Multi-agent Collective Decision-making for Foraging, Flocking and Target Search Applications

A thesis submitted in partial fulfillment of the requirements

for the award of the degree of

DOCTOR OF PHILOSOPHY

by

SWADHIN AGRAWAL

2010706



to the

DEPARTMENT OF ELECTRICAL ENGINEERING & COMPUTER SCIENCE

INDIAN INSTITUTE OF SCIENCE EDUCATION AND RESEARCH BHOPAL

Bhopal - 462 066

September 24, 2026

CERTIFICATE

The undersigned have examined the Ph.D. thesis entitled:

Multi-agent Collective Decision-making for Foraging, Flocking and Target Search Applications

presented by Mr.Swadhin Agrawal, a candidate for the degree of Doctor of Philosophy in the department of Electrical Engineering & Computer Science, and hereby certify that it is worthy of acceptance.

Date:16-04-2026
IISER, Bhopal



Prof. Arpita Sinha
External member



Prof. P.B. Sujit
Supervisor



Dr. Venkateshwar Rao
Expert member



Prof. Santanu Talukdar
Member



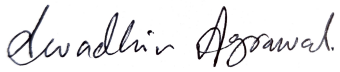
Dr. Haroon Lone
Member

ACADEMIC INTEGRITY AND COPYRIGHT DISCLAIMER

I hereby declare that this thesis is my own work and, to the best of my knowledge, it contains no materials previously published or written by any other person, or substantial proportions of material which have been accepted for the award of any other degree or diploma at IISER Bhopal or any other educational institution, except where due acknowledgment is made in the thesis.

I certify that all copyrighted material incorporated into this thesis is in compliance with the Indian Copyright (Amendment) Act, 2012, and that I have received written permission from the copyright owners for my use of their work, which is beyond the scope of the law. I agree to indemnify and save harmless IISER Bhopal from any and all claims that may be asserted or that may arise from any copyright violation.

September 24, 2026
IISER, Bhopal



Swadhin Agrawal

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to Prof. Sujit P. B. for giving me the opportunity to pursue my PhD under his supervision. His constant guidance, encouragement, and active involvement in the experiments have been invaluable throughout my research journey. My sincere thanks go to Dr. Andreagiovanni Reina for co-advising me during my doctoral studies and for his insightful suggestions that greatly enriched my work. I am deeply grateful to my collaborators, whose diverse perspectives and technical expertise were vital to my doctoral journey. I am profoundly thankful to all my teachers, who put their sweat and blood into nurturing me and guiding me, so that I could realize and use my full potential.

I also extend my gratitude to my sponsors—IISER Bhopal, MPCST Bhopal, F.R.S.-FNRS, University of Konstanz (Zukunftskolleg), and the Max Planck Institute of Animal Behavior—for their generous funding and support throughout my doctoral degree. I am grateful to the Department of Electrical Engineering and Computer Sciences, IISER Bhopal, for providing the necessary infrastructure, facilities, and a supportive environment to carry out my research. Special thanks are due to my colleagues, whose discussions and collaborations have played an important role in shaping my research.

I would like to express my deepest gratitude to my family. To my parents, Anita Agrawal and Sajan Kumar Agrawal, who stood by my side and faced countless hardships so that I could pursue my dream; to my sisters, Nisha Agrawal, Pooja Agrawal, Rinky Agrawal, Pinky agrawal, and Reena Agrawal and my brothers Sourav Agrawal, Jay Agrawal, Bimal Agrawal, and Kamal Agrawal for their constant love, guidance, and understanding; to my paternal aunts (Mrs. Rupa Bhalotia, Mrs. Meena Agrawal, and Mrs. Sangeeta Agrawal) and paternal uncles (Mr. Rajesh Agrawal and Mr. Pradeep Kumar Agrawal), who took care of my parents during difficult times and to my grand parents Late Mrs. Draupdi Agrawal and Mr. Ram Kumar Agrawal for their persistent affection, allowing me to be curious and focus on my career as the first doctorate in my family. I am equally grateful to my friends—Shivam, Anubhav, Razaul, Radhika, Ayush, Garima, Suyash, and my IISER Bhopal BS-MS batch of 2015—for their unwavering support, for keeping me cheerful, and for making the right decisions throughout this journey.

Lastly, I would like to thank nature (The Stars, The Moon, The Sun, The Sky, The

Land, The Air, The Water, and The Fire) that constantly supply energy, passion, and ideas, and cooperate with me to draw inspiration from it. I will be forever grateful for the love, support, and trust in me, without which my doctoral expedition would have been impractical.

ABSTRACT

Conventional intelligent and sophisticated robotic systems have proven effective in performing targeted tasks. However, high cost and limited scalability restrict their applicability for applications such as in-vivo medical treatment, wildfire management, military operations, and other safety-critical domains. Large-scale multi-robot systems offer a promising alternative; nevertheless, their design and deployment remain challenging due to the added complexity of pre-planning missions, ensuring coordination while performing tasks, and controlling accessibility. Animal and insect groups in the wild, in contrast, manage to perform complex tasks, such as foraging for food, building nests, avoiding predators, and constructing living structures, like the living bridge formation in army ants. Despite their physiological simplicity, limited intelligence, and local interactions, animals and insects outperform their individual capabilities through exploiting emergent phenomena when working in groups. This motivates the development of scalable animal behavior-inspired multi-robotic systems known as robot swarms, capable of autonomous operation, offering robustness to external disturbances, scalability via decentralization, and efficient human teaming.

Swarm robotics, inspired by the collective behaviors of animal groups, offers a compelling paradigm for addressing real-world challenges. The thesis proposes a framework grounded in physical constraints for deploying large-scale robotic swarms for various applications. I develop swarm intelligence–inspired solutions grounded in biological principles of collective behavior, utilizing statistical and game-theoretic approaches. The framework abstracts any application into four categories: (i) collective decision-making, (ii) collective navigation, (iii) collective exploration, and (iv) collective manipulation.

For collective decision-making, the thesis proposes a novel response threshold-based adaptive approach for resource-constrained robots to solve the best-of-n problem. For collective navigation, the thesis introduces a novel mix-strategy-based model for seamless robot flocking with anonymous and emergent leadership. For collective exploration, this thesis addresses the problem of intruder detection using multi-robot systems, which are constrained by a limited sensing range, as a case study for the applicability of robot swarms in the real world. The rigorous analysis made in studies outlines exciting possibilities, and the outcomes project how to overcome complex challenges in the domain of

multi-robotic platforms for real-world applications in simple yet efficient ways.

TABLE OF CONTENTS

Certificate iii

Academic Integrity and Copyright Disclaimer iv

Acknowledgement v

Abstract vii

Table of Contents ix

List of Figures xiii

List of Tables xviii

List of Publications xviii

Chapter 1: Introduction 1

1.1 Motivation 2

1.2 Applications and Challenges 3

1.2.1 Collective decision-making 3

1.2.2 Collective motion 5

1.2.3 Collective exploration 5

1.2.4 Collective manipulation 7

1.3 Novelty and Contributions 9

1.3.1 Contribution 1: Collective decision-making 10

1.3.2 Contribution 2: Collective motion 12

1.3.3 Contribution 3: Collective exploration 13

1.4 Outline of the Thesis 13

Chapter 2: Response Threshold Based Best-of-n Decision-making in Minimalistic Robot Swarms **15**

2.1 Overview . . . 15

2.2 Introduction . . . 15

2.3 Characterization of the Problem . . . 17

2.3.1 Highest average rate of success (HARS) . . . 18

2.4 Finding the Best Mean for Response Threshold Distribution . . . 19

2.4.1 Accuracy of our prediction of best response threshold mean . . . 22

2.5 Finding the Best Standard Deviation for Response Thresholds Distribution 22

2.6 Conclusion . . . 28

Chapter 3: Adaptive Response Threshold Mechanism for Multi-Robots to Solve Best-of-n Decision-making **31**

3.1 Overview . . . 31

3.2 Introduction . . . 32

3.3 Model . . . 33

3.3.1 Optimal Threshold Placement . . . 35

3.3.2 Clarity in majority . . . 36

3.4 Adaptive Threshold Control . . . 37

3.4.1 Desired Rejection Rate . . . 38

3.4.2 Threshold Update Rule . . . 38

3.4.3 Discussion of adaptation rate . . . 38

3.5 Simulation Experiments . . . 39

3.5.1 Convergence of Thresholds . . . 39

3.5.2 Effect of Adaptation Rate . . . 42

3.5.3 Dynamic Environments . . . 43

3.5.4 Robustness to Noise and Environmental Change . . . 44

3.6 Conclusion . . . 46

Chapter 4: Impact of Individual Defection on Collective Motion **47**

4.1 Overview . . . 47

4.2 Introduction . . . 47

4.3 Methods . . . 51

4.4 Defection in the Pairwise Interaction Model 52
4.5 Defection in the Couzin Model 53
4.6 Conclusion 62

Chapter 5: Flocking Phase Transition and Anonymous Emergent Leadership via Individual Defection 67

5.1 Overview 67
5.2 Introduction 67
5.3 Flocking model of probabilistically defecting agents 69
5.3.1 Model parameters and notation 72
5.4 Results 74
5.4.1 Non-monotonic emergence of polarization under probabilistic defection 74
5.4.2 Attraction-range dependence of collective response to defection . . 75
5.5 Robustness of defection-induced transitions across group sizes 76
5.5.1 Probabilistic defection reshapes the likelihood of collective states . 78
5.5.2 Defection acts as an effective reduction in individual turning rate 79
5.5.3 Testing alternative explanations for defection-induced polarization 82
5.5.4 Reversibility of collective states under varying defection probability 84
5.6 Neighborhood statistics and local composition do not explain defection-induced polarization 86
5.7 Diagnostic tests for recruitment and influence of defectors 87
5.7.1 Null model 88
5.7.2 Comparison with null expectations 90
5.8 Conclusion 92

Chapter 6: Multi-robot Searching with Limited Sensing Range for Static and Mobile Intruders 93

6.1 Overview 93
6.2 Introduction 93
6.3 Related Work 94
6.4 Problem Formulation 95
6.5 Hardness Results 96

6.6 Algorithmic Methods 98
6.6.1 Space-filling-curves approach 99
6.6.2 Random search approach (RS) 102
6.6.3 Cooperative random search approach (CRS) 103
6.7 Simulation Results 104
6.7.1 Simulation setup 105
6.7.2 Baseline 105
6.7.3 Effect of increase in the number of search robots 106
6.7.4 Effect of change in the shape of the polygon 108
6.7.5 Effect of increase in area of the polygon 109
6.7.6 Effect of increase in the number of spikes 111
6.8 Conclusion 111

Chapter 7: Conclusion and Future Work **113**

References **116**

List of Figures

Figure 1.1: Challenging robot applications 2

Figure 1.2: Examples of animal collectives. 2

Figure 1.3: Voter model for collective decision-making in initially unbiased robot swarms. 4

Figure 1.4: Foraging in ants. 6

Figure 1.5: Foraging in honeybees through discovery of food source, recruitment of colony members, and exploitation until the source is depleted. 7

Figure 1.6: Collective motion 8

Figure 1.7: Collective exploration 9

Figure 1.8: Collective construction in ants 10

Figure 1.9: Collective lifting in ants through synchronized movements. 11

Figure 2.1: Heatmap of success rate as a function of μ_q and μ_h 19

Figure 2.2: Generalized extreme value theorem and predicting the placement of μ_h^* 20

Figure 2.3: Curve fitting on optimal success rate region 23

Figure 2.4: Comparing heatmaps of success rate as a function of μ_h and μ_q across various types of option quality distributions 24

Figure 2.5: Comparison of the $\mathcal{D}_h = \mathcal{N}$ (top half) and $\mathcal{D}_q = \mathcal{N}$ (bottom half) for increasing ratio of σ_h to σ_q when $\mu_h = \mu_q$ (dashed distribution) and when $\mu_h = \mu_h^*$ (solid distribution). The green and blue dots show the best (q') and second-best (q'') value out of 50 sampled points in 50 trials. The red dashed line is μ_h^* computed using Eq. (2.1). 25

Figure 2.6: Curve fitting on optimal success rate region in μ_q versus μ_h heatmap as a function of individuals per option 25

Figure 2.7: Curve fitting on optimal success rate region in σ_q versus σ_h heatmap as a function of number of options 26
Figure 2.8: Curve fitting on optimal success rate region in σ_q versus σ_h heatmap as a function of number of individuals per option 26
Figure 2.9: Curve fitting on σ_h versus σ_q for optimal success rate 27
Figure 2.10: Success rate as a function of changing number of options and comparison with naive choice for μ_h^* 29

Figure 3.1: Schematic of best-of-n problem using response threshold method . 34
Figure 3.2: Clarity in majority through placement of μ_h^* 36
Figure 3.3: Block diagram of individual robot behavior 39
Figure 3.4: Adaptive response of one robot using $h-$controller 40
Figure 3.5: Collective convergence across various types of options quality distribution 41
Figure 3.6: Comparison of the decision success rate ρ between the proposed adaptive mechanism and fixed response thresholds for different numbers of options n. We consider fixed response thresholds whose distribution matches the environmental distribution [1], i.e., $\mathcal{N}_h(\mu_h, \sigma_h) = \mathcal{N}_q(\mu_q, \sigma_q)$, and fixed thresholds whose distribution matches the optimal distribution(Chapter. 2), i.e., $\mathcal{N}_h(\mu_h, \sigma_h) = \mathcal{N}_h^*(\mu_q^*, \sigma_q^*)$. Results are computed from 100 simulation runs per n and $m = 5000$. 41
Figure 3.7: Evolution of the clarity in majority 42
Figure 3.8: Effect of adaptation rate α on (a) the threshold distribution and (b) the decision success rate. Smaller α values lead to slow adaptation and narrow distributions, while large values induce heterogeneity and instability. Decision success ρ peaks at intermediate α values where the swarm balances exploration and exploitation for $n = 35$, $\mathcal{N}_q$, $v = 10$, $m = 5000$. 43
Figure 3.9: Evolution of the response threshold (green line) of an agent with $\log v = 10$ in an environment with dynamically changing n when (a) $\sigma_e \neq 0$, and when (b) $\sigma_n \neq 0$. The orange dashed line is the μ_q overlaid on the shaded region of σ_q. The red lines depict μ_h^*. . . . 44

Figure 3.10: Evolution of the response threshold distribution $\mathcal{D}_h$ (green line and shaded region depict average and standard deviation) for a swarm of $m = 5000$, note n changes and thus, $|g_i|$ agents in a dynamic environment. The orange steps and shaded region depict the stimulus quality distribution $\mathcal{D}_q$. The red dotted line depicts μ_h^*. 45

Figure 4.1: Illustration of probabilistic defection 50

Figure 4.2: Polarization time and energy cost as a function of probability to defect 53

Figure 4.3: The time lapse of a collective of size $N = 50$ consisting of agents that defect with a certain probability of defection, P_D, and interact through a pairwise interaction model. In this figure, 1.0 steps correspond to 1.0 *sec*. 54

Figure 4.4: Types of defection in the Couzin model 55

Figure 4.5: Probability of group getting polarized as a function of probability to defect in the Couzin model 57

Figure 4.6: The polarization time T and the average energy cost E across four types of defection for a group of $N = 20$ agents. The violin plot shows the distribution of values from 100 simulation runs. 57

Figure 4.7: The polarization time T and the average energy cost E across four types of defection for a group of $N = 50$ agents. The violin plot shows the distribution of values from 100 simulation runs. 58

Figure 4.8: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_{all})$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec*. 60

Figure 4.9: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_a)$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec*. 61

Figure 4.10: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_o)$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec.* 62
Figure 4.11: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_r)$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec.* 63
Figure 4.12: The plot shows the probability of $p(T) > 0.8$ as a function of the probability of defection. The solid lines are for $N = 20$, the dashed lines are for $N = 50$, and the dotted lines are for $N = 80$. 64
Figure 4.13: The polarization time (T) and the average group energy cost . . . 64

Figure 5.1: Individual defection in collective motion 70
Figure 5.2: Alignment and milling time series 74
Figure 5.3: Alignment and milling time series heatmaps for changing R_a . . . 76
Figure 5.4: Alignment and milling time series heatmaps for changing flock size 77
Figure 5.5: Probability of alignment and milling 80
Figure 5.6: Probability to defect induces reduction in effective turning rate . . 83
Figure 5.7: Alignment time series for probability to defect and ω 84
Figure 5.8: Reversibility of collective states under gradual variation of the defection probability P_D. 85
Figure 5.9: Density of neighbors in zones 88
Figure 5.10: Fraction of neighbors in different interaction zones as a function of defection probability. 89
Figure 5.11: Defector influence diagnostics compared against null models. . . . 91

Figure 6.1: Schematic of multi-robot intruder searching problem 97
Figure 6.2: Comb structure for simple orthogonal polygon 97
Figure 6.3: (a) Rectangular decomposition of simply connected rectilinear polygons, (b) Splitting of rectangle for computing space-filling curves . 100
Figure 6.4: Illustration of space-filling curves approach 101
Figure 6.5: Snapshot of cost-map in random search approach 104

Figure 6.6: Spikes counting in simple orthogonal polygon 106
Figure 6.7: Snapshot from cooperative random search for a static intruder. Some grids appear dark due to the accumulation of revisiting costs during the waiting time of search robots until all search robots reach their goal. 107
Figure 6.8: Effect of increase in the number of searchers on search steps for different search strategies on 4-spiked polygon. (a) for static intruder case and (b) for moving intruder case. 108
Figure 6.9: Effect of change in the shape of the polygon (area 4400 $unit^2$ and $\beta = 5$) on search steps for different search strategies and $k = 13$ search robots, (a) for static intruder case and (b) for moving intruder case. 109
Figure 6.10: Effect of change in the area of the polygon (fixed shape and $\beta = 5$) on search steps for different search strategies and $k = 10$ search robots, (a) for static intruder case and (b) for moving intruder case. 110
Figure 6.11: Effect of increase in the number of spikes on search steps for different search strategies and $k = 25$ search robots, (a) for static intruder case and (b) for moving intruder case when there was 160 total number of cells. 110

List of Tables

Table 4.1: Parameters used in the Couzin model for highly parallel formation. 56

Table 4.2: Mean μ_e and standard deviation σ_e of the average energy cost E for D-Z_r, same data as in Fig. 4.7(h). 58

Table 4.3: Mean μ_e and standard deviation σ_e of the average energy cost E for $D - Z_a$, same data as in Fig. 4.7(d). 59

Table 4.4: Mean μ_t and standard deviation σ_t of the polarization time T for $D - Z_{all}$, same data as in Fig. 4.7(a). 59

Table 5.1: List of fixed individual parameters (for highly aligned formation based on [2]) used in this study: 72

List of Publications

1. Swadhin Agrawal, Sujit Pedda Baliyarasimhuni, and Andreagiovanni Reina (2022). Response Threshold Distributions to Improve Best-of-N Decisions in Minimalistic Robot Swarms. In: Dorigo, M., et al. Swarm Intelligence. ANTS 2022. Lecture Notes in Computer Science, vol 13491. Springer, Cham. https://doi.org/10.1007/978-3-031-20176-9_32

2. Swadhin Agrawal, Jitesh Jhawar, Andreagiovanni Reina, Sujit Pedda Baliyarasimhuni, Heiko Hamann, and Liang Li (2024). Impact of Individual Defection on Collective Motion. In: Hamann, H., et al. Swarm Intelligence. ANTS 2024. Lecture Notes in Computer Science, vol 14987. Springer, Cham. https://doi.org/10.1007/978-3-031-70932-6_10

3. Swadhin Agrawal, Sujoy Bhore, Joseph S. B. Mitchell, P. B. Sujit, and Aayush Gohil, "Multi-robot Searching with Limited Sensing Range for Static and Mobile Intruders," 2025 International Conference on Emerging Technology in Autonomous Aerial Vehicles (ETAAV), Bangalore, India, 2025, pp. 1-7. https://doi.org/10.1109/ETAAV66793.2025.11213328.

4. Swadhin Agrawal, Sujit Pedda Baliyarasimhuni, and Andreagiovanni Reina. Adaptive Response Threshold Mechanism for Multi-Robots to Solve Best-of-n Decision-making. Submitting to Swarm Intelligence

5. Swadhin Agrawal, Jitesh Jhawar, Andreagiovanni Reina, Sujit Pedda Baliyarasimhuni, Heiko Hamann, and Liang Li; Emergence of leadership in collective motion with mixed strategy. Submitting to Journal of Royal Society Interface

CHAPTER 1

INTRODUCTION

The use of robots to respond to diseases and disasters—both natural and man-made—and to serve social and commercial needs has become increasingly realistic [3–5]. Advances in sensing, actuation, computation, and autonomy are positioning robots as a central technological backbone for hazard mitigation and the execution of repetitive or high-risk tasks in the modern world. This shift reduces human exposure to danger while enhancing productivity, accessibility, and operational efficiency [6, 7].

A wide range of sophisticated robotic platforms operating on land, in water, and in the air are already available across different domains and scales. Examples include large-scale systems such as aerial drones by DJI [8], humanoid robots by Neo [9], Figure [10], and Unitree [11], quadrupeds by Boston Dynamics [12] and Unitree [11], fish-like robots [13], autonomous underwater and surface vehicles [14–16], and industrial manipulators by KUKA [17]. At the opposite end of the spectrum, miniature and nano-scale systems include RoboBees [18], Xenobots [19], Black Hornet micro-drones [20], RoboFly [21], micro-scallop robots [22], and ingestible robotic pills developed by Rani Therapeutics [23]. Despite their diversity and sophistication, individual robots are inherently constrained by their size, energy capacity, sensing range, and computational resources [4, 5, 21, 24]. Moreover, physical damage or unforeseen environmental conditions can abruptly compromise an entire mission. Consequently, deploying a single robot for high-priority and complex tasks—such as those illustrated in Fig. 1.1—remains unreliable.

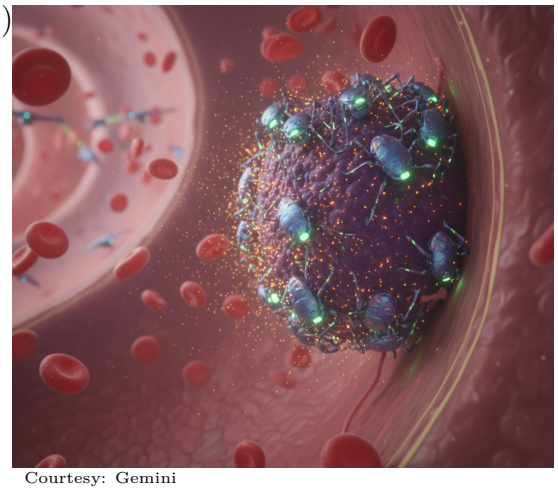
Motivated by these limitations and risks, this thesis adopts a physically grounded multi-agent paradigm, commonly referred to as swarm intelligence [4, 25, 26], inspired by collective behaviors observed in eusocial animal groups. By leveraging redundancy, decentralization, and local interactions, swarm systems offer enhanced robustness, scalability, and adaptability. This thesis, with a focus on addressing real-world robotic challenges, introduces novel mechanisms for collective decision-making in foraging, flocking, and target search applications using robot swarms.

(a)



Courtesy: Gemini

(b)



Courtesy: Gemini

Figure 1.1: Applications of human-robot teams at various length scales for (a) disaster management, and (b) in-vivo medical treatment.

(a) Ant foraging (b) Honey Bee (c) Fish shoal (d) Starling flock



Courtesy: Gemini Courtesy: Gemini Courtesy: armiblue Courtesy: James Crombie

Figure 1.2: Examples of animal collectives.

1.1 Motivation

Animals in the wild continually adapt to changing environmental conditions to survive, and many species have evolved to live and operate in groups. Notable examples include ants [27, 28], fireflies [29], honeybees [30], fish [31], birds [32–34], and wasps [35]. Despite the physiological simplicity, limited intelligence, and constrained capabilities of individuals, collective organization enables these animals to robustly accomplish complex tasks essential for survival. Such tasks include nest and structure construction [36–40], collective foraging [28, 30, 41], coordinated motion [2, 32, 34, 42–45], cooperative transport and lifting [46, 47], predator avoidance [48–50], information dissemination for danger avoidance [51], and task allocation and division of labor [52–54]. Extensive studies have shown that these group-level capabilities emerge from local interactions among individuals, without centralized control [55]. Such phenomena are collectively studied under the framework of swarm intelligence. The efficiency, robustness, and adaptability of these natural collectives have strongly motivated roboticists to draw inspiration from social

insects and animal groups exhibiting swarm intelligence [4, 56].

In swarm robotics, relatively simple embodied robots [4], equipped with basic sensing, computation, and communication capabilities, are designed to interact locally according to prescribed rules, giving rise to desired self-organized collective behaviors. These robots operate autonomously and in a decentralized manner, collectively forming a robot swarm that exhibits swarm intelligence [4]. The resulting emergent behaviors are inherently resilient to external disturbances and environmental uncertainties [7,28,57–59]. Moreover, the decentralized architecture facilitates scalability [60–63], while redundancy ensures robustness against individual failures. Beyond improved reliability, such redundancy enables task parallelism, reduces collective execution time, and enhances overall system efficiency [6, 28, 64]. Motivated by the simplicity, modularity, and resilience inherent to swarm intelligence, this thesis aims to advance and promote the deployment of robot swarms across a broad spectrum of real-world applications.

1.2 Applications and Challenges

A wide range of repetitive and hazardous applications exists for robotic systems, including forest fire management, in vivo drug delivery, disaster response, commercial and social services, and defense operations. In swarm-intelligence–based systems, collective behavior arises through self-organization; however, a systematic methodology for identifying local interaction rules that yield a desired emergent behavior remains an open challenge [5]. Many such applications can be decomposed into subtasks that closely resemble collective behaviors observed in animal groups, for which local interaction mechanisms are often well understood. This correspondence enables the classification of subtasks into distinct categories of collective behaviors, thereby facilitating the principled design of individual robot behaviors. These categories include collective decision-making and consensus formation, coordinated collective motion, collective exploration, and cooperative manipulation.

1.2.1 Collective decision-making

Consider a representative application, such as a forest fire response. A fire-affected region can be partitioned into multiple sub-regions of varying criticality: some may contain

trapped animals, while others may exhibit rapidly spreading fire fronts. Each subregion can be regarded as an option for the robot swarm to act upon, with differing priorities. The subregion containing trapped animals, for example, would constitute the highest-priority option. More generally, complex missions can be decomposed into multiple sub-tasks, each associated with a distinct level of importance. When a mission requires the swarm to prioritize and execute the most critical sub-task, individual agents must collectively identify the highest-value option and reach consensus (see Fig. 1.3). This problem is commonly referred to as the best-of-n decision-making problem [1,65–68], and is widely observed in insect societies such as ant and honeybee colonies [27,28,30,62,69,70].

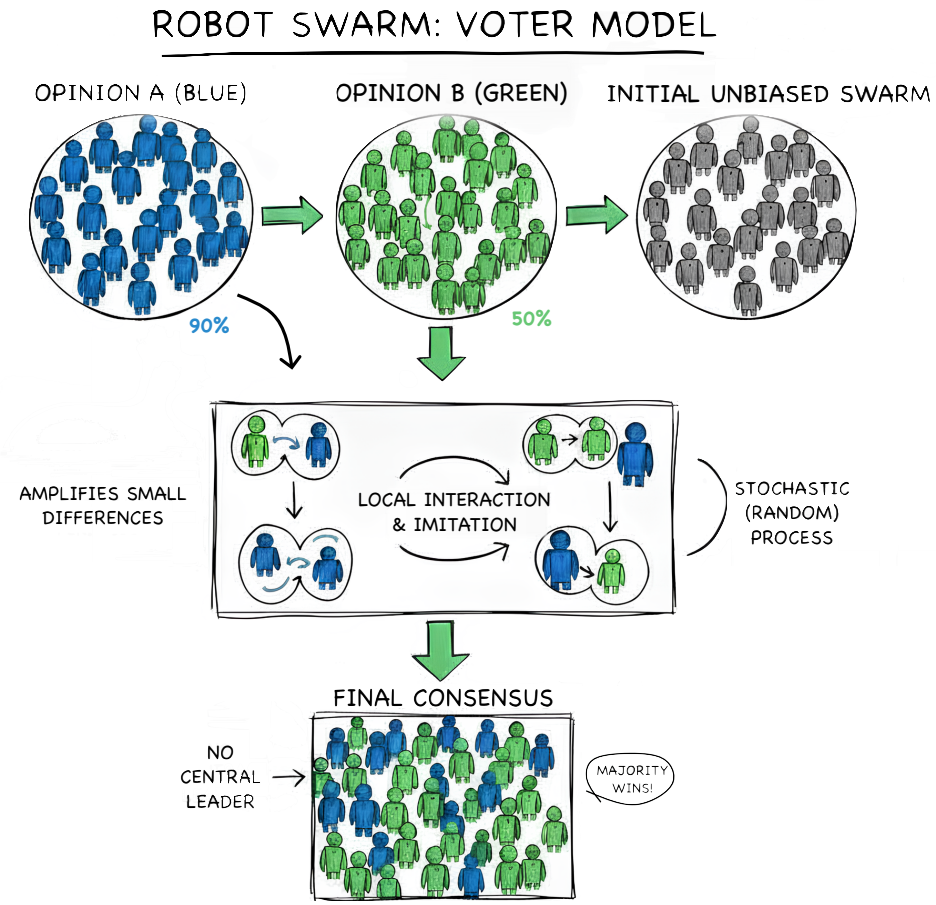


Figure 1.3: Voter model for collective decision-making in initially unbiased robot swarms.

During foraging, ants deposit pheromone trails that guide nest-mates toward food sources (Fig. 1.4), where the pheromone concentration encodes both the quality and distance of the resource. Similarly, honeybees communicate discovered food sources through waggle dances at the nest (Fig. 1.5), with higher-quality and closer sources recruiting

workers more rapidly. Such mechanisms allow colonies to reliably converge on the most rewarding option. Analogously, in applications such as disaster response or in vivo drug delivery, robot swarms must collectively identify and act upon the region or object of highest utility.

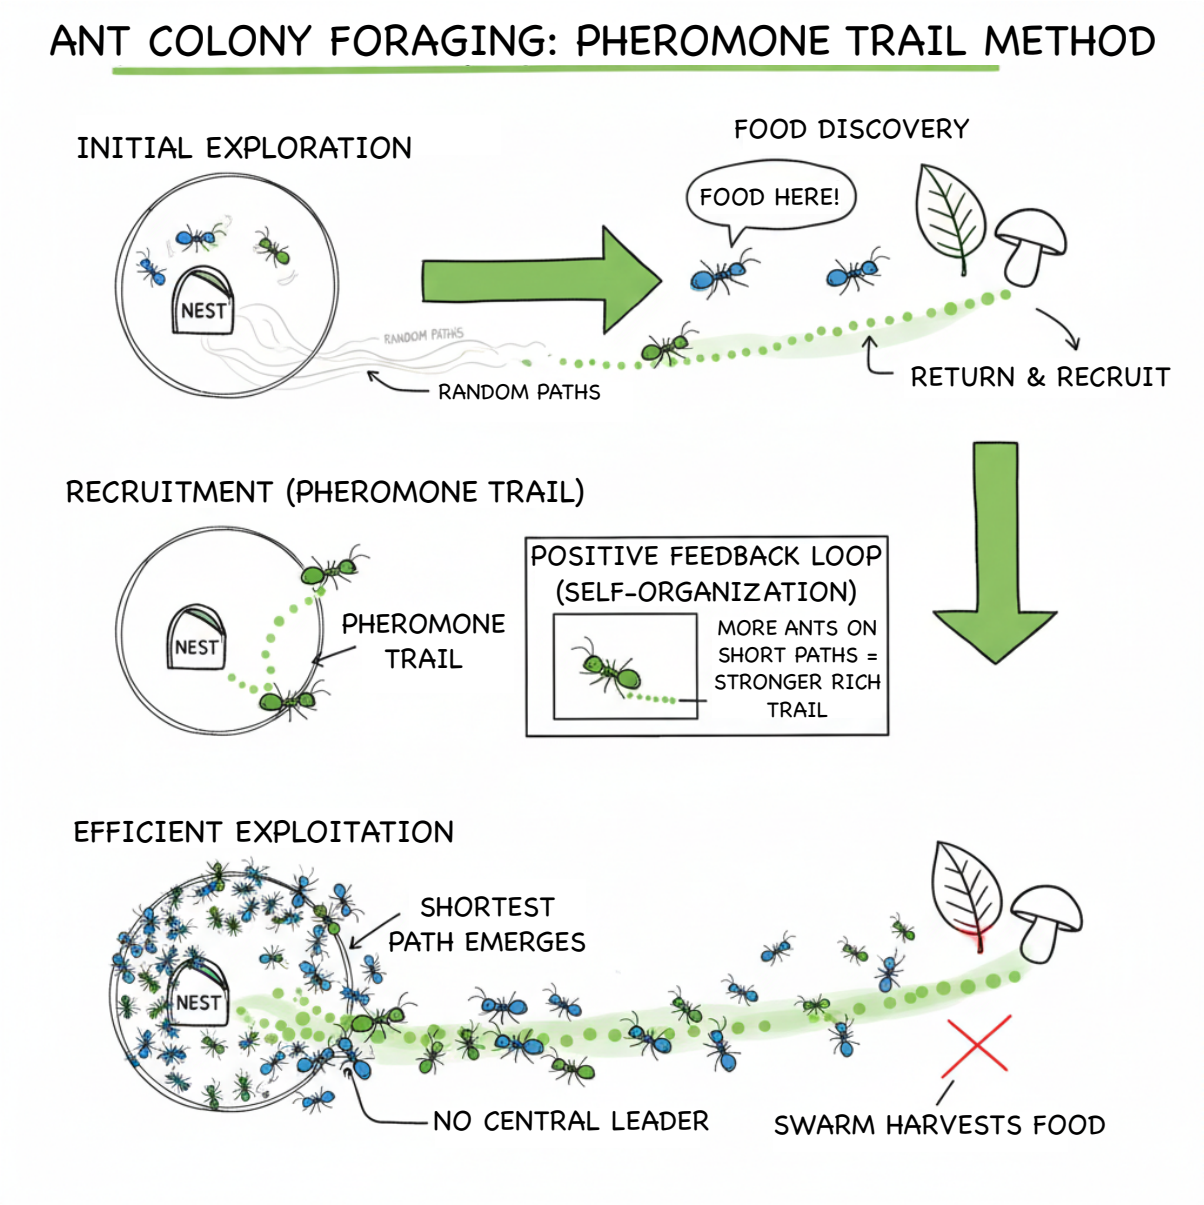
Foraging has been extensively studied across biology, social sciences, economics, and computer science [1, 65, 71, 72]. Among the various models, from the robotic application perspective, the voter model is one of the most widely adopted in swarm robotics [62, 65, 73, 74]. However, at extremely small length scales—such as in nano-robotic systems—the limited computational, communication, and operational resources of individual agents significantly restrict the implementation of existing decision-making mechanisms [75, 76]. Furthermore, most existing approaches do not explicitly account for the integration of human operators into the collective decision-making process.

1.2.2 Collective motion

In many applications, a robot swarm must collectively navigate toward a target region. Rather than relying on computationally intensive, AI-driven control mechanisms at the individual level, simple local interaction rules—such as distance-based attraction and repulsion (Fig. 1.6) have shown promise to achieve coordinated motion. Such flocking behaviors are widely observed in bird and fish schools during migration and foraging [2, 42, 77, 78]. In disaster scenarios, communication infrastructure is often damaged or unavailable, while in medical applications, robots may be too small to carry reliable communication devices. These settings necessitate minimal and intermittent involvement of a human operator, limited to essential monitoring and intervention. Existing approaches, such as shepherding [44, 79–81] and leadership-based models [7, 43, 82, 83], typically rely on continuous guidance from a human operator or pre-assigned leader/predator robots. Such reliance introduces vulnerabilities, including susceptibility to adversarial attacks targeting designated leaders/predators [84].

1.2.3 Collective exploration

Many applications require searching for static or dynamic objects of interest within a given region. Examples include security patrolling, detection of concealed explosives, border

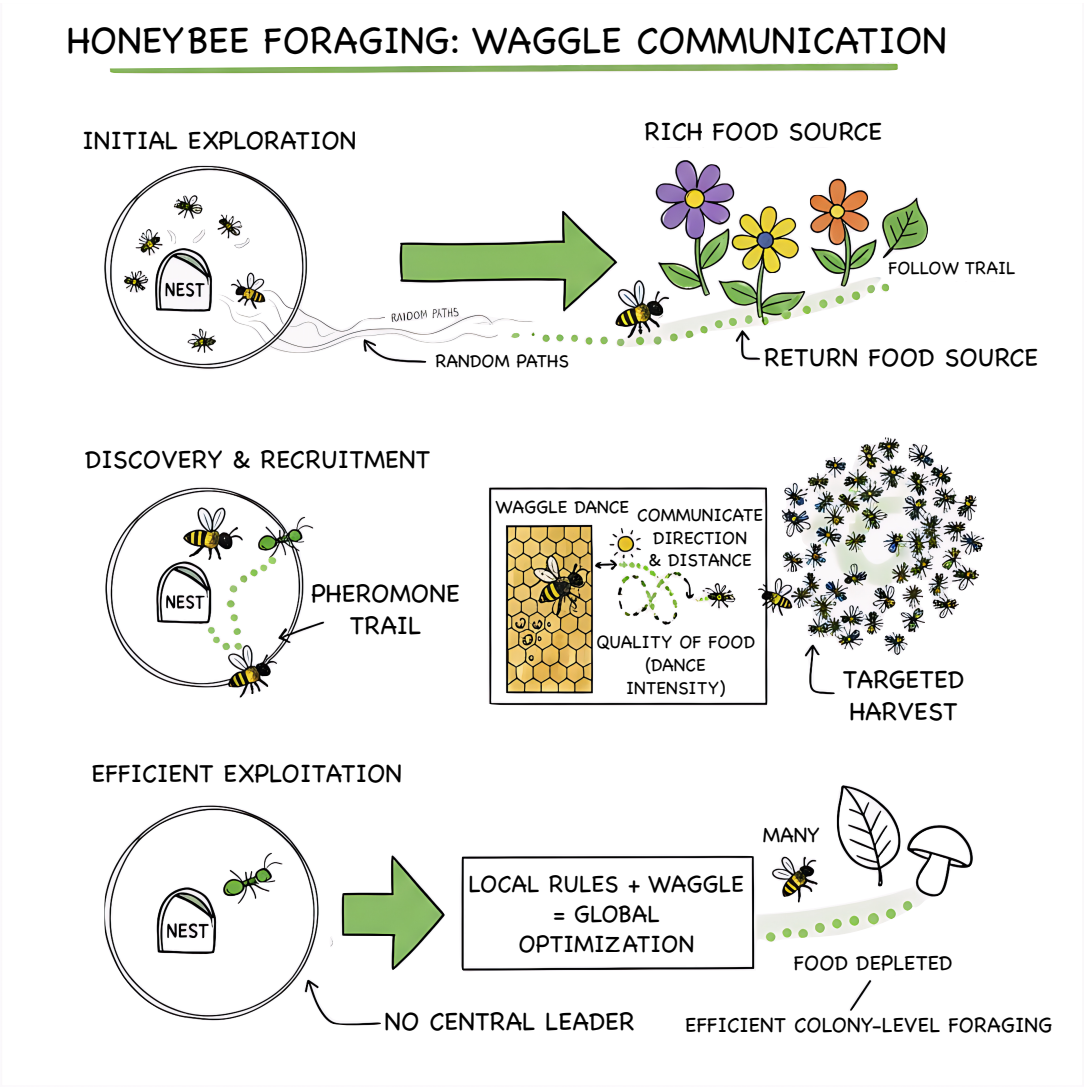


Courtesy: Gemini

Figure 1.4: Foraging in ants.

defense, and medical applications such as locating affected tissues. In such tasks, robots must collectively explore, scan, and monitor an environment. Exploration behaviors are commonly observed in social insects such as ants and honeybees during foraging and nest relocation [69, 85, 86], where individuals rely on random walks combined with stigmergic cues and local social interactions [51].

Most existing robotic exploration strategies assume unrealistically large or infinite sensing ranges and employ systematic scanning patterns akin to lawnmower coverage [87]. In practice, robots possess limited sensing capabilities [88], and environmental factors, such as narrow passages and complex geometries, can significantly influence exploration



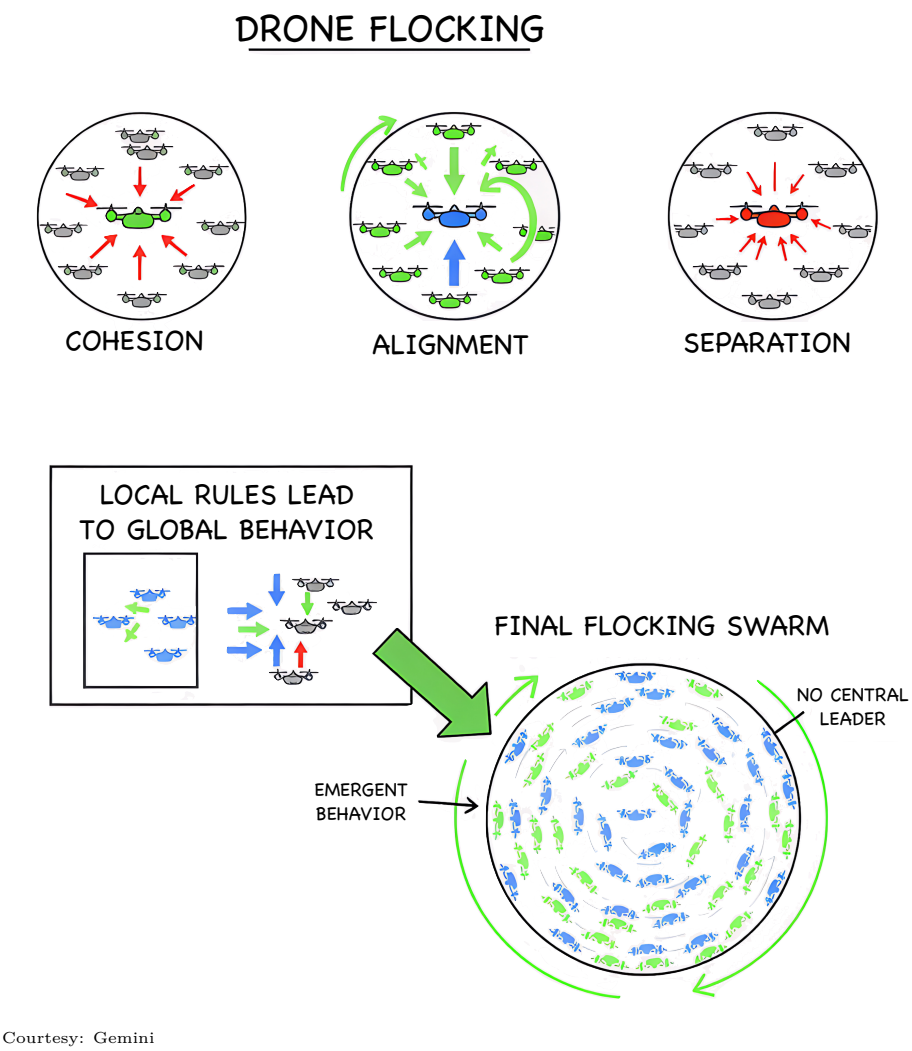
Courtesy: Gemini

Figure 1.5: Foraging in honeybees through discovery of food source, recruitment of colony members, and exploitation until the source is depleted.

efficiency and time. The effectiveness of such classical approaches, relative to animal-inspired exploration strategies, remains insufficiently understood under these realistic constraints.

1.2.4 Collective manipulation

Collective manipulation represents another important class of swarm behaviors. Ants are the only known insect species capable of cooperative lifting of objects (Fig. 1.9), and both ants and honeybees exhibit sophisticated collective construction behaviors, including nest building and the formation of living structures such as bridges (Fig. 1.8). Applications in disaster response, construction, and infrastructure repair can be mapped onto the collec-

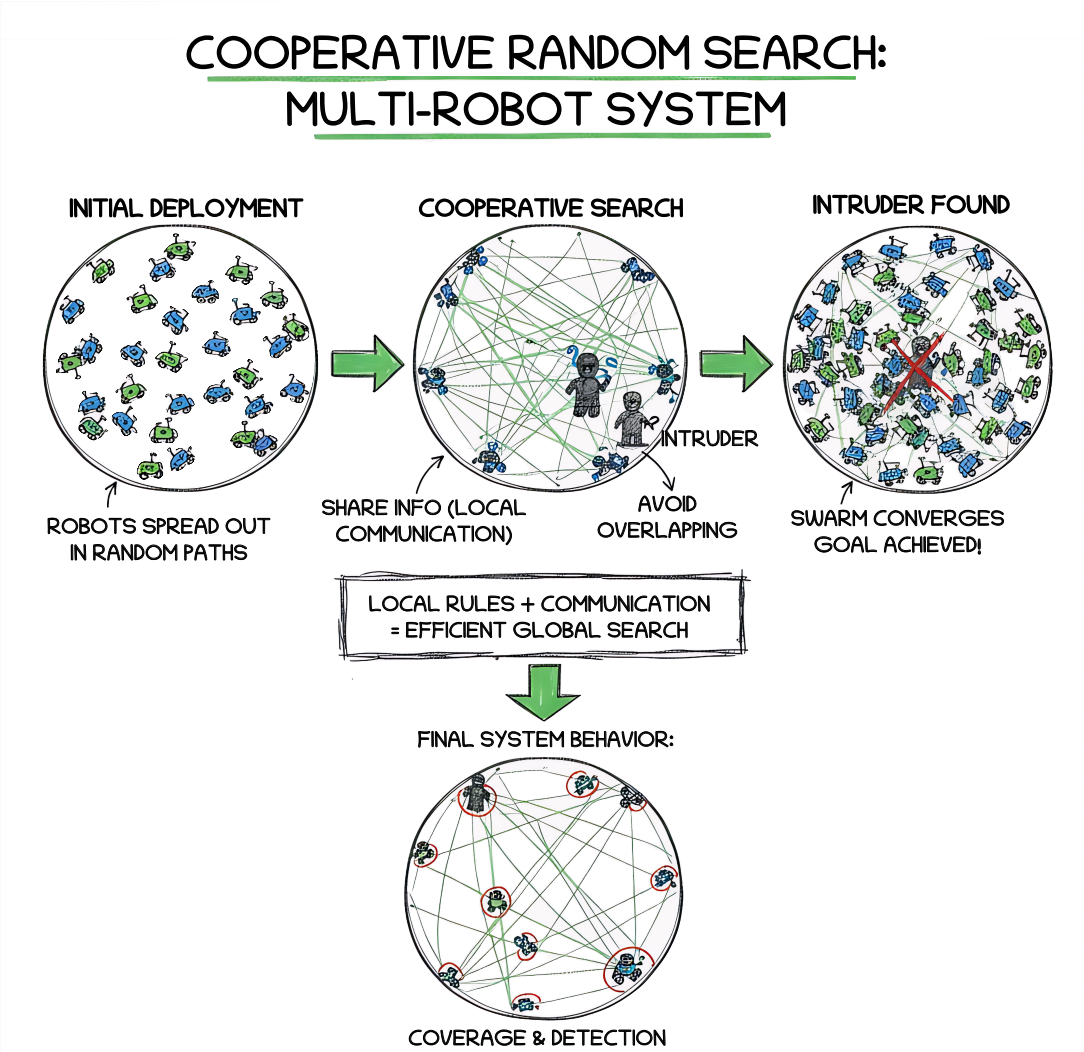


Courtesy: Gemini

Figure 1.6: Collective motion

tive manipulation behavior. In particular, strategies inspired by living bridge formation offer promising avenues for assisting humans in unstructured and extreme environments.

The diversity of real-world applications introduces significant uncertainties and imposes stringent constraints on individual agents. Some scenarios require the operation of robots in communication-deprived environments, while others demand extremely small robots, such as nano-scale agents for in vivo drug delivery. Although mapping subtasks to animal-inspired collective behaviors provides valuable design insights, not all subtasks are feasible within a swarm-robotic framework. Moreover, complex and unplanned missions often necessitate optional collaboration between robot swarms and human operators. In communication-limited and small-scale settings, such human–swarm teaming becomes particularly challenging due to limited accessibility for direct command and control.



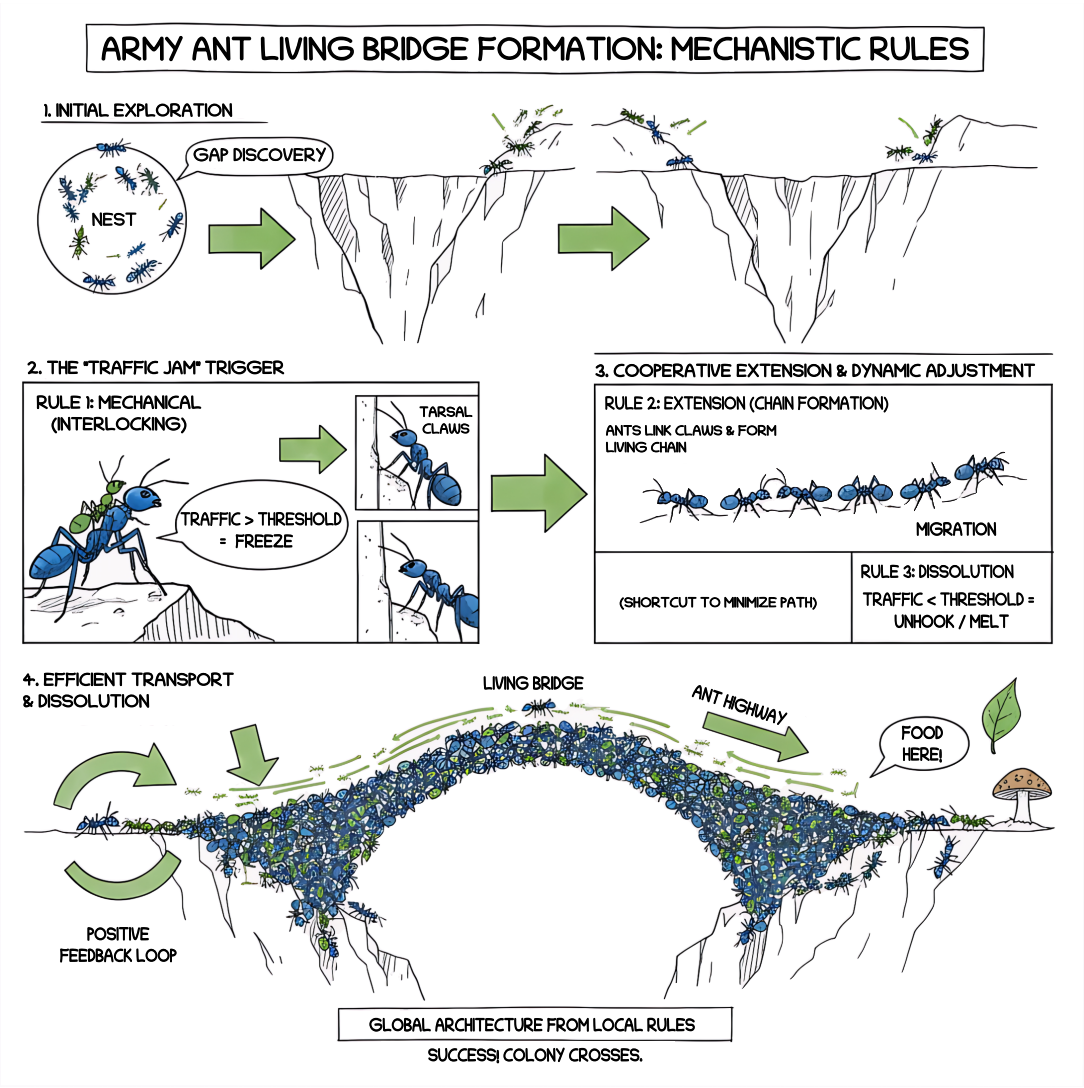
Courtesy: Gemini

Figure 1.7: Collective exploration

These challenges call for robotic platforms that are robust to single-point failures and noise, scalable through decentralization, and capable of autonomous operation in the absence of continuous human supervision. Motivated by these requirements, this thesis presents novel, physically grounded approaches for collective decision-making, collective motion, and collective exploration using robot swarms, enabling flexible and robust human–swarm teaming on demand.

1.3 Novelty and Contributions

The contributions of this thesis primarily lie in proposing solutions to the best-of-n collective decision-making problem in miniature robot swarms, introducing a novel mixed-strategy-based flocking model for collective motion, and systematically investigating mul-



Courtesy: Gemini

Figure 1.8: Collective construction in ants

tiple multi-robot target-search strategies to inform operators about their applicability under diverse operational scenarios. The main contributions on each of the above mentioned problems are as follows in the sequel.

1.3.1 Contribution 1: Collective decision-making

Existing approaches to collective decision-making in robot swarms predominantly rely on regulating the frequency or duration of information exchange among neighboring agents. Such mechanisms are resource-intensive for miniature robotic platforms. They also require agents to explicitly explore multiple options, estimate and memorize their qualities, and perform pairwise comparisons.

In this thesis, we propose a response-threshold–based collective decision-making mech-

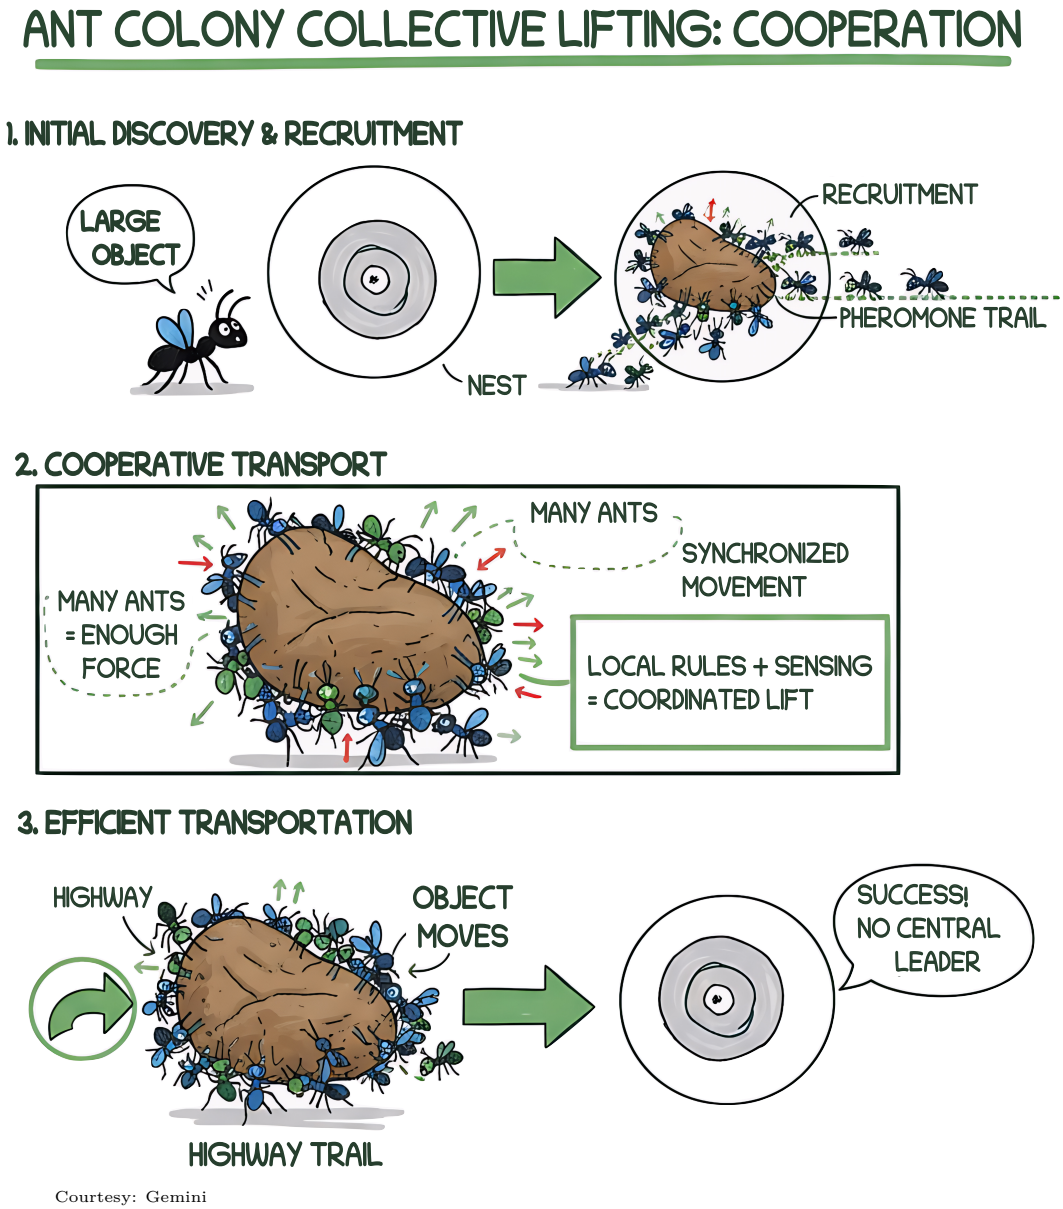


Figure 1.9: Collective lifting in ants through synchronized movements.

anism in which individual robots commit and react to any first discovered option using a binary response—*accept* or *reject*—by comparing their internal response threshold with the perceived external stimulus. Thereby removing the need to explore multiple options. We define the response threshold distribution for heterogeneous robots such that the collective accuracy is near maximum. We also define the measure of clarity in majority and propose a decentralized adaptive mechanism for the group of robots to converge to the desired best response threshold distribution. Since initial relative majority leads to the consensus on the majority opted option. By achieving high clarity in majority via adapting the response threshold distribution, we assume that the collective consen-

sus subsequently emerges through a decentralized voting process, whereby the option receiving the highest number of affirmative responses is selected.

We analytically and empirically investigate the response-threshold distributions that enable a swarm to reliably discriminate the best option from a set of alternatives. Building on this insight, we introduce a computationally minimal and fully decentralized adaptive behavioral mechanism that allows individual robots to collectively attain the desired threshold distribution. In the proposed mechanism, each robot adapts its response threshold using only a short memory of its own recent responses, thereby minimizing sensing, computation, and communication requirements while enabling scalable best-of-n decision-making.

1.3.2 Contribution 2: Collective motion

Prior studies on collective motion have modeled local interaction rules that give rise to flocking behaviors in animal groups such as fish and birds [2, 42, 77, 78], as well as in abstract particle systems [89]. These models establish fundamental principles governing group cohesion, alignment, and phase transitions, typically assuming that all individuals consistently obey the interaction rules.

To enable human-in-the-loop control, several approaches inspired by sheepdog shepherding or leader–follower paradigms introduce designated predator or leader agents that guide the swarm [7, 44, 79]. While shepherding-based methods are cognitively demanding for human operators—requiring coordinated control of multiple guiding agents—leader–follower schemes allow simpler interaction. However, both approaches rely on predefined leaders, rendering the system vulnerable to adversarial attacks or single-point failures.

In contrast, this thesis introduces a novel mixed-strategy–based flocking model in which robots dynamically alternate between cooperative and defection behaviors. Earlier studies have shown that the formations can be changed when the radius of interactions vary [2]. However, our approach allows the swarm to switch between various formations such as milling, and highly aligned state only by modulating the probability to defect. Since changing physical parameters like interaction ranges increases hardware complexities and expenses, the ability to switch formation by only via modulating the probability to defect renders the flocking mechanism efficient from engineering perspective. More-

over, as the agents do not interect while defecting, they can act as a leader or predator. Hence, our approach inherently enables a human operator to guide the collective motion without relying on pre-assigned leaders, thereby enhancing robustness, safety, and resilience to adversarial interference.

1.3.3 Contribution 3: Collective exploration

Finally, this thesis investigates the multi-robot intruder-search problem in a closed polygonal environment using a team of robots with limited sensing ranges. By devising a comb structure that generalizes to any orthogonal simply connected polygon, we proved that the multi-robot intruder searching problem (static and dynamic intruder cases) is an NP-hard problem on a grid graph. We then systematically study the relationship between intruder discovery time and swarm size across several widely used search strategies, including space-filling-curve–based exploration, random-walk searching, and cooperative random-search mechanisms.

Both static and dynamic intruder scenarios are considered to evaluate the robustness and scalability of these methods under realistic constraints and geometric properties of the search region. Through comparative analysis, this work provides actionable insights into the applicability and limitations of different multi-robot exploration strategies, informing the human operators about their suitability for deployment under varying environmental and operational conditions.

1.4 Outline of the Thesis

This thesis investigates emergent behaviors in robot swarms and develops physically grounded mechanisms for collective decision-making, motion, and exploration. Each chapter begins with a brief overview, outlines the motivation and limitations of existing approaches, formally defines the problem, and concludes with key observations and insights.

Chapter 2 formulates the best-of-n collective decision-making problem and derives the response-threshold distributions that maximize collective success. The chapter presents analytical expressions for the mean and variance of the optimal best-response threshold distribution.

Chapter 3 draws inspiration from animal foraging behaviors to propose an adaptive, computationally minimal response-threshold mechanism for robot swarms. The effectiveness of the proposed mechanism is validated through extensive Monte Carlo simulations under diverse operating conditions.

Chapter 4 examines flocking behaviors in animal groups and introduces a mixed-strategy–based modification of classical flocking models. The chapter demonstrates the advantages of incorporating defection strategies through representative parameter studies.

Chapter 5 extends this analysis to a broader parameter space, and illustrates how defectors affect the group dynamics, showing that the benefits of defection are not scenario-specific. The chapter further reveals how the defection strategy enables formation switching with the perspectives of enabling robust and anonymous human-guided collective motion.

Chapter 6 evaluates existing multi-robot target-search strategies under realistic sensing constraints to inform human operators of their suitability across different mission requirements. This chapter lays the foundation for principled design of efficient collective exploration mechanisms.

The thesis concludes by summarizing the key contributions and discussing open challenges and future research directions toward deploying robot swarms in real-world applications in Chapter 7.

CHAPTER 2

RESPONSE THRESHOLD BASED BEST-OF-N DECISION-MAKING IN MINIMALISTIC ROBOT SWARMS

2.1 Overview

For applications that require collective decision-making in a robot swarm, we aim to design algorithms that allow robots to solve the best-of-n problem using as little resources as possible. In response threshold based mechanism, each robot estimates the quality of a random option and responds with a binary response — accept or reject — depending on whether the quality is above or below its threshold. A threshold-based response approach can eliminate the need to explore all available options and quality-dependent communication frequency modulation, thereby reducing resource requirements. In this study, we aim to develop response-threshold-based solutions for the best-of-n problem, suitable for simple robots with lower memory and computational requirements than those of state-of-the-art algorithms. While the long-term goal is to implement decentralized algorithms for best-of-n decision-making based on heterogeneous response thresholds, here we focus on determining the threshold distribution that allows the swarm to best distinguish between options' qualities, thereby selecting the option with the highest quality. In this chapter, we first investigate to determine the optimal distribution of thresholds that maximizes the probability that the majority of the swarm favors the best alternative. We perform Monte-Carlo analysis across various types of environments by considering different options' quality distributions and the number of options. Our results form the basis for the Chapter. 3, which proposes a decentralized algorithm for swarms of reactive binary robots that can adapt and make best-of-n decisions.

2.2 Introduction

The design of systems composed of minimalistic units can be advantageous to operate in application scenarios where there are limitations on energy and equipment [18,19,90]. For

example, future nanorobots that operate in blood vessels must follow behaviors based on minimalistic computation due to limitations on their hardware. Similarly, environmental monitoring through biodegradable robots with a limited operational time benefits from a minimalistic design for affordable large-scale production. We aim to design minimalistic solutions for a basic form of coordination in robot swarms, i.e., best-of-n decision making, where the swarm must reach a consensus on the best option among n alternatives.

Several works investigated swarm robotics solutions for best-of-n problems [65, 91]. Existing solutions that we believe have the fewest requirements on the individual robots, in terms of communication, computation, and memory, are based on simple voting algorithms combined with quality−based frequency of communication [66, 68, 70, 85, 92, 93]. Through these methods, robots search, estimate, and memorize the quality of available nearby options, and once they find the one with the highest quality among the discovered options, they use the estimated quality to regulate their communication frequency. Robots keep a single option (and its quality) in memory and broadcast only the chosen option to nearby robots (without its quality). Robots update their opinions based on other robots' messages and reach an agreement in favor of the best option by sending messages with a frequency proportional to the self-estimated quality of their own options. Despite the robots' estimates being subject to measurement errors, this strategy allows the swarm to filter out noise and achieve high accuracy levels [94, 95].

Our hypothesis is that methods based on heterogeneous response thresholds can suffice to solve best-of-n problems, eliminating the need to explore more than one option and modulate communication based on quality, thereby removing the need to process and memorize quality measurements. Behaviors based on response thresholds have a reactive binary (yes/no) response determined by the option quality being above or below a threshold, and can be observed in several eusocial insect species [30, 54, 59, 96–98]. Despite individuals' simplicity, systems composed of response threshold units can display accurate and rational collective behavior [1, 96, 99, 100], which is enabled by a crucial element of such systems, the heterogeneity of their individuals, each having a different threshold [54, 97, 101].

The division of labor in ants, regulated by heterogeneous response thresholds, has inspired the design of several multi-robot systems to tackle task allocation problems

[53, 86, 102–107]. However, despite its potential, only limited attention has been devoted to apply response thresholds to the design of consensus decision making systems [1, 62]. While the long-term objective is the deployment of minimalist binary robots for best-of-n decisions, this work focuses solely on the relationship between threshold distribution and environmental stimuli to enhance the collective ability to distinguish between options and select the best. Understanding this relationship is the basis of the Chapter. 3 aimed at developing a decentralized mechanism for autonomous binary robots that adapt their thresholds to what is best for the given decision conditions.

2.3 Characterization of the Problem

Consider a team of robots that operate in a symmetric world with n distinct independent spatially uniformly distributed options, each characterized by a quality $\{q_1, q_2, \ldots, q_n\}$. Since options are distinct, the quality q_i of each option is represented by a random variable q associated with a probability distribution function $\mathcal{D}_q$; i.e., we assume that in a given environment the n options' quality are randomly distributed according to $\mathcal{D}_q$. We assume the robots explore randomly to discover the options. Each robot commits to the first random option it discovers as it uniformly explores the symmetric world, and then remains committed to that option. We assume that the swarm is composed of $S = n \times m$ robots. At the end of the random search, the robots are equally likely to discover any of the n options and distribute in n subgroups of similar size $\approx m$. Hence, the subgroup size m is another random variable; however, for simplicity, we fix the size of each subgroup to exactly m. After committing to an option, robots in each subgroup estimate the quality of the respective options. Each robot is characterized by a response threshold h, which is different for each robot (heterogeneous robots). The response threshold of the robots is therefore another stochastic variable with a Gaussian probability distribution $\mathcal{D}_h = \mathcal{N}(\mu_h, \sigma_h)$. The robots compare the estimated option quality $\tilde{q}_i = q_i \pm \sigma_e$ with its response threshold to either accept or reject the option (accepts if $\tilde{q}_i \geq h$, otherwise, rejects the option). Where σ_e is the standard deviation of the quality measurement error. For simplicity, but without loss of generality, we assume that the estimation errors $\sigma_e = 0$, that is, $(\tilde{q}_i = q_i)$. From the observer's perspective, in this study, we consider a decision process as successful if the most widely accepted option (i.e., the one accepted by the largest number of robots) matches the actual highest-quality option. Whenever

more than one option has the maximum number of accepting robots, we select one of the options in the tie at random. When no robots accept any option, i.e., all robots have their thresholds above the estimated quality, we consider the process as unsuccessful, as no robot will be able to vote for any option. Our vision is that robots that accept an option will engage in voting to support that option, while robots that reject the option will abstain from voting to quickly achieve a consensus on the best option in the absence of an external observer. Previous studies have shown that simple local voting mechanisms (e.g., the voter model [108,109]) can consistently lead to a consensus in favor of the option voted for by the initial relative majority [74]. While this study does not include the voting phase, we aim to obtain a proportion of acceptance across the n available options that will allow the voting process to select the best option.

In the above set-up for the best-of-n problem, our goal is to identify what values of μ_h and σ_h (i.e., which response threshold distribution $\mathcal{D}_h$) maximize the probability of decision success given a known number of options n and distribution of options' quality $\mathcal{D}_q$. Mathematically, given the number of options n and the options quality distribution, $\mathcal{D}_q$, what should be the response threshold distribution $\mathcal{D}_h$ such that the probability of decision success is maximized? Past experiments only considered Gaussian distributions of the options' qualities [1, 96, 99], however, environments with other distributions may exist. In our analysis, we consider three types of quality distributions $\mathcal{D}_q \in \{\mathcal{U}, \mathcal{N}, \mathcal{K}\}$: the uniform distribution $\mathcal{U}(\mu_q, \sigma_q)$, the Gaussian distribution $\mathcal{N}(\mu_q, \sigma_q)$, and the bimodal distribution $\mathcal{K}(\mu'_q, \mu''_q, \sigma'_q, \sigma''_q)$. The bimodal distribution $\mathcal{K}(\mu'_q, \mu''_q, \sigma'_q, \sigma''_q)$ models environments with subgroups of good or bad options, that we implement as the sum of two Gaussian functions with equal standard deviation $\sigma_q = \sigma'_q = \sigma''_q$ and mean $\mu'_q = \mu_q + \delta$ and $\mu''_q = \mu_q - \delta$ for a fixed $\delta = 2.5$. Therefore, hereafter we indicate $\mathcal{K}(\mu_q, \sigma_q)$ in terms μ_q and σ_q only. It is outside the scope of this study to implement the voting algorithm or let the robots autonomously set their thresholds; here, we focus solely on understanding which threshold distribution enhances the ability to distinguish options.

2.3.1 Highest average rate of success (HARS)

In order to study the relationship between the probability distribution functions (PDFs) of the options' qualities $\mathcal{D}_q$ and of the robots' thresholds $\mathcal{D}_h$, we perform Monte-Carlo simulations for a large set of combinations of the two PDFs (i.e., by varying their mean

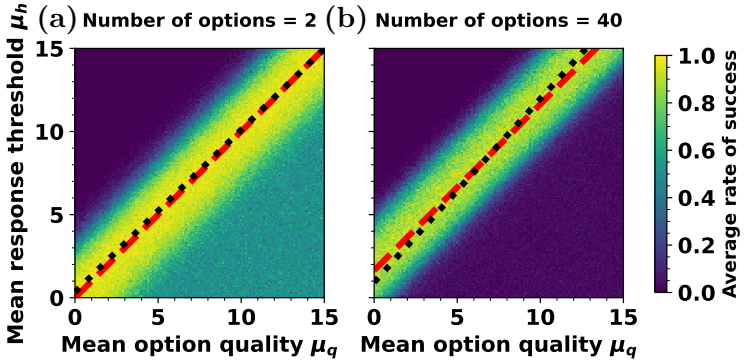


Figure 2.1: The color map shows the average rate of success (500 runs per combination) for all combinations of the mean options' quality $\mu_q \in [0, 15]$ and the mean response threshold $\mu_h \in [0, 15]$, for $\mathcal{D}_h = \mathcal{N}(\mu_h, \sigma_h)$ and $\mathcal{D}_q = \mathcal{N}(\mu_q, \sigma_q)$. The swarm size for this test was $m = 100n$, and the standard deviations were $\sigma_h = \sigma_q = 1$. We considered test cases when there were (a) $n = 2$ options and (b) $n = 40$ options. The diagonal bright band represents the region of high success. The black dotted line indicates the highest average rate of success (HARS) line, and the red dashed line is the predicted best mean μ_h^*, obtained using Eq. (2.1). Both the lines show a linear relationship between μ_q and μ_h with slope ≈ 1. The area below the HARS line in each plot indicates an average rate of success, $\approx 1/n$, as all of the n options receive a similar number of votes, which are therefore indistinguishable. Instead, the area above the HARS line in each plot indicates an average rate of success of zero as all of the n options are rejected by all the robots and no decision is made.

$\{\mu_q, \mu_h\}$ and standard deviation $\{\sigma_q, \sigma_h\}$) and computing the average rate of success (i.e., the proportion of successful decisions) for each combination. The average rate of success is displayed as color maps in Fig. 2.1. To identify which value of μ_h (or σ_h) maximizes the average rate of success for a given μ_q (or σ_q), we compute the *highest average rate of success* (HARS), which is the simplest curve (in all considered cases, a straight line) that traverses the region with the highest success rate. We compute the HARS line using a standard differential evolution method (from the SciPy library [110]) that ranks each line with the sum of success rate in all points crossed by the line (normalized by the line length) and returns the highest score line (black dotted lines in Fig. 2.1). All our simulation codes are available at [111].

2.4 Finding the Best Mean for Response Threshold Distribution

We investigate how μ_h^*—the best mean of the probability distribution of the robots' response thresholds $\mathcal{D}_h = \mathcal{N}(\mu_h, \sigma_h)$—varies for different numbers of options n and different

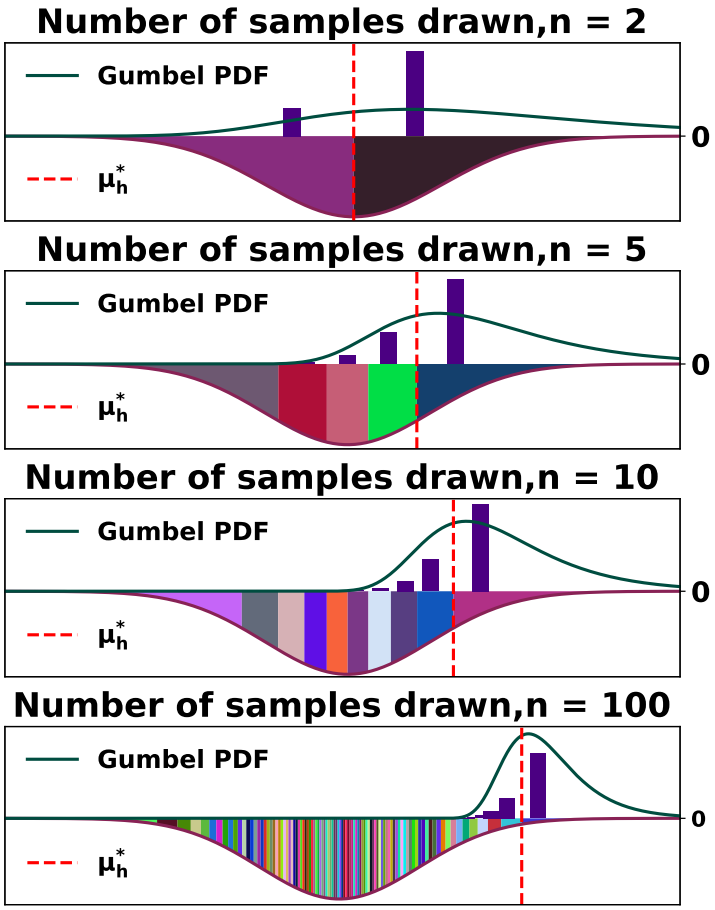


Figure 2.2: Increasing the number of samples n drawn from a PDF increases the expected value of q' (the sample with the highest value) [112]. The bottom half of each panel shows a Gaussian distribution $\mathcal{N}(\mu = 10, \sigma = 1)$ sliced in n slices with equal area in terms of CDF. Each panel's top half shows, through a histogram, the proportion of times (out of 10^3 runs) q' lies in each of the n bottom slices. The parameterized Gumbel distribution (solid line on top halves) follows the generalized extreme value distribution (GEVD) theory [112]. This is the reasoning, which led to Eq. (2.1) (red dashed line), is in good agreement with results from simulations and GEVD theory.

options' qualities distributions $\mathcal{D}_q \in \{\mathcal{U}(\mu_q, \sigma_q), \mathcal{N}(\mu_q, \sigma_q), \mathcal{K}(\mu_q, \sigma_q)\}$. While our first intuition suggested that the best results would be obtained when the thresholds' mean is equal to the qualities' mean, i.e. $\mu_h^* = \mu_q$, we find that this is true for binary ($n = 2$) decision problems (e.g. see Fig. 2.1a) but it is not the case when the number of options increases, $n > 2$, e.g. see Fig. 2.1b. Therefore, we find that both the number of options and the options quality distribution are highly relevant parameters in setting the best mean μ_h^* of the response threshold distribution.

This result, which at first glance may look counterintuitive, can be explained using the central limit theorem and the generalized extreme value distribution (GEVD) theory [112] in probability theory and statistics. When drawing a large number of quality values n from the distribution $\mathcal{D}_q$, we can expect that each of these values, on average, will be distributed according to the PDF of $\mathcal{D}_q$. Therefore, if we slice $\mathcal{D}_q$ into n sections of equal area $\frac{1}{n}$ in terms of the cumulative distribution function (CDF), on average, we expect that each of the n drawn values lies in a distinct slice. Following this reasoning, we expect that, on average, the maximum quality value (among the n qualities) lies in the last slice and that the second-to-best value lies in the second-to-last slice. Figure 2.2 illustrates this mechanism numerically for a representative case of a Gaussian $\mathcal{D}_q$; however, it also generalizes to other PDFs, in agreement with results from the generalized extreme value distribution theory [112]. Most times, the highest value among n random draws belongs to the slice with the highest value range (the right-most slice). As we are interested in distinguishing the best option from the others, ideally, a large proportion of thresholds should be between the highest quality value, which we identify with the letter q', and the second highest quality value, which we identify with q''. Therefore, the best mean μ_h^* of $\mathcal{D}_h$ is between the expected value of q' and q'' (see Fig. 2.2 and Fig. 2.5a), which we can compute through the inverse of the CDF of $\mathcal{D}_q$, as the μ_h^* that satisfies the following equation:

$$\int_{-\infty}^{\mu_h^*} \mathcal{D}_q(x|\mu_q, \sigma_q)dx = 1 - \frac{1}{n}\,, \tag{2.1}$$

where $\mathcal{D}_q(x|\mu_q, \sigma_q)$ is the PDF of $\mathcal{D}_q$ at x given μ_q and σ_q. Eq. (2.1) is dependent on both the number of options n and the options quality distribution $\mathcal{D}_q$. We hypothesize that, both the expected highest value $E(q')$ and the best threshold mean μ_h^* should increase

when n and σ_q increases, since the area of the slices decreases as n increases, whereas the range of the slices spread as σ_q increases. This result becomes more accurate as the number of options increases [112].

2.4.1 Accuracy of our prediction of best response threshold mean

Through simulations, we show that the prediction of Eq. (2.1) matches the highest average rate of success (HARS lines in Fig. 2.1 and Fig. 2.4). Fig. 2.3 shows that the obtained results generalize to a large set of conditions, for different PDFs and for different swarm sizes. The intercept of the HARS lines (Fig. 2.3a) quickly increases for low n, and it then asymptotically saturates to a constant value for large n.

Fig. 2.3b shows that the slope of the HARS lines remains constant to ≈ 1 for all tested types of quality PDFs and values of n and S. Thus, we can consider the predicted and fitted lines parallel to each other, and use the distance between them (inset of Fig. 2.3a) as the measure for the accuracy of Eq. (2.1), showing low absolute difference and good accuracy in all tested combinations. Eq. (2.1) generalizes to systems with different variability of the quality values, σ_q, as shown in Fig. 2.3c for $n = 5$, where the mean μ_h^* increases with σ_q as predicted in our hypothesis.

2.5 Finding the Best Standard Deviation for Response Thresholds Distribution

Determining the optimal standard deviation of the response threshold distribution is as important as determining its mean, because the right amount of variability can optimize the number of robots required to deal with the stochastic nature of the best-of-n decision-making process. When the thresholds' mean is far from optimal (e.g. $\mu_h = \mu_q$ for $n \gg 2$), σ_h can play an important role in discerning the highest-quality option. When σ_h is much smaller than σ_q, the probability that robots will have their response thresholds between q' and q'' is almost null, and all high-quality options will be indistinguishable during the voting phase (Fig. 2.5). Having $\sigma_h > \sigma_q$ can mitigate this problem because the thresholds are more spread out; however, it is also a waste of resources, as several thresholds are set to values much lower than necessary, and only a small percentage of robots have a

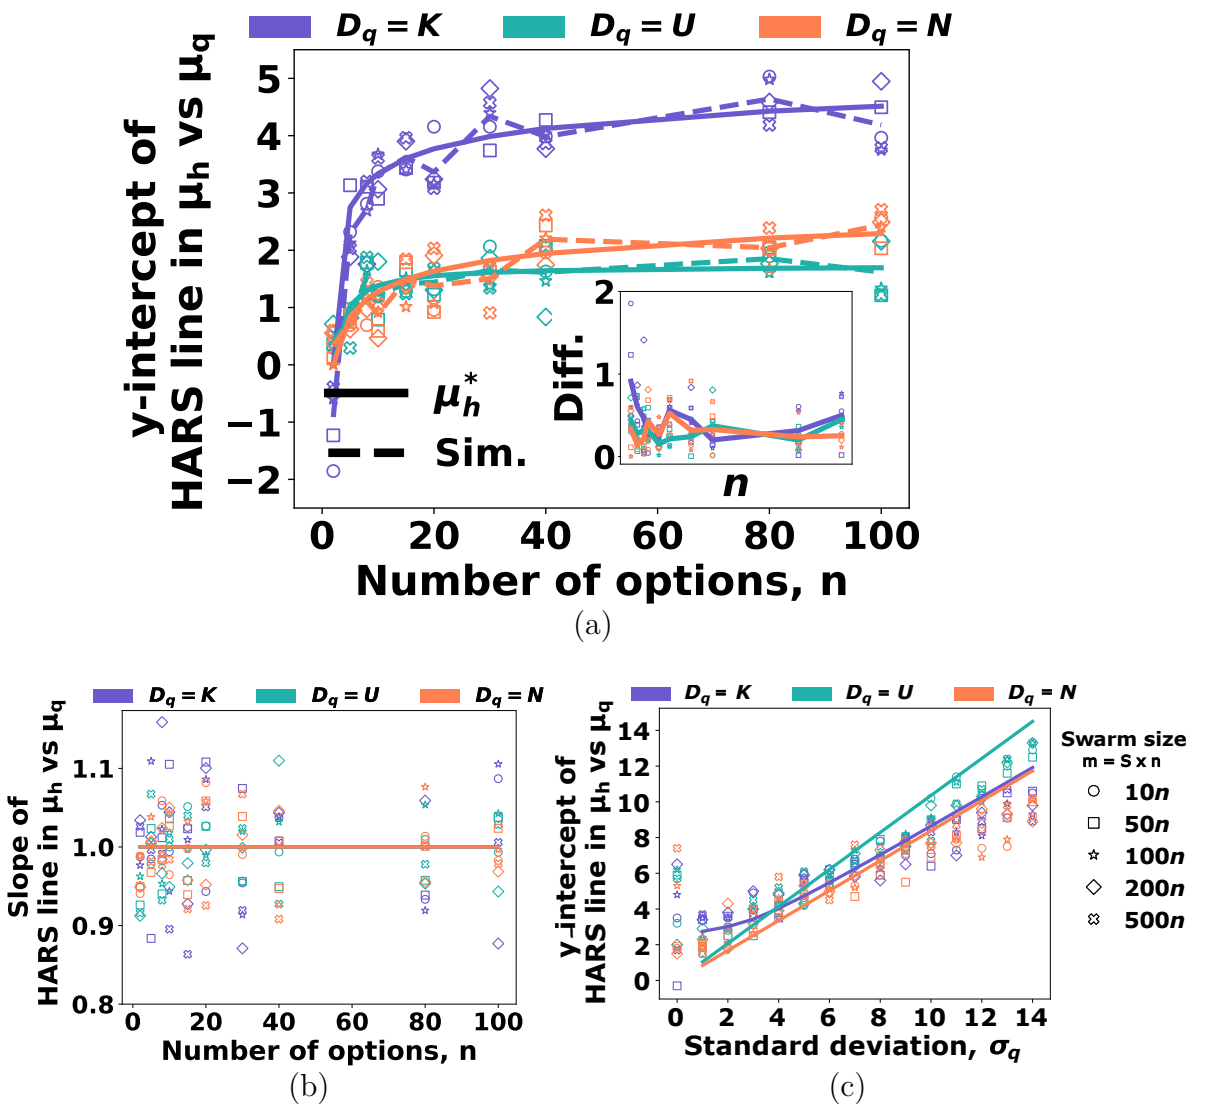


Figure 2.3: (a) $y-$intercept, and (b) slope of HARS lines for μ_h^* computed as a function of μ_q, for varying n (on x-axis) and m (markers, $S = m/n = |g_i|$ is an integer multiplier), for $\sigma_h = \sigma_q = 1$, and (c) $y-$intercept for varying std. dev. σ_q (on x-axis) for $n = 5$, $\sigma_h = 1$. In (a), the dashed lines show the average for all m and in all plots, the solid lines depict the predicted μ_h^* using Eq. (2.1). The inset of (a) shows the absolute difference between predicted and fitted HARS lines.

significant role in the collective decision-making. Differently, when μ_h is close to the optimal value μ_h^*, the influence of σ_h is much reduced. Some variability among response thresholds is always necessary; however, the standard deviation can have relatively low values ($\sigma_h < \sigma_q$), and still cover the relevant quality range, as confirmed by the fitted HARS lines of Fig. 2.9a, which always have a slope smaller than 1, for $\mu_h = \mu_h^*$.

In past studies on the division of labour [113], it has been observed that the optimal heterogeneity often takes a highly skewed form. Furthermore, as noted by [43], effective

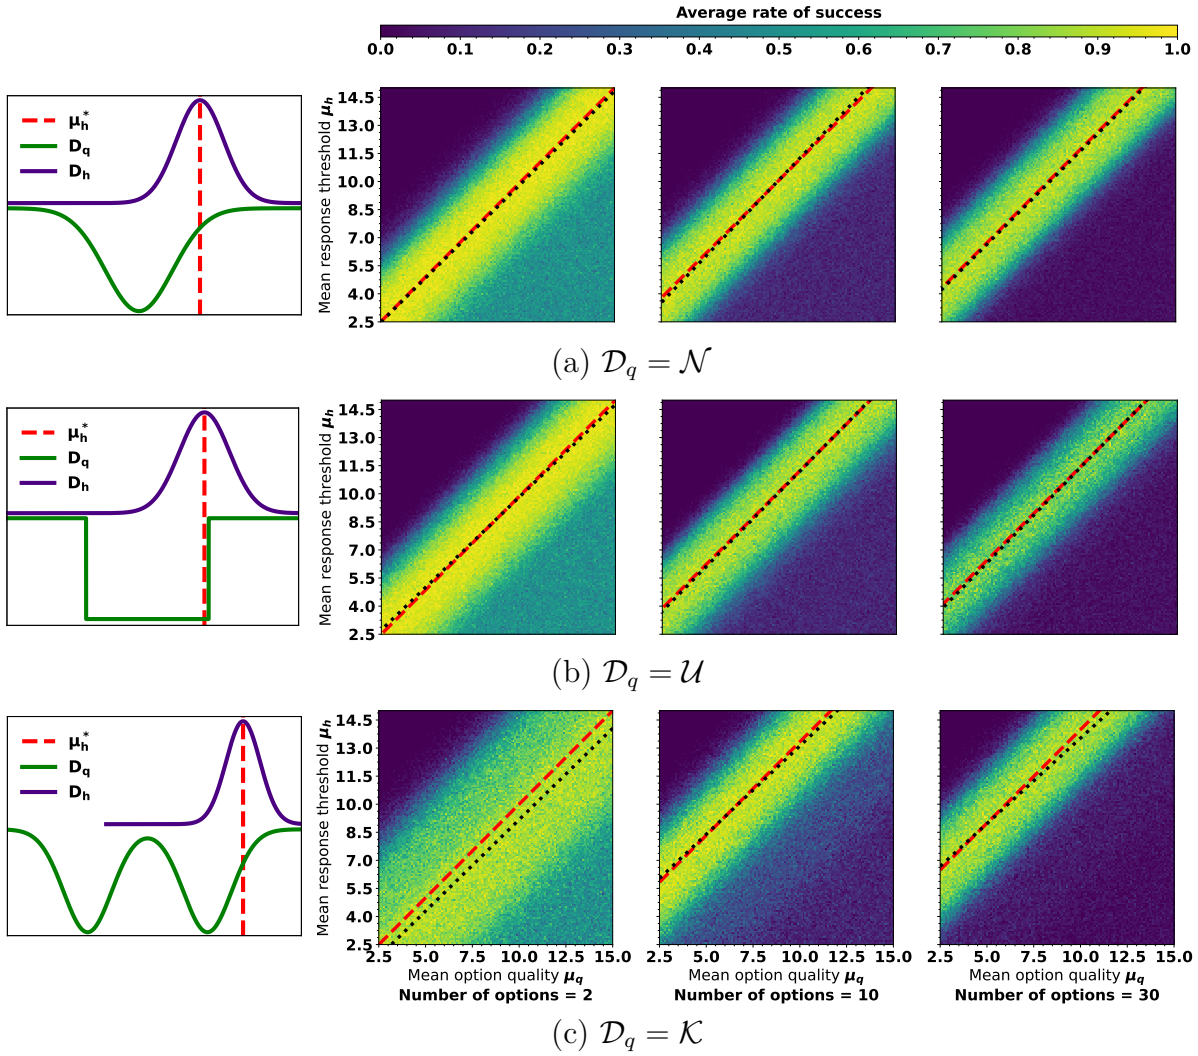


Figure 2.4: We tested three alternative distributions of options' qualities $\mathcal{D}_q$: (a) Gaussian $\mathcal{N}(\mu_q, \sigma_q)$, (b) uniform $\mathcal{U}(\mu_q, \sigma_q)$, and (c) bimodal distribution $\mathcal{K}(\mu_q, \sigma_q)$. We test the combination of μ_h vs μ_q in the value range $[2.5, 15]$ with a step size of 0.1, for distributions with equal fixed standard deviations $\sigma_h = \sigma_q = 1$ and a swarm size of $S = 100n$. The color maps show the average rate of success (500 runs per condition), the dotted black line is the fitted HARS line, and the dashed red line is the best threshold mean μ_q^* predicted with Eq. (2.1). The two lines match well for all tested types of quality distributions and the number of options.

heterogeneity may also vary with population size. While extremely low heterogeneity may trap the swarm in local maxima, large heterogeneity might suppress the majority on the best option due to noisy votes. We derived Eq. (2.1) for μ_h^* from first principles; however, due to the dependency of the optimal heterogeneity on multiple factors such as the options quality distribution, number of options, and the population size, statistically, the optimal amount of heterogeneity required for distinguishing an extreme value from non-extreme values is not evident. These challenges prevented us from finding the best σ_h^*.

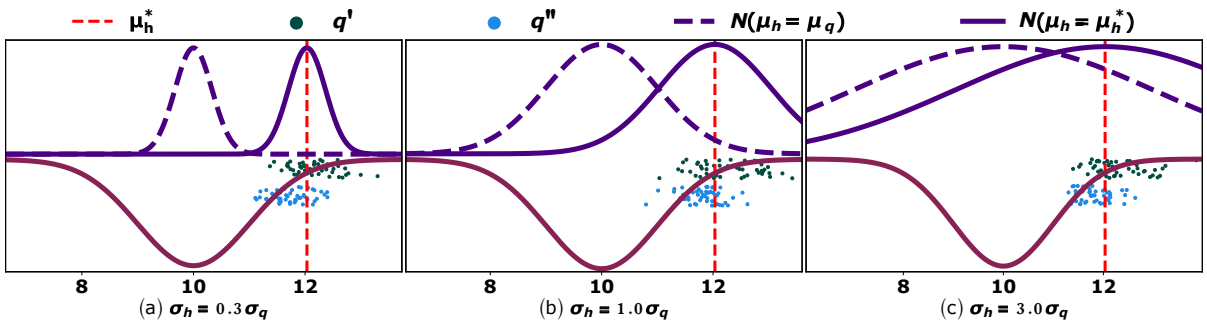


Figure 2.5: Comparison of the $\mathcal{D}_h = \mathcal{N}$ (top half) and $\mathcal{D}_q = \mathcal{N}$ (bottom half) for increasing ratio of σ_h to σ_q when $\mu_h = \mu_q$ (dashed distribution) and when $\mu_h = \mu_h^*$ (solid distribution). The green and blue dots show the best (q') and second-best (q'') value out of 50 sampled points in 50 trials. The red dashed line is μ_h^* computed using Eq. (2.1).

Instead, we numerically obtained the curves fitted along the HARS region on heatmaps of the standard deviation of the options quality distribution versus the standard deviation of the response threshold distribution, and reported the results in Figs. 2.9b-c and the mathematical equations representing the fitted curves in Eq. (2.2). Differently from μ_h^*, where the slope was approximately constant and intercept largely varied (see Fig. 2.7), for σ_h vs σ_q the intercept has negligible values close to zero (except for small $n \leq 3$), while the slope varies and is sufficient to determine the relationship between σ_h and σ_q as a function of n (see Figs. 2.9b-c). Similar to the analysis of μ_h vs μ_q, the HARS lines of σ_h vs σ_q do not show noteworthy changes with varying number of robots in both cases of $\mu_h = \mu_q$ and $\mu_h = \mu_h^*$ in σ_h vs σ_q, also when $\sigma_h = \sigma_q$ in μ_h vs μ_q (see respectively Fig. 2.8 and Fig. 2.6).

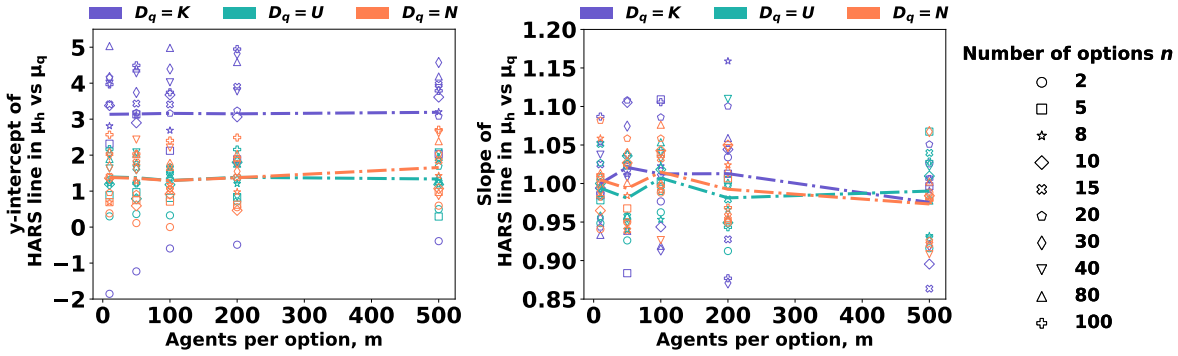


Figure 2.6: We show that the slope and the $y-$intercept of the HARS line in μ_h vs μ_q are invariant to changing the number of agents for all the cases of varying number of options in the various types of environment, when the standard deviation of the response thresholds' distribution and the options' quality distribution were fixed at $\sigma_h = \sigma_q = 1$.

We find mathematical equations, through standard polynomial fitting techniques

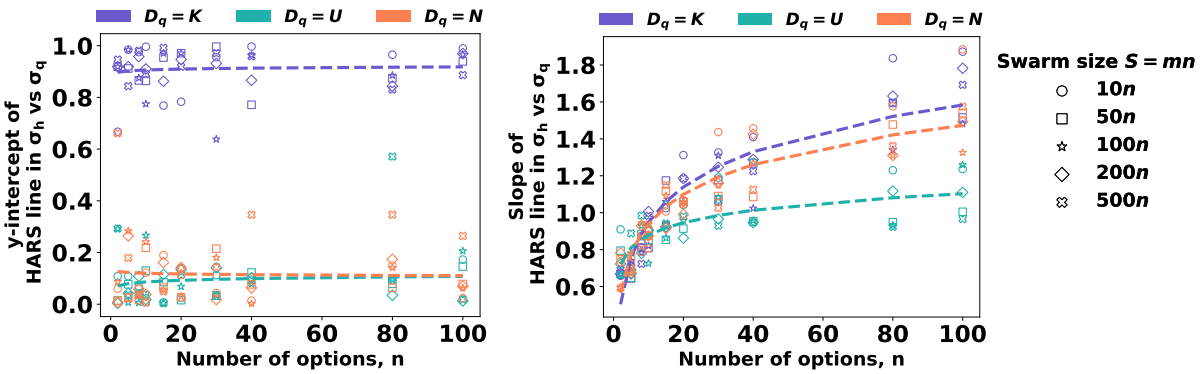


Figure 2.7: We show that the slope of the HARS line in σ_h vs σ_q varies significantly, whereas the $y-$intercept remains unaffected by changing the number of options in the different types of environments. We run these simulations by naively choosing the response threshold distribution with the same mean as that of the options' quality distribution.

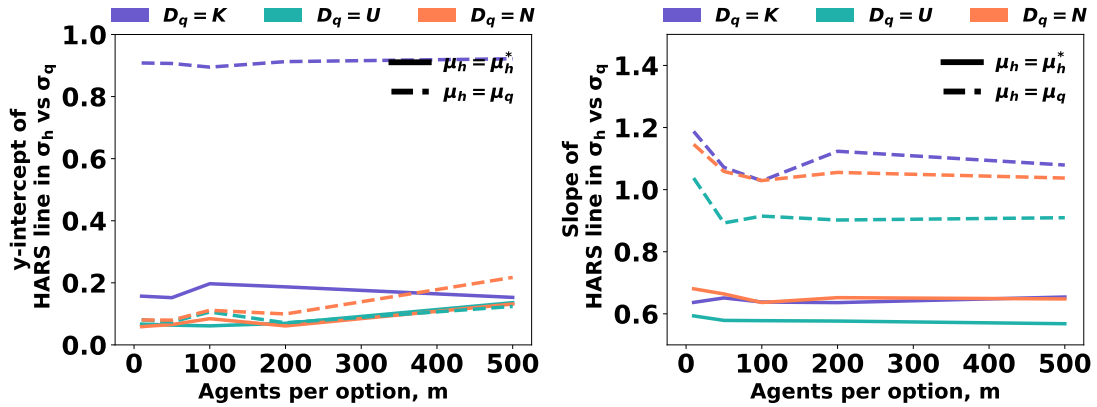


Figure 2.8: Dependency of the slope and the $y-$intercept of the HARS lines on σ_h vs σ_q plots. We compare the two conditions when $\mu_h = \mu_q$ (dashed lines) and $\mu_h = \mu_h^*$ (solid lines), to estimate how the slope and the $y-$intercept are affected by changing the number of agents assigned to each option. We found that both the slope and the $y-$intercept remain invariant to changing the number of agents, m. This remains valid for all three types of environments considered in this study.

(NumPy polyfit function [114]), to describe the change in the slopes, with a different equation for each of the considered quality distributions:

$$\sigma_h^* = \begin{cases} (-0.07 \log_{10} n + 0.67)\sigma_q & \text{if } \mathcal{D}_q = \mathcal{U}(\mu_q, \sigma_q) \\ (0.07 \log_{10} n + 0.57)\sigma_q & \text{if } \mathcal{D}_q = \mathcal{N}(\mu_q, \sigma_q) \\ (0.05 \log_{10} n + 0.58)\sigma_q & \text{if } \mathcal{D}_q = \mathcal{K}(\mu_q, \sigma_q)\ . \end{cases} \tag{2.2}$$

It is interesting to note that in environments where the options' qualities are bounded to a maximum value (e.g. $\mathcal{D}_q = \mathcal{U}$), the value of σ_h^* decreases with larger n, as best and second-best options have similar values in an increasingly narrow range. Instead, when

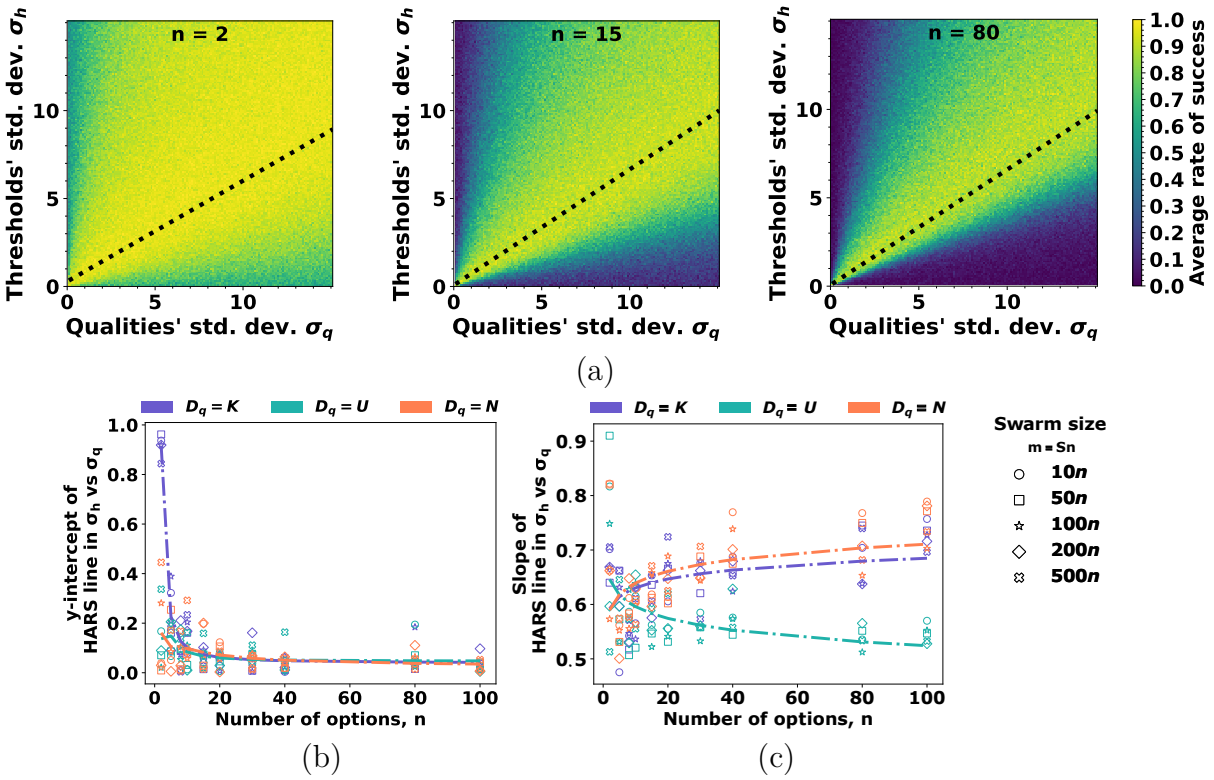


Figure 2.9: We test the relationship between the standard deviation of the distributions of the option's quality $\mathcal{D}_q = \mathcal{N}(\mu_q = 7.5, \sigma_q)$ and the response thresholds $\mathcal{D}_h = \mathcal{N}(\mu_h = \mu_h^*, \sigma_h)$, i.e., σ_q vs σ_h. The results are the average of 500 simulations per condition in a swarm of $S = 100n$ agents. (a) The color maps show the average rate of success for a Gaussian quality distribution. The black dotted HARS lines show that the highest success rate can be obtained with a threshold's standard deviation smaller than that of the qualities. Panels (b) and (c) show the $y-$intercept and slope, respectively, of the fitted HARS line for different swarm sizes and different types of quality distributions.

quality values are unbounded (e.g. $\mathcal{D}_q = \mathcal{N}, \mathcal{K}$), the probability of having high-quality outliers increases with n, and it seems that a higher standard deviation helps in making a correct decision in such cases.

2.6 Conclusion

Minimalistic robots, such as organic nanorobots with basic functionalities, have the potential to disrupt several fields, including medicine, agriculture, and environmental preservation [18, 19, 90]. However, designing solutions for robotic systems limited in memory, communication, and computation is challenging. In our research, we investigate the potential application of heterogeneous response thresholds in making best-of-n decisions. We envision the possibility of simplifying the existing algorithms used in collective decision making in the context of best-of-n [62, 66, 68, 70, 85, 92–94], by removing the (currently necessary) robots' ability of (i) scaling the estimated option's quality in a normalized quality range, (ii) memorize the option's quality in order to modulate the communication frequency (i.e. weighted voting), and (iii) exploring multiple options before committing to any. Our vision consists of building binary reactive robots that engage in an unweighted voting without keeping track of options' quality [108, 109]. Each robot uses a simple binary response threshold to accept/reject the sensed option [1, 96], and only accepting robots begin the voting process. In order to allow the swarm to converge towards the best of the n alternatives, the (unweighted) voting must begin from a state of relative majority in favor of the best option [74]. This work investigates which response threshold distribution to choose to increase the probability of having such a condition. We investigated the relationship between $\mathcal{D}_h$ (PDF) and various types of environments, characterized by different $\mathcal{D}_q$ (PDF) and the number of options. Our key finding suggests that the relative position of the mean of $\mathcal{D}_h$ should be at the left end of the last slice out of n slices of the $\mathcal{D}_q$. Whereas the trend in optimal heterogeneity still remains an open question. The collective accuracy obtained simply via placing the μ_h at μ_h^* improved almost by 20% for all the tested PDFs, compared against the naive setting of $\mathcal{D}_h = \mathcal{D}_q$, especially for large number of options (Fig. 2.10a), and small swarm sizes (Fig. 2.10b). The relative improvement decreases with increasing swarm size because, in very large systems, accuracy approaches 100% for both optimal and suboptimal response threshold distributions. We believe that this study is a necessary preliminary step towards the development of

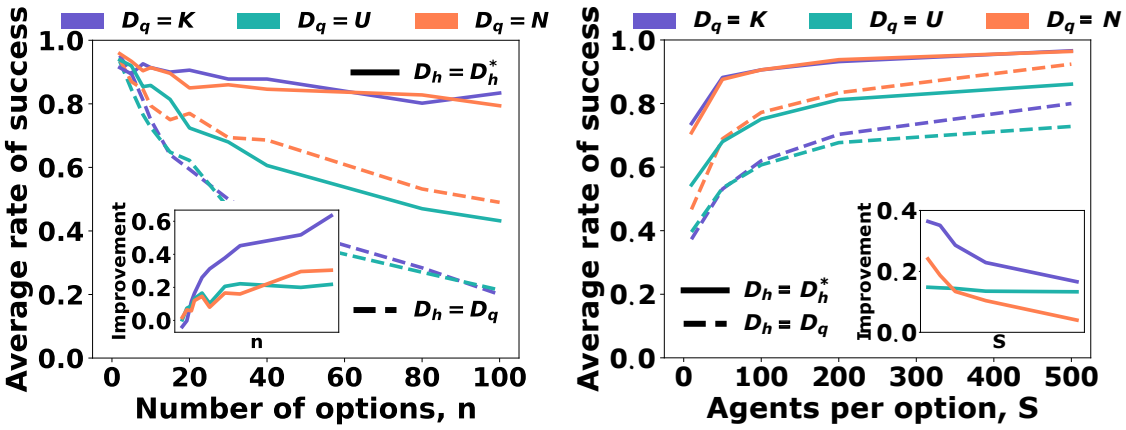


Figure 2.10: Accuracy for $\mu_q = 5$ and $\sigma_q = 2$, averaged (a) over different $m = Sn$ with $S \in \{10, 50, 100, 200, 500\}$, and (b) over different $n \in \{2, 5, 8, 10, 15, 20, 30, 40, 80, 100\}$. The solid lines depict the case when $\mathcal{D}_h = \mathcal{D}_h^*$ and the dashed lines depict the case when $\mathcal{D}_h = \mathcal{D}_q$. The two insets show the benefits of employing μ_h^* and σ_h^* as (a) the number of options n increases and (b) the swarm size m decreases. Note that the σ_h^* was obtained using Eq. (2.2).

minimalistic robot swarms based on adaptive response thresholds, capable of solving the best-of-n decision problem. In the next chapter, we propose a decentralized adaptive mechanism that enables the robots to collectively adapt an initial response threshold distribution to the desired best response threshold distribution.

CHAPTER 3

ADAPTIVE RESPONSE THRESHOLD MECHANISM FOR MULTI-ROBOTS TO SOLVE BEST-OF-*N* DECISION-MAKING

3.1 Overview

Advanced individual information processing is not always necessary to obtain a rational collective behavior. Collective cognition can also emerge from the minimal computations by aggregation of individuals, even when they have limited capabilities in sensing the world and communicating with others, as we observed in the previous chapter. In this chapter, we propose a decentralized adaptive mechanism to adjust response thresholds for a swarm of minimal-computation units that display a binary response (yes/no or accept/reject) to an external stimulus. By leveraging individual differences in the thresholds that trigger responses, the swarm can collectively compute the best alternative from a set of possible stimuli. Achieving a clear majority (the fraction of robots voting for the best option relative to the rest is significantly large) is essential to reaching consensus; however, existing solutions do not always guarantee a clear majority on the best option. The lack of a clear majority prolongs the voting phase, resulting in increased resource consumption. In this adaptive approach, a small set of prior responses during the iterative decision process forms the individual robots' belief about their rejection rate, which they use to adaptively modify the response threshold. When the number of options changes, the robots update their target rejection rate. We demonstrate that noisy environmental information is sufficient to enable individuals to adaptively adjust their response thresholds in a decentralized manner and maximize collective performance in dynamic environments while obtaining a clear majority in favor of the best option. We consider environments with varying numbers of stimuli and different distributions of stimulus strengths, which we illustrate as two factors under a fixed swarm size that critically influence the best-performing distribution of the individuals' response thresholds.

3.2 Introduction

Complex collective behavior does not necessarily require sophisticated individuals [98, 115, 116]. Groups of ants, bees, and bacteria use simple response rules to allocate tasks, sense quorums, and select food sources [59, 96, 117, 118]. Inspired by such systems, swarm robotics investigates how large numbers of simple robots can achieve tasks through local interactions, including response-threshold-based task allocation [102–104, 119]. Many application scenarios impose strict constraints on each robot's computation, sensing, and communication resources; examples include nanorobots in the bloodstream and disposable environmental sensors [120, 121]. In these settings, individual robots cannot reliably distinguish the best among a set of environmental stimuli on their own. However, diversity in how robots respond to stimuli can lead to group-level processing that exceeds any single agent's capability [122].

Response-threshold models have long described the emergence of division of labour in social insects and quorum sensing in bacteria [117, 118, 123]. These models assign each agent a threshold that determines whether it becomes active when exposed to a stimulus. Previous work has used fixed threshold distributions to enable swarms to perform best-of-n decisions; aligning the mean threshold with the mean option quality can improve discrimination between alternatives [1, 97, 99]. Yet robotic swarms operate in dynamic environments where the number and quality of stimuli may change over time [92]. In such environments, fixed distributions become suboptimal, and sensing noise or incomplete knowledge about the number of options further complicates the design. Biology suggests that animals modulate their sensitivity depending on context [30, 57, 124], hinting at the need for an adaptive mechanism [107].

In this paper, we propose a decentralized threshold adaptation rule that allows each robot to adjust its response threshold based on a short memory of recent interactions. By reflecting on how often it rejects stimuli, a robot can infer whether its threshold is too low (too permissive) or too high (too selective) relative to the current environment. Robots update their thresholds toward a target rejection rate that ensures clear separation between the number of robots that accept each option and those that reject it. The resulting distribution of thresholds evolves over time so that the mean threshold tracks an optimal value derived from the distribution of stimulus qualities. Although we illustrate

this mechanism in the context of best-of-n decision problems, it generalises to any task in which binary responses should magnify differences between alternatives.

Our contributions are:

1. We formalise a minimal model of a multi-agent system where each agent samples environmental stimuli and responds with a binary accept or reject decision based on an internal threshold.

2. We derive the optimal mean of the threshold distribution that maximises the difference between the subgroups of agents that accept the highest-quality option and those that accept any other, yielding a clear majority.

3. We design a lightweight, decentralised update rule that allows agents to adjust their thresholds toward this optimum using only local observations and a short-term memory.

4. Through extensive simulations we show that swarms using the threshold adaptation rule can adapt to changes in the number and distribution of stimuli, converge to near-optimal threshold distributions, and maintain high decision success rates under noise.

The remainder of the chapter is organized as follows. Section 3.3 introduces the model and derives the desired threshold distribution. Section 3.4 presents the threshold adaptation rule. Section 3.5 reports simulation results under various environmental conditions. Section 3.6 summarizes the chapter.

3.3 Model

We consider an environment containing n distinct stimuli (also referred to as options), $\mathcal{O} = \{o_i\}_{i=1}^{n}$. Each stimulus o_i has an associated quality q_i, drawn from a probability distribution $\mathcal{D}_q(\mu_q, \sigma_q)$. Without loss of generality, we order stimuli so that $q_1 \leq q_2 \leq \cdots \leq q_n$, with o_n being the best stimulus (i.e., q_n has the highest quality). A swarm of $m > n$ agents, denoted by $\mathcal{T} = \{a_j\}_{j=1}^{m}$, examines the environmental stimuli. At a decision instance s, agents partition into subgroups by sampling options; we denote by $g_i(s) = \{a_j \mid a_j \text{ samples option } o_i\}$ the set of agents that evaluate option o_i. Note that

we assume all the robots discovers one option and then stop exploring; however at each decision instance, robots collectively discover all the available options.

Each agent a_j has an internal response threshold h_j. After measuring the quality of its assigned option (with a measurement subject to sampling noise), agent a_j emits a binary response $d_{ij}(s) = 1$ if the estimated quality $q_{ij}(s)$ meets or exceeds its threshold $h_j(s)$, and $d_{ij}(s) = 0$ otherwise, that is,

$$d_{ij} = \begin{cases} 1 & \text{if } q_{ij} \geq h_j, \\ 0 & \text{otherwise.} \end{cases} \tag{3.1}$$

Within each subgroup, accumulated votes are given by $g_i^1 = \sum_{j \in g_i} d_{ij}$ and $g_i^0 = \sum_{j \in g_i} (1 - d_{ij})$. We define the swarm's majority choice as the option with the highest number of acceptances:

$$M(s) = o_k, \quad \text{where} \quad k = \arg\max_i g_i^1(s), \tag{3.2}$$

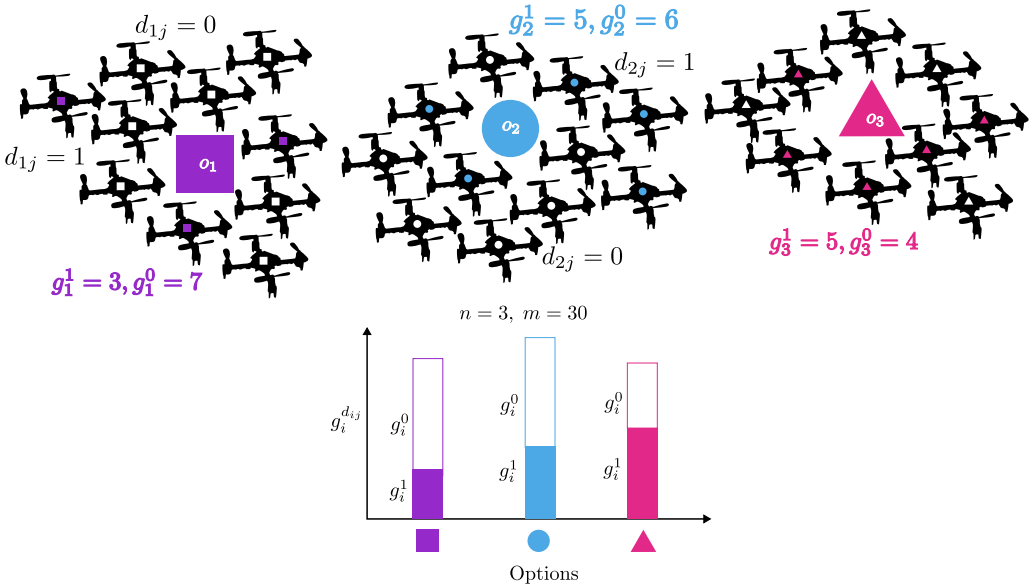


Figure 3.1: Schematic of best-of-n decision-making with response thresholds. There are three options, $\mathcal{O} = \{o_1, o_2, o_3\}$ ($n = 3$), and $m = 30$ agents distributed into subgroups of approximately equal size. Each individual either accepts (activates, colored shape) or rejects (deactivates, white shape) the option to which it is committed.

with ties broken at random. When the mechanism is used to select the best option, the team succeeds at sensing iteration s if it chooses o_n. Here, we considered a centralized external observation method to interpret the aggregated responses. However, in more general settings, the binary responses serve to amplify differences between alternatives,

and subsequent processes (e.g., decentralized voting, ranking, or centralized external observation) interpret the aggregated responses.

Figure 3.1 shows an example with three subgroups of agents sampling three options in the environment. Each agent obtains a noisy quality estimate $q_{ij} \sim \mathcal{D}_{q_i}(\mu_{q_i}, \sigma_{q_i})$ and compares it with its threshold to produce a binary response.

3.3.1 Optimal Threshold Placement

To achieve a clear separation between acceptances and rejections, we seek the mean threshold μ_h that maximizes the difference between the fraction of agents that accept the strongest option and those that accept any other option. For a fixed distribution of option qualities D_q, in Chapter. 2 we showed that the decision success probability is maximized when the mean threshold μ_h of a normal distribution $\mathcal{N}(\mu_h, \sigma_h)$ equals a critical value μ_h^*. This optimum can be derived from the cumulative distribution function of D_q [112]; for a normal distribution, it satisfies

$$\int_{-\infty}^{\mu_h^*} \mathcal{D}_q(x \mid \mu_q, \sigma_q)\, dx = 1 - \frac{1}{n}. \tag{3.3}$$

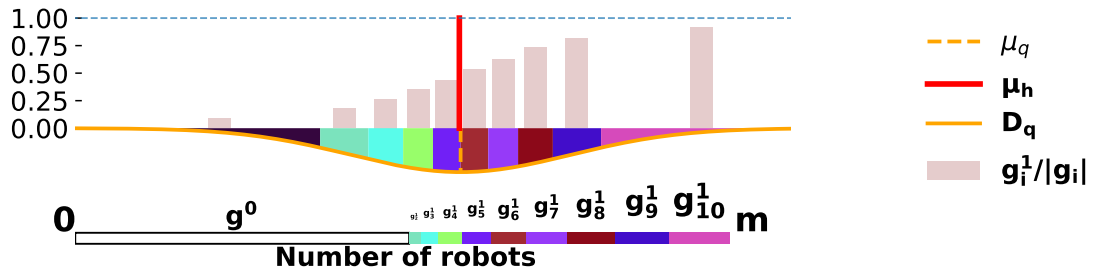
Intuitively, when the mean threshold equals μ_h^*, roughly the thresholds lie above the second maximum option quality, ensuring that most of the agents evaluating only the top stimulus respond positively. See Figure 3.2 for the description of this concept.

To validate the clarity of separation, we define a metric *clarity in majority decision*:

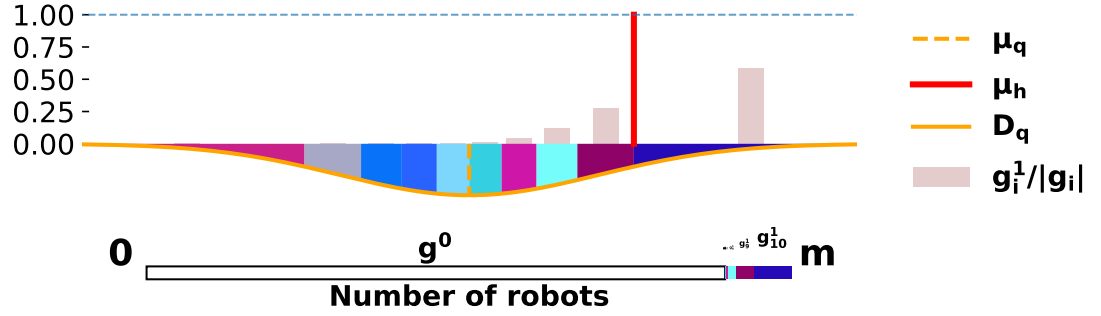
$$\mathrm{C_M} = \frac{g_n^1 - \sum_{i \neq n} g_i^1}{m}, \tag{3.4}$$

where higher values of $\mathrm{C_M}$ indicate a clearer majority. To illustrate the usefulness of the metric, we consider a fixed option quality distribution,$\mathcal{N}_q$, and robots with a randomly sampled fixed threshold. We perform Monte Carlo simulations under two scenarios: (a) $\mu_h = \mu_q$ [1], and (b) $\mu_h = \mu_h^*$ (see Eq. 3.3 and Fig. 3.2). The sliced regions in Fig. 3.2 represent the probability of at least one option belonging in each slice, and the horizontal bars show the mean number of acceptances g_i^1 versus rejections $g^0 = \sum_i g_i^0$ across 100 trials. Vertical bars indicate the acceptance fraction $g_i^1/|g_i|$.

In scenario Fig. 3.2(a), the fractions satisfy $g_9^1/|g_9| \sim g_{10}^1/|g_{10}|$, with $g^0 = 2527$ and



(a)



(b)

Figure 3.2: Simulations of $N = 100$ trials using $m = 5000$ agents, (a) when the $\mu_h = \mu_q$, and (b) when the $\mu_h = \mu_h^*$. The sliced region shows $\mathcal{D}_q = \mathcal{N}_q$ divided into $n = 10$ slices of equal area. The vertical bars represent the fraction $(g_i^1/|g_i|)$ of agents that accept the i^{th} option o_i. The horizontal bar in the bottom compares acceptance g_i^1 against rejection $g^0 = \sum_i g_i^0$. Clearly, in (b) $g_n^1 > g_i^1$ whereas in (a) $g_n^1 \approx g_{i\in\{n-1,n-2,n-3\}}^1$.

$\mathcal{C}_M = -0.3118$. Whereas, in scenario Fig. 3.2(b), we observe $g_9^1/|g_9| \ll g_{10}^1/|g_{10}|$, with a larger $g^0 = 4482$ and $\mathcal{C}_M = 0.0134$. Thus, simulations confirm that placing the mean threshold at μ_h^* yields clearer majorities (in Fig. 3.2(b), $g_{10}^1 > g_i^1 \, \forall i, i \neq 10$) than aligning it with the mean quality μ_q.

3.3.2 Clarity in majority

By definition, C_M is maximal when the rest of the options don't receive a positive response, and the best option receives a positive response from all committed individuals, that is, $C_M = 1/n$. Whereas C_M is minimal when the best option does not receive a positive response at all, but the rest of the options receive a positive response from all committed individuals, that is, $C_M = \frac{1}{n} - 1$. In essence, when $n = 2$, the range of C_M is $-0.5 \leq C_M \leq 0.5$, and as $n \to \infty$, $-1 \leq C_M \leq 0$.

Lemma 1. *Clarity in the majority is achieved when* $C_M > 0$.

Proof. A majority, by definition, means $g_n^1 > g_{i\neq n}^1$. However, in the voter model, the probability of consensus [125,126] (exit probability [66]) is defined by the initial fraction of the population responding in favor of an option:

$$\mathbb{P}(\text{consensus on } i^{th} \text{ option}) = \frac{\text{Number of response in favor of } i^{th} \text{ option}}{\text{Total number of responses}},$$

which in the best-of-n case is $\mathbb{P}(\text{consensus on } n^{th} \text{ option}) = g_n^1/m$. Therefore, to maximize the probability of consensus on the best option, the difference between the fraction of positive responses for the best option as compared to the fraction of the remaining positive responses for the remaining other options, that is, C_M, must be maximized or atleast $C_M > 0$. There are only two possibilities when $C_M = 0$:

- **Case I:** The best option and the second best option have exactly the same number of positive responses, and the rest of the options have negative responses.
- **Case II:** $g_n^1 = \sum_{i\neq n} g_i^1$.

When $g_n^1 > \sum_{i\neq n} g_i^1$, the probability of consensus on n^{th} option is maximized. Whereas, when $g_n^1 < \sum_{i\neq n} g_i^1$, the probability of consensus on n^{th} option is lower than consensus on other options. Therefore, the critical criterion for a clearer majority is $C_M > 0$. □

Theorem 1. $-1 \leq \mathbb{P}(\textit{consensus on } n^{th} \textit{ option}) - \sum_{i\neq n} \mathbb{P}(\textit{consensus on } i^{th} \neq n^{th} \textit{ option}) \leq 0.5$

Proof. As $2 \leq n \leq \infty$, $-1 \leq C_M \leq 0.5$ □

3.4 Adaptive Threshold Control

Inspired by social insects, each agent maintains a short-term memory of its recent responses [127,128]. Agent a_j stores the outcomes of its last v decisions in a log $L_j = \{\ell_{jb}\}$, where ℓ_{jb} records whether the agent a_j accepted or rejected the stimulus in the b-th most recent interaction. After every v evaluations, the agent resets the log and computes its rejection rate

$$r_j(s) = \frac{1}{v}\sum_{b=1}^{v}\Big(1 - \ell_{jb}\Big). \tag{3.5}$$

This quantity measures how selective the agent has been: a small value of r_j indicates a low threshold (the agent accepts most stimuli), whereas a large value indicates a high threshold.

3.4.1 Desired Rejection Rate

For the swarm to achieve the target mean threshold μ_h^* derived above, each agent aims for a desired rejection rate

$$r_j^*(s) = 1 - \frac{1}{n_j(s)}, \tag{3.6}$$

where $n_j(s)$ is the agent's current estimate of the number of stimuli. In practice, agents can share coarse information on how many distinct options they detect in their local neighborhood, and uncertainty in this estimate is addressed in Section 3.5.4.

3.4.2 Threshold Update Rule

After computing its current rejection rate $r_j(s)$ and desired rejection rate $r_j^*(s)$, agent a_j adjusts its threshold by a simple rule:

$$h_j(s+v) = h_j(s) + \alpha\Big[r_j^*(s) - r_j(s)\Big], \tag{3.7}$$

where $\alpha > 0$ is the adaptation rate (Sec. 3.4.3). If the agent has rejected too often relative to the desired rate, it lowers its threshold; if it has been too permissive, it raises the threshold. Because each agent updates its threshold only after filling its log of length v, this mechanism preserves decentralization and requires minimal computation and storage. Figure 3.3 shows a schematic overview of the update process.

3.4.3 Discussion of adaptation rate

The adaptation rate α regulates how strongly agents react to deviations in their rejection rate. A small value leads to gradual adjustment and low variance in the threshold distribution but slows convergence toward μ_h^*. A large value accelerates convergence but increases heterogeneity and oscillations around the target. Simulations reveal an intermediate range of α that balances responsiveness and stability, maximizing decision success (Sec. 3.5). The log length v influences how quickly the agent updates its threshold:

shorter logs result in rapid but noisy adjustments, whereas longer logs yield smoother but slower responses.

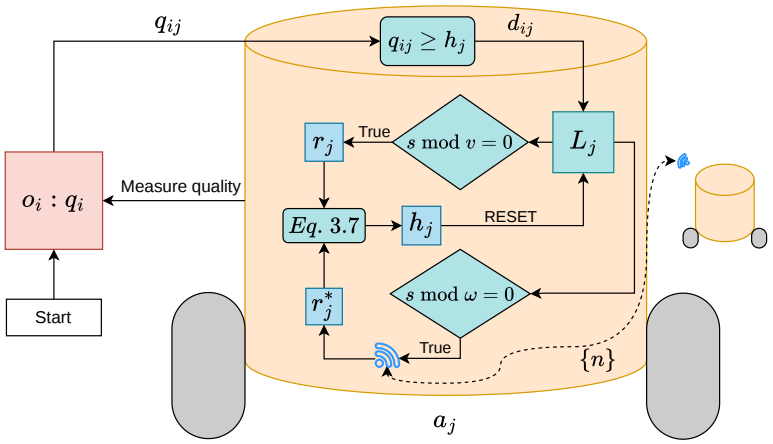


Figure 3.3: An agent a_j measures the quality q_{ij} of an option o_i and compares it with its threshold h_j to produce a binary response, d_{ij} through Eq. (3.1). Responses are recorded in a short-term log L_j to update the threshold through Eq. (3.7).

3.5 Simulation Experiments

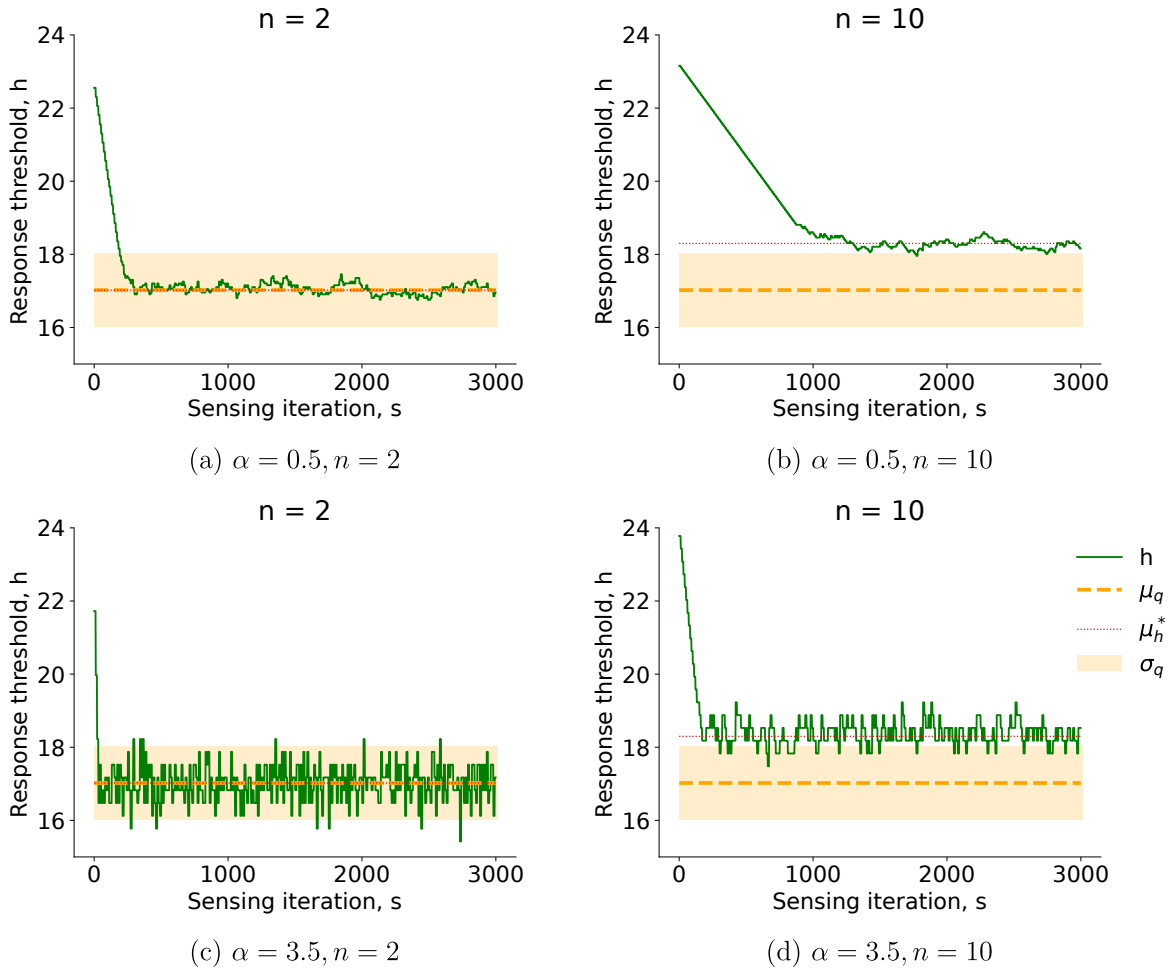
We evaluate the adaptive threshold mechanism through Monte-Carlo simulations in static and dynamic environments. Each experiment consists of N sensing iterations, and the decision success rate ρ is measured as the fraction of iterations where the subgroup with the largest amount of accepting agents is the subgroup g_n estimating the highest-quality stimulus o_n. Although we use best-of-n success as a quantitative metric, the same analyses apply to tasks where clear majority signals (e.g., C_M) rather than correctness per se are desired.

3.5.1 Convergence of Thresholds

First we analyse whether the mean threshold

$$\mu_h(s) = \frac{1}{m} \sum_{j=1}^{m} h_j(s) \tag{3.8}$$

converges toward the theoretical optimum μ_h^*. Figure 3.4 shows that a single agent evaluating repeated samples of the best stimulus adapts its threshold to oscillate around μ_h^*. Increasing the adaptation rate α speeds convergence but amplifies deviations from μ_h^*.



(a) $\alpha = 0.5, n = 2$

(b) $\alpha = 0.5, n = 10$

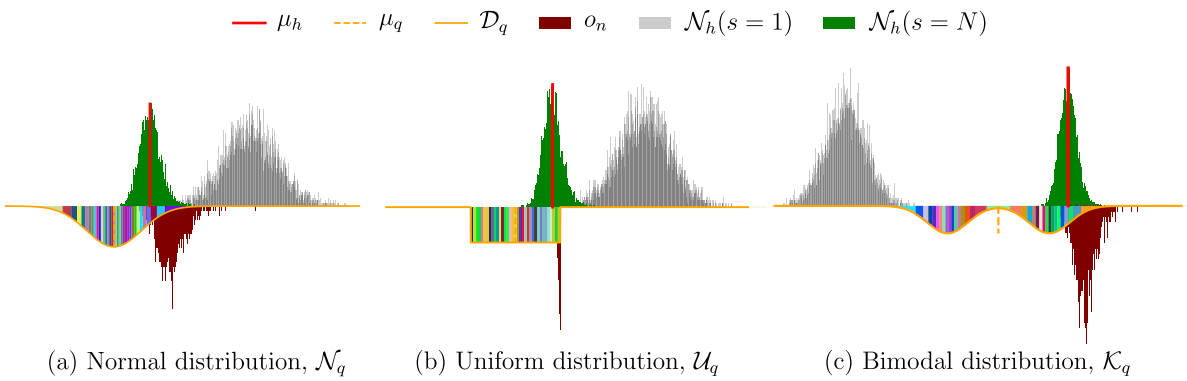
(c) $\alpha = 3.5, n = 2$

(d) $\alpha = 3.5, n = 10$

Figure 3.4: Evolution of response threshold h for (a) $\alpha = 0.5$ and $n = 2$, (b) $\alpha = 0.5$ and $n = 10$, (c) $\alpha = 3.5$ and $n = 2$, and (d) $\alpha = 3.5$ and $n = 10$. The options quality are resampled from $\mathcal{N}_q(\mu_q = 17, \sigma_q = 1)$ at each iteration (μ_q is depicted with an orange dashed line and σ_q as shaded area). The red dotted line represents μ_h^*. The logger size was $v = 10$.

In a large swarm ($m = 5000$) with memory length $v = 10$, the threshold distribution evolves from an initially random Gaussian distribution to one centered near μ_h^*, regardless of whether the stimuli are drawn from a normal, uniform, or bimodal distribution (see Fig. 3.5).

Consequently, the swarm's decision success rate ρ aligns with the optimal distribution over time and effectively amplifies stimulus differences, as shown in Fig. 3.6. The amplification of stimulus differences using the adaptive mechanism, in fact, surpasses that obtained with the best-response threshold distribution $\mathcal{D}_h^*$. We believe that this outperformance in amplifying stimulus differences could be due to our initial assumption that $\mathcal{D}_h^*$ is a Gaussian distribution. When used to select the best option, the adaptive



(a) Normal distribution, $\mathcal{N}_q$ (b) Uniform distribution, $\mathcal{U}_q$ (c) Bimodal distribution, $\mathcal{K}_q$

Figure 3.5: Evolution of the response-threshold distribution under three option-quality distributions. Grey histograms show the initial distribution of thresholds, green histograms show the adapted distribution after $N = 1000$ decision instances, and the red vertical line marks the mean threshold μ_h that maximises the clarity of the majority decision for $n = 35$.

mechanism maintains high decision success across a wide range of n, outperforming the designed fixed-threshold distributions like $\mathcal{D}_h^*$.

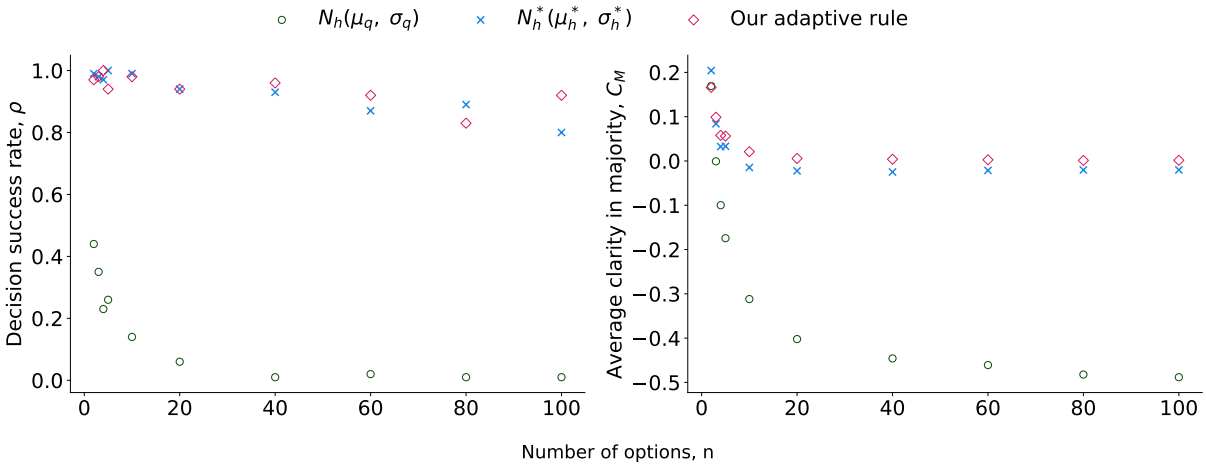


Figure 3.6: Comparison of the decision success rate ρ between the proposed adaptive mechanism and fixed response thresholds for different numbers of options n. We consider fixed response thresholds whose distribution matches the environmental distribution [1], i.e., $\mathcal{N}_h(\mu_h, \sigma_h) = \mathcal{N}_q(\mu_q, \sigma_q)$, and fixed thresholds whose distribution matches the optimal distribution(Chapter. 2), i.e., $\mathcal{N}_h(\mu_h, \sigma_h) = \mathcal{N}_h^*(\mu_q^*, \sigma_q^*)$. Results are computed from 100 simulation runs per n and $m = 5000$.

Furthermore, we consider a world with $\mathcal{D}_q(\mu_q = 17, \sigma_q = 1)$ and $n = 20$, and initialize a swarm of robots with $\mathcal{D}_h(\mu_h = 17, \sigma_h = 1)$. With changing sensing iterations, we observe in Fig. 3.7 that when $\mu_h = \mu_q$, the clarity in majority was low, and as μ_h converges to the μ_h^*, the clarity in majority approaches zero. In contrast to the initially high positive responses for the rest of the options, the count of positive responses for the best option

dominates those for the second-best and the rest of the options. The decision success rate also increases with increasing clarity in the majority.

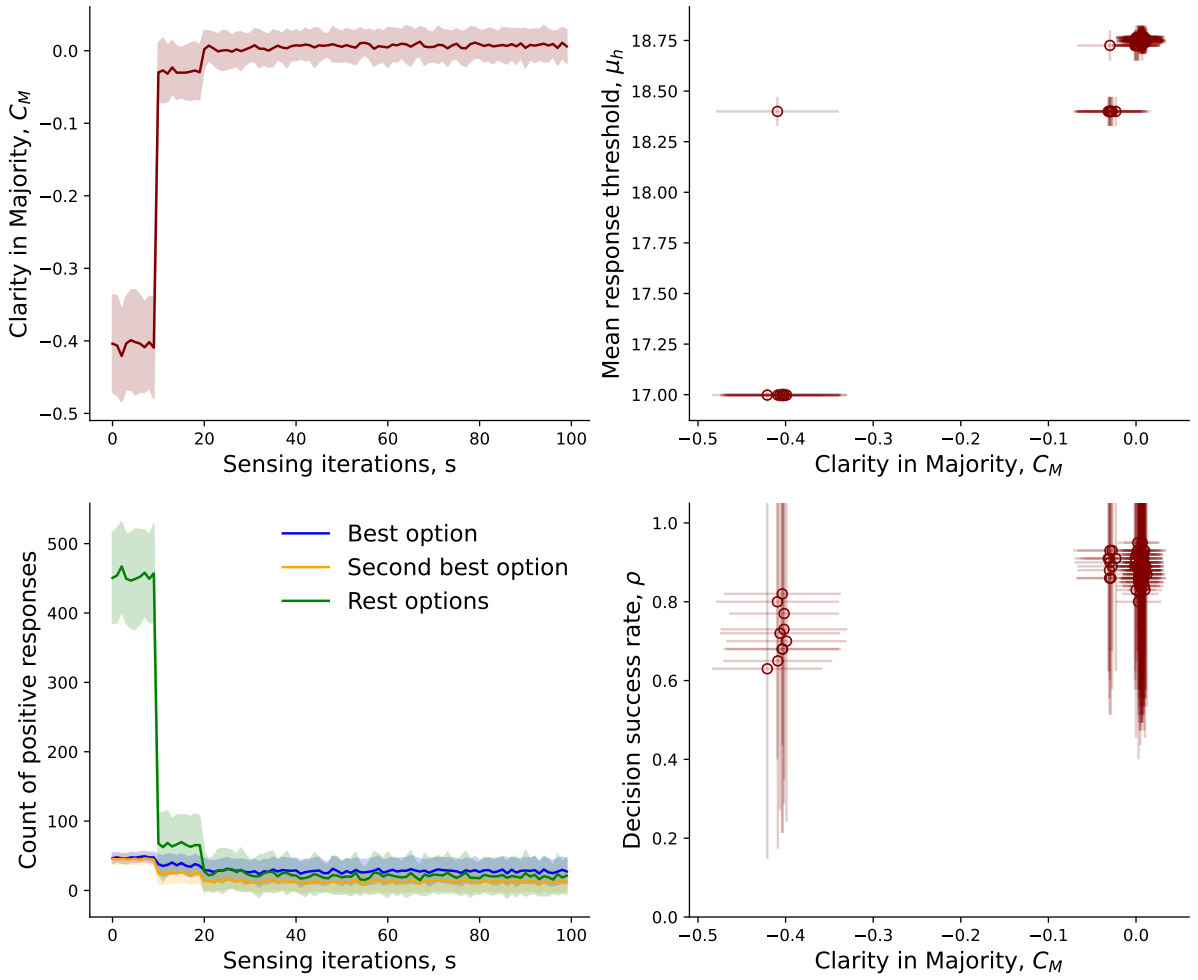
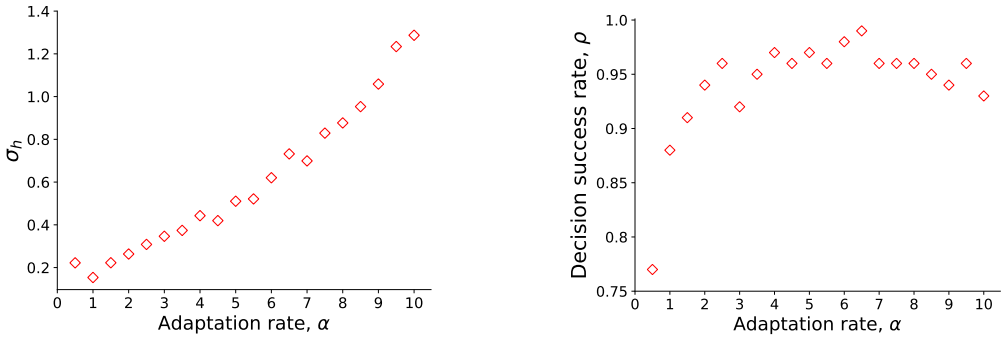


Figure 3.7: Evolution of the clarity in majority of $m = 1000$ robots initialized with $\mathcal{D}_h(\mu_h = 17, \sigma_h = 1)$ that adapt for $N = 100$ sensing iterations in a world with $n = 20$ options whose quality are distributed as $\mathcal{D}_q(\mu_q = 17, \sigma_q = 1)$. The extreme values of clarity in the majority for $n = 20$ are $-0.95 \leq C_M \leq 0.05$. The simulation was repeated for 100 times. The shaded region and horizontal/vertical lines depict standard deviations across 100 repetition.

3.5.2 Effect of Adaptation Rate

We next investigate how the adaptation rate α influences the steady-state variance of the thresholds and the decision success rate (Fig. 3.8). Simulations show that increasing α leads to a wider distribution of thresholds in the steady state, reflecting greater heterogeneity. However, decision success rate ρ does not monotonically increase with α. When α is too small, thresholds adjust too slowly, preventing the swarm from aligning with

the desired rejection rate. When α is too large, thresholds change so quickly that agents overreact to local fluctuations, causing instability. The optimal performance emerges for intermediate values of α where the swarm maintains sufficient diversity for exploration while still coordinating toward the correct majority. These trends highlight the trade-off between responsiveness and stability in the proposed adaptive mechanism.



(a) **(b)**

Figure 3.8: Effect of adaptation rate α on (a) the threshold distribution and (b) the decision success rate. Smaller α values lead to slow adaptation and narrow distributions, while large values induce heterogeneity and instability. Decision success ρ peaks at intermediate α values where the swarm balances exploration and exploitation for $n = 35$, $\mathcal{N}_q$, $v = 10$, $m = 5000$.

3.5.3 Dynamic Environments

We test the mechanism in dynamic environments where both the number of options and the quality distribution change periodically every $\omega = 1000$ decision instances (Fig. 3.10a-b). Within each block of ω decisions, the number of options $n(s)$, the option quality distribution $\mathcal{D}_q$, and the mean quality $\mu_q(s)$ remain constant, and only options' qualities are resampled at every instance. At the boundary of a block ($s \bmod \omega = 0$), they update according to

$$n(s) = \begin{cases} n(s-1), & s \bmod \omega \neq 0, \\ \max\Big(2,\, n(s-\omega) + \lceil \mathcal{U}(-5,5) \rceil\Big), & s \bmod \omega = 0, \end{cases} \tag{3.9}$$

and

$$\mu_q(s) = \begin{cases} \mu_q(s-1), & s \bmod \omega \neq 0, \\ \mu_q(s-\omega) + \mathcal{U}(-5,5), & s \bmod \omega = 0, \end{cases} \tag{3.10}$$

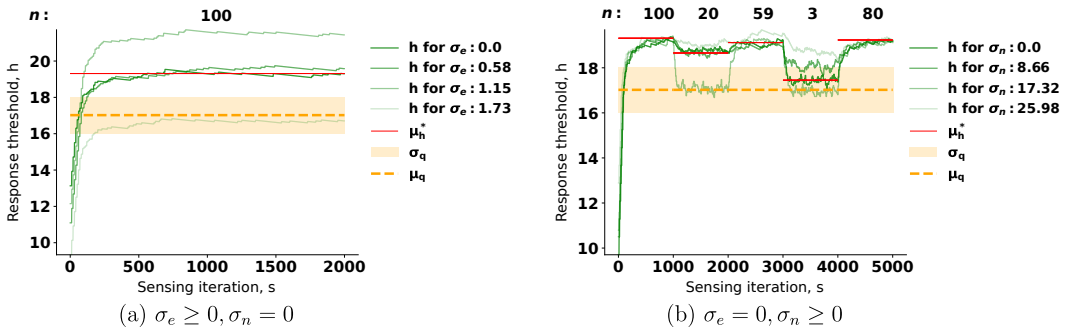
where $\mathcal{U}(-5, 5)$ is the uniform distribution between -5 and 5. At each update step ($s \bmod \omega = 0$), the type of the quality distribution $\mathcal{D}_q$ is also randomly selected from the set $\{\mathcal{N}_q, \mathcal{K}_q, \mathcal{U}_q\}$; the standard deviation is fixed at $\sigma_q = 1$. Figures 3.10a-b show that even in these highly dynamic settings, the adaptive mechanism continuously tracks the shifting $\mu_h^*(s)$ with acceptable lag.

3.5.4 Robustness to Noise and Environmental Change

Real agents operate with noisy observations and uncertain knowledge of how many options exist. We examine the effects of quality estimation noise (standard deviation σ_e), uncertainty in the estimated number of options (σ_n), and different log sizes v on both single-agent and swarm behavior.

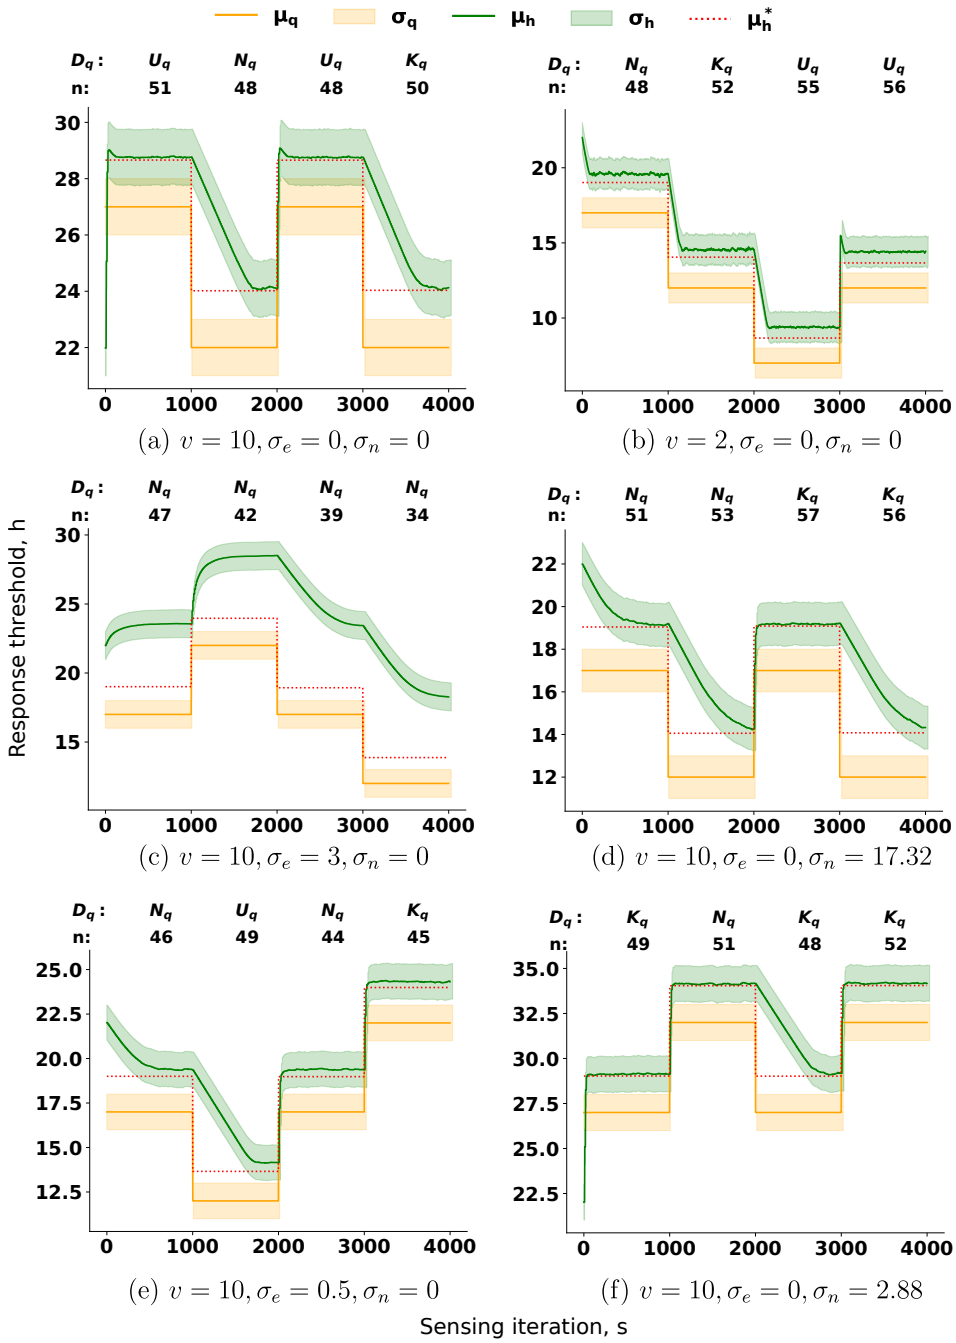
Single-Agent Analysis

When a single agent repeatedly evaluates the options in a dynamic environment with changing numbers of options n and quality distribution $\mathcal{D}_q$, the presence of quality estimation noise increases both the steady-state error $|\mu_h^* - h_j|$ and the time to reach this value (Fig. 3.9). Uncertainty in the number of options also increases the steady-state error but has little impact on convergence time. These observations highlight that noisy measurements bias the estimated rejection rate and thus the threshold update.



(a) $\sigma_e \geq 0, \sigma_n = 0$ (b) $\sigma_e = 0, \sigma_n \geq 0$

Figure 3.9: Evolution of the response threshold (green line) of an agent with $\log v = 10$ in an environment with dynamically changing n when (a) $\sigma_e \neq 0$, and when (b) $\sigma_n \neq 0$. The orange dashed line is the μ_q overlaid on the shaded region of σ_q. The red lines depict μ_h^*.



(a) $v = 10, \sigma_e = 0, \sigma_n = 0$ (b) $v = 2, \sigma_e = 0, \sigma_n = 0$

(c) $v = 10, \sigma_e = 3, \sigma_n = 0$ (d) $v = 10, \sigma_e = 0, \sigma_n = 17.32$

(e) $v = 10, \sigma_e = 0.5, \sigma_n = 0$ (f) $v = 10, \sigma_e = 0, \sigma_n = 2.88$

Figure 3.10: Evolution of the response threshold distribution $\mathcal{D}_h$ (green line and shaded region depict average and standard deviation) for a swarm of $m = 5000$, note n changes and thus, $|g_i|$ agents in a dynamic environment. The orange steps and shaded region depict the stimulus quality distribution $\mathcal{D}_q$. The red dotted line depicts μ_h^*.

Swarm-Level Behaviour

Figure 3.10 shows that, at the swarm level, three factors shape the adaptive dynamics: the memory length v, measurement noise σ_e, and uncertainty in option count σ_n. Short memories ($v = 2$) allow rapid adjustments but result in less accurate convergence;

longer memories ($v = 10$) delay updates yet yield more accurate threshold distributions. Increasing the quality-estimation noise reduces the accuracy of the converged threshold and slows adaptation, consistent with single-agent results. Interestingly, large uncertainty in the estimated number of options has negligible impact on the steady-state mean threshold when many agents interact. Only under extremely high noise levels or in very small swarms (or when a large number of options cause the subgroup size to shrink due to a fixed swarm size) does performance degrade noticeably as shown in Fig. 3.9 and 3.10.

3.6 Conclusion

We have presented a decentralized adaptive response-threshold mechanism that enables swarms of minimalist agents to adjust their sensitivity to environmental stimuli. In this approach, agents maintain a short memory of their binary responses and update their thresholds to match a target rejection rate derived from the estimated number of options. This simple rule leads the swarm to self-organize a threshold distribution whose mean tracks the optimal value for amplifying differences among options, producing clear majority signals without global knowledge of stimulus qualities. Our simulations show that the mechanism adapts to changes in the number and distribution of stimuli, tolerates measurement noise and uncertainty, and supports tasks such as best-of-n decision-making and more general difference-amplification computations. In dynamic environments and general settings, where external evaluation and placement of relative response thresholds cannot be performed with respect to the unknown quality estimates, the swarm's performance can be monitored by directly measuring clarity in the majority using individual responses. Our study indicates that minimal computation at the individual level is sufficient to achieve robust collective behavior, providing a foundation for future explorations of adaptive sensitivity modulation in both artificial swarms and biological collectives.

Supplementary Code

All our simulation code is available at https://github.com/swadhinagrawal/h-controller.git.

CHAPTER 4

IMPACT OF INDIVIDUAL DEFECTION ON COLLECTIVE MOTION

4.1 Overview

In the context of collective navigation, collective motion modeling has attracted significant attention for gaining insights into the mechanisms of collective behavior and its potential to inspire control strategies for swarm robotics. Most of the existing models assume that individuals within a group strictly adhere to the interaction rules. However, individuals in both artificial and natural collectives may occasionally fail to follow the interaction rules, which differs from noisy actions. In this chapter, we analyze how the presence of individuals who occasionally defect affects the ordered phase of the group during collective motion. Using Monte Carlo simulations, we study two collective motion models: a non-spatial (pairwise interaction) and a spatial (Couzin) model. In the non-spatial model, when individuals defect with a higher probability, both the time required by the agents to reach directional consensus (polarized group motion) and the average energy cost of the group to maintain such directional consensus (average rotational energy consumption per individual in highly polarized groups) increase. In the spatial model, there are conditions under which the presence of defecting agents can simultaneously reduce the time required for the collective to reach high polarization and the average energy cost in the polarized state. These findings not only enhance our understanding of probabilistic defective behavior in biological systems but can also inspire innovative, efficient, and controllable approaches in swarm robotics.

4.2 Introduction

Collective motion [60] is a universal phenomenon observed in nature, ranging from macromolecules [129], to the simplest multicellular organisms [130], insects [131,132], fish [133], birds [134], herds [45], and even humans [135]. Over the past decades, mathematical models [2,42,77,78] have provided insights into how individuals, by interacting with their

local neighbors, can collectively move in a coordinated manner. Most of these models assume that every individual adheres to the rules defined in these models with a certain degree of variability, which is normally included as random noise. While individuals in natural systems, such as fish shoals or bird flocks, typically adhere to interaction rules [134, 136], they may also occasionally deviate from this behavior in ways that differ from noisy actions. For instance, individuals may choose not to interact and instead take selfish actions to exploit their peers' collaboration efforts without paying their fair share of social coordination costs [33, 137]. Another reason why certain individuals may occasionally not follow the rules of collective motion is the fact that individuals have complex behaviors where collective motion is only one of its components. It is reasonable to assume that animals may concurrently forage, look for mates, or avoid dangers and predators; therefore, other individual goals than motion coordination could prevail in determining animals' actions [43, 138–141]. Competition can also play a role in causing a temporary interruption in following collective motion rules [31, 142]. We refer to the behavior where individuals deviate from the collective motion rules as *defection*.

There are previous works that modeled intermittent, or occasional, defective behavior in a variety of ways. Such studies considered variable frequencies of interaction, *i.e.*, agents always follow the cooperative rules; however, each of them has a personal frequency of interaction [143, 144]. In such an asynchronous interaction system, certain agents update their behaviors (or states) more frequently than others. Slower agents (with lower interaction frequency) may appear as defectors to quicker agents that make (unilateral) updates of their behavior without a direct response. Other studies considered individuals that selectively interact only with a subset of neighbors [145]. This research has provided interesting insights into the evolution of cooperation in collective systems [146, 147]. Investigating the impact of occasional defection on collective motion remains an open question [148, 149]. In this work, we investigate whether having individuals with a given probability of defection (where defection refers to ignoring the collective motion rule) can be beneficial for the group, and if so, what level of defection maximizes the collective benefits. Unlike previous studies that employed frequency-based or selective-interaction approaches to model defection at approximately regular time intervals [78, 150], our models introduce a probability for every agent to defect, thereby breaking the periodic patterns of defection events across the number of time steps.

Considering the impact of defecting agents in decentralized groups can also be relevant for the design of robotic systems, such as robot swarms [4]. While robots would be programmed not to defect and always follow the collaborative rules, their defective behavior could be the consequence of malfunctioning or noise. For instance, due to communication failures [151,152], noisy sensory inputs, and partial observability of the environment, robots may not always follow the prescribed algorithm [153]. Adversarial tampering is another example where a robot, temporarily under the control of a hacker, may behave erratically [28, 154]. It is thus important to design algorithms that are robust to occasional defection by part of the robots, or that would even exploit defective behavior to improve certain aspects of the group dynamics. This work takes an initial step towards this research direction by analyzing the conditions under which probabilistic defection can improve or harm group dynamics.

To do so, we analyze two prominent collective motion models, the pairwise model [78] (a non-spatial model) and the Couzin model [2] (a spatially-explicit model). We extend these two models to include the possibility that agents occasionally defect, investigating how this impacts group dynamics. At each time step, an agent defects with probability P_D and otherwise cooperates by following the interaction rules. Figure 4.1(a) presents a schematic illustration of individuals defecting in the alignment interaction rule with probability $P_D = 1/3$. That is, an agent cooperates by aligning its orientation with its neighbors in two-thirds of the time steps and, instead, in the remaining one-third of time steps, it defects and maintains its original orientation, ignoring its neighbors. In this study, the agent's decision to cooperate or defect is independent of its neighbors' strategy and solely based on P_D, as illustrated in Fig. 4.1(d).

We measure the impact of defecting agents in terms of two metrics, polarization time, T, and average energy cost, E, (defined in Sec.4.3). Polarization time is the time the group needs to agree on the direction of motion, i.e., to have most of the agents aligned in the same direction. The average energy cost is the amount of energy the agents use to maintain a high level of polarization (i.e., how much rotational energy each agent, on average, consumes to remain aligned in the same direction). For high levels of defection, the group may be unable to achieve a sufficient level of motion alignment, or it may become polarized only briefly before reverting to a disordered state (an unpolarized group).

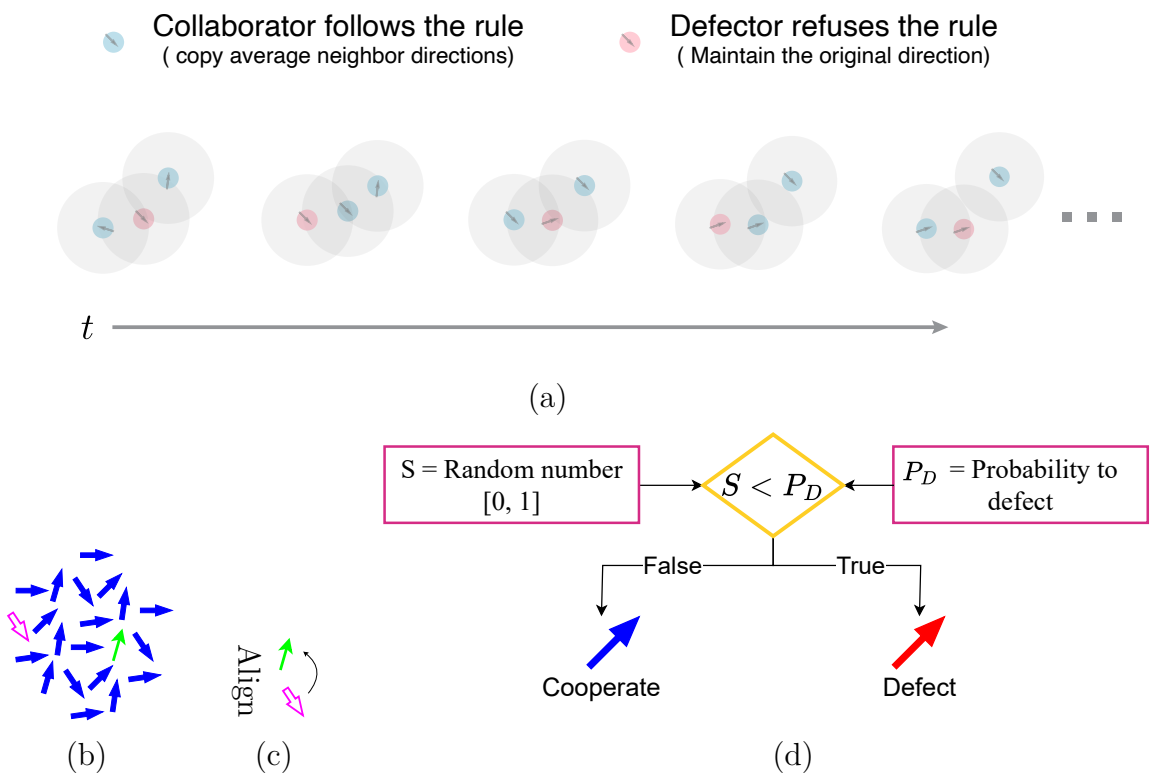


Figure 4.1: (a) Illustration of the collective behavior of individuals that probabilistically cooperate by following the collective motion rules or defect by maintaining their original directions. The pairwise model: (b) A pair of interacting agents in a group of fully connected agents, (c) Focal agents' (hollow pink) interaction rule, (d) Mechanism to update the agent's strategy (cooperate/defect) via the defection probability P_D.

In Sec. 4.4, we analyze the non-spatial pairwise model and observe that increasing defection probability P_D causes an increase in both polarization time and average energy cost. In Sec. 4.5, we analyze the spatially-explicit Couzin model [2]. While the pairwise model only includes an alignment rule, the Couzin model is more complex as it comprises three rules. We find that when individuals defect in certain collective motion rules, as the defection probability increases (up to a certain point), both the polarization time and average energy cost decrease. Note that, in this study, we do not compare the dynamics of the pairwise model with the Couzin model, as they are too different for a fair comparison, which is outside the scope of this study.

4.3 Methods

In the simulation, the agents move at a constant speed $\|v\|$. The kinematic motion of the agents is implemented using Dubin's car model [155] described by

$$\begin{aligned} \dot{x} &= \|v\| \cos(\psi) \,, \\ \dot{y} &= \|v\| \sin(\psi) \,, \\ \dot{\psi} &= \omega \,. \end{aligned} \tag{4.1}$$

Here, $\dot{x}$ and $\dot{y}$ are the velocities of the agent along each coordinate axis in the $2D$ plane, and ψ is the yaw (heading) angle of the agent. The turning rate ω, is constrained by a maximum of $\omega_{\max}$, *i.e.*, $|\omega| < \omega_{\max}$. The desired turning rate for an individual is obtained from the interaction model (Secs. 4.4 and 4.5).

Given N agents in a group, we define the group polarization $p(t)$ and the group energy cost $e(t)$ at a given time t. Group polarization $p(t)$ is the average of the unit velocity vectors $\hat{v}_i$ of individual agents i at time t, given by

$$p(t) = \frac{|\sum_{i=1}^{N} \hat{v}_i(t)|}{N} \,, \tag{4.2}$$

and group energy cost $e(t)$ at time t is measured as the sum of the rotational energy of each individual, given by

$$e(t) = -\sum_{i=1}^{N} (\omega_i(t))^2 \,. \tag{4.3}$$

We perform Monte Carlo simulations [156, 157], 100 independent runs for each condition (set of parameters). Each run lasts $t_{\max} = 200s$, with a time step $dt = 0.1$ sec. We measure the polarization time T required to achieve a group polarization $p(T)$ beyond a threshold value $p^* = 0.8$, *i.e.* $p(T) > p^*$, and the average energy cost E per individual after the group gets polarized:

$$E = \frac{1}{N(t_{\max} - T)} \sum_{t=T}^{t_{\max}} e(t). \tag{4.4}$$

The agents move in an unbounded region, and at the start of each run, they are

distributed uniformly at random inside a circular region of radius R_h with random heading angles. The initial states differed across the 100 different runs and across each model. In the case of the Couzin model, the same initial states were used across each type of defection.

4.4 Defection in the Pairwise Interaction Model

In the pairwise interaction model, the interaction network of the group is fully connected (all individuals can communicate and interact with all other agents); hence, every agent can interact with any other agent. We use the same simulation algorithm of Jhawar et al. [78]. At each time step t, we chose at random $N/2$ agents as the focal agents that we pair with the rest of the group (non-focal agents). Each focal agent i (pink and hollow arrow in Fig. 4.1(b)) pairs with another unpaired agent j (thin and green quiver in Fig. 4.1(b)) chosen at random, irrespective of their Euclidean distance in space to copy $j's$ yaw angle, as shown in Fig. 4.1(c). Hence, only half of the population, the focal agent group, changes its motion direction (*i.e.*, the rate of change of direction is 0.5, as in [78]). At every simulation step ($dt = 0.1$ s), the focal agents are re-sampled and paired with other random non-focal agents. Mathematically, the desired yaw for a focal agent is described by:

$$\psi_i(t+1) = \begin{cases} \psi_j(t), & \text{if cooperating} \\ \psi_i(t), & \text{if defecting} \end{cases} . \tag{4.5}$$

Hence, the turning rate of a focal agent is described by

$$\omega_i = \begin{cases} \min\left(\frac{\psi_i(t+1)-\psi_i(t)}{dt}, \omega_{\max}\right), & \text{if cooperating} \\ 0, & \text{if defecting} \end{cases} . \tag{4.6}$$

Under the above setup, we perform Monte Carlo simulations with $N = 50$ agents. We find that regardless of the defection probability value P_D, the group ultimately gets polarized (*i.e.*, $p(T) > p^*$), except when individuals always defect (hence, for $P_D = 1$ we do not compute polarization time nor energy cost). Both polarization time (Fig. 4.2(a)) and average energy cost per individual (Fig. 4.2(b)) increase with increasing P_D, see time lapse of the simulations in Fig. 4.3. Such an increase is much more pronounced for high defection probabilities, *i.e.*, $P_D > 0.7$.

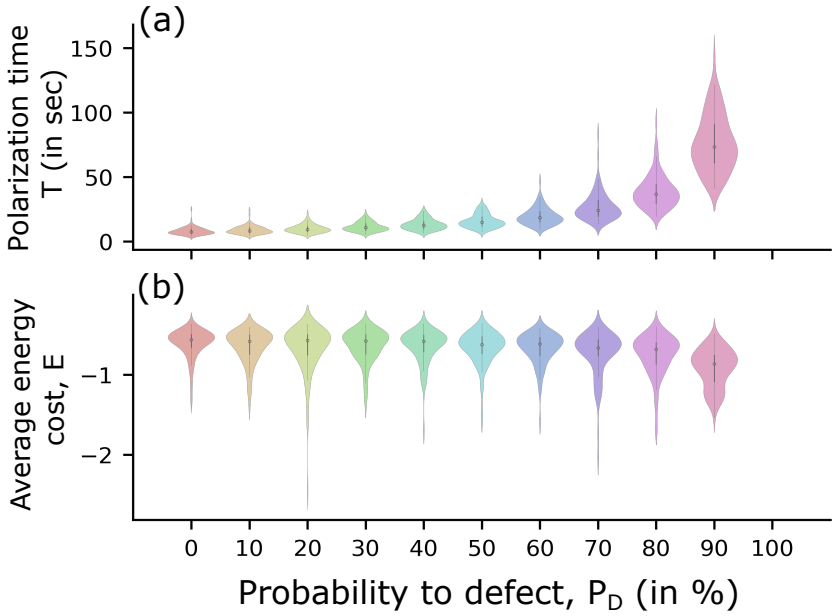


Figure 4.2: (a) The polarization time T and (b) the average energy cost E as a function of the probability to defect (P_D) in the pairwise model for $N = 50$ agents. The simulations were repeated 100 times per condition. We only report the data when at least 80% of the 100 simulation runs reached a polarization greater than $p^* = 0.8$. Hence, data for $P_D = 1$ are omitted as the population does not reach polarization.

4.5 Defection in the Couzin Model

In the Couzin model [2], an agent interacts with its neighbors within interaction range; therefore, the interaction network results from Euclidean distances between agents. The interaction range is categorized into three zones: the zone of repulsion (Z_r), the zone of orientation (Z_o), and the zone of attraction (Z_a). Different physical distances and orientations from a focal individual define each of its zones, as shown in Fig. 4.4(a). When neighbors are present in the interaction zone, the focal agent interacts by getting attracted, repelled, or orienting towards its neighbors as long as they are not in the blind zone (α). In this model, at each simulation step ($dt = 0.1$ s), every agent acts as a focal agent and updates its yaw angle as a combination of the interactions in the three zones as:

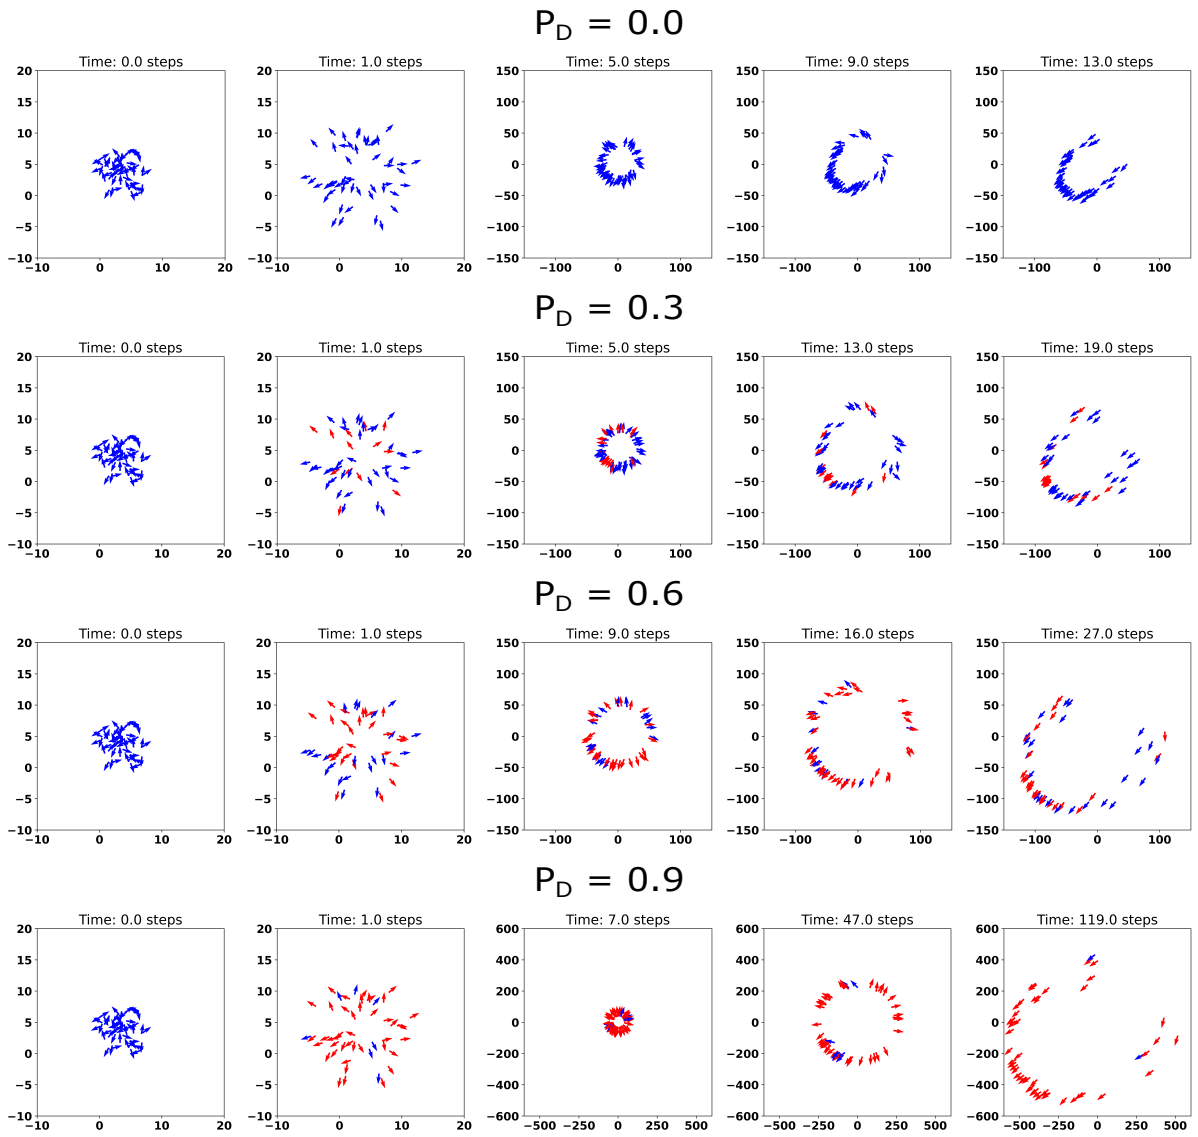


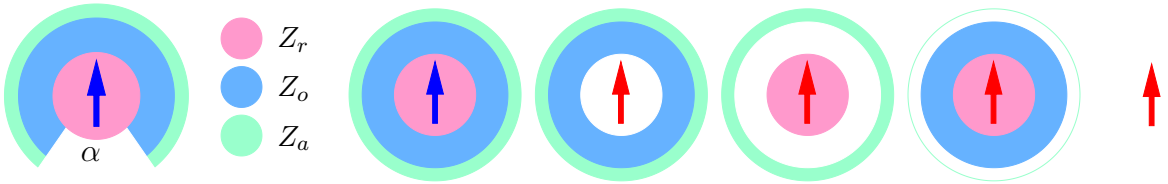
Figure 4.3: The time lapse of a collective of size $N = 50$ consisting of agents that defect with a certain probability of defection, P_D, and interact through a pairwise interaction model. In this figure, 1.0 steps correspond to 1.0 *sec*.

$$\psi_i(t+1) = \begin{cases} \tan^{-1}\left(\mathbf{d}_r^i(t)\right) & \text{if } Z_r^i \neq \phi \\ \tan^{-1}\left(\mathbf{d}_o^i(t)\right) & \text{if only } Z_o^i \neq \phi \\ \tan^{-1}\left(\mathbf{d}_a^i(t)\right) & \text{if only } Z_a^i \neq \phi \\ \tan^{-1}\left(\frac{1}{2}(\mathbf{d}_o^i(t) + \mathbf{d}_a^i(t))\right) & \text{if both } (Z_o^i, Z_a^i) \neq \phi \\ \psi_i(t) & \text{otherwise} \end{cases} . \tag{4.7}$$

Here, $\mathbf{d}_r^i = -\sum_{j \in Z_r} \hat{\beta}_{ij}$, $\mathbf{d}_o^i = \sum_{j \in Z_o} \hat{v}_j$, and $\mathbf{d}_a^i = \sum_{j \in Z_a} \hat{\beta}_{ij}$, where the symbol $\hat{\beta}_{ij}$ represents a unit vector pointing from the position of focal agent i towards the position of the

neighbor j (in the coordinate system of agent i) and $\hat{v}_j$ is the unit velocity vector. Hence, the desired ω_i of a focal agent at a given simulation step is obtained by:

$$\omega_i = \begin{cases} \min\left(\frac{\psi_i(t+1)-\psi_i(t)}{dt}, \omega_{\max}\right), & \text{if } i \text{ is cooperating} \\ 0, & \text{if } i \text{ is defecting} \end{cases} . \tag{4.8}$$



(a) Couzins model (b) Cooperator (c) $D - Z_r$ (d) $D - Z_o$ (e) $D - Z_a$ (f) $D - Z_{\text{all}}$

Figure 4.4: (a) Focal agents' interaction zones in the Couzin model depicting the zone of orientation (Z_o), zone of attraction (Z_a), zone of repulsion (Z_r), and the blind region (α). Types of defection in the Couzin model: (b) standard Couzin model with repulsion, orientation, and attraction zones, (c-f) types of defection, white regions indicate no interaction (*i.e.*, defective interaction) for each interaction zone.

In contrast to the pairwise model, we conduct Monte Carlo simulations by defining defection in four different ways based on the zones of behavioral interaction: (1) defection only in the zone of repulsion (D-Z_r, Fig. 4.4(c)), (2) defection only in the zone of orientation (D-Z_o, Fig. 4.4(d)), (3) defection only in the zone of attraction (D-Z_a, Fig. 4.4(e)), and (4) defection in all the three zones at the same time (D-Z_{all}, Fig. 4.4(f)). Given the modified Couzin model defined in Eqs.(4.7) and (4.8), we perform simulations with $N = 20$ and $N = 50$ agents using similar parameter values to the ones used in the seminal work introducing the Couzin model [2], as indicated in Table 4.1 in detail.

First, for each case of N ($N = 20$ and $N = 50$), we examine and report the probability of the group becoming highly polarized by the end of the experiment, *i.e.*, $p(t_{\max}) > p^*$. This probability is computed by computing the proportion of the 100 simulation runs (per condition) that are in a polarized state at time $t_{\max} = 200s$. The results in Fig. 4.5 show that in the case of defection of type $D - Z_r$ and $D - Z_a$, both groups of $N = 20$ and $N = 50$ agents consistently reach high polarization levels irrespectively of the defection probability P_D, including the case of $P_D = 1$. However, in the case of $D - Z_{\text{all}}$, the probability of $p(t_{\max}) > p^*$ stays below our cut-off line set at 0.8 (horizontal dotted grey line) when defection probabilities are higher than $P_D > 0.6$ for both group sizes. In the

Parameter name	Symbol	Value	Unit
Radius of repulsion	R_r	1	units
Radius of orientation	R_o	18	units
Radius of attraction	R_a	20	units
Linear speed	$\|v\|$	5	units/s
Maximum turning rate	$\omega_{\max}$	$\pi/3$	rad/s
Field of invisibility	α	0	rad
Radius of home	R_h	$\sqrt{N}$	units

Table 4.1: Parameters used in the Couzin model for highly parallel formation.

case of $D - Z_o$, the system does not reliably reach a polarized state ($p(t_{\max}) < p^*$) in 200 seconds for $P_D > 0.7$ when $N = 20$ and $P_D > 0.3$ when $N = 50$. It means the collective ability to reach a polarized state with defection in zone Z_o decreases with an increasing number of agents.

Figures 4.6 and 4.7 show the results of the polarization time T and the average energy cost E for simulations of groups of $N = 20$ and $N = 50$ agents, respectively. The analysis of these two metrics reveals significant differences depending on the zone where agents probabilistically defect. For both group sizes, defection in the orientation zone Z_o leads to the negative effects of an increase in the polarization time T and higher average energy costs E in maintaining a polarized state (see panels (e-f) of Figs. 4.6 and 4.7). We observe the opposite trend for defection in the repulsion zone Z_r (see panels (g-h) of Figs. 4.6 and 4.7, and Table 4.2) where, for both group sizes, the polarization time T and average energy costs E decrease as the defection probability P_D increases. This trend is monotonic in all cases except for E with $N = 50$, where the minimum energy cost is for $P_D = 0.6$, see Fig. 4.7(h) and Table 4.2. For defection in the attraction zone Z_a, we find different results for the two considered swarm sizes. In small swarms, $N = 20$ in Figs. 4.6(c-d), the inclusion of defection probability does not impact T but worsens average energy cost E. Instead, in larger swarms, $N = 50$ in Figs. 4.7(c-d), polarization time T decreases for moderate values of defection probability $P_D \leq 0.3$ after plateauing at a constant T.

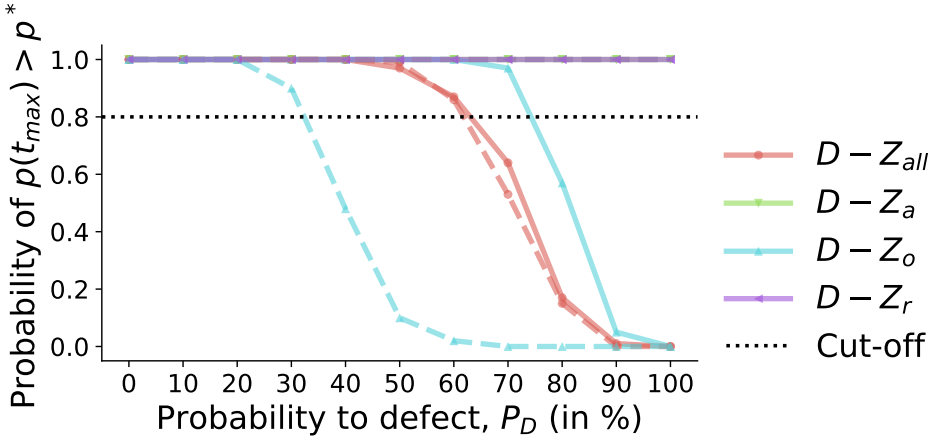


Figure 4.5: Probability of having a polarized group (*i.e.*, $p(t_{\max}) > p^*$ with $p^* = 0.8$) as a function of the agent's probability P_D to defect. This probability is computed as the proportion of 100 simulation runs where, at time $t_{\max}$, the group had polarization $p(t_{\max})$ above $p^* = 0.8$. The solid lines are for $N = 20$, and the dashed lines are for $N = 50$. In the results of Figs. 4.6 and 4.7, we do not include data for conditions where the probability of $p(t_{\max}) > p^*$ is below our cut-off line of 0.8 (dotted horizontal black line).

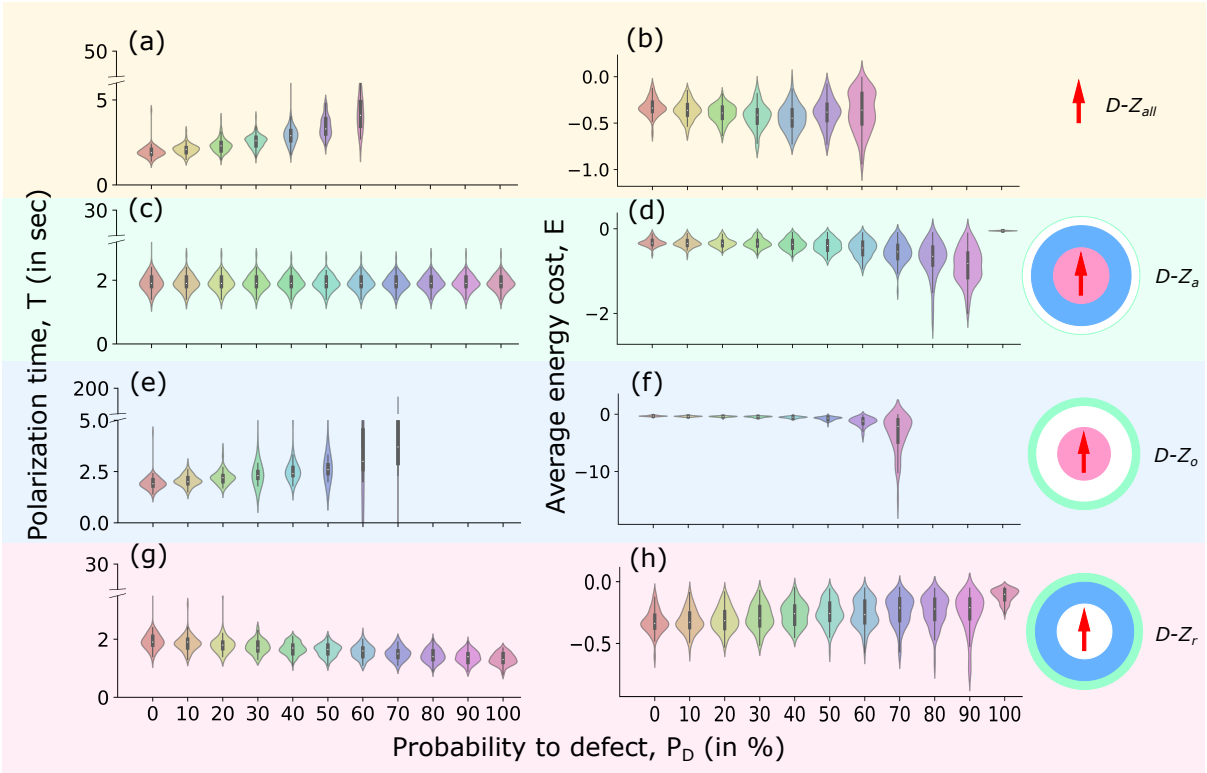


Figure 4.6: The polarization time T and the average energy cost E across four types of defection for a group of $N = 20$ agents. The violin plot shows the distribution of values from 100 simulation runs.

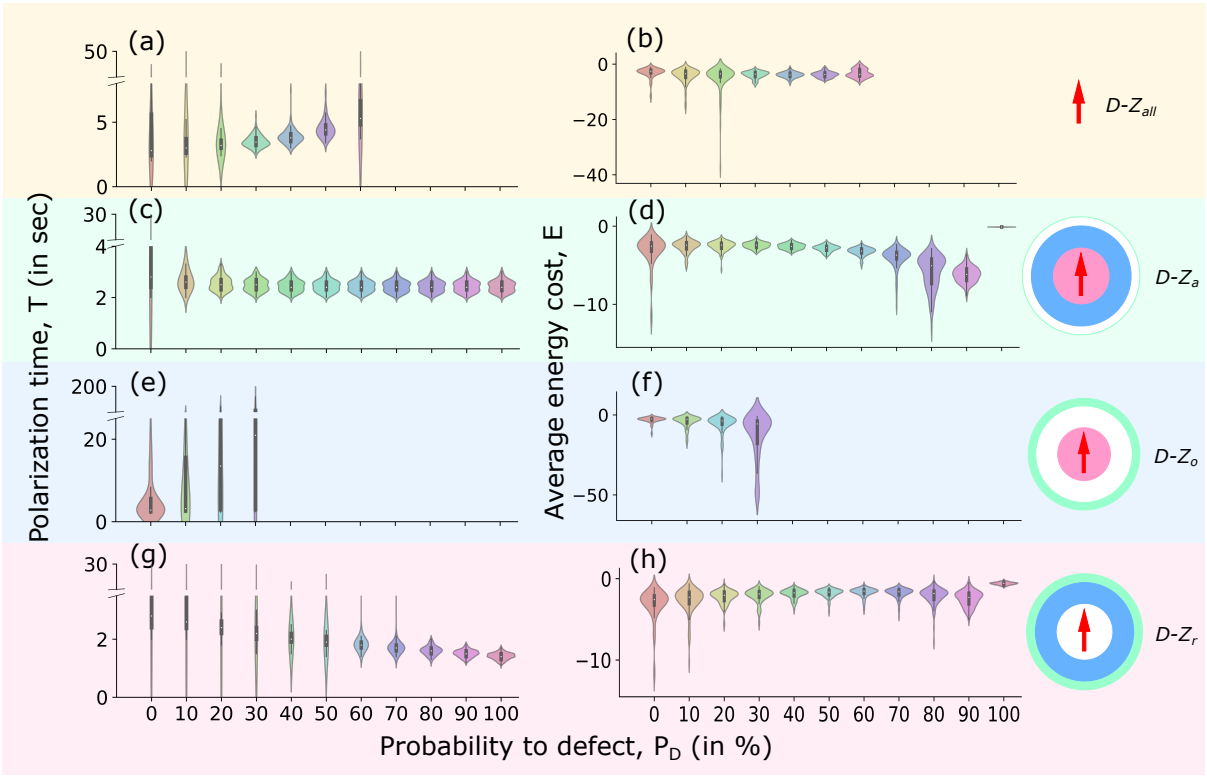


Figure 4.7: The polarization time T and the average energy cost E across four types of defection for a group of $N = 50$ agents. The violin plot shows the distribution of values from 100 simulation runs.

	0	10	20	30	40	50	60	70	80	90	100
μ_e	-3.109	-2.648	-2.234	-2.116	-1.877	-1.776	-1.697	-1.775	-2.140	-2.482	-0.595
σ_e	1.944	1.522	0.82	0.875	0.580	0.536	0.562	0.624	1.028	0.939	0.167

Table 4.2: Mean μ_e and standard deviation σ_e of the average energy cost E for D-Z_r, same data as in Fig. 4.7(h).

Average energy cost E improves for $P_D < 0.3$, is minimum at $P_D = 0.3$, and increases for higher P_D (see Table 4.3). Finally, when agents defect in all three zones Z_{all}, the dynamics are quite complex as they are a combination of the trends we observed in the three areas. Polarization time increases for $N = 20$ and has a minimum at $P_D = 0.3$ for $N = 50$ (see Fig. 4.7(a) and Table 4.4). Average energy cost E has opposite trends for $N = 20$ and $N = 50$. In the former case, Fig. 4.6(b) shows a u-shape trend, with a maximum energy cost for $P_D = 0.3$, and in the latter case, Fig. 4.7(b) shows a bell-shape trend with minimum energy cost for $P_D = 0.3$.

Note that in Figure 4.8, the polarization time required for agents to get polarized

	0	10	20	30	40	50	60	70	80	90
μ_e	-3.109	-2.586	-2.538	-2.459	-2.595	-2.841	-3.235	-4.021	-5.891	-6.199
σ_e	1.944	0.769	0.59	0.391	0.371	0.404	0.51	1.294	2.129	0.993

Table 4.3: Mean μ_e and standard deviation σ_e of the average energy cost E for $D - Z_a$, same data as in Fig. 4.7(d).

	0	10	20	30	40	50	60
μ_t	5.501	4.773	3.665	3.522	3.877	4.564	7.408
σ_t	5.713	6.671	3.044	0.482	0.592	0.851	7.167

Table 4.4: Mean μ_t and standard deviation σ_t of the polarization time T for $D - Z_{all}$, same data as in Fig. 4.7(a).

decreases at a certain probability of defection, P_D. As the value of P_D increases, the group becomes more prone to splitting into a larger number of clusters at earlier stages of the alignment process. Further, in Figure 4.9, the polarization time required for agents to get polarized decreases with increasing probability of defection, P_D. With increasing values of P_D, the group takes a formation more similar to that of the pairwise interaction model. In Figure 4.10, the polarization time required for agents to get polarized increases with increasing probability of defection, P_D. There are certain values of P_D at which the group partially transitions to the milling formation of the Couzin model. Further, in Figure 4.11, the polarization time required for agents to get polarized decreases with increasing probability of defection, P_D. As the value of P_D increases, the group becomes densely packed.

Figure 4.12 shows the impact of defectors on the collective motion when the group size is $N = \{20, 50, 80\}$. When the agents defect in all zones simultaneously, the collective resistance to the defectors decreases slowly with increasing group size. Whereas for the defection of type $D - Z_a$ and $D - Z_r$, the resistance to the defectors decreases at a smaller probability to defect when the group size is large. However, in the defection of type $D - Z_o$, the resistance to defectors decreases fast with increasing group size. Hence, we conclude that the impact of defectors is amplified in unexpected ways at larger group sizes. In

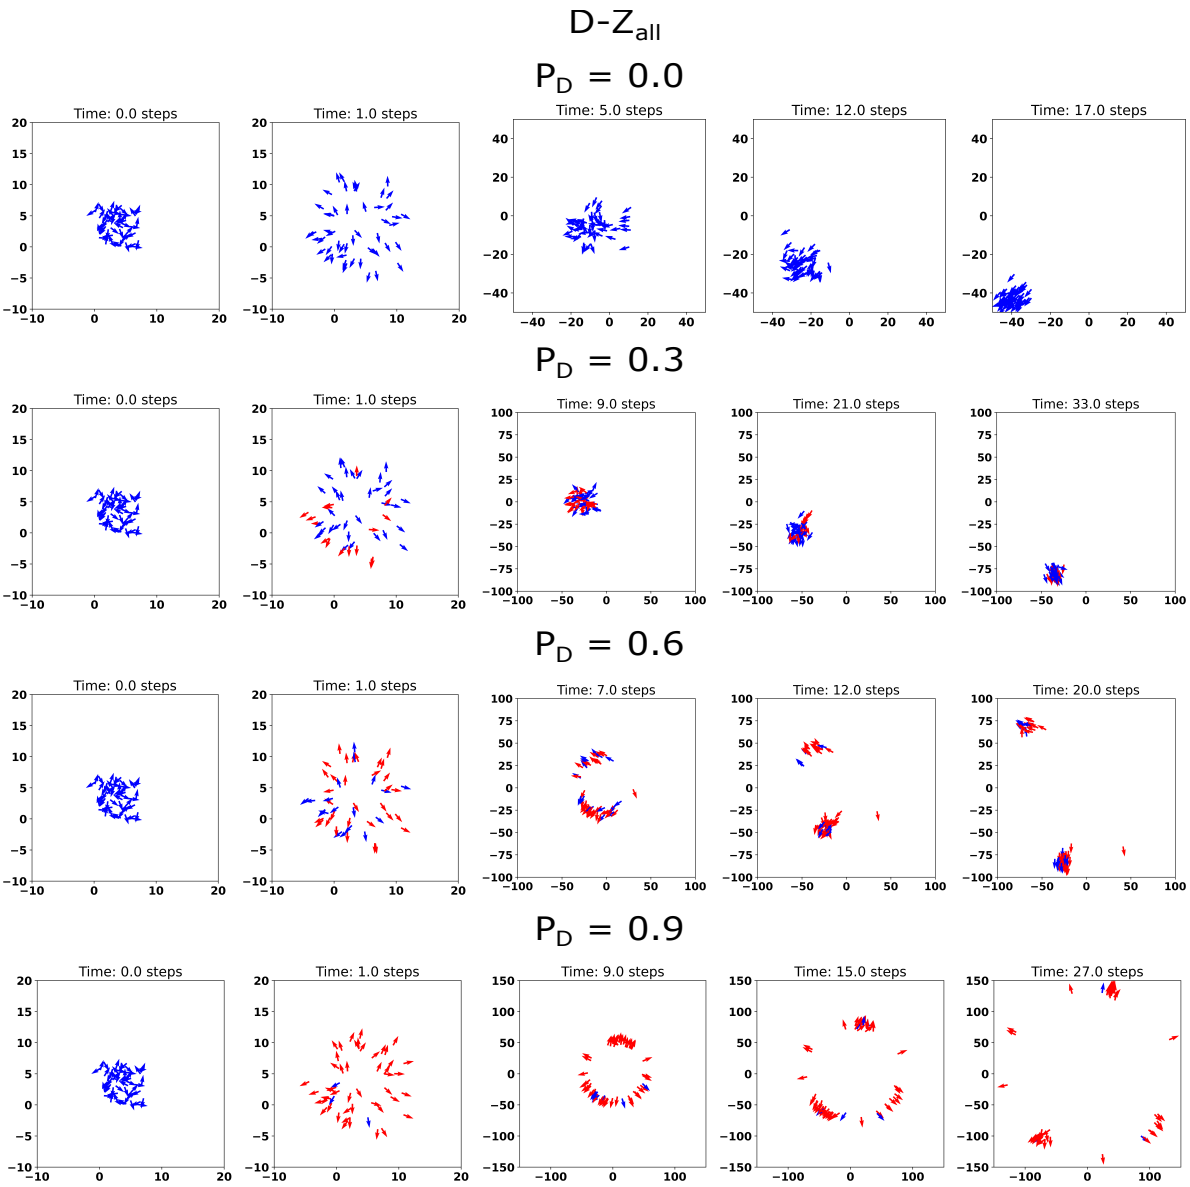


Figure 4.8: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_{all})$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec*.

the case of group size, $N = 80$, and the considered set of parameters (Table 4.1), the group takes longer than our simulation time of 200 seconds [32] to transit to the ordered phase, see Fig. 4.13 for the case of $N = 80$. The polarization time decreases when defection is of type $D - Z_a$ and $D - Z_r$ (Figs. 4.13(c, e)) with the increase in probability of defection. Whereas, for the defection of type $D - Z_{\text{all}}$, it decreases and then increases with a minimum at $P_D = 0.3$. In accordance with that, the average energy cost decreases for the defection of type $D - Z_{all}$ and $D - Z_r$ (Figs. 4.13(b, f)) with increasing probability of defection. On the other hand, in the case of $D - Z_{\text{a}}$ (Fig. 4.13(d)), the energy cost

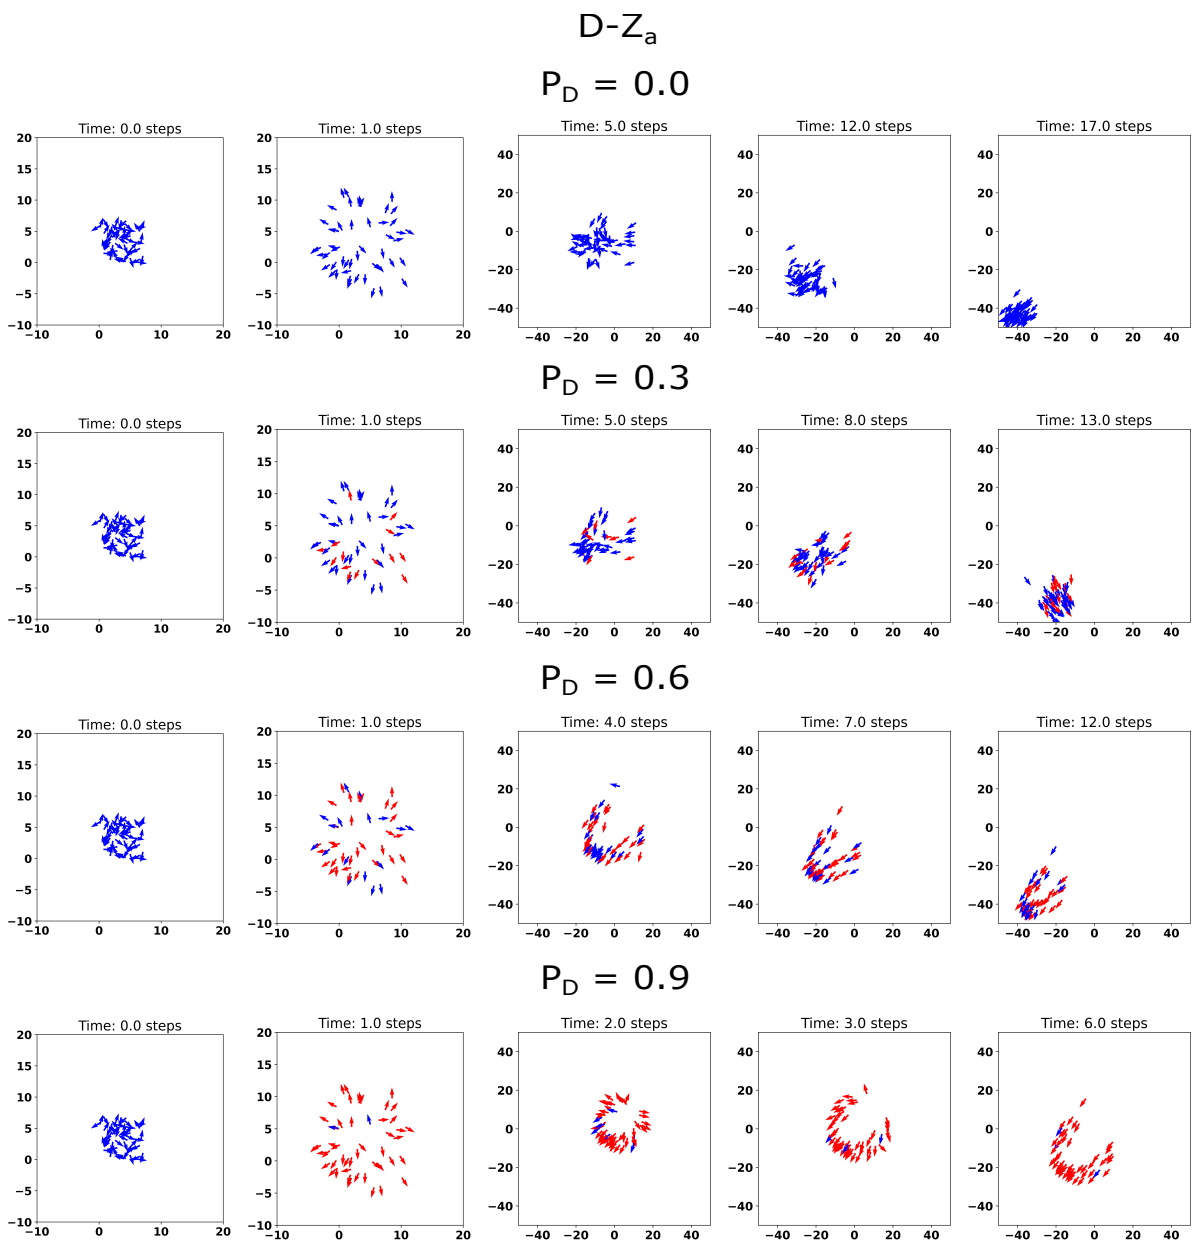


Figure 4.9: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_a)$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec*.

decreases and then increases (with a minimum around $P_D = 0.4$) with the increasing probability of defection. In case of the defection of type $D - Z_o$, the collective fails to get polarized above the threshold probability of $p(T) > 0.8$ irrespective of whether the agents defect or not. Counterintuitively, in $D - Z_r$, both the polarization time and the average energy cost decrease, indicating the benefits of incorporating defective behavior with $P_D = 0.4$ (Figs. 4.13(e, f)).

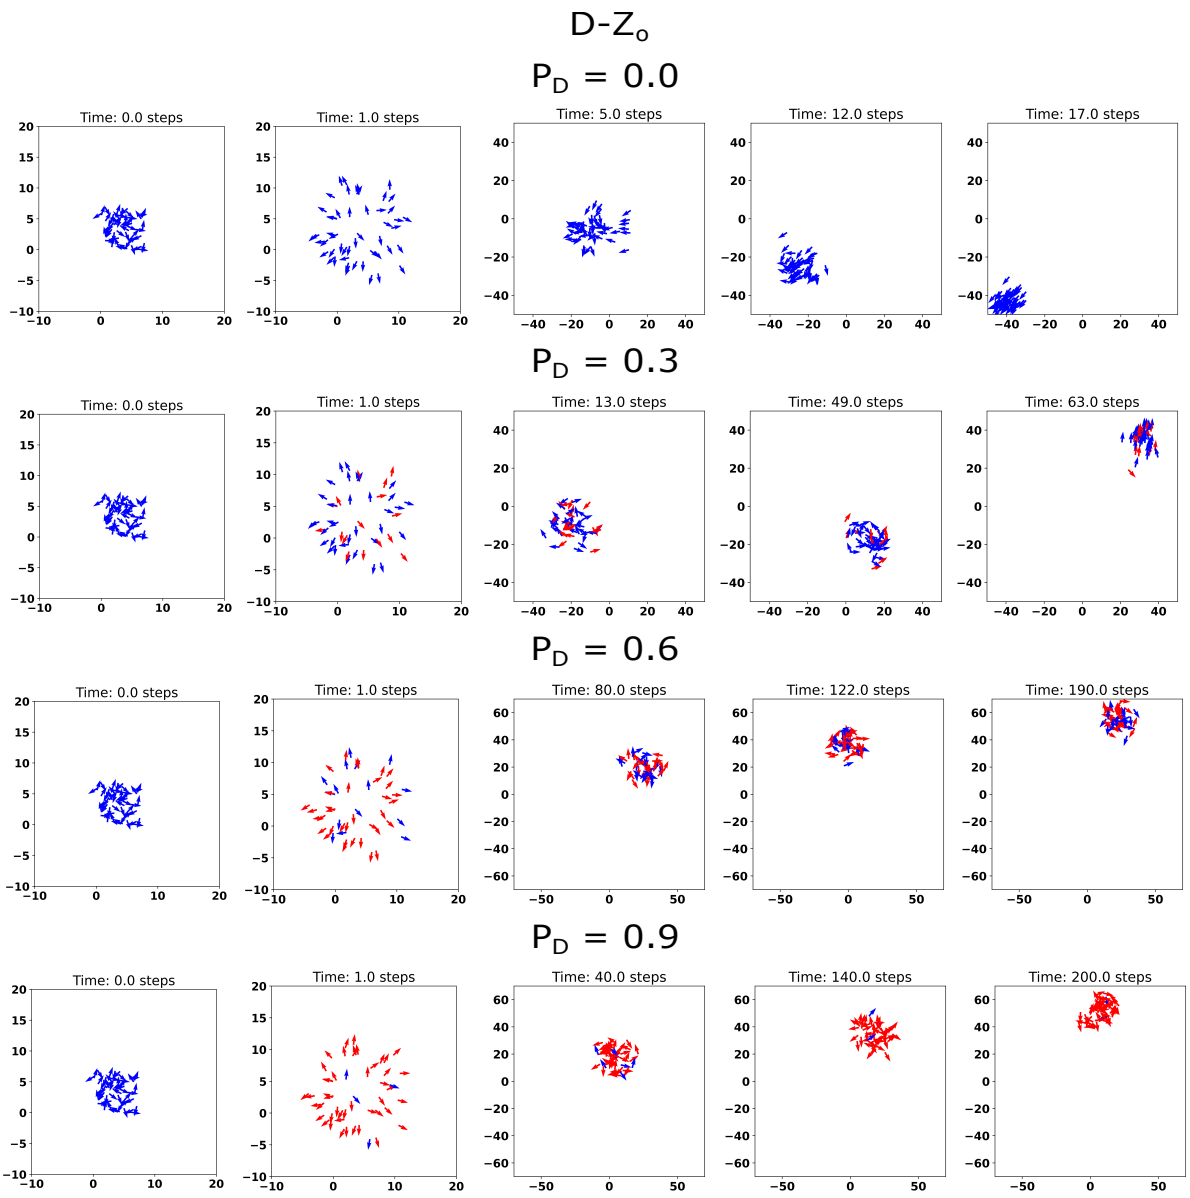


Figure 4.10: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_o)$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec*.

4.6 Conclusion

While most of the existing studies on collective motion models assume that individuals always obey the interaction rules without fail, in real-world natural and robotic systems, agents may defect and not follow the interaction rules for various reasons, such as selfish behavior, intra-group competitions, and adversarial attacks. In this chapter, we introduced probabilistically defecting agents in two—non-spatial and spatial—models of collective motion: the pairwise model [78] and the Couzin model [2]. We analyzed the

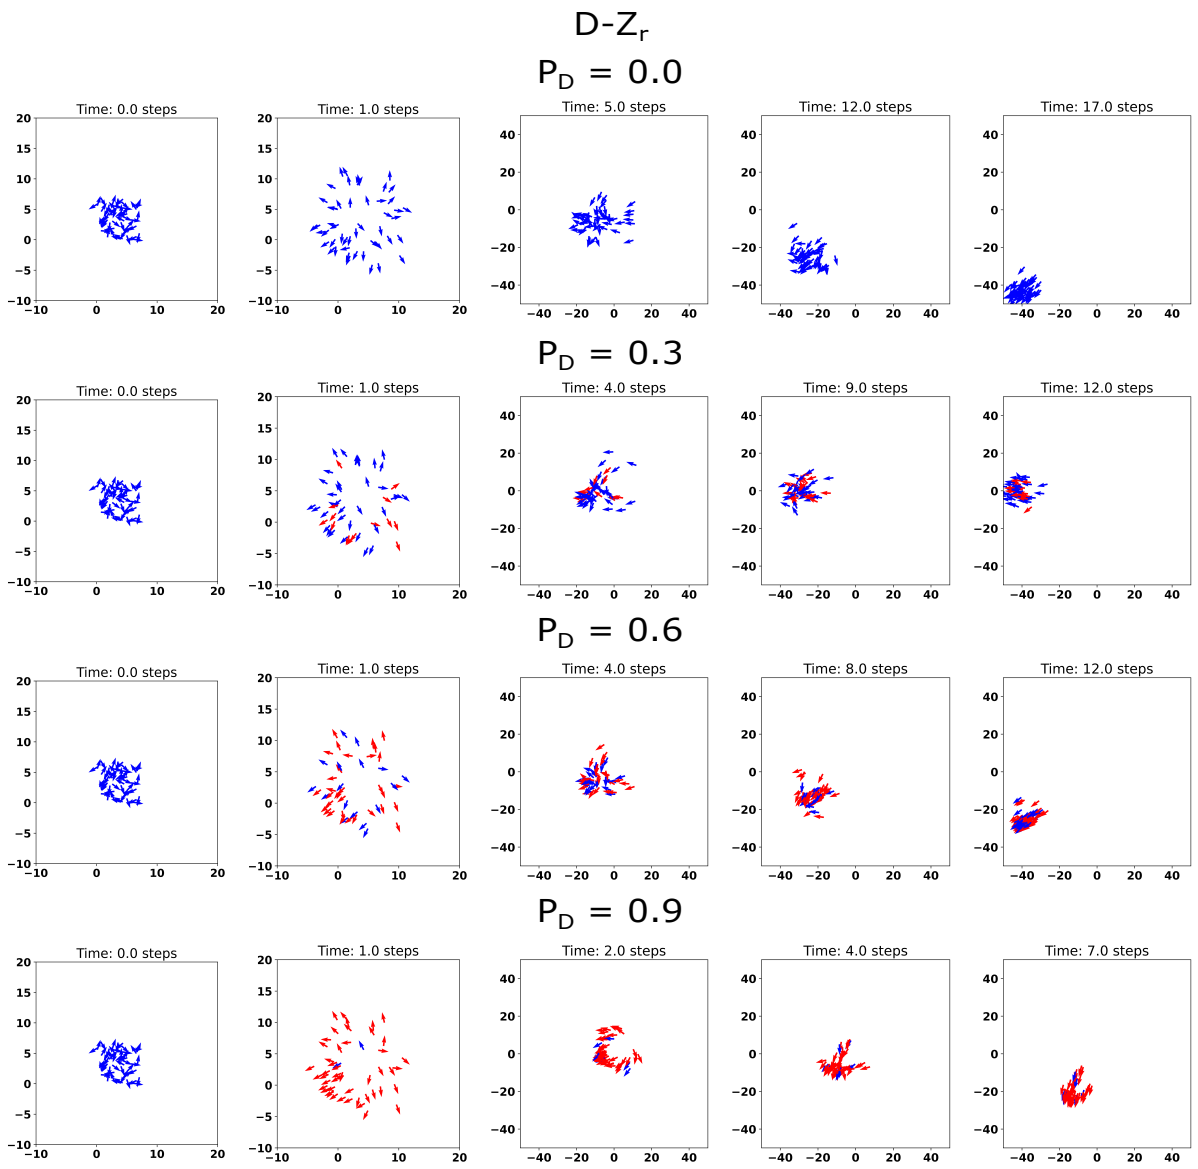


Figure 4.11: The time lapse of a collective of size $N = 50$ consisting of agents $(D - Z_r)$ that defect with a certain probability of defection, P_D. The agents interact as dictated by the Couzin model. In this figure, 1.0 steps correspond to 1.0 *sec*.

impact of probabilistically defecting agents on the collective motion dynamics by measuring the polarization time T (time needed to reach polarization $p(t)$ above $p^* = 0.8$, $p(T) > p^*$) and the average energy cost E per individual (energy that each agent spends on average to make rotational movements once the group has reached a highly polarized state).

Our analysis of the pairwise interaction model shows that both the polarization time and the average group energy cost increase with increasing defection probability. The Couzin model is a spatial model; thus, we analyzed four different types of defection

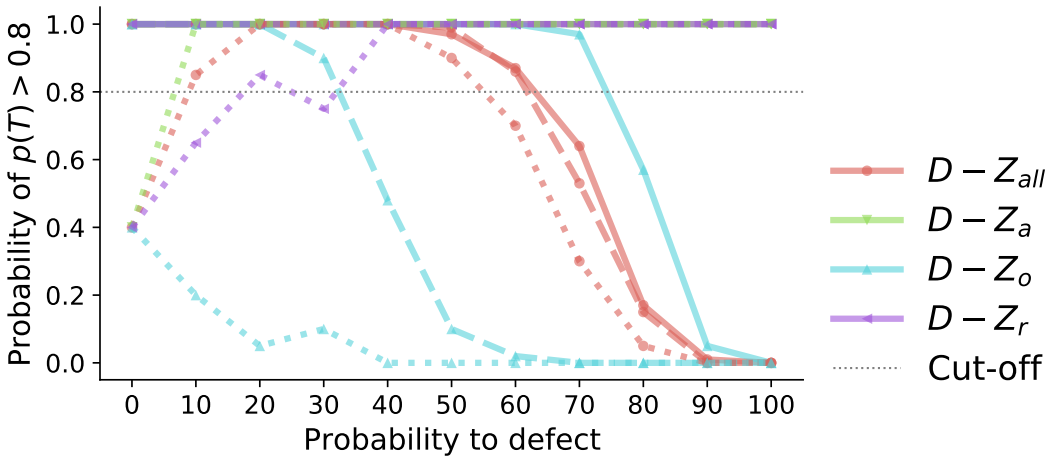


Figure 4.12: The plot shows the probability of $p(T) > 0.8$ as a function of the probability of defection. The solid lines are for $N = 20$, the dashed lines are for $N = 50$, and the dotted lines are for $N = 80$.

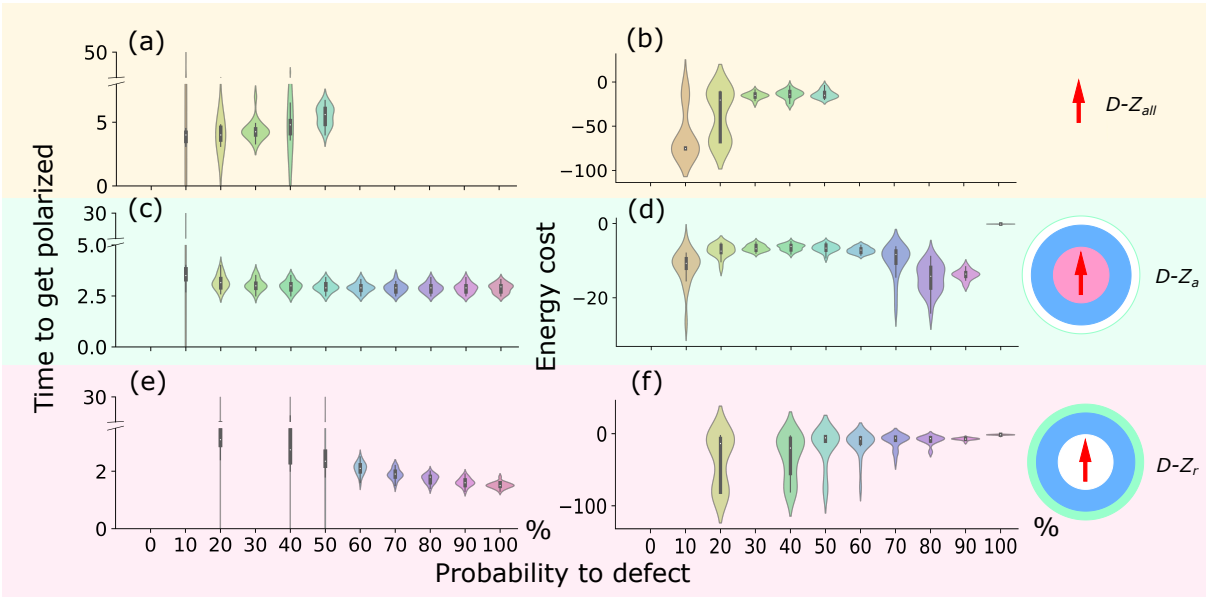


Figure 4.13: The polarization time (T) and the average group energy cost ($\Sigma_{t,N} e(t > T)$) across three types of defection for a group of size $N = 80$. Each violin depicts a distribution of values from 100 iterations.

depending on the interaction zones where agents occasionally defect. If agents defect in either the zone of attraction or repulsion, the polarization time reduces with increasing defection probability P_D. With defection in these two zones, the average group energy cost E also has a minimum for defection probability $P_D > 0$. Differently, for defection in the zone of orientation, any value of $P_D > 0$ reduces the performance in terms of both polarization time and average energy cost. In summary, in the considered spatial model (the Couzin model), there are conditions (*e.g.*, defection in the repulsion zone only) that

can provide group-level benefits and conditions (*e.g.*, defection in the orientation zone only) that are harmful to the group. The system dynamics and their relationship with the investigated control parameters (defection probability, swarm size, defection zones) are relatively complex. Additional complexity is introduced into the spatial model by the possibility that agents split into subgroups moving in different directions, each with local polarization. In the next chapter, we further extend our study of the spatial model of flocking to determine whether the outcomes of defection extend across the model's parameter space and to deduce the mechanism underlying the observed effects in the collective behavior.

CHAPTER 5

FLOCKING PHASE TRANSITION AND ANONYMOUS EMERGENT LEADERSHIP VIA INDIVIDUAL DEFECTION

5.1 Overview

Existing collective motion models assume that each individual obeys the interaction rule at all times. However, in animal collectives, one does not expect every individual to always react in accordance with the interaction rules. In Chapter. 4, we defined the individuals who occasionally disobey the interaction rule as defector individuals. We observed that when individuals defect, they help conserve collective group energy and reduce the time required for group alignment. In this chapter, we proposed a framework that promotes effortless human-swarm interaction with flocks through anonymous leadership, enabling direct manipulation of individuals' probability to defect. We considered a group of individuals whose interaction rules are described by the metric-based zonal model (the Couzin model). The individuals, while interacting probabilistically, defect with a fixed probability to defect. Prior studies have led researchers to believe that collective phases change when the model parameters, such as the radius of interaction, forward speed, or turning rate, are altered. However, from engineering point-of-view, regulating all these parameters might require a much more complex decision-making process. Our simulations show that individuals flocking with contrasting strategies, such as cooperation and defection, can switch between collective formations, such as milling or highly parallel formations, only by regulating their probability to defect while maintaining the same interaction model parameters. Further investigations shows that the observed switching of collective behavior is an outcome of reduction in effective turning rate of the individuals.

5.2 Introduction

Collective motion arises when individuals interact locally and yet produce coherent group-level organization. Such systems display a diversity of collective states, including polar-

ized flocks, rotating mills, and disordered swarms, across taxa ranging from bacteria to birds and fish [2, 77, 158]. Understanding how these patterns emerge and how groups transition between them remains a central problem in the study of collective behavior.

A wide range of theoretical and computational models have been developed to capture collective motion, typically prescribing interaction rules based on repulsion, alignment, and attraction [2, 77, 159]. In these frameworks, transitions between collective states are commonly induced by changing interaction ranges, turning rates, noise amplitudes, or densities [160, 161]. Such approaches have provided valuable insight into the mechanisms that promote order, disorder, and pattern formation.

Empirical studies increasingly emphasize that individuals within a group differ in their sensory capabilities, internal states, and behavioral tendencies, and that such heterogeneity can influence collective outcomes [162, 163]. Real biological and robotic collectives therefore consist of individuals whose responses may be intermittent, owing to sensory limitations, attention lapses, internal states, or stochastic decision processes. Moreover, animals often switch between collective states without obvious changes in group size or environmental conditions, suggesting that additional, internal sources of variability may play an important role in shaping group-level dynamics [164]. Incorporating simple forms of individual-level variability into collective motion models provides a route to exploring how such flexibility may arise.

An implicit assumption shared by many models is that individuals continuously and faithfully execute the prescribed interaction rules. Probabilistic formulations of collective behaviour have been considered previously, including models in which interactions occur stochastically or noise induces novel collective modes [165–167]. Most existing approaches, however, introduce stochasticity by perturbing headings, speeds, or interaction strengths, or by randomly sampling interactions in time [2, 77]. A complementary formulation is one in which interaction rules are occasionally not executed, so that individuals temporarily move independently of their neighbours. How such intermittent suspension of interactions shapes collective dynamics remains less well understood.

Understanding how collective behaviour responds to intermittent or unreliable interactions is relevant for both biological and engineered systems [168]. In animal groups, sensory limitations, internal states, and fluctuating attention may lead to intermittent

interactions that shape collective outcomes. Similarly, in engineered swarms, communication delays, sensor failures, and hardware faults are unavoidable. These considerations motivate the need for modelling frameworks that explicitly examine how intermittent interactions influence collective dynamics.

Here, we consider a minimal extension of a classical zonal model [2] in which individuals occasionally ignore interaction rules with a specified probability. This formulation allows us to isolate the effects of intermittent rule-following while keeping all other aspects of the model unchanged. By systematically varying the probability of rule violations, we examine how collective patterns are altered and how transitions between different modes of organization arise. We find that intermittent rule-following can substantially reshape collective dynamics and induce transitions between distinct collective states, even in the absence of changes to interaction ranges or speeds. Our results highlight intermittent interactions as an additional dimension of collective behaviour.

5.3 Flocking model of probabilistically defecting agents

Consider a group of individuals $G = \{b_i | i = 1, 2, 3, \ldots, N\}$ of size N, where b_i is the i^{th} member of the group G, at position $\mathbf{q}_i(t) = (x_i(t), y_i(t))$, with heading direction $\psi_i(t)$ at time instance, t. We assume that when the individuals move, they move with an equal and constant forward speed, $\|\mathbf{v}_i\|$, and a turning rate, $\omega_i(t)$ (see Fig. 5.1a), constrained with a maximum turning rate of, $\omega_{\max}$, as described by the Dubin's car model [155]. The equations of kinematic motion of the individuals at every time step dt are given by:

$$\begin{aligned} \dot{x}_i &= \|\mathbf{v}_i\| \cos(\psi_i(t)) \,, \\ \dot{y}_i &= \|\mathbf{v}_i\| \sin(\psi_i(t)) \,, \\ \dot{\psi}_i &= \omega_i(t), \ \|\omega_i\| \leq \omega_{\max} \,. \end{aligned} \tag{5.1}$$

Each individual has three zones of interaction, namely the zone of repulsion (Z_r), the zone of orientation (Z_o), and the zone of attraction (Z_a), as well as a blind region in Z_o and Z_a subtending an angle α. The zone Z_r spans a circular region of radius R_r, whereas the zones Z_o and Z_a cover sectorial ring regions with outer radii R_o and R_a, respectively.

A schematic diagram in Fig. 5.1b shows the Z_r (in red), the Z_o (in blue), and the Z_a (in green).

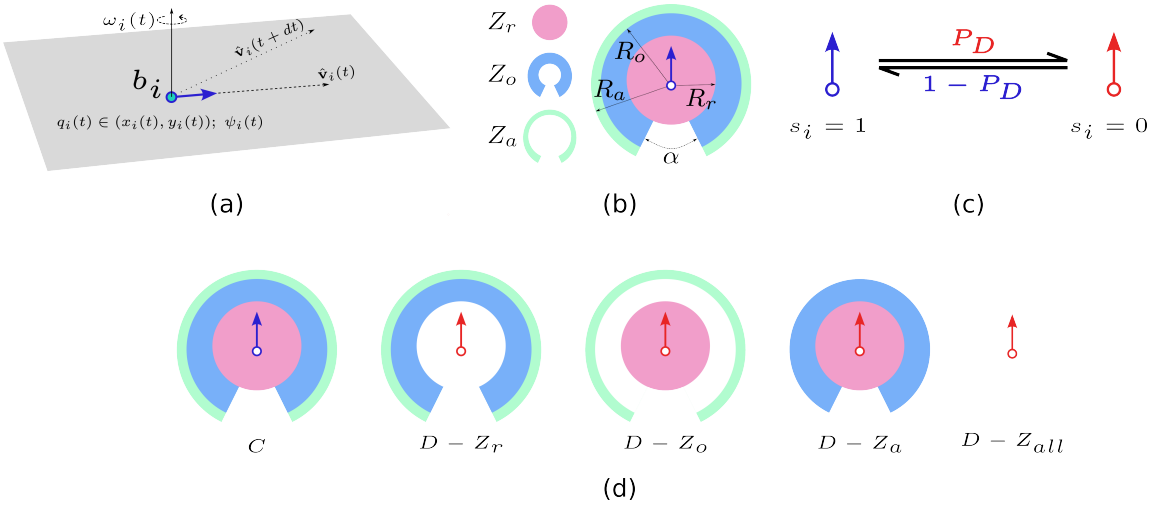


Figure 5.1: (a) An individual b_i, at time t, has a position $q_i(t) = (x_i(t), y_i(t))$, a heading direction $\psi_i(t)$, and a constant speed $|\vec{v}_i(t)|$. At any given time, the new direction $\psi_i(t + \Delta t)$, *i.e.*, after time step dt, is determined by a combination of the individual's previous direction and the social forces. Hence, while moving at the constant speed, the individual may switch its direction resulting in an angular velocity $\omega_i(t)$, as described by the Dubin's car model [155]. (b) Each individual has three zones of interaction, namely the zone of repulsion Z_r, the zone of orientation Z_o, and the zone of attraction Z_a. A sectorial region subtending an angle of α represents a blind zone, in Z_o and Z_a, due to which the individual does not interact with neighbors present in the blind zone of the focal individual. (c) The blue and red colored arrows represents the two strategies or states in the model. A focal individual can be in any of these two states. The first one, shown in blue, is a responsive or cooperative state, where the focal individual takes the social force into consideration for determining its next heading. The second state, shown in red, is a non-responsive or a defective state, where the focal individual avoids the social force for determining its next heading direction. Every individual can switch its state bidirectionally, *i.e.*, from $s = 1$ (Cooperator (C)) to $s = 0$ (Defector (D)) with a probability P_D, and from D to C with probability, $1 - P_D$. The switching probability is fixed at the start of a given simulation, for e.g., $P_D = 0.5$. (d) A focal individual can defect in multiple ways, such as if it ignores neighbors only in Z_r, then the defector is of type $D - Z_r$, and so on. When an individual ignores neighbors in all zones simultaneously, then the defector individual is of type $D - Z_{all}$. There are only one type of defectors in a given simulation run.

Each individual has a strategy, $s_i(t) \in \{0, 1\}$. The strategy of an individual at a given instance represents the state of the individual's decision to interact (cooperate, $s_i = 1$) or not (defect, $s_i = 0$). The individuals can switch their strategy from cooperation to defection with a fixed probability of P_D, as shown in Fig. 5.1c. As the individual has three zones of interaction, it can be defective in any of the three zones or in all zones simultaneously. If the individuals defect in Z_r, we annotate it as $D - Z_r$; if the individual

defects in all zones altogether, we annotate it as $D - Z_{all}$, and so on (see Fig. 5.1d). We annotate the type of defection in strategy symbol as $s_i^z(t)$, where each letter in superscript $z \in \{r, o, a\}$ points to the zone of defection. For instance, if the individual defects in all three zones, we denote it as $s_i^{roa} = 0$. At each time instance, t, individuals interact with their neighbors within the interaction zones as described by the interaction rules:

$$\mathbf{v}_i(t + \mathrm{d}t) = \begin{cases} -\sum_{j \in Z_r^i} s_i^r s_i^{roa} \hat{\mathbf{d}}_{ij}(t), & \text{if } Z_r^i \neq \varnothing, \\ \frac{1}{2}\left[\sum_{j \in Z_a^i} s_i^a \hat{\mathbf{d}}_{ij}(t) + \sum_{j \in Z_o^i} s_i^o \hat{\mathbf{v}}_j(t)\right] s_i^{roa}, & \text{if } (Z_o^i \text{ or } Z_a^i) \neq \varnothing, \\ \hat{\mathbf{v}}_i(t), & \text{otherwise.} \end{cases} \tag{5.2}$$

Here, $\hat{\mathbf{v}}_j$ is the unit velocity vector of individual j, and $\hat{\mathbf{d}}_{ij}(t)$ is the unit distance vector between individual i and j, given by:

$$\hat{\mathbf{d}}_{ij}(t) = \frac{\mathbf{q}_j(t) - \mathbf{q}_i(t)}{\|\mathbf{q}_j(t) - \mathbf{q}_i(t)\|}. \tag{5.3}$$

Based on the interaction forces, the desired heading angles (with a Gaussian noise) of each individuals is obtained using:

$$\psi_i(t + \mathrm{d}t) = \tan^{-1}\left[\frac{v_y(t + \mathrm{d}t)}{v_x(t + \mathrm{d}t)}\right] \pm \mathcal{N}(0, 0.05), \tag{5.4}$$

where v_x and v_y are the components of the unit velocity vector $\hat{\mathbf{v}}_i$. The resulting angular velocity, $\omega_i(t)$, is obtained as:

$$\omega_i(t) = \frac{\psi_i(t + \mathrm{d}t) - \psi_i(t)}{\mathrm{d}t}. \tag{5.5}$$

At time $t = 0$, let the group, G, be situated in a circular start region, H, of radius R_H. Let the individuals be distributed uniformly in H at random positions $q_i(t = 0)$

with random heading directions $\psi_i(t = 0)$, and strategies $s_i(t = 0)$. Our aim in this study is to determine whether the group, G, still exhibits flocking behavior when individuals occasionally defect. We investigate if the group is still able to switch flocking formations as described by the original Couzins zonal flocking model [2]. Lastly, determine whether either the cooperators or the defectors may act as the leaders. The fixed flocking model parameters used throughout this study (unless mentioned mainly for the individual kinematics and obstacle avoidance) are listed in Table 5.1.

Table 5.1: List of fixed individual parameters (for highly aligned formation based on [2]) used in this study:

Parameter name	Symbol	Value	Unit
Radius of repulsion	R_r	1	units
Linear forward speed	$\|\mathbf{v}_i\|$	3	units/s
Maximum turning rate	$\omega_{\max}$	$2\pi/9$	rad/s
Field of invisibility	α	$\pi/2$	rad

5.3.1 Model parameters and notation

To investigate and report the changes in group behavior resulting from the defection strategy, we consider the following metrics and perform Monte Carlo simulations with N_{rep} repetitions of the simulations. In each repetition we make an observation of the metrics at T (200 sec) time period.

Alignment order parameter, $\Phi(t)$: Alignment order parameter refers to the mean of the heading direction of all individuals in the group; it is a standard measure of collective alignment in models of collective motion [2, 77], obtained as:

$$\Phi(t) = \frac{|\sum_{i=1}^{N} \hat{\mathbf{v}}_i(t)|}{N} \tag{5.6}$$

Milling order parameter, $\Lambda(t)$: It refers to the magnitude of the sum of the vector cross product of the instantaneous individual heading directions, $\hat{\mathbf{v}}_i(t)$, and the instan-

taneous relative unit position vector of the individuals from the center of mass of the group, $\hat{\mathbf{r}}_i^{rel}(t)$.

$$\Lambda(t) = |\sum_{i=1}^{N} \hat{\mathbf{r}}_i^{rel}(t) \times \hat{\mathbf{v}}_i(t)| \tag{5.7}$$

Probability of alignment, $P(\Phi(T))$: It refers to the frequency of repetitions out of the N_{rep} repetitions where the group aligns with $\Phi(T) > 0.8$.

Probability of milling, $P(\Lambda(T))$: Similarly, this refers to the frequency of repetitions out of the N_{rep} repetitions where the group mills with $\Lambda(T) > 0.8$.

Variance ratio, R_{var}: Temporal fluctuations in alignment can provide insight into whether one behavioral class exhibits more stable collective motion. We therefore computed the ratio of temporal variances

$$R_{\text{var}} = \frac{\text{Var}[\Phi_{\text{coop}}(t)]}{\text{Var}[\Phi_{\text{def}}(t)]}. \tag{5.8}$$

Values deviating strongly from unity would indicate systematic differences in the stability of alignment between cooperators and defectors.

Spatial segregation: To test whether cooperators and defectors occupy different regions of space, we measured the ratio between the mean pairwise distances of individuals of the same type and those of different types. Similar spatial segregation metrics have been used to quantify mixing patterns in collective systems and animal groups [43].

Defector–group alignment, A: To quantify whether defectors preferentially align with cooperators or move independently of them, we computed the mean directional alignment between defectors and cooperators,

$$A = \langle \hat{\mathbf{v}}_d \cdot \hat{\mathbf{v}}_c \rangle \,, \tag{5.9}$$

where $\hat{\mathbf{v}}_d$ and $\hat{\mathbf{v}}_c$ denote normalized velocity vectors of defectors and cooperators, respectively. This metric measures whether defectors tend to follow the prevailing group direction.

Recruitment through local influence: Finally, to examine whether defectors locally influence neighboring cooperators, we measured whether cooperators within the interac-

tion radius became aligned with the defector's heading in subsequent time steps. Such local alignment responses have been used as indicators of leader–follower interactions in collective behavior studies [162].

All quantities were averaged over individuals and time (except for time-series visualizations) and subsequently averaged across independent simulation realizations.

5.4 Results

5.4.1 Non-monotonic emergence of polarization under probabilistic defection

Figure 5.2 shows representative time series of the polarization $\Phi(t)$ (top row) and angular momentum $\Lambda(t)$ (bottom row) for increasing values of the defection probability P_d. In the absence of defection ($P_d = 0$), the system resides in a stable milling state, characterized by persistently low polarization and large angular momentum. As P_d is increased, milling is progressively weakened, with $\Lambda(t)$ decreasing monotonically across all values of P_d.

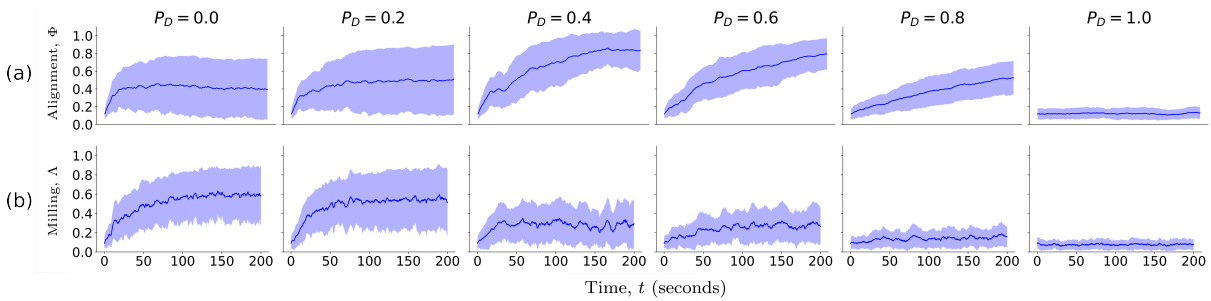


Figure 5.2: Representative time series of polarization $\Phi(t)$ (top row) and angular momentum $\Lambda(t)$ (bottom row) for increasing defection probability P_d for the case when individuals behavior and parameters were $D - Z_{all}, N = 50, R_o = 6$, and $R_a = 12$. Columns correspond to $P_d = 0, 0.2, 0.4, 0.6, 0.8$, and 1.0 (left to right). At $P_d = 0$, the system exhibits stable milling with low polarization and high angular momentum. As P_d increases, milling weakens monotonically, while polarization increases non-monotonically and peaks at intermediate values of P_d (0.4–0.6). For large P_d, both order parameters decay, indicating a transition to a disordered state.

Strikingly, polarization exhibits a non-monotonic dependence on P_d. While $\Phi(t)$ remains small for $P_d \leq 0.2$, it rises sharply for intermediate values of P_d and attains its largest values for $P_d = 0.4$ and 0.6. In this regime, the milling state collapses and is replaced by a strongly polarized collective motion. For larger defection probabilities

($P_d \geq 0.8$), polarization decreases again, and large temporal fluctuations dominate, indicating the loss of both milling and coherent polarization.

Thus, introducing probabilistic defection destabilizes milling monotonically, but induces a non-monotonic emergence of polarization, which is maximal at intermediate defection probabilities. We next examine how this defection-induced reorganization depends on the interaction range, focusing on the role of the attraction radius.

5.4.2 Attraction-range dependence of collective response to defection

To examine how the collective response to probabilistic defection depends on the interaction range, we compute heat maps of the polarization $\Phi(t)$ (top row) and angular momentum $\Lambda(t)$ (bottom row) as functions of defection probability P_d and time for different attraction radii R_a. Figure 5.3 summarizes the collective dynamics across defection probability P_d and time for different values of the attraction radius R_a. Each heat map shows the time evolution of the polarization $\Phi(t)$ (top row) or angular momentum $\Lambda(t)$ (bottom row), averaged over 50 independent realizations. The orientation radius is fixed at $R_o = 6$ for all cases.

For small attraction radii ($R_a = 8$ and 10), milling is absent for all values of P_d, and the dynamics are dominated by polarized motion. In this regime, polarization remains high over time for $P_d \lesssim 0.8$ and collapses only at very large defection probabilities, indicating that probabilistic defection has little effect on collective order unless interactions are almost entirely suppressed.

In contrast, for intermediate attraction radii ($R_a = 12$ and 15), the system exhibits stable milling at low defection probabilities, as evidenced by large angular momentum and negligible polarization. As P_d is increased, milling progressively weakens and disappears, while polarization emerges over a broad time window at intermediate values of P_d. This regime thus displays a clear transition from rotational to translational collective motion driven by probabilistic defection.

Finally, for large attraction radius ($R_a = 20$), milling remains robust across most values of P_d, with angular momentum remaining large and polarization remaining low. Only at very high defection probabilities ($P_d \gtrsim 0.6$) does milling begin to degrade, and

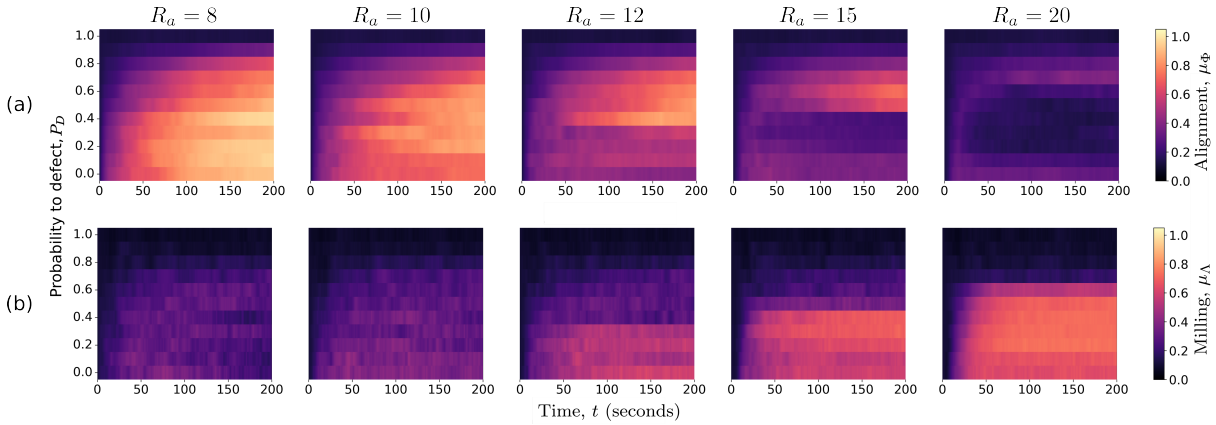


Figure 5.3: Heat maps showing the time evolution of polarization $\Phi(t)$ (top row) and angular momentum $\Lambda(t)$ (bottom row) as functions of defection probability P_d and simulation time for different attraction radii $R_a = 8, 10, 12, 15$, and 20 (left to right). Color intensities represent averages over 50 independent realizations. The orientation radius is fixed at $R_o = 6$. For small R_a, the system is polarized, and milling is absent for all P_d.For intermediate R_a, increasing P_d suppresses milling and induces polarization. For large R_a, milling remains robust, and polarization does not emerge even at high defection probabilities.

even in this limit, polarization does not emerge. Therefore, the impact of probabilistic defection depends strongly on the attraction radius: defection induces polarization only in regimes where milling is present but not overwhelmingly stable.

These results show that probabilistic defection can destabilize milling and promote polarization across a range of attraction radii. We next examine whether this defection-induced transition depends on the number of individuals in the group.

5.5 Robustness of defection-induced transitions across group sizes

To examine whether the defection-induced transition between collective states depends on the number of individuals in the group, we repeated the simulations for several different group sizes while keeping all other model parameters fixed. In particular, we considered groups of size $N = 25, 50, 100$, and 200, and measured the temporal evolution of the polarization $\Phi(t)$ and angular momentum $\Lambda(t)$ across a range of defection probabilities P_D.

Figure 5.4 shows heat maps of $\Phi(t)$ (top row) and $\Lambda(t)$ (bottom row) as functions of the defection probability and simulation time for the different group sizes. Color

intensities represent averages over 50 independent realizations. For most group sizes, the system exhibits qualitatively similar behavior: at low defection probabilities, the group predominantly displays milling, while intermediate values of P_D destabilize the milling state and promote the emergence of polarized motion. At higher defection probabilities, both polarization and milling weaken as interactions become increasingly intermittent.

For the smallest group size ($N = 25$), milling is already weak even in the absence of defection due to the limited number of interacting neighbors. Consequently, the defection-induced suppression of milling is less pronounced in this case. Nevertheless, intermediate values of P_D still promote alignment, consistent with the behavior observed for larger groups.

Overall, although the precise range of P_D values that maximize polarization varies slightly with group size, the qualitative transition from milling to polarization persists across all group sizes considered. These results indicate that the effect of probabilistic defection is not restricted to a specific group size but represents a robust feature of the collective dynamics.

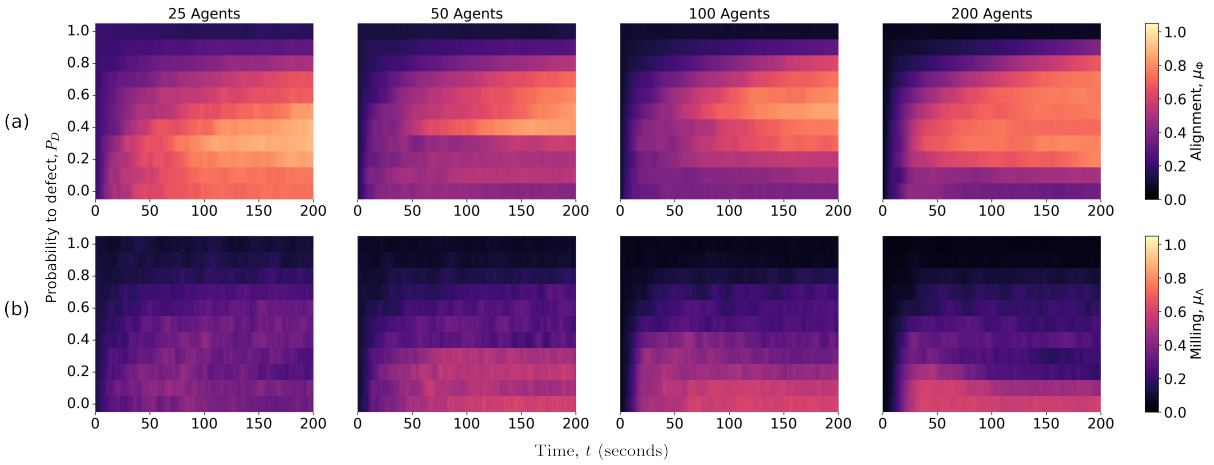


Figure 5.4: Heat maps showing the time evolution of polarization $\Phi(t)$ (top row) and angular momentum $\Lambda(t)$ (bottom row) as functions of defection probability P_D and simulation time for different group sizes $N = 25, 50, 100$, and 200 (left to right). Color intensities represent averages over 50 independent realizations. Interaction parameters are fixed at $R_o = 6$ and $R_a = 12$. Across group sizes, intermediate values of P_D suppress milling and promote the emergence of polarized motion, indicating that the defection-induced transition between collective states persists across different group sizes.

5.5.1 Probabilistic defection reshapes the likelihood of collective states

We quantified the robustness of collective order by measuring the probability that the group reaches a polarized or milling state at a fixed observation time T, averaged over multiple independent realizations. Polarization was defined by the condition $\Phi(T) > 0.8$, while milling was identified using an analogous threshold on the angular momentum order parameter.

When the attraction radius was varied at fixed orientation radius ($R_o = 6$), the probability of polarization exhibited a pronounced non-monotonic dependence on the defection probability P_D (Fig. 5.5a,b). In interaction regimes that support milling at low P_D, polarization emerged most reliably at intermediate defection probabilities, while the probability of milling decreased monotonically with increasing P_D. This confirms that probabilistic defection destabilizes rotational order and promotes alignment, but only within a finite window of defection rates.

In contrast, when attraction interactions were too weak to support milling, polarization occurred with high probability even in the absence of defection, and increasing P_D primarily acted to suppress alignment at very high values. Conversely, for sufficiently large attraction radii where milling was strongly stable, defection reduced milling only at high P_D and failed to induce polarization altogether (Fig. 5.5a,b). These regimes mirror the behavior observed in the heat maps, indicating that defection-induced polarization requires a baseline state that is neither fully aligned nor excessively robust.

A similar pattern was observed when varying the orientation radius at a fixed attraction radius ($R_a = 12$; Fig. 5.5c,d). Intermediate defection probabilities maximized the likelihood of polarization in regimes where milling was otherwise dominant, while both low and high values of P_D favoured either persistent milling or disordered motion, respectively. Notably, the effect of varying the orientation radius was considerably weaker than that of varying the attraction radius, although the non-monotonic dependence of polarization on the defection probability persisted. For $R_o = 8$, the polarization order parameter increased from $\Phi \approx 0.8$ to near-perfect alignment as P_D increased, while milling was progressively suppressed. At higher defection probabilities, polarization declined

again, indicating a weak but detectable non-monotonic dependence on P_D.

For $R_o = 10$, milling was entirely absent, and the group remained strongly polarized over almost the full range of defection probabilities (Fig. 5.5c,d). In this regime, polarization collapsed only at the extreme case $P_D = 1$, where all individuals continuously defect and collective order is lost. Thus, while changes in the orientation radius preserve the non-monotonic influence of defection, they compress it into a narrow high-polarization regime and do not induce qualitative transitions between milling and alignment.

Together, these results demonstrate that probabilistic defection acts as a control parameter that shifts the balance between collective states, but only when the underlying interaction structure places the system near a transition between milling and alignment.

5.5.2 Defection acts as an effective reduction in individual turning rate

The non-monotonic influence of probabilistic defection on collective order suggests that defection modifies an effective individual-level control parameter. In the Couzin model, the maximum turning rate $\omega_{\max}$ governs how rapidly individuals can reorient in response to their neighbours and therefore controls the stability of rotational versus aligned motion. We hypothesized that probabilistic defection effectively reduces this turning ability by intermittently removing interaction-based reorientation.

When an individual defects, it continues moving along its current heading without responding to social forces. As a result, the individual accumulates heading changes more slowly over time than a fully interacting agent, because periods of non-interaction interrupt its capacity to turn. Increasing the defection probability P_D therefore decreases the fraction of time during which an individual can actively reorient, leading to an effective reduction in turning rate even when the intrinsic $\omega_{\max}$ is unchanged.

To establish a correspondence between probabilistic defection and turning dynamics, we analysed the trajectory of an isolated individual under two complementary conditions. In the first case, the individual moved with a constant turning rate ω while $P_D = 0$, and ω was varied from 46 to 10. In the second case, the turning rate was fixed at $\omega = 46$ while the defection probability P_D was varied. Then, we quantified the turning behaviour of an isolated individual. Comparing the resulting trajectories from the two

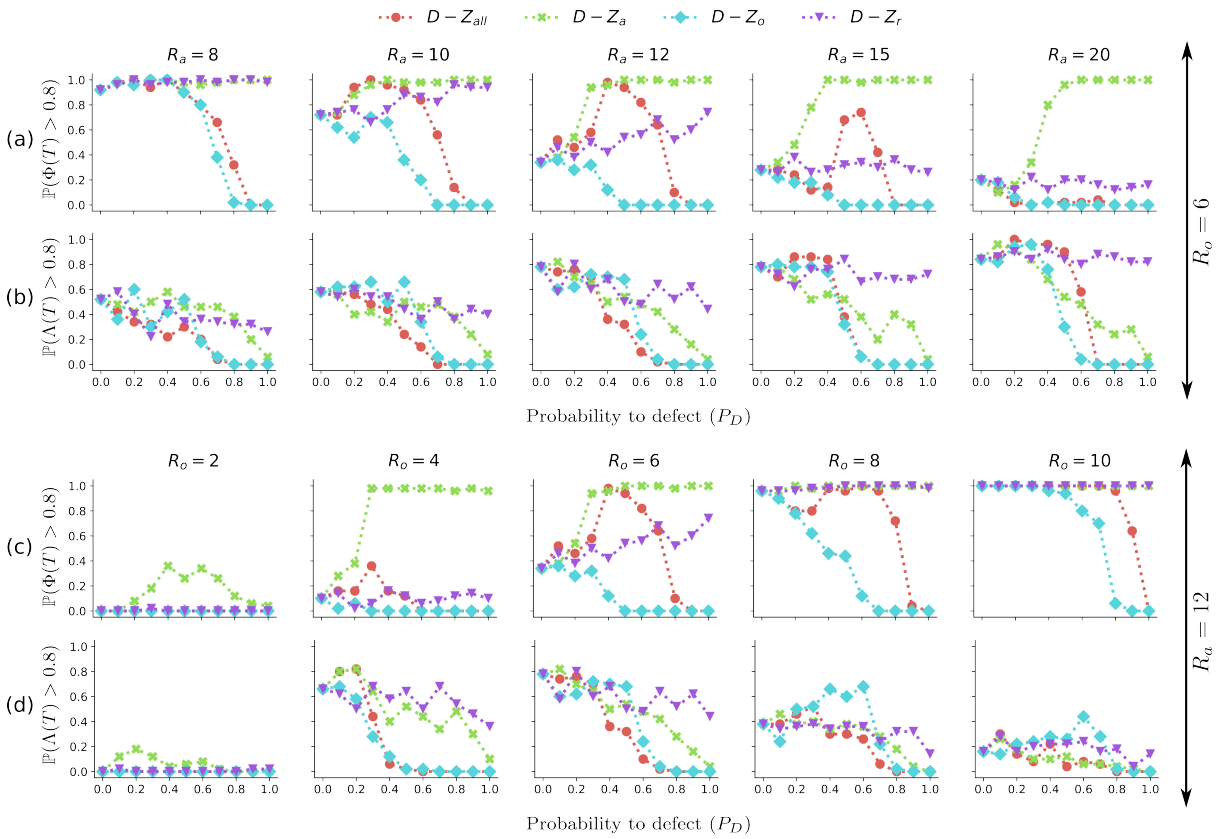


Figure 5.5: Probability of polarization and milling as functions of defection probability. Probability that the group reaches a polarized state, defined as $\mathbb{P}(\Phi(T) > 0.8)$, and the probability that the group reaches milling, defined analogously using the angular momentum order parameter $\mathbb{P}(\Lambda(T) > 0.8)$, shown as functions of the defection probability P_D. **(a,b)** Probability of polarization and milling, respectively, for varying attraction radius R_a at fixed orientation radius $R_o = 6$. **(c,d)** Corresponding probabilities for varying orientation radius R_o at fixed attraction radius $R_a = 12$. Across interaction regimes where milling is present at low P_D, polarization occurs most reliably at intermediate defection probabilities, while milling is progressively suppressed as P_D increases. In contrast, when milling is either absent or excessively stable, defection has little effect on the likelihood of collective alignment.

cases allowed us to relate the turning radii generated by reduced turning rates to those produced by increasing defection probability. As the defection probability P_D increased, the individual executed progressively wider turns, reflected in an increasing *mean turning radius* (Fig. 5.6**a**, see Mats. and Meth.). This trend is identical to the effect of directly reducing the maximum turning rate $\omega_{\max}$, which also produces wider turns (Fig. 5.6**b**). Thus, probabilistic defection and reduced turning ability generate the same qualitative change in individual trajectories.

To establish a quantitative correspondence, we compared the turning radii produced by varying P_D with those obtained by varying ω. Figure 5.6**c** shows the resulting rela-

tionship using the turning radius as a common reference variable. The horizontal axis represents the mean turning radius, while the two vertical axes correspond to the parameters P_D (left axis) and ω (right axis). For defection probabilities below $P_D \approx 0.6$, the turning radii generated by increasing P_D closely match those produced by decreasing ω, indicating a clear correspondence between the two parameters. However, for larger defection probabilities ($P_D \gtrsim 0.6$), this correspondence breaks down, and the turning radii produced by defection no longer match those obtained by reducing the turning rate. We next constructed an explicit mapping between the defection probability P_D and the turning rate ω. For each value of P_D, we identified the value of ω that produced a similar mean turning radius in the single-individual simulations. The correspondence was retained only when the difference between the two turning radii was smaller than a tolerance ϵ. Figure 5.6**d** shows the resulting relationship between P_D and ω. A clear correspondence exists for defection probabilities up to approximately $P_D \approx 0.6$, where increasing defection probability maps onto progressively lower turning rates. Beyond this regime the matching condition is no longer satisfied, indicating that the turning dynamics produced by high defection probabilities cannot be captured by a simple reduction in turning rate.

Finally, we examined whether this effective reduction in individual turning rate provides a mechanistic explanation for the non-monotonic effect of probabilistic defection on collective polarization. We simulated groups with $P_D = 0$ while systematically varying $\omega_{\max}$, *i.e.*, the maximum turning rate allowed in the collective simulation model. The probability of polarization ($\mathbb{P}(\Phi(T) > 0.8)$) exhibited a pronounced maximum at intermediate turning rates (Fig. 5.6**e**). The range of $\omega_{\max}$ values for which polarization was most likely (18–24) corresponds, via the turning-rate mapping, to defection probabilities $P_D \sim 0.4$ (Fig. 5.6**d**). Notably, this is also the range of P_D that maximized polarization in Fig. 5.5. The close correspondence between the polarization time series for mapped pairs of P_D and $\omega_{\max}$ further supports this mechanistic explanation linking probabilistic defection to an effective reduction in individual turning rate.

To further examine this correspondence, we compared the temporal evolution of group polarization under probabilistic defection with that obtained by varying the maximum turning rate in the absence of defection. Figure 5.7 presents representative polarization time series for several values of the defection probability P_D at fixed $\omega_{\max} = 40$, alongside

simulations with $P_D = 0$ in which $\omega_{\max}$ was progressively reduced. For each parameter combination, the trajectories represent the mean polarization across 50 independent realizations, with shaded regions indicating the variability across simulations.

Strikingly, the polarization dynamics produced by intermediate defection probabilities closely resemble those obtained by lowering the intrinsic turning rate. In particular, the case $P_D = 0.4$ with $\omega_{\max} = 40$ exhibits a gradual increase in polarization that closely matches the trajectory observed for $P_D = 0$ with $\omega_{\max} = 24$. In contrast, simulations with lower defection probabilities or higher turning rates rapidly saturate at substantially lower polarization levels ($\Phi \approx 0.4$). This temporal correspondence further supports the interpretation that probabilistic defection effectively reduces the turning rate of individuals and thereby alters the collective dynamics.

5.5.3 Testing alternative explanations for defection-induced polarization

Having established that probabilistic defection effectively reduces individual turning rates, we next asked whether the resulting increase in polarization could instead be explained by more direct mechanisms such as recruitment, leadership, or spatial reorganization.

The increase in polarization at intermediate defection probabilities could, in principle, arise from several alternative mechanisms, including recruitment of cooperators by defectors, persistent leadership by defectors, or changes in spatial organization and neighborhood composition. We tested these alternative hypotheses using a suite of diagnostics that quantify asymmetric influence, temporal ordering, and spatial structure between cooperators and defectors. A summary of these analyses is presented below in Sec. 5.6.

Specifically, we quantified the temporal ordering between defector and cooperator polarization using cross-correlation lags, measured defector–group alignment and local influence on neighboring cooperators, and examined spatial organization through segregation metrics and neighborhood composition statistics. Cross-correlation lags remained near zero, indicating no systematic leading by defectors. Defector–group alignment and local influence were weak and close to null expectations, arguing against persistent leadership or recruitment. Moreover, both the number and fraction of cooperative and defector neighbors varied monotonically with P_D, and showed no non-monotonicity that could

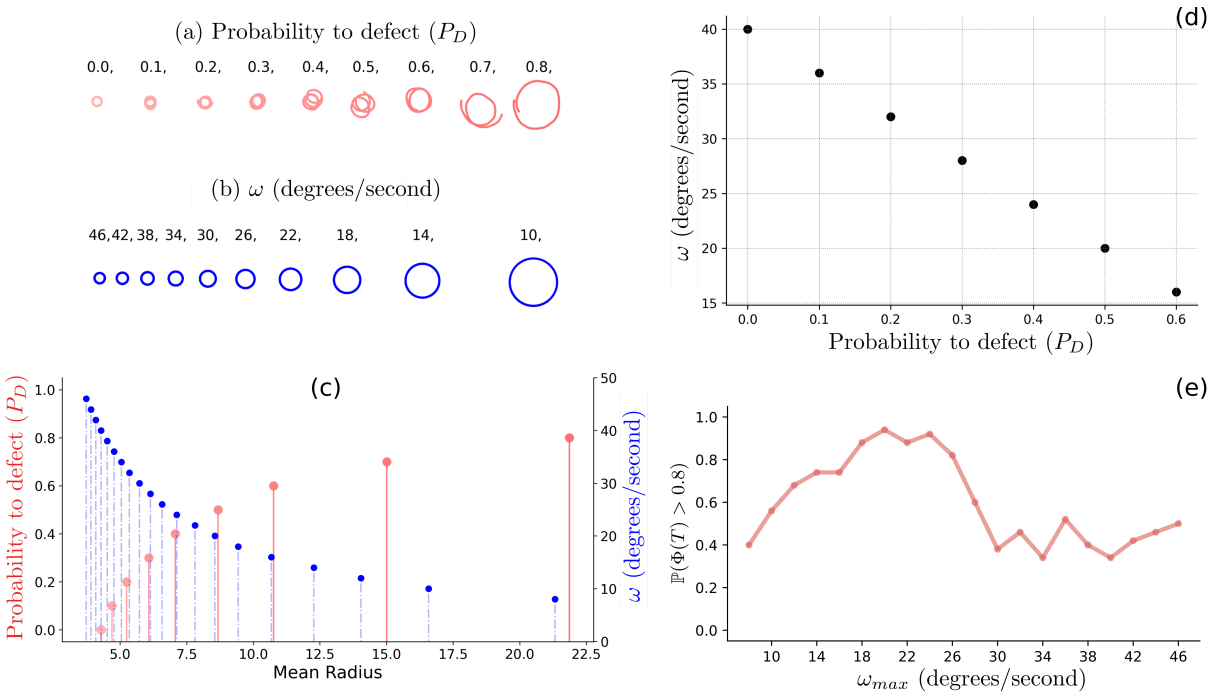


Figure 5.6: Defection promotes polarization by reducing the effective turning rate of individuals. **(a)** Trajectories of an isolated individual obtained by fixing $\omega = 46$ and varying the defection probability P_D. **(b)** Trajectories obtained by varying the constant turning rate ω from 46 to 10 while fixing $P_D = 0$. **(c)** Mean turning radius $\langle R \rangle$ obtained from single-individual trajectories plotted against the parameters P_D (left vertical axis) and ω (right vertical axis). Each point represents the mean turning radius measured from simulated trajectories. For $P_D \lesssim 0.6$, the turning radii produced by increasing defection probability closely match those generated by decreasing turning rate, indicating an approximate mapping between the two parameters. Beyond $P_D \approx 0.6$, this correspondence breaks down. **(d)** Relationship between P_D and ω obtained by matching the mean turning radii from the two sets of single-individual simulations. Points are shown only when the difference between the turning radii is smaller than a tolerance ϵ. A clear correspondence between increasing defection probability and decreasing turning rate exists up to $P_D \approx 0.6$, beyond which the mapping breaks down. **(e)** Probability that the group reaches a polarized state, $\mathbb{P}(\Phi(T) > 0.8)$, as a function of the maximum turning rate $\omega_{\max}$ for simulations with $P_D = 0$. Polarization occurs most frequently at intermediate turning rates. Through the correspondence between turning rate and defection probability established in panels (c–d), this regime corresponds to defection probabilities $0.3 \lesssim P_D \lesssim 0.6$. All measurements were performed for $R_o = 6$ and $R_a = 12$, to match the regime analyzed in Fig. 5.5.

account for the peak in polarization.

Hence, across all diagnostics, we find no evidence that defectors act as persistent leaders, recruit cooperators, or generate polarization through changes in local composition or spatial segregation. These results rule out recruitment-, leadership-, and composition-driven explanations for the observed non-monotonic dependence of polarization on P_D

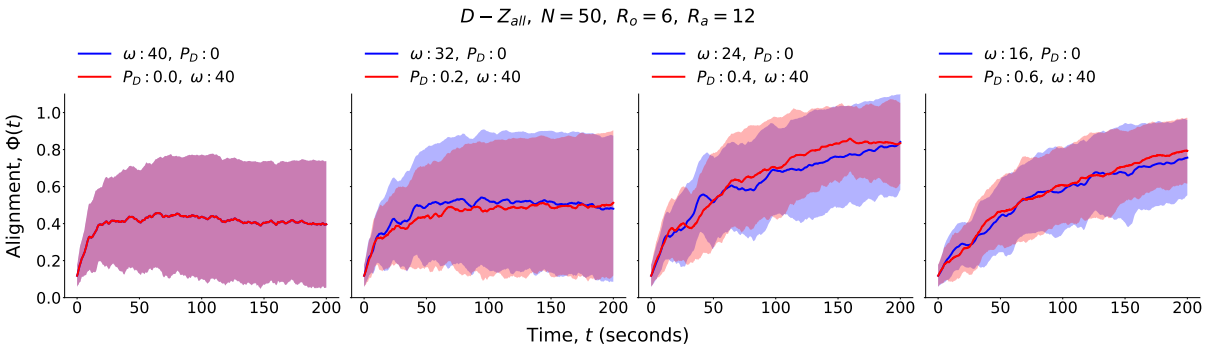


Figure 5.7: Alignment time series averaged over 50 simulations for groups of size $N = 50$, comparing dynamics obtained by varying the defection probability P_D at fixed $\omega_{\max} = 40$ with those obtained by decreasing the maximum turning rate $\omega_{\max}$ in the absence of defection ($P_D = 0$). Across parameter pairs the trajectories closely overlap, indicating that increasing P_D produces collective dynamics similar to those arising from reductions in $\omega_{\max}$.

(see Supplementary Information). Instead, they point to a collective dynamical mechanism, motivating our analysis of how probabilistic defection alters effective individual turning rates as we illustrated in the previous section.

5.5.4 Reversibility of collective states under varying defection probability

Having established that increasing the defection probability can drive the system from a milling to a polarized state, we next asked whether this transition is reversible. In particular, if probabilistic defection acts as a control parameter for collective organization, reducing P_D after polarization has formed should allow the system to return to a milling state.

To examine this, we performed simulations in which the defection probability P_D was gradually increased and subsequently decreased within a single realization. Starting from $P_D = 0$, the defection probability was increased in increments of 0.1 every 200 time steps until $P_D = 1$, and then reduced in the same manner back to $P_D = 0$. Figure 5.8 shows the resulting dynamics, with polarization (Φ) in the top panel and angular momentum (Λ) in the bottom panel.

During the forward sweep ($P_D : 0 \to 1$), the system exhibited the expected transition from milling to polarization. Polarization began to increase around $P_D \approx 0.3$ and reached a maximum near $P_D \approx 0.4$, after which it gradually declined as the defection

probability increased further; whereas milling reduce monotonically after $P_D = 0.2$. In the reverse phase, when the defection probability was reduced ($P_D : 1 \rightarrow 0$), polarization again increased at intermediate values of P_D before decreasing at lower defection probabilities. However, this behaviour was not fully consistent across realizations: in 4 out of 10 simulations the group remained in a highly polarized state during the retreat phase instead of transitioning back to low polarization.

Despite this variability, the milling state was partially recovered during the retreat phase. The angular momentum order parameter Λ (Fig. 5.8, bottom panel) increased again as P_D decreased, indicating a re-emergence of rotational order. However, the recovered milling state was weaker than the initial one, reaching a maximum value of $\Lambda \approx 0.6$, compared with $\Lambda \approx 1$ at the beginning of the simulation. This partial recovery suggests that while probabilistic defection can destabilize milling and promote alignment, reducing defection allows the system to regain rotational structure, highlighting the potential of defection probability as a tunable control parameter for switching between collective states.

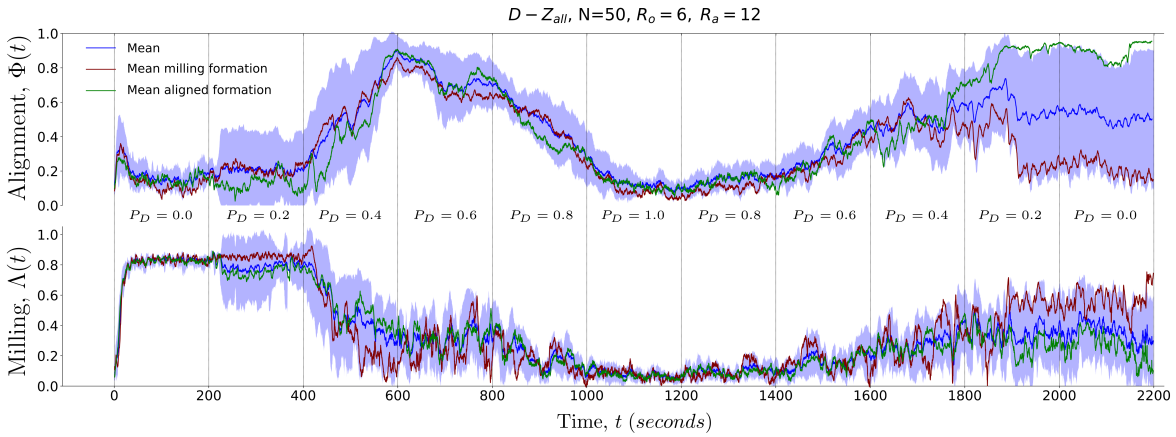


Figure 5.8: Reversibility of collective states under gradual variation of the defection probability P_D. In a single realization, P_D was increased from 0 to 1 in steps of 0.1 every 200 time steps and subsequently reduced back to 0. Top panel shows the polarization order parameter Φ, and the bottom panel shows the angular momentum order parameter Λ.

5.6 Neighborhood statistics and local composition do not explain defection-induced polarization

The non-monotonic dependence of polarization on the defection probability P_D could in principle arise from spatial effects. For example, if defectors move differently from cooperators, they might diffuse outward more rapidly or occupy different regions of space. Such spatial biases could lead to systematic differences in the local composition of neighbors around a focal individual, potentially altering interaction patterns and thereby affecting collective alignment.

One possibility we considered is that occasional defection may allow defectors to move more independently of their neighbors, causing them to disperse more than cooperators. If this were the case, the local proportion of defectors around a focal individual might differ from the global fraction expected from P_D. However, in our model the behavioral identity of individuals switches probabilistically over time, so that cooperators and defectors continually interchange roles. Under this mechanism, one would expect the local fractions of cooperators and defectors to remain close to the null expectation determined by P_D, provided no spatial segregation emerges.

To test whether the non-monotonic dependence of polarization on P_D could be explained by such spatial effects, we quantified both the absolute number of neighbors and the fraction of neighbors of each type as functions of P_D for a range of attraction radii R_a at fixed orientation radius $R_o = 6$ (Figs. 5.9–5.10).

For each value of P_D, the reported quantities were averaged over all individuals and over time, and finally averaged across independent simulation realizations. Specifically, for a quantity $X_i(t)$ measured for individual i at time t, we compute

$$\langle X \rangle = \frac{1}{R} \sum_{r=1}^{R} \left(\frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} X_{i,r}(t) \right), \tag{5.10}$$

where N is the number of individuals, T is the number of time steps considered, and R is the number of independent simulation realizations. Two types of dashed error bars indicate the standard deviation across individuals (solid dashes) and across time (dotted dashes).

Across all interaction ranges, the mean number of cooperative neighbors decreases approximately monotonically with increasing P_D, while the number of defector neighbors increases correspondingly. A similar monotonic trend is observed for neighbor fractions: the fraction of cooperative neighbors decreases nearly linearly with P_D, whereas the fraction of defector neighbors increases.

These trends closely follow the trivial expectation that, as the probability of defection increases, individuals encounter defectors more frequently and cooperators less frequently in their local neighborhoods. Importantly, however, both the number and the fraction of neighbors vary smoothly and monotonically with P_D, without exhibiting any intermediate peak or non-monotonic behaviour.

This contrasts with the collective polarization observed in the main text, which exhibits a pronounced maximum at intermediate defection probabilities. If the polarization transition were driven by changes in neighborhood size, local density, or the mixing ratio between cooperators and defectors, one would expect these quantities to display a corresponding non-monotonic dependence on P_D. The absence of such behavior indicates that the enhanced alignment at intermediate defection probabilities cannot be explained by simple spatial or compositional effects.

As a simple control, when $P_D = 0.5$ the mean number and fraction of cooperative and defector neighbors around a focal individual are approximately equal when averaged over individuals and time (Figs. 5.9–5.10). This balanced composition corresponds to the trivial expectation under random mixing, i.e., in the absence of spatial assortment between the two behavioral states.

Importantly, this value of P_D lies within the regime where polarization is maximal in the collective simulations. Thus, even when cooperators and defectors are locally present in equal proportions, strong collective alignment still emerges. This demonstrates that the enhanced polarization observed at intermediate defection probabilities cannot be explained by changes in local composition or neighborhood structure.

5.7 Diagnostic tests for recruitment and influence of defectors

The emergence of polarization at intermediate defection probabilities raises the possibility that defectors might actively influence or recruit cooperators during collective motion.

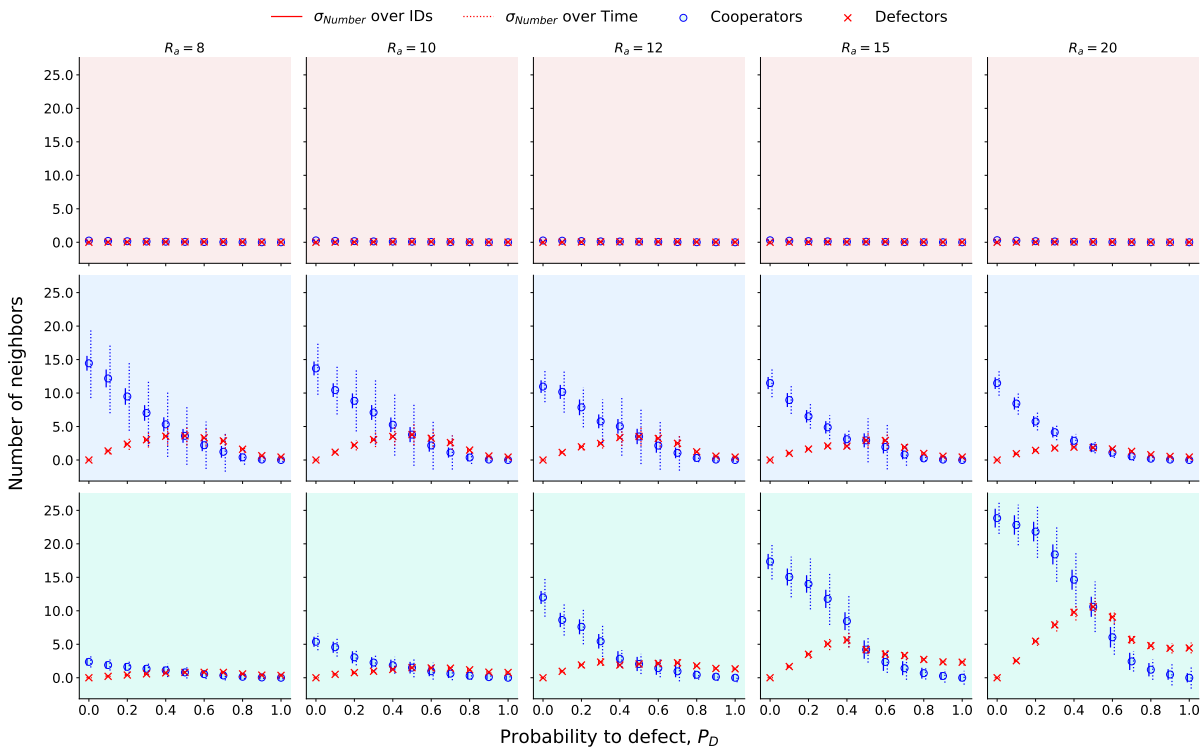


Figure 5.9: Mean number of neighbors in different interaction zones as a function of defection probability. Mean number of cooperative neighbors (blue circles) and defector neighbors (red crosses) as functions of the defection probability P_D for different attraction radii R_a at fixed orientation radius $R_o = 6$. Rows correspond to the zone of repulsion (top), zone of orientation (middle), and zone of attraction (bottom). Values represent averages over all individuals and over time, further averaged across independent simulation realizations. Solid dashed error bars denote the standard deviation across individuals, while dotted dashed error bars denote the standard deviation across time. For all interaction zones and values of R_a, the number of cooperative neighbors decreases monotonically with P_D, while the number of defector neighbors increases monotonically.

For example, defectors might locally bias the headings of neighbouring cooperators, transiently lead alignment events, or segregate spatially in a manner that promotes collective order. To examine these possibilities, we computed a set of diagnostic metrics that quantify spatial mixing, temporal influence, and alignment relationships between cooperators and defectors.

These metrics have been widely used in studies of collective behavior to characterize alignment, segregation, and leader–follower dynamics in moving groups [2, 60, 162].

5.7.1 Null model

To determine whether the observed patterns could arise trivially from the instantaneous mixture of cooperators and defectors, we constructed a null model by randomizing be-

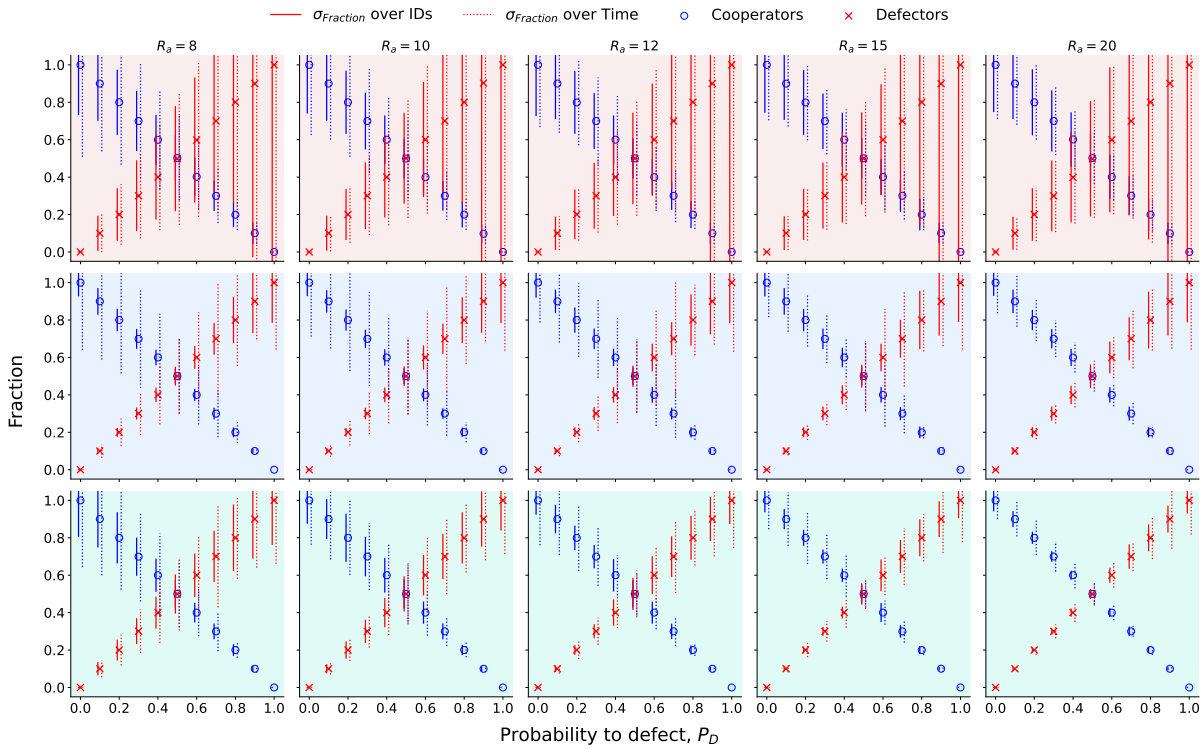


Figure 5.10: Fraction of neighbors in different interaction zones as a function of defection probability. Fraction of cooperative (blue circles) and defector (red crosses) neighbors as functions of the defection probability P_D for different attraction radii R_a at fixed orientation radius $R_o = 6$. Rows correspond to the zone of repulsion (top), zone of orientation (middle), and zone of attraction (bottom). Values represent averages over all individuals and over time, further averaged across independent simulation realizations. Solid dashed error bars denote the standard deviation across individuals, while dotted dashed error bars denote the standard deviation across time. Neighbor fractions vary approximately linearly with P_D and show no non-monotonic behavior.

havioural identities while preserving the spatial configuration and velocities of individuals.

Specifically, at each time step the labels identifying cooperators and defectors were randomly permuted across individuals while keeping their positions and velocity vectors unchanged. This procedure preserves the global fraction of cooperators and defectors as well as the spatial structure of the system, while removing any systematic relationship between behavioral identity and collective dynamics.

For each value of P_D, this randomization procedure was repeated $R = 200$ times to generate a null distribution for each metric. The mean and standard deviation of these null realizations are shown in the figures as dashed lines and shaded bands.

Because spatial positions and velocities are preserved in the null model, the resulting variability reflects only the random reassignment of behavioural identities. Consequently,

the null envelopes are relatively narrow and represent the expected range of metric values under random mixing of cooperators and defectors.

5.7.2 Comparison with null expectations

Observed metrics were compared with the null expectations to determine whether defectors exert a detectable influence on collective dynamics beyond that expected from random mixing. Because these analyses serve as diagnostic controls, we report graphical comparisons between observed values and the variability of the null ensemble, rather than performing formal hypothesis tests.

The diagnostic metrics reveal little evidence that defectors systematically recruit cooperators or act as persistent leaders of collective motion (Fig. 5.11). The mean polarization of defectors (Fig. 5.11**a**) closely follows the behavior of the group as a whole and remains near the null expectation across the range of defection probabilities. Similarly, the variance ratio between cooperator and defector polarization time series remains close to unity (Fig. 5.11**b**), indicating comparable temporal fluctuations in alignment for both behavioral classes.

The defector–group alignment metric (Fig. 5.11 **c**) remains small for low values of P_D, reflecting weak alignment between defectors and the collective direction when defectors are rare. However, this trend remains close to the null expectation, suggesting that defectors do not systematically bias the group direction.

Spatial segregation between cooperators and defectors also remains weak (Fig. 5.11**d**), indicating that the two behavioural states remain well mixed within the group. The cross-correlation lag between cooperator and defector polarization time series (Fig. 5.11**e**) fluctuates around zero, showing no consistent evidence that defectors lead or lag collective alignment dynamics.

Finally, the local influence metric exhibits a modest peak at intermediate defection probabilities (Fig. 5.11**f**), coinciding with the regime where polarization is strongest. However, this increase remains within the variability of the null ensemble, suggesting that it arises from stochastic fluctuations rather than systematic recruitment.

Together, these diagnostics indicate that defectors do not act as persistent leaders or recruiters of cooperators. Instead, the emergence of strong polarization at intermediate

defection probabilities is more consistent with a change in the effective turning dynamics of individuals, as described in the main text.

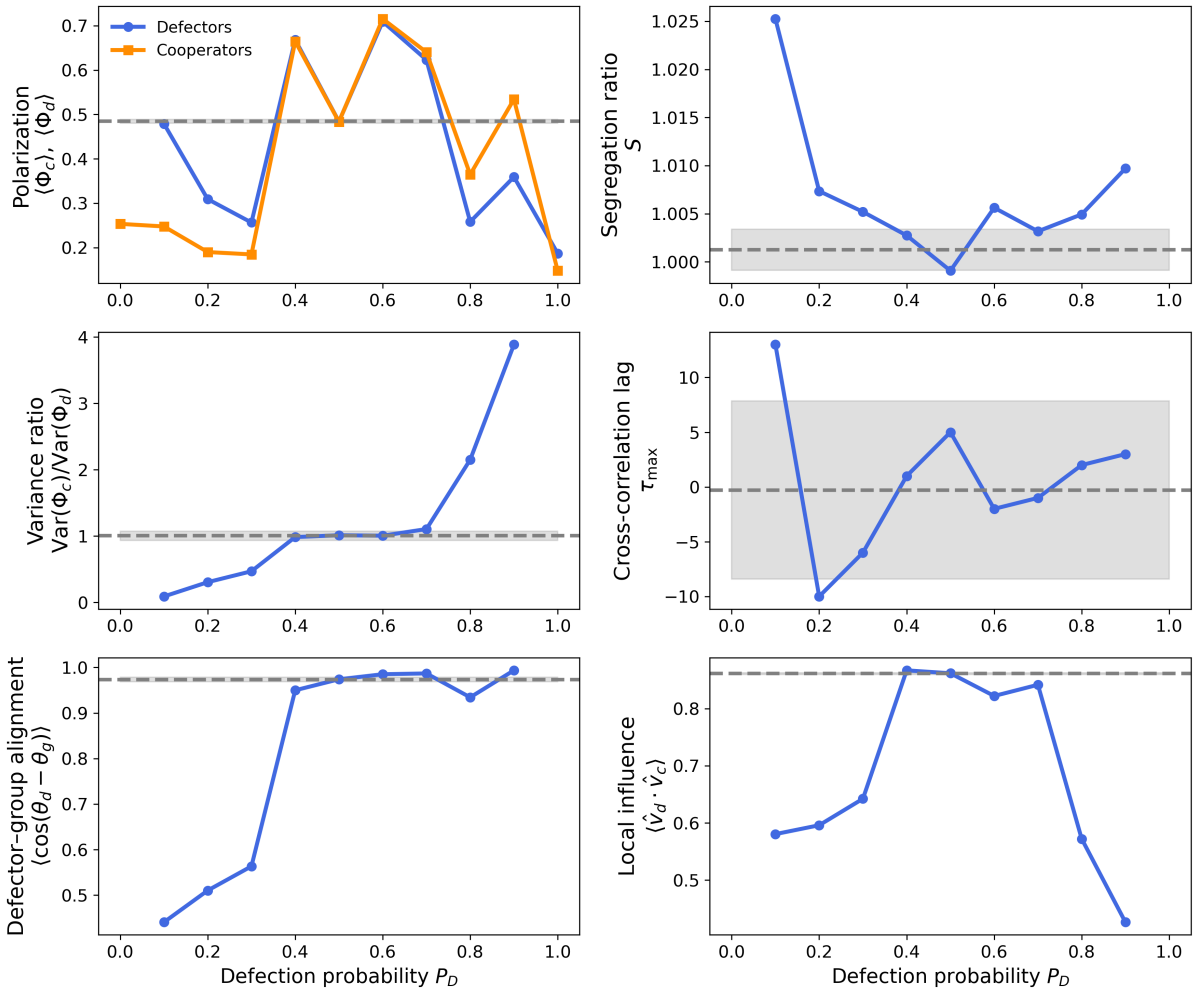


Figure 5.11: Defector influence diagnostics compared against null models. Defector influence diagnostics as functions of the defection probability P_D, compared against null expectations. Solid blue lines show mean values averaged over independent simulation realizations. Grey dashed lines and shaded bands indicate the mean and $\pm 2\sigma$ range obtained from label-shuffled null models that preserve the instantaneous proportions of cooperators and defectors while destroying dynamical correlations. **(a)** Mean polarization of defectors, $\langle\Phi_d\rangle$. **(b)** Spatial segregation ratio S, defined as the ratio of mean intra-type to inter-type distances. **(c)** Variance ratio $\mathrm{Var}(\Phi_c)/\mathrm{Var}(\Phi_d)$, quantifying relative fluctuations in collective alignment. **(d)** Cross-correlation lag $\tau_{\max}$ between defector and cooperator polarization time series; positive values indicate defectors leading. **(e)** Defector–group alignment, $\langle\cos(\theta_d - \theta_g)\rangle$. **(f)** Local influence, measured as the mean alignment between defector and neighboring cooperator velocities. Across all metrics, deviations from null expectations are weak or absent, indicating that defectors do not act as persistent leaders or recruiters. Notably, local influence peaks at intermediate P_D, while defector–group alignment remains low at small P_D, supporting the interpretation that enhanced polarization at intermediate defection rates arises from collective dynamical effects rather than direct recruitment or leadership by defectors.

5.8 Conclusion

Existing flocking models assume that the individuals always move according to the interaction rule. However, animals do not always follow the local interaction rules. They often violate the rules due to selfish behavior. Still, the group collectively manages to remain coherent. We referred to such individuals who stochastically violate interaction rules as defectors. Here, we investigated the impact of individuals who probabilistically switch their interaction strategy on the group motion. We explore the parameter space by varying the interaction radii in the Couzin model. Our simulation results show that the group benefits from defecting individuals as they help the group to transition its phase from the milling state to a highly aligned state. We reasoned that the phase change occurs due to the reduced effective turning rate of defecting individuals.

Since the group can transition its phase alone by modulating the probability to defect, a human operator can simply regulate it while interacting with a robot swarm flocking, as per the proposed model, without having to change the interaction radii of each individual. As each individual acts as a defector randomly, an operator can designate defectors as the flock's anonymous leaders. Furthermore, whether defectors, cooperators, or some combination of both are the primary catalysts of the faster alignment process (as we observed in Chapter. 4) remains unclear.

CHAPTER 6

MULTI-ROBOT SEARCHING WITH LIMITED SENSING RANGE FOR STATIC AND MOBILE INTRUDERS

6.1 Overview

In this chapter, we consider the problem of searching for an intruder in a geometric domain by utilizing multiple search robots. We consider the case in which the domain is a simply connected orthogonal polygon, with edges parallel to the coordinate axes. Each of the k robots has a limited range sensing capability, within a square of fixed size, assumed to be the same range for all robots. This study shows that even for such a simple setup, the problem is NP-hard both for a static and for a mobile intruder. Hence, we turn our attention towards comparing various heuristic algorithms to elucidate their benefits and limitations. We compared three methods, namely methods based on space-filling curves, random search, and cooperative random search by evaluating the trade-off between the number, k, of search robots and the time required for the robots to complete the search by varying the geometrical complexity of the search region.

6.2 Introduction

A central problem for the surveillance and monitoring of a domain is that of devising efficient strategies for a team of searchers to intercept an intruder (a target) within the domain. The problem is well motivated by applications in securing sensitive installations, including military bases, nuclear power plants, border patrol, natural habitat conservation, etc. The use of robots is ubiquitous for these types of purposes. This class of problems has been investigated extensively in the literature over the decades [87, 169–174]. Moreover, many practical algorithms have been developed that are used in real-life scenarios [175, 176].

With recent developments in sensing and electronic technology, the deployment of multi-robot systems is becoming increasingly commonplace. However, the sensing capa-

bilities are still limited. The limited sensing range raises critical concerns, particularly in the context of intruder interception. Therefore, it is evident that local geometric properties, such as the shape and structure of a region, can be used to calibrate the effectiveness of the search process.

Motivated by this, we consider the following problem in this work: considering a closed search space, we model the region as a simply connected orthogonal polygon in $\mathbb{R}^2$. The objective is to deploy a set of k robots with limited but uniform sensing capacities, such that if there is a stationary or mobile intruder inside the region, then the robots will be guaranteed to intercept the intruder as quickly as possible. Each of the k robots can be deployed to an arbitrary location in the domain; this location serves as the starting point for a path that the robot will execute. Our goal is to compute these starting points and the corresponding paths so that the length of the longest path is minimized while assuring that the paths intercept the intruder within ϕ-steps, thereby leading to the discovery of the intruder. Clearly, if one deploys a very large number of robots, enough so that the sensing regions of the robots at their deployment locations completely cover the domain, then any intruder is detected instantaneously. However, in a more realistic scenario, one has a limited number of robots, and therefore, the robots must be moved strategically in order to detect an intruder, so the objective is to minimize the "makespan" (maximum length of the robot paths) so that the search is concluded as soon as possible without having to search the entire domain (in most of the cases) while using only the k robots that are available to the user. We seek to determine the trade-off between the number, k, of robots provided and the total time, T, it takes to execute the k search paths in order to assure the interception of the intruder and determine how the trade-off changes with the considered geometric properties of the orthogonal search area.

6.3 Related Work

The problem of placing a small number of static sensors within a geometric domain in order to achieve visibility coverage (in which a point is "covered" by the sensors if it is within line of sight of any sensor, regardless of the distance to the sensor) is commonly called an "art gallery problem" or a "guarding problem"; this class of problems has been extensively studied, particularly in the field of computational geometry, e.g., [177–180]. Most variants of this class of geometric set cover problems that are intractable (NP-hard

and often also hard to approximate). The usual model assumes no limit on the range of sensing, a constraint we impose in our setting in this work. In the case that there is a mobile "guard", a single guard is enough to cover a connected domain, and the *watchman route problem* [181–183] seeks to minimize the length of a route in order that all points of the domain are seen by the guard. (In the k-watchman route problem, one seeks to optimize the routes of k mobile guards, either to minimize the longest of the routes or to minimize the sum of the route lengths.)

Another related problem that is well studied is the *cops-robber* game on graphs, where a robber is stationed at different vertices of an input graph, and the objective is to determine the minimum number of cops that is sufficient to catch the robber. Both cops and robbers move around the graph, and the goal of the cops is to intercept the robber. The problem is NP-hard on general graphs, and there exists an approximation scheme [184]. However, it is not known if the problem remains hard on special graph families, e.g., grid graphs within a simply connected domain (see also [185–188]). In a continuous domain, another problem related to *cops-robber* is the class of *pursuit evasion* problems, including the visibility-based pursuit-evasion problem, for which bounds have been studied on the number of pursuers needed in order to ensure capture of a mobile intruder [173, 189]. Similar work that focuses on mobile intruders (see [190]) provides scalable solutions. However, the performance of such solutions can not be guaranteed when the intruder's motion model is unknown or when the intruder is adversarial. This raises the need to develop robust multi-agent searching algorithms for worst-case and unknown target motion models.

6.4 Problem Formulation

Consider a simply connected orthogonal polygon $\mathcal{P}$ as shown in Fig. 6.1(a) having integral coordinates. The polygon $\mathcal{P}$ can be considered to be the union of a finite number of unit square pixels (together considered as a grid-graph, $G(\mathcal{P})$) (black squares in Fig. 6.1(b)). Let $\mathcal{I}$ be an intruder (red dot in Fig. 6.1(a)) that is present in the region $\mathcal{P}$. Let $\mathcal{R} = \{r_1, \ldots, r_k\}$, for some $k \in \mathbb{N}$, be a set of robots (green dots in Fig. 6.1(a)), each of which is modelled as a point mass and deployed to search the intruder, $\mathcal{I}$ in $\mathcal{P}$. Each robot $r_i \in \mathcal{R}$ moves from one-pixel center to another, at each integral time step t moving one unit vertically or horizontally (along the edges of the dual graph $\mathcal{D}_G$ (red dotted lines

in Fig. 6.1(b)) of $G(\mathcal{P})$). We consider the sensing range of a robot to be 1/2 in the L_∞ metric so that when a robot is at the center of a pixel, it senses the unit area of the square pixel, detecting and capturing the intruder if it is within the pixel as depicted in Fig. 6.1(a) with the red square region around the robots. Due to the limited sensing range in this formulation, multiple robots can be on the same pixel at any given instance of time. In what follows, we define the problems formally.

Problem 1. *Multi-Robot Search (MRS): Given the starting locations of k robots in a simply connected orthogonal polygonal region $\mathcal{P}$ having n vertices, the goal is to intercept an intruder I in ϕ steps, where I is stationed at any unknown point in $\mathcal{P}$.*

Problem 2. *Dynamic Multi-Robot Search (DMRS): Given the starting locations of k robots in a simply connected orthogonal polygonal region $\mathcal{P}$ having n vertices, the goal is to decide if it is possible to intercept an intruder I in ϕ steps, where I is initially stationed at any unknown point in $\mathcal{P}$ and is moving inside $\mathcal{P}$.*

We assume that the intruder has perfect knowledge of where every robot is at every moment (the intruder has no limit on its sensing range). While the intruder is assumed to have positional knowledge of the robots, it is assumed not to have the power to learn the movement patterns of robots in a manner that would allow it to predict future locations and trajectories of robots.

Given the search scenario, it is NP-hard to determine an algorithm that can guarantee the solution to *Problem 1* since the length of the paths can increase exponentially with respect to the number of cells. Further, in the case of *Problem 2*, a similar argument can be extended by considering the emergence and disappearance of the cells based on the ones that are not yet searched and the ones that are previously searched (such searched cells can re-emerge as recontaminated cell as the time progresses), respectively. Thus, in the following section, we describe the NP-hardness nature of the problem formally by reducing the *Problem 1* and *Problem 2* to the previously known NP-hard, 3-*Partition*, problem [191].

6.5 Hardness Results

Let $\mathcal{P}$ be a simple orthogonal polygon, and $G_\mathcal{P}$ be the underlying grid graph on $\mathcal{P}$. We are given k robots that move with the same constant speed, i.e., at each time step, each

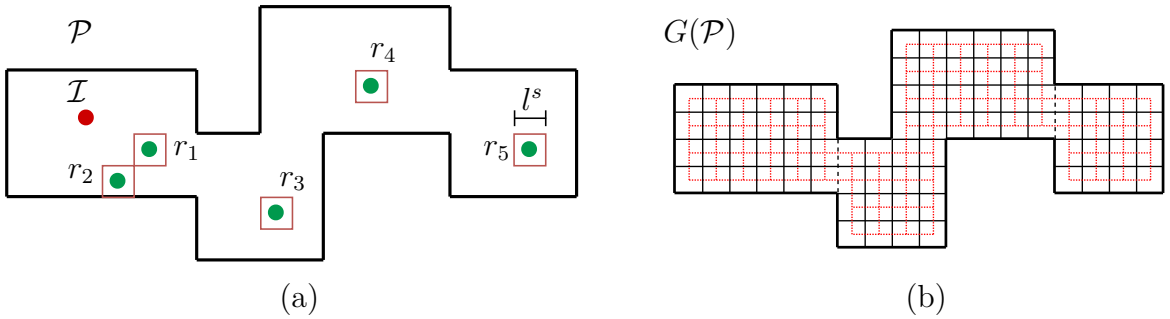


Figure 6.1: (a) Polygonal region $\mathcal{P}$ with intruder ($\mathcal{I}$) in red color and the search robots ($\mathcal{R}' \subseteq \mathcal{R}$) in green. The red square boxes around the robots is the robot sensing range, l^s. (b) The grid graph $G(\mathcal{P})$ shown in black solid lines of the region. The dual of $G(\mathcal{P})$, $\mathcal{D}_G$, is shown with red dotted lines. The robots move (horizontally/vertically) from the center of each black pixel (square) to the center of another at each time step t.

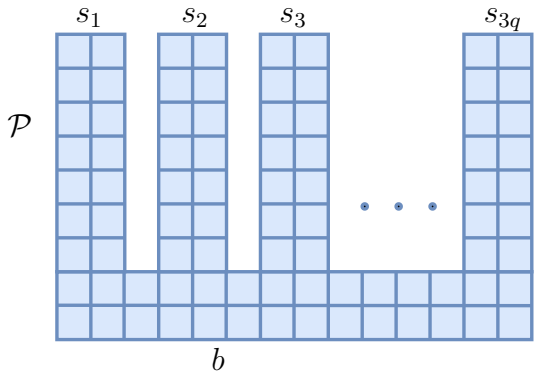


Figure 6.2: Comb structure for the polygon $\mathcal{P}$. The base b, of the structure has a diameter $O(q)$. The spikes $s_0, s_1, \ldots, s_{3q}$ represent sub-regions that are connected to the base. Each spike has a width that is equal to the shortest edge of the polygon that appears in that spike.

robot can move from the current vertex to any of its four neighbouring vertices. There is an intruder $\mathcal{I}$ that is present on a vertex of $G_{\mathcal{P}}$. The objective is to discover $\mathcal{I}$ within a minimum number of steps. In what follows, we show that the *MRS* is NP-hard, by showing a reduction from the 3-*Partition* problem.

Theorem 2. *MRS on a grid graph is NP-hard.*

Proof. Let $\mathcal{F}$ be an instance of the 3-*Partition* problem, where the input is a multiset of positive integers, S, with $|S| = 3q$. The objective is to find if there exists a partition of S into q triplets $S_1, S_2, \ldots, S_q$ such that the sum of the numbers in each one is equal to T. Hence, the sum of all triplets is qT. Moreover, the $S_1, S_2, \ldots, S_q$ must form a disjoint partition of S, and their union result in the whole set S. This 3-*Partition* problem is known to be strongly NP-Complete.

Now, consider an instance $\mathcal{W}$ of *MRS*. The orthogonal polygon $\mathcal{P}$ can be mapped onto a comb-like structure (see Fig. 6.2) that consists of a base and $\beta = 3q$ number of spikes. The base b is a rectangle with the diameter $O(q)$. Let each spike s_i be represented by a distinct number n_i that corresponds to the time taken by k robots to clear the i^{th} spike. The objective is to determine the triplets ($n_i + n_j + n_k = T$) of such spikes so that the total time taken by k robots to clear all the triplets is qT.

In the forward direction, we show that if there is a solution to the 3-*Partition* problem, then we can obtain a solution for the *MRS* problem in polynomial time. Consider a triplet $S_i = \{n_x, n_y, n_z\}$ from the solution of 3-*Partition*. We assign the k robots to traverse the spikes $\{s_x, s_y, s_z\}$ corresponding to $\{n_x, n_y, n_z\}$, respectively. The traversal follows the order of the spikes from left to right. For each triplet, we follow the same procedure. Since the sum of $\{n_x, n_y, n_z\}$ is T, it will take qT time for a robot to clear the corresponding spikes. We shall obtain the solution to the MRS problem in qT steps as searching is done simultaneously with traversal in the spikes.

Conversely, let there exists a solution of *MRS* that can be obtained in polynomial time, which means that the robots can be assigned to exactly three spikes such that the sum of their search time is qT. We can consider the corresponding numbers as a triplet and report all such triplets as a solution to the 3-*Partition* problem. This concludes the proof. □

Remark In the proof, we did not explicitly prohibit multiple robots from being on the same location at any particular time. However, it is easy to adjust the proof even for the case, when it is not allowed to have two robots at the same location at the same time. For each robot which is assigned to clear a spike (say, s_i), we enforce that the robot has to wait for n_i amount of time.

Remark We conjecture that the dynamic version (DMRS), with a mobile intruder, is also NP-hard, perhaps using a similar reduction.

6.6 Algorithmic Methods

Given the hardness status of the problem, in this section, we propose three algorithmic approaches, namely, (i) Space-filling curves, (ii) Random search, and (iii) Cooperative

random search, for robots to search a static/mobile intruder in $\mathcal{P}$.

6.6.1 Space-filling-curves approach

A space-filling curve $\mathcal{C}$ is a mathematical curve on grid graphs spanning a simply connected closed polygonal region without intersecting or repeating itself. It is not known how to compute a space-filling curve for any simply connected non-convex orthogonal polygonal region in $\mathbb{R}^2$ [192]. Hence, in this approach, we rectangulate the input polygon $\mathcal{P}$ to decompose it into disjoint rectangles $d_1, \ldots, d_q$, for some $q \in \mathbb{N}$. Subsequently, we compute the space-filling curve (whose edges form a subset of the edge set of the dual-graph, $\mathcal{D}_G$) for each of the rectangles by utilizing the grid graph of the rectangular sub-region. Then, we distribute robots proportionately to the area of the rectangles. Each robot is assigned to a pre-computed fraction (the robot moves from its start position until the start position of its neighbouring adjacent robot) of the space-filling curve. Finally, the robots simultaneously search in all the rectangles by moving along the assigned paths.

Rectangulation A simple method for rectangulation is to shoot vertical and horizontal rays from each boundary point of an orthogonal polygon. However, this process can create a quadratic number of cells, which would require at least quadratic many robots to continue with a parallel search. Therefore, we use a different method for the rectangulation.

We assume that $\mathcal{P}$ is provided as a list of vertices along the boundary. Consider a bounding rectangle $\mathcal{B}$ that contains $\mathcal{P}$. We overlay a grid of side length l^b on $\mathcal{B}$, where l^b is the shortest edge of the boundary of $\mathcal{P}$. We build a grid graph, $\mathcal{T}$, by considering the grid cells that are inside the polygon $\mathcal{P}$. In order to select grids that lie inside $\mathcal{P}$, for each grid cell, we shoot vertical and horizontal rays in all four directions from the center of the cell. We count the number of times each line intersects the boundary of $\mathcal{P}$. If the count is odd for all four lines, then the cell lies inside the polygon. Further, we sample random grids from the grid graph and traverse the neighbour grids in all four directions to extract the rectangle with the largest area possible. This results in a set of disjoint rectangles $d_1, \ldots, d_q$, such that the union of d_i's forms the polygon $\mathcal{P}$, see Fig. 6.3(a).

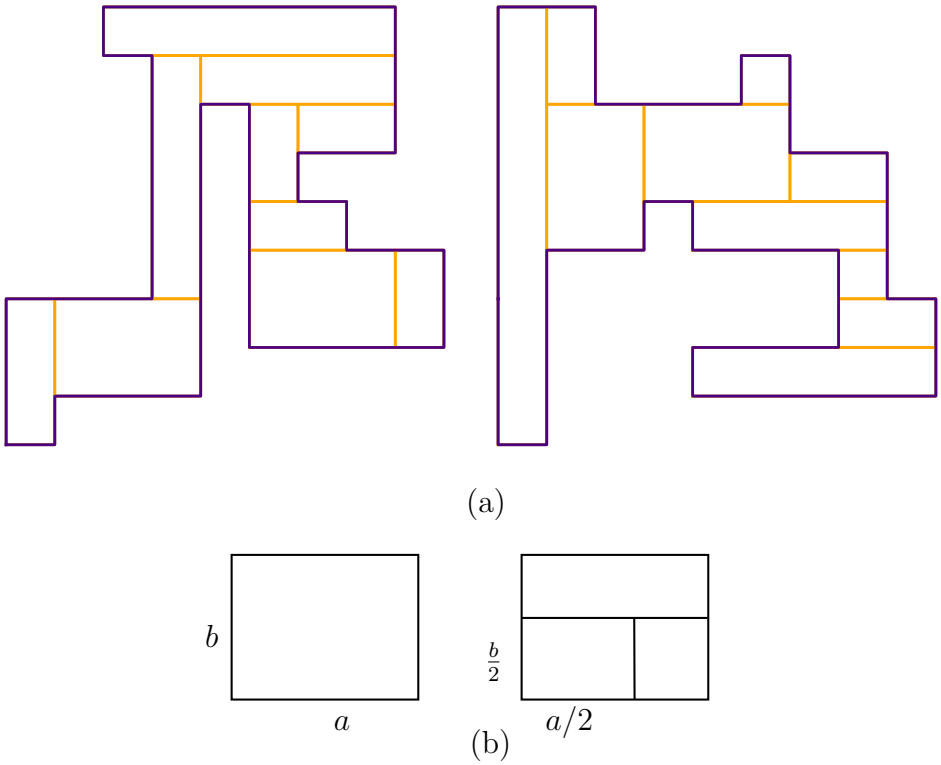


Figure 6.3: (a) Rectangular decomposition of simply connected rectilinear polygons, (b) Splitting of rectangle for computing space-filling curves

Computing Space-filling curves We utilize the extension [193] of [194] to obtain the generalized rectangular space-filling curves for rectangles that has a side of length which is not a power of two. In this approach, a rectangle (with horizontal length: a and vertical length: b) is spilt into three regions 'up', 'right', and 'down', recursively until a trivial path is produced; see Fig. 6.3(b).

Let the sensing range of the robot be $l^s/2$. For each rectangular region d_i, we obtain the grid graph $G_i = G(s_i)$ using smaller square grids g_i of dimensions double the size of the robot's sensing range l^s. Then, we obtain the dual graph $\mathcal{D}_{G_i}$ of G_i, where the dual graph is obtained by connecting the centroids (nodes of $\mathcal{D}_{G_i}$) of each g_i to its neighbouring grids' centroids through an edge (edge of $\mathcal{D}_{G_i}$). Using the nodes of $\mathcal{D}_{G_i}$, we obtain the modified Hilbert's rectangular space-filling curves [193] c_i, for each sub-region d_i. When two sub-regions d_i and d_j are adjacent to each other, they share a common edge. We define these common edges as junctions J_{ij}, as shown by dashed black lines in Fig. 6.1(b). The curve $\mathcal{C}$ ensures complete spanning of all the regions' d_i in $\mathcal{P}$. However, each c_i is disconnected at J_{ij} (see Fig. 6.4). The robots utilize these disjoint curves to move along it and search for the intruder while staying inside the same d_i for the entire time.

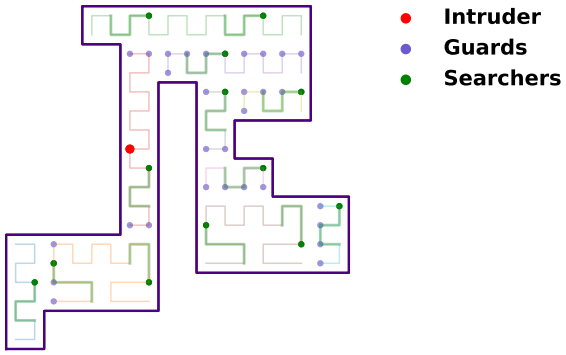


Figure 6.4: Illustration for the arrangement and distribution of search robots in Space-filling curve approach. The polygonal region is decomposed into many sub-rectangular regions (d_i) and each sub-rectangular region is filled with a space-filling curve c_i. The trail behind each agent shows the trajectory during its past few time steps.

Assigning robots and search process The robots are divided into two groups: search robots and guard robots of size k_s (green dots in Fig. 6.4) and k_g (purple dots in Fig. 6.4), respectively. When $k_g = 0$, we define this case as a space-filling curve approach (SFC) whereas if $k_g \neq 0$, then we define it as a Guarded space-filling curve approach (SFC-G). The number of robots assigned to each rectangular area, d_i, is equal to the proportion of the area of the sub-rectangular with respect to the total area of the polygon. The search robots keep patrolling the assigned path, c_i. However, the guard robots remain stationary at the junction J_{ij} to prevent re-contamination. For a detailed illustration of the arrangement, see Fig. 6.4.

Consider the grid graph $G(\mathcal{P}) = \bigcup G(d_p)$, that contains a total of $n \times m$ nodes (line 1 in Alg. 1). In order to determine the time required to search the static/moving intruder using k robots, where $k = |\mathcal{R}'|$, we determine the number of steps that the robots are required to move to find the intruder. We start with a total number of robots that is equal to the total number of grids (nm) in the polygon and continue until k becomes equal to the sum of the number of d_i regions and the number of guarding robots, i.e. $k = |\{d_i | d_i \epsilon \mathcal{P}\}| + k_g$. The simulation finishes if the searcher and the intruder share the same cell. The pseudo-code in Alg. 1 (line 5-8) shows the logical execution of each robot's task.

Runtime: If there are k robots. The algorithm for space-filling curves approach has a worst-case time complexity $\mathcal{O}(k)$.

Algorithm 1 Space filling approach

```
1: grids ← n × m nodes
2: search_state ← False
3: while not search_state do
4:     for r in robots do
5:         if 'r' location = intruder's location then
6:             search_state ← True
7:         end if
8:         Move the robot 'r' along SFC
9:     end for
10:    if search_state then
11:        store search number with the time taken
12:    end if
13: end while
```

6.6.2 Random search approach (RS)

The space-filling approach is partially deterministic in nature. If either end of c_i is revealed to the intruder, it may successfully avoid the searchers. Therefore, we propose a randomized approach, the random search algorithm (RS). Here, k robots are distributed randomly uniformly in the region $\mathcal{P}$. Each robot r_i randomly uniformly samples one target grid cell (say, g_t) inside $\mathcal{P}$, and plans a path (x_i) using A* algorithm [195] connecting its current position to g_t.

The random search approach has an added cost due to the possibility of repeated visits to previously explored cells. For instance, when two larger spaces are connected through a narrow channel, robots will search the narrow region multiple times more often. To avoid this, we redesign the heuristic function of the A* algorithm by accommodating the retraversal cost during the path planning. Whenever a robot visits a grid, it raises its cost (L_{g_i}) by 0.05. So, any robot(s) that plans a path in the future will avoid travelling through the previously visited grids, reducing the frequency of revisits. Such paths are obtained by minimizing the path length to the target grid and the cost (the total sum, ΣL_{g_i}) of revisiting the grids along the path. If a robot reaches its target grid, irrespective of other robots' status, it re-samples a new location to continue searching for the intruder.

In this approach, since the searchers are not constrained to move inside sub-regions of $\mathcal{P}$, guarding robots are not required. Thus reducing the extensive need for robots, and hence, the algorithm works for any number of robots. The pseudo-code in Alg. 2 shows

Algorithm 2 RS and CRS

```
1: search_state ← False, grids ← n × m nodes
2: CRS ← True
3: costs ← array of cost of traversing through a grid
4: while not search_state do
5:     Update costs of each grid
6:     for r in robots do
7:         if r reached its destination then
8:             Sample new location uniformly randomly.
9:             Plan path
10:        end if
11:        if all robots at destination and CRS == True then
12:            Allocate new targets via Hungarian algorithm
13:            Plan path
14:        end if
15:        Update 'r' location along its path
16:    end for
17:    if intruder is found then
18:        search_state ← True
19:    end if
20: end while
```

the logical execution of the random search algorithm.

Runtime: Since in the worst case, each robot can have a path length equal to the total number of grids in the grid graph. The A* itself will have a time complexity of $\mathcal{O}(nm)$. Therefore, if there are k robots, this approach has a worst-case time complexity of $\mathcal{O}(knm)$.

6.6.3 Cooperative random search approach (CRS)

In this approach, the search process is the same as RS however, the robots undertake target-matching once all the robots reach their respective destinations. The robots use the Hungarian task matching algorithm [196] to assign a target among from the newly generated set of target locations. The assignment reduces the total path length of the agent paths.

Say, there are k robots that sample uniformly randomly k grid cells. If any robot discovers the intruder along its path then the algorithm terminates. Otherwise, after all the robots reach their random targets, they sample the next k grid cells uniformly randomly. In Alg. 2, see the lines 11–14 that include the cooperative aspect to the RS

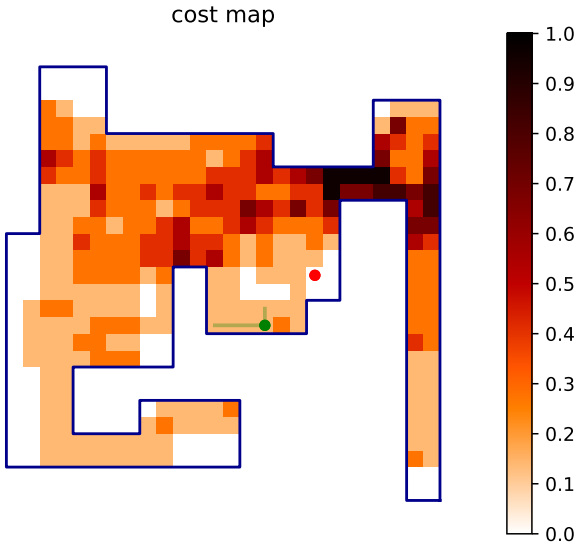


Figure 6.5: A snapshot of the cost map from the random search algorithm. The intruder and the robot are depicted in red and green circles, respectively. The trajectory of the searcher is shown with a green trail and the intruder with a red trail. The grid colours indicate the cost of re-traversal. The higher (brighter) the intensity, the lower the cost.

method.

Runtime: If there are k robots, the Hungarian algorithm has a worst-case time complexity of $\mathcal{O}(k^3)$. However, to minimize the path length, robots first plan paths using A* which has a worst-case time complexity of $\mathcal{O}(k^2nm)$. Therefore, the worst-case time complexity of this approach becomes $\mathcal{O}(k^3 + k^2nm)$. Since we assume that $k \ll nm$, the worst-case time complexity for the cooperative random search approach is $\mathcal{O}(k^2nm)$.

6.7 Simulation Results

In this Section, we evaluate the performance of the proposed methods through Monte-Carlo simulations. We analyze the change in the intruder discovery time by increasing the number of searchers and changing the geometric property by regulating (a) the number of spikes, β, (b) the area and (c) the shape of the polygon $\mathcal{P}$. We assume that the search robots have a limited sensing range and do not have apriori information (position) about the intruder. Also, we model the intruder to move randomly in the domain, $\mathcal{P}$. The codes for all three methods can be found at [197].

6.7.1 Simulation setup

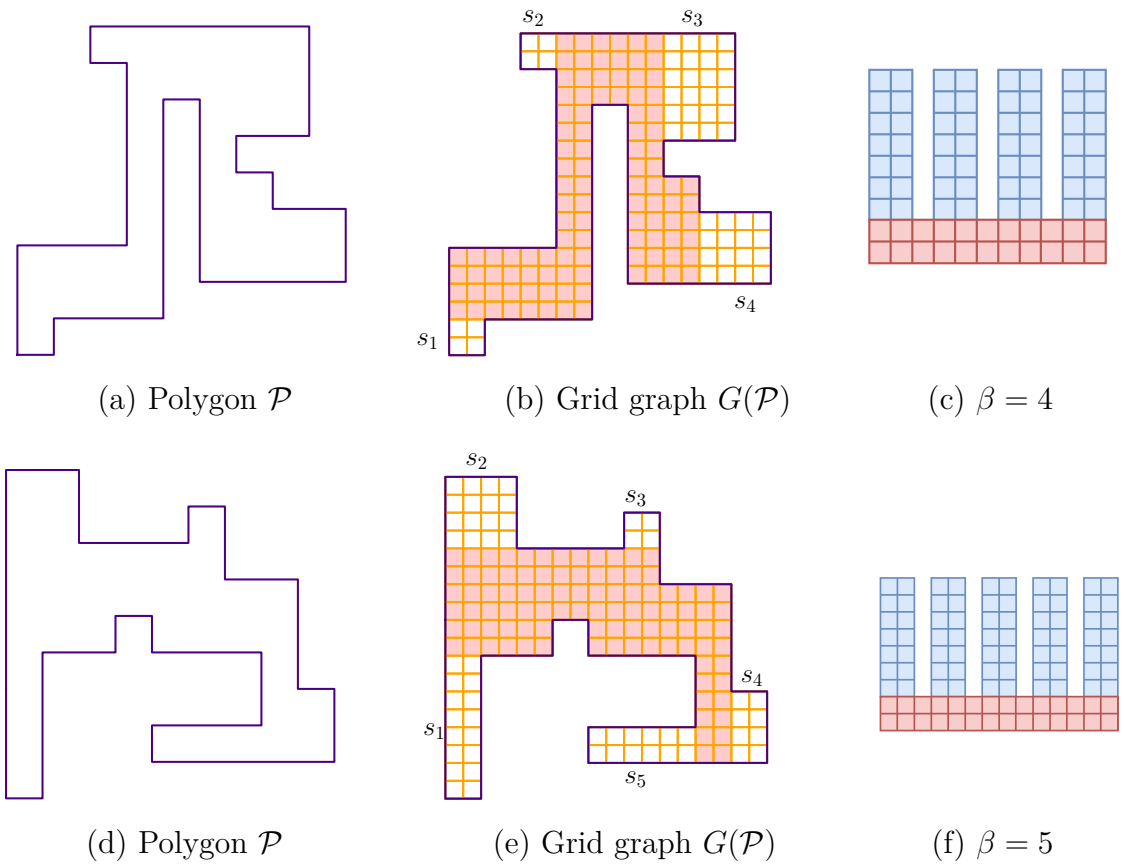
The comb structure defined above in Sec. 6.5 spans the whole space of simply connected orthogonal polygons via the number of spikes, β. Therefore, we model the environment's complexity in terms of β; the higher β, the more complex the environment. The orthogonal polygons of equal area are obtained through the polynomial time Inflate-Cut algorithm [198], as shown in Fig. 6.6(a), & (d). The polygon $\mathcal{P}$ has minimum edge length varied from $l^b = 5m$ to $l^b = 30m$ across various experiments, and each robot has a fixed sensing range of $l^s/2 = 2.5m$. We discretize the domain into cells. Each cell is of $l^s = 5m$ length. The robots move from the centroid of one cell to another in a unit time step. To our knowledge, no known methods exist to determine or regulate the number of spikes in a given simple orthogonal polygon. Therefore, we limit our exploration to manually counting the number of spikes for any random simple orthogonal polygon through visual inspection.

Counting spikes, β We classify a part of the polygon as a spike only if two adjacent edges of a spike form a subset of the polygon's boundary edge set and any of the third edge of the spike must be a fractional part or complete edge of the next adjacent edge in the boundary edge set. Whereas the fourth side of the spike must be open (see Fig. 6.6).

The performance of the proposed heuristics is validated on two types of scenarios - static intruder (S-I) and moving intruder (M-I). The performance of the algorithms can vary based on the number of search robots, the number of spikes β, and the shape of the polygon with an equal number of spikes. Therefore, we perform Monte-Carlo simulations (100 simulations) for each case while varying the number of searching robots, k, the number of spikes β, and the shape of the polygon for each scenario (S-I and M-I) to compare the average performance of each approach.

6.7.2 Baseline

We compare the performance of the proposed algorithms with a baseline approach where the searchers know the position of the intruder at every instant of time, and they compute the shortest path (Dijkstra algorithm [199]) to intercept the intruder. This forms the lower-bound solution. Ideally, we would like the heuristic solution to be as close as



(a) Polygon $\mathcal{P}$ (b) Grid graph $G(\mathcal{P})$ (c) $\beta = 4$

(d) Polygon $\mathcal{P}$ (e) Grid graph $G(\mathcal{P})$ (f) $\beta = 5$

Figure 6.6: The region $\mathcal{P}$ with $\beta = 4$ spikes shape (top row), and $\beta = 5$ spikes (bottom row). Both the polygons have a minimum edge length of $l^b = 10m$. The orange grid shows the grid graph, $G(\mathcal{P})$, with grid cells of the size, $l^s = 5m$. The light red shaded region is the base of the comb structure both in the original polygon and the schematic illustration beside the corresponding polygons. The un-shaded regions in the polygons are the spikes s_i.

possible to the optimal solution (baseline solution).

6.7.3 Effect of increase in the number of search robots

Consider the scenario as shown in Figure 6.6(a), where the polygon $\mathcal{P}$ has $\beta = 4$ spikes depicted with uncoloured grids in Fig. 6.6(b). The polygon has a total area of 152 cells. Given 4 spikes, we vary the number of robots from 2 (or minimum possible for each approach) to 152 (or maximum possible corresponding to the number of cells) as $2, 5, 8, \ldots, 60$ with an interval of three robots. For a given number of robots, we iterate 100 simulations with randomly generated positions of the searchers and the intruders. For each simulation, we evaluate the performance of the four methods – SFC, SFC-G, RS and CRS with respect to the baseline strategy.

Figure 6.8 shows the average number of search steps required to intercept the intruder

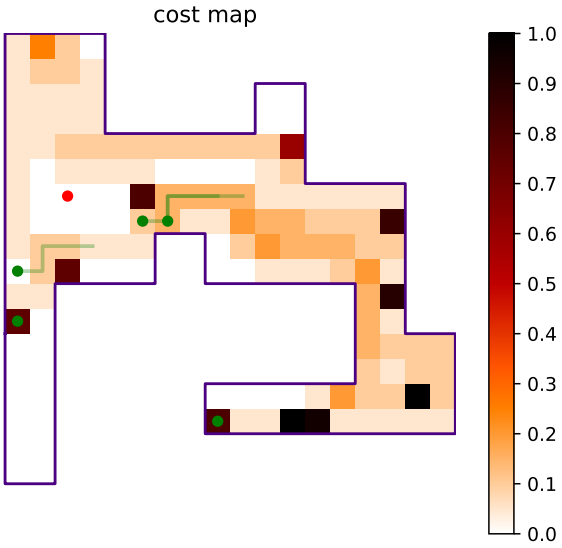


Figure 6.7: Snapshot from cooperative random search for a static intruder. Some grids appear dark due to the accumulation of revisiting costs during the waiting time of search robots until all search robots reach their goal.

for the different number of robotic searchers under S-I and M-I scenarios. From the figure, we can see that in the case of RS, CRS and Baseline strategies, solutions exist for the minimal number of searchers (2). However, for the SFC strategy, for S-I case, a minimum of 10 searchers are required to determine a solution and 44 searchers to determine a solution in the case of M-I problem. The number of searchers required in the case of SFC is high because of the rectangulation of the polygonal area as shown in Fig. 6.3(a). For each rectangle, we need at least one robot. While in the case of SFC-G, guard robots are also deployed for each junction between the rectangles. This further increases the number of robots required.

A natural hypothesis between RS and CRS is that the CRS will perform better than RS. However, the results show a different trend, where RS is performing better than CRS in the S-I case, while CRS is performing better than RS in the M-I case. For S-I, this happens because, in the case of RS, the robots are randomly given goals and upon reaching the goal they are immediately given another random goal. However, in the case of CRS, the robot that has reached a goal early has to wait for all the other robots to reach their goal and then generate the next set of goals. Due to this, there is a significant loss of search time steps for the agents as shown in Fig. 6.7. This loss is reflected in the results in terms of increased search time. In the M-I case, the motion of the intruder helps the searchers to capture it quickly, given that the number of searchers is much higher.

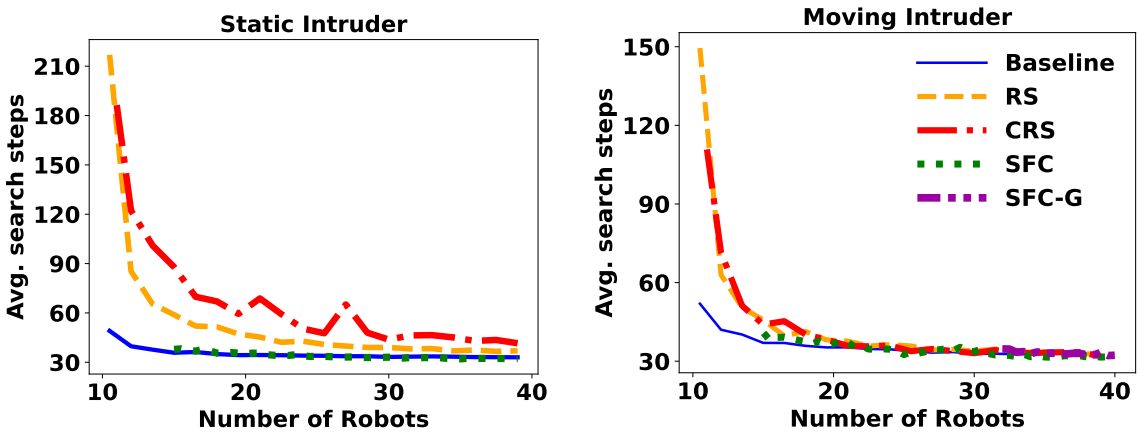


Figure 6.8: Effect of increase in the number of searchers on search steps for different search strategies on 4-spiked polygon. (a) for static intruder case and (b) for moving intruder case.

When the number of searchers is high, in the CRS approach, the robots have to wait for longer at their target grids for other searchers to reach soon. However, by the time all robots reach their target grids, the randomly moving intruder gets discovered by the distributed waiting searchers.

When the minimum number of robots available is equal to the number of rectangles in the arena, then SFC performs far better than RS, CRS and SFC-G. The results are very close to the baseline strategy. This is because the robots need to search fewer number of cells. Note that, in the case of static intruder case, we do not include the SFC-G method because guards are not required in this case.

6.7.4 Effect of change in the shape of the polygon

Further, changing the shape of the polygonal region is similar to re-arranging the spikes, and it should not affect the search time required by the k robots. To validate the same, we consider 13 search robots in three polygons of varying shapes and equal area (4400 $unit^2$) that fall in the category of $\beta = 5$. As shown in Fig. 6.9, for each proposed approach, the search time remains statistically indifferent with changing shapes, irrespective of whether the intruder is static or moving. In the case of a static intruder, regardless of the shape of the polygon, the search time closest to the Baseline approach is in the case of SFC. Whereas, RS performs better than CRS and worse than SFC. Moreover, the intruder searching is faster in the case of the Baseline and SFC approach when the intruder is

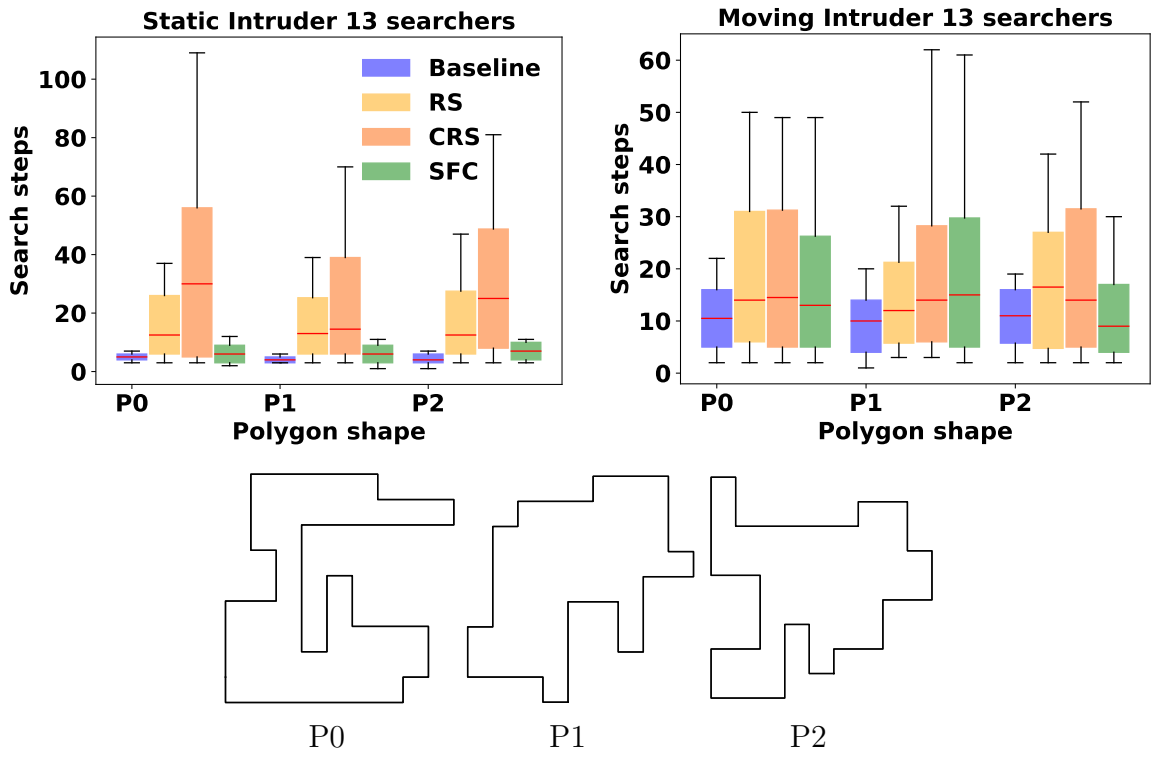


Figure 6.9: Effect of change in the shape of the polygon (area 4400 $unit^2$ and $\beta = 5$) on search steps for different search strategies and $k = 13$ search robots, (a) for static intruder case and (b) for moving intruder case.

static. The RS performs indifferently irrespective of the intruder being static or moving. However, CRS is faster in searching for the intruder when it is moving. Hence, we conclude that the shape of the polygon does not affect the performance of the approaches proposed in this work.

6.7.5 Effect of increase in area of the polygon

Next, we consider a polygon with a fixed shape in the category of $\beta = 5$. When there were 10 searchers, and the area of the search region was increased, the search time increased irrespective of the intruder being static or moving in all the different approaches as shown in Fig. 6.10. However, the performance of the different algorithms remained similar, as in the case of changing shapes of the polygon with a static intruder in it. In contrast, with the increasing area of the region, all the approaches took longer to search for the moving intruder than for the static intruder.

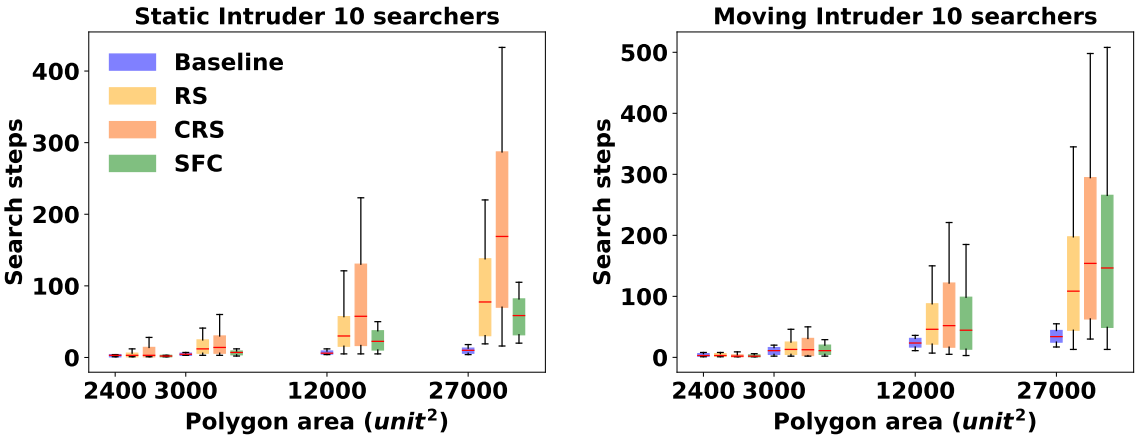


Figure 6.10: Effect of change in the area of the polygon (fixed shape and $\beta = 5$) on search steps for different search strategies and $k = 10$ search robots, (a) for static intruder case and (b) for moving intruder case.

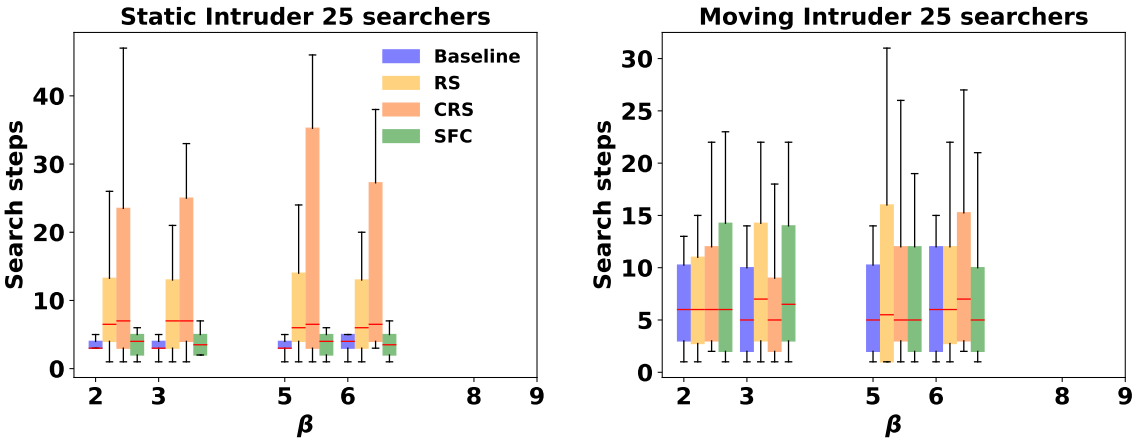


Figure 6.11: Effect of increase in the number of spikes on search steps for different search strategies and $k = 25$ search robots, (a) for static intruder case and (b) for moving intruder case when there was 160 total number of cells.

6.7.6 Effect of increase in the number of spikes

The geometric complexity of the search region is captured by the number of spikes. To understand the effect of geometric complexities on search time in the given scenario, we increase the number of spikes from 2 to 6 and see the change in performance as shown in Figure 6.11. The total area of the polygon is fixed to 4000 $unit^2$; however, the shape of the polygon changes with the changing β. We can observe that when there were 25 search robots, both for static and moving intruders, the number of search steps required to accomplish the task remains unaffected as the number of spikes, β, increase from 2 to 6 in all four approaches. This simple illustration shows that with an increase in geometric complexity, the search time is unaffected. However, we believe that increased search overlap percentage (resulting from increasing β) should significantly reduce the effective search in new regions and hence increase the intruder discovery time. To validate our hypothesis, more work is required to generate polygons with desired β, which is beyond the scope of this work. Furthermore, we also observe that the searching of a moving intruder becomes faster with the RS and CRS approach as compared to the static intruder. Meanwhile, searching becomes slower for the moving intruder with the Baseline and SFC approach.

6.8 Conclusion

In this work, we study the MRS and DMRS problems inside a simply connected orthogonal polygon. We have shown that MRS is NP-hard, and conjectured that DMRS is also NP-hard. On the positive side, we proposed three algorithmic methods to provide solutions to the problem. We performed Monte-Carlo simulations and evaluated the performance of these methods and also compared them with a baseline. The RS and CRS methods are effective even with a minimal number of search robots; however, they need longer search time. The space-filling methods provide solutions close to the baseline. However, it requires a minimum number of robots equal to the decomposed number of rectangular regions in $\mathcal{P}$. With an increased number of robots, the performance is almost near the baseline. In most scenarios, finding a moving intruder via RS and CRS approaches is quicker. In contrast, finding a static intruder via Baseline and SFC approaches is quicker. Also, we conclude that a few geometric properties like β and area

can severely affect the searching, whereas in other cases, for instance, the shapes with the same geometric properties do not affect the searching significantly.

CHAPTER 7

CONCLUSION AND FUTURE WORK

In this thesis, we developed physically grounded and computationally minimal mechanisms for three fundamental classes of swarm behaviors: (i) collective decision-making, (ii) collective navigation, and (iii) collective exploration. Across these domains, the central objective was to demonstrate that robust and scalable collective intelligence can emerge from simple local rules, even under severe sensing, communication, and computational constraints.

For collective decision-making, the problem was formulated as a best-of-n task (Chapters 2 and 3). This thesis shows that a response-threshold–based mechanism, requiring only binary responses and minimal individual memory, is sufficient for a robot swarm to reliably identify and agree upon the best available option. Analytical expressions were derived for the mean and variance of the optimal response-threshold distribution that maximizes collective success. Building on this theoretical foundation, a decentralized and lightweight adaptive controller, termed the h-controller, was proposed to enable miniature robots to autonomously adapt their thresholds and prioritize high-value tasks. The resulting framework is scalable, resilient to noise and failures, and particularly well suited to resource-constrained robotic platforms.

For collective navigation, this thesis introduced a contrasting-strategy–based extension of classical flocking models (Chapters 4 and 5). By allowing agents to probabilistically defect from local alignment rules, the proposed model achieves faster collective alignment while simultaneously reducing overall energy expenditure. A systematic exploration of the parameter space demonstrates that these benefits are not limited to specific scenarios but arise as a general property of the interaction dynamics. Furthermore, the analysis reveals how controlled defection leads to the spontaneous emergence of anonymous leadership, enabling flexible and safe human guidance of the swarm without reliance on predefined leaders or centralized control.

For collective exploration, the thesis examined multi-robot target-search problems in

closed polygonal environments under realistic sensing limitations (Chapter 6). Several existing search strategies were analyzed and compared to provide actionable guidance to human operators regarding their applicability across different mission constraints. The study establishes that the Multi-Robot Search (MRS) problem in simply connected orthogonal polygons is NP-hard and conjectures that its dynamic variant (DMRS) is also NP-hard. Through geometric analysis, this work further clarifies the extent to which target-search tasks—and, more broadly, real-world robotic missions—can be meaningfully reformulated within a swarm-intelligence paradigm.

Collectively, these contributions reinforce the viability of swarm intelligence as a practical and principled approach for deploying large-scale robot systems in uncertain, constrained, and safety-critical environments.

Future Work

Several promising research directions naturally follow from this thesis across each class of swarm behaviors.

Collective Decision-making

First, the present rule only tunes the mean of the threshold distribution; dynamically controlling the variance could further improve performance in tasks that require balancing exploration and exploitation. Second, more sophisticated methods of estimating the number of options or sharing this estimate among neighbors may reduce uncertainty and accelerate adaptation. Third, integrating the threshold mechanism with voting schemes or consensus protocols (e.g., [200–202]) may yield hybrid strategies that combine the low computational cost of response thresholds with the group cohesion of voting. Finally, deploying the threshold adaptation rule on physical robots would help bridge the gap between our theoretical, simulation-based study and practical real-world relevance, by exposing the method to real sensing, actuation, and communication constraints.

Collective Navigation

One could investigate the causal relationship between spatially selective defection and emergent group dynamics, as well as its dependence on swarm size. For example, we hypothesize that regulating the density of defecting agents within each interaction zone

could provide more insights into the observed dynamics or help explain the size-dependent dynamics of defectors in the zone orientation (Fig. 4.5). Zone-dependent probabilistic defection strategies may further enhance energy efficiency, alignment speed, and robustness against adversarial interference [203, 204]. One may study in the future to infer the actual catalyzer of alignment, whether defectors alone naturally emerge as leaders or some combination of cooperator-defector leads to the emergence of leadership. Integrating the proposed motion models with evolutionary game-theoretic and probabilistic response threshold frameworks may also yield deeper insights into the emergence of leadership, leader–follower dynamics, and collective splitting or merging behaviors [7, 82, 83, 205, 206].

Collective Exploration

Future extensions include multi-intruder scenarios, intelligent intruders with evasive strategies, and environments with increased geometric complexity, such as polygons with holes. While space-filling-curve-based methods exhibit strong performance, their efficiency depends on the underlying decomposition of the environment. Developing algorithms that minimize such decompositions—or eliminate them altogether through generalized space-filling curves remains an open and mathematically challenging problem.

Overall, this thesis establishes a foundation for resource-constrained (minimalistic), scalable, and resilient robot swarms. Our framework may inspire future engineering of emergent swarms in a way that enables a human operator to interact with them efficiently and anonymously. Lastly, the study opens several pathways for deploying robot swarms in uncertain, safety-critical real-world environments.

REFERENCES

[1] Eisuke Hasegawa, Nobuaki Mizumoto, Kazuya Kobayashi, Shigeto Dobata, Jin Yoshimura, Saori Watanabe, Yuuka Murakami, and Kenji Matsuura. Nature of collective decision-making by simple yes/no decision units. *Scientific Reports*, 7:14436, 2017.

[2] Iain D. Couzin, Jens Krause, Richard James, Graeme D. Ruxton, and Nigel R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 2002.

[3] Robin R. Murphy, Vignesh B.M. Gandudi, Trisha Amin, Angela Clendenin, and Jason Moats. An analysis of international use of robots for covid-19. *Robotics and Autonomous Systems*, 148:103922, 2022.

[4] Heiko Hamann. *Swarm robotics: A formal approach*. Springer, Cham, Switzerland, 2018.

[5] Marco Dorigo, Guy Theraulaz, and Vito Trianni. Swarm robotics: Past, present, and future [point of view]. *Proceedings of the IEEE*, 109(7):1152–1165, 2021.

[6] Andreas Kolling, Phillip Walker, Nilanjan Chakraborty, Katia Sycara, and Michael Lewis. Human interaction with robot swarms: A survey. *IEEE Transactions on Human-Machine Systems*, 46(1):9–26, 2016.

[7] Michael A. Goodrich, Brian Pendleton, Sujit P. Baliyarasimhuni, and José Pinto. Toward human interaction with bio-inspired robot teams. In *2011 IEEE International Conference on Systems, Man, and Cybernetics*, pages 2859–2864. IEEE, 2011.

[8] Dji drones. https://www.dji.com/global.

[9] Neo home robot. https://www.1x.tech/neo.

[10] Figure. https://www.figure.ai/.

[11] Unitree. https://www.unitree.com/.

[12] Boston dynamics. https://bostondynamics.com/.

[13] Robotswim. http://www.robotswim.com/?lang=English.

[14] Blue robotics. https://bluerobotics.com/store/rov/bluerov2/.

[15] Searobotics. https://searobotics.com/.

[16] Ocean power technologies. https://oceanpowertechnologies.com/.

[17] Kuka. https://www.kuka.com/.

[18] Noah T. Jafferis, E. Farrell Helbling, Michael Karpelson, and Robert J. Wood. Untethered flight of an insect-sized flapping-wing microscale aerial vehicle. *Nature*, 570(7762):491–495, 2019.

[19] Sam Kriegman, Douglas Blackiston, Michael Levin, and Josh Bongard. Kinematic self-replication in reconfigurable organisms. *Proceedings of the National Academy of Sciences*, 118(49):e2112672118, 2021.

[20] Bss holland. https://bssholland.com/product/black-hornet-4/.

[21] Yogesh M Chukewad, Johannes James, Avinash Singh, and Sawyer Fuller. Robofly: An insect-sized robot with simplified fabrication that is capable of flight, ground, and water surface locomotion, 2020, 2001.02320.

[22] Tian Qiu, Tung-Chun Lee, Andrew G. Mark, Konstantin I. Morozov, Raphael Münster, Otto Mierka, Stefan Turek, Alexander M. Leshansky, and Peer Fischer. Swimming by reciprocal motion at low reynolds number. *Nature Communications*, 5(1):5119, Nov 2014.

[23] Rani therapeutics. https://www.ranitherapeutics.com/technology/.

[24] Bitcraze. https://www.bitcraze.io/.

[25] Konstantinos V. Katsikopoulos and Andrew J. King. Swarm intelligence in animal groups: When can a collective out-perform an expert? *PLOS ONE*, 5(11):1–5, 2010.

[26] Eric Bonabeau, Guy Theraulaz, Jean-Louls Deneubourg, Serge Aron, and Scott Camazine. Self-organization in social insects. *Trends in Ecology & Evolution*, 12(5):188–193, 1997.

[27] B. Hölldobler, E.O.W. Bert Hölldobler, R.J.C.S.I.R.U.P.L.S.R.P.B. Hölldobler, E.O. Wilson, H.C.E.U.R.P.E.E.O. Wilson, and F.P.B.B. Holldobler. *The Ants*. Belknap Press of Harvard University Press, 1990.

[28] Ashay Aswale, Antonio López, Aukkawut Ammartayakun, and Carlo Pinciroli. Hacking the colony: On the disruptive effect of misleading pheromone and how to defend against it. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '22, page 27–34, Richland, SC, 2022. International Foundation for Autonomous Agents and Multiagent Systems.

[29] Anders Lyhne Christensen, Rehan OGrady, and Marco Dorigo. From fireflies to fault-tolerant swarms of robots. *IEEE Transactions on Evolutionary Computation*, 13(4):754–766, 2009.

[30] Thomas D Seeley. Social foraging in honey bees: How nectar foragers assess their colony's nutritional status. *Behavioral Ecology and Sociobiology*, 24(3):181–199, 1989.

[31] Neil B. Metcalfe and Bruce C. Thomson. Fish recognize and prefer to shoal with poor competitors. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 259(1355):207–210, 1995.

[32] Bernard Chazelle. The convergence of bird flocking. *Journal of the Association for Computing Machinery*, 61(4), 2014.

[33] Iain D. Couzin. Collective animal migration. *Current Biology*, 28(17):R976–R980, 2018.

[34] Andrea Flack, Máté Nagy, Wolfgang Fiedler, Iain D. Couzin, and Martin Wikelski. From local collective behavior to global migratory patterns in white storks. *Science*, 360(6391):911–914, 2018.

[35] Shivani Krishna, Apoorva Gopinath, and Somendra M. Bhattacharjee. Ordering and topological defects in social wasps' nests. *Scientific Reports*, 12(1):12901, Jul 2022.

[36] Jason M. Graham, Albert B. Kao, Dylana A. Wilhelm, and Simon Garnier. Optimal construction of army ant living bridges. *Journal of Theoretical Biology*, 435:184–198, 2017.

[37] Simon Garnier, Tucker Murphy, Matthew Lutz, Edward Hurme, Simon Leblanc, and Iain D. Couzin. Stability and responsiveness in a self-organized living architecture. *PLOS Computational Biology*, 9(3):1–10, 03 2013.

[38] Scott Camazine. Self-organizing pattern formation on the combs of honey bee colonies. *Behavioral Ecology and Sociobiology*, 28(1):61–76, Jan 1991.

[39] ISTVÁN KARSAI and GÁBOR BALÁZSI. Organization of work via a natural substance: Regulation of nest construction in social wasps. *Journal of Theoretical Biology*, 218(4):549–565, 2002.

[40] Guy Theraulaz and Eric Bonabeau. Coordination in distributed building. *Science*, 269(5224):686–688, 1995.

[41] David W. Stephens. Decision ecology: Foraging and the ecology of animal decision making. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4):475–484, 2008.

[42] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.

[43] Iain D. Couzin, Jens Krause, Nigel R. Franks, and Simon A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature 2004 433:7025*, 433(7025):513–516, 2005.

[44] Shripad Gade, Aditya A. Paranjape, and Soon-Jo Chung. Herding a flock of birds approaching an airport using an unmanned aerial vehicle. In *AIAA Guidance, Navigation, and Control Conference*. American Institute of Aeronautics and Astronautics, Jan 2015.

[45] Luis Gómez-Nava, Richard Bon, and Fernando Peruani. Intermittent collective motion in sheep results from alternating the role of leader and follower. *Nature Physics*, 18(12):1494–1501, 2022.

[46] Ofer Feinerman, Itai Pinkoviezky, Aviram Gelblum, Ehud Fonio, and Nir S. Gov. The physics of cooperative transport in groups of ants. *Nature Physics*, 14(7):683–693, Jul 2018.

[47] Daniele Carlesso, Madelyne Stewardson, Donald James McLean, Geoffrey P. F. Mazué, Simon Garnier, Ofer Feinerman, and Chris R. Reid. Leaderless consensus decision-making determines cooperative transport direction in weaver ants. *Proceedings of the Royal Society B: Biological Sciences*, 291(2028):20232367, 08 2024.

[48] Palina Bartashevich, James E. Herbert-Read, Matthew J. Hansen, Félicie Dhellemmes, Paolo Domenici, Jens Krause, and Pawel Romanczuk. Collective anti-predator escape manoeuvres through optimal attack and avoidance strategies. *Communications Biology*, 7(1):1586, Nov 2024.

[49] Zhijie Zhao, Qichen Shi, and Yong Liu. Modeling and simulation of the fish collective behavior with risk perception and startle cascades. *Physica A: Statistical Mechanics and its Applications*, 659:130337, 2025.

[50] Ugo Lopez, Jacques Gautrais, Iain D. Couzin, and Guy Theraulaz. From behavioural analyses to models of collective motion in fish schools. *Interface Focus*, 2(6):693–707, 10 2012.

[51] Guy Theraulaz and Eric Bonabeau. A brief history of stigmergy. *Artificial Life*, 5(2):97–116, 04 1999.

[52] Yun Kang and Guy Theraulaz. Dynamical models of task organization in social insect colonies. *Bulletin of Mathematical Biology*, 78(5):879–915, May 2016.

[53] Thomas H Labella, Marco Dorigo, and Jean-Louis Deneubourg. Division of labor in a group of robots inspired by ants' foraging behavior. *ACM Transactions on Autonomous and Adaptive Systems*, 1(1):4–25, 2006.

[54] Eisuke Hasegawa, Yasunori Ishii, Koichiro Tada, Kazuya Kobayashi, and Jin Yoshimura. Lazy workers are necessary for long-term sustainability in insect societies. *Scientific Reports*, 6(1):20846, 2016.

[55] SCOTT CAMAZINE, JEAN-LOUIS DENEUBOURG, NIGEL R. FRANKS, JAMES SNEYD, GUY THERAULAZ, and ERIC BONABEAU. *Self-Organization in Biological Systems*, volume 38. Princeton University Press, 2025/12/29 2001.

[56] G. C. H. E. de Croon, J. J. G. Dupeyroux, S. B. Fuller, and J. A. R. Marshall. Insect-inspired ai for autonomous robots. *Science Robotics*, 7(67):eabl6334, 2022.

[57] Gene E. Robinson. Modulation of alarm pheromone perception in the honey bee: evidence for division of labor based on hormonall regulated response thresholds. *Journal of Comparative Physiology A*, 160(5):613–619, 1987.

[58] Jacob W. Crandall, Nathan Anderson, Chace Ashcraft, John Grosh, Jonah Henderson, Joshua McClellan, Aadesh Neupane, and Michael A. Goodrich. Human-swarm interaction as shared control: Achieving flexible fault-tolerant systems. In *Engineering Psychology and Cognitive Ergonomics: Performance, Emotion and Situation Awareness*. Springer International Publishing, 2017.

[59] Anja Weidenmüller. The control of nest climate in bumblebee (*Bombus terrestris*) colonies: interindividual variability and self reinforcement in fanning response. *Behavioral Ecology*, 15(1):120–128, 2004.

[60] Tamás Vicsek and Anna Zafeiris. Collective motion. *Physics Reports*, 517(3):71–140, 2012.

[61] Michal Pluhacek, Simon Garnier, and Andreagiovanni Reina. Decentralised construction of a global coordinate system in a large swarm of minimalistic robots. *Swarm Intelligence*, Jul 2025.

[62] Andreagiovanni Reina, Gabriele Valentini, Cristian Fernández-Oto, Marco Dorigo, and Vito Trianni. A design pattern for decentralised decision making. *PLOS ONE*, 10(10):e0140950, 2015.

[63] Justin Werfel, Kirstin Petersen, and Radhika Nagpal. Designing collective behavior in a termite-inspired robot construction team. *Science*, 343(6172):754–758, 2014.

[64] Vivek Shankar Varadharajan, David St-Onge, Bram Adams, and Giovanni Beltrame. Swarm relays: Distributed self-healing ground-and-air connectivity chains. *IEEE Robotics and Automation Letters*, 5(4):5347–5354, 2020.

[65] Gabriele Valentini, Eliseo Ferrante, and Marco Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4, 2017.

[66] Gabriele Valentini, Heiko Hamann, and Marco Dorigo. Self-organized collective decision making: The weighted voter model. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS 2014, page 45–52, Richland, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

[67] Masao Kubo, Hiroshi Sato, Nhuhai Phung, and Akihiro Yamaguchi and. Collective decision-making algorithm for the best-of-n problem in multiple options. *SICE Journal of Control, Measurement, and System Integration*, 15(2):71–88, 2022, https://doi.org/10.1080/18824889.2022.2077588.

[68] Chris A. C. Parker and Hong Zhang. Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics*, 14(2):240–251, 2009.

[69] Ryusuke Fujisawa, Shigeto Dobata, Ken Sugawara, and Fumitoshi Matsuno. Designing pheromone communication in swarm robotics: Group foraging behavior mediated by chemical substance. *Swarm Intelligence*, 8(3):227–246, Sep 2014.

[70] Raina Zakir, Marco Dorigo, and Andreagiovanni Reina. Robot swarms break decision deadlocks in collective perception through cross-inhibition. In *Swarm Intelligence (ANTS 2022)*, volume 13491 of *LNCS*. Springer, Cham, 2022.

[71] Judith Pérez-Velázquez, Meltem Gölgeli, and Rodolfo García-Contreras. Mathematical modelling of bacterial quorum sensing: A review. *Bulletin of Mathematical Biology*, 78:1585–1639, 2016.

[72] Larissa Conradt. Models in animal collective decision-making: information uncertainty and conflicting preferences. *Interface Focus*, 2, 2012.

[73] Gabriele Valentini, Heiko Hamann, and Marco Dorigo. Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, page 1305–1314, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

[74] Sidney Redner. Reality-inspired voter models: A mini-review. *Comptes Rendus Physique*, 20(4):275–292, 2019.

[75] Swadhin Agrawal, Sujit P. Baliyarasimhuni, and Andreagiovanni Reina. Response threshold distributions to improve best-of-n decisions in minimalistic robot swarms. In *Swarm Intelligence*, pages 350–359, Cham, 2022. Springer International Publishing.

[76] Fabio Oddi, Andreagiovanni Reina, and Vito Trianni. Minimalist protocols for quorum sensing in robot swarms. In Heiko Hamann, Marco Dorigo, Leslie Pérez Cáceres, Andreagiovanni Reina, Jonas Kuckling, Tanja Katharina Kaiser, Mohammad Soorati, Ken Hasselmann, and Eduard Buss, editors, *Swarm Intelligence*, pages 141–154, Cham, 2024. Springer Nature Switzerland.

[77] Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226–1229, 1995, cond-mat/0611743.

[78] Jitesh Jhawar, Richard G. Morris, U. R. Amith-Kumar, M. Danny Raj, Tim Rogers, Harikrishnan Rajendran, and Vishwesha Guttal. Noise-induced schooling of fish. *Nature Physics*, 16(4):488–493, 2020.

[79] Daniel Strömbom, Richard P. Mann, Alan M. Wilson, Stephen Hailes, A. Jennifer Morton, David J.T. Sumpter, and Andrew J. King. Solving the shepherding problem: Heuristics for herding autonomous, interacting agents. *Journal of the Royal Society Interface*, 2014.

[80] Ivan D Chase. Cooperative and Noncooperative Behavior in Animals. *The American Naturalist*, 115(6):827–857, 1980.

[81] Alyssa Pierson and Mac Schwager. Controlling noncooperative herds with robotic herders. *IEEE Transactions on Robotics*, 34(2):517–525, 2018.

[82] Sara Bernardi, Raluca Eftimie, and Kevin J. Painter. Leadership through influence: What mechanisms allow leaders to steer a swarm? *Bulletin of Mathematical Biology*, 83(6):69, 2021.

[83] Rohan Tiwari, Puneet Jain, Sachit Butail, Sujit P. Baliyarasimhuni, and Michael A. Goodrich. Effect of leader placement on robotic swarm control. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, pages 1387–1394, Richland, SC, 2017. International Foundation for Autonomous Agents and Multiagent Systems.

[84] Swadhin Agrawal, Jitesh Jhawar, Andreagiovanni Reina, Sujit P. Baliyarasimhuni, Heiko Hamann, and Liang Li. Impact of individual defection on collective motion. In *Swarm Intelligence: 14th International Conference, ANTS 2024, Konstanz, Germany, October 9–11, 2024, Proceedings*, page 127–140, Berlin, Heidelberg, 2024. Springer-Verlag.

[85] Till Aust, Mohamed S. Talamali, Marco Dorigo, Heiko Hamann, and Andreagiovanni Reina. The hidden benefits of limited communication and slow sensing in collective monitoring of dynamic environments. In *Swarm Intelligence (ANTS 2022)*, volume 13491 of *LNCS*. Springer, Cham, 2022.

[86] Alexander Scheidler, Daniel Merkle, and Martin Middendorf. Stability and performance of ant queue inspired task partitioning methods. *Theory in Biosciences*, 127(2):149–161, 2008.

[87] SM LaValle and JE Hinrichsen. Visibility-based pursuit-evasion: The case of curved environments. *IEEE Trancsactions on Robotics and Automation*, 17(2):196–202, APR 2001.

[88] Swadhin Agrawal, Sujoy Bhore, Joseph S.B. Mitchell, P.B. Sujit, and Aayush Gohil. Multi-robot searching with limited sensing range for static and mobile intruders. In *2025 International Conference on Emerging Technology in Autonomous Aerial Vehicles (ETAAV)*, pages 1–7, 2025.

[89] William M. Spears, Diana F. Spears, Rodney Heil, Wesley Kerr, and Suranga Hettiarachchi. An overview of physicomimetics. In Erol Şahin and William M. Spears, editors, *Swarm Robotics*, pages 84–97, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

[90] Immihan Ceren Yasa, Hakan Ceylan, Ugur Bozuyuk, Anna-Maria Wild, and Metin Sitti. Elucidating the interaction dynamics between microswimmer body and immune system for medical microrobots. *Science Robotics*, 5(43), 2020.

[91] Andreagiovanni Reina, Eliseo Ferrante, and Gabriele Valentini. Collective decision-making in living and artificial systems: editorial. *Swarm Intelligence*, 15(1):1–6, 2021.

[92] Mohamed S. Talamali, Arindam Saha, James A. R. Marshall, and Andreagiovanni Reina. When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics*, 6(56):eabf1416, 2021.

[93] Gabriele Valentini, Eliseo Ferrante, Heiko Hamann, and Marco Dorigo. Collective decision with 100 Kilobots: speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, 30(3):553–580, 2016.

[94] Mohamed S. Talamali, James A. R. Marshall, Thomas Bose, and Andreagiovanni Reina. Improving collective decision accuracy via time-varying cross-inhibition. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9652–9659, 2019.

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

[96] Elva J. H. Robinson, Nigel R. Franks, Samuel Ellis, Saki Okuda, and James A. R. Marshall. A simple threshold rule is sufficient to explain sophisticated collective decision-making. *PLOS ONE*, 6(5):e19981, 2011.

[97] Naoki Masuda, Thomas A. O'Shea-Wheller, Carolina Doran, and Nigel R. Franks. Computational model of collective nest selection by ants with heterogeneous acceptance thresholds. *Royal Society Open Science*, 2(6):140533, 2015.

[98] Takao Sasaki and Stephen C. Pratt. Emergence of group rationality from irrational individuals. *Behavioral Ecology*, 22(2):276–281, 2011.

[99] Tatsuhiro Yamamoto and Eisuke Hasegawa. Response threshold variance as a basis of collective rationality. *Royal Society Open Science*, 4(4):170097, 2017.

[100] James A R Marshall, Gavin Brown, and Andrew N Radford. Individual confidence-weighting and group decision-making. *Trends in Ecology & Evolution*, 32(9):636–645, 2017.

[101] Fernanda Sánchez-Puig, Octavio Zapata, Omar K. Pineda, Gerardo Iñiguez, and Carlos Gershenson. Heterogeneity extends criticality. *Frontiers in Complex Systems*, Volume 1 - 2023, 2023.

[102] Michael J.B. Krieger and Jean-Bernard Billeter. The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, 30(1-2):65–84, 2000.

[103] William Agassounon and Alcherio Martinoli. Efficiency and robustness of threshold-based distributed allocation algorithms in multi-agent systems. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems (AAMAS 2002)*, pages 1090–1097, New York, USA, 2002. ACM Press.

[104] Wenguo Liu, Alan F. T. Winfield, Jin Sa, Jie Chen, and Lihua Dou. Towards energy optimization: Emergent task allocation in a swarm of foraging robots. *Adaptive Behavior*, 15(3):289–305, 2007.

[105] Paulo R. Ferreira, Felipe S. Boffo, and Ana L. C. Bazzan. Using swarm-GAP for distributed task allocation in complex scenarios. In *Massively Multi-Agent Technology*, volume 5043 of *LNAI*, pages 107–121. Springer Berlin, Heidelberg, 2008.

[106] Anshul Kanakia, John Klingner, and Nikolaus Correll. A response threshold sigmoid function model for swarm robot collaboration. In *Distributed Autonomous Robotic Systems, Springer Tracts in Advanced Robotics (STAR, volume 112)*, pages 193–206, Tokyo, 2016. Springer Japan.

[107] Eduardo Castello, Tomoyuki Yamamoto, Fabio Dalla Libera, Wenguo Liu, Alan F T Winfield, Yutaka Nakamura, and Hiroshi Ishiguro. Adaptive foraging for simulated and real robotic swarms: the dynamical response threshold approach. *Swarm Intelligence*, 10(1):1–31, 2016.

[108] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60:581–588, 1973, https://academic.oup.com/biomet/article-pdf/60/3/581/576759/60-3-581.pdf.

[109] Richard A. Holley and Thomas M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, 3:643 – 663, 1975.

[110] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[111] Swadhin Agrawal, Sujit P. Baliyarasimhuni, and Andreagiovanni Reina. Supplementary materials of the article "Response threshold distributions to improve best-of-n decisions in minimalistic robot swarms". https://github.com/zorawar12/yesnounits.git.

[112] Alex Hansen. The three extreme value distributions: An introductory review. *Frontiers in Physics*, 8, 2020.

[113] Hyunjoong Kim, Zachary Kilpatrick, and Kresimir Josic. Daring few, patient many: division of labor in decentralized foraging collectives, 2026, 2602.08840.

[114] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.

[115] Aaron E. Hirsh and Deborah M. Gordon. Distributed problem solving in social insects. *Annals of Mathematics and Artificial Intelligence*, 31:199–221, 2001.

[116] Iain D. Couzin. Collective intelligence in animals and robots. *Nature Communications*, 16:65814, 2025.

[117] Sampriti Mukherjee and Bonnie L. Bassler. Bacterial quorum sensing in complex and dynamically changing environments. *Nature Reviews Microbiology*, 17:371–382, 2019.

[118] Noa Pinter-Wollman, Julia Hubler, Jo-Anne Holley, Nigel R. Franks, and Anna Dornhaus. How is activity distributed among and within tasks in Temnothorax ants? *Behavioral Ecology and Sociobiology*, 66:1407–1420, 2012.

[119] Eduardo Castello, Yutaka Nakamura, Tomoyuki Yamamoto, and Hiroshi Ishiguro. Task allocation for a robotic swarm based on an adaptive response threshold model. In *2013 13th International Conference on Control, Automation and Systems (ICCAS 2013)*, pages 259–266, 2013.

[120] Matimba Swana, Johanna Blee, Namid Stillman, Jonathan Ives, and Sabine Hauert. Swarms: The next frontier for cancer nanomedicine. In *Cancer, Complexity, Computation*, volume 46 of *Emergence, Complexity and Computation*, pages 269–288. Springer, 2022.

[121] Can Dincer, Richard Bruch, Estefanía Costa-Rama, María Teresa Fernández-Abedul, Arben Merkoçi, Andreas Manz, Gerald Anton Urban, and Firat Güder. Disposable sensors in diagnostics, food, and environmental monitoring. *Advanced Materials*, 31(30):1806739, 2019.

[122] Thomas A. O'Shea-Wheller, Edmund R. Hunt, and Takao Sasaki. Functional heterogeneity in superorganisms: Emerging trends and concepts. *Annals of the Entomological Society of America*, 114(5):562–574, 2021.

[123] E. O. Wilson. Behavioral discretization and the number of castes in an ant species. *Behavioral Ecology and Sociobiology*, 1:141–154, 1976.

[124] Ralf H J M Kurvers, Max Wolf, and Jens Krause. Humans use social information to adjust their quorum thresholds adaptively in a simulated predator detection experiment. *Behavioral Ecology and Sociobiology*, 68(3):449–456, 2014.

[125] PETER CLIFFORD and AIDAN SUDBURY. A model for spatial conflict. *Biometrika*, 60(3):581–588, 12 1973.

[126] James Cruise and Ayalvadi Ganesh. Probabilistic consensus via polling and majority rules. *Queueing Systems*, 78(2):99–120, Oct 2014.

[127] Guy Theraulaz, Eric Bonabeau, and Jean-Louis Deneubourg. Response threshold reinforcement and division of labour in insect societies. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 265(1393):327–332, 1998.

[128] Shaowu Zhang, Frank Bock, Aung Si, Jürgen Tautz, and Mandyam V. Srinivasan. Visual working memory in decision making by honey bees. *Proceedings of the National Academy of Sciences of the United States of America*, 102(14):5250–5255, 2005.

[129] Volker Schaller, Christoph Weber, Christine Semmrich, Erwin Frey, and Andreas R. Bausch. Polar patterns of driven filaments. *Nature*, 467(7311):73–77, 2010.

[130] Mircea R Davidescu, Pawel Romanczuk, Thomas Gregor, and Iain D Couzin. Growth produces coordination trade-offs in trichoplax adhaerens, an animal lacking a central nervous system. *Proceedings of the National Academy of Sciences*, 120(11):e2206163120, 2023.

[131] Pavan Ramdya, Pawel Lichocki, Steeve Cruchet, Lukas Frisch, Winnie Tse, Dario Floreano, and Richard Benton. Mechanosensory interactions drive collective behaviour in drosophila. *Nature*, 519(7542):233–236, 2015.

[132] J. Buhl, D. J. T. Sumpter, I. D. Couzin, J. J. Hale, E. Despland, E. R. Miller, and S. J. Simpson. From disorder to order in marching locusts. *Science*, 312(5778):1402–1406, 2006.

[133] Kolbjørn Tunstrøm, Yael Katz, Christos C. Ioannou, Cristián Huepe, Matthew J. Lutz, and Iain D. Couzin. Collective states, multistability and transitional behavior in schooling fish. *PLOS Computational Biology*, 9(2):1–11, 2013.

[134] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237, 2008.

[135] Noam Miller, Simon Garnier, Andrew T Hartnett, and Iain D Couzin. Both information and social cohesion determine collective decisions in animal groups. *Proceedings of the National Academy of Sciences*, 110(13):5263–5268, 2013.

[136] Yael Katz, Kolbjørn Tunstrøm, Christos C. Ioannou, Cristián Huepe, and Iain D. Couzin. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725, 2011.

[137] João A. Moreira, Jorge M. Pacheco, and Francisco C. Santos. Evolution of collective action in adaptive social structures. *Scientific Reports*, 3:1521, 2013.

[138] Bertrand Collignon, Axel Séguret, Yohann Chemtob, Leo Cazenille, and José Halloy. Collective departures and leadership in zebrafish. *PLoS ONE*, 14(5):e0216798, 2019.

[139] Wen-Chi Yang and Thomas Schmickl. Collective motion as an ultimate effect in crowded selfish herds. *Scientific Reports*, 9:6618, 2019.

[140] E.M.A. Hensor, J.-G.J. Godin, D.J. Hoare, and J. Krause. Effects of nutritional state on the shoaling tendency of banded killifish, fundulus diaphanus, in the field. *Animal Behaviour*, 65(4):663–669, 2003.

[141] Egbert G. Leigh. How does selection reconcile individual advantage with the good of the group? *Proceedings of the National Academy of Sciences*, 74(10):4542–4546, 1977.

[142] Hannah E. A. MacGregor and Christos C. Ioannou. Collective motion diminishes, but variation between groups emerges, through time in fish shoals. *Royal Society Open Science*, 8(10):210655, 2021.

[143] Nikolai W. F. Bode, Jolyon J. Faria, Daniel W. Franks, Jens Krause, and A. Jamie Wood. How perceived threat increases synchronization in collectively moving animal groups. *Proceedings of the Royal Society B: Biological Sciences*, 277(1697):3065–3070, 2010.

[144] Nazim Fatès. A tutorial on elementary cellular automata with fully asynchronous updating. *Natural Computing*, 19(1):179–197, 2020.

[145] Vivek Jadhav, Vishwesha Guttal, and Danny Raj Masila. Randomness in the choice of neighbours promotes cohesion in mobile animal groups. *Royal Society Open Science*, 9(3):220124, 2022.

[146] Aming Li, Lei Zhou, Qi Su, Sean P. Cornelius, Yang-Yu Liu, Long Wang, and Simon A. Levin. Evolution of cooperation on temporal networks. *Nature Communications*, 11(1):2259, 2020.

[147] Liang Li, Chen Chen, and Aming Li. Autonomy promotes the evolution of cooperation in prisoner's dilemma. *Physical Review E*, 102, 2020.

[148] Alberto Antonioni and Alessio Cardillo. Coevolution of synchronization and cooperation in costly networked interactions. *Physical Review Letters*, 118(23):238301, 2017, 1607.03186.

[149] Feng You, Han-Xin Yang, Yumeng Li, Wenbo Du, and Gang Wang. A modified vicsek model based on the evolutionary game. *Applied Mathematics and Computation*, 438:127565, 2023.

[150] Daniel Strömbom, Tasnia Hassan, W. Hunter Greis, and Alice Antia. Asynchrony induces polarization in attraction-based models of collective motion. *Royal Society Open Science*, 6(4):190381, 2019.

[151] Siddharth Mayya and Magnus Egerstedt. Safe open-loop strategies for handling intermittent communications in multi-robot systems. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5818–5823, Piscataway, NJ, 2017.

[152] Jennifer Gielis, Ajay Shankar, and Amanda Prorok. A critical review of communications in multi-robot systems. *Current Robotics Reports*, 3(4):213–225, 2022.

[153] M. van Otterlo. *The Logic of Adaptive Behavior: Knowledge Representation and Algorithms for Adaptive Sequential Decision Making Under Uncertainty in First-order and Relational Domains*. Frontiers in artificial intelligence and applications. IOS Press, Amsterdam, The Netherlands, 2009.

[154] Yogesh Kumar, Aditya A. Paranjape, Supratim Ghosh, and Sujit P. Baliyarasimhuni. Adversarial fragmentation of robotic teams operating under reynolds' rules with bounded communication radius. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 2809–2814. IEEE, 2023.

[155] L. E. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):497–516, 1957.

[156] Swadhin Agrawal, Jitesh Jhawar, Andreagiovanni Reina, Sujit P. Baliyarasimhuni, Heiko Hamann, and Liang Li. Supplementary materials of the article "Impact of individual defection on collective motion". https://github.com/swadhinagrawal/defectorsInCollectiveMotion.git.

[157] Daniel T Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, 1976.

[158] David J. T. Sumpter. *Collective Animal Behavior*. Princeton University Press, 2010.

[159] Andrea Cavagna and Irene Giardina. Bird flocks as condensed matter. *Annual Review of Condensed Matter Physics*, 5:183–207, 2014.

[160] Guillaume Grégoire and Hugues Chaté. Onset of collective and cohesive motion. *Physical Review Letters*, 92(2):025702, 2004.

[161] Andrea Cavagna, Andrea Cimarelli, Irene Giardina, Giorgio Parisi, Raffaele Santagati, Fabio Stefanini, and Massimiliano Viale. Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences*, 107(26):11865–11870, 2010.

[162] James E Herbert-Read, Andrea Perna, Richard P Mann, Timothy M Schaerf, David J. T. Sumpter, and Ashley J. W. Ward. Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, 108(46):18726–18731, 2011.

[163] Damien R Farine and Iain D Couzin. Leadership and social information use in animal groups. *Trends in Ecology & Evolution*, 30(6):347–357, 2015.

[164] Andrea Attanasi, Andrea Cavagna, Lorenzo Del Castello, Irene Giardina, Andrija Jelic, Stefano Melillo, Leonardo Parisi, Oliver Pohl, Edward Shen, and Massimiliano Viale. Information transfer and behavioural inertia in starling flocks. *Nature Physics*, 10:691–696, 2014.

[165] Vishwesha Guttal and Iain D Couzin. Social interactions, information use, and the evolution of collective migration. *Proceedings of the National Academy of Sciences*, 107(37):16172–16177, 2010.

[166] Jitesh Jhawar and Vishwesha Guttal. Noise-induced modes in collective behaviour. *Journal of The Royal Society Interface*, 2020.

[167] Akanksha Rathore et al. Probabilistic interaction models in collective motion. *Journal of The Royal Society Interface*, 2022.

[168] Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7:1–41, 2013.

[169] Torrence D Parsons. Pursuit-evasion in a graph. In *Theory and Applications of Graphs: Proceedings, Michigan May 11–15, 1976*, pages 426–441. Springer, 2006.

[170] Andreas Kolling and Stefano Carpin. Pursuit-evasion on trees by robot teams. *IEEE TRANSACTIONS ON ROBOTICS*, 26(1):32–47, FEB 2010.

[171] Ichiro Suzuki and Masafumi Yamashita. Searching for a mobile intruder in a polygonal region. *SIAM Journal on Computing*, 21(5):863–888, 1992, https://doi.org/10.1137/0221051.

[172] S Sachs, SM LaValle, and S Rajko. Visibility-based pursuit-evasion in an unknown planar environment. *International Journal of Robotics Research*, 23(1):3–26, JAN 2004.

[173] Leonidas J Guibas, Jean-Claude Latombe, Steven M LaValle, David Lin, and Rajeev Motwani. A visibility-based pursuit-evasion problem. *International Journal of Computational Geometry & Applications*, 9:471–493, 1999.

[174] Sang-Min Park, Jae-Ha Lee, and Kyung-Yong Chwa. Visibility-based pursuit-evasion in a polygonal region by a searcher. In Fernando Orejas, Paul G. Spirakis, and Jan van Leeuwen, editors, *Automata, Languages and Programming*, pages 456–468, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.

[175] Chao Zhang, Albert Xin Jiang, Martin B Short, P Jeffrey Brantingham, and Milind Tambe. Defending against opportunistic criminals: New game-theoretic frameworks and algorithms. In *International Conference, GameSec, Los Angeles, CA, USA*, pages 3–22. Springer, 2014.

[176] Timothy C Haas and Sam M Ferreira. Optimal patrol routes: interdicting and pursuing rhino poachers. *Police Practice and research*, 19(1):61–82, 2018.

[177] Jorge Urrutia. Art gallery and illumination problems. In *Handbook of computational geometry*, pages 973–1027. Elsevier, 2000.

[178] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is $\exists\mathbb{R}$-complete. *ACM Journal of the ACM (JACM)*, 69(1):1–70, 2021.

[179] D Lee and Arthurk Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[180] Frank Hoffmann. On the rectilinear art gallery problem. In *Automata, Languages and Programming: 17th International Colloquium Warwick University, England, July 16–20, 1990 Proceedings 17*, pages 717–728. Springer, 1990.

[181] Wei-Pang Chin and Simeon Ntafos. Optimum watchman routes. In *Proceedings of the second annual symposium on Computational geometry*, pages 24–33, 1986.

[182] Moshe Dror, Alon Efrat, Anna Lubiw, and Joseph S. B. Mitchell. Touring a sequence of polygons. In *Proc. 35th ACM Symposium on Theory of Computing*, pages 473–482, 2003.

[183] Joseph SB Mitchell. Approximating watchman routes. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 844–855. SIAM, 2013.

[184] Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. Cops, robbers, and threatening skeletons: Padded decomposition for minor-free graphs. *SIAM Journal on Computing*, 48(3):1120–1145, 2019.

[185] Béla Bollobás, Gábor Kun, and Imre Leader. Cops and robbers in a random graph. *Journal of Combinatorial Theory, Series B*, 103(2):226–236, 2013.

[186] Anthony Bonato. *The game of cops and robbers on graphs*. American Mathematical Soc., 2011.

[187] Ge na Hahn. Cops, robbers and graphs. *Tatra Mt. Math. Publ*, 36(163):163–176, 2007.

[188] Anthony Bonato, Petr Golovach, Gena Hahn, and Jan Kratochvíl. The capture time of a graph. *Discrete Mathematics*, 309(18):5588–5595, 2009.

[189] Kyle Klein and Subhash Suri. Catch me if you can: Pursuit and capture in polygonal environments with obstacles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, pages 2010–2016, 2012.

[190] Geoffrey Hollinger, Sanjiv Singh, Joseph Djugash, and Athanasios Kehagias. Efficient multi-robot search for a moving target. *The International Journal of Robotics Research*, 28(2):201–219, 2009.

[191] M. R. Garey and D. S. Johnson. " strong " np-completeness results: Motivation, examples, and implications. *J. ACM*, 25(3):499–508, jul 1978.

[192] H. Sagan. *Space-Filling Curves*. Universitext. Springer New York, 2012.

[193] Jakub Červený. Generalized space-filling curves, 2019. https://github.com/jakubcerveny/gilbert.

[194] Jian Zhang, Sei-ichiro Kamata, and Yoshifumi Ueshige. A Pseudo-hilbert Scan Algorithm for Arbitrarily-Sized Rectangle Region. *Lecture Notes in Computer Science*, pages 290–299, 2006.

[195] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[196] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.

[197] Swadhin Agrawal, Sujoy Bhore, Joseph S.B. Mitchell, P.B. Sujit, and Aayush Gohil. Supplementary materials of the article "Multi-robot searching with limited sensing range for static and mobile intruders". https://github.com/swadhinagrawal/Intruder_searching.

[198] Ana Paula Tomás and António Leslie Bajuelos. Generating random orthogonal polygons. In Ricardo Conejo, Maite Urretavizcaya, and José-Luis Pérez-de-la Cruz, editors, *Current Topics in Artificial Intelligence*, pages 364–373, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

[199] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.

[200] Andreagiovanni Reina, Raina Zakir, Giulia De Masi, and Eliseo Ferrante. Cross-inhibition leads to group consensus despite the presence of strongly opinionated minorities and asocial behaviour. *Communications Physics*, 6:236, 2023.

[201] Andreas Scheidler, Arne Brutschy, Eliseo Ferrante, and Marco Dorigo. k-unanimity rule for decision making in robot swarms. *IEEE Transactions on Cybernetics*, 46(5):1175–1188, 2016.

[202] Michael Crosscombe and Jonathan Lawry. Collective preference learning in the best-of-n problem: From best-of-n to ranking n. *Swarm Intelligence*, 15:145–170, 2021.

[203] Bruno L. Mendívez Vásquez and Jan Carlo Barca. Adversarial scenarios for herding uavs and counter-swarm techniques. *Robotica*, 41(5):1436–1451, 2023.

[204] Luderic Van Calck, Alexandre Pacheco, Volker Strobel, Marco Dorigo, and Andreagiovanni Reina. A blockchain-based information market to incentivise cooperation in swarms of self-interested robots. *Scientific Reports*, 13:20417, 2023.

[205] Gustavo A. Cardona, Kevin Leahy, and Cristian-Ioan Vasile. Temporal logic swarm control with splitting and merging. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12423–12429, Piscataway, NJ, 2023.

[206] Ariana Strandburg-Peshkin, Danai Papageorgiou, Margaret C. Crofoot, and Damien R. Farine. Inferring influence and leadership in moving animal groups. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1746):20170006, 2018.