. Focusing on the exploitation of a single source may be advantageous under certain conditions, such as when a sufficient number of robots must be aggregated to support collective localization or when exploitation requires multiple robots at the source. Alternatively, excessive robot density can lead to congestion, and partitioning the swarm can be beneficial depending on the environment's layout.

1.2 From simple communication to structured language

As previously discussed, communication is a cornerstone of swarm robotics, enabling coordination of activities. Throughout this thesis, we explore a range of communication methods, from simpler forms to more complex ones, and explore how the resulting decision-making processes influence both the swarm's performance in foraging tasks and its ability to interpret the environment it navigates. Communication methods can be grouped within three broad categories: indirect communication (stigmergy), direct interactions (such as pulling/pushing forces between robots), and direct communication (messages sent and received).

Collective behaviours in robot swarms typically rely on straightforward communication processes, which consist of simple and direct exchanges of a limited amount of information. This is because swarm robots are often assumed to be simple and limited in their computational and communication abilities. Hence, basic communication mechanisms often utilize pre-defined signals or messages to facilitate cooperation among swarm agents. While such an approach is effective in certain situations, it offers limited adaptability to variations in the task or alterations in the working environment, posing a constraint on the swarm's autonomy.

One approach to expand communication capabilities in swarm robotics is to examine the communication strategies found in biological systems. For example, animal societies often exhibit complex communication systems that involve various modalities, such as visual, auditory, and chemical cues. By studying and emulating these systems, researchers can develop innovative communication methods for swarm robots.

Similarly, neural networks can play a significant role in enhancing communication within robot swarms. By employing artificial neural networks, robots in a swarm can learn and adapt their communication strategies based on the data they receive from their environment and interactions with other robots. Neural networks can be trained to recognize patterns and associations, allowing swarm robots to generate more complex and context-dependent communication signals.

Another approach relies on behavior modules, as a means to enriching communication systems in swarm robotics. These modules are predefined sets of behaviors that can be combined and activated to suit a specific situation or task. By implementing behavior modules, swarm robots can adapt their communication strategies based on the current context, allowing for more efficient cooperation and coordination.

In the cases presented above, the rules of communication are often designed specifically for the task at hand. In order to enhance the adaptability and performance of swarm robotics, it is crucial to explore alternative communication systems that go beyond basic predefined signals or messages. To address the inherent limitation of a local sensory system, while still aligning with the principles of swarm robotics, adaptability can be enhanced by incorporating self-organization into the communication process layer. Developing more sophisticated communication capabilities in swarm robotics would allow swarms to engage in advanced negotiation dynamics, particularly in changing environments where conditions may significantly vary over time. This would help swarms better address the challenges posed by changing environments or complex tasks, ultimately improving their autonomy and cooperative abilities.

One potential approach to achieving this objective draws inspiration from the field of linguistics, specifically through the concept of *language games*. These games are designed for agents or robots to engage in turn-based interactions, simulating real-world dynamics that contribute to the emergence of structured language. Language games serve as minimal algorithms that exhibit key characteristics of comprehensive languages while also demonstrating their adaptability.One potential approach to achieving this objective draws inspiration from the field of linguistics, specifically through the concept of *language games*. These games are designed for agents or robots to engage in turn-based interactions, simulating real-world communication dynamics that contribute to the emergence of structured languages. They are shaped by cultural evolution principles, reflecting the collective and local interactions seen in natural language development. Language games serve as minimal algorithms that exhibit key characteristics of comprehensive languages while also demonstrating their adaptability.

Through language games, robots can develop and adapt their communication strategies, enhancing their coordination and efficiency in task execution, particularly in dynamic foraging scenarios. Language games thus play a crucial role in demonstrating the capability of robot swarms to evolve complex, context-specific communication methods.

1.3 Main Contributions

In this thesis, we aim to examine the relationship between a specific task (foraging) and various communication methods (all falling within the broad category of direct communication methods) used by robot swarms. In particular, we focus on how task and communication influence each other.

First, this thesis aims to understand how different ways of communication among robots in a swarm affect their ability to work together while searching for and gathering resources. In particular, we will look at how different communication protocols impact the swarm's ability to forage efficiently, the way the swarm's topology changes over time, and how robots make decisions and use available resources.

Additionally, the thesis explores how the execution of a specific task impacts the way in which robots exchange information. In this context, communication protocols are not examined in isolation but in relation to the swarm's foraging dynamics. This allows us to study communication protocols over complex and dynamic topologies resulting from the swarm's foraging. On top of that, when language games are used to create words, we observe how these are anchored in their environment, linked with the task at hand. This addresses the question of how language can be used as a tool for describing the environment in relation to a given task.

The main contributions of this thesis are highlighted in the following.

- The initial communication protocol we study is social odometry, by which robots share with each others information based on their proprioception to correct odometry errors (Gutiérrez et al., 2009). We introduce three novel mechanisms that do not require any specific parameter and have varying levels of tolerance for conflicting information (Miletitch et al., 2013b).

  We analyzed the impact of each mechanism on the swarm's resource exploitation, environmental navigation, and resulting topology and its link with communication between robots. Furthermore, we tested these approaches in more intricate scenarios that involve object manipulation (Miletitch et al., 2013a). The obtained results, particularly those related to social odometry and congestion, serve as a foundation for my following research. This work is discussed in Chapter 4.

- Building on top of the previous study, we studied how an algorithm inspired by the nest site selection behaviour (NSS) of honeybees can best exploit communication about the areas of interest to support foraging (Miletitch et al., 2018). Unlike the information aggregation mechanisms exploited within social odometry, as mentioned above, which are inflexible with respect to changes in the environmental conditions, NSS enables sustainable foraging in a multiple resource context and is flexible to varying contingencies. This contribution is developed in Chapter 5.

- Lastly, we introduce two language games, namely the Minimal Language Game and the Category Game (Miletitch et al., 2022), performed by the robots on top of the NSS algorithm introduced in the above research. We present how varying and evolving topologies influence the way language games are played, and how this affects the convergence of the vocabulary within the swarm and its subgroups. Additionally, we study how language games can be used to provide a correct and complete description of the swarm's environment. This study is presented in Chapter 6.

- In addition to the research on social odometry and NSS algorithms, we propose a framework for advanced communication in swarm robotics (Cambier et al., 2020) with a focus on language games and more complex cognitive representations and behaviors. This framework highlights the potential for interdisciplinary research using swarm robotics as a test bed, particularly in the fields of linguistics and sociology. Ultimately, this could lead to the development of more sophisticated swarm robotics systems that can adapt to complex environments and tasks. This framework is discussed in the conclusions, with some preliminary results as examples.

Relevant publications:

- R. Miletitch, V. Trianni, A. Campo, and M. Dorigo. Information aggregation mechanisms in social odometry. In *Proceedings of the 20th European Conference on Artificial Life (ECAL 2013)*, pages 102–109. MIT Press, Cambridge, MA, 2013b

- R. Miletitch, M. Dorigo, and V. Trianni. Balancing exploitation of renewable resources by a robot swarm. *Swarm Intelligence*, 12(4):307–326, 2018

- N. Cambier, R. Miletitch, V. Fr éemont, M. Dorigo, E. Ferrante, and V. Trianni. Language evolution in swarm robotics: A perspective. *Frontiers in Robotics and AI*, 7:12, 2020

- R. Miletitch, A. Reina, M. Dorigo, and V. Trianni. Emergent naming conventions in a foraging robot swarm. *Swarm Intelligence*, 16(3):211–232, 2022

1.3.1 Additional contributions

Contributions to additional studies were provided during the timeframe of the PhD studies. I participated in the research aimed at the creation of a novel tool to design automatically the control software for robot swarms (Francesca et al., 2014, 2015). This tool, named AutoMoDe-Vanilla and later developed in another version called AutoMoDe-Chocolate, combines simple behavioral blocks and optimizes the resulting behavior in order to create a control software for a robot swarm. The aim of the studies I participated in was to test this tool against behaviors designed by human experts of the field under a constraint of time. I was one of the programmers recruited for the study. Results revealed that in both articles, the automatic tool outperformed on average the human expert designers.

Additionally, I participated to a research study about quantifying the link between the local scale and the global one in the context of foraging and decision making (Reina et al., 2015a). In this context, I provided code for simulated robots and contributed to develop the overall swarm behaviour that have led to the experimental validation.

Finally, I participated in a study that explored the cultural evolution of language in swarm robotics, specifically examining the human self-domestication (HSD) hypothesis (Cambier et al., 2022). This hypothesis posits that the evolution of modern languages may be partially attributable to the self-domestication of the human species. To investigate this hypothesis and the process of language evolution in general, we created an embodied model using swarm robots, simulating the effects of prosociality on language formation. Our model featured robots in multiple nests engaged in a foraging task while playing a naming game, with novel features such as robot individuation and parametrizable prosociality. The results demonstrated the formation of an "in-group bias", an increased efficiency in resource collection with higher prosociality values, and the modulation of the effect of physical distance on lexical convergence by prosociality.

Relevant publications:

- G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pinciroli, et al. An experiment in automatic design of robot swarms: Automode-vanilla, evostick, and human experts. In *Swarm Intelligence: 9th International Conference, ANTS 2014, Brussels, Belgium, September 10-12, 2014. Proceedings 9*, pages 25–37. Springer, 2014

- Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pinciroli, et al. Automode-chocolate: automatic design of control software for robot swarms. *Swarm Intelligence*, 9 (2-3):125–152, 2015

- A. Reina, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro-macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intelligence*, 9(2-3):75–102, 2015a

- N. Cambier, R. Miletitch, A. B. Burraco, and L. Raviv. Prosociality in swarm robotics: A model to study self-domestication and language evolution. In *Joint Conference on Language Evolution (JCoLE)*, pages 98–100. Joint Conference on Language Evolution (JCoLE), 2022

1.4 Structure of the thesis

This thesis is organized into seven chapters, including this introduction. Two chapters are dedicated to background knowledge, followed by three chapters that discuss specific experimental studies in detail. Chapter 7 concludes the thesis and outlines related works that build upon it. In the following, we provide a more detailed overview of the content of each chapter.This thesis is organized into seven chapters, including this introduction. The next two chapters are dedicated to background knowledge, followed by three chapters that discuss specific experimental studies in detail. Chapter 4 lays the groundwork by exploring information aggregation mechanisms in social odometry, a foundational concept for the subsequent chapters. From then, we consider that positioning is a solved task. Chapter 5 builds up on this work, as it proposes a more refined load-balancing exploitation then the one presented in the previous chapter, setting up the groundwork for Chapter 6. Chapter 6 keeps the same exploitation behavior, and introduces an additional communication layer through language games. This advancement is key for moving towards a linguistic representation of the entire space and potentially enables more intricate information exchanges. This progression illustrates an increasing complexity in communication techniques. Each chapter builds upon the previous one, progressively enhancing the integration of task execution and communication within the swarm. Chapter 7 concludes the thesis and outlines related works that build upon it. In the following, we provide a more detailed overview of the content of each chapter.

Chapter 2 — State of the art Following this introduction, we present a review of the topics addressed in this thesis: foraging and exploitation dynamics, navigation and exploration, collective decision making, and language games in the context of swarm robotics.

Chapter 3 — Tools All experiments in this thesis are conducted using the ARGoS simulator with either marXbot or e-puck robots. In this chapter, we introduce these tools and describe the basic algorithms employed for creating the robots' behaviors, such as random walk exploration, collision avoidance, and object grabbing. The chapter also discusses the rationale behind the choice of these tools and algorithms and highlights their advantages for the experiments conducted.

Chapter 4 — Information Aggregation Mechanisms in Social Odometry In this chapter, we examine three distinct information aggregation mechanisms applied to social odometry. The focus is on understanding their impact on the swarm decision process concerning the choice of splitting or converging on a single resource. This analysis provides valuable insights into optimizing swarm performance, highlighting the importance of effective information aggregation mechanisms for adapting to diverse foraging situations.

Chapter 5 — Balancing Exploitation of Renewable Sources Building upon the insights from Chapter 4, we implement an algorithm inspired by the nest site selection abilities of honeybees. This algorithm involves exchanging more complex signals among robots, introducing recruitment and cross-inhibition feedbacks. The result is a more refined load-balancing approach between resources, which allows for better adaptation to changes in the environment and in task requirements.

Chapter 6 — Emergent Naming Conventions in a Foraging Robot Swarm Lastly, on top of to the load-balanced foraging, we introduce a language game played by the robots in which they assign names to resources and reach consensus on these names. This language game allows the swarm to converge on an accurate representation of its environment, enhancing its ability to make informed decisions. This chapter emphasizes the importance of communication in swarm robotics and demonstrates its potential for improving the swarm's performance and overall knowledge.

Chapter 7 — Conclusion and Future Perspectives In the concluding chapter, we synthesize the principal findings from preceding chapters and emphasize the potential of integrating language games with swarm robotics tasks. This includes exploring various language games and more advanced linguistic/social frameworks, and drawing inspiration from mammalian and primate social systems (such as the self-domestication hypothesis) for more intricate language dynamics.

Chapter 2

State of the art

Swarm robotics is a growing field within robotics and artificial intelligence that explores the design and control of decentralised multi-robot systems capable of self-organisation. This chapter highlights the state of the art in the field, limited to aspects relevant for the work presented in this dissertation. We start with **collective navigation** (see Section 2.1), stating the challenges in developing effective navigation strategies to enable exploration of an unknown environment. In Section 2.2, the focus shifts to **foraging and exploitation of sources**, addressing the study of how robots find valuable resources, often drawing inspiration from nature. Section 2.3 addresses **collective decision making** and examines how individual decisions by robots can lead to meaningful group actions. **Communication systems** are considered in Section 2.4, which reviews the different ways robots can communicate with each other, including both simple and more complex methods. The chapter then moves on to the discussion of **language games** (see Section 2.5), an area that explores how artificial systems can mimic the evolution of natural languages. Finally, in Section 2.6, we discuss the **coupling between task and language**, that is, how the connection between language and behavior can become an essential aspect to consider when designing a swarm robotics system.

2.1 Collective Navigation

In swarm robotics, computational and sensorimotor limitations do not allow the exploitation of advanced self-localisation and navigation strategies, e.g., based on SLAM approaches (Thrun, 2008). Instead, random walks are often exploited (Dimidov et al., 2016). Such a stochastic exploration of the environment can be inefficient, as robots tend to re-visit previously visited areas multiple times.

Several studies in swarm robotics implement navigation and exploration algorithms without any sharing of structured information, sometimes exploiting robots as physical landmarks. Rekleitis et al. (2001) divided the swarm into two teams, one moving and the other stationary, serving as a reference for navigation.

The teams alternate between stationary and moving states. Nouyan et al. (2008, 2009) exploit robots to form complex structures such as chains, in which one end of the chain connects to a central place while the other end explores the environment. Once the goal location is reached, the chain can be exploited by other robots for navigation purposes, or a bucket brigade method can be used to transport objects along the chain (Ostergaard et al., 2001).

On the other hand, there are various ways to improve navigation through information-sharing within a swarm (Martinelli et al., 2005). Ducatelle et al. (2011a) model a swarm as a communication network that propagates relevant information. Each robot in the swarm maintains a table with navigation information about all known robots, similar to how nodes in a mobile ad hoc network maintain routing tables. Then, the robots propagate the available information and use the table to find the best path to reach a target robot within the swarm. Sperati et al. (2011) also study navigation in a swarm robotics context. In this case, communication is performed through visual signals only and therefore the information exchanged is much less structured. For this reason, they used artificial evolution to synthesize effective navigation strategies.

Alternatively, information-sharing in groups of robots can be used as a means to collectively reduce the overall odometric error, for instance by sharing the estimated position of specific landmarks (Martinelli et al., 2005), or directly sharing the position of target areas. This is a straightforward mechanism that easily lends itself to implementation on very simple robots, scaling well in big swarms. This mechanism was first introduced by Gutiérrez et al. (2009) and is referred to as *social odometry.*Alternatively, information-sharing in groups of robots can be used as a means to collectively reduce the overall odometric error. Some methods rely on sharing the estimated position of specific landmarks (Martinelli et al., 2005) and using Kalman filters to fuse both proprioceptive and exteroceptive sensor data to correct the robot's position information. Although Kalman filters are effective as recursive filters, they necessitate external data and are computationally demanding. Another approach referred to as *social odometry* and firstly introduced in Gutiérrez et al. (2009) relies on directly sharing the position of target areas with a simpler information aggregation process (detailed in Chapter 4). This is a straightforward mechanism that easily lends itself to implementation on very simple robots, scaling well in big swarms. In this approach, the robots estimate the navigation path between two target areas in the environment (*i.e.*, home and goal locations) using odometry and attach to this estimate a confidence level that decreases with the distance travelled. At the same time, the robots share their navigation information within the swarm in a local peer-to-peer manner. Thanks to this process, information about target areas spreads gradually within the swarm, contributing to reduce the error in the position estimation. Overall, this decentralized process results in an increased efficiency in the swarm navigation abilities.

2.2 Foraging and exploitation of resources

Exploration of the environment and exploitation of valuable resources represent problems commonly studied in robotics, and in particular in the multi-robot systems and swarm robotics sub-fields (Winfield, 2009; Ducatelle et al., 2014; Trianni and Campo, 2015). In (Winfield, 2009), robots have to retrieve items (preys) spread around the environment or in specific goal areas (resources) and to bring them back to a specific location (nest). Such exploitation patterns are often found in biological systems. Among others species, ants display complex foraging behaviours through which they are able to adapt to a dynamic environment and retrieve preys (Camazine, 2003). The abilities of a robot to move in space and to autonomously identify locations of interest make this problem particularly relevant for a number of different application scenarios, from search and rescue to mining, from precision agriculture to space exploration (Murphy et al., 2008; Cheein and Carelli, 2013; Yoshida, 2009; Trianni and Dorigo, 2005).

The exploration and resource exploitation problem has been previously approached in swarm robotics, mainly for non-renewable resources. Several studies present adaptive foraging algorithms inspired by the well-known response tresholds model (Bonabeau et al., 1996). Krieger et al. (2000) studied the effects of heterogeneities in the individual response thresholds and of additional recruitment mechanisms to adapt the size of the foraging group to the features of the available resources. Labella et al. (2006) tested an extended response threshold model with individual learning abilities in a group of robots that foraged for sparse resources, and observed the adaptation of the robot activities between foraging and idleness, linking it to hardware differences among robots. Liu et al. (2007) employed a similar adaptation mechanism to allocate workers for a foraging task, and later presented a macroscopic probabilistic model that predicts the robotic system dynamics (Liu and Winfield, 2010). An adaptive response treshold model was presented by Castello et al. (2015), tailored to fast adaptations to changing environmental conditions. In the above mentioned studies, task allocation resulted from the adaptivity of the individual behaviour, which balances the foraging rates on the basis of information collected about the resource availability.

When resources are clustered in specific areas in space, recruitment of robots to areas in which resources are likely to be found becomes important (Krieger et al., 2000). Gutiérrez et al. (2010) studied the collective behaviour of robots foraging from static resources and sharing information about the resources position, eventually leading to the exploitation of the closest one thanks to a positive feedback given by a larger number of robots promoting the closer alternative. Hecker and Moses (2015) developed a foraging algorithm based on a delicate balance between individual search and recruitment from peers, and optimised the system parameters through a genetic algorithm to fit different environmental conditions, including clustered resources. Similarly, Pitonakova et al. (2016) considered foraging of resources possibly clustered in various deposits, also taking into account dynamic conditions where the quality of the deposit abruptly changed, to evaluate the plasticity of the proposed behaviour. In similar conditions, the need to select the most profitable resource among many available can lead to collective decision-making problems (Valentini et al., 2017).

The studies mentioned above did not deal with sustainable resource exploitation, but instead optimised the foraging efficiency, either by choosing the most profitable resource or by switching to different resources when the current one gets depleted. Sustainable or continuous foraging was instead focused on the optimisation of the foraging rate, and in the maintenance of resources in lieu of their depletion (Song and Vaughan, 2013; Liemhetcharat et al., 2015). The "maximum sustainable yield model" introduced by Song and Vaughan (2013) prescribes that resources characterised by a logistic growth should be maintained at the level of maximum regeneration rate. An algorithm was proposed to allocate a slightly higher number of robots to each resource, where each robot adapted its foraging rate to maintain the resource around the optimal size for regeneration. Maximisation of the foraging rate was also studied by Liemhetcharat et al. (2015), who however employed an heterogeneous system in which some agents could have an overview of the resource exploitation level, and helped the other foraging agents to adjust their activity rate so as to maximise the system efficiency.

When robots forage from the same resource, interferences arise as congestion builds up. Rybski et al. (2007) showed in their work that the introduction of communication in real foraging experiments does not always increase the performance of the system because of an increase in interference.

2.3 Collective decision making

Collective decision-making deals with how the sum of robots' local decisions can result in a meaningful decision at the swarm level. In its simplest form, it aims to reach a common agreement between robots: a consensus, defined by having enough robots converging towards a single possibility among different alternatives. This common agreement is often linked with maximised performance of the swarm, as it focuses its resource appropriately. Consensus can be made about destinations, foraging patches, words in a vocabulary, aggregation areas, traveling paths, etc. A consensus is often difficult to reach due to the limited sensing and information sharing range of the robots. Information leading to the best possible alternative is often hard to reach and spreading that information to the rest of the swarm can be costly and lengthy.

Finding the best alternative among many is studied under the umbrella term of best-of-n problem, an abstraction capturing the structure and logic of discrete consensus achievement problems. The best-of-n problem requires a swarm of robots to make a collective decision over which option, out of n available options, offers the best alternative to satisfy the current needs of the swarm. Each option is characterized by a quality and by a cost that are function of one or more attributes of the target environment (Reid et al., 2015). In this specific definition of the problem, the aim is for the swarm to take into account cost and quality of each option and to reach a decision, defined as having a large majority of the robots favoring the same option (threshold set by the designer of the experiment).

For instance, Valentini et al. (2016) propose a collective perception scenario in which the floor of a closed arena is covered with tiles of different colors. They compare three different strategies in which the swarm aims to determine which color is the most frequent in the environment. In this experimental setup, the color is the feature, the quality is defined as the total arena surface of a particular color, and the cost is null (instant perception of the color once a robot is on it).

Montes de Oca et al. (2011) study robots that perform multiple parallel executions of collective transport in groups of three from a nest area to a goal area. Both area are connected by two paths, one longer than the other. Each robot starts with a preferred path and votes in their group of three for which path to take. Once a path is decided upon, it is updated as the preferred path for each robot of the group. As robots choosing the short path return more quickly to the nest area, they are more often joining new groups and influencing other robots. This leads to a convergence toward the shorter path.

Parker and Zhang (2011) proposed a consensus achievement behavior based on quorum sensing. The algorithm is inspired by how social insects choose the best nest over multiple alternatives (as described in (Seeley et al., 2012b)). When a robot finds a new potential nest, it evaluates its quality and advertises it by sending recruiting messages. The frequency of these messages is proportional to the perceived quality of the alternative which influences over time the overall convergence of the swarm toward the best alternative. A similar result can be achieved by keeping track of the number of encounters of robots (in order to estimate an encounter rate) as seen in Pavlic et al. (2021) where memory-less robots modify their environment in order to keep track of such information.

While voting works by having robots directly comparing information, robots can converge on a similar behavior by mimicking each other. Both mechanics are compared in Wessnitzer and Melhuish (2003) where robots try to reach two moving targets. They decide which target to reach first and then focus on the second target. Two mechanics are compared: in the first, the robots simply follow the robot closest to a target, resulting in a decision based on the spatial distribution of the swarm; in the second, the robots vote, using a majority rule, to decide which target to follow.

Methods of information gathering are particularly important as a first step toward collective decision. As such, through collective perception, sparsely distributed agents aim to form a shared global view of a spatially distributed problem with access to only locally perceived information. Such information is subject to spatial correlations and depends on the positions of the robots. This rises the question on how to share and combine the collected information among robots in order to achieve a precise global estimate efficiently.

Bartashevich and Mostaghim (2021) compare the performance of multiple common belief combination operators from evidence theory, extending from previously introduced benchmark for a collective perception scenario (Bartashevich and Mostaghim, 2019). Here, the authors consider several possible environments that differ in the type of noise experienced by individuals: from uniform noise throughout the entire environment, to highly heterogeneous levels of noise across different locations. Their results show that swarms of limited size benefit the most by employing the fusion operator of "proportional conflict redistribution" (Smarandache and Dezert, 2005).

2.4 Communication systems

As mentioned in Section 1.2, three forms of communication are often considered in swarm robotics: indirect communication (stigmergy), direct interactions (pulling/pushing forces), and direct communication. Direct communication is implementable with radio links, simple devices like infra-red transceivers or through visual signals (e.g., coloured LEDs), whereas indirect communication and direct interactions require much more specific sensors and actuators (e.g., UV-light emitters coupled with a bespoke floor material, see Alers et al. (2014), or force/torque sensors, see Groß et al. (2006)). In this thesis, we focus on direct communication, which was successfully employed in a variety of tasks, like self-organised aggregation (Soysal and Sahin, 2005), morphogenesis (O'Grady et al., 2009), foraging (Ducatelle et al., 2011c), flocking (Ferrante et al., 2014), collective exploration and navigation (Ducatelle et al., 2011b; Nouyan et al., 2009; Campo et al., 2010; Ferrante et al., 2013). In the cases presented above, the rules of communication are designed specifically for the task at hand, offering little to no adaptation toward variations of the task or changes in the working environment, which may become a strong limitation for successful deployment. Consider, for example, the case of self-assembly and coordinated motion, which have been implemented in robot swarms to negotiate obstacles such as hills or holes during exploration of complex unstructured environments (O'Grady et al., 2010). If the conditions necessary to trigger the self-assembly in a given shape are predefined by the experimenter and are specific for a few types of obstacles, no adaptation is possible when the robots have to navigate in an heterogeneous environment with obstacles of different kinds. Therefore, for robot swarms to cope with uncertain environments with unknown obstacles, the set of shapes in which they can self-assemble should be part of the robots' action and communication space. Conversely, the rules of communication should be sufficiently adaptable to induce the type of self-assembly required for said obstacles. This adaptability is even more important when the nature of the tasks itself might change over time.

To date, however, there has been limited work on providing robot swarms with adaptive forms of communication that can scale up with the task complexity. Except for a subset of studies that use artificial systems as a means to explain linguistic features, efforts in producing adaptive communication for multi-robot systems have mostly focused on automatic design (Nolfi and Mirolli, 2009), often within the framework of evolutionary robotics (Nolfi and Floreano, 2000). The main challenge in evolving a functional communication system resides in the need to concurrently determine both the signal and a suitable response to the signal. Either of the two traits, if taken individually, may be either maladaptive or neutral and can easily be selected out by the automatic design process. Hence, specific conditions must be met to observe the emergence of communication.

The situation grows in complexity when a non-trivial swarm behaviour must be synthesised. As a matter of fact, research on evolving communication that is useful for robot swarms represents only a small fraction of the available studies. In a prominent study on this topic, small colonies of robots were evolved within a particular scenario that did not especially encourage communication (Floreano et al., 2007). In this experiment, the robots were assigned a foraging task. However, the environment also hosted poison sources indistinguishable from food sources. At the end of an evolutionary optimisation process, the robot swarms equipped with a visual communication system showed significantly better performance with respect to communication-less swarms. Specifically, two types of signals emerged in different populations, whereby agents either shared the position of the food sources to attract teammates, or they signalled the poison sources to repel other robots.

Other works used evolutionary robotics to evolve signalling for categorisation of environmental features (Ampatzis et al., 2008), expecting these signalling systems to produce more adaptable behaviours, especially when porting controllers evolved in simulation to the real world. In successfully evolved controllers, signalling emerged without any incentive as a cue to distinguish between two different environments, that the robots could recognise only after some exploration. Minimal instances of communications have also been evolved to allow the synchronisation of a swarm (Trianni and Nolfi, 2009), as well as to coordinate the activities in robot pairs (Tuci, 2009; De Greeff and Nolfi, 2010; Uno et al., 2011). Besides evolutionary approaches, other automatic design methods have been proposed that are capable of producing efficient communication for behaviours such as aggregation, coordination and categorisation (Hasselmann et al., 2018).

By looking at these research studies, some important lessons can be learned. Indeed, communication systems may emerge spontaneously, even if there is no explicit reward (e.g., no selective pressure from the fitness function in the evolutionary robotics approach). Furthermore, they evolve to provide an advantage to evolved populations compared to those that evolved without means to communicate.

Except for a handful of examples that present features of communication necessary for the emergence of human languages (i.e., compositionality and joint attention Tuci, 2009; Uno et al., 2011), the studies available in the swarm robotics literature obtained communication systems limited to simple instances of signalling, very far away from the complex communication schemes that characterise animal societies and, of course, humans. Indeed, the emerged communication systems are defined as a limited set of fixed signals triggering a pre-determined and hard-wired reaction in conspecifics, akin to early definitions of signalling (Owren et al., 2010). This is especially obvious in several studies within the swarm robotics literature (Trianni and Dorigo, 2006; Ampatzis et al., 2008; Trianni and Nolfi, 2009; Tuci, 2009), wherein the signal produced by an individual robot yields reactions from its neighbours as well as from itself.

Automatic design methods present several disadvantages for the evolution of communication. Indeed, these design methods rely on simplistic building blocks (e.g., neurons, as in Trianni (2008), or predetermined behaviour modules, as in

Francesca et al. (2014)), allowing for little variety in the resulting communication processes. Moreover, the emergence and evolution of the communication rules is strictly confined to the training step and thus the evolved rules remain identical after deployment. The current use of automatic design methods therefore limits the adaptivity necessary for communication in uncertain environments.

2.5 Language games

Multi-agent artificial systems have been widely used to study the evolution of natural languages *in silico*. For instance, the "talking heads" experiment (Steels, 2015) was the first to showcase self-organisation of language in a complex system within a population of artificial agents.

This approach does not rely on evolutionary computation or any other automatic design methods to grow the language mechanics. Instead, inspiration is from linguistics, whereby the self-organisation of natural languages is classified as cultural evolution, in opposition to biological evolution. Cultural evolution is the result of the sum of local interactions based on confirmation and agreement rules. In computer science, such interactions are often modelled by simple games played by a population of agents, seen as potential ways for them to cooperate (Ackley and Littman, 1994). When linguistic interactions are considered, a *language game* can be defined to be played between agents/robots, turn by turn, with the purpose of mimicking real-world dynamics leading to the emergence of a structured language. Language games make direct reference to concepts developed in the philosophy of language (Wittgenstein, 1953). If, for Wittgenstein, language games were simple abstractions of a real-world language, in computer science language games are minimal algorithms that display, in an artificial context, the salient characteristics of a whole language.

In such context, linguistic conventions can provide compact ways of identifying relevant aspects of the environment (e.g., different terms to identify different resources from which to forage), which can evolve to adapt to a changing landscape (e.g., assigning new terms to newly discovered resources, or dropping terms associated with depleted resources), hence maximising the communication efficiency. Moreover, an evolving language can be useful to represent sequences of terms, providing swarms the ability to decide on the most useful course of action (e.g., a sequence of resources from which to forage).

Various kinds of language games have been proposed to date, such as the imitation game which deals with vowel vocalisation (De Boer, 2000). In this work, agents are equipped with an articulatory synthesizer, a module for calculating the distances between different vowels (according to human perception) and a repertoire for storing vowel prototypes. Then, two agents (among many) are selected randomly and start the game interaction. The first agent (*the initiator*) selects a random vowel from its repertoire and utters it. The second agent (*the imitator*) then tries to imitate this vowel by uttering the closest vowel in its own repertoire. The initiator subsequently has to find the closest vowel to the one uttered by the imitator in its own repertoire, the goal being to thus find the initial vowel. Depending on the issue of previous games and on the success of the current one, both agents then either "merge" their vowels (they shift their vowel in the articulatory space towards the one they perceived) or add a new one. This protocol, coupled with some communication noise, causes the emergence of vowel systems that are strikingly similar to those found in actual human languages because the agents self-organise in order to produce vowels that are as distinguishable from each other as possible.

The imitation game requires to separate the agents into initiator and imitator categories. Other language games instead rely on a speaker and a hearer, which are interchangeable roles for the agents playing the game. This is the case for the *guessing game* (Steels, 2001), where the speaker chooses a concept within a context (physical or abstract), and communicates the corresponding word to the hearer. The latter has to guess which topic was chosen based on the communicated word. If it fails, both speaker and hearer update their inner representation of the concept. The guessing game can be seen as an implementation of the *Gavagai* thought experiment (Quine, 2013), addressing the question of the inscrutability of reference in a computational context, that is, the fact that one word can never have exactly the same meaning for different agents. Similar to the guessing game, the *category game* (Puglisi et al., 2008; Baronchelli et al., 2010) aims to self-organise discrete sub-intervals of one or many perceptual channels through negotiation dynamics. The agents start without any predefined category, and develop a pattern of categories shared among the agents via repeated interactions. Eventually, a global agreement emerges within the population. The negotiation dynamics lead to a *communication grounding* among all agents (as detailed in Clark et al., 1991), assuring a matching signified/signifier link between words and concepts to be exploited in future communications.

The category game can be simplified into the *naming game* (Steels, 1995, 2003) where categories are provided from the beginning, shifting the emphasis of the game on the negotiation dynamics and the emergence of an agreement. In this game, two or more robots interact to assign a unique name to a set of objects. At each interaction, one robot is chosen as a speaker and another as a listener. The speaker chooses a referring object and an associated word from its vocabulary—or invents one when no word is available—and then transmits it to the listener. If the listener knows the word, then the game is a success, and both agents remove all other words associated to the chosen object from their vocabulary, keeping only the shared word. If instead the listener does not know the received word, then the game fails, and the listener adds this new word to its vocabulary. We use in our study a specific version of the naming game, the minimal naming game (MNG, see Baronchelli et al., 2006b; Baronchelli, 2016). In this version, focus is given only to reaching consensus on a single word within a population of communicating agents. In a specific variant, speakers broadcast their word to all agents in their neighbourhood, while the listener is the only agent that updates the vocabulary upon success or failure of a game (Baronchelli, 2011).

2.6 Coupling between task and language

Various degrees of coupling are possible between language games and the behaviour of a robot swarm. First, at the lowest level of complexity, the robot behaviour is not affected at all by the language game, which is simply played by the robots upon repeated encounters (Trianni et al., 2016b). Second, the language game can affect the behaviour of the robots, but the latter has no direct influence on the way in which language evolves (Cambier et al., 2017). Finally, in the third case, the behaviour of the robot affects the evolving language, resulting in a strong coupling between the two (Cambier et al., 2018, 2021). As the strength of the coupling between language and robot behavior increases, so does the complexity of the emergent swarm behavior. This heightened coupling allows the swarm to leverage the entire spectrum of language complexity (Group" et al., 2009)). In the first two types of coupling, language merely conveys information useful for the swarm designer, serving as descriptors of the environment or task. However, in the third scenario, the language carries an emergent semantics intrinsically relevant for the task execution, hence enabling the robots to purposefully utilize language themselves. This capacity becomes a crucial aspect in the development of grounded symbols (Harnad, 1990).

Trianni et al. (2016b) studied the consensus dynamics generated by the MNG in a dynamic network formed by robots moving about in a bounded arena, without any interaction between language game and robot behaviour. In this research, the communication network was shaped by the encounters between the robots, each independently performing a simple random walk. This work concluded that the collisions between wireless transmitted messages, due to the simple communication protocol and the relatively high density of robots used in the experiments, led to the abortion of a significant portion of games. This turned to be a positive fact as the strain on the robots' memory was thus reduced, which is advantageous considering the limited capacities assumed in swarm robotics (Brambilla et al., 2013). Moreover, the embodiment of the robots and their collisions led to the formation of aggregates of robots that do not easily disband, leading to a reduced interaction rate in the population and a slower convergence with respect to simulated agents. This second phenomenon impacts the capacity of information transfer within the swarm, but does not impair the ability of the swarm to reach consensus, albeit with longer delays.

Recent studies focused on the effects that a self-organised behaviour and the MNG can have on each other. In Cambier et al. (2017), a swarm of robots performed self-organised aggregation and concurrently played a MNG where the exchanged words were used to identify the aggregate to which robots would belong to. Under specific density conditions, robots split into a controllable quantity of coalitions, each characterised by a different word used as identifier.

Cambier et al. (2018, 2021) considered a further improvement, as the words used within the MNG encode the parameter of the aggregation controller, directly impacting the quality of the self-organised aggregation behaviour. As a matter of fact, in the MNG, words supported by highly-connected agents propagate more (Baronchelli, 2011). This means that agents that better aggregate are able to widely propagate their words, and, thus, the aggregation parameters. This creates a positive-feedback loop that selects and maintains parameters promoting the formation of stable and large aggregates. Variations of the available words are introduced in the swarm from errors in the communication, which is modelled as a stochastic process in which some bits are flipped during the transmission. Such variations allow to explore new aggregation parameters, which can lead to changes in the way robots behave. The dynamics of the MNG therefore leads to the cultural evolution of the aggregation behaviour itself.

Chapter 3

Tools

Throughout this thesis, the aim is to have results from foraging experiments run in controlled environments that still replicate to some extent real life conditions. For that, the best appraoch would be to use real robots and tangible objects to be retrieved. However, this can be time consuming and can introduce issues that do not relate with the main research questions addressed in this thesis, and potentially skew the appreciation of the design of the experiment. For instance, the quality of sensors and actuators of the specific robots being used can have a strong impact on the success of their behavior (reaching a destination, grabbing an object, and so forth). The capacity of their battery limits the length of the experiments, and the number of physical robots available limits the size of the swarm studied. Working in simulation alleviates these issues, enabling quicker prototyping and frequent feedback loops on the design of the robots' behavior.

In the present work, all experiments were done in simulation, using the ARGoS simulator (Pinciroli et al., 2012), which has been widely exploited for swarm robotics experimentation and proved sufficiently realistic to enable a smooth transfer of robot behaviours observed in simulation to the real-world.

3.1 The ARGoS Simulator

ARGoS is an open source multi-robot simulator (Pinciroli et al., 2012) designed to simulate complex physical experiments involving large swarms of robots of different types. It allows simulated robots and simulated objects to share a 3D environment, and lets them interact through an array of sensors (e.g., light sensors, IR sensors detecting the ground color, omnidirectional cameras, proximity sensors and many others) and actuators (e.g., grippers, coloured LEDs, wheels, to name a few). ARGoS allows to simulate in detail real robotics platforms (e.g., the marXbot shown in Figure 3.1a or the e-puck shown in 3.1b). The fidelity of the simulated model (plus potential noise on sensors/actuators) aims at reducing the simulation-to-reality gap, allowing to substantially trust the results achieved in such simulation. To counteract the high computational power required by detailed simulations, ARGoS was designed with a specific focus on both efficiency (performance with many robots) and flexibility (customisation for specific experiments of the robots' model, sensors, actuators, as well as the representation of the experimental environment).

In this work, the overall aim of each robot is to start moving from a home location (often referred to as *nest*), explore the environment (which can be open or closed) and locate sources from which to forage valuable items. The foraging activity consists in navigating between nest and source to bring home as many items as possible. Within ARGoS, the home location is represented as a black disk painted on the ground, recognisable by the robots through ground sensors when they are on top of it. The rest of the environment has instead a white ground. Sources can be either intangible (represented as a gray disk on the ground) or tangible (a cluster of physical objects representing resources scattered around a specific location). A open environment is modelled within ARGoS as an experimental arena without boundaries, while a closed environment is limited by surrounding walls, which can be perceived by robots through their proximity sensors.

3.2 Robots

Two different types of robots where used in the experiments presented in this work. First, the *marXbots* (Bonani et al., 2010b) (Figure 3.1a)—used mostly in the first part, see Chapter 4 and 5—is a 17 cm high and 17 cm diameter robot developed for swarm robotics research (Bonani et al., 2010a; Dorigo et al., 2013). In the final experiment presented in Chapter 6, the *e-puck* was used instead. The e-puck is a 4.5 cm high and 7.4 cm diameter robot, similar in the basic functionalities to the marXbot (see Figure 3.1b). This later change of robot was made to simplify comparison with previous studies.

Both robots have a differential drive motion, and their speed is measured by a wheel encoder (with noise modeled using a Gaussian distribution), which supports simple odometry and local positioning/navigation. Both robots have a differential drive motion that supports simple odometry and local positioning/navigation. Their speed is measured by a wheel encoder. In order to approach the characteristics of a real robot, noise is added on the wheel actuation as follows:

$$W = N(0, \sigma)A \tag{3.1}$$

The actual actuated value W is the result of the ideal wheel actuation A (as set in the controller) multiplied by a random factor represented as a Gaussian distribution of standard deviation σ. When present in our experiments, we use $\sigma = 0.05$.

The e-puck has a maximum linear speed of $v = 0.1\,\mathrm{m/s}$, and the marXbot of $v = 0.3\,\mathrm{m/s}$. Avoidance of both obstacles (when present) and other robots is done at short range ($\approx 10\,\mathrm{cm}$) with infra-red proximity sensors. Robots can also avoid each other at a longer range ($\approx 1\,\mathrm{m}$) exploiting the infrared range and

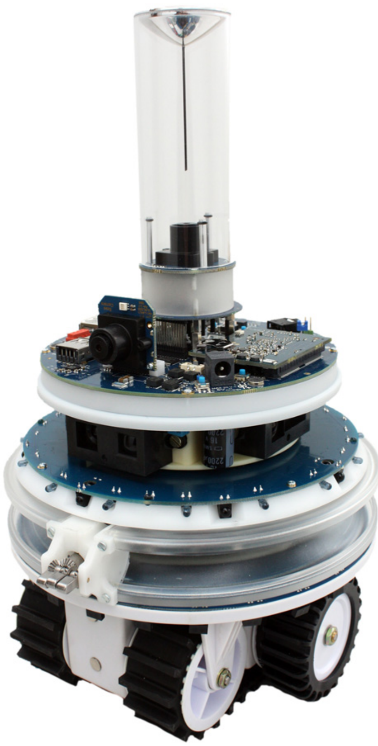

(a) Picture of the marXbot robot. Retrieved from https://www.swarmanoid.org/swarmanoid_hardware.php.html (2023)

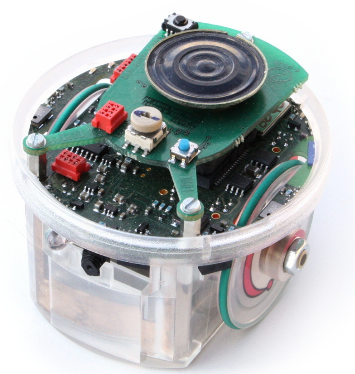

(b) Picture of the e-puck robot. Retrieved from Wiki Commons, by Stéphane Magnenat, CC BY-SA 3.0.

bearing system (Roberts et al., 2009; Gutierrez et al., 2009), which is also used for communication among robots.

Both types of robots perceive nest and sources only when they are located in the corresponding areas by means of infrared ground sensors, that robots use to differentiate between the white colour of the floor, the grey colour of the sources and the black colour of the nest. In each of our experiments, the robots start from (or near) the nest without any knowledge about the number and position of available sources, which need to be located through exploration.

Robots can locally broadcast short messages (e.g, 10 bytes) through the infrared range and bearing system within a range that is limited to 0.2 m for the e-puck, and 0.7 m for the marXbot. Robots can broadcast a message at regular intervals of 0.1 s with no re-broadcast of information received (no multi-hop communication). They keep track of the position of the nest and of the known sources through odometry. The error on positioning produced through this tracking method can be efficiently compensated through social odometry (Gutiérrez et al., 2010; Miletitch et al., 2013b), as discussed in Chapter 4. Owing to this, in the later experiments using the e-pucks, we neglect odometry errors.

When present in the experiments, marXbots can localize the items to be collected using the omnidirectional camera, which is used to perform a simple blob detection, and are able to recognise items up to a distance of 1 m. Items can be reached and then grabbed with a specific claw that rotates around the marXbots, making it easier to navigate while holding an item.

The control loop of the robots is executed 10 times per second.

3.3 Navigation

In order to achieve their goals, robots need to navigate the open arena. This includes searching for resources through a random walk, navigating toward a specific location (nest or known resource), or approaching specific objects, while avoiding robots and objects present in their surroundings. Each of these sub-tasks results in a force vector characterising the movement direction of the robots which is translated into wheel velocity according to the differential drive model.

3.3.1 Goal Vector

In our experiments, robots are aiming to move toward a specific direction at each step, defined as a goal vector. The goal vector is defined by the position of a source (or the nest), depending on the current state in the foraging behaviour. When they are exploring, a random unit vector is computed, with orientation drawn from a uniform distribution. This vector (changing every τ timestep in the case of experiments with marXbot, and every 5 τ for experiments with e-pucks) represents the next random move the robot will perform, to be aggregated with other force vectors from obstacle and robot avoidance.

3.3.2 Targeted Object Approach

When tangible resources are considered, we implemented a grabbing behavior for the marXbot. Once an object is in sight (located through its camera), the marXbot approaches the object using its location as the goal vector. Once close enough, it enters a grabbing behavior, opening its claw and closing it on the object. Once finished, the robot rotates its claw to put the object in its back, and then executes its next behavior (e.g., coming back to the nest to bring back the object). Postured this way, the object does not cause any meaningful occlusion for the robot sensors, and does not impede movements, allowing the robot to navigate in the same ways as when not holding any object.

3.3.3 Avoidance of robots and objects

Both marXbots and e-pucks exploit information from their long range sensors to compute a repulsive virtual force that pushes them away from their neighbours (Borenstein and Koren, 1989):

$$\vec{V} = \sum_{i \in \mathcal{N}} \frac{D_M - |\vec{v}_i|}{D_M} e^{-\mathrm{i}\angle \vec{v}_i} \tag{3.2}$$

The virtual repulsive force $\vec{V}$ in Equation 3.2 is computed for all neighbours $\mathcal{N}$ found within a maximum distance D_M. The expression considers the relative positions of neighboring robots, $\vec{v}_i$, and the force contribution from each robot is inversely proportional to its distance, growing stronger as robots come closer. The directional component, represented by the exponential term with $\angle \vec{v}_i$, ensures that the force acts in the opposite direction to the relative position vector, pushing the robots away from each other. The sum over all robots within the set $\mathcal{N}$ provides the total repulsive force, allowing for controlled separation and collision avoidance within the swarm.

The short-range avoidance behavior is described by Equation 3.3, where $\vec{S}$ represents the total repulsive force within the short-range maximum distance D_S. The sum runs over all robots or objects in the set $\mathcal{N}_s$, considering their relative positions $\vec{v}_i$. Unlike the long-range behavior, the weights w_i are binary, with a value of 1 if proximity is detected and 0 otherwise, reflecting the low resolution of the proximity sensor. This ensures a distinct reaction to close neighbors, allowing the robots to respond effectively to immediate obstacles or other robots in their path.

$$\begin{aligned} \vec{S} &= \sum_{i \in \mathcal{N}_s} w_i \frac{D_S - |\vec{v}_i|}{D_S} e^{-\mathrm{i}\angle \vec{v}_i} \\ w_i &= \begin{cases} 1 & \text{if proximity detected} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{3.3}$$

3.3.4 Aggregation of force vectors

The final direction vector, $\vec{D}$, is a weighted combination of the goal vector $\vec{G}$, the long-range avoidance vector $\vec{V}$, and the short-range avoidance vector $\vec{S}$, as shown in Equation 3.4. The weights ω_g, ω_l, and ω_s balance the contributions of these components, allowing for a navigation strategy that takes into account both targeted movement and obstacle avoidance. This aggregation enables the robots to smoothly navigate through the environment, minimizing interference and congestion.

$$\vec{D} = \omega_g \vec{G} + \omega_l \vec{V} + \omega_s \vec{S} \tag{3.4}$$

This aggregation has been optimised for the marXbot to minimise the effects of robot density and congestion on the ability to navigate back and forth between resources (as studied in Chapter 4). The speed of each wheels of the robots is calculated from this direction vector as follows:

$$\omega_L = \omega_b - \omega_b \times \left(1 - \frac{\alpha_M - \alpha}{\alpha_M}\right) \tag{3.5}$$

$$\omega_R = \omega_b + \omega_b \times \left(1 - \frac{\alpha_M - \alpha}{\alpha_M}\right) \tag{3.6}$$

where α corresponds to the angle of the direction vector, ω_b to the base speed of the robots, and α_M the maximum angle allowed for turning, defined here as $\frac{\pi}{2}$.

When running experiments with the e-puck, the robots always take a small negative 30° deviation with respect to their direction vector in order to approach it from the right-hand side. In this way, during exploration, robots moving back and forth between nest and source create a loop that allows to minimise interference between robots, resulting in less congestion between robots entering/leaving the nest/resources (see Chapter 6 for details).

Chapter 4

Information aggregation mechanisms in social odometry

The goal of this chapter is to develop a cooperative exploration and resource exploitation strategy based on a peer-to-peer exchange of information between robots in a swarm. This chapter contributes to the overall goal of the thesis by showing how the communication protocols determining how robots aggregate information received from peers can influence the foraging dynamics, which in turn can have an impact on the swarm topology and the spread of information among robots.

Specifically, we introduce three novel information aggregation mechanisms built on top of the social odometry methodology. Indeed, the efficiency of social odometry as a navigation and resource exploitation mechanism—and the resulting collective dynamics of decision-making—depends heavily on the way information is shared and aggregated in the robot swarm. In particular, we found that even small variations in some parameters of the individual behaviour may lead to huge differences in the swarm dynamics. For this reason, we propose parameter-free mechanisms for information aggregation and processing that make the swarm adapt to the distribution of resources in the environment. These mechanisms have varying levels of tolerance for conflicting information. I analyze the impact of each mechanism on the swarm's resource exploitation, environmental navigation, and resulting topology and its link with communication between robots. I test these approaches in scenarios involving navigation between nest and resources as well as in more intricate scenarios that involve object manipulation. We then study the resulting dynamics of the swarm in simulation.

Section 4.1 presents the basic methodology for navigation and introduces the three proposed parameter-free mechanisms. Then, Section 4.2 focuses on navigation. Here, the only variable of the experimental setup that impacts the decision of the swarm is the distance of the source to the nest. This condition changes in Section 4.3, in which we focus on the exploitation of the source. In this case, the sources are defined both by their distance from home and by their quality (e.g., rate of regeneration of items in the source, and size of a source). When the sources vary in quality, they are not only valued based on their distance from home but also on the ease in finding/processing items from them. In this case, the swarm must adapt to the dynamics of the environment and find a balance between exploiting close sources which are easier to reach or farther sources that might be of better quality. In doing so, the swarm must continuously choose between focusing on one single source or splitting among many.

Overall, we demonstrate that the proposed information aggregation mechanisms not only find the best split among sources at a given time, but also react to the upcoming variations in quality as the sources get depleted and hence continuously exploit the environment in an efficient way. For that, it aims for a fine balance between exploration and exploitation so that the swarm can react quickly to variations while keeping a steady pace of exploitation.

4.1 Social odometry

In our experiments, the goal of the robots is to locate both a home area and a source area and then to efficiently navigate/forage between them. In all our scenarios and experiments, once a target area is discovered, its position is kept in memory and updated using odometry.

The information about target areas is shared with other robots upon encounter, following the social odometry principle. The way in which the information exchanged is shared and processed is independent of the individual behaviour of the robots, which is different in the two scenarios we present in this chapter. Therefore, we start by introducing the information processing mechanisms we have devised.

4.1.1 Information Sharing

While robots navigate between target areas, they share the information they have on the relative locations in order to counterbalance the reduction in information confidence. How and when this information is shared has a strong influence on the overall quality of the information in the swarm, and on its decision-making. Not all information is shared at the same time. When randomly exploring, the robots share the sole information they have. In the other cases, the robots share only the information of the last visited location.

Given that robots do not share a global coordinates system or a common reference frame, a transformation of the shared position is needed in order to fit the frame of the receiving robot (Gutiérrez et al., 2009). To that end, the robots use as a reference the communication axis defined by the usage of their range-and-bearing device. This transformation is presented in Fig. 4.1 for two robots i and j, j receiving a message from i.

For that transformation to be possible, robot i first needs to know the direction of robot j. To this end, our communication protocol follows a handshaking approach in which robots constantly broadcast their needs (nest or source) while the other robots in range answer back with relevant information only when they possess it. When receiving calls for information, the robots not only come to know which information they should share but also the direction of the robot (γ^i).

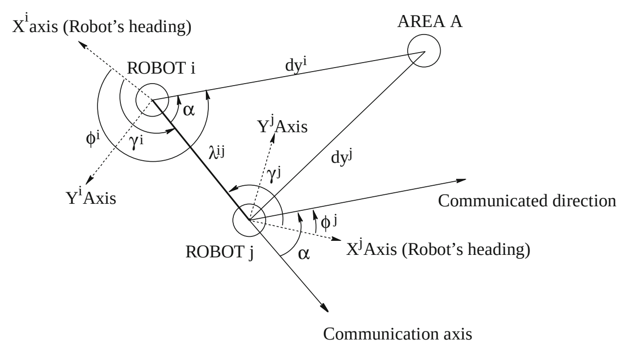

Figure 4.1: Diagram of the transformation of the shared position of area A between the frame of reference of robot i (emitting) and robot j (receiving), reprinted from Gutiérrez et al. (2009)

Once robot i receives a call for information from robot j, it shares back the distance (dy^i) to the relative area as well as its direction (α), the latter in the new frame of reference defined by the axis of communication: $\alpha = \phi^i - \gamma^i$

Now j needs to transform the received data in its own frame of reference. For that it must first find the communicated direction of robot i: $\phi^j = \gamma^j + \alpha - \pi$. It can then calculate the position of the target area in its own coordinate system:

$$dy_x^i = \lambda^{ij} \cdot \cos(\gamma^j) + dy^i \cdot \cos(\phi^j)$$
$$dy_y^i = \lambda^{ij} \cdot \sin(\gamma^j) + dy^i \cdot \cos(\phi^j)$$

with λ^{ij} being the distance between two robots, provided by the range-and-bearing device.

4.1.2 Information Processing

Once the information is received by robot i, it is aggregated with the robot's own knowledge. The way this aggregation is performed depends on the information processing mechanism implemented. Let $\mathbf{p}_i$, $\mathbf{p}_j$ be the estimated position of an area (either home or source) for robots i and j, and c_i, c_j be the confidence over their respective estimation. The result of any aggregation performed by robot j when receiving information from robot i is the updated pair $\langle \mathbf{p}_j, c_j \rangle$.

Here, we first describe the information aggregation mechanism used by Gutiérrez et al. (2009), and then we introduce our contributed mechanisms.

Fermi distribution The aggregation mechanism used by Gutiérrez et al. (2009) is based on a Fermi distribution. A weight is calculated from the difference in confidence in order to make a linear combination of the positions:

$$\langle \mathbf{p}_j, c_j \rangle \leftarrow \langle k \cdot \mathbf{p}_j + (1-k) \cdot \mathbf{p}_i, k \cdot c_j + (1-k) \cdot c_i \rangle$$

$$k = \frac{1}{1 + e^{-\beta(c_j - c_i)}}$$

The parameter β measures the importance of the relative confidence levels in the information aggregation. For low values, the aggregation is close to an average, ignoring the confidence. For higher values, the aggregation is stiff: only the information with highest confidence is kept.

Finding the right value of β is often a process of trial and error. Our contribution in this chapter is the introduction of three parameter-free aggregation mechanisms: *Hard Switch* (*HS*), *Random Switch* (*RS*) and *Weighted Average* (*WA*).Finding the right value of β is often a process of trial and error. Our desired outcome in this chapter is to find a workaround to this issue. To this end, we provide parameter-free algorithms that behave in a similar way as the Fermi mechanism, given different β value: *Hard Switch* (*HS*), *Random Switch* (*RS*) and *Weighted Average* (*WA*).

Hard Switch (*HS*) In this winner-take-all mechanism, the robots keep the information with highest confidence (either the current information or the received one) and discard the other one. In case of equal confidence, the current information is kept. This mimics the Fermi mechanism with a high β.

$$\langle \mathbf{p}_j, c_j \rangle \leftarrow \langle \mathbf{p}_x, c_x \rangle, \qquad x = \underset{k \in \{i,j\}}{\arg\max}\, c_k$$

Random Switch (*RS*) As in the mechanism above, here the robots keep one piece of information and discard the other. In this case, however, the switch is stochastic: the higher the confidence, the higher the probability of accepting the information. In practice, this mechanism is a stochastic version of the *HS*.

$$P(\langle \mathbf{p}_j, c_j \rangle \leftarrow \langle \mathbf{p}_i, c_i \rangle) = \frac{c_i}{c_j + c_i}$$

Weighted Average (*WA*) This mechanism consists in a linear combination of both estimated positions with their confidence as weight. On the one hand this implies no loss of information; on the other, when information about different sources is aggregated, the new position may not coincide with a real source location, leading to the apparition of artefacts. While the Fermi mechanism focuses on the difference between the two confidences, here we directly use each of them as weights.

$$\langle \mathbf{p}_j, c_j \rangle \leftarrow \left\langle \frac{c_j \cdot \mathbf{p}_j + c_i \cdot \mathbf{p}_i}{c_j + c_i}, \frac{c_j + c_i}{2} \right\rangle$$

4.2 Navigation Task

In this chapter, we focus on the navigation ability of the swarm as supported by the social odometry navigation mechanism. To this purpose, robots have to navigate between target areas represented as grey circles painted on the ground. We will study the influence that the different information aggregation mechanisms described in Section 4.1.2 have on navigation efficiency and collective decisions.collective decisions and navigation efficiency, defined as the number of round trips of the robots between nest and source.

First of all, we introduce the individual behaviour of the robots. Then, we introduce the experimental setup and finally we discuss the obtained results.

4.2.1 Individual Behaviour

The behaviour of the robot is defined by a finite state automaton with five states: *Explore*, *Go Home*, *Go to Source*, *Leave Home*, *Leave Source* (Fig. 4.2). Robots start in the *Explore* state and return to it whenever they lack relevant information. The other four states form a loop that corresponds to the robot navigating back and forth between the target areas: go to a target area, enter and leave it, then go to the next one. On top of these control states, both short and long range collision avoidance are implemented.

The robots start without any knowledge about the location of the target areas. Therefore, they first have to explore the arena. When in the *Explore* state, the robots perform a random walk until they discover the position of both target areas (home and source). This can happen in two ways: either they receive relevant information from team-mates or they stumble upon a target location (*Got(Area)* becomes true, with $Area \in \{Home, Source\}$). In both the *Go to Source* and *Go Home* states, the robots move straight to the target location, possibly avoiding other robots and obstacles. Along their way, they update the target areas location using odometry and update their confidence in the information. The confidence is defined as the inverse of the distance that the robot had travelled from the target area. Therefore, a straight path results in a higher confidence than a curved one.

Once a robot reaches an area (*i.e.*, *In(Area)* is true), it traverses it in a straight line (possibly dodging other robots to avoid collisions) and stores the area location. In order to get an estimated position closer to the center of the area, the robot averages its entering and its exiting positions. No matter how many sources there are in the arena, the robots always memorize only one home and one source (the last seen or agreed upon).

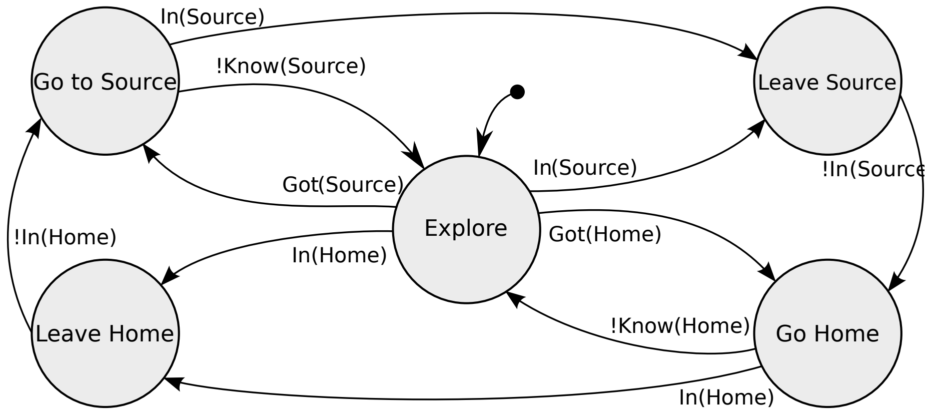

Figure 4.2: Robot's finite-state automaton. The circles define the states while the arrows define the transitions. *In(Area)*, $Area \in \{Home, Source\}$, is true when the robot senses the grey level of the area, *Know(Area)* is true when the robot knows the position of the area, *Got(Area)* is true when it just gets this estimation. The robots start in the *Explore* state.

4.2.2 Experiments

We used an experimental setup with as few variables as possible: a circular arena (radius: 11 m) with the home in the center and the sources scattered around (Fig. 4.3). The sources are characterized by their distance from the home (d_i) and the angle between each other ($\alpha_{ij} \in [\pi/3, \pi]$). Both source and home have a radius of 50 cm, and are coloured with grey levels to be distinguished by the robots. Unless stated otherwise, we used 75 robots spawned randomly.

By varying the number of sources, we study different aspects of the collective behaviour, such as the impact of the density of robots on their navigation abilities, the collective decision made by the swarm in a two sources setup, and how this generalizes in multiple sources setups.

Single Source When a single source is present, we expect that all robots will converge on the same path. The more robots in the arena, the harder it is for them to avoid each other. As density rises, the robots have to handle more and more congestion on their path, which leads them to travel bigger distances and to accumulate more error. This also corresponds to fewer round trips between the home and the source, hence lowering the efficiency of the swarm. We define the density on a path as the number of robots on it divided by its length.

In order to study the impact of density on navigation, we devised an experimental setup in which we vary both the distance between the home and the source and the number of robots. All three information processing mechanisms are tested and compared with a benchmark condition in which the robots are provided with perfect information (PI) about the source and home locations. In

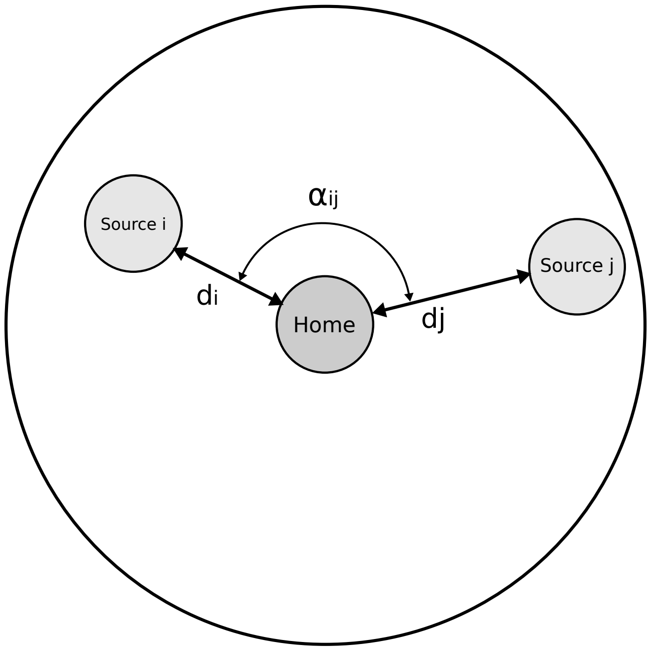

Figure 4.3: Setup of the experimental arena. The home area is placed in the center of a circular arena of 11 m radius, surrounded by walls. The sources are characterised by their distance to the home d_i, d_j and the angles they form with each other α_{ij}.

each experiment, we measure the navigation speed, computed as the number of round trips over time. We study its evolution for values of density between 2 and 40 robot/m. For each density value, we run 100 trials in which we randomly draw the distance between the home and source in the interval [3,8] m, and we compute the corresponding number of robots to obtain the specified density value (which is in the range [6,320]).

Two Sources When there is more than one source, a decision has to be made about how to spread the robots among the available paths. In this setup, we study if and how the robots converge on a single path as well as the implications of such a convergence on efficiency. In order to study this decision-making process, we count the number of robots committed to each source, as well as the uncommitted ones. Given that robots do not distinguish between different sources and only store one estimated position $\mathbf{p}_g$, a robot is considered to be committed to a source i among n possible if it has information about both source ($c_g \neq 0$) and home ($c_h \neq 0$), and if source i is the closest one to the robot's estimated source position $\mathbf{p}_g$.

In this setup, we have two sources which can either be at a short distance

(5 m) or a long distance (8 m). We run experiments with both equal and different distances for the sources: Short/Short (*SS*), Short/Long (*SL*) and Long/Long (*LL*). For each condition, we perform 1000 replications by randomly varying the angle between the sources with $\alpha_{ij} \in [\pi/3, \pi]$ (cf. Fig. 4.3).

Multiple Sources The environment in which a swarm evolves is rarely as simple as in the two sources setup. Through a multiple sources setup, we enquire about the scalability of the previously gathered results. M sources are uniformly distributed around the home location, with an angular separation between adjacent sources of π/M, where $M \in [3, 6]$. To investigate both the navigation and the decision-making abilities, we test three different conditions. Either all sources are at the same distance, short (*SSS*) or long (*LLL*), or a single source is closer to home (*SLL*). For each condition, we performed 250 trials.

4.2.3 Results

Each trial in all the previous setups lasts 20 minutes of simulated time. We use the same random initialization in all the runs for the different opinion processing. For each run we compute the number of robots on each path to study the dynamics of collective decisions, the number of round trips to study the navigation efficiency and the error made by the robots on the estimated position of the nest to gauge the quality of information in the swarm

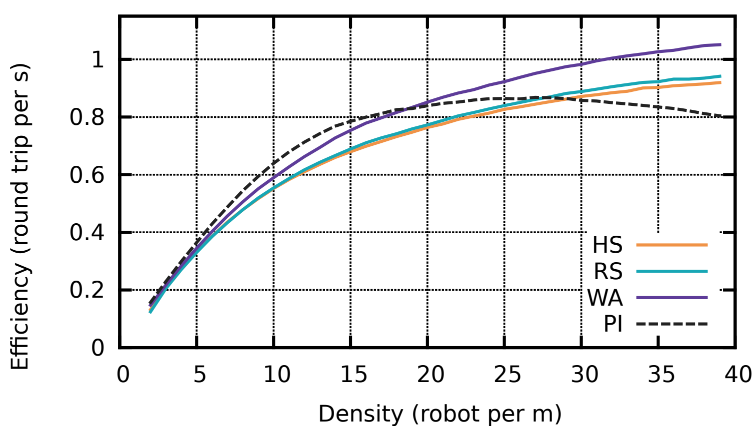

Figure 4.4: Impact of density on navigation efficiency for each mechanism and in the perfect information control condition. Each line is the mean over 100 trials.

Congestion

As we can see in Fig. 4.4, all the proposed mechanisms follow the same tendency. For low densities, we can observe a linear increase in the number of round trips.

With higher densities, the growth slows down. As expected, robots with perfect information are the most efficient at first, but their efficiency reaches a peak because of the artefacts created by perfect information. With *PI*, since all robots aim for the center of the target areas (either home or source), as the density rises they have increased difficulties avoiding collisions and entering or exiting the target areas.

Congestion has a lower impact on navigation efficiency with social odometry. In this case, *WA* proves to be more resilient to congestion than *HS* and *RS*. This is due to a smoother navigation in the surrounding of the home and sources, where robots try to enter small and densely populated area. First, since the *WA* mechanism never discards information but averages it, the precision on the estimated position is better than with *HS* or *RS*. Second, the reception of even slightly better information is smoothly integrated in the *WA* mechanisms resulting in better average information (Fig. 4.5), while in both *HS* and *RS* it may cause a large leap of the new location, which may be difficult to reach in case of high densities. Contrary to what could be expected, the quality of such information does not rise with the density of robots. Once there are enough robots to manage a steady connection between locations, the quality of information is virtually at its best. As the number of robots rises, congestion creates issues for them to reach each location, implying longer travelling distances and hence worse information kept in memory, despite enhanced communication relying on a denser network of robots.

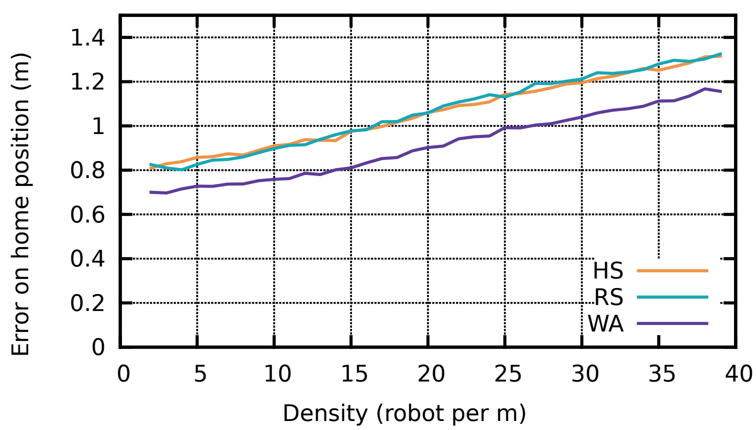

Figure 4.5: Evolution of the error of the estimated position of the centre of the nest for each mechanisms. Each line is a mean over 100 trials.

Collective Decision

Congestion explains why sometimes it is better to spread along multiple paths when there is more than one source. This decisions impacts not only the efficiency but also the spatial arrangement of the swarm and the way it reacts to changes in the environment.

Decision The decision pattern of the swarm results from the sum of local decisions made by the robots. The dynamics of the collective decision are shown in Fig. 4.6, which plots the convergence pattern generated by the *HS* and *WA* mechanisms when confronted with the SL experimental condition. In all cases, the swarm decides to focus on the closest area/source and most robots converge on the associated path. This behaviour is typical of all three social mechanisms when there is a source closer to home. We can already see a strong difference between the two mechanisms, where *HS* converges quicker, with less variations among experiments.

We can observe three different phases. At first (0-120 s), most robots are uncommitted and explore for source areas, reinforcing each as they discover them. Then (120-400 s), a competition among the two alternative paths occurs. The shorter path is reinforced more because of the improved information the robots have when encountering robots coming from the other source. Eventually, the swarm enters a maximization state in which mostly one path is exploited while uncommitted robots continue to join.

Fig. 4.7 shows the percentage of robots that choose path A (*i.e.*, the shortest path in the *SL* condition). We note that in the *SL* case, all information aggregation mechanisms lead to a single path convergence of at least 90% of the

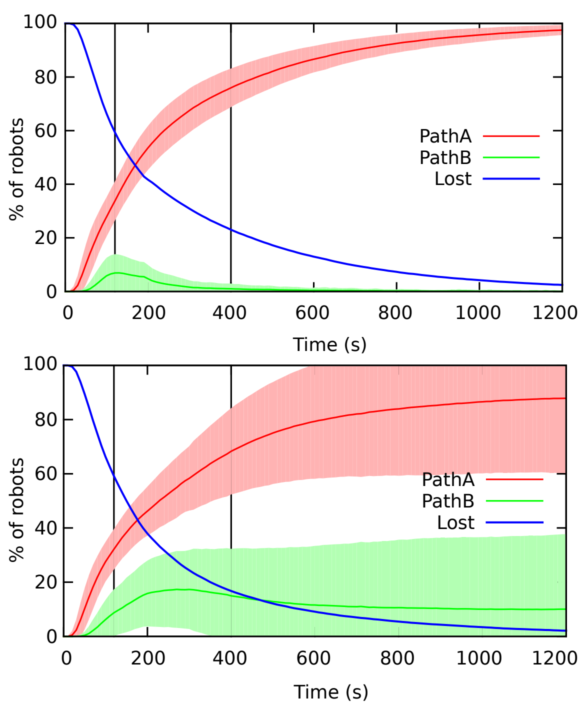

Figure 4.6: Evolution of the robots' repartition between the two target areas using *HS* (top) and *WA* (bottom) in the SL condition. Bold lines indicate the mean over 1000 repetitions, and the shaded areas indicate the standard deviation. These two figures present the two extremes in convergence pattern in case of the existence of a shorter path.

robots. Both *HS* and *RS* always lead to a convergence on the closest source. The same is the case for *WA*, which, however, also presents a low probability for the robots to converge on the distant source. This happens because with *WA* no information is discarded. When a large number of robots discovers the distant source early in the experiment, they may influence the whole swarm despite the lower confidence of their information. This cannot happen with *HS* and *RS*, because low quality information is instantly discarded. In both the *SS* and *LL* experimental conditions, when there is no better choice, *HS* and *RS* lead to a split in the swarm, and robots spread among the two paths (Fig. 4.7). In these experimental conditions, the more robots on a path, the higher the congestion, and the larger the distance the robots travel. This causes robots to have worse confidence in their information with respect to these from a less congested path. Therefore, switches to the other path are very likely. Congestion creates a sort of negative feedback that leads to an oscillating dynamic in which no decision ends up being taken. On the contrary, *WA* is not affected by such negative feedback and systematically leads to convergence (randomly on either path, the setup being symmetrical). Indeed, the poor confidence that results from congestion is counterbalanced by the larger number of robots with which the information is shared and averaged. Therefore, the swarm converges to the more populated path.

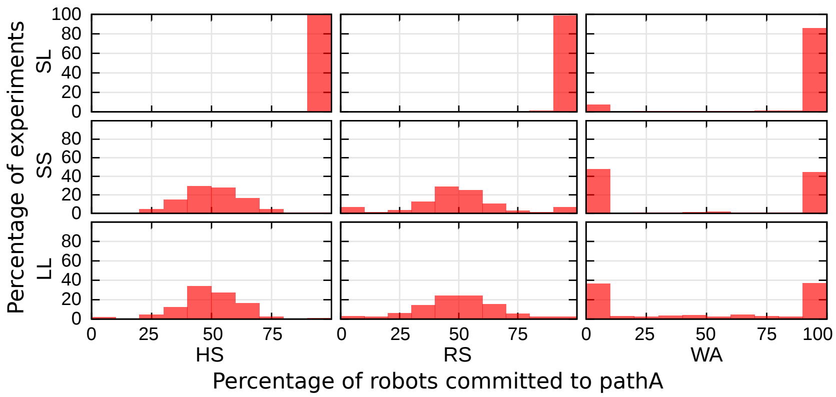

Figure 4.7: Robots repartition on path A. Each histogram shows the observed frequencies of the number of robots committed to path A (the shortest possible path).

Efficiency The robot behaviour does not explicitly encode the ability to make collective decisions. Instead, it is conceived to provide efficient navigation ability thanks to the information shared within the swarm. The decision process is an emergent result of this behaviour and so is the variation in efficiency depending on the setup and the mechanisms involved, as shown in Fig. 4.8. In the *SL* condition, all three mechanisms make the robots converge on the closest path, therefore resulting in density of 15 robot/m. As shown in Fig. 4.4, *WA* is more resilient to congestion, which is why it is the most efficient mechanism in this setup, followed by *RS* and *HS*. In the *SS* condition, both *HS* and *RS* result in the swarm splitting between the two paths as discussed above. By exploiting two paths with a low density of 7.5 robot/m (instead of one with high density of 15 robot/m) the robots create less congestion, which explains why the performance for *HS* and *RS* is slightly better than in the *WA* case. Indeed, *WA* makes the swarm converge on a single path with a high density, and navigation is slightly less efficient. Congestion has a lower impact in the *LL* conditions as both densities (9.4 robot/m on a single path, 4.7 robot/m on two paths) fall in the linear part of the congestion curve (see Fig. 4.4), explaining why the mechanisms result in the same efficiency.

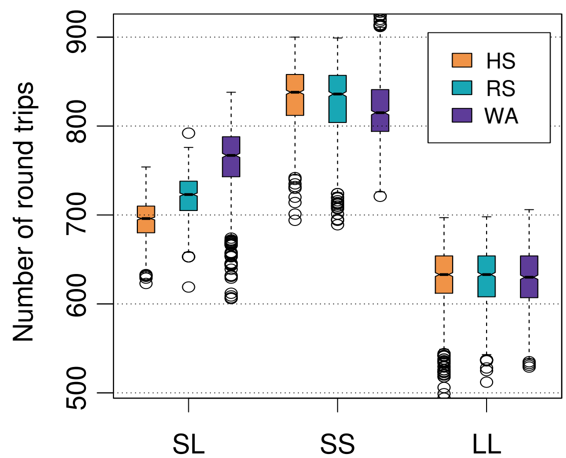

Figure 4.8: Efficiency of the swarm for two sources, for all mechanisms and conditions. Each box represents the inter-quartile range, whiskers extend to 1.5 times the corresponding quartiles, and the dots represent outliers.

Switching Patterns Social odometry and the various mechanisms studied above not only influence the efficiency of navigation and decision-making, but also the physical shape of the swarm. This can be seen not only by studying the movement of the robots, but by focusing on their switching patterns. A robot switches from one source to another when it encounters better information that directs it to another source. The switching patterns for each mechanism are displayed in Fig. 4.9, in case of both sources being at the same distance, or in the presence of a closer source.

First, we note that most of the switches occur in or near the nest. This is because the nest is the destination that all robots have in common, no matter their choice of source. This is where the density of information, and even more its variety, is at its highest. Furthermore, as mentioned above, robots do not share and request both pieces of information (the nest's and the source's position) at the same time. In order to switch from one source to another, a robot has first to enter the nest to request new directions. The halo of switches around the nest is the result of the range of communication allowed by the range-and-bearing device. Its shape varies for different arena setup (for instance more centred when the sources are on each side of the nest).

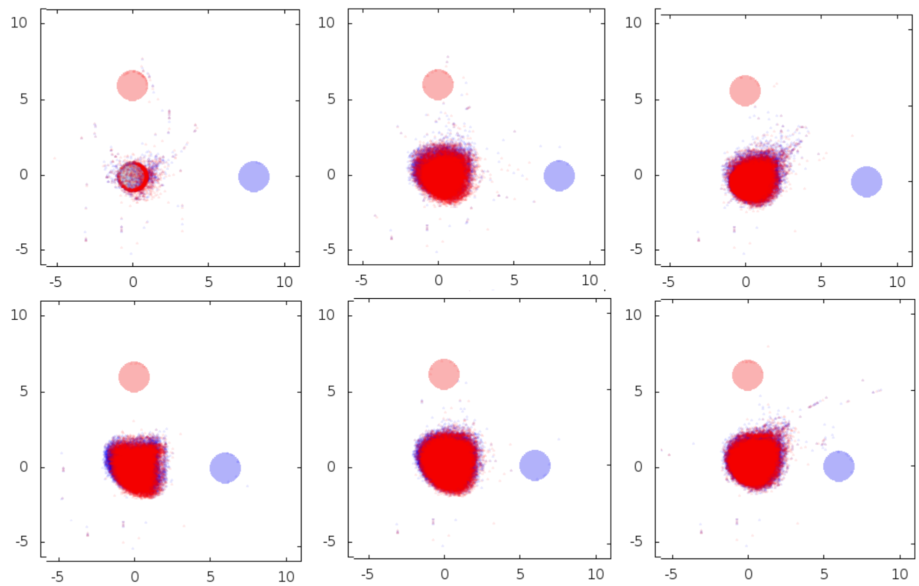

Figure 4.9: Cloud dot of the positions of the robots' switches from one source to another. In red are the switches to the red source (north, always the closest to the nest) and in blue the switches to the blue source. Red switches are drawn on top of the blue switches. Top: SL setup condition, bottom: LL setup condition - in both case the angle between the sources is $alpha = \pi/2$. From left to right: *HS*, *RS* and *WA*.

All mechanisms do not show the same pattern of switches. For instance, when there is a better solution, the *HS* mechanism only needs a few switches for all the robots to converge on the closest source. On the contrary, *RS* and *WA* show a much higher number of switches. Both observations are coherent with the speed of each mechanism's convergence. When no closer source is present, all mechanisms present a high number of switches, as robots oscillate between one possible solution and another. *WA* converges as in the previous condition, but with a higher number of switches.

In both *LL* and *SL* conditions, the switches pattern displayed have the tendency to grow toward the barycentre of both source. This effect is even stronger in the case of *WA* because of its averaging aggregation of information. This leads to the creation of a trail of switches, in which they are no longer the result of a communication with the nest, but with either path connecting the sources.

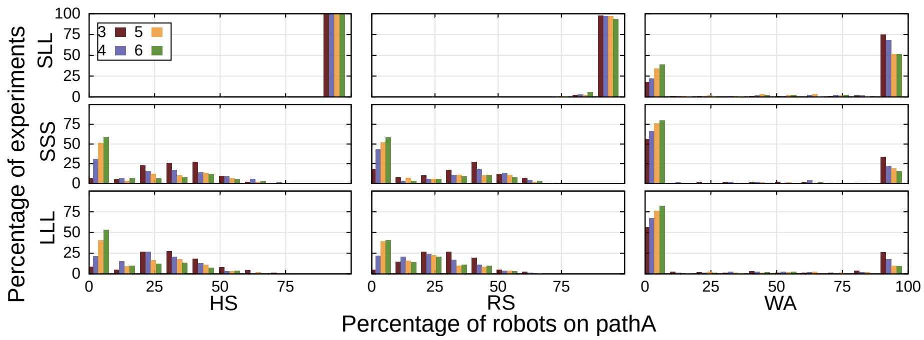

Figure 4.10: Robots repartition on path A for different number of source areas (3,4, 5 and 6). Each histogram shows the observed frequencies of the number of robots committed to path A (the shortest possible path).

Generalization to Multiple Sources

The dynamics we observe with multiple source locations are similar to the ones displayed in the two sources setup, no matter the number of added sources. Fig. 4.10 shows the percentage of robots that choose path A (*i.e.*, the shortest path in the *SLL* condition), when multiple source locations are present. All mechanisms leads to convergence in the *SLL* case, even if *WA* sometimes leads to the selection of one of the distant sources, for the same reasons discussed in the two sources setup. We can observe a similar splitting behaviour in the *SSS* and *LLL* conditions for both *HS* and *RS*, while convergence is observed for *WA*. When the swarm splits, the repartition of robots is no longer centred on 50% but is closer to 33%, implying that the repartition is no longer between only two paths. Nonetheless, not all are exploited at the same time, as can be inferred from the existence of paths selected by no robot. This can be explained by the oscillation dynamics discussed earlier. When the amplitude of the oscillations is greater than the number of robots on a path, all the robots on this path switch to another one. This happens in the case of multiple sources because the robots are spread among more paths, and their number on each is therefore lower.

To better understand the exploitation of the available sources, in Tab. 4.1 we report the average percentage of robots on the different paths, ordered from the most to the least exploited path. We note that the number of exploited sources is usually no greater than 3. This explains why the efficiency of the swarm does not vary with the number of available sources, as shown in Fig. 4.11. The slight increase in performance can be attributed to the fact that the more sources there are, the easier it is for uncommitted robots to join a path earlier in the experiment. Overall, we note similar patterns over efficiency between the multiple sources condition and the two sources condition.

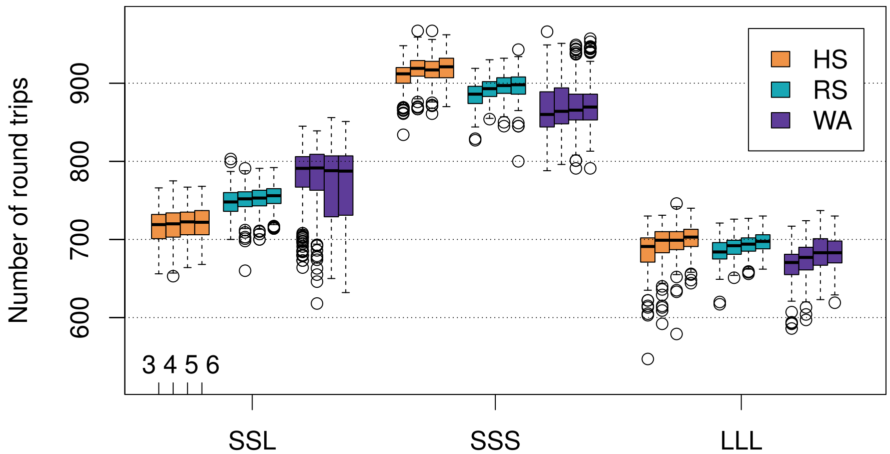

Figure 4.11: Efficiency of the swarm for multiple sources and all mechanisms and conditions. Each box represents the inter-quartile range, whiskers extend to 1.5 times the corresponding quartiles, and the dots represent outliers.

When there are multiple sources, *WA* in the *SLL* condition leads to a frequent selection of a distant source instead of the closest one, as shown in Fig. 4.10. If several distant locations are present, they end up reinforcing each other as their angular distance becomes smaller. In other words, two distant source locations that are close to each other attract more robots than a single closer location. This explains why the chance of *WA* leading to the selection of a distant source increases with the number of sources.

Discussion

The experiments above reveal the specificities of the three information aggregation mechanisms. *WA* leads to convergence to a single path in all conditions, but this is slower and error-prone (similar to the behavior resulting from a low β value for the Fermi mechanism). On the whole, *WA* leads to better cohesion of the swarm and deals better with congestion thanks to more accurate information about the target areas. *HS* and *RS* also lead to convergence when there is a shorter path to exploit, and handle better the presence of multiple distant source locations (similar to the behavior resulting from a high β value for the Fermi mechanism). When congestion results in inefficient navigation, both mechanisms lead to the exploitation of multiple paths, spreading the load of robots in a balanced way with similar dynamics, although *HS* appears to be stiffer than *RS*.

Table 4.1: Repartition in percentage of robots for 3, 4, 5 and 6 sources. The 1^{st} source is the one associated with the highest number of robots. The mean and maximum of the standard deviation is (4.7, 10.5) for *HS* and *RS* and (6.1, 16.9) for *WA*.

	SL			*SS*			*LL*		
	HS	*RS*	*WA*	*HS*	*RS*	*WA*	*HS*	*RS*	*WA*
1^{st}	98.5	98.1	96.0	48.0	52.6	93.0	48.8	47.0	90.8
2^{nd}	0.1	0.6	2.3	34.3	37.3	5.7	33.4	32.5	7.0
3^{rd}	0.0	0.0	0.0	17.2	9.7	0.0	17.4	19.7	0.1
1^{st}	98.4	97.7	95.2	50.6	54.1	92.3	44.8	43.8	89.5
2^{nd}	0.2	1.0	3.6	35.2	38.0	6.8	32.0	31.0	9.2
3^{rd}	0.0	0.1	0.0	12.5	6.9	0.0	17.6	17.2	0.1
4^{th}	0.0	0.0	0.0	2.1	0.7	0.0	5.1	7.0	0.0
1^{st}	98.6	97.3	92.4	51.1	51.1	94.8	44.9	42.6	89.4
2^{nd}	0.2	1.0	6.8	35.2	37.0	4.5	31.6	30.0	9.5
3^{rd}	0.0	0.1	0.0	12.5	10.4	0.2	17.7	17.7	0.5
4^{th}	0.0	0.0	0.0	1.1	1.1	0.0	5.1	7.3	0.0
5^{th}	0.0	0.0	0.0	0.0	0.0	0.0	0.3	1.4	0.0
1^{st}	98.6	97.3	93.6	50.1	53.0	94.7	43.7	42.4	88.5
2^{nd}	0.2	1.5	5.4	34.9	36.1	4.7	31.7	28.2	10.4
3^{rd}	0.0	0.1	0.5	13.5	9.2	0.2	17.2	17.3	0.6
4^{th}	0.0	0.0	0.0	1.4	1.3	0.0	6.0	8.3	0.0
5^{th}	0.0	0.0	0.0	0.1	0.2	0.0	0.8	2.3	0.0
6^{th}	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	0.0

4.3 Exploitation Task

In this section, we focus on the exploitation ability of the swarm as supported by the social odometry navigation mechanism. To this end, the sources are represented as a distribution of cylindrical items on the ground. By using objects, we give a topology to the source areas, and hence we can study how the different information aggregation mechanisms described in Section 4.1.2 react when confronted with physical resources, and their exploitation efficiency.

As in the previous set of experiments of this chapter, we start by introducing the individual behaviour of the robots. Then, we introduce the experimental setup and finally discuss the obtained results.

4.3.1 Individual Behaviour

The behaviour of the robot is defined by a slightly different finite state automaton (Fig. 4.12). The state *Leave Source* is now replaced by the state *Grab Item* for resources are no longer painted areas on the ground but items to be retrieved. The robots now go toward the source area, grab an item, return

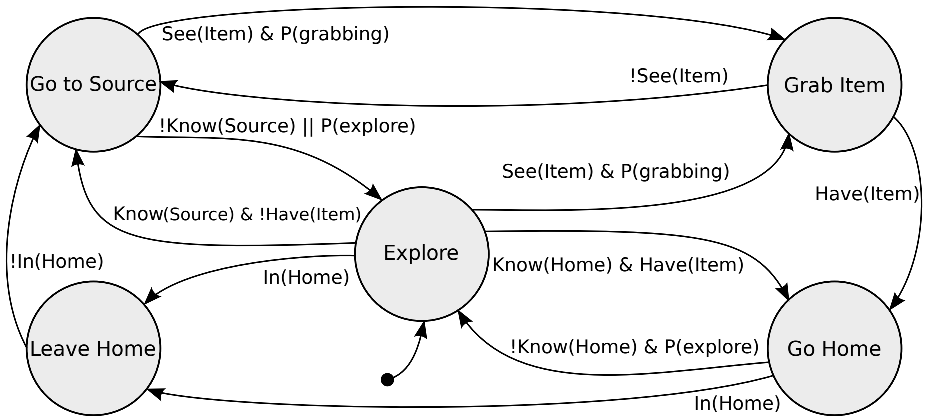

Figure 4.12: Robot's finite-state automaton. The circles define the states while the arrows define the transitions. *In(Home)* is true when the robot senses a grey colour on the ground. *Know(Area)*, $Area \in \{Home, Source\}$, is true when the robot has an estimation of the position of the area. *Have(Item)* is true when the robot is holding an item. *See(Item)* is true when the robot is able to see a grabbable item with its camera sensor. $P(grabbing)$ is the probability a robot will go grab the closest item.

home, drop the item, and start again following the foraging loop defined by the states *Go to Source*, *Grab Item*, *Go Home* and *Leave Home*. When lacking information, the robots fall back to the *Explore* state in which they start at the beginning of each experiment. On top of these control states, both short and long range collision avoidance are implemented.

In the *Go to Source* state, the robot moves straight to the target location, possibly avoiding other robots and obstacles. As in previous experiments, each robot updates its information (position and confidence) using odometry. Whenever a robot sees a resource item, it probabilistically enters the *Grab Item* state with a probability $P(grabbing)$. We wanted on average to allow the robots to cross the source, which led to $P(grabbing) = 1/(v \cdot d)$, where v is the speed of the robots and d is two times the standard deviation of the Gaussian spread of items characterising all sources. If the robot reaches the estimated location of the target source before getting to grab an item, it goes back to the *Explore* state. The robot has a small probability $P(explore)$ of going back to the explore state, which ensures that the robot does not remain idle in case it cannot reach the estimated position of the source area. This allows the swarm to reach a balance between the maximisation of currently known sources and the exploration of potential new ones.

Once in the *Grab Item* state, the robot moves toward the closest item. Then, once in contact, it grabs it and at the same time stores the source location as the average position of all grabbable items in sight. The robot always selects and goes toward the closest grabbable item, which may change over time due to robot movements or changes in the environment. Having grabbed the item, the robot enters the *Go Home* state. If for any reason no further items are in sight, the robot goes back to the *Go to Source* state.

The *Go Home* state works closely as the *Go to Source* state. In this case, the robot moves straight toward home. If it reaches the grey painted area, it enters the *Leave Home* state, and iterates the loop anew. If not, it goes back to the *Explore* state, either because it has reached its estimated position of the home location (without entering the grey area, implying that the robots had bad information memorised), or because of the probability $P(explore)$ to explore again.

Finally, when in the *Leave Home* state, the robot moves in the home area following a random walk pattern (possibly dodging other robots to avoid collisions) and probabilistically drops its item with probability $P(dropping)$. Once out, it stores the home location as the average of its entering and exiting positions.

When a robot is out of the foraging loop, it is in the *Explore* state. In this state, it performs a random walk, either searching for grabbable objects or for the home location. If the robot sees an item, it enters the *Grab Item* state with probability $P(grabbing)$. Otherwise, the robot exits from the *Explore* state only when it obtains the position of either the source or home, either from its own sensors or through social interaction.

4.3.2 Experiments

At the beginning of an experiment, the robots are spread inside an arena containing a home (circular grey area painted on the ground) and one or more sources of varying quality, as depicted in Fig. 4.13. The sources are regions with items to be grabbed and brought back home. Using real objects lets us shape these regions and define their topology through the items themselves, as opposed to regions painted on the ground, which are defined symbolically (allowing for only abstract interactions). Sources are Gaussian distribution of items around their centre with a fixed standard deviation of 0.5. Source regions were intended not as a dense bulk of numerous items (forcing interaction only on its edge) but as a wide spread of objects between which the robots can manage to get around. For this purpose, we introduced a minimum distance d_{min} between the cylinders equal to 5 times the robot's radius.

Furthermore, using real objects has the added effect of making the sources more complex, allowing for greater variations. The main motivation for using real objects (despite growing closer to real life conditions) was to integrate a notion of quality in the new source areas. The quality characterises the number of items present in a source at a given time. A source is defined by the maximum number of items and their rate of replenishment, expressed in items per second. This way, the quality of a source is grounded in reality and shares common proprieties with real life conditions. In the following experiment, we focus on the study of the rate of replenishment: the maximum number of items in a source is fixed at 35, which means that higher qualities are provided by higher rates of replenishment.

In all the following experiments we measure when possible: the number of items brought back home per second from each source, the number of objects in each source, the number of robots exploiting each source, the robots' switches among sources, and the quality of their localization information.

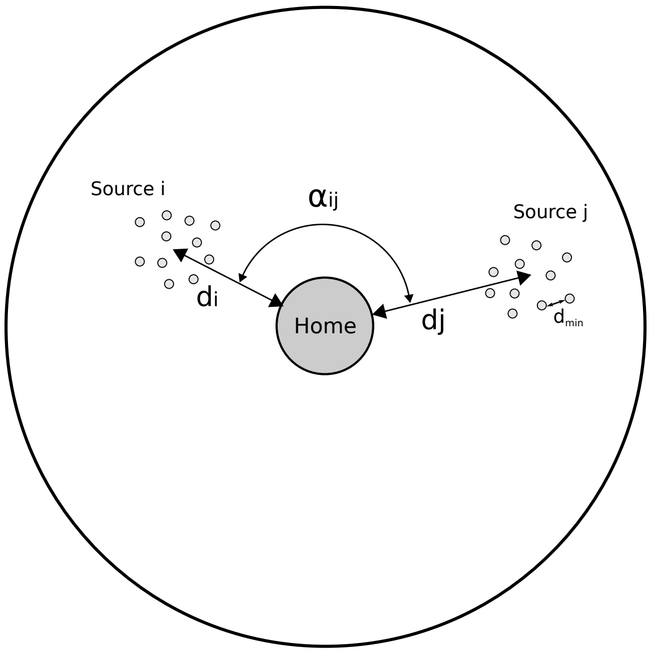

Figure 4.13: Setup of the experimental arena for the exploitation task. The home area is placed in the center of a circular arena of an 11 m radius, surrounded by walls. In this setup, the sources are defined as a Gaussian spread of items with a minimum distance between them of d_{min} (five times the robots' radius). They are characterised by their distance from home d_i, d_j, the angles they form with each other α_{ij} and their respective quality.

Congestion in the presence of items Our first objective is to understand the swarm behaviour with respect to navigation between the home and the sources, and the impact of the presence of physical items on previous congestion results. In order to compare painted area sources and spread item sources, sources should always have enough items for robots to grab in this setup. For that, the sources need to have a high constant number of grabbable items, implying an infinite quality. When an object is grabbed, it is immediately replaced by another one in the same position, so that the source size remains constant over time. This way, we focus on the navigation dynamics only.

Similar to the preceding congestion experiment, here we study the impact of the density of robots on congestion and hence navigation. For that, we use a setup with only one source. We vary both its distance from home and the number of robots to reach the wanted robot density values.

The following two sets of experiments focus on the exploitation of sources of varying quality, and the way a swarm of robots using *social odometry* reacts to a dynamic environment.

Optimal exploitation of a single source As mentioned above, the main interest of using real objects is the ability to study the impact of the quality of a source on the decision-making process of the swarm. In this experimental setup, we redo the same experience as above (one source of items, varying density of robots), but with a finite rate of replenishment. We vary its rate (0.1 $item/s$, 0.5 $item/s$ and 1.0 $item/s$) and study how various density of robots (5 $robot/m$ to 23 $robot/m$) manage to exploit a source bearing this rate. While the density variations are made by varying the number of robots, in this setup the distance between the source and the nest is constant, at the average of both the previously defined short and long distance (*i.e.*, $d = 6.5\ m$). Through these experiences, we search for optimum rates of a source's exploitation and their link with density and the rate of replenishment.

Optimal exploitation of two sources Finally, we study how the swarm decides and adapts in the presence of two sources. We study both the impact of the distance among sources as well as the impact of the rate of replenishment. For that, we choose among two possible distances ($d_{short} = 5\ m$ and $d_{long} = 8\ m$) and two possible rates ($rate_{min} = 0.1\ item/s$ and $rate_{max} = 1\ item/s$) for the sources. We will study each possibility. First, same rate and distance and same rate but different distance (to compare with the previous results). Then same distance but different rate (to study the impact of the rate). And last, different rate and different distance with the further source having the best replenishment rate (to compare the effect of the rate and the distance from home). Through these experiments, we explore the dynamics of the swarm and its ability to balance between the distance and quality of a source, and switch dynamically among sources in order to maximise its efficiency.

4.3.3 Results

In this section we present the current results over each experimental setup described above and compare them to the results presented in Section 4.2.3. We kept the same duration for the trials (20 minutes of simulated time) and the same random seeds. For each run we compute the number of robots on each path to study the dynamics of collective decisions. We also compute the number of round trips to study the navigation efficiency, as well as the error made by the robots on the estimated position of the nest to gauge the quality of information in the swarm.

Congestion

As can be seen in Fig. 4.14 left, all mechanisms follow a commonly shared tendency (sharp rise in low value of density, stalling for higher values). We can not make a direct comparison with results on density from Section 4.2.3 because the sources are not defined in the same way. For instance, the actual perceived distance can be much smaller for sources represented through items because the spread can grow closer to the nest than a static painted ground area would be. Furthermore, in the updated individual behaviour, each robot has the probability $P(grabbing)$ to stop exploiting the current source and explore. The density values output in Fig. 4.14 are starting densities. In previous experimentations, the robots had no possibility to go back to explore. We observed then a much erratic curve, a quick stall and even a drop in efficiency as density rose. Allowing the robots to explore again when they are stuck on the exploitation path not only give the swarm an opportunity to find better source, but helps the swarm exploiting the current source at an optimal rate by reducing the interferences between robots. In this experimental setup, the swarm self-organises to find a balance in the number of robots: too few would be a loss of potential, too many would make navigation non-practical.

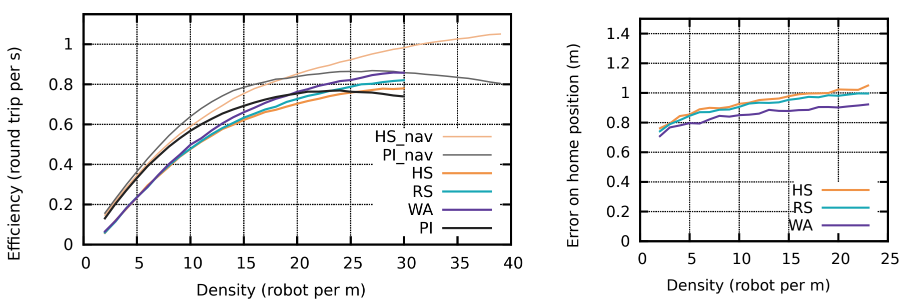

Figure 4.14: Left: Impact of density on navigation efficiency for each mechanism and in the perfect information control condition. Each line is the mean over 100 trials. *HS_Nav* and *PI_Nav* are taken from our previous results from Section 4.2.3 on congestion with simple painted sources, for reference. Right: Evolution of the error of the estimated position of the centre of the nest for each mechanism. Each line is a mean over 100 trials.

If we cannot actually compare the efficiency in absolute value, we can compare the evolution of this efficiency. *HS_Nav* and *PI_Nav* references of previous results shows us that the tendencies of results in each arena setup are similar, with a slightly stronger *RS* compared to both *WA* and *HS*. As in the previous setup, congestion has too a lower impact on navigation efficiency with social odometry. We note that *WA* is more resilient to congestion than *HS* and *RS* for the same reasons mentioned in section 3.3.1. The same trends as with the previous arena setup can be seen for the evolution of error (Fig. 4.14 right): error grows with density and *WA* is doing better through its averaging process.

Finally, if both physical setups differ a lot, the end results are similar. This proves first that adding physical interaction with items and the resulting updated individual behaviour do not change dramatically the higher level behaviour of the swarm. Second, such similarity indicates that our first abstraction of the physical setup is pertinent in a simple case with one nest and one source.

Optimal Rate of Exploitation

Using infinite rate allows us to study only the navigation aspect of the exploitation task. If we want to understand the dynamics of the swarm while exploiting a source, we need to study how the swarm reacts to the presence of a source of various finite rates (Fig. 4.15).

As in previous subsection, we plotted the efficiency of the swarm over its density. The three chosen rates show the three archetypical results. The top figure corresponds to a low rate (0.1 $item/s$) and display a virtually constant efficiency, no matter the number of robots. Indeed, in this case the rate is so low that even a small number or robots is enough to deplete the source and hence exploit it in an optimal manner. In the bottom figure, the rate is high (1 $item/s$) and the trends of each mechanism displayed in the plot are similar to the ones displayed in the previous subsection. In this case, for all density values tested, the source does not get depleted. Last, the middle figure has an in between rate (0.5 $item/s$). If at first increasing the density increases the efficiency, the curve reaches quickly a plateau around a density of 13 $robot/m$. After that, the rate is not high enough to withstand so many robots; increasing the density would only increase the number of exploring robots.

Exploitation of Two Sources

In this section, we study in a similar way as in section 3.3.2 the swarm dynamics and the collective decision when sources have varying distances and rates.

Decision In this section, we find on average similar tendencies than in the previous experimental setup with two source area painted on the ground (Fig. 4.16). *WA* always converges, even if sometimes on the longer path. *HS* and *RS* converge when a closer source exists. If such a source does not exist, then the swarm is split over the possible paths.

The plotted histograms reveal a few differences compared to previous experiments with two sources. First, the convergences are not as strong as previously observed. This is a result of the possibility for the robots to go back to explore when they are already on a path. Second, we see that when there is a competition between a closer source and a source with a better replenishment rate, the latter is the one toward which the swarm converges. Last, we note that in the perfectly symmetrical setup (LL with equal rate for both sources), the results are not symmetrical. The reason for this asymmetry is the now significant number of uncommitted robots. Each path bear in average less robots, the effect of which is that all the graphs are translated toward the left. Since the graphs are not

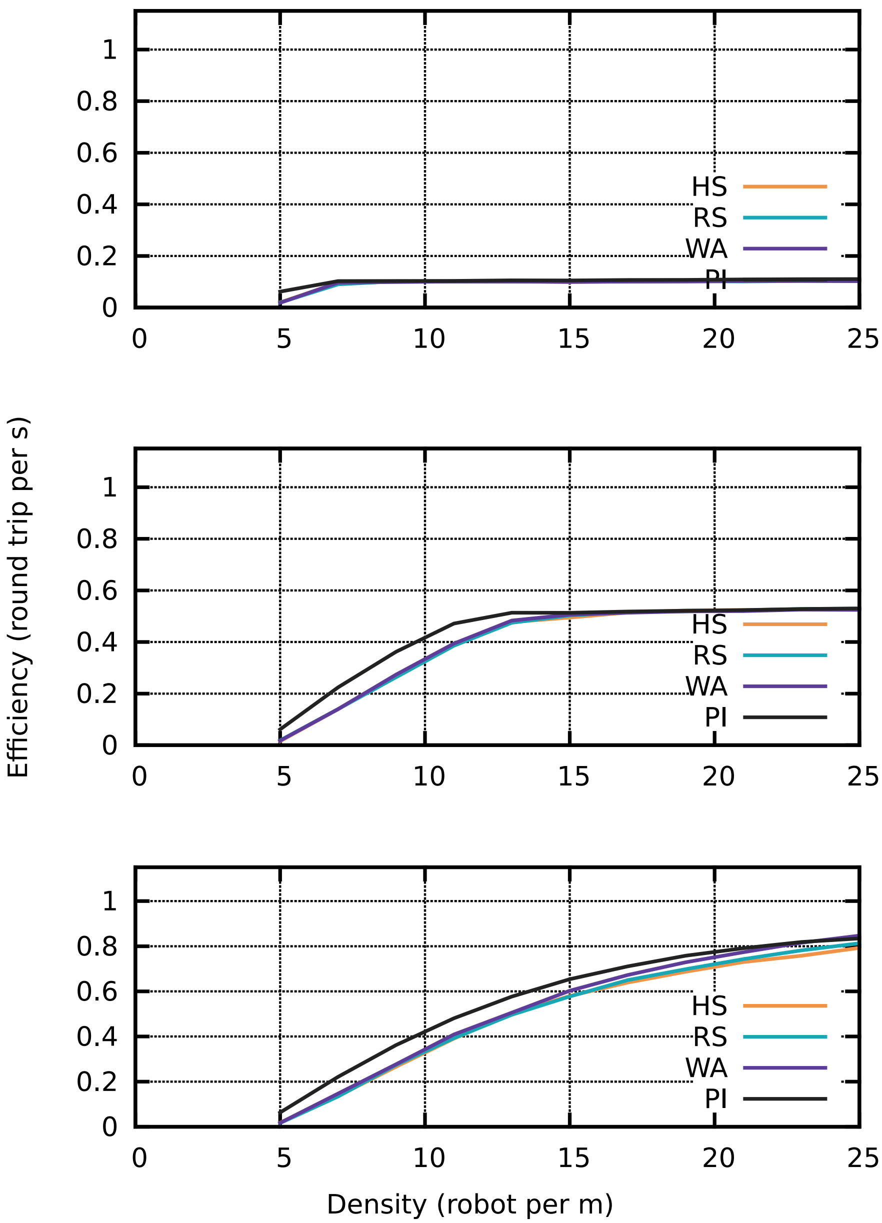

Figure 4.15: Impact of density on navigation and exploitation efficiency for each mechanism and for sources of various replenishment rate. From top to bottom, the rate is 0.1, 0.5 and 1.0. Each line is the mean over 100 trials.

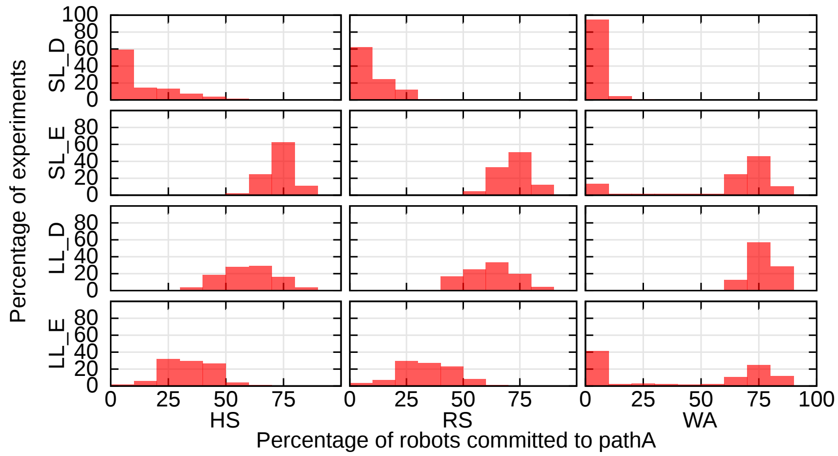

Figure 4.16: Observed frequencies of the number of robots committed to path A (the shortest possible). S: short, L: long, E: equal rate, D: different rate.

symmetrical anymore, it can be hard to spot a convergence just by looking at the histogram. Fig. 4.17 makes this convergence clearer by showing the evolution of the number of robots.

The black separators (at 120 s and 400 s) present in Fig. 4.17 correspond to the three different phases previously observed in Section 4.2.3 (exploration, competition and maximization). They show that despite both experimental setup having different swarm dynamics, their resulting evolution of the number of robots follow a same pattern.

Despite a clear spread of robots in the histograms, we observe in Fig. 4.17 top a clear convergence on the path with a better rate of replenishment for the *HS* mechanism. The remaining robots are not committed to the other path but mainly exploring the environment. In this experimental setup, *HS* converges even in the case of sources at similar distance, but only if their rate of replenishment differs. Here the swarm self-organise and proves that not only it can value a source on its distance from home but also on its rate by the amount of robots returning to an explore state when the source is being over exploited.

The Middle figure corresponds to the same conditions, but for the *WA* mechanism. The evolution of the number of robots follows a similar trend with the *HS* mechanism. It also converges more strongly on the path linked to the best replenishment rate source. In previous results, when sources were painted on the ground in Section 4.2.3, the number of exploring robots was strictly decreasing over time. Here, this is not the case anymore. After the number of robots on the longer path reaches its peak, the number of exploring robots starts rising again. It happens closely at the average time at which the source with lower replenishment rate gets depleted, inclining the robots committed to this

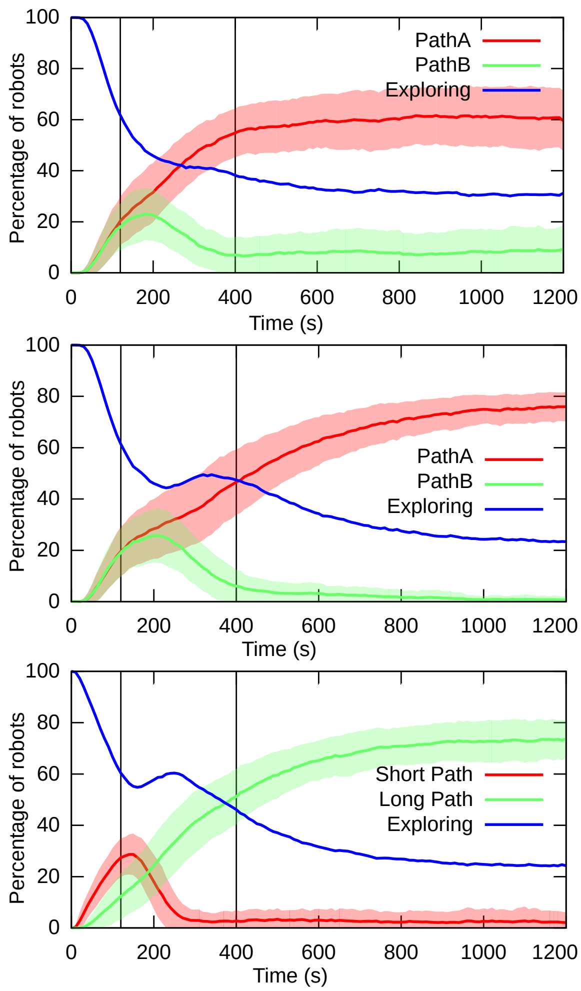

Figure 4.17: Evolution of the robots' repartition with two sources of varying distance (5 m and 8 m) and rate (0.1 $item/s$ and 1.0 $item/s$). Bold lines indicate the mean over 250 repetitions, and the shaded areas indicate the standard deviation. These three figures present the three archetypical convergence patterns with sources of finite replenishment rate.

Top: *HS*, *LL* sources with different rate.
Middle: *WA*, *LL* sources with different rate.
Bottom: *WA*, *SL* sources, the further source having the best rate.

source's path to go explore. Such exploring robots are then integrated in the better path.

Finally, the bottom figure presents a competition between a higher replenishment rate and a closer source while using the *WA* mechanism. We observe that if at first robots converge on the closest source, the latter quickly gets depleted. Then, the robots previously on the shortest path go back to explore and finally join the further but substantial source.

Efficiency We saw that the decision process is not only influenced by the distance of the sources from home, but also by the sources' replenishment rate. Fig. 4.18 displays the variation in overall efficiency over all mechanism and experimental setup when two sources defined as a spread of items are present.

When both sources' replenishment rate are equal, the boxplots describing the efficiency of the swarm are similar to those found in section 4.3.3. The main difference is an overall lower efficiency (a slower swarm), due to the physical interaction with the sources' items. When sources bear different rates, they are on average doing worse than when they have the same rate. This can be explained by the fact that the overall rate is higher when the two sources have the same rate. Another reason is that a setup with same rate will incline the swarm to split among two paths, which makes the swarm more efficient as we proved in section 4.3.3.

Last, in the *SL_D* condition, the swarm is less efficient than in the *SL_E* condition. The reason for that is that the closest (but with low rate) source is regularly rediscovered and depleted, distracting the robots from the high-quality source. This creates a cycle in which the closest source is regularly depleted and abandoned until it grows back enough for exploring robots to discover it again. These robots then come back home with information in which they are very confident, and hence recruit even more robots. This cycle is even more pronounced for the *HS* mechanism when just one single robot can spread its information to many others as long as its confidence is better.

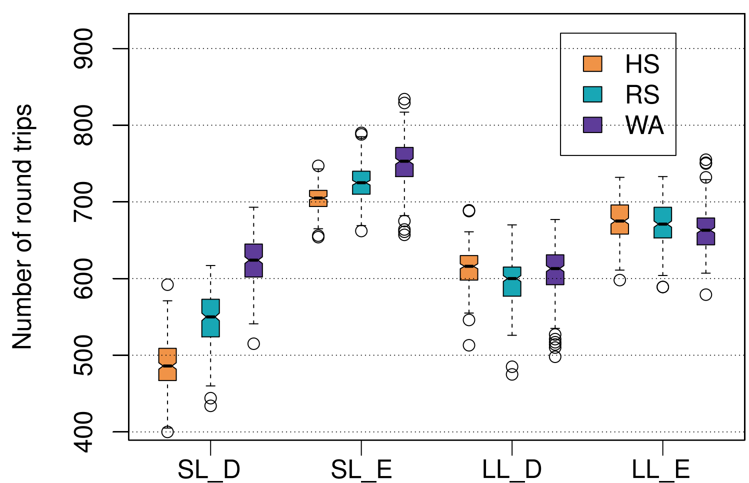

Figure 4.18: Efficiency of the swarm for two sources, for all mechanisms and conditions. Each box represents the inter-quartile range, whiskers extend to 1.5 times the corresponding quartiles, and the dots represent outliers. S: short, L: long, E: equal rate, D: different rate.

4.4 Conclusions

In this chapter, we presented an extensive analysis of three parameter-free information processing mechanisms for social odometry with abstract source locations and with sources made up of a distribution of physical items. We studied the impact of these mechanisms on both the navigation and exploitation efficiency and on the dynamics of the swarm. In particular, we observed how the information processing mechanism can either lead to a convergence on the exploitation of a single path, or to a split over multiple comparable options. These results are meant to give future designers a guideline of which mechanism they should choose depending on the situation and objectives considered.

For instance, if the cohesion of the swarm is important, then the *WA* mechanism should be selected as it favors a group that does not split over multiple sources. As for the navigation efficiency, we observed that it highly depends on the congestion of the selected paths. As a consequence, *HS* and *RS* lead to the exploitation of multiple paths whenever congestion results in inefficient navigation. When physical objects are present, the sources' rate of replenishment influence strongly both the efficiency of the swarm and its dynamics. We observed that variations in this rate has an even stronger impact on the efficiency than the distance alone.

These results show similar trends among both kinds of experimental setups, with more realistic interactions in the case of sources defined through items. In all setups, the swarm displays a behaviour in which it balances the robots' load over the possible paths (splitting when necessary), implementing a sort of *load-balancing* mechanism. In Chapter 5, we will investigate this issue further in order to provide an optimal load-balancing behaviour, which can maximize the exploitation of different paths to relevant areas/sources. In this context, the swarm favors the best distribution of robots among the available paths and reacts in real time to changes in its environment.

A number of possible extensions to the presented mechanisms can be envisaged. The first straightforward extension is to provide for our social odometry mechanisms a way to deal with more complex paths, for instance in the presence of obstacles. Robots may also be given the ability to memorize multiple source locations, implying that the competition among paths would not be only at the swarm level but also at the individual robots level.

Chapter 5

Balancing exploitation of renewable sources

This chapter further contributes to the overall goal of the thesis by digging into the interplay between foraging dynamics and communication. We propose an adaptive strategy to balance exploitation of renewable sources by a robot swarm, inspired by the nest site selection behaviour (NSS) of honeybees. Unlike the communication protocols studied in the previous chapter (which are unflexible with respect to changes in the environmental conditions), the NSS behaviour enables sustainable foraging in a multiple resource context and is flexible to varying contingencies, thanks to a richer communication protocol. Indeed, while in Chapter 4 communication was limited to sharing the position of a source where other robots could navigate to—hence fulfilling a recruitment function—, with NSS communication also serves to inhibit teammates from choosing a given source (e.g., because it could become depleted). This chapter demonstrates that, by enriching the communication system, it is possible to also observe richer foraging dynamics.

Individual limitations, such as the inability of robots to be aware of each available source and of its profitability, as well as the fact that the behaviour may be constrained to following simple reactive rules, entail that achieving a correct balance is not trivial. In such conditions, a balanced exploitation should result from a collective self-organising process in which information about the availability of sources is shared among the robots to achieve a correct allocation, preventing source depletion and maximising the flow of goods. Despite being conceived for collective decision making, NSS indicates how the swarm dynamics can switch between convergence toward the exploitation of a single source (when its quality is good enough to sustain a large swarm) and balancing between the available sources (when no source has sufficient quality to sustain the whole swarm). We start from a source exploitation problem whereby items have to be collected from multiple sources, which can replenish at a fixed, unknown rate. We borrow and adapt the collective decision-making algorithm from Reina et al.

(2015b), and we improve on it by studying for the first time how this algorithm adapts to dynamic environmental conditions that result from continuous source replenishment after foraging. Indeed, the variation of the number of items within a single source corresponds to a variation of its perceived quality, which has an impact on the macroscopic dynamics of the proposed algorithm that has never been studied to date. To this end, we isolate the different processes determining the collective dynamics and study their impact. This contributes to the identification of a parameterisation that can lead to a balanced exploitation not only with respect to the regeneration rate of the sources—hence avoiding source depletion—but also with respect to their distance from the central place where foraged items need to be retrieved. Integrating all these aspects proves particularly challenging, and we present here a large-scale study on the most important parameters determining the system behaviour.

5.1 Experimental setup

The experimental setup is similar to the one presented in Sec. 4.3. Robots have to search for items scattered in an open environment and retrieve them to a home location (hereafter referred to as "nest", in analogy to foraging by social insects). Differently from the previous experiment, in this case the arena has no boundaries, introducing a central place foraging problem in a open space. The nest is a circular area (radius: 0.8 m) at the center of the robot arena represented by a black disk painted over an otherwise white floor (see Figure 5.1). Retrievable items are cylinder-shaped objects (radius: 0.05 m) clustered together to form a "source". In this chapter, we focus on a simplified exploitation problem in which only two sources are present—labbelled A and B—although the proposed solution can be easily generalised to larger numbers of sources. Each source $i \in \{A, B\}$ contains at most $M_i = 30$ items positioned according to a 2D Gaussian distribution around the source centre ($\sigma_R = 0.35\,\mathrm{m}$), keeping a minimum distance $d_{min} = 0.14\,\mathrm{m}$ between items. Sources are characterised by quality and position. The source quality r_i is defined by the rate of creation of new items ($r_i \in \{0.01, 0.03, 0.05, 0.1\}$ items/second). The source position is defined by the distance from the nest ($d_i \in \{4, 6, 8, 10\}\,\mathrm{m}$). The relative angle α_{AB} between sources is chosen at random with a minimum angle of $\frac{\pi}{3}$ between different sources to ensure separation. When the number of items within source i is lower than M_i, new items are generated with the given regeneration rate r_i, practically implementing a Bernoulli model for source regeneration as described by Liemhetcharat et al. (2015).

5.1.1 Individual and Collective Behaviour

The overall goal of the robot swarm is to maximise the retrieval rate, that is, the number of items per unit time that are successfully transported to the nest. We assume that robots have no a priori knowledge about the position and profitability of sources, neither they have any map of the environment to support

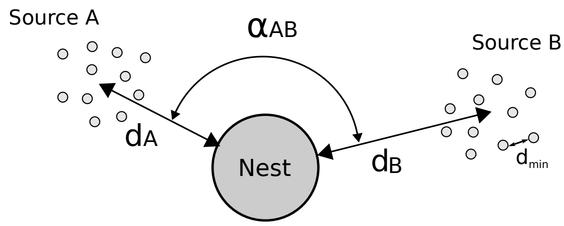

Figure 5.1: The source exploitation problem requires robots to bring cylindrical items back to the nest. The experimental setup defines two sources—labelled as source A and B—positioned at distance d_A and d_B from the nest, and separated by the angle $\alpha_{AB} > \frac{\pi}{3}$. Each source i contains at most M_i items, scattered around the source center according to a 2D Gaussian distribution, keeping a minimum distance d_{min} between items. The nest is a circular area (radius: 0.8 m) painted in black to be recognised through the robots infrared ground sensors.

navigation. Hence, exploration is required to gather information on the available sources. We also assume that robots can track—e.g., through odometry—the position of the nest and of at most one source at the time, which is the one most recently visited. In this way, we ensure that robots have up-to-date information about the state of the sources they have found, avoiding to memorize the location of sources that may be unprofitable or completely depleted. Finally, we assume that robots can carry only one item at a time, hence multiple robots can forage from the same source at the same time to maximise the exploitation rate.

The desired swarm behaviour requires a recruitment process, so that robots can spend less time in exploration, focus on the available sources and exploit them in parallel. When multiple sources are available, exploitation can be focused on one of them if its quality is high enough to sustain the whole swarm. Otherwise, a balanced exploitation of the two available sources is preferred. This collective behaviour has properties similar to value-sensitive decision-making studied in house-hunting honeybees (Pais et al., 2013; Reina et al., 2017). Indeed, when engaged in a collective decision, a honeybee swarm may arrive at consensus when the quality of the option is sufficiently high, otherwise remain in a “undecided” state when the quality is low, in the hope that a better alternative is discovered later. As a matter of fact, such undecided state corresponds to the swarm being split in sub-populations committed to the low-quality alternatives that they could find. This “undecided” state can be seen as a load-balancing state, because the house-hunting swarm is split among potential nest-sites, much as the foraging swarm is split among the available sources. This highlights the usefulness of framing the load balancing problem studied here in terms of a value-sensitive collective decision problem.

Starting from this observation, we decided to synthesise the individual robot behaviour taking inspiration from a design pattern derived from the honeybee nest-site selection behaviour (Reina et al., 2015b). The design pattern provides guidelines to implement the individual behaviour as a probabilistic finite state machine (PFSM), where any robot can be in two macro states: committed to exploit a known source, or uncommitted and exploring. Additionally, upon robot-robot encounters, local information exchange can lead to changes in the commitment state. Overall, four concurrent processes need to be implemented in the individual behaviour: (i) spontaneous discovery of any source i with probability $P_{D,i}$; (ii) spontaneous abandonment of commitment to source i with probability $P_{L,i}$; (iii) recruitment of uncommitted agents following interaction with a robot committed to source i with probability $P_{R,i}$; (iv) inhibition of commitment, whereby an agent committed to source i becomes uncommitted after interaction with a robot committed to source $j \neq i$, with probability $P_{I,j}$ (cross-inhibition). These probabilities are either completely defined by the problem itself (e.g., $P_{D,i}$ for discovery of source i, see Section 5.2) or are parameters defined at design time in order to tune the collective behaviour and to achieve the desired exploitation of the available sources. Cross-inhibition is particularly relevant, as it can determine the switch from the parallel exploitation of multiple sources to full convergence towards a single source. This mechanism has been observed in house-hunting honeybees (Seeley et al., 2012b), and is used to adaptively select nest sites of high quality, quickly abandoning those of low value (Reina et al., 2017). Indeed, through cross-inhibition, agents committed to a source can return uncommitted, explore for other—possibly better—alternatives, or get recruited by other agents.

In this work, we have implemented the individual behaviour as the PFSM represented in Fig. 5.2, which is executed every $\Delta t = 0.1\,\mathrm{s}$. Here, boxes represent macro states corresponding to the commitment state of a robot, while circles represent micro states corresponding to basic behaviours executed until some (probabilistic) transition is triggered. The robot is considered to be committed to a source when it knows its location, otherwise it is considered uncommitted. The actual movements of the robot are governed by the following basic behaviours:

- **Explore**: in this state, the robot explores the arena performing a correlated random walk (Dimidov et al., 2016). Whenever sufficient information becomes available (e.g., location of nest and sources), either through exploration or following interactions with other robots, a different behaviour may be triggered.

- **To source**: the robot moves toward the location of a known source to search for more items to retrieve.

- **Pick up**: when some item is in close range, the robot navigates toward it and picks it up. Should the grasping procedure fail, the robot tries again or chooses another item to pick up, if available.

- **To nest**: the robot navigates back to the nest, possibly bringing back an item to deposit.

- **In nest**: when in the nest, the robot deposits the item it is carrying—if any—and then performs a random walk until it moves out of the nest.

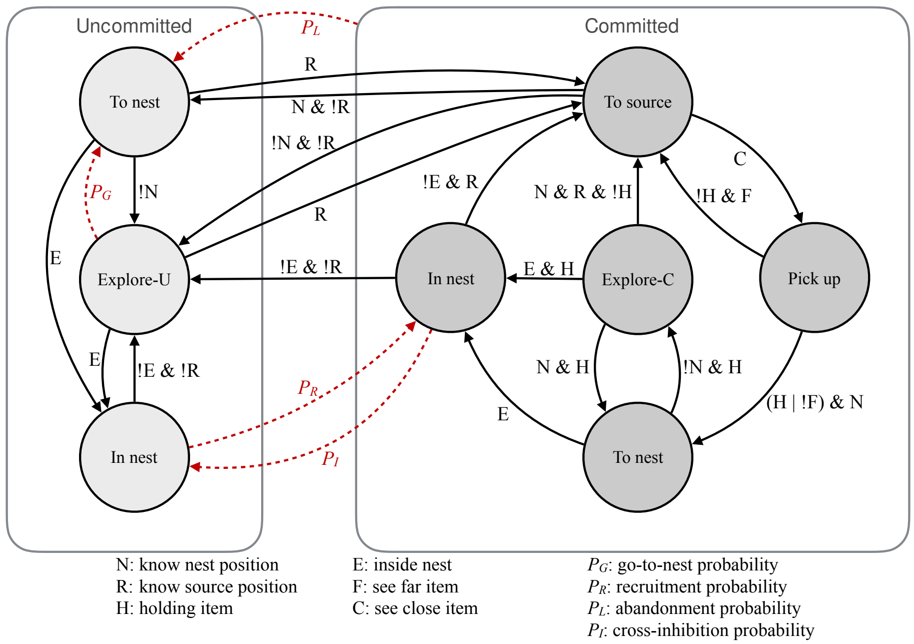

Figure 5.2: Probabilistic finite state machine (PFSM) representing the individual robot behaviour. The two boxes represent macro states for the committed and uncommitted robot. Circles inside the macro states represent PFSM states in which a robot executes a basic behaviour. The "Explore" behaviour is executed in different conditions depending on the commitment state, hence states are named differently to avoid confusion. Arrows represent transitions between states, and are triggered when a certain Boolean expression is verified (AND: '&'; OR: '|'; NOT: '!'; see legend at the bottom). When o̊bot is uncommitted (left box), it has no knowledge of any source and searches for it, periodically returning to the nest. When committed (right box), the robot knows where a source is and tries to retrieve items from it. The red dashed arrows represent probabilistic transitions.

Robots start from the nest at the beginning of each experiment and keep track of their positions through individual and social odometry (see Gutiérrez et al., 2010, more details below).

When a robot is uncommitted, it explores the environment to gather information about the location of the nest (if unavailable) and of the sources (see left box in Fig. 5.2). With a fixed probability P_G, a robot stops exploring and returns to the nest, where it has a high probability of interacting with other robots. With this mechanism, exploration is constrained around the nest and robots do not wander away for too long.

When a robot is committed, it moves back and forth between the nest and a known source to retrieve some item (see right box in Fig. 5.2). If it loses track of the nest, it explores the neighbourhood until either it finds the nest or it receives its location by other robots in its neighbourhood. When in the nest, it deposits the carried item and starts another exploitation trip to the known source. At any time—and from any state in the committed macro state—a robot can abandon commitment for source i with probability $P_{L,i}$. This corresponds to erasing the information about the source and returning back to the nest, from where to retrieve exploration.

Robots interact locally through infrared communication, broadcasting at regular intervals their knowledge about the position of the nest and of the known source, if available. Such information is used for two purposes. On the one hand, it is used by neighbours to update their own knowledge about the same locations, following the social odometry paradigm with weighted average aggregation mechanism (see Chapter 4 and Gutiérrez et al., 2010; Miletitch et al., 2013b, for details). This assures that the swarm maintains through time a good overall knowledge of the nest and source positions. When robots are located in the nest, the same message can lead to recruitment and cross-inhibition of uncommitted and committed robots, respectively. The uncommitted robot can get recruited with a probability $P_{R,i}$ by another robot committed to source i upon reception of a message. Similarly, a robot committed to source i can get inhibited with probability $P_{I,j}$—and turns uncommitted—upon reception of a message from a robot committed to source $j \neq i$ (see Fig. 5.2).

In this chapter, the probability of discovering a source results from the random exploration that robots perform, and is dependent on the distance d_i of source i from the nest: the closer the source, the higher the probability of discovering it. On the other hand, the other probabilities introduced are control parameters that determine the overall macroscopic behaviour of the robots. Here, we use fixed probabilities independent of the source quality, hence $P_{L,i} = P_L$, $P_{R,i} = P_R$ and $P_{I,i} = P_I$, and we perform a thorough analysis to uncover the effects of the control parameters on the emergent swarm behaviour.

As mentioned above, our goal is to study the macroscopic behaviour resulting from the rules defined in Sect. 5.1.1 for different values of the control parameters we identified. We want to obtain different types of macroscopic behaviour, from exploitation of a single good source to a load balancing between two different sources. Additionally, we want to maximise the exploitation efficiency of the robot swarm by optimising the rate of retrieved items, either from one or from multiple sources. To understand the effects of the different processes determining the collective dynamics, we performed a set of experiments to isolate the contribution of each component of the developed behaviour. In Sect. 5.2, we analyse the exploration efficiency when robots are uncommitted, while in Sect. 5.3 we focus on the exploitation efficiency when robots are committed to a given source. We analyse the effects of recruitment in determining the tradeoff between exploration and exploitation in Sect. 5.4. Then, we report on the effects of cross-inhibition on the ability of a swarm to converge to the exploitation of a single source or split among multiple ones, and we assess the exploitation efficiency of the swarm when dealing with multiple sources (see Sect. 5.5).

5.2 Baseline exploration efficiency

Resource exploration is the activity that agents perform when uncommitted. As mentioned above, robots perform a correlated random walk (Dimidov et al., 2016) until they find an item to be picked up. Random walk continues either until a source is found, or until the robot is triggered (with probability P_G) to return home. When in the nest, a robot can share information about the source found, or interact with other robots.

To evaluate the exploration efficiency of the swarm, we run a series of experiments to measure (i) the average rate of discovery of a source with respect to the distance, and (ii) the average percentage of robots that are found in the nest. The former gives us an idea of the probability of discovery P_D as a function of the distance of a source from the nest: the higher this probability, the sooner the swarm can start exploiting a source. The latter gives us an idea of the ability of robots to interact with each other when uncommitted, and has a bearing on the ability to be recruited by other robots to a known source. Both metrics depend on the probability P_G and on the distance d_i of source i from the nest.

Experimental Setup In this set of experiments, we use $N = 40$ robots that are constrained to remain in the uncommitted state: whenever a source is found, a robot goes back to the nest, but does not store the source location. In this way, when in the nest a robot starts again an exploration trip. We provide only one source to be found with high regeneration rate $r = 0.1\,\text{items/s}$, so that it remains close to the maximum number of items (i.e., 30 items). This source is placed at a fixed distance $d \in \{4, 6, 8, 10\}\,\text{m}$ from the nest. Additionally, we vary the probability to spontaneously go back to the nest $P_G \in \{0.0001, 0.0005, 0.001, 0.005, 0.01\}$. These values have been chosen to provide a sufficient exploration time before returning to the nest. Considering that the probabilistic choice is taken 10 times per second, the average exploration time corresponds to $(10P_G)^{-1}\,\text{s}$. We simulate the exploration for $T = 2000\,\text{s}$, and we measure the rate of discovery and the percentage of robots found in the nest, in average.

Results Figure 5.3 summarises the average results from 100 runs performed in each condition, varying source distance and probability to go back to the nest. It is possible to note an expected pattern for which, the higher the distance of the source, the smaller is the discovery rate (see the color shades of the different points). Similarly, the lower the probability P_G, the smaller the percentage of robots in the nest. The distance of the source also has an impact, although relatively small, on the percentage of robots found in the nest, because robots take less time to travel from the source to the nest once the source is found. Indeed, such a shift in the percentage of robots within the nest is visible especially for higher discovery rates.

Overall, we note that sources that are 10 m away from the nest are difficult to discover, and only small enough values for P_G ensure a non-null rate. However,

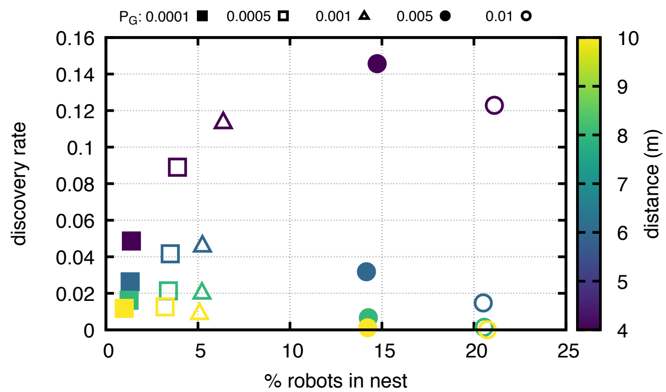

Figure 5.3: Exploration efficiency: We highlight the tradeoff between the average rate of discovery of a source and the average percentage of robots found within the nest. Lighter points correspond to larger distances, varying between 4 m and 10 m. Different point types correspond to different values of P_G.

P_G should not be too small, in order to grant a sufficient percentage of robots within the nest. A suitable tradeoff is given by $P_G = 0.001$, and we choose this value for the following experiments. With this value, the rate of discovery is non-null also for large distances, and at the same time the percentage of robots within the nest remains reasonably high, allowing for sufficient robot-robot interactions.

5.3 Baseline efficiency in source exploitation

Once the exploration efficiency of uncommitted robots has been determined, we can evaluate the exploitation efficiency of committed robots. This way, we study separately both sub-behaviors (committed and uncommitted) of the implementation. Generally speaking, we can define the exploitation efficiency of a swarm as the overall rate of retrieval of items, that is, the number of items retrieved by all robots per second independently of the source from which the items were collected. The retrieval rate will depend on the quality of the source, the distance from the nest, and the number of robots committed to the source and actively collecting items from it. Clearly, the way in which the individual behaviour is implemented could lead to interferences and congestion that have an impact on the overall retrieval rate. In order to evaluate the maximum efficiency of a swarm given the implemented behaviour (i.e., navigate back and forth from sources and pick up and deposit collected items, see Sect. 5.1.1), we perform a series of experiments largely varying the experimental conditions. We then introduce a model of exploitation of multiple sources which provides a baseline to evaluate the overall efficiency of the swarm when decision making and load balancing will be introduced.

Experimental Setup We consider the case in which a fixed number $N \in [1, 40]$ of robots exploit a single non-depletable source (i.e., a source with the maximum regeneration rate $r = 0.1$ items/s containing at most 30 items), placed at a fixed distance from the nest ($d \in \{6, 8\}$ m). We measure the retrieval rate of a group of robots that continuously exploit the source. To this end, we force robots to stay committed to the given source and we provide them with perfect information about the source location (i.e., robots never lose track of the source and can always navigate back and forth between source and nest). Under these conditions, we measure the overall rate of returned items per second once its evolution reaches a plateau.

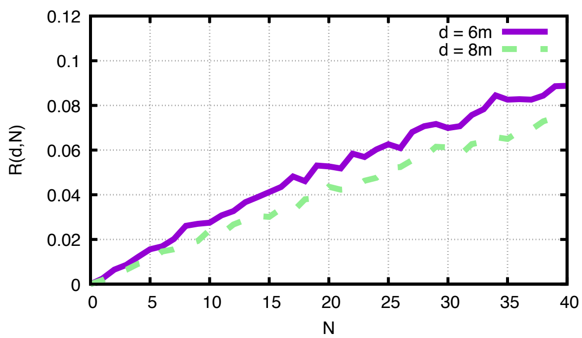

Figure 5.4: Exploitation rate $R(d, n)$, computed as the number of items per second retrieved from a non-depletable source at varying distance d and with varying group size N.

Results Figure 5.4 and Figure 5.5 summarise the average results obtained from 100 independent runs in each experimental condition, varying N and d_i. Each run lasts $T = 2000$ s during which robots continuously exploit the known source. In Fig. 5.4, the retrieval rate $R(d, N)$ is shown, indicating a linear dependency between group size N and rate of retrieval. This implies that, for the group sizes and distances considered, there is no negative impact from interferences or congestion, which would instead result in a sub-linear growth. We also note, as expected, that the efficiency is higher for closer sources, due to the fact that robots need to travel shorter distances. This will have an impact on the collective behaviour when multiple sources are presented at the same time. Indeed, the swarm may have to face the choice between exploiting a close-but-poor source, or a farther-but-rich one. The correct balance between the two must emerge from the different expected efficiencies in the exploitation.

To evaluate the efficiency in presence of two sources with different quality and distance, we provide a compact visualisation built on top of the maximum efficiency computed for single sources and fixed groups. We consider here a total group size of $N = 40$ robots, and we compute the expected efficiency for all possible allocations of robots to sources A and B, under the assumption that

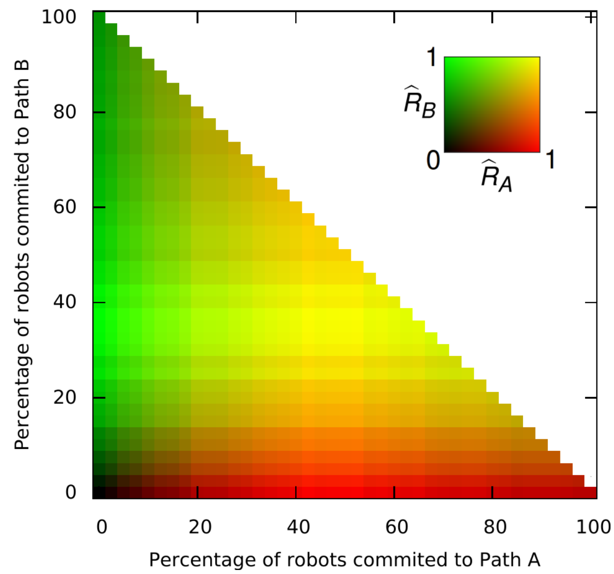

Figure 5.5: Combined efficiency in the exploitation of two sources with different quality and distance. In this example, we consider the case for $d_A = 6, d_B = 8, r_A = 0.03, r_B = 0.05$. The normalised exploitation efficiency $\hat{R}$ is represented in red for source A and in green for source B. The combination of both —as indicated by the inset— represents the total efficiency: dark zones imply low overall efficiency and bright ones imply high efficiency. Yellow zones (as a mix of green and red) imply that both zones are exploited in parallel.

$N = N_A + N_B + N_U$, where N_U is the number of uncommitted robots, which therefore do not contribute to the exploitation. Given the maximum retrieval rate $R(d, N)$, experimentally obtained for non-depletable sources at different distance shown in Fig. 5.4, we compute the normalised exploitation efficiency of a source of quality r as follows:

$$\hat{R}(d, N, r) = 1 - \left| \frac{R(d, N) - r}{R(d, N) + r} \right| \tag{5.1}$$

which has its maximum when $R(d, N) = r$, corresponding to a source that can completely support exploitation from N robots without being depleted. The normalised exploitation efficiency slightly decreases when $R(d, N) > r$, corresponding to the over-exploitation of the source, leading to complete depletion. We use this simple model to visualise the combined efficiency in foraging from two sources in parallel. We show the combined efficiency as an heatmap on a ternary plot (Fig. 5.5). Here, each point $\langle N_U, N_A, N_B \rangle$ corresponds to a given allocation of robots to the two available sources A and B. We color-code the normalised retrieval rate $\hat{R}_A$ in shades of red (see the horizontal axis for $N_B = 0$), while $\hat{R}_B$ is visualised in shades of green (see the vertical axis for $N_A = 0$). The combined efficiency is rendered as the sum of the two colours, hence bright yellow for the optimal values given the sources' quality and distance (see Fig. 5.5). This visualisation allows to indicate if a certain allocation of robots corresponds to the balanced exploitation of both sources, and will be used to evaluate the actual efficiency of the swarm when two sources are present, as discussed in Sect. 5.5.

5.4 Exploration vs. exploitation of a single source

The tradeoff between exploiting a given source and exploring in search of other possibilities is the result of a delicate balance of multiple forces, and needs to be carefully studied. With the implemented behaviour, robots commit to a source either upon discovery through random search, or upon recruitment from an already committed robot. While individual discoveries are always occurring with a constant probability, as shown in Sect. 5.2, the probability of a robot to be recruited grows with the size of the recruiting population. This creates a positive feedback loop for which, the more a swarm exploits a specific source, the more it recruits to it. On the other hand, source depletion following excessive foraging provides a negative feedback that works in the opposite direction and tends to stabilise the system, because those robots that do not find an item when they reach a depleted source turn uncommitted and stop recruiting once back to the nest. To evaluate the coupled effects from recruitment and over-exploitation, we analyse the dynamics of a swarm presented with a single source of varying quality.

Experimental Setup We consider the case of a single source of quality $r \in \{0.01, 0.03, 0.05, 0.1\}$ items/s, placed at distance $d = 8\,\mathrm{m}$ from the nest. Here, the number of items within a source can decrease so that sources can get depleted upon high exploitation. Robots execute the complete behaviour discussed in Sect. 5.1.1, although cross-inhibition is not present as there is only a single source. We consider a constant probability to go back to the nest $P_G = 0.001$ as resulting from the experiments discussed above. Coherently, we fix the probability of abandonment to $P_L = 0.001$ to have a similar rate of abandonment in both exploration and exploitation. When in the nest, robots can only recruit each other with a probability $P_R \in \{0.01, 0.02, 0.03\}$. Also in this case, every run is executed for $T = 2000\,\mathrm{s}$ using $N = 40$ robots, and we perform 100 experimental runs for each experimental condition.

Results To appreciate the macroscopic dynamics resulting from different parameterisations, we show the average percentage of robots committed to the source, the fraction of robots that switch commitment state per second, and the fraction of items left in the source (see Figure 5.6). When the quality is high ($r = 0.1$), nearly all robots get committed to the source and exploit it, and the number of items remaining in the source stays very high ($> 80\%$). Small recruitment probabilities correspond to a slow increase of the committed

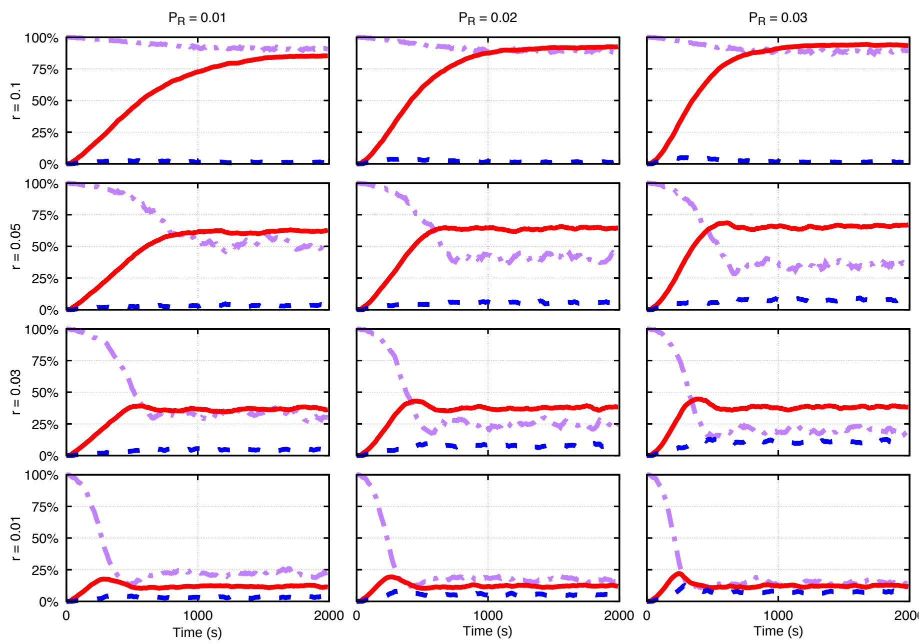

Figure 5.6: Dynamics of exploitation vs exploration with varying source quality and recruitment probability. Each graph shows the variation through time of the percentage of (i) committed robots (solid red lines), (ii) robots switching commitment state (dashed blue lines) and (iii) items left in the source (dot-dashed purple lines). Each graph represents a different setup with a distance of $d = 8\,\mathrm{m}$, recruitment probability $P_R \in \{0.01, 0.02, 0.03\}$ and source quality $r \in \{0.01, 0.03, 0.05, 0.1\}$

population, until a plateau is reached. The larger values we tested lead to a much quicker increase of the committed population, which stabilises earlier (see the top row in Fig. 5.6). Lower qualities of the source ($r \leq 0.05$) lead to a balance between positive and negative feedbacks that stabilises the committed population to a value that strongly depends on the source quality r, and to a much lesser extent also on the recruitment probability P_R. Interestingly, we also observe a higher rate of change in the commitment state in correspondence of higher values of P_R, which suggests that the macroscopic dynamics oscillate around the average values displayed in Fig. 5.6. Overall, the main impact of P_R is on the speed of growth of the committed population. Fast growth is useful for high-quality sources, but not so much for low quality ones, as the risk to quickly over-exploit the source may lead to fast depletion of the source and strong instabilities and oscillations due to massive abandonment. Hence, we consider a value of $P_R = 0.02$ as suitable for balancing quick growth with stability of exploitation.

5.5 Balancing source exploitation

Whenever two or more sources are present, robots need to choose which source to exploit, and a competition between the sub-populations committed to the one or the other source is observable due to committed robots recruiting uncommitted ones and cross-inhibiting each other. To evaluate the extent to which such competition leads to a balanced exploitation, we run a set of experiments to understand what is the average allocation of robots among committed populations as a function of varying sources' quality and distance, and for different values of the cross-inhibition probability P_I.

Experimental Setup In this experiment, we consider two available sources which can be at varying distance $d_i \in \{6, 8\}$ m and varying quality $r \in \{0.01, 0.03, 0.05, 0.1\}$ items/s. Robots execute the complete behaviour presented in Sect. 5.1.1, with $P_G = P_L = 0.001$ and $P_R = 0.02$, in accordance to the experiments presented above. Here, robots committed to different sources can cross-inhibit each other with probability $P_I \in \{0.01, 0.02, 0.03\}$. We perform 100 runs that last each $T = 2000$ s, and we look at the final allocation of robots committed to the different sources, or uncommitted. Additionally, we discuss the overall efficiency of the exploitation of the two sources in parallel, following the empirical model introduced in Sect. 5.3.

Results Overall, cross-inhibition defines how tolerant the swarm is of having a segmented population: the smaller the cross-inhibition probability, the lower the negative interaction between committed populations, the higher the probability that sub-populations committed to different source can coexist (see also similar results from the macroscopic models by Pais et al., 2013; Reina et al., 2017). The resulting dynamics can lead to convergence to a single source or balancing among many. This can be understood by looking at the distribution of the commitment state of the robots at the end of each run. The histogram in Fig. 5.7 represent such distributions, specifically for the percentage of robots committed to any source (that is, $\frac{N_A + N_B}{N}$) and for the percentage of committed robots that chose source A (that is, $\frac{N_A}{N_A + N_B}$). The former informs us about the ability of robots to successfully exploit sources, in average, given the experimental conditions. The latter informs us about the tendency of committed robots to choose source A (and conversely to not choose source B), hence revealing the collective choice or load balancing achieved by the swarm. In our experiments, we observe the full range of macroscopic dynamics for varying experimental conditions, which we discuss in the following (see Fig. 5.7).

Symmetric case, high-quality sources ($d_A = d_B = 6$ m, $r_A = r_B = 0.1$ items/s, first row in Fig. 5.7). In this condition, both sources could support the whole swarm, hence a collective decision in which the swarm achieves convergence on the exploitation of a single source may lead to the

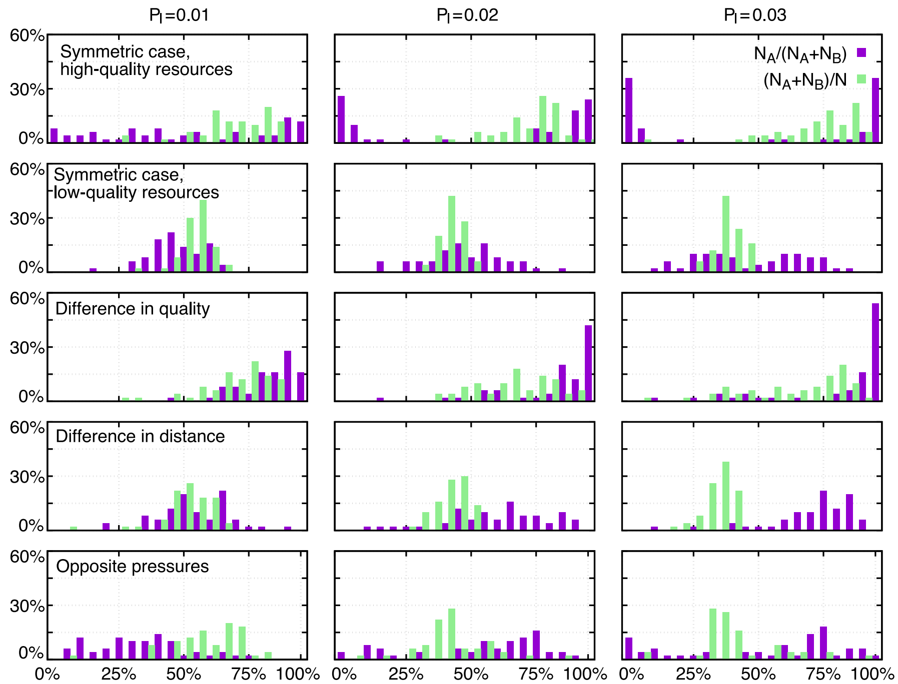

Figure 5.7: Allocation of robots to different sources. Each histogram represents the distribution of observables across $N = 100$ independent runs. Light green bars represent the percentage of committed robots $(N_A + N_B)/N$. Dark violet bars represent the percentage of committed robots that have chosen source A: $N_A/(N_A + N_B)$. Each plot represents a different setup, defined by the value of P_I and the quality and distance of the two sources. The value of P_I is reported on top of each column. For additional results in different experimental conditions, see the figures in the supplementary material. The experimental setup characterising each row is detailed as follow, from top to bottom: Symmetric case, high-quality sources ($d_A = d_B = 6\,\text{m}$, $r_A = r_B = 0.1\,\text{items/s}$); Symmetric case, low-quality sources ($d_A = d_B = 6\,\text{m}$, $r_A = r_B = 0.03\,\text{items/s}$); Difference in quality ($d_A = d_B = 6\,\text{m}$, $r_A = 0.1\,\text{items/s}$, $r_B = 0.03\,\text{items/s}$); Difference in distance ($d_A = 6\,\text{m}$, $d_B = 8\,\text{m}$, $r_A = r_B = 0.03\,\text{items/s}$); Opposite pressures] ($d_A = 6\,\text{m}$, $d_B = 8\,\text{m}$, $r_A = 0.03\,\text{items/s}$, $r_B = 0.1\,\text{items/s}$.

best results, as in this condition there are practically no robots uncommitted (i.e., light green bars are shifted to high percentages). We observe that low values of P_I do not grant convergence, resulting in a uniform repartition of robots over the two sources in the different runs (see the dark violet histograms). Higher values of $P_I \geq 0.02$ result instead in a collective decision, as observable from the bi-modal distribution of $N_A/(N_A + N_B)$, indicating full commitment for either A or B.

Symmetric case, low-quality sources ($d_A = d_B = 6\,\mathrm{m}$, $r_A = r_B = 0.03\,\mathrm{items/s}$, second row in Fig. 5.7). In this condition, no source can sustain the whole population, and we therefore observe a somewhat equal allocation of the committed robots among the two sources, especially for low cross-inhibition values ($P_I \leq 0.02$) where a unimodal distribution is present (see dark violet bars). Stronger cross-inhibition leads to a large competition and the appearance of a bi-modal distribution, although less pronounced than in the previous case. The number of uncommitted robots is in general high due to the low quality of the sources (corresponding to light green bars centred around low percentages of committed robots), and increases for larger values of P_I.

Difference in quality ($d_A = d_B = 6\,\mathrm{m}$, $r_A = 0.1\,\mathrm{items/s}$, $r_B = 0.03\,\mathrm{items/s}$, third row in Fig. 5.7). In this condition, sources differ only in the rate of replenishment, leading to a bias towards the choice of the most profitable one (A in this case, see the dark violet bars shifted towards high percentages). The higher the cross-inhibition probability, the stronger the shift of the distribution toward the high-quality source.

Difference in distance ($d_A = 6\,\mathrm{m}$, $d_B = 8\,\mathrm{m}$, $r_A = r_B = 0.03\,\mathrm{items/s}$, fourth row in Fig. 5.7). In this case, sources are both of somewhat low quality, but one is farther away than the other, leading to an exploitation balancing biased towards the closer source. We can observe here the high number of uncommitted robots, due to the low quality and the large distance of one source (see the distribution of committed robots centered around low percentages). The effects of the cross-inhibition probability are largely similar to the symmetric case with low-quality sources (second row), but the distribution is biased toward the closer source (A) especially for larger values of P_I.

Opposite pressures ($d_A = 6\,\mathrm{m}$, $d_B = 8\,\mathrm{m}$, $r_A = 0.03\,\mathrm{items/s}$, $r_B = 0.1\,\mathrm{items/s}$, last row in Fig. 5.7). This condition represents the most difficult case for the algorithm as the asymmetries in distance and quality oppose and may compensate each other. Indeed, the distribution of the committed robots (dark violet bars) is very wide, indicating that both sources are selected from time to time as they present advantages and disadvantages. For $P_I = 0.01$, source quality seems to matter, as the distribution is shifted towards the exploitation of source B. Higer values of the cross-inhibition probability lead to a larger number of runs ending with a balanced exploitation of both sources, with 50% to 75% of the committed agents exploiting source A and the remaining ones exploiting source B. Nevertheless, some runs end up with full commitment for the high quality source (B).

Overall, these results confirm that the implemented strategy for exploration and balanced exploitation results in expected distributions of robots among the available sources, giving preference to the most profitable one by allocating more

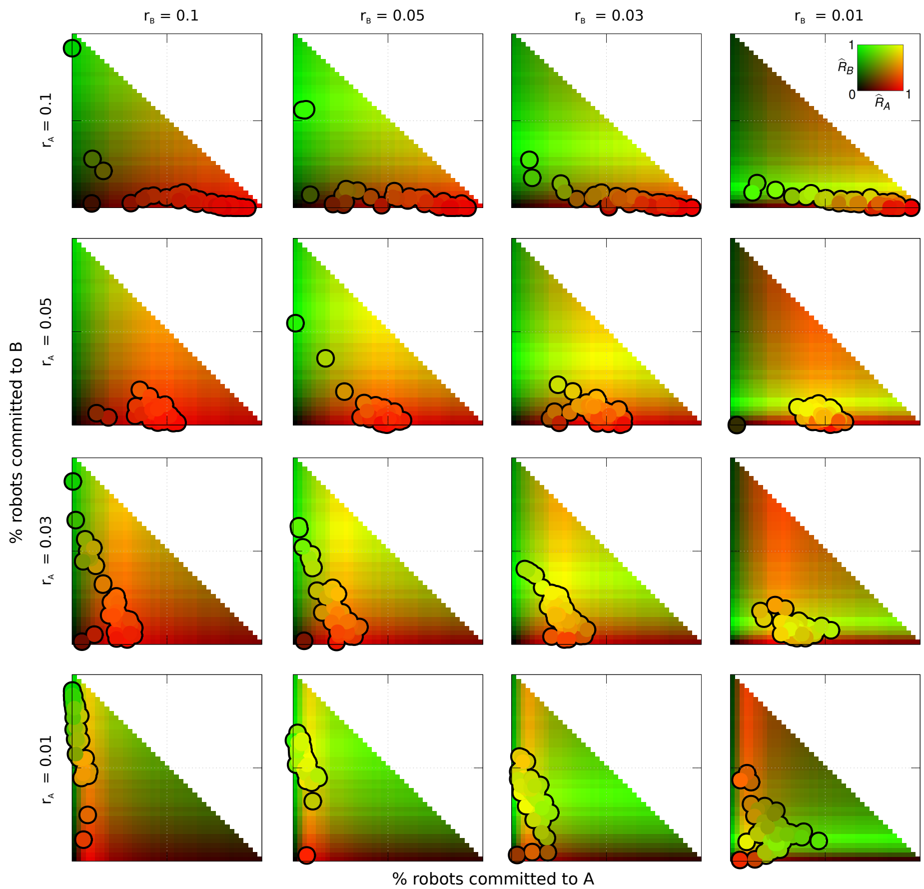

Figure 5.8: Efficiency in the exploitation of two sources, placed at different distances ($d_A = 6\,\mathrm{m}$, $d_B = 8\,\mathrm{m}$) and for all combination of qualities ($r_A, r_B \in \{0.01, 0.03, 0.05, 0.1\}$ items/s). Each panel shows a ternary plot where the underlying heatmap represents the theoretical efficiency in exploitation following the model introduced in Sect. 5.3, while the scatter plot represents the results of 100 experiments, each point indicating the final allocation of robots to different sources and the observed efficiency in exploitation, represented with the same colour-coding of the heatmap (see Sect. 5.3 for details).

robots on it. To evaluate the efficiency of the system, we compare the achieved allocation and retrieval rate with the model empirically obtained in Sect. 5.3. Figure 5.8 shows the heatmaps representing the ideal efficiency for the case of two sources with different distances ($d_A = 6\,\mathrm{m}, d_B = 8\,\mathrm{m}$), and for all possible combination of source qualities ($r_A, r_B \in \{0.01, 0.03, 0.05, 0.1\}$ items/s). On top of the heatmaps, we show a scatter plot corresponding to the results from 100 experimental runs obtained with $P_I = 0.02$. Each point corresponds to the final allocation of robots to committed or uncommitted populations, and the color corresponds to the experimentally observed efficiency, using the same color coding as for the heatmap. It is possible to note that the scatter plot generally matches the areas where efficiency is high or maximal, especially when the closer source is not the one with the highest quality ($r_A \neq 0.1$ items/s). Indeed, when A is also a high-quality source, the distribution is strongly biased towards its exploitation, because the source is discovered earlier and can sustain a large number of robots. Allocating other robots to the second source is therefore less probable. When B is the most profitable source, the distribution is biased towards its exploitation and is more balanced. Also note that the efficiency in the exploitation is matched between model and experiments, as the colors of the points in the scatter plot closely correspond to the underlying heatmap, hence confirming the suitability of the model we have introduced to evaluate the system efficiency.

5.6 Discussion and conclusions

In this chapter, we have implemented a strategy for exploration and balanced exploitation of renewable sources inspired by the honeybee value-sensitive decision making abilities. We have performed a large-scale simulation analysis to identify the effects of the different parameters governing the individual behaviour on the macroscopic, swarm-level dynamics. The results obtained confirm that our approach is suitable to provide the ability to adaptively balance exploitation of sources at the collective level, without requiring individuals to compare the profitability of different sources, and without a central planner with global knowledge of the environmental conditions.

The decentralised approach we adopt naturally leads to generalisations in the number of sources to be considered and different kinds of exploitation dynamics (as also studied in Reina et al., 2017; Miletitch et al., 2013b). Knowledge from modelling studies could be integrated in order to provide parameterisations suitable to deal with more complex working conditions, for instance dealing with a large number of sources in parallel (Reina et al., 2017). However, macroscopic models that consider at the same time the dynamics of the swarm and of the renewable sources are not available to date and require an important analytical effort. Work in this direction can provide means to obtain a precise micro-macro link between the robotics implementation and the modelling predictions, as obtained elsewhere for collective decision-making problems (Reina et al., 2015a,b). The study presented in this chapter can be considered a step towards the definition of a decentralised algorithm capable of optimally dealing with complex and dynamic environmental conditions. Thanks to a wide-ranging analysis of the parameter space, we have demonstrated that the macroscopic dynamics correspond to the expectations, opening up the possibility to develop swarm robotics solutions that appropriately balance the exploitation of sources.

A possible amelioration would be the further development of the proposed behaviour to obtain a more robust implementation. For instance, we found that fine tuning the probability P_G of returning to the nest is complex if one wants to deal with a large range of distances, as P_G strongly influences the average distance from the nest covered by robots while searching. To deal with a large range in the expected distances of sources, the exploration ability of robots should be changed, possibly exploiting recent results on the usage of Lévy walks, which are more suitable for searching in open environments in which the encounter of sources is an episodic event (Dimidov et al., 2016; Schroeder et al., 2017). Advancements in the exploration and exploitation abilities can be obtained also by allowing robots to share information widely while they move, instead of limiting interactions within the nest (Gutiérrez et al., 2010). To this end, it is necessary to understand how the mobility pattern of robots influences their network of interactions, and what is the bearing of an heterogeneous interaction network on the macroscopic dynamics. Additionally, it is worth considering the ability of agents to provide a motion bias to neighbours, thereby including in the study reinforced random walks (Perna and Latty, 2014; Schroeder et al., 2017). The characterisation of the interaction network resulting from the given mobility pattern can be done in terms of degree distribution and other properties relevant from a network theory point of view (Holme and Saramäki, 2012), while the macroscopic analysis of the effects of the interaction topology on the collective outcome needs to take into account heterogeneous mean field approximations (Moretti et al., 2013).

The way in which different sources are taken into account within the swarm is also worth further investigation. In this work, we limited robots to store only one source location at the time, therefore constraining the space of possible actions. Different experimental and modelling studies include additional mechanisms that better exploit individual knowledge, such as keeping memory of multiple sources and revisiting previously depleted ones (Dornhaus et al., 2006; Bailis et al., 2010; Granovskiy et al., 2012). Thanks to these mechanisms, higher adaptability is possible against variable environmental conditions.

Another direction to explore would be the study of the possibility of preserving information of known sources even when robots are uncommitted, also linking this aspect with the possibility to develop an emergent language used by the robots to refer to each different option. This is useful especially in case the number and position of sources is not known a priori, so that arriving at consensus on a single label for each source could be useful to let the robot balance exploitation by interacting in terms of labels and associated features (e.g., estimated regeneration rate or profitability). An especially interesting aspect, then, is to study is the potential interactions between exploitation dynamics and language dynamics (Steels and Belpaeme, 2005; Loreto et al., 2011), which can lead to synergies between the two processes if these are designed in the correct way (e.g., assign a different label only to the most profitable source, and the same label to those sources that should not be considered by the swarm). This topic is approached in the next chapter, in which we add a language game on top of the NSS algorithm studied in this chapter, aiming for the robots to evolve a language grounded on their experience (Spranger, 2013).

Chapter 6

Emergent naming conventions in a foraging robot swarm

In this chapter, we relax the assumption of a simple communication system that informed the two previous chapters, and link the foraging task to an emergent language, with the goal of describing the environment in which foraging takes place. Hence, this chapter provides a further demonstration of the interrelation between communication and behaviour, this time by taking the perspective of how language can evolve to represent features relevant for the behaviour execution. Specifically, our aim here is to demonstrate how language games can be grounded onto the execution of a task useful for the swarm. In particular, we introduce two language games, namely the Minimal Language Game and the Category Game (Miletitch et al., 2022), performed by the robots on top of the NSS algorithm introduced in the previous chapter. We present how varying and evolving topologies influence the way language games are played, and how this affects the convergence of the vocabulary within the swarm and its subgroups. Additionally, we study how language games can be used to provide a correct and complete description of the swarm's environment.

In Section 6.1, we discuss how language games can be meaningfully played by a robot swarm engaged in a source exploitation task. In Section 6.2 we present the experimental setup. In Section 6.3, we show how the dynamics of the interaction network can lead to emergent linguistic conventions that correlate with the actual task execution, but only if the language game is played by robots actually exploiting a source. Then, in Section 6.4, we analyse the properties of the interaction network, suggesting that it meaningfully supports the evolution of useful linguistic conventions. Finally, in Section 6.5 we exploit the gained understanding to define a category game tailored to better represent the different sources distributed in space, as long as these are relevant to the foraging task. We discuss the relevance of this study in Section 6.6, which concludes the chapter.

6.1 Language games in a foraging robot swarm

A popular approach to the study of language dynamics is represented by language games played by a population of agents/robots, with the purpose of mimicking real-world linguistic interactions leading to the emergence of a structured language. Various kinds of language games have been proposed to date, from imitation games (Billard and Hayes, 1997) to guessing games (Steels, 2001) and category games (Puglisi et al., 2008; Baronchelli et al., 2010). One game in particular has received a lot of attention: the *naming game* (Steels, 1995, 2003). In this game, two or more robots interact to assign a unique name to a set of objects. At each interaction, one robot is chosen as a speaker and another as a listener. The speaker chooses a referring object and an associated word from its vocabulary—or invents one when no word is available—and then transmits it to the listener. If the listener knows the word, then the game is a success, and both agents remove all other words associated to the chosen object from their vocabulary, keeping only the shared word. If instead the listener does not know the received word, then the game fails, and the listener adds this new word to its vocabulary. We use in our study a specific version of this game: the minimal naming game (MNG, see Baronchelli et al., 2006b). Here, focus is given only to reaching consensus on a single world within a population of communicating agents. Specifically, we consider an implementation in which the speaker broadcasts its word to all agents in his neighbourhood, while the listener is the only agent that updates the vocabulary upon success or failure of a game (Baronchelli, 2011).

As naming games are based on interactions between pairs of speaker and listener agents, the time to achieve consensus and the underlying dynamics are directly linked to the topology of the interaction network. In non-embodied implementations, the link between topology and language dynamics have been extensively studied (e.g., fully-connected regular, small-world or random geometric networks, see Baronchelli et al., 2007; Lu et al., 2008). Embodied implementations can be divided in two cases. On the one hand, a population of virtual agents can use a small number of robots (sometimes reduced to two, as in Spranger, 2013) to play the naming game, so that at each iteration, agents are selected and assigned to robots in order to record physical interactions among them. On the other hand, the naming game can be played among a population of embodied mobile agents (Baronchelli and Díaz-Guilera, 2012; Trianni et al., 2016b) that interact locally with each other according to a topology of interactions that is the direct result of the mobility pattern of the agents. This may lead to a strong heterogeneity among agents' neighbourhoods—and therefore different influences on the language game outcomes—which is crucially determined by the tasks that the agents/robots are engaged in.

In this study, the MNG is played on top of a self-organised foraging task. When foraging, a swarm needs to explore the environment, identify and evaluate the available sources and make decisions on which source to exploit, going through different transitory states before reaching an equilibrium (e.g., convergence on one single source to exploit or split/load-balance among many, as in Miletitch et al., 2018). Similar behaviours provide a complex and time varying interaction network among robots, which can be exploited to support linguistic interactions among agents. Our main goal is to study whether the dynamics of the interaction network are sufficient to determine language dynamics that represent features of the task execution (e.g., choice of one or the other source), of the environment (e.g., the presence of more than one sources, each associated to a different word), or both. To this end, we run experiments with two versions of the MNG. Beside the classic MNG, we play a version where the creation of words is linked with the discovery of sources by exploring robots. In this setup, we study how well the robots manage to have an accurate description of their surroundings, that is both complete (a word for each source) and correct (no misnomer) for as long as each source is relevant to the swarm, where relevance is measured as the number of robots actively foraging from the source (see Section 6.2). Our goal is to understand how the swarm interaction topology influences the language dynamics, and how the creation of words is correlated with the robots foraging from a source.

6.2 Experimental setup

In this chapter, we study a foraging task similar to the one presented in Chapter 4. Instead of being defined as a spread of physical items, the sources are now represented as circular colored areas on the floor. This allows us to focus on the language dynamics without interference from the need to identify and grasp an object to retrieve, hence simulating infinite quality sources. We also simulate a swarm 50 e-puck robots instead of marXbots, as grasping of objects is not required, allowing us to exploit a previously developed foraging and decision-making behaviour (Reina et al., 2015a). Here, the goal is to play a MNG while identifying and exploiting either of two sources (referred to as source A and source B) placed at the opposite side of a home area (referred to as nest, see Figure 6.1). The environment is a 2D infinite plane without obstacles, and both nest and sources have circular shape with radius $R = 0.3\,\text{m}$. Each source is located at the same distance $d = 2.5\,\text{m}$ from the nest.

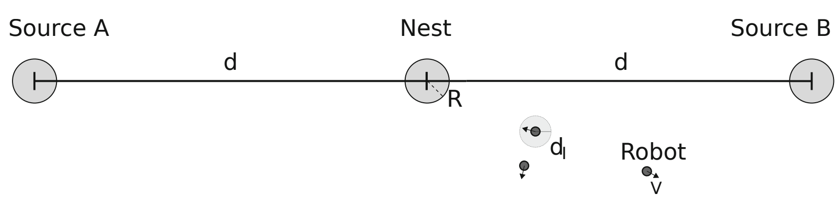

Figure 6.1: Graphical representation of the environment. sources A and B are each located at the same distance $d = 2.5\,\text{m}$ from the nest. All the three areas have radius $R = 0.3\,\text{m}$. Robots move at constant speed $v = 0.1\,\text{ms}^{-1}$ and can communicate with neighbours within a range $d_I = 0.2\,\text{m}$.

At the beginning of the experiment, robots are uniformly distributed within a 0.8 m side square centered on the nest. During the first 200 s, robots perform a blind random walk during which they do not communicate or search for sources. This allows us to neglect the initial transitory phase in which robots are too densely distributed around the nest, allowing us to study the system dynamics after the robots spread out in the environment according to their search pattern. This assures that—whatever the experimental condition—the initial distribution of robots does not severely impact the final outcome. In the following experiments, unless mentioned otherwise, we perform 100 runs for each experimental setup. These runs last until language convergence, which, depending on internal parameters, can take up to 12000 s.

6.2.1 Individual and collective behaviour

Source exploitation

The desired swarm behaviour (localization and exploitation of sources) takes inspiration from the decision-making process displayed by house-hunting honeybees—also know as nest-site selection (NSS, see Pais et al., 2013; Seeley et al., 2012a; Reina et al., 2017). The spatial dynamics during foraging resulting from the NSS process have been studied by Reina et al. (2015a) and Miletitch et al. (2018). Here, we make use of the individual robot behaviour from the former (Reina et al., 2015a), which was designed for the e-puck robots following a design pattern based on the NSS process (Reina et al., 2015b). According to this design pattern, a robot is considered to be committed to a source when it knows its location, and hence moves back and forth between the source and the nest. Otherwise, a robot is considered uncommitted and explores the arena searching for a source. Robots committed to source A (B) are considered to belong to the population $\mathcal{P}_A$ ($\mathcal{P}_B$), while uncommitted robots belong to the population $\mathcal{P}_U$, all summing up to N robots: $|\mathcal{P}_A| + |\mathcal{P}_B| + |\mathcal{P}_U| = N$.

Similarly to the behaviour discussed in Chapter 5, four concurrent processes determine the individual behaviour, two for transitions between uncommitted and committed states, and two for the opposite. An uncommitted robot turns committed either through **discovery** or through **recruitment**. The former takes place when the robot enters the area of a source. The latter takes place with probability P_ρ when a robot receives the information about a source known by a committed neighbour. Conversely, a committed robot turns uncommitted either through **abandonment** or through **cross-inhibition**. The former takes place anytime with a fixed probability P_α per time-step. The latter takes place with probability P_σ upon iteraction with a neighbouring robot committed to a different source. Cross-inhibition introduces a negative feedback loop that helps the system break the symmetry and leads to a choice between two identical sources (see Reina et al., 2015a,b, for more details). In our study, recruitment and cross-inhibition happen only upon communication with other robots when located into the nest. Differently from Reina et al. (2015a), we set the probability of abandonment P_α to zero, so that the only way for robots to become uncommitted is through cross-inhibition. This favours the attainment of a consensus state in

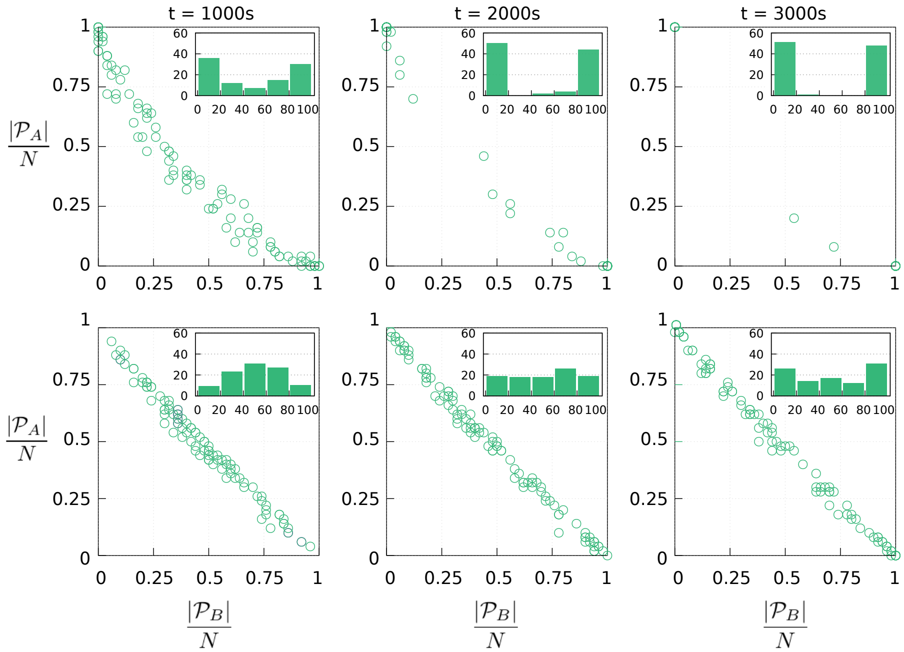

Figure 6.2: Distribution of robots in a swarm as a percentage of robots committed to source A (y axis) and B (x axis) for 100 independent runs. Each column displays the distribution at different time steps. The insets show the histogram of the frequencies of runs with respect to the percentage of robots committed to A. Top row: **strong cross-inhibition** with $P_\rho = 0.7$ and $P_\sigma = 0.7$, robots can change commitment and eventually the swarm converges toward either source A or B. Bottom row: **weak cross-inhibition** with $P_\rho = 0.7$ and $P_\sigma = 0.1$, the dynamic is much slower. Over the duration of our experiments, each run ends up with a different distribution of robots among sources, with points close to the diagonal representing low number of uncommitted robots.

which all robots within the swarm are committed to the one or the other source (Reina et al., 2015b).

The actual movements of the robot are governed by the following basic behaviours. When uncommitted, the robots explore the arena, performing a correlated random walk (Dimidov et al., 2016), and have a fixed and small probability at every control step to return to the nest. When committed, the robots enter an exploitation loop where they move back and forth between the known source and the nest (see Reina et al., 2015a, for a detailed description).

Depending on the value of P_ρ and P_σ, the swarm displays different dynamics and different final distributions of robots among the populations $\mathcal{P}_U$, $\mathcal{P}_A$ and $\mathcal{P}_B$. In this study, we focus on two specific cases: **strong cross-inhibition** and **weak cross-inhibition**. In the strong case ($P_\sigma = 0.7$, Figure 6.2 top row) the swarm rapidly converges to a consensus for the one or the other source, whereas the weak case ($P_\sigma = 0.1$, see Figure 6.2 bottom row) leads to slower dynamics (Reina et al., 2016). Given enough time the swarm would end up converging to a consensus for a single source. However, over the duration of our experiments, the swarm did not break the symmetry but splits between the two sources (see Figure 6.2, bottom row). At any time, with or without consensus, we define the source with the highest number of committed robots (relative majority) as the "selected" source. We define $O \in \{A, B\}$ as the selected source and $X \in \{A, B\}$ as the non-selected source, and $\mathcal{P}_O$ and $\mathcal{P}_X$ as the respective populations, with $\mathcal{P}_O \geq \mathcal{P}_X$.

Minimal naming game

The language game played by the robots in our study is an implementation of the minimal naming game (MNG) for mobile agents/robots (Baronchelli et al., 2006b; Baronchelli and Díaz-Guilera, 2012; Trianni et al., 2016b). Each robot starts with an empty inventory. At each time step (of length $\tau_c = 100\,\text{ms}$), each robot has a probability P_s of becoming a speaker (here, $P_s \in \{0.0003, 0.0006, 0.001, 0.002\}$). These values of P_s were selected so that foraging dynamics and language dynamics would share comparable time scales. The language game is played as follows: the speaker robot selects a word from its inventory and broadcasts it to its neighbours. At each time step, if a robot receives at least one message, it becomes a hearer robot. The hearer selects one (and only one) word at random among those received and checks it against its own inventory. If the hearer finds the selected word in its inventory, the hearer keeps only that word in the inventory while deleting all the others. If instead the hearer does not find the selected word in its inventory, it updates its inventory by adding the word (see Trianni et al., 2016a, for more details).

In this study, we consider two variants of the MNG, which differ in the way in which words are generated. In one case (referred to as **classic game**), the robots create a new word when becoming speaker with an empty vocabulary. In the other (referred to as **spatial game**), the robots create a new word when encountering a source with an empty vocabulary. In both cases, we associate each word with the closest source to the robot at the time of the word creation, and we define W_A (W_B) the set of words associated with source A (B). Note that, by construction, $W_A \cap W_B = \emptyset$. Robots having in their inventory any word $w \in W_A$ (W_B) constitute population $\mathcal{P}_{W_A}$ ($\mathcal{P}_{W_B}$). Robots with no words constitute population $\mathcal{P}_{W_O}$. In Figure 6.3, we depict a possible partition of robots among different populations, both with respect to the commitment state and to their vocabulary. Since a robot can have at a given time an inventory with words originating in both source A and B, the propriety $\mathcal{P}_{W_A} \cap \mathcal{P}_{W_B} = \emptyset$ is not always verified. Similarly, through exchanges of words and robots between the different populations, at a given time the inventory of robots committed to one source might contain a word associated with the other source (resulting in $\mathcal{P}_A \neq \mathcal{P}_{W_A}$). At any time, we can look at the population of robots that know words associated with the source they are committed to, that is:

$$\mathcal{P}_M = (\mathcal{P}_{W_A} \cap \mathcal{P}_A) \cup (\mathcal{P}_{W_B} \cap \mathcal{P}_B). \tag{6.1}$$

Conversely, we can define the population of committed robots that know words from a non-matching source:

$$\mathcal{P}_S = (\mathcal{P}_A \cap \mathcal{P}_{W_B}) \cup (\mathcal{P}_B \cap \mathcal{P}_{W_A}). \tag{6.2}$$

Corresponding to the collectively selected source O (see definition above), we define the set of matching words W_O and non-matching words W_X as follows:

$$W_O = \{w|(w \in W_A \wedge \mathcal{P}_A > \mathcal{P}_B) \vee (w \in W_B \wedge \mathcal{P}_B > \mathcal{P}_A)\} \tag{6.3}$$

$$W_X = \{w|(w \in W_A \wedge \mathcal{P}_B > \mathcal{P}_A) \vee (w \in W_B \wedge \mathcal{P}_A > \mathcal{P}_B)\} \tag{6.4}$$

We define:

- **polarisation**, the condition in which committed robots know only words associated with the source they are committed to, that is, when $\mathcal{P}_S = \emptyset$;
- **vocabulary matching**, the condition in which only words associated with the selected source are retained within the swarm vocabulary, that is $W_X = \emptyset$ and $W_O \neq \emptyset$;
- **vocabulary completeness**, the condition in which exactly one word associated with each source is retained within the swarm vocabulary, that is $|W_O| = 1$ and $|W_X| = 1$.

Given a sufficiently connected swarm, the MNG dynamics ensure that the swarm will eventually converge to a final single-word vocabulary, albeit after a very long time (Baronchelli et al., 2006b; Baronchelli and Díaz-Guilera, 2012; Trianni et al., 2016a). According to the previous definitions, the final vocabulary can be matching or not the selected source.

6.3 Correctness and completeness of the swarm vocabulary

In this section, we focus on the evolution of the swarm's vocabulary, looking in particular to the provenance of the last words and their relation to the selected source. As already discussed (see Figure 6.2), the foraging dynamics lead to either the quick selection of a single source, or to the swarm being split between the two sources, possibly for a long time. This means that, apart for a few cases and random fluctuations, there will always be a source that is selected—albeit temporarily—by the swarm. In certain settings, the swarm may forage from both sources for a long time, hence vocabulary completeness may be observed. In other cases, the swarm will quickly converge to exploit a single source, and vocabulary matching is expected. In any case, interactions between different

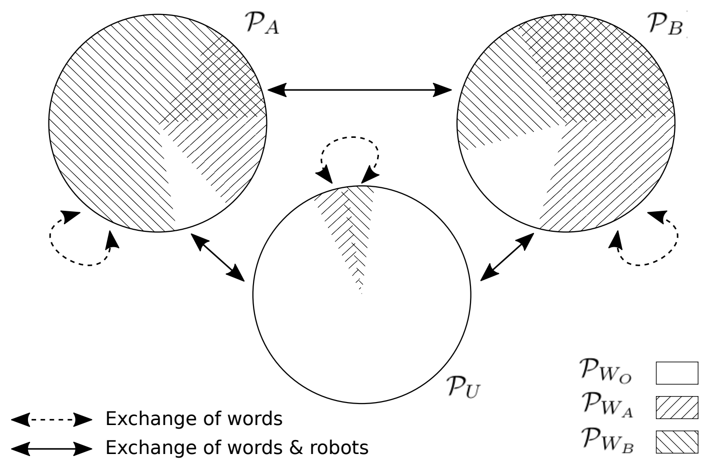

Figure 6.3: Diagram representing how the swarm can be split in different sub-populations with respect to the robots' commitment state and the word distribution. The circles represents the three populations with respect to the commitment state: ($\mathcal{P}_U$, $\mathcal{P}_A$ and $\mathcal{P}_B$). The fill patterns represent populations with respect to the robots' inventory ($\mathcal{P}_{W_O}$, $\mathcal{P}_{W_A}$ and $\mathcal{P}_{W_B}$). Note that, in general, $\mathcal{P}_{W_A} \cap \mathcal{P}_{W_B} \neq \emptyset$. Depending on the experimental setup, populations can exchange robots and words among themselves.

populations of robots are frequent, ensuring that the language dynamics always converge to a single-word vocabulary.

The complex interplay between foraging and language dynamics makes it difficult to observe a clear emergence of vocabulary matching or completeness during a run. It is possible that matching or completeness is achieved at some point, but the frequent interactions among sub-populations through the exchange of robots and words (as depicted in Figure 6.3) make the analysis of the transitory phases complex. Hence, we first focus on the patterns observed when the vocabulary converges to one or two words, to determine if matching and completeness are achieved. First, we analyse the provenance of the final word w_f to determine if it matches the selected source or not (i.e., $w_f \in W_O$). As the distribution of robots among sub-populations may sometimes change even after convergence to a single-word dictionary (e.g., if the language dynamics are much faster than the source selection dynamics), the final selected source may also change. Hence, we consider the source selected at the time of convergence to the final word w_f, no matter what happens later to the population distribution. Similarly, we consider also the second-last word w_e, to determine whether it was also matching the selected source or not at the time in which only two words remained within the whole swarm. Given such definitions, every run can end up

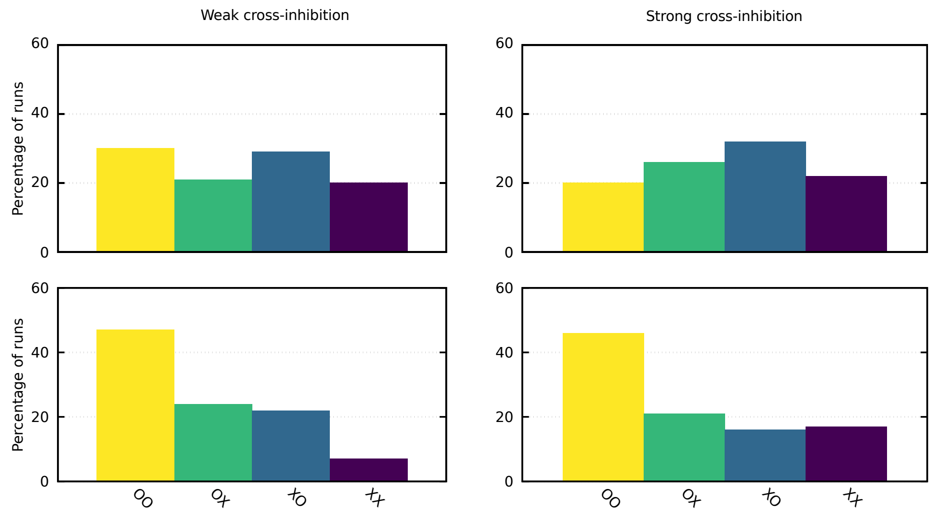

Figure 6.4: Empirical distribution over 100 runs of the occurrences of the last two words in the vocabulary within the four identified classes (OO, OX, XO and XX) representing words matching or not the selected source. The graph refers to the case with $P_s = 0.001$. All other tested values of P_s produce similar results. Top row: classic game. Bottom row: spatial game.

in one of the following four possibilities:

$$\text{OO} : w_f \in W_O \wedge w_e \in W_O \tag{6.5}$$
$$\text{OX} : w_f \in W_O \wedge w_e \in W_X \tag{6.6}$$
$$\text{XO} : w_f \in W_X \wedge w_e \in W_O \tag{6.7}$$
$$\text{XX} : w_f \in W_X \wedge w_e \in W_X \tag{6.8}$$

In case OO or OX is observed, the swarm has identified a final word that matches the currently-selected source, although in the OX case the second-last word was associated with the non-selected source. The XO case represents a missed opportunity of matching, as a matching word was still existing in the vocabulary and could have been chosen. The XX case instead suggests that the association of words to source does not reflect the current state of the source selection. Both middle cases (OX and XO) indicate a complete vocabulary up until convergence on one word.

Given these definitions, we study the influence of the language game and the foraging dynamics over the provenance of the last two words of the vocabulary. Figure 6.4 shows the frequency of each case out of the 100 runs performed for each different experimental condition. When playing the classic game (top row in Figure 6.4), the swarm shows no tendency to favor a specific provenance for the final two words, and a distribution close to uniform across the four possible cases is observed. On the other hand, when playing the spatial game (bottom row in Figure 6.4), the swarm favours words that match the selected source, both for the last and second-last word. In particular, the OO state is strongly favoured for both weak and strong cross-inhibition, and the XX state is especially disfavoured when the weak cross-inhibition leads to slower decision dynamics. In conclusion, we clearly find that the spatial game, by making the creation of words conditional to the discovery of sources, determines a strong tendency to converge towards words that represent the source that is ultimately selected. The naming process is "correct" as it best represents the source that is the most relevant for the swarm. In about 40% of the cases (OX + XO), the naming is "complete" as the last two words represent "names" for both the available sources. This remains valid for different values of the probability of speaking P_s, suggesting that the spatial game is resilient to variations in the timescale of the language game.

To better understand the relationship between source selection and naming dynamics, in Figure 6.5 we show how the distribution of agents between sources relates with the provenance of the last two words in the swarm vocabulary. Indeed, there is a large difference between a swarm that forages from a single source and one that instead is evenly split between the two sources. In the former, we expect vocabulary matching, that is, only words from the selected source are retained (hence, case OO and to some extent OX). In the latter, we instead expect vocabulary completeness, that is, words coming from both sources are present (hence, cases OX and XO) because both sources are still exploited by the swarm and the selected source can change over time. Indeed, the swarm does not clearly favor the exploitation of any source, to the point of possibly changing its selected source overtime, and multiple times.[1]

When the classic game is played, the distribution of robots across sources has little to no impact on the provenance of the last two words (top row of Figure 6.5). For the spatial game, instead, vocabulary matching is observed when the swarm has clearly selected one of the sources. Conversely, vocabulary completeness is more often observed with swarms that are still exploiting two sources. This is evident in case of weak cross-inhibition that entails slower dynamics in the source selection process. With strong cross-inhibition, the swarm quickly converges to exploiting a single source, and the cases in which the swarm is exploiting both sources at the time of convergence are very rare. Only when the language dynamics are particularly fast we can observe cases of vocabulary completion for strong cross-inhibition.

From this analysis we can conclude that the spatial game leads to language dynamics that correctly represent the sources relevant to the swarm, that is, those from which the swarm is currently foraging. This is obtained solely by the creation of words, which is strongly correlated with the source discovery. The interplay between language and foraging dynamics preserves such correlation

[1] Recall that the distribution of robots can change over time, and always converges to the selection of one source, although after a very long time as discussed in Section 6.2.1. Here, we consider the distribution at the time of convergence of the naming dynamics, which is determined by the probability of speaking P_s. Hence, an even distribution of robots among the sources is observable not only with weak cross-inhibition ($P_\rho = 0.1$, see Figure 6.2), but also for strong cross-inhibition when high values of P_s cause a quick convergence of the vocabulary.

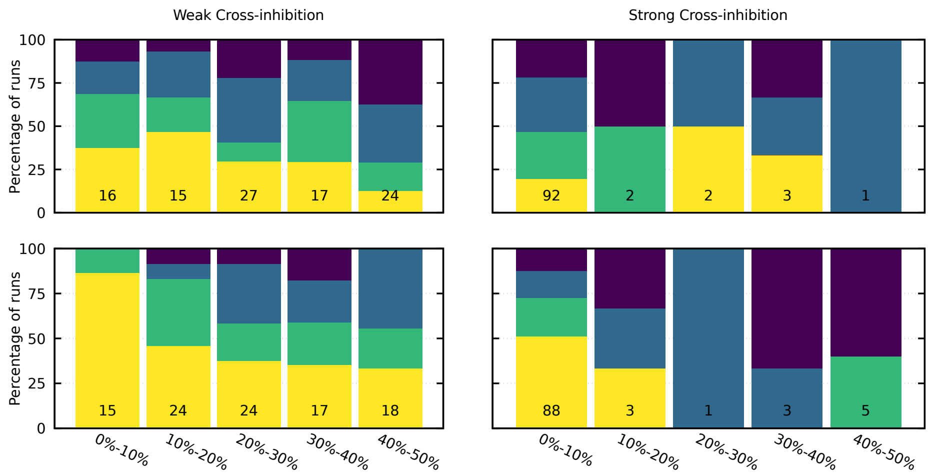

Figure 6.5: Empirical distribution over 100 runs of the occurrence of the last two words in the vocabulary (see Figure 6.4) detailed for different distribution of the foraging swarm across the two sources, computed at the time of vocabulary convergence with $P_s = 0.001$. Each stacked histogram corresponds to a specific distribution of robots over the non-selected source ($\frac{\mathcal{P}_X}{\mathcal{P}_O + \mathcal{P}_X}$). Bars are colour-coded as in Figure 6.4. Over each histogram, the number of runs that resulted in the specified range is displayed. All tested values of P_s present similar results, shown in Figure S2. In the rare case of an equally split swarm ($\mathcal{P}_O = \mathcal{P}_X$), there is no notion of matching an non-matching words. In that case, we redistribute AA and BB equally between OO and XX (one half each). Similarly, AB and BA are redistributed equally to OX and XO. Top row: classic game. Bottom row: spatial game.

despite the high number of interactions between robots from different populations and with different vocabularies. In the next section, we study how this is possible by looking at the interaction patterns between robots.

6.4 A study of the swarm's spatial characteristics

There are two extremes for the swarm to reach convergence on a final word. Either the swarm converges as a whole—homogeneously—on this final word, or sub-populations foraging from different sources first converge toward a word representing their source, and then a competition between these two words determines the final outcome. The explanations of how a swarm's vocabulary converges, and to what states it converges, lay in the spatial characteristics of the swarm, a direct consequence of the exploitation task. In particular, we look at how robots create and share their words, and how they exchange words within

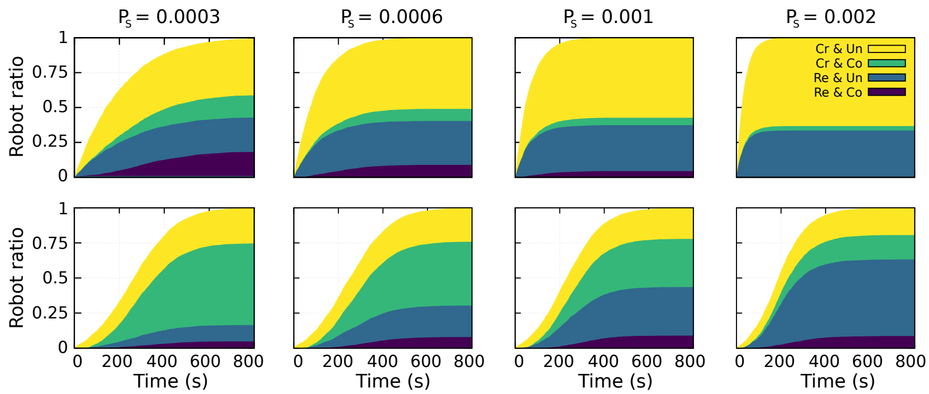

Figure 6.6: Evolution over time of the origin of each robot's first word (weak cross-inhibition). The value of the y axis correspond to the ratio of robots having a word in their vocabulary. This word can be either created independently by a robot (Cr) or received from another robot (Re); and either while the robot is uncommitted (Un) or committed (Co). Similar dynamics are displayed in the case of strong cross-inhibition. Top row: classic game. Bottom row: spatial game.

and across foraging sub-populations.

6.4.1 Impact of spatial word creation

First of all, we look at the initial phases of the naming game, when robots create and share new words. Indeed, the difference between the classic and the spatial game is solely related to this phase. Besides word creation, robots can fill their vocabulary with words shared by others. To better understand how robots obtain their first word, we plot in Figure 6.6 the cumulative number of robots with at least one word in their vocabulary for the case of weak cross-inhibition.[2] We highlight whether the first word was created by the robot itself or received from other robots upon playing the naming game. Finally, we distinguish between robots being uncommitted and exploring, or robots committed and exploiting one source. Uncommitted robots are particularly relevant, as they can get committed to any source, despite having a word associated with one or the other: they do carry a naming information that may not correspond to the source they will become committed to.

For the classic game (top row in Figure 6.6), we note that the word creation dynamics is rather fast and solely depends on the probability of speaking P_s. Additionally, uncommitted robots represent the large majority, meaning that word creation is strongly uncorrelated from source selection: even if a word is created closer to a source, it is generally associated to an uncommitted robot that

[2]Results for strong cross-inhibition are very similar and are displayed in Figure S3.

may eventually get committed to any source, due to recruitment or discovery.[3]

In the spatial game instead (see bottom row in Figure 6.6), the dynamics of word creation is independent of P_s because it is determined by robots encountering a source. Specifically, P_s does not impact the number of robots that create a word when uncommitted, as these robots individually discover a source following the foraging dynamics. However, P_s determines the share of robots that create a word when committed or that receive a word when uncommitted. The former is higher when P_s is small, as the foraging dynamics are faster than the language game dynamics, meaning that several robots get recruited first and encounter a source while still having an empty vocabulary. These robots have a naming information that is strongly correlated with the source they are exploiting. Conversely, with high P_s the number of uncommitted robots that receive a word from other robots grows. These robots potentially have a naming information that differs from the source they will exploit, leading to lower spatial correlation. As a matter of fact, matching and completeness are slightly worse for this case.

6.4.2 Communication topology and interactions within the swarm

Once words have been generated, the MNG imposes a selection process until a single one is selected. This process takes place through speaker-hearer interactions, and can be strongly influenced by the communication topology (Baronchelli et al., 2006a; Moretti et al., 2013). The latter is determined by the distribution of robots in space, which is a result of the foraging task the robots carry out. To understand how the different sub-populations of the swarm interact, we performed an experiment with locked-size populations, forcing all robots in a pre-defined committed state. We measure the size of the neighbourhood $\mathcal{N}$ with which robots can potentially interact anytime, and we further distinguish between neighbours belonging to the same or to a different population. In Figure 6.7, the probability of observing a neighbourhood of a given size is displayed for each possible partition $\mathcal{P}_s$ between sub-populations, where $\mathcal{P}_s = p$ indicates that $\mathcal{P}_A = p$ and $\mathcal{P}_B = N - p$ (in these tests, $\mathcal{P}_U = 0$). Additionally, we also consider the case in which for $\mathcal{P}_U = N$, where robots are forced in the random exploration state.

For small values of $\mathcal{P}_s$, one of the sub-populations is large and interactions within sub-population dominate (see Figure 6.7, left panels). The neighbourhood size can take large values (e.g., more than 5 robots), even larger than the case of randomly exploring robots (see Figure 6.7, bottom-left panel). Contrarily, interactions between sub-populations are practically absent, the typical neighbourhood size being $|\mathcal{N}| = 0$ (see Figure 6.7, top-right panel). The more the partition among sub-populations is even, the more frequent the interactions

[3] Recall that robots periodically return to the home location, where they can get recruited by any other robot, or they can start a new exploration trip in a totally different direction from the previous one. Hence, an uncommitted robot that creates a word near one source may get recruited to the other source or discover it in the following exploration trip.

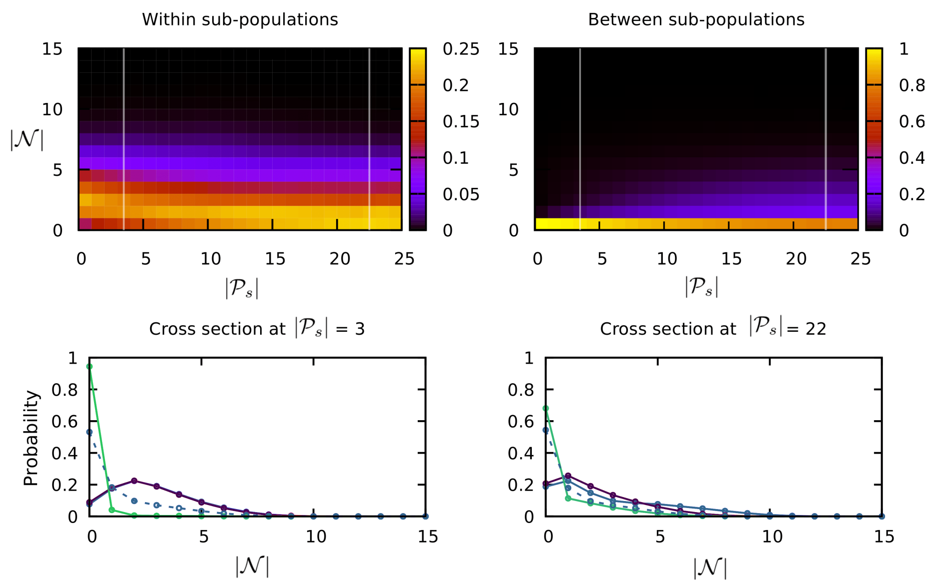

Figure 6.7: Top row: the heatmaps represent the probability distribution P_Σ of each robot's neighbourhood's size ($|\mathcal{N}|$, y axis) for each possible partition in sub-populations ($|\mathcal{P}_s|$, x axis), limited to interactions occurring within a sub-population (top left) or between sub-populations (top right). Vertical lines indicate the cross-sections displayed in the bottom panels. Bottom row: probability of occurrence of each robot's neighbourhood's size for $|\mathcal{P}_s| = 3$ (bottom left) and $|\mathcal{P}_s| = 22$ (bottom right). The plots represent the probability P_Σ of observing a neighbourhood size considering interactions within the whole swarm (blue), within sub-populations (purple), and between sub-populations (green). The dotted-blue line represents the case of the whole swarm forced to remain in the exploring state (sub-population $\mathcal{P}_U$).

among sub-populations become. Still, robots more likely interact within the same population, and only few cross-population interactions are recorded (see Figure 6.7, bottom-right panel). This confirms that, if the swarm leans towards selecting a single source, the language dynamics are played mostly within the same population, reinforcing the correlation between words and sources in favour of matching. At the same time, the small number of interactions between sub-populations also favour completeness, with each sub-population having the chance to converge on its own word.

It is worth recalling that, besides communications between sub-populations, a mismatching word can enter a sub-population also when it is physically carried by a robot changing from one to the other population. In order to understand how relevant the movements of robots between sub-populations are for the spreading of words, we measured the rate at which these movements take place, and compared it with the rates of interactions within and between

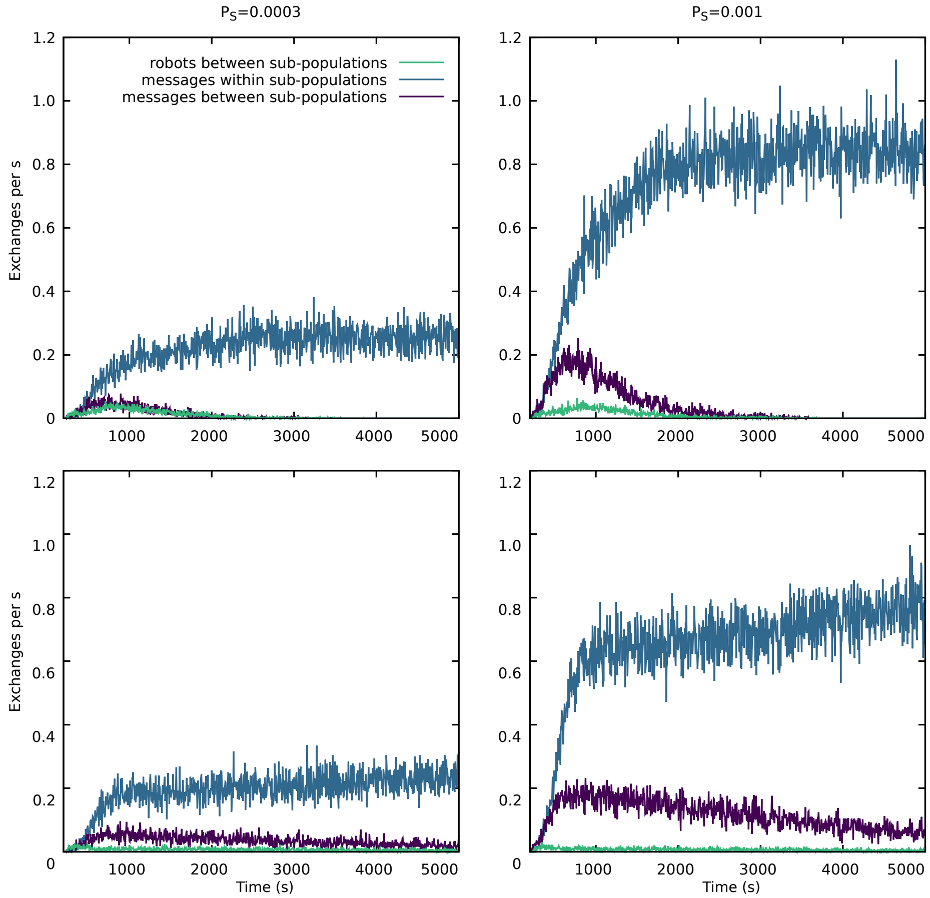

Figure 6.8: Evolution over time of the rate of communications within and between sub-populations exploiting different sources, and of the rate of robot movements between sub-populations. Each graph has been plotted for the spatial game. Similar dynamics are displayed by the classic game. Top row: strong cross-inhibition. Bottom row: weak cross-inhibition.

populations during a standard experiment (see Figure 6.8). The results indicate that movements between sub-populations are not as frequent as the interactions via message exchange, especially when the probability of speaking P_s is high (see also Figure S4). Indeed, the rate at which messages are exchanged within and between populations increases with P_s, and is generally larger for intra-population interactions, confirming our previous analysis. Conversely, the rate at which robots move from one population to the other does not depend on P_s, and is higher when cross-inhibition is strong. We infer that the movements of robots between sub-populations do not have a relevant impact on the language dynamics in this specific experimental setup.

In the light of the presented results, we can conclude that the pattern of interactions between robots favours the segregation between sub-populations. This means that different words are likely selected within each sub-population, resulting in the vocabulary completeness. At the same time, vocabulary matching is possible thanks to the strong correlation between word creation and source exploitation by committed robots, as discussed above. While the vocabularies well represent the environmental features and their relevance for the swarm, we note that completeness is a transient property. Indeed, the MNG dynamics determine the convergence towards a single word shared by the swarm, loosing information about previously exploited sources. To avoid this, we present in the next section a proof of concept of a language game to preserve matching and complete vocabularies.

6.5 Emergence of spatial categories for foraging swarms

With both the classic and spatial game, the swarm vocabulary always converges toward a single word, losing the completeness of information as one source is never represented. Keeping a complete description of the environment with all its sources requires the ability to distinguish between different regions in space, leading to the construction of spatial categories. We consider a spatial category as a set of possible words, associated to an area representing the region covered by the category (here, a circle defined by its radius and its center, the latter determining the prototype location of the spatial category). Speaking in general terms, any location in space can belong to one category, to multiple ones (in case of overlapping categories) or to none (in the case of a non exhaustive partition of the space). The same robot can potentially hold multiple words (synonyms) referring to a given category. As a consequence, the set of the categories known to a robot—and, by extension, to the swarm—results in a kind of thesaurus. In this section, we propose a language game based on word-location pairs with the goal of representing the landscape of available sources. The language game is now first played on categories and then on words, making it more similar to a category game (Baronchelli et al., 2010).This revised language game unfolds in two distinct phases. Initially, the game focuses on categories, determining the relevant spatial area under consideration. Subsequently, within the identified category, a minimal language game is played, making the overall process more similar to a category game (Baronchelli et al., 2010). This approach facilitates spatial segregation of the resulting vocabularies, with the aim of maintaining distinct naming conventions for different spatial areas.

6.5.1 Experimental SetupImplementation of the Category Game

Similarly to the spatial game discussed above, categories are spontaneously created when a robot encounters a source at a location that is not represented by any available category. Even if a category exists for the same source, a robot may enter from a location that is not covered by the current category description. This leads to an initial proliferation of categories, which are subsequently pruned by a merging mechanism (see below).

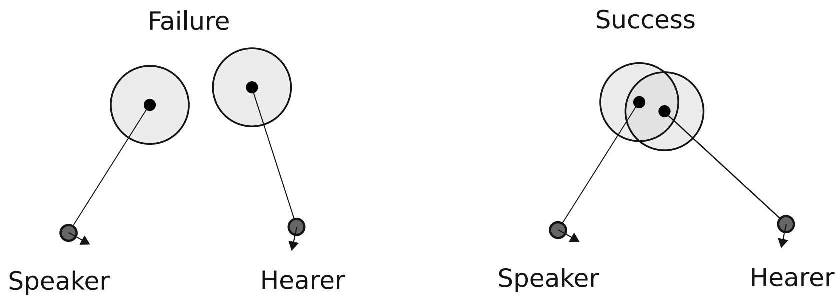

Figure 6.9: Graphical representation of a played category game. On the left, the speaker is sharing a category that is not know by the hearer, resulting in a failed game. On the right, the speaker is sharing a category that is in the range of one of the hearer categories, resulting in a successful game.

With probability P_s, a robot knowing at least one category takes the speaker role: it first selects one of its known categories, followed by a word belonging to this category's inventory. The speaker will share with the neighbours the selected word paired with the category prototype's location. In order to maintain a correspondence between the foraging behaviour and the language game, the selection of the category is determined by the commitment status: the speaker always selects the category corresponding to the sources it is foraging from. For uncommitted robots, the category is selected randomly. On the hearer side, first a match of the received word-location pair must be found with the known categories. If the location does not belong to any known category (left of Figure 6.9), the hearer creates a category centered on that location, with a default starting radius of $r_0 \in \{0.2, 0.3, 0.4\}$, and add the received word to this category. If the location belongs to only one category (right of Figure 6.9), the MNG is played as previously described (see Section 6.2.1) with respect to the matching category's inventory. If the word is fitting multiple categories, these are merged into one (see Figure 6.10), and then the MNG is played with respect to the resulting category's inventory. Categories are merged two by two, with the resulting category being the smallest possible circle containing each original category's circle. The merged vocabulary is the union of each category's vocabulary.

To evaluate the ability of the swarm to generate shared spatial categories that correctly represent the available source landscape, we performed a series of experiments varying both the probability of speaking P_s and the value of the initial category radius r_0. We introduce no change in the physical layout of the arena (see Figure 6.1). Experiments are run for longer times, and are stopped once convergence is reached on both categories and number of words in each

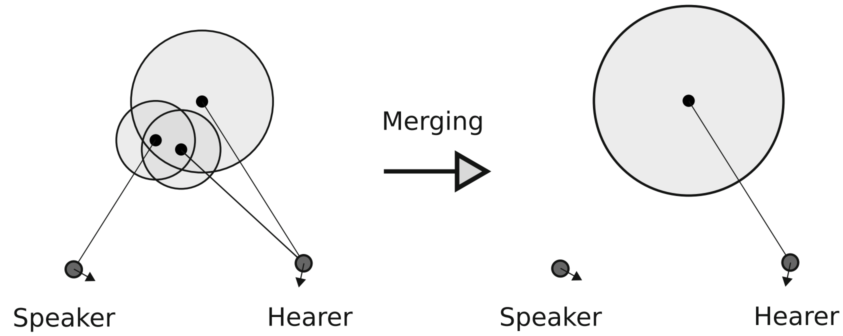

Figure 6.10: Graphical representation of the merging of categories. On the left, the hearer receive a word from the speaker that is fitting multiple categories. As a result, the hearer merges them into an updated category whose area encompasses all categories that were at play in this game.

category. The additional complexity introduced by categories entails a slower language dynamics with respect to the simple naming game described before. To study the ability of the foraging swarm to correctly represent both sources, we prevent the selection of a single one by forcing $P_\sigma = 0$. In this way, the robots will find and exploit both sources (possibly with an uneven distribution across the two), and no robot will ever change source. As we observed in Section 6.4, the effects on the language dynamics of robots physically moving from one to the other source are anyway negligible.

6.5.2 Results

The evolution over time of the number of words and of categories is shown in Figure 6.11 for $P_s = 0.001$ (see Figure S6 for other values). Both words and categories follow a similar pattern, with an initial fast proliferation and a following convergence toward the minimum number of elements: one single category for each source, and one single word per category. The radius r_0 determines the likelihood that a new category is created: when the radius is large enough, the initial category easily covers the whole source, and creation of new categories for the same source is unlikely. As a consequence, also the number of words generated is lower, because different words are generated for different categories, and the vocabularies are preserved by the category merging. In any case, the system tends to converge to the minimum number of words/categories for each value of r_0. We note that the actual convergence on two categories (and hence two words) is not always permanent, as new categories can emerge. These rare events are unlikely to have long lasting impact as the swarm can recover quickly. Under these conditions, we define as time of convergence (over two categories or two words) the first time the whole swarm reaches the minimum number of words/categories.

Both category and word convergence times depends heavily on r_0, but also

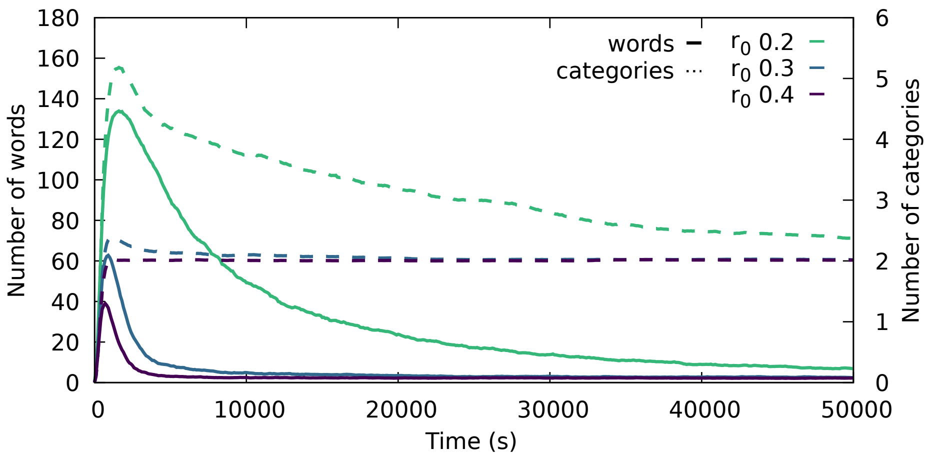

Figure 6.11: Average number of different words (solid lines) and different categories (dotted lines) present within the swarm. The dynamics over time are plotted for different values of r_0, and for a fixed probability of speaking $P_s = 0.001$.

on the probability of speaking P_s (see the top-left panel in Figure 6.12). When r_0 is intermediate-small, the large proliferation of categories requires several merging operations, and having more variability in each category does not give an advantage. On the other hand, for large r_0 few categories are formed, and a high P_s helps in quickly converging. These dynamics are confirmed also by the time of convergence to a single word per category (Figure 6.12 top right), which always decreases when P_s increases, with a larger effect for larger r_0.

Apart from the speed of convergence, another relevant aspect concerns the accuracy with which the emerging categories describe the sources to which they are associated. To measure this, we consider the position error as the distance between the center of the category and the center of the source (see Figure 6.12 bottom left) and the average radius of the final category (Figure 6.12 bottom right). When the initial radius is smaller, the error in the position of the prototype is very small, as it results from the average of many categories defined all around the source. With larger r_0, the position error increases because fewer categories are generated. Large values of the probability of speaking P_s result in even fewer generated categories: as robots receive their initial category from other robots, larger errors are made. For what concerns the final radius of the emergent categories, smaller values are observed for small r_0. However, the relative increase of the final radius with respect the initial r_0 is much larger for small r_0 than for large r_0 because many different categories are merged together.

6.6 Conclusion

In this chapter, we studied how the language game dynamics are influenced by the evolving topology of a swarm engaged in a decision-making and foraging

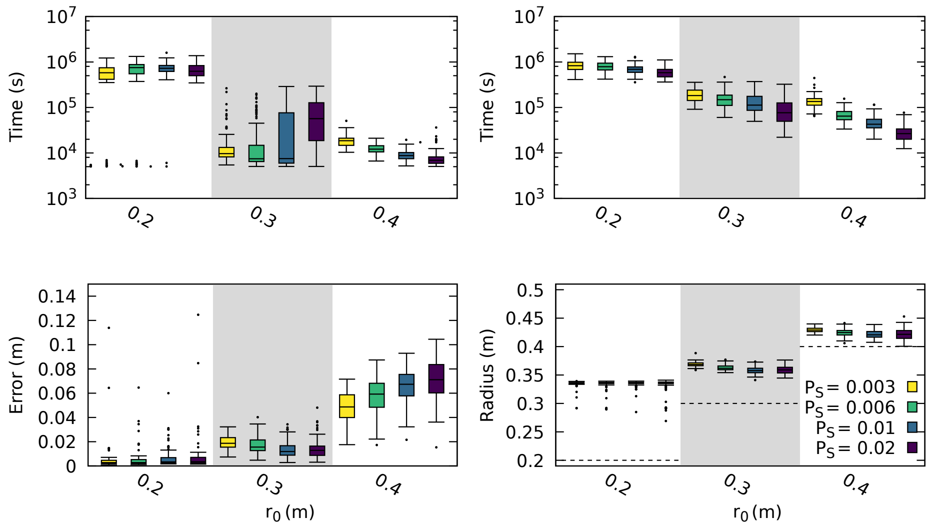

Figure 6.12: Effects of the initial category radius r_0 and of the probability of speaking P_s. Top left: categories' convergence time. Top right: words' convergence time. Bottom left: average error of the final category prototype with respect to the center of the associated source. Bottom right: average final radius of each category compared with the initial value of r_0 (dotted line).

task. In particular, we studied how well the swarm could maintain a description of its whole environment that is at the same time correct and complete, with the vocabulary containing only words that are relevant to the swarm, that is, those associated to sources under exploitation. We focus on such a compelling research question, without questioning properties commonly studied in swarm robotics such as robustness or scalability. Such properties have been largely studied for foraging, language dynamics and decision-making in previous studies and in conditions very close to the ones discussed here (Trianni and Campo, 2015; Reina et al., 2015a). Hence, they are no further debated, allowing us to focus on the interplay between language and decision dynamics.

We began by comparing two variations of the MNG. One that binds the creation of words with the sources available in the environment (spatial game), the other without such spatial correlation (classic game, where words are used as simple tokens). The differences in word creation between these two language games resulted in a large difference in the final outcome, despite also in the classic game words were created at locations that are always closer to one of the sources. The stronger correlation between creation of words and source location granted by the spatial game is not the only reason for the better matching and completeness. We observed that a major difference is given by the role of uncommitted, exploring robots into the creation and sharing of words. These robots can end up choosing any source, bringing words created near one source to the population exploiting the other. Additionally, we observed that the topology of the robot's interaction network—determined by the robot's movements during the foraging activity—consists of two almost segregated sub-populations, with sporadic interactions constrained to the central nest area. Such segregation creates the conditions for the maintenance of one word for each source, supporting completeness of the evolving vocabulary. In order for the swarm to maintain a complete description of the environment even when sources are not relevant any more, we proposed as a proof of concept a simple version of a category game embedded in space. In this setup, the swarm creates different categories for each source, and ends up retaining an exhaustive description that can also be sufficiently precise to potentially support the foraging activities.

One potential drawback of language evolution as observed in our experiments is related to the time required for emergent conventions to settle, which can be very large if interactions are sporadic, as well as the possibility that new conventions enter the population and destabilize the language dynamics. In this respect, it is important to note that linguistic conventions do not have an intrinsic value, but are more valuable when they are largely shared within a population, favouring coordination and avoiding misunderstandings. Hence, it is possible to speed up convergence toward a shared convention within a population by means of positive feedback mechanisms that favour the conventions more commonly found within the population. For instance, the simple rules of the naming game could be enhanced with estimates of the frequency of words in the population, allowing to favour the selection of more frequent words when speaking, hence speeding up convergence. Additionally, decentralised quorum sensing approaches can be exploited to determine a final convention, avoiding that noise is added by new alternatives when a largely shared one is already present. These and similar mechanisms can reduce the number of interactions required to achieve language convergence within a population, making language games practicable in realistic settings beyond the abstract scenario studied in this chapter.

In future studies, besides describing the relevant features of the environment, linguistic conventions can be exploited also to agree on the best course of action for the swarm. For instance, robots would share short term plans described as a sequence of linguistic elements, creating and merging them following shared compositional strategies. In this sense, the possibilities offered by language evolution are vast, allowing robot swarms to autonomously find sentence-like solutions to complex tasks made of several spatially-distributed and temporally-dependent sub-tasks.

Overall, we believe that merging language dynamics with the self-organising behaviour of robot swarms can be extremely useful, as the behaviour can exploit the emergent descriptions of the environment in a way that is dependent on the features relevant for the swarm behaviour. The link between language and behaviour was relegated here to the creation of words/categories. However, stronger links can be built if behavioural decisions can be determined by the evolving language, leading to the emergence of behaviours that vary as the relevant descriptions of the environment gets more precise and shared within the swarm. This also allows to adapt the language to the environmental contingencies encountered, possibly enabling more flexibility in the swarm behaviour with respect to changing environmental conditions (Cambier et al., 2021).

Chapter 7

Conclusions

This thesis studies the interaction between foraging and communication dynamics in swarm robotics, progressively introducing more complex layers of communication, culminating with the implementation of a version of language games linked with the task at hand.

We initially engaged in an in-depth study of the communication protocol to sustain social odometry, which enabled an effective navigation and foraging activity by the swarm when confronted with both abstract goal locations and tangible resources. The study revealed the profound effects that the proposed information processing mechanisms have on navigation and exploitation efficiency, as well as the global dynamics of the swarm. Specifically, we observed the potential for these mechanisms to either induce convergence on a singular path exploitation or initiate a distribution over several comparably efficient paths. These insights serve as a valuable guide for future designers, helping them to select the most suitable information aggregation mechanism depending on the task and goals of the swarm they design. This initial study, presented in Chapter 4, contributed to the understanding of the basic building blocks necessary for an effective foraging behaviour, which became the basis of the research presented in the following chapters. In particular, we aimed to better understand how the topology of the swarm is affected by the foraging task, and hence its role with respect to communication among robots.

Subsequently, in Chapter 5, we developed a more complex strategy for exploration and balanced exploitation of renewable sources inspired by the honeybee value-sensitive decision making abilities. A large-scale simulation analysis was conducted to identify the effects of the different parameters governing the individual behaviour on the macroscopic, swarm-level dynamics. The results showed the effectiveness of our approach in enabling adaptive exploitation of sources at the collective level, without the need for individual profitability comparisons or a centralized planner with global environmental knowledge. This approach and its subsequent validation served as a significant step towards the overarching goal of the thesis: to better understand the interplay between the task dynamics, and the social dynamics of a swarm. The balance between exploration and source exploitation sets a foundation for the complex communication systems that we next introduced.

Lastly, in Chapter 6, we studied the interplay between language dynamics and the evolving topology of a swarm engaged in decision-making and foraging. Notably, we demonstrated the ability of a swarm to sustain an accurate and comprehensive environmental description, with a vocabulary restricted to relevant terms, specifically those associated with sources being exploited. This investigation into the interconnected dynamics of language games and swarm behaviors marks a significant achievement of the thesis' main objective: to advance swarm robotics through the integration of task related linguistically-inspired communication mechanisms, and the study of the links between both tasks and language. Demonstrating the swarm's capacity to maintain an accurate and useful linguistic representation of its environment underscores the potential of this approach to enhance the swarm's foraging adaptability. This confirms the foundational belief driving this thesis: that the synergy between swarm robotics and language games can produce a new level of operational sophistication. Indeed, the transition towards a richer language-based communication follows a belief that the compatibility between language games and swarm robotics can yield great results, both in enhancing the efficiency and adaptivity of the communication between the robots in the swarm and in providing new means to study the evolution of language. We further discuss these aspects in the following sections, also highlighting possible directions for future work.

7.1 A perspective for language evolution in swarm robotics

A key ingredient for both swarm robotics and language evolution is self-organisation in a population of agents resulting from local interactions. This is a dynamic process that comes about within the population in response to the contingencies experienced by the agents while displaying their behaviour, resulting in a dynamic process that can react to changes in the environment, with new behaviours and/or concepts/words arising when needed.

Indeed, it has been shown that natural languages are the result of a self-organising process. Natural languages arose (at least partly) from the need to purposely share information (Noble and Davidson, 1991). As a matter of fact, the rise of a complex communication system is linked to specific tasks to be accomplished, and ultimately, to survival of the population. Language games and swarm robotics would find value being combined as both have decentralisation and self-organisation at their roots, allowing to produce communication systems that evolve online, that are exploited to represent a dynamic and uncertain environment and that can be used to improve performance in task execution. Self-organising languages should therefore enable robots to describe and tackle new challenges as they come. They are useful to build truly adaptive robot swarms.

In addition to their focus on self-organisation, language games address several challenges related to the development of complex artificial systems. First is the symbol grounding problem, i.e., how to associate a word to experiences of the real world (Harnad, 1990). Most studies on language games rely on linking perceptions of the environment to names. Indeed, as robots have different bodies, they can experience the environment with different modalities and from different points of view, so that the sensory data they associate with a shared word is usually not exactly the same. This association mechanism only requires simple feature extraction algorithms and a bidirectional memory (mapping between words and meanings, as in Steels and Loetzsch, 2012). The consensus dynamics of language games ensures the symbol grounding at the population level, although each robot has its own internal representation of each word's meaning. This can cause issues in communication early on, but can result in a very efficient system once a robust association between concepts and experiences has been established. This favours especially "uneducated" robots that join the swarm later on, as they can quickly acquire new concepts from few interactions with their peers.

Alternative combinations where swarms of robots learn to link actions to verbs—hence, to actions—are possible and could provide interesting new abilities. Some work has already been performed in that direction using language games (Steels, 2008), and additional efforts on symbol grounding could also take advantage from *developmental language acquisition* models (Rasheed and Amin, 2016). Such models use neural networks to incrementally teach agents (in a teacher/learner scenario) new meanings, starting from observable objects before moving onto more abstract concepts or actions.

Relevant studies in the evolution of language have shown that a simple grammar can be generated exploiting the plasticity of the learned language (Spranger et al., 2010). The emergence of a simple compositionality in the language can lead to more complex expressions, paving the way to full sentences. This is of great relevance for swarm robotics studies in order to address complex tasks that require multiple actions to be scheduled and coordinated among robots. An efficient scheduling and coordination plan can emerge from local knowledge to be shared among robots, without any pre-defined structure or plan. By exploiting the compositionality of language, a sequence of tasks can be defined and eventually executed, leading to the emergence of swarm behaviours that are far more complex than the current state of the art. Through the development of *fluid construction grammars* (Steels and De Beule, 2006), language games can evolve grammatical structures as well as lexicons (Beuls and Steels, 2013). The dynamics of such games have, however, not been as thoroughly studied as, e.g., the naming game, and are not ready for implementation in swarm robotics contexts. The robotics community should therefore attempt to facilitate the development of such language games in order to be able to benefit from them.

7.2 Future work

Swarm robotics is heavily influenced by the social insects metaphor, embodying certain characteristics often observed in such systems such as simple behavior, simple and often memory-less agents, and homogeneous swarms. On the other hand, mammals, particularly primates, have developed more complex communication skills than insects, which in some cases bear a closer resemblance to human communication. Consequently, it may be beneficial to transition from an approach solely inspired by insects' biological characteristics to one that also incorporates more complex forms of (social) cognition inspired by mammals and primates.

In this perspective, swarm robotics can serve as an ideal platform also to deepen our understanding of the drives and constrains underlying social structure and behavior in mammals and humans. For instance, current research activities are focused on applying swarm robotics to model the human self-domestication hypothesis, thanks to a multidisciplinary research group that includes both swarm roboticists and linguistics researchers. This hypothesis states that the evolution of present-day languages might have resulted in part from the self-domestication of the human species (see Thomas and Kirby, 2018). This evolutionary process, which is similar to animal domestication (Hare, 2017), would result in less aggressive individuals, more prone to interact with one another (and particularly, with their kin), promoting more social contacts within a community, and supporting the emergence of more sophisticated forms of language. To conduct this research, it was necessary to incorporate two concepts that are typically not associated with swarm robotics: differentiation and prosociality. In this context, differentiation refers to the capability of robots to identify one another and display distinct behaviors in interaction with various peers. This differentiation can affect two key aspects: the way in which robots collaborate on tasks, and the manner they share information with one another. On the other hand, prosociality is a trait that favors behavior characterized by actions intended to benefit others. In this study, it is represented as a factor influencing the likelihood of communication between robots. When the prosociality level between robots is high, they tend to interact more frequently. Robots that start with a high base level of prosociality are considered to be more sociable because they are more open to interactions. Both differentiation and prosociality are examples of elements that contribute to the formation of distinctive interactions among agents. When agents possess the capacity to recognize specific teammates and exhibit differential behavior with them, the emergence of tight social groups becomes possible. This not only fosters the emergence of cultural in-groups within the swarm but also leads to the observation of more intricate language dynamics (Cambier et al., 2022).

Future research endeavors could explore the utilization of other existing language games beyond the naming game or propose new and more sophisticated language games coming for the more complex social layer introduced above. Furthermore, the work presented in this thesis very much assumes some de facto protocols agreed upon in advance, innate to the robot. Future work could aim to allow these frames to evolve organically during preliminary experiments conducted in controlled environments. For instance, Cambier et al. (2023) propose an emergent naming system for robotic swarms to facilitate collective navigation and decision-making in unstructured environments. The robots collectively name landmarks they discover, using them as beacons for navigation and scoring them based on relevance to the task. Comparisons with non-communicating swarms and swarms with prior knowledge show that the naming-based approach performs similarly to the latter, enabling robots to find a topological path without individually mapping the environment. We plan to engage in a similar path, in which the language game is exploited to identify sequences of actions to be executed (e.g., the sequence of decisions to be taken to explore and exit from a maze with the shortest path). This can leverage language by evolving references to places (names) and actions (verbs), hence going beyond the evolution of a shared vocabulary that informs much of the research on language games. We plan to move in this way towards the emergence of a grammar structure and simple compositionality in language, as mentioned above. This is of profound relevance for swarm robotics studies aiming to tackle complex tasks requiring coordinated and scheduled actions among robots. By capitalizing on the compositionality of language, a sequence of tasks can be defined, planned, shared, and eventually executed, leading to the emergence of swarm behaviors far more complex than the current state of the art.

In conclusion, this thesis underscores the promising potential inherent in the fusion of swarm robotics and language games. We have barely scratched the surface of this new field, with numerous directions awaiting exploration. The breadth and depth of possibilities that lie ahead are vast, spanning from more sophisticated language games to nuanced behavioral dynamics. As this work propels us forward, it is our hope that the momentum will continue, prompting further inquiry and driving the field to new heights. The field of swarm robotics is evolving rapidly and this is an encouraging time for the development of new concepts and methods. It is with a sense of optimism that we look forward to seeing how future research will further shape and enhance this domain.

Bibliography

D. H. Ackley and M. L. Littman. Altruism in the evolution of communication. In *Artificial life IV*, pages 40–48, Cambridge, MA, USA, 1994.

S. Alers, K. Tuyls, B. Ranjbar-Sahraei, D. Claes, and G. Weiss. Insect-inspired robot coordination: foraging and coverage. In *Artificial Life Conference Proceedings 14*, pages 761–768, Cambridge, MA, USA, 2014. MIT Press.

C. Ampatzis, E. Tuci, V. Trianni, and M. Dorigo. Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16(1):5–26, 2008.

P. Bailis, R. Nagpal, and J. Werfel. Positional communication and private information in honeybee foraging models. In *International Conference on Swarm Intelligence 2010*, pages 263–274. 2010.

G. Baldassarre, D. Parisi, and S. Nolfi. Distributed coordination of simulated robots based on self-organization. *Artificial Life*, 12(3):289–311, 2006.

A. Baronchelli. Role of feedback and broadcasting in the naming game. *Physical Review E*, 83(4):046103, 2011.

A. Baronchelli. A gentle introduction to the minimal naming game. *Belgian Journal of Linguistics*, 30(1):171–192, 2016.

A. Baronchelli and A. Díaz-Guilera. Consensus in networks of mobile communicating agents. *Physical Review E*, 85(1):016113, Jan. 2012.

A. Baronchelli, L. Dall'Asta, A. Barrat, and V. Loreto. Topology-induced coarsening in language games. *Physical Review E*, 73(1):015102, 2006a.

A. Baronchelli, M. Felici, V. Loreto, E. Caglioti, and L. Steels. Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(06):P06014, 2006b.

A. Baronchelli, L. Dall'Asta, A. Barrat, and V. Loreto. The role of topology on the dynamics of the naming game. *The European Physical Journal Special Topics*, 143(1):233–235, 2007.

A. Baronchelli, T. Gong, A. Puglisi, and V. Loreto. Modeling the emergence of universality in color naming patterns. *Proceedings of the National Academy of Sciences*, 107(6):2403–2407, 2010.

P. Bartashevich and S. Mostaghim. Ising model as a switch voting mechanism in collective perception. In *EPIA Conference on Artificial Intelligence*, pages 617–629. Springer, 2019.

P. Bartashevich and S. Mostaghim. Multi-featured collective perception with evidence theory: tackling spatial correlations. *Swarm Intelligence*, 15(1): 83–110, 2021.

K. Beuls and L. Steels. Agent-based models of strategies for the emergence and evolution of grammatical agreement. *PLoS ONE*, 8(3):e58960, 2013.

A. Billard and G. Hayes. Learning to communicate through imitation in autonomous robots. In *International Conference on Artificial Neural Networks*, pages 763–768. Springer, 1997.

E. Bonabeau, G. Theraulaz, and J.-L. Deneubourg. Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 263 (1376):1565–1569, 1996.

M. Bonani, V. Longchamp, S. Magnenat, P. Rétornaz, D. Burnier, G. Roulet, F. Vaussard, H. Bleuler, and F. Mondada. The marXbot, a miniature mobile robot opening new perspectives for the collective-robotic research. In *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4187–4193. IEEE Press, 2010a.

M. Bonani, V. Longchamp, S. Magnenat, P. Rétornaz, D. Burnier, G. Roulet, F. Vaussard, H. Bleuler, and F. Mondada. The MarXbot, a miniature mobile robot opening new perspectives for the collective-robotic research. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 4187–4193. IEEE Press, Piscataway, NJ, 2010b.

J. Borenstein and Y. Koren. Real-time obstacle avoidance for fast mobile robots. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5):1179–1187, 1989.

M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.

S. Camazine. *Self-organization in biological systems*. Princeton University Press, 2003.

N. Cambier, V. Frémont, and E. Ferrante. Group-size regulation in self-organised aggregation through the naming game. In *International Symposium on Swarm Behavior and Bio-Inspired Robotics (SWARM 2017)*, Kyoto, Japan, Oct 2017. URL https://hal.archives-ouvertes.fr/hal-01679600.

N. Cambier, V. Frémont, V. Trianni, and E. Ferrante. Embodied evolution of self-organised aggregation by cultural propagation. In *Swarm Intelligence: 11th International Conference, ANTS 2018, Rome, Italy, October 29–31, 2018, Proceedings 11*, pages 351–359. Springer, 2018.

N. Cambier, R. Miletitch, V. Frémont, M. Dorigo, E. Ferrante, and V. Trianni. Language evolution in swarm robotics: A perspective. *Frontiers in Robotics and AI*, 7:12, 2020.

N. Cambier, D. Albani, V. Frémont, V. Trianni, and E. Ferrante. Cultural evolution of probabilistic aggregation in synthetic swarms. *Applied Soft Computing*, 113:108010, 2021.

N. Cambier, R. Miletitch, A. B. Burraco, and L. Raviv. Prosociality in swarm robotics: A model to study self-domestication and language evolution. In *Joint Conference on Language Evolution (JCoLE)*, pages 98–100. Joint Conference on Language Evolution (JCoLE), 2022.

N. Cambier, A. Eiben, and E. Ferrante. Emergent naming system in an unstructured environment: a shortest-path discovery case study. In *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*. MIT Press, 2023.

A. Campo, Á. Gutiérrez, S. Nouyan, C. Pinciroli, V. Longchamp, S. Garnier, and M. Dorigo. Artificial pheromone for path selection by a foraging swarm of robots. *Biological Cybernetics*, 103(5):339–352, 2010.

E. Castello, T. Yamamoto, F. D. Libera, W. Liu, A. F. T. Winfield, Y. Nakamura, and H. Ishiguro. Adaptive foraging for simulated and real robotic swarms: the dynamical response threshold approach. *Swarm Intelligence*, 10(1):1–31, 2015.

H. Çelikkanat, A. E. Turgut, and E. Şahin. Guiding a robot flock via informed robots. In *Distributed Autonomous Robotic Systems 8*, pages 215–225. Springer, 2009.

F. A. A. Cheein and R. Carelli. Agricultural robotics: Unmanned robotic service units in agricultural tasks. *IEEE Industrial Electronics Magazine*, 7(3):48–58, 2013.

H. H. Clark, S. E. Brennan, et al. Grounding in communication. *Perspectives on Socially Shared Cognition*, 13(1991):127–149, 1991.

B. De Boer. Self-organization in vowel systems. *Journal of Phonetics*, 28(4): 441–465, 2000.

J. De Greeff and S. Nolfi. Evolution of implicit and explicit communication in mobile robots. In *Evolution of communication and language in embodied agents*, pages 179–214. Springer, Berlin, Germany, 2010.

J.-L. Deneubourg, S. Aron, S. Goss, and J. M. Pasteels. The self-organizing exploratory pattern of the argentine ant. *Journal of Insect Behavior*, 3(2): 159–168, 1990.

C. Dimidov, G. Oriolo, and V. Trianni. Random walks in swarm robotics: An experiment with kilobots. In M. Dorigo, M. Birattari, X. Li, M. López-Ibáñez, K. Ohkura, C. Pinciroli, and T. Stützle, editors, *Proceedings of the 10th International Conference on Swarm Intelligence (ANTS 2016)*, volume 9882 of *LNCS*, pages 185–196. Springer New York, 2016.

M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, et al. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2):223–245, 2004.

M. Dorigo, D. Floreano, L. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A. L. Christensen, A. Decugniere, G. Di Caro, F. Ducatelle, E. Ferrante, A. Forster, J. Martinez Gonzales, J. Guzzi, V. Longchamp, S. Magnenat, N. Mathews, M. Montes de Oca, R. O'Grady, C. Pinciroli, G. Pini, P. Rétornaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stützle, V. Trianni, E. Tuci, A. E. Turgut, and F. Vaussard. Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics & Automation Magazine*, 20 (4):60–71, 2013.

M. Dorigo, G. Theraulaz, and V. Trianni. Reflections on the future of swarm robotics. *Science Robotics*, 5(49):eabe4385, 2020.

M. Dorigo, G. Theraulaz, and V. Trianni. Swarm Robotics: Past, Present, and Future. *Proceedings of the IEEE*, 109(7):1152–1165, 2021.

A. Dornhaus, F. Klügl, C. Oechslein, F. Puppe, and L. Chittka. Benefits of recruitment in honey bees: effects of ecology and colony size in an individual-based model. *Behavioral Ecology*, 17(3):336–344, 2006.

F. Ducatelle, G. Di Caro, C. Pinciroli, F. Mondada, and L. Gambardella. Communication assisted navigation in robotic swarms: Self-organization and cooperation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4981–4988. IEEE Computer Society Press, Los Alamitos, CA, 2011a.

F. Ducatelle, G. A. Di Caro, C. Pinciroli, and L. M. Gambardella. Self-organized cooperation between robotic swarms. *Swarm Intelligence*, 5(2):73, 2011b.

F. Ducatelle, G. A. Di Caro, C. Pinciroli, F. Mondada, and L. Gambardella. Communication assisted navigation in robotic swarms: self-organization and cooperation. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4981–4988, San Francisco, California, USA, 2011c. IEEE.

F. Ducatelle, G. A. Di Caro, A. Forster, M. Bonani, M. Dorigo, S. Magnenat, F. Mondada, R. O'Grady, C. Pinciroli, P. Rétornaz, V. Trianni, and L. M. Gambardella. Cooperative navigation in robotic swarms. *Swarm Intelligence*, 8(1):1–33, 2014.

E. Ferrante, M. Brambilla, M. Birattari, and M. Dorigo. Socially-mediated negotiation for obstacle avoidance in collective transport. In *Distributed Autonomous Robotic Systems*, pages 571–583. Springer, Berlin, Germany, 2013.

E. Ferrante, A. E. Turgut, A. Stranieri, C. Pinciroli, M. Birattari, and M. Dorigo. A self-adaptive communication strategy for flocking in stationary and non-stationary environments. *Natural Computing*, 13(2):225–245, 2014.

D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17(6):514–519, 2007.

G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pinciroli, et al. An experiment in automatic design of robot swarms: Automode-vanilla, evostick, and human experts. In *Swarm Intelligence: 9th International Conference, ANTS 2014, Brussels, Belgium, September 10-12, 2014. Proceedings 9*, pages 25–37. Springer, 2014.

G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pinciroli, et al. Automode-chocolate: automatic design of control software for robot swarms. *Swarm Intelligence*, 9 (2-3):125–152, 2015.

S. Garnier, J. Gautrais, and G. Theraulaz. The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31, 2007.

B. Granovskiy, T. Latty, M. Duncan, D. J. T. Sumpter, and M. Beekman. How dancing honey bees keep track of changes: the role of inspector bees. *Behavioral Ecology*, 23(3):588–596, 2012.

R. Groß, E. Tuci, M. Dorigo, M. Bonani, and F. Mondada. Object transport by modular robots that self-assemble. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 2558–2564, Orlando, Florida, USA, 2006. IEEE.

F. G. Group", C. Beckner, R. Blythe, J. Bybee, M. H. Christiansen, W. Croft, N. C. Ellis, J. Holland, J. Ke, D. Larsen-Freeman, et al. Language is a complex adaptive system: Position paper. *Language Learning*, 59:1–26, 2009.

A. Gutierrez, A. Campo, M. Dorigo, J. Donate, F. Monasterio-Huelin, and L. Magdalena. Open e-puck range bearing miniaturized board for local communication in swarm robotics. In *2009 IEEE International Conference on Robotics and Automation*, pages 3111–3116, May 2009.

A. Gutiérrez, A. Campo, F. Santos, F. Monasterio-Huelin Maciá, and M. Dorigo. Social odometry: imitation based odometry in collective robotics. *International Journal of Advanced Robotic Systems*, 6(2):129–136, 2009.

A. Gutiérrez, A. Campo, F. Monasterio-Huelin, L. Magdalena, and M. Dorigo. Collective decision-making based on social odometry. *Neural Computing & Applications*, 19(6):807–823, 2010.

B. Hare. Survival of the friendliest: Homo sapiens evolved via selection for prosociality. *Annual Review of Psychology*, 68:155–186, 2017.

S. Harnad. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3):335–346, 1990.

K. Hasselmann, F. Robert, and M. Birattari. Automatic design of communication-based behaviors for robot swarms. In *Swarm Intelligence: 11th International Conference, ANTS 2018, Rome, Italy, October 29–31, 2018, Proceedings 11*, pages 16–29. Springer, 2018.

J. P. Hecker and M. E. Moses. Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intelligence*, 9(1): 1–28, 2015.

P. Holme and J. Saramäki. Temporal networks. *Physics Reports*, 519(3):97–125, 2012.

M. J. B. Krieger, J.-B. Billeter, and L. Keller. Ant-like task allocation and recruitment in cooperative robots. *Nature*, 406(6799):992–995, 2000.

C. R. Kube and E. Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1-2):85–101, 2000.

T. H. Labella, M. Dorigo, and J.-L. Deneubourg. Division of labor in a group of robots inspired by ants' foraging behavior. *ACM Transactions on Autonomous Adaptive Systems*, 1(1):4–25, 2006.

S. Liemhetcharat, R. Yan, and K. P. Tee. Continuous foraging and information gathering in a multi-agent team. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1325–1333. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2015.

W. Liu and A. F. T. Winfield. Modeling and optimization of adaptive foraging in swarm robotic systems. *The International Journal of Robotics Research*, 29 (14):1743–1760, 2010.

W. Liu, A. F. T. Winfield, J. Sa, J. Chen, and L. Dou. Towards energy optimization: Emergent task allocation in a swarm of foraging robots. *Adaptive Behavior*, 15(3):289–305, 2007.

V. Loreto, A. Baronchelli, A. Mukherjee, A. Puglisi, and F. Tria. Statistical physics of language dynamics. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(04):P04006, Apr. 2011.

Q. Lu, G. Korniss, and B. K. Szymanski. Naming games in two-dimensional and small-world-connected random geometric networks. *Physical Review E*, 77(1): 016111, 2008.

A. Martinelli, F. Pont, and R. Siegwart. Multi-robot localization using relative observations. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2797–2802. IEEE Press, New York, NY, 2005.

R. Miletitch, M. Dorigo, and V. Trianni. Social dynamics for an exploration and exploitation task in swarm robotics. *Rapport d'avancement des recherches*, 2013a.

R. Miletitch, V. Trianni, A. Campo, and M. Dorigo. Information aggregation mechanisms in social odometry. In *Proceedings of the 20th European Conference on Artificial Life (ECAL 2013)*, pages 102–109. MIT Press, Cambridge, MA, 2013b.

R. Miletitch, M. Dorigo, and V. Trianni. Balancing exploitation of renewable resources by a robot swarm. *Swarm Intelligence*, 12(4):307–326, 2018.

R. Miletitch, A. Reina, M. Dorigo, and V. Trianni. Emergent naming conventions in a foraging robot swarm. *Swarm Intelligence*, 16(3):211–232, 2022.

M. A. Montes de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo. Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3):305–327, 2011.

P. Moretti, A. Baronchelli, M. Starnini, and R. Pastor-Satorras. Generalized voter-like models on heterogeneous networks. In A. Mukherjee, M. Choudhury, F. Peruani, N. Ganguly, and B. Mitra, editors, *Dynamics On and Of Complex Networks, Volume 2: Applications to Time-Varying Dynamical Systems*, pages 285–300. Springer New York, 2013.

R. R. Murphy, S. Tadokoro, D. Nardi, A. Jacoff, P. Fiorini, H. Choset, and A. M. Erkmen. Search and rescue robotics. In *Springer Handbook of Robotics*, pages 1151–1173. Springer, 2008.

W. Noble and I. Davidson. The evolutionary emergence of modern human behaviour: Language and its archaeology. *Man*, 26(2):223–253, 1991.

S. Nolfi and D. Floreano. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines.* MIT press, Cambridge, MA, USA, 2000.

S. Nolfi and M. Mirolli. *Evolution of communication and language in embodied agents*. Springer Science & Business Media, Berlin, Germany, 2009.

S. Nouyan, A. Campo, and M. Dorigo. Path formation in a robot swarm: Self-organized strategies to find your way home. *Swarm Intelligence*, 2:1–23, 2008.

S. Nouyan, R. Groß, M. Bonani, F. Mondada, and M. Dorigo. Teamwork in self-organized robot colonies. *IEEE Transactions on Evolutionary Computation*, 13(4):695–711, 2009.

R. O'Grady, A. L. Christensen, and M. Dorigo. Swarmorph: multirobot morphogenesis using directional self-assembly. *IEEE Transactions on Robotics*, 25 (3):738–743, 2009.

E. H. Ostergaard, G. S. Sukhatme, and M. J. Matari. Emergent bucket brigading: a simple mechanisms for improving performance in multi-robot constrained-space foraging tasks. In *Proceedings of the Fifth International Conference on Autonomous Agents*, AGENTS '01, pages 29–30, New York, NY, 2001. ACM. ISBN 1-58113-326-X. doi: 10.1145/375735.375825. URL http://doi.acm.org/10.1145/375735.375825.

M. J. Owren, D. Rendall, and M. J. Ryan. Redefining animal signaling: influence versus information in communication. *Biology & Philosophy*, 25(5):755–780, 2010.

R. O'Grady, R. Groß, A. L. Christensen, and M. Dorigo. Self-assembly strategies in a group of autonomous mobile robots. *Autonomous Robots*, 28(4):439–455, 2010.

D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. R. Marshall. A mechanism for value-sensitive decision-making. *PLoS ONE*, 8(9): e73216, 2013.

C. A. Parker and H. Zhang. Biologically inspired collective comparisons by robotic swarms. *The International Journal of Robotics Research*, 30(5):524–535, 2011.

T. P. Pavlic, J. Hanson, G. Valentini, S. I. Walker, and S. C. Pratt. Quorum sensing without deliberation: biological inspiration for externalizing computation to physical spaces in multi-robot systems. *Swarm Intelligence*, 15(1): 171–203, 2021.

A. Perna and T. Latty. Animal transportation networks. *Journal of The Royal Society Interface*, 11(100):20140334–20140334, 2014.

C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, F. Ducatelle, M. Birattari, L. M. Gambardella, and M. Dorigo. ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4):271–295, 2012.

L. Pitonakova, R. Crowder, and S. Bullock. Information flow principles for plasticity in foraging robot swarms. *Swarm Intelligence*, 10(1):33–63, 2016.

A. Puglisi, A. Baronchelli, and V. Loreto. Cultural route to the emergence of linguistic categories. *Proceedings of the National Academy of Sciences*, 105 (23):7936–7940, 2008.

W. V. O. Quine. *Word and object*. MIT press, Cambridge, MA, USA, 2013.

N. Rasheed and S. H. Amin. Developmental and evolutionary lexicon acquisition in cognitive agents/robots with grounding principle. *Computational Intelligence and Neuroscience*, 2016:16, 2016.

C. R. Reid, S. Garnier, M. Beekman, and T. Latty. Information integration and multiattribute decision making in non-neuronal organisms. *Animal Behaviour*, 100:44–50, 2015.

A. Reina, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro-macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intelligence*, 9(2-3):75–102, 2015a.

A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A design pattern for decentralised decision making. *PLoS ONE*, 10(10):e0140950–18, 2015b.

A. Reina, T. Bose, V. Trianni, and J. A. Marshall. Effects of spatiality on value-sensitive decisions made by robot swarms. In *Proceedings of 13th International Symposium on Distributed Autonomous Robotic Systems (DARS 2016)*, pages 1–8, Natural History Museum in London, UK, 2016.

A. Reina, J. A. R. Marshall, V. Trianni, and T. Bose. Model of the best-of-n nest-site selection process in honeybees. *Physical Review E*, 95(5):052411–15, May 2017.

I. Rekleitis, G. Dudek, and E. Milios. Multi-robot collaboration for robust exploration. *Annals of Mathematics and Artificial Intelligence*, 31(1-4):7–40, 2001.

J. Roberts, T. S. Stirling, J.-C. Zufferey, and D. Floreano. 2.5D infrared range and bearing system for collective robotics. In *Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3659–3664. IEEE Press, 2009.

M. Rubenstein, A. Cornejo, and R. Nagpal. Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198):795–799, 2014.

P. E. Rybski, A. Larson, H. Veeraraghavan, M. LaPoint, and M. a. Gini. Communication strategies in multi-robot search and retrieval: Experiences with mindart. In *Distributed Autonomous Robotic Systems 6*, pages 317–326. Springer, 2007.

E. Şahin. Swarm robotics: From sources of inspiration to domains of application. In *International Workshop on Swarm Robotics*, pages 10–20, Santa Monica, CA, USA, 2004. Springer.

A. Schroeder, S. Ramakrishnan, M. Kumar, and B. Trease. Efficient spatial coverage by a robot swarm based on an ant foraging model and the lévy distribution. *Swarm Intelligence*, 11(1):39–69, 2017.

T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–111, 2012a. ISSN 0036-8075.

T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall. Stop Signals Provide Cross Inhibition in Collective Decision-Making by Honeybee Swarms. *Science*, 335(6064):108–111, 2012b.

I. Slavkov, D. Carrillo-Zapata, N. Carranza, X. Diego, F. Jansson, J. Kaandorp, S. Hauert, and J. Sharpe. Morphogenesis in robot swarms. *Science Robotics*, 3(25):eaau9178, 2018.

F. Smarandache and J. Dezert. Information fusion based on new proportional conflict redistribution rules. In *2005 7th International Conference on Information Fusion*, volume 2, pages 8–pp. IEEE, 2005.

Z. Song and R. T. Vaughan. Sustainable robot foraging: Adaptive fine-grained multi-robot task allocation for maximum sustainable yield of biological resources. In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3309–3316. IEEE Press, 2013.

O. Soysal and E. Sahin. Probabilistic aggregation strategies in swarm robotic systems. In *Swarm Intelligence Symposium, 2005. SIS 2005. Proceedings 2005 IEEE*, pages 325–332, Piscataway NJ, United States, 2005. IEEE.

V. Sperati, V. Trianni, and S. Nolfi. Self-organised path formation in a swarm of robots. *Swarm Intelligence*, 5(2):97–119, 2011.

M. Spranger. Evolving grounded spatial language strategies. *Künstliche Intelligenz*, 27(2):97–106, 2013.

M. Spranger, S. Pauw, and M. Loetzsch. Open-ended semantics co-evolving with spatial language. In *The evolution of language*, pages 297–304. World Scientific, 2010.

L. Steels. A self-organizing spatial vocabulary. *Artificial life*, 2(3):319–332, 1995.

L. Steels. Language games for autonomous robots. *IEEE Intelligent Systems*, 16 (5):16–22, 2001.

L. Steels. Evolving grounded communication for robots. *Trends in Cognitive Sciences*, 7(7):308–312, 2003.

L. Steels. The symbol grounding problem has been solved. so what's next. *Symbols and embodiment: Debates on meaning and cognition*, pages 223–244, 2008.

L. Steels. *The Talking Heads experiment: Origins of words and meanings*, volume 1. Language Science Press, Berlin, Germany, 2015.

L. Steels and T. Belpaeme. Coordinating perceptually grounded categories through language: A case study for colour. *The Behavioral and Brain Sciences*, 28(04):1–61, 2005.

L. Steels and J. De Beule. A (very) brief introduction to fluid construction grammar. In *Proceedings of the Third Workshop on Scalable Natural Language Understanding*, pages 73–80, New York City, New York, 2006. Association for Computational Linguistics.

L. Steels and M. Loetzsch. The grounded naming game. *Experiments in Cultural Language Evolution*, 3:41–59, 2012.

M. S. Talamali, T. Bose, M. Haire, X. Xu, J. A. Marshall, and A. Reina. Sophisticated collective foraging with minimalist agents: a swarm robotics test. *Swarm Intelligence*, 14(1):25–56, 2020.

J. Thomas and S. Kirby. Self domestication and the evolution of language. *Biology & Philosophy*, 33(1):1–30, 2018.

S. Thrun. Simultaneous localization and mapping. Springer Tracts in Advanced Robotics, pages 13–41–41. Springer Berlin Heidelberg, 2008.

V. Trianni. *Evolutionary swarm robotics: evolving self-organising behaviours in groups of autonomous robots*. Springer, 2008.

V. Trianni and A. Campo. Fundamental collective behaviors in swarm robotics. In J. Kacprzyk and W. Pedrycz, editors, *Springer Handbook of Computational Intelligence*, pages 1377–1394. Springer, 2015.

V. Trianni and M. Dorigo. Emergent collective decisions in a swarm of robots. In *Proceedings of the 2005 IEEE Swarm Intelligence Symposium (SIS 2005)*, pages 241–248, 2005.

V. Trianni and M. Dorigo. Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, 95(3):213–231, 2006.

V. Trianni and S. Nolfi. Self-organising sync in a robotic swarm. A dynamical system view. *IEEE Transactions on Evolutionary Computation*, 13(4):722–741, 2009.

V. Trianni, D. De Simone, A. Reina, and A. Baronchelli. Emergence of Consensus in a Multi-Robot Network: From Abstract Models to Empirical Validation. *IEEE Robotics and Automation Letters*, 1(1):348–353, 2016a.

V. Trianni, D. De Simone, A. Reina, and A. Baronchelli. Emergence of consensus in a multi-robot network: from abstract models to empirical validation. *IEEE Robotics and Automation Letters*, 1(1):348–353, 2016b.

E. Tuci. An investigation of the evolutionary origin of reciprocal communication using simulated autonomous agents. *Biological Cybernetics*, 101(3):183–199, 2009.

R. Uno, D. Marocco, S. Nolfi, and T. Ikegami. Emergence of protosentences in artificial communicating systems. *IEEE Transactions on Autonomous Mental Development*, 3(2):146–153, 2011.

G. Valentini, D. Brambilla, H. Hamann, and M. Dorigo. Collective perception of environmental features in a robot swarm. In *Swarm Intelligence: 10th International Conference, ANTS 2016, Brussels, Belgium, September 7-9, 2016, Proceedings 10*, pages 65–76. Springer, 2016.

G. Valentini, E. Ferrante, and M. Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:1–43, 2017.

J. Wessnitzer and C. Melhuish. Collective decision-making and behaviour transitions in distributed ad hoc wireless networks of mobile robots: Target-hunting. In *European Conference on Artificial Life*, pages 893–902. Springer, 2003.

A. F. Winfield. Foraging robots. In *Encyclopedia of Complexity and Systems Science*, pages 3682–3700. Springer New York, 2009.

L. Wittgenstein. Philosophical investigations, trans. *GEM Anscombe*, 261:49, 1953.

K. Yoshida. Achievements in space robotics. *IEEE Robotics & Automation Magazine*, 16(4):20–28, 2009.