FACULTY OF ENGINEERING AND ARCHITECTURE

Designing Swarm Intelligence for the Conservation of Animal Collectives: A Step towards Animal-Robot Societies

Stef Van Havermaet

Doctoral dissertation submitted to obtain the academic degree of
Doctor of Computer Science Engineering

Supervisors
Prof. Pieter Simoens, PhD - Prof. Yara Khaluf, PhD
Department of Information Technology
Faculty of Engineering and Architecture, Ghent University

June 2025

GHENT UNIVERSITY

ISBN 978-94-93464-01-8
NUR 980
Wettelijk depot: D/2025/10.500/61

Members of the Examination Board

Chair

Prof. Filip De Turck, PhD, Ghent University

Other members entitled to vote

Prof. Tony Belpaeme, PhD, Ghent University
Prof. Peter Goethals, PhD, Ghent University
Mary Katherine Heinrich, PhD, Université libre de Bruxelles
Prof. Tim Landgraf, PhD, Freie Universität Berlin, Germany
Andreagiovanni Reina, PhD, Universität Konstanz, Germany

Supervisors

Prof. Pieter Simoens, PhD, Ghent University
Prof. Yara Khaluf, PhD, Ghent University

Acknowledgments

Although these are the final words I have written for this dissertation, they appear first — and for good reason. Every meaningful journey, whether in life or in stories, is shaped not just by the person who walks it, but by those who walk alongside them. My PhD journey has been no different. The pages that follow bear my name, but they are also the product of the support, encouragement, and generosity of many others. To them, I owe my deepest gratitude.

First and foremost, I owe a special debt of gratitude to my supervisor, Prof. Yara Khaluf. There is no one more fitting to acknowledge first. After all, had you not reached out to me (tenaciously) to collaborate on a paper after I had already left university and moved into industry, I most likely would never have started this PhD journey. I believe it has now been seven years since we first met, and for six of those, you have guided me not only academically, but also supported me personally. I am determined that it will not be the last we see of each other, but I want you to know how truly grateful I am for all the time, effort, wisdom, and kindness you have shared with me so far.

My thanks also go to my other supervisor, Prof. Pieter Simoens. Your support in allowing me to follow my own research interests meant a great deal to me and made the journey far more enjoyable. You gave me the space to pursue a topic I was passionate about, but you also helped ground it within the expectations of a PhD, something I found genuinely helpful. I appreciated your openness to discussions, even the impromptu ones over your coffee and my protein shake, which often led to insightful ideas.

I would also like to thank in advance my jury members: Prof. Filip De Turck, Prof. Tony Belpaeme, Prof. Tim Landgraf, Dr. Mary Katherine Heinrich, Prof. Peter Goethals, and Dr. Andreagiovanni Reina. Thank you for dedicating your time and expertise to reviewing my dissertation. I hope my work offers you some new insights, and I look forward to talking about these with you.

From the jury members, I want to express my further gratitude to Prof. Tim Landgraf. In fact, it was your work that I got to know through Yara, which introduced me to the world of robots and animals. Beyond being a jury member, you have also been a co-author, and your feedback has always been invaluable in enhancing the quality of my work. Even more importantly, I want to thank you for your generosity in hosting my research visit at your lab. Without it, a significant part of this dissertation would not have been possible.

As another co-author, my sincere thanks go to Prof. Deni Mazrekaj. I doubt either of us could have predicted that we would end up writing a paper together when we first met on a video game where we quickly became friends over our shared passion for large Egyptian birds. I'm also grateful for your constant availability and valuable advice on navigating

the PhD journey and academia in general.

Before starting my PhD, I already met and collaborated with my first future colleague: Johannes Nauta. Thank you for all the random talks, the many ideas we bounced back and forth, and the support you have offered throughout the years — even on the final days, helping me sort out the numbering in this dissertation.

When I first joined the office, I met one of my colleagues: Inês Terrucha. From the beginning, I had a feeling we would become best friends, as there was no topic we could not laugh about or complain about (which usually ended in laughter as well). You made life both inside the office (not only by keeping the room cool) and outside of it much more enjoyable. Thank you for passing along your thesis template (and Johannes's), and for generously sharing your supervisor-like advice when I needed it.

I would also like to thank my colleague and friend Amrapali, who wholeheartedly joins in my delusions. From my office, I also thank all my other colleagues who, over the years, have offered both welcome distractions and genuinely helpful tips. From IDLab, I would like to express my gratitude in particular to Sai Roberts. You made it possible for me to complete my PhD by rescuing my laptop from failure on several occasions. I also extend my thanks to everyone at the Landgraf Lab, and especially to Andreas Gerken, for being such a welcoming host and collaborator. You not only made my research visits in Berlin more productive, but also much more enjoyable.

Away from work, there are also some amazing people who have made this journey truly worthwhile and for whom I want to express my gratitude. Thanks to the Surmontre crew: Wouter, Tresor, Jasper, Nick, and Hendrik. Thanks for board games, video games, and the laughs, especially during the isolation of the COVID days. Thanks to Paula, Alessandra, Francesca, and Umberto for many days and nights in Ghent and beyond. Thanks to Tony for his generosity and friendship. From helping me put up a light, to late night gym sessions, casino trips, and more. Thanks to Stephen for all the kikis, and for reminding me to finish the PhD every single time we met.

Thanks to my roommate Celia in Brussels for all the support, care, consideration, talks, and laughs when I was writing this dissertation. I could not have asked for a better person to live with. Thanks to Floran for designing a prototype of the dissertation book cover, the long calls while working, and spending almost every weekend together in the last year. Thanks to Sam for listening and always being able to cheer me up for every problem I have come to face, for all the biscuits and for all the crazy things we have done. Like you always say; it feels like I have known you forever. A special thanks goes to Mauricio. At this point, it feels like there is nothing we do not know about each other (even though there are a few things I wish I never knew). Thank you for being there literally every day. Thank you for all the time we have spent together, your friendship means a lot to me.

Finally, a deep gratitude goes to my family. Thank you to my brother and sister-in-law for your kindness and for always going the extra mile to support me. Thank you to my father for never judging me in my adult life, for being able to fix anything, and for always being ready to help. And last but not least, thank you to my mother. Your kindness and love are immeasurable, and the way you take care of not only me but so many others around you has always been truly inspiring. I do not need to say that I hope I made you proud, because I would already know your answer. So instead, let me say that you make me proud to be your son.

Stef (March 2025)

Table of contents

List of Figures

List of Tables

Summary

The loss of biodiversity is of great concern to all life on Earth, including humans, as it threatens food security, increases the spread of zoonotic diseases, diminishes essential ecosystem services, and so on. To this end, nature conservation tasks have been undertaken, with aim to protect and restore ecosystems around the world. Among these tasks is exploration, which involves documenting both known and unknown biodiversity, as a significant portion of Earth's biodiversity has yet to be described, with many species potentially facing threats before they are even identified. Intervention tasks such as the protection of endangered species are equally important. Current nature conservation interventions often rely heavily on human involvement or conventional robotic technologies such as Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs), both of which face significant limitations. Human-led efforts are constrained by accessibility, cost, and the potential to inadvertently harm ecosystems, while ROVs and AUVs can struggle in complex environments, may still disturb wildlife, and lack the capabilities for close animal interactions. To address these issues, researchers are increasingly turning to biomimetic robots, which are designed to imitate the structure, behavior, or function of animals. Swarms of these bio-inspired robots offer a promising, sustainable, and minimally invasive approach to conservation, capable of accessing remote ecosystems, monitoring wildlife without disruption, and even interacting with animals in ecologically beneficial ways.

This thesis presents a series of algorithms developed for swarms of biomimetic robots to support animal collectives through exploration and protection tasks, using fish as a case study. From a research perspective, fish are among the most commonly studied organisms in animal-robot interactions, with recent years witnessing significant theoretical and practical advancements in the development of bio-inspired fish robots. From an ecological perspective, fish play a crucial role in aquatic ecosystems but are increasingly threatened by anthropogenic pressures such as illegal fishing, pollution, and invasive species, highlighting the need for conservation efforts. To develop these algorithms, we draw upon the paradigm of swarm intelligence, in which collective behavior emerges from individuals operating based on simple behavioral rules derived from local interactions with their neighbors and their environment. This decentralized approach allows swarm intelligence algorithms to address complex challenges while

maintaining key properties essential for conservation tasks, including robustness to individual robot failures or environmental disturbances, scalability, and adaptability to changing environmental conditions or animal behaviors.

First, we investigate how swarm robots can explore and navigate like fish within dynamic environments. We identify three essential collective structures for effective exploration and navigation, characterized by the swarm's spatial coverage and the degree of alignment among individuals. The ability of traditional collective motion models of fish schooling to generate these structures is assessed by varying the width of behavioral zones. Two behavioral zonal models are considered, differing in how influential neighbors are selected: metric-based and topological-based approaches. Our results indicate that these models are unable to produce the full range of required structures. To address this limitation, we propose an extension of the metric-based zonal model. By varying both the alignment zone width and converting the alignment zone to a neutral zone, all three structures can be generated. Furthermore, we show that a robot swarm following the proposed model can dynamically transition between these structures in dynamic environments, where different stimuli appear over time and are perceived by few individuals.

Second, we consider how swarm robots can support fish schools in protecting them from dynamic dangers. To preemptively detect threats before they interact with the fish, and without prior knowledge of their presence, the robots are tasked with constructing a caging formation around the school. This approach relies on a certain level of school cohesion. Therefore, we first examine the conditions under which different collective motion models maintain this cohesion. Since changes in cohesion affect spatial coverage, we also measure the resulting coverage in each case and computationally determine the minimum number of robots required to successfully construct a caging formation. The results show that the proposed algorithm enables robots to effectively cage the school when the constraints of school cohesion and robot swarm size are met. To guide fish away from danger, robots adaptively switch between two states based on their local position relative to the desired direction: one that induces a repulsive effect on nearby fish and one that remains neutral. The effectiveness of this approach is tested in dynamic environments across two tasks: (i) avoiding a stationary danger that appears stochastically, and (ii) staying within a shrinking circular safe zone. The findings demonstrate that the robot swarm successfully prevents any fish from being affected by danger.

Third, we consider a more complex guidance problem in which the robot swarm must navigate the school along a time-varying safe path while ensuring the fish remain within its boundaries. Although this constraint imposes stricter movement requirements, it simultaneously eliminates the need for consensus decision-making among robots, which was essential in the previous approach focused on avoiding dynamic dangers. Instead,

each robot derives its guidance from an artificial potential field that encodes path information, enabling the swarm to guide the fish without requiring inter-robot communication. To validate the algorithm, paths with varying turn sharpness and space between boundaries are considered, with effectiveness measured as the percentage of fish remaining within the boundaries until reaching the target location. The results show that performance tends to decline in sharp turns, particularly when these turns occur within narrow segments of the path. This decline is influenced by multiple factors, including the width of the margins around the path, the steepness of margin narrowing, the number of robots relative to the number of fish, and the maximum velocity of the fish. At higher speeds, the robots struggle to maintain a coherent caging formation, which makes it more difficult to contain the group during abrupt directional changes.

Fourth, we transition from model-based assumptions to experimental validation, investigating how robot behavior influences fish avoidance responses. The previously developed protection algorithms were based on theoretical models of fish-robot interactions; however, before these can be tested in real-world scenarios, it is crucial to first assess the validity of these underlying assumptions. To this end, a conspecific-like robot repeatedly approached individual guppies, varying its speed across trials within a medium to high range. Our findings suggest that as fish experienced repeated encounters with the robot, their likelihood of exhibiting avoidance responses increased, indicating a form of behavioral adaptation. In some instances, they displayed typical anti-predator reactions, such as freezing or executing rapid escape maneuvers, suggesting that the robot was perceived as a potential threat. In particular, escape responses were not only dictated by proximity, but were also influenced by the robot's approach speed, challenging conventional models that assume distance alone determines avoidance behavior. Moreover, avoidance speed increased with higher robot speeds and decreased proximity, highlighting the dynamic nature of fish responses.

The work presented in this dissertation advances our understanding of how swarm robotics can support animal collectives through exploration and protection tasks. By integrating principles of swarm intelligence, we demonstrate how decentralized robotic systems can navigate dynamic environments, safeguard fish schools from dangers, and guide them along safe paths. Furthermore, by experimentally validating fish avoidance responses to robotic stimuli, we provide empirical insights that refine existing models of fish-robot interactions. These contributions lay the foundation for future research in animal-robot mixed societies and highlight the potential of swarm robotics as a non-invasive, scalable tool for conservation efforts.

Samenvatting
– Summary in Dutch –

Het verlies aan biodiversiteit is een grote zorg voor al het leven op aarde, inclusief de mens, omdat het de voedselzekerheid bedreigt, de verspreiding van zoönosen vergroot, essentiële ecosysteemdiensten vermindert, enzovoort. Daarom zijn er natuurbehoudstaken ondernomen om ecosystemen over de hele wereld te beschermen en te herstellen. Een van deze taken is exploratie, waarbij zowel bekende als onbekende biodiversiteit wordt gedocumenteerd, aangezien een aanzienlijk deel van de biodiversiteit op aarde nog moet worden beschreven, waarbij veel soorten mogelijk worden bedreigd voordat ze zelfs maar zijn geïdentificeerd. Interventietaken zoals de bescherming van bedreigde soorten zijn net zo belangrijk. De huidige interventies voor natuurbehoud zijn vaak sterk afhankelijk van menselijke betrokkenheid of conventionele robottechnologieën zoals ROV's (Remotely Operated Vehicles) en AUV's (Autonomous Underwater Vehicles), die beide aanzienlijke beperkingen hebben. Menselijke inspanningen worden beperkt door toegankelijkheid, kosten en de mogelijkheid om ecosystemen onbedoeld te beschadigen, terwijl ROV's en AUV's het moeilijk hebben in complexe omgevingen en nog steeds wilde dieren kunnen verstoren. Om deze problemen aan te pakken, richten onderzoekers zich steeds meer op biomimetische robots, die ontworpen zijn om de structuur, het gedrag of de functie van dieren te imiteren. Zwermen van deze bio-geïnspireerde robots bieden een veelbelovende, duurzame en minimaal invasieve benadering van natuurbehoud, die in staat is om toegang te krijgen tot afgelegen ecosystemen, om wilde dieren te monitoren zonder ze te verstoren en zelfs om met dieren te communiceren op ecologisch gunstige manieren.

Deze dissertatie presenteert een reeks algoritmen die zijn ontwikkeld voor zwermen biomimetische robots om dierencollectieven te ondersteunen bij exploratie- en beschermingstaken, met vissen als casestudy. Vanuit een onderzoeksperspectief behoren vissen tot de meest bestudeerde organismen in dier-robot interacties, waarbij de laatste jaren aanzienlijke theoretische en praktische vooruitgang is geboekt in de ontwikkeling van bio-geïnspireerde visrobots. Vanuit ecologisch perspectief spelen vissen een cruciale rol in aquatische ecosystemen, maar worden ze steeds meer bedreigd door antropogene druk zoals illegale visvangst, vervuiling en

invasieve soorten. Om deze algoritmen te ontwikkelen, maken we gebruik van het paradigma van zwermintelligentie, waarin collectief gedrag ontstaat uit individuen die werken op basis van eenvoudige gedragsregels die zijn afgeleid van lokale interacties met hun buren en hun omgeving. Met deze gedecentraliseerde aanpak kunnen zwermintelligentie-algoritmen complexe uitdagingen aanpakken met behoud van belangrijke eigenschappen die essentieel zijn voor natuurbehoudstaken, zoals robuustheid tegen individuele robotstoringen of omgevingsverstoringen, schaalbaarheid en aanpassingsvermogen aan veranderende omgevingsomstandigheden of gedrag van dieren.

Eerst onderzoeken we hoe zwermrobots als vissen kunnen verkennen en navigeren in dynamische omgevingen. We identificeren drie essentiële collectieve structuren voor effectieve exploratie en navigatie, gekarakteriseerd door de ruimtelijke dekking van de zwerm en de mate van afstemming tussen individuen. Het vermogen van traditionele collectieve bewegingsmodellen van scholende vissen om deze structuren te genereren wordt beoordeeld door de breedte van gedragszones te variëren. Twee gedragszonale modellen worden bekeken, die verschillen in hoe invloedrijke buren worden geselecteerd: metrisch-gebaseerde en topologisch-gebaseerde benaderingen. Onze resultaten geven aan dat deze modellen niet in staat zijn om alle vereiste structuren te produceren. Om deze beperking aan te pakken, stellen we een uitbreiding van het metrisch-gebaseerde zonemodel voor. Door zowel de breedte van de uitlijningszone te variëren als de uitlijningszone te converteren naar een neutrale zone, kunnen alle drie de structuren worden gegenereerd. Verder laten we zien dat een robotzwerm die het voorgestelde model volgt, dynamisch kan wisselen tussen deze structuren in dynamische omgevingen, waar verschillende stimuli in de loop van de tijd verschijnen en door weinig individuen worden waargenomen.

Ten tweede onderzoeken we hoe zwermrobots visscholen kunnen beschermen tegen dynamische gevaren. Om bedreigingen preventief te detecteren voordat ze interactie hebben met de vissen, en zonder voorkennis van hun aanwezigheid, worden de robots belast met het bouwen van een kooiformatie rond de school. Deze aanpak is afhankelijk van een bepaald niveau van schoolcohesie. Daarom onderzoeken we eerst de voorwaarden waaronder verschillende collectieve bewegingsmodellen deze cohesie behouden. Omdat veranderingen in de cohesie van invloed zijn op de ruimtelijke dekking, meten we ook de resulterende dekking in elk geval en bepalen we rekenkundig het minimale aantal robots dat nodig is om met succes een kooiformatie te bouwen. De resultaten laten zien dat het voorgestelde algoritme robots in staat stelt om de school effectief in te kooien wanneer aan de beperkingen van de schoolcohesie en de grootte van de robotzwerm wordt voldaan. Om de vissen weg te leiden van het gevaar, schakelen de robots adaptief tussen twee toestanden op basis van hun lokale positie ten opzichte van de gewenste richting: een toestand die een

afstotend effect heeft op nabijgelegen vissen en een toestand die neutraal blijft. De effectiviteit van deze aanpak wordt getest in dynamische omgevingen voor twee taken: (i) het vermijden van een stationair gevaar dat stochastisch verschijnt, en (ii) binnen een krimpende cirkelvormige veilige zone blijven. De bevindingen tonen aan dat de robotzwerm met succes voorkomt dat een vis getroffen wordt door gevaar.

Ten derde beschouwen we een complexer geleidingsprobleem waarbij de robotzwerm de school langs een in de tijd variërend veilig pad moet navigeren en er tegelijkertijd voor moet zorgen dat de vissen binnen de grenzen blijven. Hoewel deze beperking strengere bewegingsvereisten oplegt, elimineert het tegelijkertijd de noodzaak voor besluitvorming op basis van consensus tussen robots, wat essentieel was in de vorige benadering die gericht was op het vermijden van dynamische gevaren. In plaats daarvan ontleent elke robot zijn begeleiding aan een kunstmatig potentiaalveld dat padinformatie codeert, waardoor de zwerm de vis kan leiden zonder dat communicatie tussen de robots nodig is. Om het algoritme te valideren, worden paden met variërende bochtscherpte en ruimte tussen de grenzen bekeken, waarbij de effectiviteit wordt gemeten als het percentage vissen dat binnen de grenzen blijft tot de doellocatie is bereikt. De resultaten tonen aan dat de prestaties afnemen bij scherpe bochten, vooral wanneer deze bochten voorkomen in smalle segmenten van het pad. Deze achteruitgang wordt beïnvloed door meerdere factoren, waaronder de breedte van de marges rond het pad, de steilte van de vernauwing van de marges, het aantal robots ten opzichte van het aantal vissen en de maximale snelheid van de vissen. Bij hogere snelheden hebben de robots moeite om een coherente kooivorming te behouden, wat het moeilijker maakt om de groep in bedwang te houden tijdens abrupte richtingsveranderingen.

Ten vierde gaan we over van modelgebaseerde aannames naar experimentele validatie, waarbij we onderzoeken hoe het gedrag van de robots de vermijdingsreacties van de vissen beïnvloedt. De eerder ontwikkelde beschermingsalgoritmen waren gebaseerd op theoretische modellen van vis-robot interacties; voordat deze echter getest kunnen worden in echte scenario's, is het cruciaal om eerst de geldigheid van deze onderliggende aannames te beoordelen. Daartoe benaderde een robot, die op een soortgenoot leek, herhaaldelijk individuele guppy's, waarbij de snelheid tijdens de proeven varieerde binnen een gemiddeld tot hoog bereik. Onze bevindingen suggereren dat naarmate de vissen de robot meer en meer ontmoetten, de kans toenam dat ze vermijdingsreacties vertoonden, wat wijst op een vorm van gedragsaanpassing. In sommige gevallen vertoonden ze typische anti-predator reacties, zoals bevriezen of het uitvoeren van snelle ontsnappingsmanoeuvres, wat suggereert dat de robot als een potentiële bedreiging werd gezien. Vluchtreacties werden niet alleen bepaald door nabijheid, maar werden ook beïnvloed door de naderingssnelheid van de robot, wat conventionele modellen die ervan uitgaan dat afstand alleen bepalend is voor vermijdingsgedrag, in twijfel trekt. Bovendien nam de

vermijdingssnelheid toe met hogere robotsnelheden en verminderde nabijheid, wat de dynamische aard van visresponsen benadrukt.

Het werk gepresenteerd in dit proefschrift bevordert ons begrip van hoe zwerm robotica diercollectieven kan ondersteunen bij exploratie- en beschermingstaken. Door principes van zwermintelligentie te integreren, laten we zien hoe gedecentraliseerde robotica dynamische omgevingen kan navigeren, visscholen kan beschermen tegen gevaren en ze langs veilige paden kan leiden. Door experimenteel de vermijdingsreacties van vissen op robotstimuli te valideren, bieden we bovendien empirische inzichten die bestaande modellen van vis-robot interacties verfijnen. Deze bijdragen leggen de basis voor toekomstig onderzoek naar gemengde dier-robot samenlevingen en benadrukken het potentieel van zwermrobots als een niet-invasief, schaalbaar hulpmiddel voor natuurbehoud.

1 Introduction

"In nature, nothing exists alone. The balance of nature is not a status quo; it is fluid, ever shifting, in a constant state of adjustment. Humans, too, are part of this balance, and they are as dependent upon it as any other life form."

Rachel Carson, Silent Spring (1962)

NATURE is becoming increasingly shaped by human activities. Disruption of the *balance of nature* has lead to the extinction of many animal species, and consequently loss of biodiversity. As technology continues to advance, and given that *nothing in nature exists alone*, we should strive to develop technologies that not only coexist with nature but actively contribute to restoring the balance *on which we all depend*. The development of animal-robot mixed societies is not aimed at replacing natural organisms with machines, but rather fostering a symbiotic relationship between them. In such societies, robots integrate seamlessly with animal groups, assisting them in tasks like navigation, protection, and adaptation to changing environments. This thesis focuses on fish schools as a case study, proposing algorithmic solutions for swarm robots to support fish in their daily lives. The first objective is to explore how robots can explore and navigate like fish in the dynamic environments that fish inhabit. The second objective is to investigate how robots can guide fish away from dangers and along safe trajectories, offering protection in critical situations. Finally, the third objective is to experimentally examine how robot behavior influences fish avoidance responses, providing insights into how robots can leverage these behaviors to steer fish movement.

1.1 Towards animal-robot mixed societies

Drawing inspiration from biological systems has led to significant advancements in a wide range of domains, including computer science and robotics. For instance, state-of-the-art artificial intelligence models, such as ChatGPT, are based on artificial neural networks, which are computational models inspired by the structure and function of the neural architecture of the human brain [1]. In robotics, bioinspired design has facilitated the development of advanced systems capable of complex autonomous behaviors. This approach laid the foundation for the field of "biorobotics", which focuses on creating robots that mimic the structure, function, or behavior of biological organisms. Robots designed in this way are also known as "biomimetic" or bioinspired robots (see Figure 1.1). A notable example is SoFi, a soft robotic fish developed by MIT roboticists, designed to resemble and move like a real fish [2]. Its design draws inspiration from the flexible body structure and swimming dynamics of fish, particularly the tail propulsion used for efficient aquatic movement. In addition to imitating the anatomy and biomechanics of living organisms, scientists and engineers have drawn inspiration from the social and collective behaviors observed in nature. For example, swarms of robots have been developed to mimic the schooling behavior of fish and the flocking patterns of birds [3]. By designing the intelligence of individual robots based on the simple behavioral principles that govern interactions in these animals, the robots appear to function as a single collective unit, even in the absence of centralized control. This approach has revolutionized swarm robotics, with applications ranging from environmental monitoring [4] to autonomous construction [5].

1.1.1 Biorobotics: from inspiration to interaction

Initially, the field of biorobotics focused on improving robotic technologies inspired by principles observed in nature. Over time, this focus has expanded to include the development of technologies that allow robots to interact directly with living animals [6]. This emerging branch of biorobotics extends beyond traditional bioinspiration, paving the way for the merging of natural and artificial systems into synergistic ecosystems [7–10]. A distinction is typically made between animal-robot technologies based on whether the interaction occurs at the individual (physiological) level or the group (behavioral) level. When natural and artificial components are integrated at the individual level, the resulting entity is referred to as a "biohybrid organism", also known as cyborgs, bio-robots, or animal-robots. In such entities, either an artificial component is incorporated into an animal or a biological organ is integrated into a robot, enabling direct interaction between the nervous or muscular system and electronic components. For instance, Kobayashi et al. [11] artificially induced swimming behavior in a goldfish (*Carassius auratus*) by stimulating its brain through an electrode.

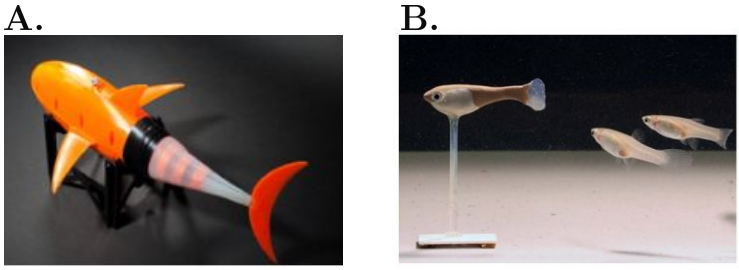

Figure 1.1: Examples of biomimetic fish robots. (A) A soft robotic fish utilizing oscillatory propulsion and a compliant tail, designed to operate without disturbing or damaging underwater flora and fauna. Adapted from Van Den Berg et al. [12]. (B) A 3D replica of a guppy, attached to a robot (not shown), interacting with live guppies, as part of an animal-robot mixed society. Reproduced with permission from Maxeiner et al. [13].

Bio-hybrid organisms combine the adaptability of animals to dynamic environments with the possibility of controlling them as robots, offering unique advantages for various potential applications, such as search-and-rescue operations [6]. Moreover, bio-hybrid organisms serve as valuable tools for advancing neuroscience and biological research.

When animal-robot interactive technologies are integrated at the group level, an "animal–robot mixed society"[1] is established [10] (see Figure 1.1B). In this context, living animals and robots interact with one another through their behaviors as separate entities. These robots are typically biomimetic [14], and can be either static or mobile [8]. Depending on the goal of the interaction, robots can mimic specific features of the animal they are interacting with to appear as a conspecific (i.e., a member of the same species) or even replicate the traits of the animal's natural predators [15]. Such robots provide numerous advantages for studying animal behavior and behavioral ecology. They are easier to handle compared to real animals and allow precise control of their position in the environment, enabling highly standardized and reproducible experimental designs. Furthermore, the ability to deliver distinct cues through robots interacting with animals makes it possible to dissect the relative importance of various stimuli in shaping animal responses. These cues can be programmed into the robot and delivered in various ways, including visual cues based on morphophysiological features (e.g., color patterns [16], size [17], and shape [18]), spatial cues (e.g., position within the environment [19]), locomotory cues (e.g., movement dynamics [20]), as well as auditory [21], light [22], and olfactory [23] cues. For instance, Marras and Porfiri [24]

[1] In this text, we adopt the terminology and categorization of animal-robot interactive technologies proposed by Romano et al. [6]. It is worth noting that the term "society" can also encompass one-to-one animal-robot interactions.

demonstrated that locomotion serves as a determinant cue for attracting conspecifics by exposing golden shiners (*Notemigonus crysoleucas*) to different tail-beat frequencies (i.e., the rate at which the tail oscillates over time) of a robotic replica. Similarly, Landgraf et al. [25] showed that for Trinidadian guppies (*Poecilia reticulata*), the combination of realistic eye dummies and natural motion patterns significantly increased the robot's acceptance as a conspecific.

Even more intriguing is the concept of a closed-loop mixed society, where robots and living animals interact and mutually influence each other according to the principles of social behavior [14]. These mixed societies function as dynamic systems, with the actions of animals and robots being interdependent. Robots can trigger specific reactions or behaviors in animals and dynamically adjust their own behavior in response to the animals' actions [26]. In this way, robots are active participants in shaping the group's social dynamics. As such, biorobotics serves as a sophisticated tool for studying social and collective behaviors in animals.

1.1.2 Robots in social and collective animal behaviors

In closed-loop mixed societies, a biomimetic robot can be used to influence an animal's social behaviors [27] such as learning in chicks [19], courtship in lizards [28] and frogs [21], or foraging in birds [29]. Moreover, they can be used to observe and modify collective animal behaviors [30], such as decision-making in cockroaches [31], or coordinated movement in fish [32]. While biomimetic robots have been used to study social and collective behaviors in a variety of species, most research has focused on fish influenced by a robotic conspecific [6]. Fish are easier to maintain in laboratory conditions compared to other vertebrates and provide a valuable model for investigating behavioral interactions, such as the influence of swimming speed and orientation.

Kim et al. [33] found that zebrafish (*Danio rerio*) not only followed the robotic replica but also adjusted their swimming patterns to align with its movements, displaying both attraction and mimicry. This behavior was attributed to the interactive nature of the robot, which dynamically responded to the fish's movements in real time. Similarly, Maxeiner et al. [13] showed that conspecific-like robots that adjust their behavior in response to Trinidadian guppies can more effectively elicit following behavior and reduce avoidance responses. At the collective level, zebrafish have also been shown to exhibit schooling behaviors with a robotic conspecific, where the speed of the robot serves as a determinant cue for group cohesion [34]. In addition to studies involving many-fish-to-one-robot interactions, research has also examined one-fish-to-many-robots interactions. For instance, Butail et al. [35] found that zebrafish spent more time in the proximity of robots when the robots swam farther apart, compared to when they swam closer together.

In a study with golden shiners, a robotic conspecific was shown to modulate the risk-taking behavior of the fish [36]. Risk-taking was quantified by measuring the time it took for the fish to leave its shelter after a simulated predator attack. The fish appeared to mimic the robot's behavior, exhibiting increased boldness (i.e., higher risk-taking) when the robot was bold, and decreased boldness when the robot was cautious. Related to this is the concept of leadership, wherein one or more individuals initiate a new direction of locomotion and the others follow. Faria et al. [37] developed a robotic replica to initiate new swimming directions in both individuals and groups of sticklebacks (*Gasterosteus aculeatus*). In a similar context, Ward et al. [38] demonstrated that in a Y-maze, the binary choice of direction for a solitary stickleback was strongly influenced by a single robotic conspecific. However, for larger groups (four or more fish), an additional robot was required to influence the decision.

In addition to robots mimicking conspecifics, researchers have developed robotic predators to study predator-prey dynamics. These robots are designed to replicate the characteristics of the natural predators of the animals. Cord-Cruz et al. [39] found that zebrafish exhibited similar behavioral responses to both a live predator (red tiger fish; *Hydrocynus goliath*) and a mobile 3D-printed predator replica. Anti-predator behaviors included freezing (i.e., remaining stationary as a fight-or-flight response) [40], and escape movements (i.e., sudden sharp turns and rapid bursts of acceleration) [41]. These findings were further supported by Spinello et al. [42] and Ladu et al. [43], who showed that robotic replicas induced significant fear responses in zebrafish, as evidenced by the fish distancing itself from the robot.

1.1.3 Applications of animal-robot interaction

Although interactive robots are a relatively novel method for studying animal behavior, they have already proven to be highly effective tools for gaining valuable insights [27, 44, 45]. These findings can lead to a wide range of practical applications [46], such as improving animal welfare in captivity [47], optimizing animal husbandry practices [48], and developing more effective diagnostic tools and treatments for behavioral disorders [49]. Beyond their use in behavioral research, interactive biomimetic robots can also have direct practical applications involving animals, such as managing animal populations in agriculture [50–52], and wildlife conservation [53–55].

In agriculture, robots have been developed for tasks such as feeding and disease control [52]. In addition, mobile robots have been shown to effectively encourage animal movement in free-range poultry farms through interactive engagement [56]. Physical activity is crucial for the health and well-being of birds, as prolonged inactivity or sedentary behavior can cause significant health problems [57]. More interestingly, robots have also

been developed to interact socially with chicks in poultry farming [19]. Chicks hatched in incubators are separated from their mother hens, missing natural stimuli that influence their behavior and social learning [58]. Using robots as social partners to replace the mother hen offers a promising approach to promote healthy behavior in chicks [50]. Controlling animal movement is another key task in agriculture [59]. For example, by precisely managing cattle movement, farmers can optimize pasture land use, thereby reducing material and labor costs while avoiding human-livestock conflict [60]. Instead of relying on conventional fencing, a collar-mounted robotic device controls the animal's movement using stimuli such as sound or low electrical currents [61].

A more specific example of animal movement control is "shepherding", which involves guiding a group of animals to a desired location (e.g., water, pasture, shelter) while maintaining the group's cohesion [62]. In a broader sense, the shepherding problem can be defined as agents ("shepherds") guiding a group of other autonomous agents ("herd")[2]. This concept is applicable beyond agriculture, where agents could represent animals, humans, or robots. Indeed, robots have been successfully demonstrated in herding animals such as sheep [63–65], cattle [66], and ducks [67], as well as in other contexts such as keeping birds away from aircraft [68, 69], managing human crowds [70, 71], and facilitating the safe movement of people [72, 73]. In most works, the intelligence of these robots is inspired by natural shepherds, namely sheepdogs. Sheepdogs are particularly skilled at herding, with a single dog capable of managing more than 80 sheep [74]. These dogs exploit the sheep's collective behavior of aggregating and escaping from threats [75], a phenomenon widely recognized as the selfish herd theory [76]. Similarly, biomimetic robots that trigger an aversive response in the herd by delivering specific cues can rely on the same shepherding mechanisms as sheepdogs [77].

At the intersection of agricultural improvement and wildlife conservation, sensor systems have been developed for real-time monitoring of honey bees (*Apis mellifera*) [78]. Although these insects are vital to the natural ecosystem, a widespread decline in bee colonies has been reported in recent years [79]. Honey bees not only produce honey and other industrial products, but they are also fundamental to plant pollination, playing a key role in the proliferation of both wild and cultivated flora [80]. Thus, monitoring the health of bee colonies in real-time is crucial, as it allows for timely interventions to support their survival. Alongside sensor systems, interactive robots have been employed to benefit both agriculture and wildlife. Conflicts arising from wildlife preying on farm animals and damaging crops result in both agricultural losses and declines in wildlife populations. Therefore, balancing the protection of livestock and wildlife

[2]In literature, the term "group," "herd," "flock," or "school" is used depending on the species under study. Here, "herd" is used, as the shepherding problem is traditionally inspired by dogs herding sheep.

in shared landscapes is essential. Breck et al. [54] demonstrated that adaptive mobile robots significantly improved livestock survival rates without harming wildlife. Their findings highlight how interactive robots can enhance agricultural protection and support the development of nonlethal wildlife management tools. Non-invasive monitoring of animal health and non-lethal human-animal conflict resolution are among the crucial tools in wildlife conservation, and robots offer a promising approach to achieving this [81, 82]. In the following section, we explore in greater detail how specifically biomimetic robots can contribute to nature conservation.

1.2 Robots inspired by nature, for nature

Humanity has reached a critical point in its relationship with the natural world. Activities such as unsustainable resource extraction, large-scale development, and widespread pollution are severely damaging the health of the planet and jeopardizing the survival of its ecosystems [83]. Over the past five decades, we have witnessed a dramatic loss in biodiversity, with nearly 70% of the planet's species lost [84]. Biodiversity forms the foundation for all life on Earth, including humans, and its loss is therefore of great concern to us [85–90]. For instance, biodiversity loss heightens the risk of zoonotic diseases (e.g., COVID-19 [91–93]), diminishes access to critical medicinal resources (many of which are derived from plants and animals) [94–96], and weakens the human immune system by destabilizing ecosystems [97–99]. Additionally, it threatens food security by reducing crop and pollinator diversity, which directly impacts agricultural productivity [100–102]. The decline also undermines natural defenses against environmental disasters, such as floods and droughts [103–105]. Furthermore, biodiversity plays a vital role in climate regulation by capturing and storing carbon, and its loss accelerates climate change, affecting human livelihoods globally [106–108]. Given the critical role biodiversity plays in the sustainability of life on Earth, it is imperative that we take action to preserve it and repair the damage it has sustained. This reflects the core principle of nature conservation.

1.2.1 The conservation of nature

Nature conservation efforts are essential to protect and restore ecosystems around the world and to prevent a sixth mass extinction event [109, 110]. These efforts can be categorized into exploration, data collection, monitoring, intervention, and maintenance tasks [55]. Exploration tasks involve documenting both known and unknown biodiversity, including previously unexplored environments. Data collection focuses on gathering information to address specific questions concerning the ecosystem. Monitoring includes regular surveillance to assess ecosystem health, such as population trends, presence of invasive species, and human-wildlife inter-

actions. Intervention tasks aim to promote ecosystem recovery and improve its health, such as resolving human-wildlife conflicts, protecting endangered species, and removing invasive species. Maintenance tasks focus on preserving the ecosystem's natural state and ensuring the sustainability of ecosystem services, such as protecting against natural disasters and mitigating pollution.

Each of these tasks is essential to maintain and restore biodiversity, and a lot of work is still to be done [111]. For example, studies suggest that a significant portion of Earth's biodiversity has yet to be described, with many species potentially facing threats before they are even identified [112, 113]. Despite significant progress, substantial knowledge gaps persist in ecosystem data collection. The IPBES has identified numerous areas where data is lacking, hindering effective conservation strategies [114]. Intervention tasks such as the protection of endangered species are equally important, as recent findings suggest that 12.7% of marine fish species are threatened, as defined by the International Union for the Conservation of Nature (IUCN) [115]. Given the United Nations Convention on Biological Diversity's target to conserve 30% of both land and marine areas by 2030 [116], there is an urgent need for innovative and efficient approaches to conservation. As we have long drawn inspiration from nature to enhance our technologies, it may now be time to return the favor by using these innovations to help preserve the natural world. Among these technologies, biomimetic robots stand out as a promising candidate, with the potential to become effective, sustainable, and non-invasive tools for nature conservation.

1.2.2 The role of biomimetic robots

Current nature conservation efforts rely heavily on human involvement, which introduces several limitations, as highlighted by Chellapurath et al. [55]. Many biodiversity-rich areas are inaccessible to humans and therefore remain unexplored, including the deep ocean [117], underground cave systems [118], polar regions [119], and high-altitude areas [120]. Additionally, traditional monitoring methods, which often use stationary devices or periodic sampling, can miss critical information [121]. While human-driven conservation actions aim to protect ecosystems, they can inadvertently cause harm, contributing to pollution, disturbing wildlife behavior, and causing injury or stress to animals [122–127]. Robots, such as UAVs, AUVs, and ROVs, have addressed some of these challenges by enabling surveys in previously inaccessible or risky areas [128–130]. However, these technologies still face limitations, such as difficulty navigating uneven terrain, disturbing wildlife, and lacking the capabilities for close animal interactions or intervention tasks [131–133].

Biomimetic robots offer a promising solution to these challenges. By mimicking the structure or function of animals, they can access ecosystems

that were previously difficult or impossible to explore [134]. As previously discussed, research on the interactions between biomimetic robots and animals has provided valuable insights into how these robots can engage with animals up close. This includes influencing their behavior in beneficial ways, such as guiding them to move in a different direction, or observing them without causing stress or altering their natural behavior in undesirable ways [135]. So far, a number of biomimetic robots have been designed and demonstrated for a variety of conservation tasks. For instance, drones mimicking natural predators, such as the peregrine falcon (*Falco peregrinus*), have been developed to reduce bird visits in areas like airports, landfills, and farms, without causing harm to the birds [68, 136, 137]. Promising biomimetic robots have also been developed for aquatic tasks. Undulatory fish-inspired robots, such as MIT's RoboTuna [138], demonstrate superior agility and maneuverability compared to propeller-driven systems in complex or unstable flows [139, 140]. Their ability to generate stable propulsion in turbulent conditions and perform tight turns with minimal environmental disturbance makes them well-suited for navigating complex underwater terrains such as coral reefs [141, 142]. Ocean One, a humanoid robot, is designed for collecting and manipulating objects in deep-sea environments, ranging from delicate coral pieces to heavy frames and tools [143–145]. Another underwater robot, inspired by crab locomotion, is able to walk and run more easily on the seabed while collecting samples and monitoring marine habitats [146]. Similarly, the LAURON robot, inspired by the stick insect (*Carausius morosus*), showcases the potential of bio-inspired robots to navigate rough and hazardous terrains, such as volcanic sites [147, 148].

Beyond these existing applications, Chellapurath et al. [55] highlight several promising future uses of biomimetic robots in nature conservation. One such application is the potential for robots to train captive animals before their reintroduction into the wild, providing a safer and more effective transition. In addition, mixed animal-robot societies offer a unique approach to conservation. By influencing the social behavior of animals, robots can ultimately impact the ecosystems in which these societies are embedded. Known as "ecosystem hacking", this approach has the potential to enhance ecosystem stability or, at least slow the rate of ecosystem degradation [149–151].

1.2.3 Operational considerations

When developing conservation robots, researchers must consider several operational factors, including durability, energy efficiency, and biodegradability [55]. Given that these robots are often deployed in hazardous environments, they need to be highly durable to withstand challenging conditions. The limited capacity of current energy storage systems can pose a challenge for conservation tasks. However, harvesting energy from re-

newable sources or receiving power wirelessly are promising solutions that also help reduce electronic waste [152, 153]. For example, Mohamed et al. [154] demonstrated the use of phase change materials to convert temperature differences in the ocean into usable energy, highlighting the potential for long-endurance underwater robots. In addition, the robot can become damaged, causing parts to fall off or rendering it inoperable, which would make retrieval difficult. To minimize environmental impact, it would be ideal for these robots to have a high degree of biodegradability, as they will primarily operate in fragile ecosystems. Encouragingly, sustainable materials are gaining traction in robotics, with efforts targeting all subsystems, from actuators and energy storage to electronics [155, 156]. For instance, Donatelli et al. [157] constructed a fish-like robot body using silk hydrogel embedded with fibers to mimic the structure of natural fish skin.

Environmental constraints, particularly in underwater contexts, impose additional design challenges that must be addressed alongside conservation-related objectives. In underwater contexts, sensing, actuation, and communication each present unique challenges that require specialized approaches. Fortunately, substantial progress has been made in the theoretical and practical development of underwater robotic systems [158]. Underwater sensing is complicated by poor light penetration, turbidity, signal attenuation, and limited communication bandwidth. Despite these challenges, advances in sensing technologies are enabling robots to navigate complex terrains, localize themselves, detect objects, and monitor environmental parameters. These capabilities are supported by the integration of sonar, stereo and monocular cameras, inertial measurement units (IMUs), and environmental sensors measuring temperature, pressure, and salinity [159]. To improve close-range sensing (typically less than one meter), researchers are increasingly exploring biomimetic designs, such as artificial lateral lines [160–162] and whisker-like sensors [163, 164], inspired by the sensory systems of aquatic animals.

In terms of locomotion, a particularly promising class of underwater robots are soft robots. Constructed from materials such as polymers, hydrogels, or liquid metals, these robots can bend, stretch, and adapt their shape for safe and flexible movement [165, 166]. Their soft and deformable materials allow them to squeeze through narrow spaces, avoid obstacles, and adapt to complex underwater terrains, which is more difficult for traditional rigid robots [167]. In addition, their compliant structures allow for more accurate imitation of aquatic animal movements [168]. Various actuation mechanisms have been developed to support different modes of underwater locomotion in bioinspired soft robots [169]. Examples include a fish-inspired robot using oscillatory tail motion [170], a stingray-like robot that uses undulatory swimming [171], a jellyfish-inspired robot with silent and untethered propulsion [172], and a starfish-like robot that uses crawling movements [173]. While soft robots excel in adaptability

and safe interaction, they often lack the structural rigidity, actuation precision, and force output of their rigid counterparts. These limitations can restrict their effectiveness in tasks that require high force or precise manipulation [169]. To address this, hybrid designs combining soft and rigid components are emerging, offering a promising balance between flexibility and control [174]. Ultimately, the optimal design depends on the specific requirements of the intended conservation task.

Beyond the previously mentioned considerations, researchers should also evaluate the use of multiple robots versus a single robot. Multi-robot systems offer significant advantages for many conservation tasks and can be classified into two types based on their control structure: centralized and decentralized. In a centralized system, a central unit manages communication with all robots, directing their actions. In contrast, decentralized systems rely on local observations and communication, enabling each robot to make independent decisions. Each architecture presents distinct strengths and limitations. Centralized systems are well-suited for achieving global coordination and task optimization but can become vulnerable to single points of failure unless additional fault-tolerant mechanisms are implemented. In comparison, decentralized systems inherently offer greater robustness, scalability, adaptability, and fault tolerance, as their distributed nature allows them to more easily adapt to changes in the number of robots. Robustness refers to the system's ability to maintain core functions despite challenges or environmental changes, while fault tolerance ensures that the system can continue operating even if individual robots fail. These qualities are crucial for many conservation tasks, where robots must function in hazardous or dynamic environments [55, 175]. Moreover, depending on the task, robots must be capable of exhibiting collective behavior to perform actions beyond the capability of a single robot. For example, we can enable robots to explore in a coordinated manner, similar to how fish move in a school as a cohesive unit, by designing intelligence inspired by the collective behavior of these animals [176, 177]. This type of intelligence is also known as swarm intelligence. In the next section, we discuss the principles of swarm intelligence and explore how these principles can be applied to robotics.

1.3 Swarm intelligence

Many animal species exhibit remarkable collective behavior, where simple individuals are capable of producing highly organized and complex group dynamics [178]. Examples of such behaviors include ant colonies efficiently foraging for food [179], flocks of birds synchronizing their flight patterns [180], and schools of fish evading predators with coordinated motion [181]. These complex group behaviors are not the result of a centralized control structure or global knowledge. Instead, they emerge from local interactions between individuals and their environment, guided

by simple behavioral rules. Through these interactions, large-scale patterns and coordinated behaviors arise without external control, a process known as self-organization [182–184]. Inspired by these biological phenomena, "swarm intelligence" has emerged as an artificial intelligence paradigm that models such collective behaviors [185–187]. By leveraging distributed and cooperative strategies, swarm intelligence algorithms tackle complex tasks while exhibiting robustness, scalability, and flexibility. As a result, swarm intelligence has been applied to challenges in optimization [188], network systems [189], robotics [190], and various other fields [191]. "Robot swarms" are designed under swarm intelligence principles and excel in certain scenarios where traditional AI methods may face limitations [192–196]. For instance, many teams of underwater robots rely on centralized control through explicit communication above water, which limits coordination complexity. However, recent advancements, such as the fish-inspired robot swarm developed by Berlinger et al. [197], demonstrate how implicit communication (e.g., blue light) can achieve complex and adaptive coordination without any centralized control.

1.3.1 Collective motion

A prevalent form of collective behavior is collective motion, where individuals synchronize their movements to produce large-scale coordination [198]. In nature, this coordination is observed in various contexts, as animals move together to forage, migrate, and avoid predators. Engaging in collective motion provides several advantages, such as enhanced discovery of food [199], energy efficiency [200], and improved protection from predators [201]. As a result, collective motion is widespread across biological systems, from bacterial colonies and swarming locusts to schools of fish, flocks of birds, and even groups of mammals, including humans [202–204].

Over the past decades, numerous models of collective motion have been proposed in both theoretical and empirical studies [205–207]. A common characteristic in these models is that individual behavior follows a set of simple rules, influenced by the movement of a relatively small number of neighboring individuals. Three fundamental interaction forces have been suggested as the foundation for collective motion in biological systems:

(i) Repulsion, which prevents collisions by maintaining a minimum distance from nearby individuals.

(ii) Alignment, also referred to as orientation, where individuals adjust their movement to match the direction and speed of their neighbors.

(iii) Attraction, which encourages individuals to stay close to others in the group to maintain cohesion.

These interaction forces shape the movement dynamics of individuals within a group. Mathematically, let p_i denote the position of an individual

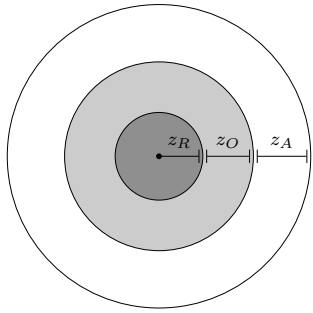

Figure 1.2: Behavioral zones of the three interaction forces in collective motion; repulsion, orientation and attraction with respective widths of z_R, z_O and z_A.

i. Furthermore, the interaction sets $\mathcal{G}_i^R$, $\mathcal{G}_i^O$, and $\mathcal{G}_i^A$ represent the groups of interaction partners associated with repulsion, orientation, and attraction, respectively. The preferred motion vector q_i is then determined based on the motion vector q_j and the relative position $u_{ij} = p_i - p_j$ of an interaction partner j, as follows:

$$q_i = -\alpha_R \sum_{j \in \mathcal{G}_i^R} \frac{u_{ij}}{\|u_{ij}\|} + \alpha_O \sum_{j \in \mathcal{G}_i^O} \frac{q_j}{\|q_j\|} + \alpha_A \sum_{j \in \mathcal{G}_i^A} \frac{u_{ij}}{\|u_{ij}\|} \tag{1.1}$$

where $\alpha_R > 0$, $\alpha_O > 0$, and $\alpha_A > 0$ are weight coefficients corresponding to repulsion, orientation, and attraction, respectively. In some models, these weight factors are not constant but instead vary as a function of the distance to an interaction partner, i.e., $\alpha = f(\|u_{ij}\|)$. For example, a power law decay function has been employed to model repulsion interactions in simulations of sheep herds [208]. Distance-dependent functional forms have also been proposed for all weight factors to simulate fish schools [209], where each interaction partner simultaneously elicits all three types of behavioral responses, albeit with varying degrees of influence. Traditionally, an additional term η_i is added to Equation 1.1 as a stochastic component, such as Gaussian white noise [210], or incorporated into the simulated motion kinematics to account for the intrinsic stochasticity and complexity of animal behavior [211].

In most models of fish schools and bird flocks, each interaction force is associated with a distinct zone surrounding the focal individual. The zone in which a neighbor is located determines the type of interaction that occurs (see Figure 1.2) [212]. Specifically, individuals repel neighbors within the repulsion zone (a disk of radius z_R), align their orientation with those in the alignment zone (an annulus of width z_O), and move toward individuals in the attraction zone (an annulus of width z_A). As a result, the interaction sets are mutually exclusive ($\mathcal{G}_i^R \cap \mathcal{G}_i^O \cap \mathcal{G}_i^A = \emptyset$), meaning each neighbor belongs to only one set at a time. The influence of these zones

has been extensively studied. Couzin et al. [213] demonstrated through simulations that varying the width of the alignment zone (z_O) caused sudden transitions in collective behavior, leading to distinct schooling patterns such as swarming, milling, and schooling, each characterized by different levels of polarization and structure. Hoare et al. [214] experimentally showed that the widths of the alignment and attraction zones are not fixed but vary in response to changes in the grouping motivation of banded killifish (*Fundulus diaphanus*). For group-living animals, individual fitness depends on the size of the group. As a result, the decision to participate in a group reflects a dynamic trade-off between the costs and benefits of membership [215]. This trade-off is regulated by adjustments in the widths of z_O and z_A, which influence the degree of group cohesion. Beyond contextual influences, physiological traits such as body size have also been linked to behavioral zone widths. Kunz and Hemelrijk [209] proposed that both repulsion (z_R) and alignment (z_O) zone widths scale with body size, as larger fish maintain greater inter-individual distances [216].

The presence of alignment forces, however, has been a subject of debate in several studies. Tien et al. [217] analyzed a mixed group of creek chubs (*Semotilus atromaculatus*) and blacknose dace (*Rhinichthys atratulus*), finding no clear behavioral response to others within the typical alignment zone. Instead, they labeled this region as a neutral zone, implying $\alpha_O = 0$. Analyses of schooling golden shiners (*Notemigonus crysoleucas*) [218] and mosquitofish (*Gambusia holbrooki*) [219] found only weak evidence that these fish align with the orientation of their interaction partners. Rather, these studies suggest that apparent alignment emerges as a byproduct of short-range repulsion and longer-range attraction. However, Lukeman et al. [220] analyzed flocks of hundreds of surf scoters (*Melanitta perspicillata*) and found that incorporating an explicit alignment zone into the behavioral model provided a better fit than models without it.

The circular modeling of behavioral zones has also been contested, as interaction behaviors in schooling fish have been shown to exhibit non-isotropic characteristics. These differences have been attributed to the various sensory systems that fish employ, such as the lateral line and visual system [221]. Attraction appears to be primarily mediated by vision, whereas alignment is believed to rely predominantly on the lateral line, suggesting that alignment is most effective with interaction partners positioned laterally. Repulsion, on the other hand, is thought to be influenced by both the lateral line and vision [209]. Similar non-isotropic and non-central features have been observed in speed and turning forces, indicating that attraction-repulsion interactions are not only radial. In the study by Katz et al. [218], the speeding force depended on the front–back distance of interaction partners, while the turning force was influenced by those positioned to the side. This finding also challenges the common assumption in most models that individuals move at similar speeds. Further research has demonstrated that variability in individual speed can significantly impact emergent col-

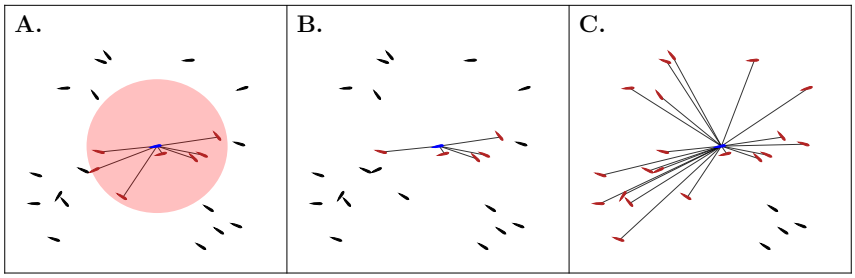

Figure 1.3: The set of influential neighbors (red) is illustrated in regards to a focal fish (blue) for the different models. (A) Metric: all individuals within a certain radius. (B) Topological: only the k nearest neighbors ($k = 5$ in this example). (c) Visual: all individuals that are visually observable, i.e. not obstructed by any other neighbor.

lective dynamics. For instance, differences in individual speed variability among guppies (*Poecilia reticulata*) have been shown to affect polarization, i.e., the degree to which a school moves in the same direction [222].

Another key distinction between models is the selection of an individual's interaction partners ($\mathcal{G}_i$), also known as influential neighbors. These can be categorized into three main types: (i) metric, (ii) topological, and (iii) visual reconstruction (see Figure 1.3). First, in metric models, an individual interacts with all neighbors located within a fixed interaction radius that encompasses the combined range of all behavioral zones [213, 223, 224]. Secondly, in topological models, the set of influential neighbors remains invariant under changes in group density, unlike metric models. As a result, 3D simulations have shown that topological models produce groups that are significantly more stable in maintaining cohesion against noise and external perturbations [225]. This approach was first proposed by Ballerini et al. [226] in an analysis of flocks of starlings (*Sturnus vulgaris*), where each bird appears to interact with a fixed number of nearby neighbors, commonly referred to as the k-nearest neighbors [180, 227]. Similar interaction patterns have also been observed in shoaling sticklebacks (*Gasterosteus aculeatus*), where the influential neighborhood has been estimated at approximately three to five individuals [228]. Thirdly, in visual-based models, the influential neighborhood is derived from an individual's field of view [229]. Analysis of golden shiners (*Notemigonus crysoleucas*) suggests that the structure of visual interaction networks differs markedly from those in metric and topological models, implying that previous assumptions may not accurately capture how information flows within a group [230]. Unlike models that assign specific sensory modalities to different interaction forces, visual-based approaches provide a more general way to incorporate sensory limitations. A simple version of this model only assumes a blind area behind the individual, meaning that neigh-

bors positioned directly behind exert no influence on its movement. This constraint is often combined with metric- or topological-based interaction models [184, 213].

As discussed, numerous model variants based on Equation 1.1 exist, influenced by various individual characteristics (e.g., species, body size) and group-level factors (e.g., group size, context). Despite these variations, Equation 1.1 maintains a balance between simplicity and the ability to capture fundamental properties of collective motion by identifying the minimal general components required for fish schooling [205].

1.3.2 External control of collective motion: from hunting to shepherding

Animals do not exist in isolation within their environment. In many cases, external agents influence and, to some extent, control their collective motion in both competitive (e.g., hunting[3]) and cooperative (e.g., shepherding) contexts. In the competitive context, the collective motion of prey is shaped by hunting predators. As previously discussed, individuals benefit in multiple ways from participating in collective motion when facing predation. For instance, larger prey groups can induce a confusion effect on the predator, reducing its attack success as group size increases [231, 232]. This effect is believed to arise because tracking a single prey among many becomes increasingly difficult, effectively overloading the predator's sensory system [233, 234]. Additionally, prey benefit from information transfer, where the detection of a predator by a few individuals is rapidly propagated throughout the group, enhancing collective awareness and responsiveness [235, 236]. Furthermore, in accordance with the selfish herd theory, prey individuals gain protection by positioning themselves within the group, thereby reducing their exposure to predation [76, 237].

What is believed to be an evolved response to the schooling defenses of prey, some predators also engage in collective motion as part of their hunting strategies [239–242]. For instance, groups of seven yellowtail amberjacks (*Seriola lalandi*) cooperate by forming specific spatial structures to systematically hunt large schools of prey. In one observation, they successfully captured approximately three hundred individuals out of a school of more than two thousand [238]. The predators initially approach prey that are swimming parallel to the coastline and align themselves in a line formation to isolate a smaller group from the main school. Once separation is achieved, they swiftly transition into an arc formation—also referred to as crescent, semicircular, or half-moon shaped—to steer the prey toward the nearby subtidal wall (see Figure 1.4A). The predators remain on one side, with the prey positioned between them and the wall. As the predators

[3]In literature, the term "foraging" is sometimes used. Here, we use "hunting" to specifically refer to predator-prey interactions, whereas foraging encompasses more general relationships between foragers and resources, such as plants.

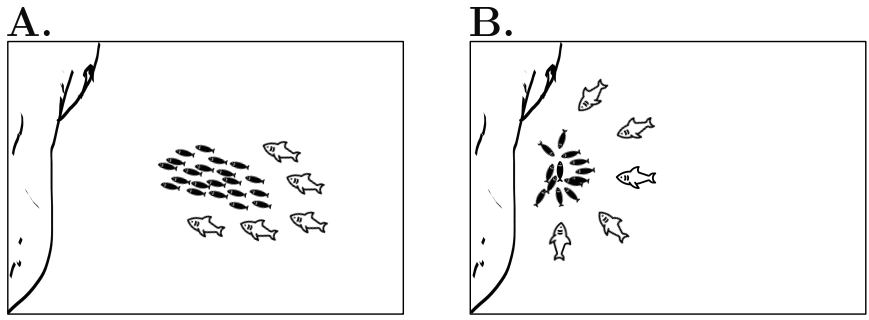

Figure 1.4: Illustrative example of the hunting strategies of yellowtail amberjacks, inspired by Schmitt and Strand [238]. (A) Predators position themselves in an arc formation to steer prey toward a nearby subtidal wall. (B) As the prey become confined near the wall, the predators evenly spread out along the arc, extending their formation to connect with the wall.

advance, the prey instinctively repel away from them, moving toward the wall. Eventually, the prey become confined near the wall, at which point the predators evenly spread out along the arc, connecting their formation to the wall and preventing escape (see Figure 1.4B).

This final spatial configuration is also observed in the group hunting strategies of other animals, where prey, instead of being driven against a wall, become entirely surrounded by predators. For example, in the pack-hunting strategies of wolves (*Canis lupus*), individuals arrange themselves in a stable regular polygon formation, encircling a single stationary prey in open space [243]. We refer to this spatial arrangement as a "caging formation", in which the prey is effectively trapped with no viable escape routes that do not bring it dangerously close to a predator. As a result, the prey remains confined within the interior of this formation. The emergence of a caging formation by wolves has been modeled using two simple decentralized rules [243]: (i) each wolf moves toward the prey until it reaches a certain distance threshold, and (ii) once near the prey, the wolf moves away from other wolves that are also close to the prey.

In the cooperative context, the collective motion of a herd is influenced by shepherds. Shepherding is typically mutually beneficial for both the herd and the shepherds. For instance, shepherds guide the herd toward new pastureland, providing food for the herd while optimizing pasture efficiency for themselves. Shepherds, such as sheepdogs, rely on the same interactive force as predators to guide their herd. When a shepherd approaches too closely, it triggers an avoidance response, causing the individual to move in the opposite direction. This shepherd-induced repulsive force is commonly incorporated into the general model of collective motion

presented in Equation 1.1, extending it as follows:

$$q_i = -\alpha_R \sum_{j \in \mathcal{G}_i^R} \frac{u_{ij}}{\|u_{ij}\|} + \alpha_O \sum_{j \in \mathcal{G}_i^O} \frac{q_j}{\|q_j\|} + \alpha_A \sum_{j \in \mathcal{G}_i^A} \frac{u_{ij}}{\|u_{ij}\|} - \alpha_S \sum_{j \in \mathcal{S}_i} \frac{u_{ij}}{\|u_{ij}\|} \quad (1.2)$$

where $\alpha_S > 0$ is the weight coefficient for shepherd influence, and $\mathcal{S}_i$ denotes the the subset of shepherds $\mathcal{S}$ that affects individual i. Some studies have proposed distance-dependent functional forms for the weight factor α_S, such as a power-law decay function [208], where the repulsive effect varies with the relative distance between the shepherd and the herd individual. In the shepherding literature, $\mathcal{S}_i$ is defined by a behavioral zone with a fixed radius;

$$\mathcal{S}_i = \{j \in S \mid \|u_{ij}\| \leq z_S\} \quad (1.3)$$

where z_S represents the shepherd avoidance-initiation distance. In other words, only shepherds within a radius of z_S induce an avoidance response in the herd individual.

Originally, the shepherding problem was defined as the task of guiding a group of animals to a desired location, where robots have been demonstrated to serve as effective shepherds [77]. Robotic shepherding has a wide range of applications, as previously discussed in Subsection 1.1.3. One of the first developments in robotic shepherding was carried out in The Robot Sheepdog Project, where a single robot successfully shepherded a flock of ducks (*Anas platyrhynchos*) within a constrained environment [244]. However, various studies have shown that as herd size increases, multiple cooperative shepherds outperform single shepherds in execution time and efficiency. In unconstrained environments, multiple shepherds are more effective at maintaining group cohesion, leading to improved control over the herd [245–247].

When using multiple shepherds, researchers have generally relied on spatial formations to guide the herd. For instance, line formations help distribute the applied force along the herd while pushing it toward the goal location [62]. Arc formations, inspired by predator hunting strategies, are effective in maintaining herd cohesion while controlling movement. To achieve arc formations, sliding mode controllers and single continuous control laws have been proposed for coordinating multiple robots [246, 248]. A caging formation has also been introduced, defined as follows:

Definition 1.1. *Robots are in a caging formation around animals if the convex hull of their spatial positions forms a polygon where each side is shorter than* $2z_S$ *and and all animals are positioned within its interior.*

When two robots are positioned less than $2z_S$ apart, they exert a combined repulsive force on the herd, preventing animals from passing between them. Varava et al. [249] proposed a centralized RRT-based (rapidly-exploring random tree) algorithm to position multiple shepherds in a

caging formation around the herd. However, their approach assumes an initially well-formed cage, and their simulated herd remains stationary unless repelled by the robots. Additionally, their simulations suggest that even a small increase in the number of shepherds significantly reduces the algorithm's performance. Other AI-based techniques, such as machine learning [250, 251] and genetic algorithms [252, 253], have also been explored. However, the learning problem is non-trivial due to the large search space required for training a model capable of effective generalization [59]. As a result, these techniques have mainly been applied to single-shepherd scenarios.

The shepherding approaches discussed so far, along with the majority of existing literature, rely on centralized coordination or global information, such as knowledge of the herd's center of mass [254]. In their comprehensive review of robotic shepherding, Long et al. [77] emphasize that such approaches lack flexibility, as they are not easily adaptable to the dynamic environments in which animals naturally move. This limitation poses a significant challenge for real-world applications, where robotic shepherds must rely solely on local information gathered by sensors with limited range [255]. Furthermore, Long et al. [77] argue that practical applications of robotic shepherding require robust and scalable systems that can withstand robot failures and adapt to varying herd sizes. Additionally, they highlight the necessity for shepherds to learn the behavioral profiles of the animals and adjust their decision-making models accordingly. Indeed, shepherds may encounter a wide range of animal behaviors. For instance, in predator-prey fish interactions, the avoidance-initiation distance z_S has been suggested to depend on individual factors such as body size and prior exposure to threats [256]. Given these challenges, swarm robotic systems offer a promising alternative for shepherding, as they provide flexibility, robustness, and scalability [77]. Designed based on swarm intelligence principles, these robots operate under simple local rules, allowing them to explicitly model and adapt to animal behavior in real-time.

1.3.3 Behavior-based design methods

In swarm intelligence, the desired behavior is defined at the collective (macroscopic) level, while the control design occurs at the individual (microscopic) level. Directly translating collective requirements into individual behaviors is inherently complex in most swarm robotics problems [189, 257, 258]. Behavior-based design methods are one of the key approaches, which draws inspiration from nature and relies on an intuitive understanding of collective behavior [190]. In this design process, developers iteratively refine individual behavior models based on how closely the emergent collective behavior aligns with the desired outcome.

A common method is to divide the full individual behavior by modeling different behavioral states. Each behavioral state causes the individual

to act in a different way, and individuals transition between states based on set of rules (e.g., probabilistic finite state machines [259, 260]). For instance, the herding of sheep by dogs has been modeled by adaptively switching between two different states of collecting the herd when they are too dispersed, and driving them once they are aggregated [261].

Another widely used method is inspired by physics, where each individual is considered as a virtual particle that is influenced by virtual forces. The models of collective motion previously described (see Equations 1.1 and 1.2) are designed in this matter[4]. Not only intelligent agents are designed to exert virtual forces, but also relevant environmental characteristics (e.g. obstacles, light, temperature) can be expressed through virtual forces, creating what is referred to as an artificial potential field [263, 264].

1.4 Research scope: fish as a case study

Animals are increasingly being faced with new challenges caused by human activity. The loss of biodiversity is detrimental not only to animals, but also us humans. Nature conservation tasks aim to restore and preserve ecosystems to counteract this biodiversity loss. Bioinspired robots form a promising tool to fulfill these tasks, overcoming the challenges that current human-led and conventional robot approaches face. More specifically, swarms of bioinspired robots provide the scalability, robustness, and flexibility required for dealing with the challenges of interacting with animal collectives in particular, and the dynamic environments they inhabit. While research is advancing the robotic technology for such robots to be used in real applications, there is an increasing need for the development of the intelligence of these robot systems.

This thesis therefore focuses an overarching aim: *Developing the intelligence of swarm robot systems to offer support and protection for animal collectives*. To explore this overarching aim, this research uses *fish as a case study* of animal collectives, motivated by the current literature and the potential environmental impact. As shown in this chapter, the majority of the literature on animal-robot interactions uses fish as study organisms. In recent years, there has also been significant progress both theoretically and practically in developing bio-inspired fish robots [141]. Additionally, fish behavior (e.g. schooling) has been extensively studied and modeled over the years by biologists [265]. Although fish are an important part of our natural environment, they are increasingly faced with anthropogenic threats such as illegal fishing [266], pollution [267], and invasive species [268]. As Rachel Carson famously noted in "The Sea Around Us" [269];

> *"It is a curious situation that the sea, from which life first arose, should now be threatened by the activities of one form of that life."*

[4] It should therefore come to no surprise that collective motion has also been extensively studied in the context of self-propelled particles [262].

Instead of threatening aquatic life, this thesis envisions a future in which human activities support these ecosystems through the deployment of intelligent bioinspired robotic systems.

1.5 Research questions

In order to address the overarching aim of this thesis, the work is divided into four research questions that include accompanying hypotheses. Below, these research questions are listed, including motivations for their selection.

For swarm robotic systems to contribute to fish conservation, the ability to explore and navigate in a manner similar to fish collectives is essential (see Section 1.2). Robot capable of traversing the dynamic environments that fish inhabit, such as coral reefs, could map their structure and identify potential threats. More interestingly, when biomimetic robots are accepted as conspecifics, they can move collectively with fish, enabling the monitoring of their health and potentially supporting them (e.g. signaling the presence of a threat that fish cannot detect through adaptive behavior). Hence, we pose the following research question:

Research question 1 (Chapter 2). *How can robot teams explore and navigate like fish in dynamic environments?*

Exploration is a critical aspect of fish behavior, particularly during foraging when fish search for food [270]. For many species, this involves continuous swimming and maneuvering, often occupying a significant portion of their active hours [270]. Fish foraging strategies are flexible, influenced primarily by prior knowledge of resource distribution [271]. In certain scenarios, fish prefer to search for food independently, spreading out as a group while maintaining proximity to some conspecifics. This proximity facilitates the exchange of information regarding food locations or nearby predators. Such loosely aggregated behavior is referred to as shoaling, where individuals swim in various directions but remain socially connected for mutual benefits like increased foraging efficiency and predator protection [272]. This collective structure can be described as high coverage (occupying a considerable area or volume in space) and low order (low alignment in movement direction).

Alternatively, fish may forage in a highly coordinated manner, collectively moving in the same direction, a behavior known as schooling. Schooling is not limited to foraging, but also occurs in other contexts, such as migration. An ordered group structure (high order) may provide additional anti-predatory advantages, such as facilitating the detection of abrupt changes in the behavior of conspecifics, which could signal potential danger [273]. Schooling can also be essential for maintaining group cohesion, particularly in fast-moving species, and can enhance the flow

of information within the group [235, 274]. During schooling, the spatial coverage of the group may fluctuate [275], with high coverage being advantageous for foraging by maximizing the exploration of the environment, while low coverage is preferable for predator avoidance [276]. Additional factors, such as local population density, can further influence the collective structure [277]. For instance, the volume of fish schools has been shown to depend on the density of the local population, and a similar relationship is presumed for the area of fish living in shallow waters, where movement occurs primarily in two-dimensional space.

The tendency to alternate between shoaling and schooling varies both between species and within species, depending on their ecological niche and motivational state [278]. Some species opportunistically shift between these behaviors to maximize their chances of survival in response to environmental conditions. To enable effective interaction between robots and fish in such dynamic environments, robot teams must be capable of adopting and adapting to the distinct collective structures outlined above: (i) high coverage with low order, (ii) high coverage with high order, and (iii) low coverage with high order. In addition, these collective structures are also crucial robot teams when performing exploration tasks independently of fish. For instance, (i) enables robots to maximize their field of observation by maintaining broad spatial coverage and maximizing inter-individual distance; (ii) facilitates simultaneous navigation and exploration; and (iii) allows for efficient movement through confined spaces, such as underwater cave systems, or coordinated transport of objects.

As previously discussed in Subsection 1.3.1, various models of collective motion have been proposed, incorporating short-range repulsion, long-range attraction, and alignment of orientation with certain neighbors. These models have been shown to generate varying degrees of group order depending on parameter settings. However, there is limited research on how these parameters influence group coverage, which is another key factor in defining the necessary collective structures for effective exploration. It remains unclear whether current models can produce the full range of required structures, combining both order and coverage appropriately. Additionally, it is uncertain whether such models can dynamically adjust the emergent collective structures in response to environmental changes. Therefore, we test the following hypothesis:

Hypothesis 1.1. *Collective motion models of fish are limited in their ability to generate the full range of collective structures relevant to exploration.*

Studies testing the assumptions underlying models of collective motion of fish have identified key similarities with several model assumptions (see Subsection 1.3.1 for a detailed overview). Specifically, they confirmed that fish respond to the position of their neighbors by following short-range repulsion and longer-range attraction rules. However, some studies only found weak evidence to support the alignment of fish with their neighbors'

orientation [218, 219]. Instead, the findings suggest that group alignment is achieved through alternative mechanisms, potentially via attraction and repulsion dynamics and by fish following individuals directly in front of them [279, 280].

Importantly, the researchers emphasize that this does not imply that alignment rules are never used by fish. For instance, during predation threats, alignment with neighbors has been shown to facilitate the rapid transmission of information throughout the group [281]. They hypothesize that fish may adapt their alignment behaviors in response to environmental stimuli or transmitted information. We adopt this hypothesis, positing that the full range of collective structures required for exploration can emerge through adaptations in alignment rules in response to local information, particularly by adjusting whether or not alignment is used and by altering which neighbors influence alignment:

Hypothesis 1.2. *Adaptations in alignment rules, in response to local information, facilitate the emergence of the full range of collective structures.*

To test this hypothesis, we examine the effects of adjusting the alignment weight factor (i.e. $\alpha_O = 0$ or $\alpha_O > 0$), and altering the widths of the behavioral zones which directly influence the set of influential neighbors in Equation 1.1. Under this hypothesis, in a mixed society of fish and robots, the robots could adapt their alignment rules, prompting a similar response in nearby fish. This behavior would then propagate through the group, dynamically altering the collective structure as needed. This approach could be particularly beneficial in situations where fish are unaware of potential feeding grounds or environmental threats, such as areas occupied by illegal fishers. Conversely, the robots should also adapt to the behavior of the fish when they modify their structure to optimize foraging strategies or detect threats. This establishes a reciprocal relationship, where both robots and fish influence each other's collective behavior, enhancing the adaptive capacity of the group.

Beyond exploring and navigating fish habitats for biodiversity documentation, habitat mapping, and threat identification, this dissertation also addresses fish protection. While exploration helps detect potential threats, and robots in mixed societies may influence the collective structure of fish schools, these factors alone do not guarantee that fish successfully avoid dangers. Therefore, this dissertation seeks to address the following question:

Research question 2 (Chapter 3). *How can robot teams effectively protect fish schools from dynamic dangers?*

In our rapidly changing world, fish populations increasingly face anthropogenic threats for which they are not evolutionarily adapted, such as

pollution and illegal fishing. In many instances, fish fail to recognize these threats as dangers until they have already been negatively impacted [282]. This research question seeks to explore the development of artificial intelligence for swarm robotic systems capable of proactively protecting fish groups and mitigating these types of risk by guiding them away from potential dangers.

As introduced in Subsection 1.3.2, the classic shepherding problem involves robots guiding animals to a specific target location, which is assumed to be known to the robots as prior information. However, in the context of dynamic environments, prior information may be unreliable or entirely unavailable due to constant environmental changes. Addressing this challenge requires adaptive strategies that account for two possible scenarios. First, in the absence of immediate danger, the robots should be positioned to detect potential threats while minimizing unnecessary disruption, allowing fish to carry out their natural activities. Second, when a threat is detected, the robots should steer the fish away from danger while continuing to monitor the surroundings for emerging risks.

To most optimally detect dangers approaching the fish from any direction, the robots should be evenly spread out surrounding the fish. This spatial configuration corresponds to the caging formation described in Definition 1.1. Since swarm robotic systems exclusively rely on local observation and communication, this formation also ensures that information about potential dangers can be effectively transmitted across the entire robot team. A key factor influencing this formation is the nearest distance r between fish and robots. To construct a caging formation, robots can leverage the cohesive properties of fish schools, particularly when the school is connected by radius r. Geometrically, this occurs when the union of circles, where the center points are the fish positions and the radii are r, forms a single connected set. This union naturally defines a simple closed contour, meaning a boundary that does not intersect itself except at its starting and ending points, which coincide.

When the group is connected by r and a sufficient number of robots are deployed to form an equilateral polygon with side lengths shorter than $2z_S$ (see Definition 1.1), we hypothesize that robots can establish a caging formation by following two simple decentralized rules inspired by animal hunting strategies (see Subsection 1.3.2). First, each robot moves along the contour defined by r. Second, each robot maintains an equal distance from its two adjacent neighbors along this contour.

Hypothesis 2.1. *Robots can construct a caging formation around a collectively moving fish school with a nearest fish-robot distance r, provided the school is connected by r and a sufficient number of robots are deployed.*

The connectivity of the school for a given r depends on its cohesion, which is influenced by various parameters in collective motion models. To test this hypothesis, a behavioral zonal model is used, where connectivity

is influenced by the widths of these zones (see Subsection 1.3.1). Additionally, different approaches for selecting influential neighbors are considered to assess their impact on school cohesion.

Once the robots have established a caging formation, they are positioned to detect surrounding dangers and preemptively steer the group away. Since threats can emerge dynamically, varying in form, location, and timing, the robots must continuously adapt without prior knowledge of their occurrence. Each robot locally observes potential dangers and shares this information with neighboring robots to ensure an informed response. When multiple threats arise, the team integrates observations from different individuals to collectively determine the most suitable direction for steering the fish. By continuously exchanging and updating this information, the robots reach a consensus on the optimal escape route.

As previously stated, the robots should maintain a caging formation while steering the fish. According to Equation 1.3, fish only respond to robots when they are within a given distance ($r < z_S$), otherwise continuing their natural behaviors. Therefore, we hypothesize that robots can steer the school while maintaining the caging formation by adaptively switching between two configurations: one that exerts a repulsive influence on the fish and one that does not. When a robot is positioned behind all nearby fish relative to the direction of movement, it follows the caging algorithm for $r < z_S$, inducing a repulsive effect. Conversely, when positioned in front of nearby fish, it applies the caging algorithm for $r > z_S$, allowing the school to move freely. This adaptive behavior essentially results in a combined caging structure consisting of two arc formations. One arc formation actively influences the collective motion of the fish by exerting a repulsive effect, similar to the hunting strategy of yellowtail amberjacks as described in Subsection 1.3.2. The other arc remains passive to the fish, creating a semi-open space that allows the fish move within the caging formation. When no threats are present, all robots maintain a distance between the nearest fish of $r > z_s$, ensuring that the fish can continue their natural behavior without unnecessary interference. Thus, we test the following hypothesis:

Hypothesis 2.2. *Robot teams can guide fish away from dangers while maintaining a caging formation by adjusting their nearest fish-robot distance.*

While steering fish away from dangers is crucial for their protection, effective conservation also requires guiding them toward safe breeding and feeding grounds. In many scenarios, simply avoiding threats is not sufficient, as fish must reach specific destinations to sustain their populations. A key challenge is ensuring that fish remain within a designated area while being directed toward a target location. This guidance follows a controlled trajectory that avoids dangers, obstacles, or unsuitable terrain,

similar to how robots establish migration corridors in disrupted environments. For example, robots can facilitate fish migration while simultaneously shielding them from hazards. This approach is particularly valuable for endangered species facing threats from invasive predators or human disturbances. Conversely, robotic guidance can also be used to manage invasive predatory species by confining them within a caging formation and directing them toward controlled capture zones. Therefore, this thesis aims to seek an answer to the following question:

Research question 3 (Chapter 4). *How can robot teams effectively guide fish along a dynamic path?*

To follow a path, the robots require only local directional information at their current position rather than global knowledge of the entire trajectory. If they can operate with just this local information, the path itself can dynamically change without compromising the robots' performance. This adaptability is particularly advantageous in the dynamic environments where fish live. The path is represented abstractly using an artificial potential field, a common approach in swarm intelligence design (see Subsection 1.3.3). In this formulation, the potential function encodes the path such that, mathematically, it reaches a local minimum at the destination. Within the path, the potential decreases in the direction leading toward the destination, while outside the path, it decreases toward the path itself. Each robot computes the gradient of the potential field based on local observations and moves accordingly, ultimately guiding the group to the destination. For real-world applications, this potential field can incorporate various environmental stimuli, such as natural light or planned landmarks. Since the robots surround the fish, they can also monitor environmental conditions—such as oxygen levels, pollutants, or temperature—which can be integrated into the potential function. This ensures that the fish remain in habitats that are both healthy and suitable for their survival.

If each robot can determine its local desired direction through a potential function, direct information exchange between robots may no longer be necessary. The consensus previously required among robots in Research Question 2 is now inherently embedded within the potential function. The constraint of a designated destination naturally shapes the function, implicitly guiding all robots toward a unified direction. This leads us to the following hypothesis:

Hypothesis 3.1. *Robot teams can guide a fish school along a dynamic path without direct communication when path information is encoded in an artificial potential field.*

Since the robots have no prior knowledge of the path, they cannot rely on preplanned, optimized control policies. Instead, they must continuously discover the trajectory as they move, which may result in suboptimal performance depending on key factors related to both the path and

the system (i.e., the fish and robots). Here, performance is measured as the percentage of fish remaining within the path over time. Detecting changes in path orientation is a local process, meaning certain path characteristics, such as sharp turns or varying widths, can pose challenges. These factors may cause delays in the system's ability to adapt its direction, affecting the effectiveness of fish guidance. Additionally, system size, defined by the number of robots and fish, plays a role in this process. When a portion of the group alters its direction, the change must propagate through the entire school, which takes time. This delay effect can be further intensified when the school moves at higher speeds, as adjustments must be made more rapidly. Therefore, we test the following hypothesis:

Hypothesis 3.2. *The effectiveness of guiding fish along a path while keeping the school within its boundaries depends on path characteristics, system size, and movement speed.*

So far, the research questions and their corresponding hypotheses have been tested through simulations. Developing artificial robotic intelligence in a simulated environment is a crucial first step for several reasons. It ensures the safety of real animals, provides a controlled testing environment, and allows for scalable adjustments to the number of fish or robots involved. Before Research Questions 2 and 3 can be explored in real-world experiments, where robots influence the collective motion of fish by leveraging their avoidance response to robots, it is essential to first examine the underlying model assumptions of such fish-robot interaction. In particular, it remains largely unexplored which aspects of robot behavior shape fish avoidance responses and how this behavior deviates from conventional models (as described in Subsections 1.3.1 and 1.3.2).

Traditionally, predator-like biomimetic robots designed to resemble natural predators of a species have been used to elicit avoidance responses in fish (see Subsection 1.1.2). In contrast, robots that mimic the appearance of conspecifics, known as conspecific-like robots, have primarily been employed to attract, lead, or swim alongside fish. However, live animals do not always respond positively to conspecifics. Depending on the social context, such as mitigating the risk of disease transmission [283] or mating dynamics [284], they may instead exhibit avoidance behavior.

This thesis argues that understanding how also conspecific-like robots influence avoidance behavior in fish is essential for advancing mixed animal-robot societies. This knowledge could enable robots to adapt their roles dynamically, alternating between evoking avoidance and attraction as needed. Such adaptability could significantly expand the potential applications of robotic systems in shaping fish behavior. A possible extension of this idea is the use of heterogeneous robot teams, where some robots act as informed agents swimming along (Research Question 1) while others

function as repulsive agents to guide fish (Research Questions 2 and 3). This dual-role approach may enhance control performance, particularly given the behavioral heterogeneity observed in animal collectives, which could limit the effectiveness of a single control strategy. To advance these capabilities, further insights are needed to refine modeling and controller design, with a particular focus on the elicitation of avoidance behavior. Therefore, we pose the final question of this thesis:

Research question 4 (Chapter 5). *How do the avoidance responses of fish to a conspecific-like robot differ from conventional models of animal-conspecific and animal-shepherd interactions?*

To answer this question, a robot was programmed to approach a guppy at a fixed speed. The robot featured a 3D-printed replica of a Trinidadian guppy (*Poecilia reticulata*) attached to its two-wheeled base, a design previously shown to be accepted as a conspecific by guppies in multiple studies [13, 25, 285–287]. Each fish was subjected to multiple consecutive trials during an experiment, with the robot sampling a different speed in the medium to high range for each trial. This setup allows for testing hypotheses on fish avoidance behavior in relation to repeated exposure, robot approach speed, and proximity.

As previously noted, studies on fish avoidance behavior have primarily focused on anxiety-related and anti-predator contexts. In response to predation threats, guppies exhibit various anti-predator behaviors, including escape movements, freezing, and thigmotaxis [288]. Escape responses typically involve rapid acceleration and a directional change (i.e. sharp turn) to increase distance from the threat [289]. Freezing is a defense mechanism in which the fish remains motionless as an acute stress response. Thigmotaxis refers to the tendency of fish to stay close to boundaries, avoiding open areas such as the center of a tank. Research has shown that prey do not necessarily escape at the earliest detection of a threat but instead at a distance dictated by the trade-off between the costs and benefits of fleeing. This decision is influenced by factors such as predator approach speed [290] and repeated exposure to the threat [256]. Although the robot in this study visually mimics a conspecific, we hypothesize that repeated approaches by the robot may lead the fish to perceive it as a threat, increasing the likelihood of avoidance behavior over successive trials.

Hypothesis 4.1. *Repeated exposure to robot approaches increases avoidance behavior in fish.*

Hypothesis 4.2. *The robot's behavior elicits characteristic anti-predator responses in fish.*

Avoidance behavior (repulsion) has been modeled both between animal conspecifics (Equation 1.1) and animal-shepherd (Equation 1.2). In

both contexts, the avoidance response has traditionally been modeled solely dependent on the distance to influential neighbors, as discussed in Subsections 1.3.1 and 1.3.2. Typically, avoidance responses are defined as an individual moving at constant speed in the opposite direction to the average position of others within a defined radius [213, 218, 291]. However, studies on fish avoidance behavior in response to a conspecific-like robot suggest that avoidance speed may not be fixed. One study indicates that fish avoidance speed is influenced by the robot's approach speed [292], while another shows that avoidance speed increases as proximity to the robot decreases [293]. Building on these studies, we hypothesize that both robot speed and proximity interact to influence fish avoidance speed;

Hypothesis 4.3. *Fish avoidance speed is influenced by the interaction between the robot's approach speed and its distance to the fish.*

In addition to speed, key components of the avoidance mechanism include reaction distance (the distance at which a fish initiates an avoidance response), as well as acceleration and turning. Traditionally, models treat reaction distance as a fixed threshold. However, more recent work challenges this assumption by incorporating decay functions that gradually reduce the strength of the repulsion force with increasing distance [294, 295]. Notably, studies in predator-prey dynamics have shown that reaction distance is modulated by the predator's approach speed [290]. We hypothesize that a similar effect may also arise in response to a robotic conspecific.

Building on this idea, we propose that the reaction distance is influenced not only by the robot's approach speed, but also by the fish's own speed at the moment of interaction. From an energy-optimization perspective, fish may dynamically adapt their avoidance strategy—whether to initiate early trajectory adjustment or to delay and execute a sudden evasive maneuver—based on their relative speed to the robot. More specifically, when the fish's relative speed to the robot is high, it may prefer to adjust its trajectory early, maintaining movement while gradually increasing distance. Conversely, when its relative speed is low, the fish may be more inclined to remain still or move minimally, delaying its response until the robot is in close proximity, at which point it may execute a sudden escape maneuver. In this context, kinematic aspects such as acceleration and turning become essential to understand how avoidance responses are formed. Gaining insight into how these are shaped by both fish and robot speed is crucial for accurately modeling fish-robot interactions. We therefore conclude with the following hypothesis:

Hypothesis 4.4. *Fish reaction distance is shaped by the interaction between their own speed and the robot's approach speed, which in turn modulates acceleration and turning as components of the avoidance mechanism.*

Delimitations. This thesis investigates the proposed research questions within two-dimensional environments. Both the simulations and experimental study are conducted in two dimensions, as guppies primarily inhabit shallow waters where movement is predominantly lateral [296]. Many other schooling fish species also reside in shallow habitats such as estuaries, coral reefs, and coastal regions [297]. Adopting a two-dimensional framework provides a structured and progressive approach to model development. Additionally, many existing models of collective behavior have been experimentally analyzed in two dimensions. By first establishing and validating interaction rules in two dimensions, we can systematically refine and extend these models to three-dimensional environments in future research. Thus, two-dimensional modeling remains both biologically relevant and a practical simplification.

1.6 Thesis outline

This dissertation aims to answer the research questions posed in the previous section, and to test the enumerated hypotheses. Each chapter corresponds to work that was published in international journals or conferences, or is currently in review for publication. In this sense, the chapters were written to be self-contained[5].

In Chapter 2, we investigate how swarm robots can explore and navigate like fish within dynamic environments. We identify three essential collective structures for effective exploration and navigation, characterized by the swarm's spatial coverage and the degree of alignment among individuals. First, the ability of traditional collective motion models of fish schooling to generate these structures is assessed by varying the width of behavioral zones. Two behavioral zonal models are considered, differing in how influential neighbors are selected: metric-based and topological-based approaches. Our results indicate that these models are unable to produce the full range of required structures. To address this limitation, we propose an extension of the metric-based zonal model. By varying both the alignment zone width and converting the alignment zone to a neutral zone, all three structures can be generated. Furthermore, we show that a robot swarm following the proposed model can dynamically transition between these structures in dynamic environments, where different stimuli appear over time and are perceived by few individuals.

Chapter 3 presents a swarm intelligence-based approach for protecting fish schools from dynamic dangers. To preemptively detect threats before they interact with the fish, and without prior knowledge of their presence, the robots are tasked with constructing a caging formation around the school. This approach relies on a certain level of school cohesion.

[5]Note that some of the text might have been changed to accommodate the reader of this dissertation.

Therefore, we first examine the conditions under which different collective motion models maintain this cohesion. Since changes in cohesion affect spatial coverage, we also measure the resulting coverage in each case and computationally determine the minimum number of robots required to successfully construct a caging formation. The results show that the proposed algorithm enables robots to effectively cage the school when the constraints of school cohesion and robot swarm size are met. To guide fish away from danger, robots adaptively switch between two states based on their local position relative to the desired direction: one that induces a repulsive effect on nearby fish and one that remains neutral. The effectiveness of this approach is tested in dynamic environments across two tasks: (i) avoiding a stationary danger that appears stochastically, and (ii) staying within a shrinking circular safe zone. The findings demonstrate that the robot swarm successfully prevents any fish from being affected by danger.

In Chapter 4, the guidance of the school is further constrained by introducing a more complex problem where the robot swarm must navigate the school along a time-varying trajectory. Although this constraint imposes stricter movement requirements, it simultaneously eliminates the need for consensus decision-making among robots, as was necessary in the approach from Chapter 3. Instead, each robot derives its guidance from an artificial potential field that encodes path information, enabling the swarm to guide the fish without requiring inter-robot communication. To validate the algorithm, paths with varying turn sharpness and space between boundaries are considered, with effectiveness measured as the percentage of fish remaining within the boundaries until reaching the target location. The results indicate that performance declines when navigating sharp turns, which is influenced by both the available space between the boundaries, and the number of robots and fish. Furthermore, the effectiveness of the algorithm depends on the maximum velocity of the fish, as maintaining caging formation becomes more challenging at higher speeds.

As the previous two chapters relied on model assumptions about fish avoidance behavior in response to robots, Chapter 5 presents experimental evidence on how robot behavior influences fish avoidance, highlighting deviations from conventional model assumptions. To investigate this, a conspecific-like robot repeatedly approached individual guppies, varying its speed across trials within a medium to high range. Our findings suggest that as fish experienced repeated encounters with the robot, their likelihood of exhibiting avoidance responses increased, indicating a form of behavioral adaptation. In some instances, they displayed typical anti-predator reactions, such as freezing or executing rapid escape maneuvers, suggesting that the robot was perceived as a potential threat. In particular, escape responses were not only dictated by proximity, but were also influenced by the robot's approach speed, challenging conventional models that assume distance alone determines avoidance behavior. Moreover, avoidance speed increased with higher robot speeds and decreased prox-

imity, highlighting the dynamic nature of fish responses.

Finally, Chapter 6 concludes the dissertation by revisiting the main research questions in light of the findings. It highlights key contributions, identifies limitations, and outlines open questions that emerged. The chapter then broadens its perspective to propose future research directions that go beyond the scope of the original research questions.

1.7 List of publications

The list below contains research that is discussed in this dissertation, as well as other research work that was developed throughout this PhD journey, which has been published in international scientific journals or presented at international conferences.

1.7.1 Publications in international journals[6]

1. **Stef Van Havermaet**, Andreas Gerken, Deni Mazrekaj, David Bierbach, Pieter Simoens, Tim Landgraf, Yara Khaluf (2025) "*The influence of a fish-like robot on the avoidance behavior of fish*" (in review). Scientific Reports.
2. **Stef Van Havermaet**, Yara Khaluf, Pieter Simoens (2024) "*Reactive shepherding along a dynamic path*". Scientific Reports, 14. https://doi.org/10.1038/s41598-024-65894-5
3. **Stef Van Havermaet**, Pieter Simoens, Tim Landgraf, Yara Khaluf (2023) "*Steering herds away from dangers in dynamic environments*". Royal Society Open Science, 10 (5). https://doi.org/10.1098/rsos.230015

1.7.2 Publications in international conferences[7]

1. **Stef Van Havermaet**, Pieter Simoens, Yara Khaluf (2022) "*An adaptive metric model for collective motion structures in dynamic environments*". ANTS2022, International conference on swarm intelligence, pp. 257-265
2. **Stef Van Havermaet**, Yara Khaluf, Pieter Simoens (2021) "*No more hand-tuning rewards: Masked constrained policy optimization for safe reinforcement learning*". AMAS2021, Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, pp. 1344-1352
3. Johannes Nauta, **Stef van Havermaet**, Pieter Simoens, Yara Khaluf (2020) "*Enhanced foraging in robot swarms using collective Lévy walks*". ECAI2020, the 24th European Conference on Artificial Intelligence, pp.171-178

[6]These publications are classified as 'A1' by Ghent University according to the following definition: "*Articles included in one of the Web of Science databases 'Science Citation Index', 'Social Science Citation Index' or 'Arts and Humanities Citation Index'. Limited to the publications document type article, review, letter, note, proceedings paper.*"

[7]These publications are classified as 'C1' by Ghent University according to the following definition: "*Articles published in proceedings of academic conferences, not included in (A1) or (P1) (full articles with exception of abstracts, unpublished presentations, posters, ...). P1: Proceedings included in one of these Web of Science indexes: 'Conference Proceedings Citation Index - Science' of 'Conference Proceedings Citation Index - Social Science and Humanities'. Limited to publications document type: article, review, letter, note, proceedings paper, with exception of publications classified A1.*"

Bibliography

[1] K. I. Roumeliotis and N. D. Tselikas. Chatgpt and open-ai models: A preliminary review. *Future Internet*, 15(6):192, 2023.

[2] R. K. Katzschmann, J. DelPreto, R. MacCurdy, and D. Rus. Exploration of underwater life with an acoustically controlled soft robotic fish. *Science Robotics*, 3(16):eaar3449, 2018.

[3] A. E. Turgut, H. Çelikkanat, F. Gökçe, and E. Şahin. Self-organized flocking in mobile robot swarms. *Swarm Intelligence*, 2:97–120, 2008.

[4] M. Duarte, J. Gomes, V. Costa, T. Rodrigues, F. Silva, V. Lobo, M. M. Marques, S. M. Oliveira, and A. L. Christensen. Application of swarm robotics systems to marine environmental monitoring. In *OCEANS 2016-Shanghai*, pages 1–8. IEEE, 2016.

[5] J. Werfel, K. Petersen, and R. Nagpal. Designing collective behavior in a termite-inspired robot construction team. *Science*, 343(6172):754–758, 2014.

[6] D. Romano, E. Donati, G. Benelli, and C. Stefanini. A review on animal–robot interaction: from bio-hybrid organisms to mixed societies. *Biological cybernetics*, 113:201–225, 2019.

[7] B. Webb. What does robotics offer animal behaviour? *Animal behaviour*, 60(5):545–558, 2000.

[8] S. Garnier. From ants to robots and back: How robotics can contribute to the study of collective animal behavior. In *Bio-inspired self-organizing robotic systems*, pages 105–120. Springer, 2011.

[9] J. Krause, A. F. Winfield, and J.-L. Deneubourg. Interactive robots in experimental biology. *Trends in ecology & evolution*, 26(7):369–375, 2011.

[10] J. Halloy, F. Mondada, S. Kernbach, and T. Schmickl. Towards bio-hybrid systems made of social animals and robots. In *Biomimetic and Biohybrid Systems: Second International Conference, Living Machines 2013, London, UK, July 29–August 2, 2013. Proceedings 2*, pages 384–386. Springer, 2013.

[11] N. Kobayashi, M. Yoshida, N. Matsumoto, and K. Uematsu. Artificial control of swimming in goldfish by brain stimulation: confirmation of the midbrain nuclei as the swimming center. *Neuroscience letters*, 452(1):42–46, 2009.

[12] S. C. Van Den Berg, R. B. Scharff, Z. Rusák, and J. Wu. Openfish: Biomimetic design of a soft robotic fish for high speed locomotion. *HardwareX*, 12:e00320, 2022.

[13] M. Maxeiner, M. Hocke, H. J. Moenck, G. H. Gebhardt, N. Weimar, L. Musiolek, J. Krause, D. Bierbach, and T. Landgraf. Social competence improves the performance of biomimetic robots leading live fish. *Bioinspiration & Biomimetics*, 18(4):045001, 2023.

[14] T. Landgraf, G. H. Gebhardt, D. Bierbach, P. Romanczuk, L. Musiolek, V. V. Hafner, and J. Krause. Animal-in-the-loop: Using interactive robotic conspecifics to study social behavior in animal groups. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:487–507, 2021.

[15] F. Mondada, A. Martinoli, N. Correll, A. Gribovskiy, J. I. Halloy, R. Siegwart, and J.-L. Deneubourg. A general methodology for the control of mixed natural-artificial societies. Technical report, Pan Stanford Publishing Singapore, 2013.

[16] F. Bonnet, S. Binder, M. E. de Oliveria, J. Halloy, and F. Mondada. A miniature mobile robot developed to be socially integrated with species of small fish. In *2014 IEEE International Conference on Robotics and Biomimetics (ROBIO 2014)*, pages 747–752. IEEE, 2014.

[17] T. Bartolini, V. Mwaffo, A. Showler, S. Macrì, S. Butail, and M. Porfiri. Zebrafish response to 3d printed shoals of conspecifics: the effect of body size. *Bioinspiration & biomimetics*, 11(2):026003, 2016.

[18] N. Abaid, T. Bartolini, S. Macrì, and M. Porfiri. Zebrafish responds differentially to a robotic fish of varying aspect ratio, tail beat frequency, noise, and color. *Behavioural brain research*, 233(2):545–553, 2012.

[19] A. Gribovskiy, J. Halloy, J.-L. Deneubourg, H. Bleuler, and F. Mondada. Towards mixed societies of chickens and robots. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4722–4728. IEEE, 2010.

[20] G. Polverino, P. Phamduy, and M. Porfiri. Fish and robots swimming together in a water tunnel: robot color and tail-beat frequency influence fish behavior. *PloS one*, 8 (10):e77589, 2013.

[21] R. C. Taylor, B. A. Klein, J. Stein, and M. J. Ryan. Faux frogs: multimodal signalling and the value of robotics in animal behaviour. *Animal Behaviour*, 2008.

[22] D. Romano, G. Benelli, E. Donati, D. Remorini, A. Canale, and C. Stefanini. Multiple cues produced by a robotic fish modulate aggressive behaviour in siamese fighting fishes. *Scientific reports*, 7(1):4667, 2017.

[23] G. Benelli, D. Romano, G. Rocchigiani, A. Caselli, F. Mancianti, A. Canale, and C. Stefanini. Behavioral asymmetries in ticks–lateralized questing of ixodes ricinus to a mechatronic apparatus delivering host-borne cues. *Acta tropica*, 178:176–181, 2018.

[24] S. Marras and M. Porfiri. Fish and robots swimming together: attraction towards the robot demands biomimetic locomotion. *Journal of The Royal Society Interface*, 9(73):1856–1868, 2012.

[25] T. Landgraf, D. Bierbach, H. Nguyen, N. Muggelberg, P. Romanczuk, and J. Krause. Robofish: increased acceptance of interactive robotic fish with realistic eyes and natural motion patterns by live trinidadian guppies. *Bioinspiration & biomimetics*, 11(1):015001, 2016.

[26] D. Romano, M. Porfiri, P. Zahadat, and T. Schmickl. Animal–robot interaction—an emerging field at the intersection of biology and robotics. *Bioinspiration & Biomimetics*, 19(2):020201, 2024.

[27] B. A. Klein, J. Stein, and R. C. Taylor. Robots in the service of animal behavior. *Communicative & integrative biology*, 5(5):466–472, 2012.

[28] E. P. Martins, T. J. Ord, and S. W. Davenport. Combining motions into complex displays: playbacks with a robotic lizard. *Behavioral Ecology and Sociobiology*, 58:351–360, 2005.

[29] E. Fernández-Juricic, N. Gilak, J. C. Mcdonald, P. Pithia, and A. Valcarcel. A dynamic method to study the transmission of social foraging information in flocks using robots. *Animal Behaviour*, 71(4):901–911, 2006.

[30] N. Horsevad, H. L. Kwa, and R. Bouffanais. Beyond bio-inspired robotics: how multi-robot systems can support research on collective animal behavior. *Frontiers in Robotics and AI*, 9:865414, 2022.

[31] J. Halloy, G. Sempo, G. Caprari, C. Rivault, M. Asadpour, F. Tâche, I. Saïd, V. Durier, S. Canonge, J. M. Amé, et al. Social integration of robots into groups of cockroaches to control self-organized choices. *Science*, 318(5853):1155–1158, 2007.

[32] M. Aureli, F. Fiorilli, and M. Porfiri. Portraits of self-organization in fish schools interacting with robots. *Physica D: Nonlinear Phenomena*, 241(9):908–920, 2012.

[33] C. Kim, T. Ruberto, P. Phamduy, and M. Porfiri. Closed-loop control of zebrafish behaviour in three dimensions using a robotic stimulus. *Scientific reports*, 8(1):657, 2018.

[34] S. Butail, T. Bartolini, and M. Porfiri. Collective response of zebrafish shoals to a free-swimming robotic fish. *PLoS One*, 8(10):e76123, 2013.

[35] S. Butail, G. Polverino, P. Phamduy, F. Del Sette, and M. Porfiri. Influence of robotic shoal size, configuration, and activity on zebrafish behavior in a free-swimming environment. *Behavioural brain research*, 275:269–280, 2014.

[36] N. Abaid, S. Marras, C. Fitzgibbons, and M. Porfiri. Modulation of risk-taking behaviour in golden shiners (notemigonus crysoleucas) using robotic fish. *Behavioural processes*, 100:9–12, 2013.

[37] J. J. Faria, J. R. Dyer, R. O. Clément, I. D. Couzin, N. Holt, A. J. Ward, D. Waters, and J. Krause. A novel method for investigating the collective behaviour of fish: introducing 'robofish'. *Behavioral Ecology and Sociobiology*, 64:1211–1218, 2010.

[38] A. J. Ward, D. J. Sumpter, I. D. Couzin, P. J. Hart, and J. Krause. Quorum decision-making facilitates information transfer in fish shoals. *Proceedings of the National Academy of Sciences*, 105(19):6948–6953, 2008.

[39] G. Cord-Cruz, T. Ruberto, D. Neri, and M. Porfiri. Zebrafish response to live predator and biologically-inspired robot in a circular arena. In *Bioinspiration, Biomimetics, and Bioreplication 2017*, volume 10162, pages 120–125. SPIE, 2017.

[40] R. El Khoury, R. B. Ventura, G. Cord-Cruz, T. Ruberto, and M. Porfiri. Interactive experiments in a robotics-based platform to simulate zebrafish response to a predator. In *Bioinspiration, Biomimetics, and Bioreplication VIII*, volume 10593, pages 134–140. SPIE, 2018.

[41] R. Gerlai. Antipredatory behavior of zebrafish: adaptive function and a tool for translational research. *Evolutionary Psychology*, 11(3):147470491301100308, 2013.

[42] C. Spinello, Y. Yang, S. Macrì, and M. Porfiri. Zebrafish adjust their behavior in response to an interactive robotic predator. *Frontiers in Robotics and AI*, 6:38, 2019.

[43] F. Ladu, T. Bartolini, S. G. Panitz, F. Chiarotti, S. Butail, S. Macrì, and M. Porfiri. Live predators, robots, and computer-animated images elicit differential avoidance responses in zebrafish. *Zebrafish*, 12(3):205–214, 2015.

[44] S. Mitri, S. Wischmann, D. Floreano, and L. Keller. Using robots to understand social behaviour. *Biological Reviews*, 88(1):31–39, 2013.

[45] J. Abdai and A. Miklosi. *An Introduction to Ethorobotics: Robotics and the Study of Animal Behaviour*. Taylor & Francis, 2024.

[46] E. Klinghammer and M. W. Fox. Ethology and its place in animal science. *Journal of animal science*, 32(6):1278–1283, 1971.

[47] E. J. Fernandez and A. L. Martin. Animal training, environmental enrichment, and animal welfare: A history of behavior analysis in zoos. *Journal of Zoological and Botanical Gardens*, 2(4):531–543, 2021.

[48] S. Waiblinger, J. Baumgartner, M. Kiley-Worthington, and K. Niebuhr. Applied ethology: the basis for improved animal welfare in organic farming. 2004.

[49] H. Sambraus. Applied ethology—it's task and limits in veterinary practice. *Applied Animal Behaviour Science*, 59(1-3):39–48, 1998.

[50] U. Özentürk, Z. Chen, L. Jamone, and E. Versace. Robotics for poultry farming: Challenges and opportunities. *Computers and Electronics in Agriculture*, 226:109411, 2024.

[51] R. Sparrow and M. Howard. Robots in agriculture: prospects, impacts, ethics, and policy. *precision agriculture*, 22:818–833, 2021.

[52] C. Cheng, J. Fu, H. Su, and L. Ren. Recent advancements in agriculture robots: Benefits and challenges. *Machines*, 11(1):48, 2023.

[53] K. Arts, R. Van der Wal, and W. M. Adams. Digital technology and the conservation of nature. *Ambio*, 44:661–673, 2015.

[54] S. W. Breck, J. T. Schultz, D. Prause, C. Krebs, A. J. Giordano, and B. Boots. Integrating robotics into wildlife conservation: testing improvements to predator deterrents through movement. *PeerJ*, 11:e15491, 2023.

[55] M. Chellapurath, P. C. Khandelwal, and A. K. Schulz. Bioinspired robots can foster nature conservation. *Frontiers in Robotics and AI*, 10:1145798, 2023.

[56] G. Li, X. Hui, Y. Zhao, W. Zhai, J. L. Purswell, Z. Porter, S. Poudel, L. Jia, B. Zhang, and G. D. Chesser. Effects of ground robot manipulation on hen floor egg reduction, production performance, stress response, bone quality, and behavior. *Plos one*, 17(4): e0267568, 2022.

[57] G. Abbas, S. Jaffery, A. Hashmi, A. Tanveer, M. Arshad, Q. Amin, M. Saeed, M. Saleem, R. Qureshi, A. Khan, et al. Prospects and challenges of adopting and implementing smart technologies in poultry production. *Pakistan Journal of Science*, 74(2), 2022.

[58] J. Edgar, S. Held, C. Jones, and C. Troisi. Influences of maternal care on chicken welfare. *Animals*, 6(1):2, 2016.

[59] J.-M. Lien and E. Pratt. Interactive planning for shepherd motion. In *AAAI Spring Symposium: Agents that Learn from Human Teachers*, pages 95–102, 2009.

[60] A. Monteiro, S. Santos, and P. Gonçalves. Precision agriculture for crop and livestock farming—brief review. *Animals*, 11(8):2345, 2021.

[61] Z. Butler, P. Corke, R. Peterson, and D. Rus. From robots to animals: Virtual fences for controlling cattle. *The International Journal of Robotics Research*, 25(5-6):485–508, 2006.

[62] J.-M. Lien, O. B. Bayazit, R. T. Sowell, S. Rodriguez, and N. M. Amato. Shepherding behaviors. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 4, pages 4159–4164. IEEE, 2004.

[63] B. Bat-Erdene and O.-E. Mandakh. Shepherding algorithm of multi-mobile robot system. In *2017 First IEEE International Conference on Robotic Computing (IRC)*, pages 358–361. IEEE, 2017.

[64] M. Evered, P. Burling, M. Trotter, et al. An investigation of predator response in robotic herding of sheep. *International Proceedings of Chemical, Biological and Environmental Engineering*, 63:49–54, 2014.

[65] K. J. Yaxley, N. McIntyre, J. Park, and J. Healey. Sky shepherds: A tale of a uav and sheep. *Shepherding UxVs for Human-Swarm Teaming: An Artificial Intelligence Approach to Unmanned X Vehicles*, pages 189–206, 2021.

[66] H. Anzai and H. Sakurai. Preliminary study on the application of robotic herding to manipulation of grazing distribution: Behavioral response of cattle to herding by an unmanned vehicle and its manipulation performance. *Applied Animal Behaviour Science*, 256:105751, 2022.

[67] R. Vaughan, N. Sumpter, J. Henderson, A. Frost, and S. Cameron. Experiments in automatic flock control. *Robotics and autonomous systems*, 31(1-2):109–117, 2000.

[68] A. A. Paranjape, S.-J. Chung, K. Kim, and D. H. Shim. Robotic herding of a flock of birds using an unmanned aerial vehicle. *IEEE Transactions on Robotics*, 34(4):901–915, 2018.

[69] S. Gade, A. A. Paranjape, and S.-J. Chung. Herding a flock of birds approaching an airport using an unmanned aerial vehicle. In *AIAA Guidance, Navigation, and Control Conference*, page 1540, 2015.

[70] M. Li, Z. Hu, J. Liang, and S. Li. Shepherding behaviors with single shepherd in crowd management. In *System Simulation and Scientific Computing: International Conference, ICSC 2012, Shanghai, China, October 27-30, 2012. Proceedings, Part I*, pages 415–423. Springer, 2012.

[71] A. Rahman, J. Jin, A. Cricenti, A. Rahman, M. Palaniswami, and T. Luo. Cloud-enhanced robotic system for smart city crowd control. *Journal of Sensor and Actuator Networks*, 5(4):20, 2016.

[72] A. Garrell, A. Sanfeliu, and F. Moreno-Noguer. Discrete time motion model for guiding people in urban areas using multiple robots. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 486–491. IEEE, 2009.

[73] D. A. Shell and M. J. Mataric. Directional audio beacon deployment: an assistive multi-robot application. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 3, pages 2588–2594. IEEE, 2004.

[74] L. Coppinger, R. Coppinger, et al. Dogs for herding and guarding livestock. *Livestock handling and transport*, 13:235–253, 1993.

[75] A. J. King, A. M. Wilson, S. D. Wilshin, J. Lowe, H. Haddadi, S. Hailes, and A. J. Morton. Selfish-herd behaviour of sheep under threat. *Current Biology*, 22(14):R561–R562, 2012.

[76] W. D. Hamilton. Geometry for the selfish herd. *Journal of theoretical Biology*, 31(2):295–311, 1971.

[77] N. K. Long, K. Sammut, D. Sgarioto, M. Garratt, and H. A. Abbass. A comprehensive review of shepherding as a bio-inspired swarm-robotics guidance approach. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 4(4):523–537, 2020.

[78] S. Cecchi, S. Spinsante, A. Terenzi, and S. Orcioni. A smart sensor-based measurement system for advanced bee hive monitoring. *Sensors*, 20(9):2726, 2020.

[79] S. G. Potts, J. C. Biesmeijer, C. Kremen, P. Neumann, O. Schweiger, and W. E. Kunin. Global pollinator declines: trends, impacts and drivers. *Trends in ecology & evolution*, 25 (6):345–353, 2010.

[80] A. J. Vanbergen, I. P. Initiative, et al. Threats to an ecosystem service: pressures on pollinators. *Frontiers in Ecology and the Environment*, 11(5):251–259, 2013.

[81] J. J. Lahoz-Monfort and M. J. Magrath. A comprehensive overview of technologies for species and habitat monitoring and conservation. *BioScience*, 71(10):1038–1062, 2021.

[82] F. Angelini, P. Angelini, C. Angiolini, S. Bagella, F. Bonomo, M. Caccianiga, C. Della Santina, D. Gigante, M. Hutter, T. Nanayakkara, et al. Robotic monitoring of habitats: The natural intelligence approach. *IEEE Access*, 2023.

[83] D. Western. Human-modified ecosystems and future evolution. *Proceedings of the National Academy of Sciences*, 98(10):5458–5465, 2001.

[84] S. E. Ledger, J. Loh, R. Almond, M. Böhm, C. F. Clements, J. Currie, S. Deinet, T. Galewski, M. Grooten, M. Jenkins, et al. Past, present, and future of the living planet index. *npj Biodiversity*, 2(1):12, 2023.

[85] S. Naeem, R. Chazdon, J. E. Duffy, C. Prager, and B. Worm. Biodiversity and human well-being: an essential link for sustainable development. *Proceedings of the Royal Society B: Biological Sciences*, 283(1844):20162091, 2016.

[86] B. J. Cardinale, J. E. Duffy, A. Gonzalez, D. U. Hooper, C. Perrings, P. Venail, A. Narwani, G. M. Mace, D. Tilman, D. A. Wardle, et al. Biodiversity loss and its impact on humanity. *Nature*, 486(7401):59–67, 2012.

[87] E. Chivian and A. Bernstein. *Sustaining life: how human health depends on biodiversity*. Oxford University Press, 2008.

[88] S. Díaz, J. Fargione, F. S. Chapin III, and D. Tilman. Biodiversity loss threatens human well-being. *PLoS biology*, 4(8):e277, 2006.

[89] A. K. Duraiappah, S. Naeem, T. Agardy, N. J. Ash, H. D. Cooper, S. Díaz, D. P. Faith, G. Mace, J. A. McNeely, H. A. Mooney, et al. Ecosystems and human well-being: biodiversity synthesis; a report of the millennium ecosystem assessment. 2005.

[90] J. Rockström, W. Steffen, K. Noone, Å. Persson, F. S. Chapin, E. F. Lambin, T. M. Lenton, M. Scheffer, C. Folke, H. J. Schellnhuber, et al. A safe operating space for humanity. *nature*, 461(7263):472–475, 2009.

[91] O. K. Lawler, H. L. Allan, P. W. Baxter, R. Castagnino, M. C. Tor, L. E. Dann, J. Hungerford, D. Karmacharya, T. J. Lloyd, M. J. López-Jara, et al. The covid-19 pandemic is intricately linked to biodiversity loss and ecosystem health. *The Lancet Planetary Health*, 5(11):e840–e850, 2021.

[92] S. Platto, J. Zhou, Y. Wang, H. Wang, and E. Carafoli. Biodiversity loss and covid-19 pandemic: The role of bats in the origin and the spreading of the disease. *Biochemical and biophysical research communications*, 538:2–13, 2021.

[93] F. Keesing and R. S. Ostfeld. Impacts of biodiversity and biodiversity loss on zoonotic diseases. *Proceedings of the National Academy of Sciences*, 118(17):e2023540118, 2021.

[94] T. Sen and S. K. Samanta. Medicinal plants, human health and biodiversity: a broad review. *Biotechnological applications of biodiversity*, pages 59–110, 2015.

[95] R. Rajeswara, K. Syamasundar, D. Rajput, G. Nagaraju, and G. Adinarayana. Biodiversity, conservation and cultivation of medicinal plants. *World*, 422000(77000):18–2, 2012.

[96] R. R. N. Alves and U. P. Albuquerque. Animals as a source of drugs: bioprospecting and biodiversity conservation. In *Animals in traditional folk medicine: implications for conservation*, pages 67–89. Springer, 2012.

[97] R. Aerts, O. Honnay, and A. Van Nieuwenhuyse. Biodiversity and human health: mechanisms and evidence of the positive health effects of diversity in nature and green spaces. *British medical bulletin*, 127(1):5–22, 2018.

[98] G. A. Rook. Regulation of the immune system by biodiversity from the natural environment: an ecosystem service essential to health. *Proceedings of the National Academy of Sciences*, 110(46):18360–18367, 2013.

[99] T. Haahtela. A biodiversity hypothesis. *Allergy*, 74(8):1445–1456, 2019.

[100] J. Bélanger and D. Pilling. *The state of the world's biodiversity for food and agriculture*. FAO;, 2019.

[101] J. W. Erisman, N. van Eekeren, J. de Wit, C. Koopmans, W. Cuijpers, N. Oerlemans, B. J. Koks, et al. Agriculture and biodiversity: a better balance benefits both. *AIMS Agriculture and Food*, 1(2):157–174, 2016.

[102] L. G. Carvalheiro, R. Veldtman, A. G. Shenkute, G. B. Tesfay, C. W. W. Pirk, J. S. Donaldson, and S. W. Nicolson. Natural and within-farmland biodiversity enhances crop productivity. *Ecology letters*, 14(3):251–259, 2011.

[103] A. M. Kerr and A. H. Baird. Natural barriers to natural disasters. *BioScience*, 57(2): 102–103, 2007.

[104] B. V. Li, C. N. Jenkins, and W. Xu. Strategic protection of landslide vulnerable mountains for biodiversity conservation under land-cover and climate change impacts. *Proceedings of the National Academy of Sciences*, 119(2):e2113416118, 2022.

[105] J. Luan, S. Li, S. Liu, Y. Wang, L. Ding, H. Lu, L. Chen, J. Zhang, W. Zhou, S. Han, et al. Biodiversity mitigates drought effects in the decomposer system across biomes. *Proceedings of the National Academy of Sciences*, 121(13):e2313334121, 2024.

[106] M. H. Daba and S. W. Dejene. The role of biodiversity and ecosystem services in carbon sequestration and its implication for climate change mitigation. *Environmental Sciences and Natural Resources*, 11(2):1–10, 2018.

[107] M. A. Huston and G. Marland. Carbon management and biodiversity. *Journal of environmental Management*, 67(1):77–86, 2003.

[108] X. Jing, S. Jiang, H. Liu, Y. Li, and J.-S. He. Complex relationships and feedback mechanisms between climate change and biodiversity. *Biodiversity Science*, 30(10), 2022.

[109] G. Ceballos, P. R. Ehrlich, A. D. Barnosky, A. García, R. M. Pringle, and T. M. Palmer. Accelerated modern human–induced species losses: Entering the sixth mass extinction. *Science advances*, 1(5):e1400253, 2015.

[110] I. E. Hendriks, C. M. Duarte, and C. H. Heip. Biodiversity research still grounded, 2006.

[111] W. Mengist, T. Soromessa, and G. L. Feyisa. A global view of regulatory ecosystem services: Existed knowledge, trends, and research gaps. *Ecological Processes*, 9:1–14, 2020.

[112] L. Ball, S. Rodríguez-Machado, D. Paredes-Burneo, S. Rutledge, D. A. Boyd, D. Vander Pluym, S. Babb-Biernacki, A. S. Chipps, R. Ç. Öztürk, Y. Terzi, et al. What 'unexplored'means: mapping regions with digitized natural history records to look for 'biodiversity blindspots'. *PeerJ*, 13:e18511, 2025.

[113] G. F. Ficetola, C. Canedoli, and F. Stoch. The racovitzan impediment and the hidden biodiversity of unexplored environments. *Conservation Biology*, 33(1):214–216, 2019.

[114] R. Watson, I. Baste, A. Larigauderie, P. Leadley, U. Pascual, B. Baptiste, S. Demissew, L. Dziba, G. Erpul, A. Fazel, et al. Summary for policymakers of the global assessment report on biodiversity and ecosystem services of the intergovernmental science-policy platform on biodiversity and ecosystem services. *IPBES Secretariat: Bonn, Germany*, pages 22–47, 2019.

[115] N. Loiseau, D. Mouillot, L. Velez, R. Seguin, N. Casajus, C. Coux, C. Albouy, T. Claverie, A. Duhamet, V. Fleure, et al. Inferring the extinction risk of marine fish to inform global conservation priorities. *PLoS Biology*, 22(8):e3002773, 2024.

[116] Convention on Biological Diversity. Kunming-montreal global biodiversity framework targets. https://www.cbd.int/gbf/targets, 2022. Accessed: 2024-01-24.

[117] E. Ramírez-Llodra and D. S. Billett. Deep-sea ecosystems: pristine biodiversity reservoir and technological challenge. 2006.

[118] D. C. Culver and T. Pipan. *Shallow subterranean habitats: ecology, evolution, and conservation*. Oxford University Press, USA, 2014.

[119] K. Bastmeijer. Protecting polar wilderness: Just a western philosophical idea or a useful concept for regulating human activities in the polar regions? *The Yearbook of Polar Law Online*, 1(1):73–99, 2009.

[120] C. Körner. Conservation of mountain biodiversity in the context of climate change. In *International Mountain Biodiversity Conference*, page 21, 2009.

[121] J. A. Zwerts, P. Stephenson, F. Maisels, M. Rowcliffe, C. Astaras, P. A. Jansen, J. van Der Waarde, L. E. Sterck, P. A. Verweij, T. Bruce, et al. Methods for wildlife monitoring in tropical forests: Comparing human observations, camera traps, and passive acoustic sensors. *Conservation Science and Practice*, 3(12):e568, 2021.

[122] C. Larrosa, L. R. Carrasco, and E. Milner-Gulland. Unintended feedbacks: challenges and opportunities for improving conservation effectiveness. *Conservation Letters*, 9(5): 316–326, 2016.

[123] O. Berger-Tal, B. Wong, U. Candolin, and J. Barber. What evidence exists on the effects of anthropogenic noise on acoustic communication in animals? a systematic map protocol. *Environmental Evidence*, 8(1):1–7, 2019.

[124] D. Breed, L. C. Meyer, J. C. Steyl, A. Goddard, R. Burroughs, and T. A. Kohn. Conserving wildlife in a changing world: Understanding capture myopathy—a malignant outcome of stress during capture and translocation. *Conservation Physiology*, 7(1):coz027, 2019.

[125] S. de Faria Lopes. The other side of ecology: thinking about the human bias in our ecological analyses for biodiversity conservation. *Ethnobiology and Conservation*, 6, 2017.

[126] J. S. Lewis, S. Spaulding, H. Swanson, W. Keeley, A. R. Gramza, S. VandeWoude, and K. R. Crooks. Human activity influences wildlife populations and activity patterns: implications for spatial and temporal refuges. *Ecosphere*, 12(5):e03487, 2021.

[127] H. J. Albers, K. D. Lee, J. R. Rushlow, and C. Zambrana-Torrselio. Disease risk from human–environment interactions: environment and development economics for joint conservation-health policy. *Environmental and Resource Economics*, 76(4):929–944, 2020.

[128] M. Álvarez-González, P. Suarez-Bregua, G. J. Pierce, and C. Saavedra. Unmanned aerial vehicles (uavs) in marine mammal research: A review of current applications and challenges. *Drones*, 7(11):667, 2023.

[129] E. Danson. The economics of scale: Using autonomous underwater vehicles (auvs) for wide-area hydrographic survey and ocean data acquisition. In *FIG XXII International Congress, Washington, Dc USA*, 2002.

[130] P. I. Macreadie, D. L. McLean, P. G. Thomson, J. C. Partridge, D. O. Jones, A. R. Gates, M. C. Benfield, S. P. Collin, D. J. Booth, L. L. Smith, et al. Eyes in the sea: unlocking the mysteries of the ocean using industrial, remotely operated vehicles (rovs). *Science of the Total Environment*, 634:1077–1091, 2018.

[131] X. Gao, J. Li, L. Fan, Q. Zhou, K. Yin, J. Wang, C. Song, L. Huang, and Z. Wang. Review of wheeled mobile robots' navigation problems and application prospects in agriculture. *Ieee Access*, 6:49248–49268, 2018.

[132] F. Christiansen, L. Rojano-Doñate, P. T. Madsen, and L. Bejder. Noise levels of multirotor unmanned aerial vehicles with implications for potential underwater impacts on marine mammals. *Frontiers in Marine Science*, 3:277, 2016.

[133] M. Chellapurath, K. L. Walker, E. Donato, G. Picardi, S. Stefanni, C. Laschi, F. Giorgio-Serchi, and M. Calisti. Analysis of station keeping performance of an underwater legged robot. *IEEE/ASME Transactions on Mechatronics*, 27(5):3730–3741, 2021.

[134] N. Savage. Bioinspired robots walk, swim, slither and fly. *Nature*, 2022.

[135] J. Abdai and Á. Miklósi. Poking the future: When should we expect that animal-robot interaction becomes a routine method in the study of behavior. *Anim Behav Cogn*, 5(4): 321–325, 2018.

[136] R. F. Storms, C. Carere, R. Musters, H. Van Gasteren, S. Verhulst, and C. K. Hemelrijk. Deterrence of birds with an artificial predator, the robotfalcon. *Journal of the Royal Society Interface*, 19(195):20220497, 2022.

[137] G. A. Folkertsma, W. Straatman, N. Nijenhuis, C. H. Venner, and S. Stramigioli. Robird: a robotic bird of prey. *IEEE robotics & automation magazine*, 24(3):22–29, 2017.

[138] J. M. Anderson and N. K. Chhabra. Maneuvering and stability performance of a robotic tuna. *Integrative and comparative biology*, 42(1):118–126, 2002.

[139] P. H. Nam Anh, H.-S. Choi, J. Huang, R. Zhang, and J. Kim. Study on oscillatory and undulatory motion of robotic fish. *Applied Sciences*, 14(8):3239, 2024.

[140] G. V. Lauder and E. D. Tytell. Hydrodynamics of undulatory propulsion. *Fish physiology*, 23:425–468, 2005.

[141] B. Sun, W. Li, Z. Wang, Y. Zhu, Q. He, X. Guan, G. Dai, D. Yuan, A. Li, W. Cui, et al. Recent progress in modeling and control of bio-inspired fish robots. *Journal of Marine Science and Engineering*, 10(6):773, 2022.

[142] J. Sattar, G. Dudek, O. Chiu, I. Rekleitis, P. Giguere, A. Mills, N. Plamondon, C. Prahacs, Y. Girdhar, M. Nahon, et al. Enabling autonomous capabilities in underwater robotics. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3628–3634. IEEE, 2008.

[143] O. Khatib, X. Yeh, G. Brantner, B. Soe, B. Kim, S. Ganguly, H. Stuart, S. Wang, M. Cutkosky, A. Edsinger, et al. Ocean one: A robotic avatar for oceanic discovery. *IEEE Robotics & Automation Magazine*, 23(4):20–29, 2016.

[144] G. Brantner and O. Khatib. Controlling ocean one: Human–robot collaboration for deep-sea manipulation. *Journal of Field Robotics*, 38(1):28–51, 2021.

[145] H. Stuart, S. Wang, O. Khatib, and M. R. Cutkosky. The ocean one hands: An adaptive design for robust marine manipulation. *The International Journal of Robotics Research*, 36 (2):150–166, 2017.

[146] G. Picardi, C. Laschi, and M. Calisti. Model-based open loop control of a multigait legged underwater robot. *Mechatronics*, 55:162–170, 2018.

[147] B. Gaßmann, K.-U. Scholl, and K. Berns. Locomotion of lauron iii in rough terrain. In *2001 IEEE/ASME International Conference on Advanced Intelligent Mechatronics. Proceedings (Cat. No. 01TH8556)*, volume 2, pages 959–964. IEEE, 2001.

[148] A. Rönnau, G. Heppner, T. Kerscher, and R. Dillmann. Fault diagnosis and system status monitoring for a six-legged walking robot. In *2011 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 874–879. IEEE, 2011.

[149] M. Stefanec, D. N. Hofstadler, T. Krajník, A. E. Turgut, H. Alemdar, B. Lennox, E. Şahin, F. Arvin, and T. Schmickl. A minimally invasive approach towards "ecosystem hacking" with honeybees. *Frontiers in Robotics and AI*, 9:791921, 2022.

[150] A. Ilgün, K. Angelov, M. Stefanec, S. Schönwetter-Fuchs, V. Stokanic, J. Vollmann, D. N. Hofstadler, M. H. Kärcher, H. Mellmann, V. Taliaronak, et al. Bio-hybrid systems for ecosystem level effects. In *Artificial Life Conference Proceedings 33*, volume 2021, page 41. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info . . . , 2021.

[151] T. Schmickl, M. Szopek, F. Mondada, R. Mills, M. Stefanec, D. N. Hofstadler, D. Lazic, R. Barmak, F. Bonnet, and P. Zahadat. Social integrating robots suggest mitigation strategies for ecosystem decay. *Frontiers in Bioengineering and Biotechnology*, 9:612605, 2021.

[152] B. Tian and J. Yu. Current status and prospects of marine renewable energy applied in ocean robots. *International Journal of Energy Research*, 43(6):2016–2031, 2019.

[153] J. Iqbal, A. Al-Zahrani, S. A. Alharbi, and A. Hashmi. Robotics inspired renewable energy developments: prospective opportunities and challenges. *IEEE Access*, 7:174898–174923, 2019.

[154] I. R. Mohamed, A. S. Shehata, W. M. El-Maghlany, and M. A. Kotb. Experiment study on harvesting ocean thermal energy using phase change material for autonomous underwater vehicle powering. In *AIP Conference Proceedings*, volume 2769. AIP Publishing, 2023.

[155] F. Hartmann, M. Baumgartner, and M. Kaltenbrunner. Becoming sustainable, the new frontier in soft robotics. *Advanced Materials*, 33(19):2004413, 2021.

[156] J. Kim, H. Park, and C. Yoon. Advances in biodegradable soft robots. *Polymers*, 14(21): 4574, 2022.

[157] C. M. Donatelli, S. A. Bradner, J. Mathews, E. Sanders, C. Culligan, D. Kaplan, and E. D. Tytell. Prototype of a fish inspired swimming silk robot. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 60–65. IEEE, 2018.

[158] Y. Cong, C. Gu, T. Zhang, and Y. Gao. Underwater robot sensing technology: A survey. *Fundamental Research*, 1(3):337–345, 2021.

[159] Z. Zhou, J. Liu, and J. Yu. A survey of underwater multi-robot systems. *IEEE/CAA Journal of Automatica Sinica*, 9(1):1–18, 2021.

[160] G. Liu, A. Wang, X. Wang, and P. Liu. A review of artificial lateral line in sensor fabrication and bionic applications for robot fish. *Applied bionics and biomechanics*, 2016(1): 4732703, 2016.

[161] B. J. Wolf, J. A. Morton, W. N. MacPherson, and S. M. Van Netten. Bio-inspired all-optical artificial neuromast for 2d flow sensing. *Bioinspiration & biomimetics*, 13(2): 026013, 2018.

[162] L. DeVries, F. D. Lagor, H. Lei, X. Tan, and D. A. Paley. Distributed flow estimation and closed-loop control of an underwater vehicle with a multi-modal artificial lateral line. *Bioinspiration & biomimetics*, 10(2):025002, 2015.

[163] H. Beem, M. Hildner, and M. Triantafyllou. Calibration and validation of a harbor seal whisker-inspired flow sensor. *Smart Materials and Structures*, 22(1):014012, 2012.

[164] J. Z. Gul, K. Y. Su, and K. H. Choi. Fully 3d printed multi-material soft bio-inspired whisker sensor for underwater-induced vortex detection. *Soft robotics*, 5(2):122–132, 2018.

[165] M. Ilami, H. Bagheri, R. Ahmed, E. O. Skowronek, and H. Marvi. Materials, actuators, and sensors for soft bioinspired robots. *Advanced Materials*, 33, 2020. URL https://api.semanticscholar.org/CorpusID:229341622.

[166] Z. Ye, L. Zheng, W. Chen, B. Wang, and L. Zhang. Recent advances in bioinspired soft robots: fabrication, actuation, tracking, and applications. *Advanced Materials Technologies*, 9(21):2301862, 2024.

[167] S. Panda, S. Hajra, P. M. Rajaitha, and H. J. Kim. Stimuli-responsive polymer-based bioinspired soft robots. *Micro and Nano Systems Letters*, 11(1):2, 2023.

[168] J. Qu, Y. Xu, Z. Li, Z. Yu, B. Mao, Y. Wang, Z. Wang, Q. Fan, X. Qian, M. Zhang, et al. Recent advances on underwater soft robots. *Advanced Intelligent Systems*, 6(2):2300299, 2024.

[169] S. M. Youssef, M. Soliman, M. A. Saleh, M. A. Mousa, M. Elsamanty, and A. G. Radwan. Underwater soft robotics: A review of bioinspiration in design, actuation, modeling, and control. *Micromachines*, 13(1):110, 2022.

[170] Z. Ye, P. Hou, Z. Chen, and I. Member. 2d maneuverable robotic fish propelled by multiple ionic polymer–metal composite artificial fins. *International Journal of intelligent robotics and applications*, 1(2):195–208, 2017.

[171] A. Cloitre, V. Subramaniam, N. Patrikalakis, and P. V. y Alvarado. Design and control of a field deployable batoid robot. In *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pages 707–712. IEEE, 2012.

[172] C. Christianson, C. Bayag, G. Li, S. Jadhav, A. Giri, C. Agba, T. Li, and M. T. Tolley. Jellyfish-inspired soft robot driven by fluid electrode dielectric organic robotic actuators. *Frontiers in Robotics and AI*, 6:126, 2019.

[173] S. Mao, E. Dong, H. Jin, M. Xu, S. Zhang, J. Yang, and K. H. Low. Gait study and pattern generation of a starfish-like soft robot with flexible rays actuated by smas. *Journal of Bionic Engineering*, 11(3):400–411, 2014.

[174] J. Zhang, Y. Chen, Y. Liu, and Y. Gong. Dynamic modeling of underwater snake robot by hybrid rigid-soft actuation. *Journal of Marine Science and Engineering*, 10(12):1914, 2022.

[175] J. S. Jaffe, P. J. Franks, P. L. Roberts, D. Mirza, C. Schurgers, R. Kastner, and A. Boch. A swarm of autonomous miniature underwater robot drifters for exploring submesoscale ocean dynamics. *Nature communications*, 8(1):14189, 2017.

[176] H. Duan, M. Huo, and Y. Fan. From animal collective behaviors to swarm robotic cooperation. *National Science Review*, 10(5):nwad040, 2023.

[177] A. Ayali and G. A. Kaminka. The hybrid bio-robotic swarm as a powerful tool for collective motion research: a perspective. *Frontiers in Neurorobotics*, 17:1215085, 2023.

[178] J. Krause, G. D. Ruxton, and S. Krause. Swarm intelligence in animals and humans. *Trends in ecology & evolution*, 25(1):28–34, 2010.

[179] A. Dussutour, V. Fourcassié, D. Helbing, and J.-L. Deneubourg. Optimal traffic organization in ants under crowded conditions. *Nature*, 428(6978):70–73, 2004.

[180] H. Hildenbrandt, C. Carere, and C. K. Hemelrijk. Self-organized aerial displays of thousands of starlings: a model. *Behavioral Ecology*, 21(6):1349–1359, 2010.

[181] D. Pavlov, A. Kasumyan, et al. Patterns and mechanisms of schooling behavior in fish: a review. *Journal of Ichthyology*, 40(2):S163, 2000.

[182] G. Beni and J. Wang. Swarm intelligence in cellular robotic systems. In *Robots and biological systems: towards a new bionics?*, pages 703–712. Springer, 1993.

[183] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraula, and E. Bonabeau. Self-organization in biological systems. In *Self-Organization in Biological Systems*. Princeton university press, 2020.

[184] C. K. Hemelrijk and H. Hildenbrandt. Schools of fish and flocks of birds: their shape and internal structure by self-organization. *Interface focus*, 2(6):726–737, 2012.

[185] E. Bonabeau. Swarm intelligence: From natural to artificial systems. *Oxford University Press google schola*, 2:25–34, 1999.

[186] S. Garnier, J. Gautrais, and G. Theraulaz. The biological principles of swarm intelligence. *Swarm intelligence*, 1:3–31, 2007.
[187] X.-S. Yang, S. Deb, Y.-X. Zhao, S. Fong, and X. He. Swarm intelligence: past, present and future. *Soft Computing*, 22:5923–5933, 2018.
[188] C. Blum and X. Li. Swarm intelligence in optimization. In *Swarm intelligence: introduction and applications*, pages 43–85. Springer, 2008.
[189] Z. Zhang, K. Long, J. Wang, and F. Dressler. On swarm intelligence inspired self-organized networking: its bionic mechanisms, designing principles and optimization approaches. *IEEE Communications Surveys & Tutorials*, 16(1):513–537, 2013.
[190] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7:1–41, 2013.
[191] C. Blum and D. Merkle. *Swarm intelligence: introduction and applications*. Springer Science & Business Media, 2008.
[192] M. Dorigo, G. Theraulaz, and V. Trianni. Swarm robotics: Past, present, and future [point of view]. *Proceedings of the IEEE*, 109(7):1152–1165, 2021.
[193] E. Şahin. Swarm robotics: From sources of inspiration to domains of application. In *International workshop on swarm robotics*, pages 10–20. Springer, 2004.
[194] M. Dorigo, G. Theraulaz, and V. Trianni. Reflections on the future of swarm robotics. *Science Robotics*, 5(49):eabe4385, 2020.
[195] I. Navarro and F. Matía. An introduction to swarm robotics. *International Scholarly Research Notices*, 2013(1):608164, 2013.
[196] Y. Tan and Z.-y. Zheng. Research advance in swarm robotics. *Defence Technology*, 9(1): 18–39, 2013.
[197] F. Berlinger, M. Gauci, and R. Nagpal. Implicit coordination for 3d underwater collective behaviors in a fish-inspired robot swarm. *Science Robotics*, 6(50), 2021.
[198] T. Vicsek. *Fluctuations and scaling in biology*. Oxford University Press New York, 2001.
[199] J. Nauta, P. Simoens, and Y. Khaluf. Group size and resource fractality drive multimodal search strategies: A quantitative analysis on group foraging. *Physica A: Statistical Mechanics and its Applications*, 590:126702, 2022.
[200] M. Saadat, F. Berlinger, A. Sheshmani, R. Nagpal, G. V. Lauder, and H. Haj-Hariri. Hydrodynamic advantages of in-line schooling. *Bioinspiration & Biomimetics*, 16(4):046002, 2021.
[201] A. E. Magurran. The adaptive significance of schooling as an anti-predator defence in fish. In *Annales Zoologici Fennici*, pages 51–66. JSTOR, 1990.
[202] A. Deutsch, G. Theraulaz, and T. Vicsek. Collective motion in biological systems, 2012.
[203] A. Okubo. Dynamical aspects of animal grouping: swarms, schools, flocks, and herds. *Advances in biophysics*, 22:1–94, 1986.
[204] T. Vicsek and A. Zafeiris. Collective motion. *Physics reports*, 517(3-4):71–140, 2012.
[205] D. J. Sumpter. The principles of collective animal behaviour. *Philosophical transactions of the royal society B: Biological Sciences*, 361(1465):5–22, 2006.
[206] U. Lopez, J. Gautrais, I. D. Couzin, and G. Theraulaz. From behavioural analyses to models of collective motion in fish schools. *Interface focus*, 2(6):693–707, 2012.
[207] I. Giardina. Collective behavior in animal groups: theoretical models and empirical studies. *HFSP journal*, 2(4):205–219, 2008.
[208] W. Lee and D. Kim. Autonomous shepherding behaviors of multiple target steering robots. *Sensors*, 17(12):2729, 2017.
[209] H. Kunz and C. K. Hemelrijk. Artificial fish schools: collective effects of school size, body size, and body form. *Artificial life*, 9(3):237–253, 2003.
[210] H.-S. Niwa. Self-organizing dynamic model of fish schooling. *Journal of theoretical Biology*, 171(2):123–136, 1994.
[211] H.-S. Niwa. Newtonian dynamical approach to fish schooling. *Journal of Theoretical Biology*, 181(1):47–63, 1996.
[212] A. Huth and C. Wissel. The simulation of fish schools in comparison with experimental data. *Ecological modelling*, 75:135–146, 1994.
[213] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.

[214] D. J. Hoare, I. D. Couzin, J.-G. Godin, and J. Krause. Context-dependent group size choice in fish. *Animal Behaviour*, 67(1):155–164, 2004.

[215] H. Pulliam and T. Caraco. Living in groups: Is there an optimal group size? in "behavioural ecology: An evolutionary approach"(jr krebs and nb davies, eds.). *Sinuaer, Sunderland, Mass*, pages 122–147, 1984.

[216] J. C. v. Olst and J. R. Hunter. Some aspects of the organization of fish schools. *Journal of the Fisheries Board of Canada*, 27(7):1225–1238, 1970.

[217] J. H. Tien, S. A. Levin, and D. I. Rubenstein. Dynamics of fish shoals: identifying key decision rules. *Evolutionary Ecology Research*, 6(4):555–565, 2004.

[218] Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, and I. D. Couzin. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725, 2011.

[219] J. E. Herbert-Read, A. Perna, R. P. Mann, T. M. Schaerf, D. J. Sumpter, and A. J. Ward. Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, 108(46):18726–18731, 2011.

[220] R. Lukeman, Y.-X. Li, and L. Edelstein-Keshet. Inferring individual rules from collective behavior. *Proceedings of the National Academy of Sciences*, 107(28):12576–12580, 2010.

[221] B. L. Partridge and T. J. Pitcher. The sensory basis of fish schools: relative roles of lateral line and vision. *Journal of comparative physiology*, 135:315–325, 1980.

[222] P. P. Klamser, L. Gómez-Nava, T. Landgraf, J. W. Jolles, D. Bierbach, and P. Romanczuk. Impact of variable speed on collective movement of animal groups. *Frontiers in Physics*, 9:715996, 2021.

[223] A. Huth and C. Wissel. The simulation of the movement of fish schools. *Journal of theoretical biology*, 156(3):365–385, 1992.

[224] H. Reuter and B. Breckling. Selforganization of fish schools: an object-oriented model. *Ecological Modelling*, 75:147–159, 1994.

[225] M. Camperi, A. Cavagna, I. Giardina, G. Parisi, and E. Silvestri. Spatially balanced topological interaction grants optimal cohesion in flocking models. *Interface focus*, 2(6): 715–725, 2012.

[226] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, et al. 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.

[227] C. K. Hemelrijk and H. Hildenbrandt. Some causes of the variable shape of flocks of birds. *PloS one*, 6(8):e22479, 2011.

[228] R. W. Tegeder and J. Krause. Density dependence and numerosity in fright stimulated aggregation behaviour of shoaling fish. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 350(1334):381–390, 1995.

[229] B. H. Lemasson, J. J. Anderson, and R. A. Goodwin. Collective motion in animal groups from a neurobiological perspective: the adaptive benefits of dynamic sensory loads and selective attention. *Journal of theoretical biology*, 261(4):501–510, 2009.

[230] A. Strandburg-Peshkin, C. R. Twomey, N. W. Bode, A. B. Kao, Y. Katz, C. C. Ioannou, S. B. Rosenthal, C. J. Torney, H. S. Wu, S. A. Levin, et al. Visual sensory networks and effective information transfer in animal groups. *Current Biology*, 23(17):R709–R711, 2013.

[231] B. G. Hogan, H. Hildenbrandt, N. E. Scott-Samuel, I. C. Cuthill, and C. K. Hemelrijk. The confusion effect when attacking simulated three-dimensional starling flocks. *Royal Society open science*, 4(1):160564, 2017.

[232] J. M. Jeschke and R. Tollrian. Prey swarming: which predators become confused and why? *Animal Behaviour*, 74(3):387–393, 2007.

[233] C. Ioannou, C. Tosh, L. Neville, and J. Krause. The confusion effect—from neural networks to reduced predation risk. *Behavioral Ecology*, 19(1):126–130, 2008.

[234] D. C. Krakauer. Groups confuse predators by exploiting perceptual bottlenecks: a connectionist model of the confusion effect. *Behavioral Ecology and Sociobiology*, 36:421–429, 1995.

[235] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.

[236] D. Sumpter, C. Buhl, D. Biro, and I. Couzin. Information transfer in moving animal groups. *Theory in biosciences*, 127(2):177–186, 2008.

[237] W.-C. Yang and T. Schmickl. Collective motion as an ultimate effect in crowded selfish herds. *Scientific reports*, 9(1):6618, 2019.

[238] R. J. Schmitt and S. W. Strand. Cooperative foraging by yellowtail, seriola lalandei (carangidae), on two species of fish prey. *Copeia*, 1982(3):714–717, 1982.

[239] E. S. Hobson. Interactions between piscivorous fishes and their prey. *Predator-prey systems in fisheries management*, pages 231–242, 1979.

[240] C. Packer and L. Ruttan. The evolution of cooperative hunting. *The American Naturalist*, 132(2):159–198, 1988.

[241] N. O. Handegard, K. M. Boswell, C. C. Ioannou, S. P. Leblanc, D. B. Tjøstheim, and I. D. Couzin. The dynamics of coordinated group hunting and collective information transfer among schooling prey. *Current biology*, 22(13):1213–1217, 2012.

[242] S. Pal, N. Pal, S. Samanta, and J. Chattopadhyay. Effect of hunting cooperation and fear in a predator-prey model. *Ecological Complexity*, 39:100770, 2019.

[243] C. Muro, R. Escobedo, L. Spector, and R. Coppinger. Wolf-pack (canis lupus) hunting strategies emerge from simple rules in computational simulations. *Behavioural processes*, 88(3):192–197, 2011.

[244] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot sheepdog project achieves automatic flock control. In *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour*, volume 489, page 493, 1998.

[245] J.-M. Lien, S. Rodriguez, J.-P. Malric, and N. M. Amato. Shepherding behaviors with multiple shepherds. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3402–3407. IEEE, 2005.

[246] A. Pierson and M. Schwager. Bio-inspired non-cooperative multi-robot herding. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1843–1849. IEEE, 2015.

[247] T. Miki and T. Nakamura. An effective simple shepherding algorithm suitable for implementation to a multi-mmobile robot system. In *First International Conference on Innovative Computing, Information and Control-Volume I (ICICIC'06)*, volume 3, pages 161–165. IEEE, 2006.

[248] M. Bacon and N. Olgac. Swarm herding using a region holding sliding mode controller. *Journal of Vibration and Control*, 18(7):1056–1066, 2012.

[249] A. Varava, K. Hang, D. Kragic, and F. T. Pokorny. Herding by caging: a topological approach towards guiding moving agents via mobile robots. In *Robotics: Science and Systems*, pages 696–700, 2017.

[250] H. T. Nguyen, T. D. Nguyen, M. Garratt, K. Kasmarik, S. Anavatti, M. Barlow, and H. A. Abbass. A deep hierarchical reinforcement learner for aerial shepherding of ground swarms. In *International Conference on Neural Information Processing*, pages 658–669. Springer, 2019.

[251] N. R. Clayton and H. Abbass. Machine teaching in hierarchical genetic reinforcement learning: Curriculum design of reward functions for swarm shepherding. In *2019 IEEE congress on evolutionary computation (CEC)*, pages 1259–1266. IEEE, 2019.

[252] G. Y. Dosieah, A. Özdemir, M. Gauci, and R. Groß. Moving mixtures of active and passive elements with robots that do not compute. In *International Conference on Swarm Intelligence*, pages 183–195. Springer, 2022.

[253] H. Singh, B. Campbell, S. Elsayed, A. Perry, R. Hunjet, and H. Abbass. Modulation of force vectors for effective shepherding of a swarm: A bi-objective approach. In *2019 IEEE congress on evolutionary computation (CEC)*, pages 2941–2948. IEEE, 2019.

[254] K. Fujioka. Effective herding in shepherding problem in v-formation control. *Transactions of the Institute of Systems, Control and Information Engineers*, 31(1):21–27, 2018.

[255] J. Hu, A. E. Turgut, T. Krajnik, B. Lennox, and F. Arvin. Occlusion-based coordination protocol design for autonomous robotic shepherding tasks. *IEEE Transactions on Cognitive and Developmental Systems*, 14(1):126–135, 2020.

[256] D. A. Feary, A. M. Fowler, and D. J. Booth. Predator-avoidance behaviour of target and non-target temperate reef fishes is lower in areas protected from fishing. *Marine Biology*, 171(3):66, 2024.

[257] H. Hamann, Y. Khaluf, J. Botev, M. Divband Soorati, E. Ferrante, O. Kosak, J.-M. Montanier, S. Mostaghim, R. Redpath, J. Timmis, et al. Hybrid societies: challenges and perspectives in the design of collective behavior in self-organizing systems. *Frontiers in Robotics and AI*, 3:14, 2016.

[258] T. Schmickl, H. Hamann, H. Wörn, and K. Crailsheim. Two different approaches to a macroscopic model of a bio-inspired robotic swarm. *Robotics and Autonomous Systems*, 57(9):913–921, 2009.

[259] S. Konur, C. Dixon, and M. Fisher. Analysing robot swarm behaviour via probabilistic model checking. *Robotics and Autonomous Systems*, 60(2):199–213, 2012.

[260] A. F. Winfield, W. Liu, J. Nembrini, and A. Martinoli. Modelling a wireless connected swarm of mobile robots. *Swarm intelligence*, 2:241–266, 2008.

[261] D. Strömbom, R. P. Mann, A. M. Wilson, S. Hailes, A. J. Morton, D. J. Sumpter, and A. J. King. Solving the shepherding problem: heuristics for herding autonomous, interacting agents. *Journal of the royal society interface*, 11(100):20140719, 2014.

[262] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Physical review letters*, 75(6):1226, 1995.

[263] C. W. Warren. Global path planning using artificial potential fields. In *1989 IEEE International Conference on Robotics and Automation*, pages 316–317. IEEE Computer Society, 1989.

[264] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *The international journal of robotics research*, 5(1):90–98, 1986.

[265] R. J. Wootton. *Fish ecology*. Springer Science & Business Media, 1991.

[266] B. Le Gallic and A. Cox. An economic analysis of illegal, unreported and unregulated (iuu) fishing: Key drivers and possible solutions. *Marine Policy*, 30(6):689–695, 2006.

[267] L. Jacquin, Q. Petitjean, J. Côte, P. Laffaille, and S. Jean. Effects of pollution on fish behavior, personality, and cognition: some research perspectives. *Frontiers in Ecology and Evolution*, 8:86, 2020.

[268] F. Essl, B. Lenzner, S. Bacher, S. Bailey, C. Capinha, C. Daehler, S. Dullinger, P. Genovesi, C. Hui, P. E. Hulme, et al. Drivers of future alien species impacts: An expert-based assessment. *Global Change Biology*, 26(9):4880–4893, 2020.

[269] R. Carson. *The sea around us*. Oxford University Press, USA, 2003.

[270] M. H. Keenleyside. *Diversity and adaptation in fish behaviour*, volume 11. Springer Science & Business Media, 2012.

[271] C. H. Ryer and B. L. Olla. Influences of food distribution on fish foraging behaviour. *Animal Behaviour*, 49(2):411–418, 1995.

[272] N. Miller and R. Gerlai. From schooling to shoaling: patterns of collective motion in zebrafish (danio rerio). *PloS one*, 7(11):e48865, 2012.

[273] R. R. Krausz. Living in groups. *Transactional Analysis Journal*, 43(4):366–374, 2013.

[274] R. L. Day, T. MacDonald, C. Brown, K. N. Laland, and S. M. Reader. Interactions between shoal size and conformity in guppy social foraging. *Animal Behaviour*, 62(5): 917–925, 2001.

[275] D. Schneider. Identifying the spatial scale of density-dependent interaction of predators with schooling fish in the southern labrador current. *Journal of Fish Biology*, 35:109–115, 1989.

[276] C. K. Hemelrijk and H. Kunz. Density distribution and size sorting in fish schools: an individual-based model. *Behavioral Ecology*, 16(1):178–187, 2005.

[277] S. Furuichi, Y. Kamimura, M. Suzuki, and R. Yukami. Density-dependent attributes of schooling in small pelagic fishes. *The Science of Nature*, 109(5):49, 2022.

[278] T. J. Pitcher, G. Greenberg, and M. Haraway. Shoaling and schooling behaviour in fishes. *Comparative psychology: a handbook*, pages 748–760, 1998.

[279] W. L. Romey. Individual differences make a difference in the trajectories of simulated schools of fish. *Ecological Modelling*, 92(1):65–77, 1996.

[280] D. Strömbom. Collective motion from local attraction. *Journal of theoretical biology*, 283 (1):145–151, 2011.
[281] D. V. Radakov. *Schooling in the Ecology of Fish*. John Wiley & Sons, New York, 1973.
[282] K. J. Parton, T. S. Galloway, and B. J. Godley. Global review of shark and ray entanglement in anthropogenic marine debris. *Endangered Species Research*, 39:173–190, 2019.
[283] J. F. Stephenson, S. E. Perkins, and J. Cable. Transmission risk predicts avoidance of infected conspecifics in trinidadian guppies. *Journal of Animal Ecology*, 87(6):1525–1533, 2018.
[284] P. Guevara-Fiore, J. Stapley, and P. J. Watt. Mating effort and female receptivity: how do male guppies decide when to invest in sex? *Behavioral Ecology and Sociobiology*, 64: 1665–1672, 2010.
[285] D. Bierbach, L. Gómez-Nava, F. A. Francisco, J. Lukas, L. Musiolek, V. V. Hafner, T. Landgraf, P. Romanczuk, and J. Krause. Live fish learn to anticipate the movement of a fish-like robot. *Bioinspiration & Biomimetics*, 17(6):065007, 2022.
[286] J. W. Jolles, N. Weimar, T. Landgraf, P. Romanczuk, J. Krause, and D. Bierbach. Group-level patterns emerge from individual speed as revealed by an extremely social robotic fish. *Biology letters*, 16(9):20200436, 2020.
[287] D. Bierbach, H. J. Mönck, J. Lukas, M. Habedank, P. Romanczuk, T. Landgraf, and J. Krause. Guppies prefer to follow large (robot) leaders irrespective of own size. *Frontiers in Bioengineering and Biotechnology*, 8:441, 2020.
[288] J. A. Fox, M. W. Toure, A. Heckley, R. Fan, S. M. Reader, and R. D. Barrett. Insights into adaptive behavioural plasticity from the guppy model system. *Proceedings of the Royal Society B*, 291(2018):20232625, 2024.
[289] P. Domenici, J. M. Blagburn, and J. P. Bacon. Animal escapology i: theoretical issues and emerging trends in escape trajectories. *Journal of Experimental Biology*, 214(15):2463–2473, 2011.
[290] P. Domenici. Context-dependent variability in the components of fish escape response: integrating locomotor performance and behavior. *Journal of Experimental Zoology Part A: Ecological Genetics and Physiology*, 313(2):59–79, 2010.
[291] S. Hubbard, P. Babak, S. T. Sigurdsson, and K. G. Magnússon. A model of the formation of fish schools and migrations of fish. *Ecological Modelling*, 174(4):359–374, 2004.
[292] M. Kruusmaa, G. Rieucau, J. C. C. Montoya, R. Markna, and N. O. Handegard. Collective responses of a large mackerel school depend on the size and speed of a robotic fish but not on tail motion. *Bioinspiration & biomimetics*, 11(5):056020, 2016.
[293] A. Pino, R. Vidal, E. Tormos, J. M. Cerdà-Reverter, R. Marín Prades, and P. J. Sanz. Towards fish welfare in the presence of robots: Zebrafish case. *Journal of Marine Science and Engineering*, 12(6):932, 2024.
[294] D. S. Calovi, A. Litchinko, V. Lecheval, U. Lopez, A. Pérez Escudero, H. Chaté, C. Sire, and G. Theraulaz. Disentangling and modeling interactions in fish with burst-and-coast swimming reveal distinct alignment and attraction behaviors. *PLoS computational biology*, 14(1):e1005933, 2018.
[295] A. D. Hartono, L. T. H. Nguyen, and T. V. Ta. A stochastic differential equation model for predator-avoidance fish schooling. *Mathematical Biosciences*, 367:109112, 2024.
[296] A. E. Magurran. *Evolutionary ecology: the Trinidadian guppy*. Oxford University Press, 2005.
[297] J. D. Parrish. Fish communities of interacting shallow-water habitats in tropical oceanic regions. *Marine ecology progress series. Oldendorf*, 58(1):143–160, 1989.

The work presented in this chapter has been published in:
ANTS, International conference on swarm intelligence (Van Havermaet et al. 2022)

2

An adaptive metric model for collective motion structures in dynamic environments

Abstract
For robot swarms to contribute to fish conservation, they must be able to explore and navigate the dynamic environments where fish reside, such as coral reefs and cave structures. This ability enables biodiversity documentation, habitat mapping for future interventions, and the identification of potential threats. Fish schools, as natural decentralized systems, have evolved collective motion strategies to efficiently move in such environments. Although various models of collective motion exist, their ability to generate adaptive collective structures in response to environmental stimuli remains limited. In this chapter, we identify three key collective structures necessary for effective exploration, characterized by the swarm's spatial coverage and degree of alignment. We find that conventional models, such as metric and topological approaches, are insufficient for generating all required formations. To address this limitation, we propose an extension to the metric model that enables swarms to transition between all three structures as needed. The proposed model is validated in a simulated dynamic environment, demonstrating the robot swarm's ability to adapt its structure in response to locally observed stimuli.

2.1 Introduction

Collective motion is one of nature's most fundamental demonstration of coordinated activity, performed by a wide spectrum of systems, such as bird flocks, fish shoals, and human crowds [2]. It is a key example of emergent behavior that results from the individuals' intensive interactions to perform particular tasks such as foraging or migrating. Collective artificial systems, such as robot swarms, have inherited the importance of collective motion. Robot swarms is a a field of study that focuses on simulating observable behavior patterns in social animals. The individual robots have basic abilities, but when they work together in groups, they can undertake more sophisticated behaviors including foraging [3–6], exploration [7, 8], and collective perception [9, 10].

Many activities in robot swarms need collective motion. Navigation from a source to a destination, forming topologies and patterns, tracking targets, and moving objects are among the examples. Depending on the particular task, specific structures need to be displayed in the swarm. For example, the swarm needs to aggregate while navigating through narrow paths, and to expand while exploring new environments. A large number of theoretical [11, 12] and empirical [13, 14] studies have proposed models to generate collective motion. Most of these models consider short-range repulsion and long-range attraction among the individuals, in addition to the alignment of velocities along with the their (nearest) neighbors [2]. These models show high efficiency in generating aligned motion based on simple individual rules. However, there is little to no evidence on whether such models can modify spatial features of the group (e.g., structure) as a response to environmental stimuli.

In this study, we use two system measures to define our target structures: swarm order (an expression of alignment degree) and swarm relative coverage (an expression of compactness). We define three target structures which the swarm displays based on the environmental stimuli: (a) high coverage, low order (HCLO), (ii) high coverage, high order (HCHO), and (iii) low coverage, high order (LCHO). In HCLO, the swarm is supposed to maximize its coverage, and individuals have different orientation (low order). Such structures are desired for exploration tasks, where robots need to observe the environment in all possible directions and maximize their inter-individual distance. In HCHO, the swarm is displaying a high order, while maintaining a high coverage. This structure is suitable for navigating tasks while preserving a maximized coverage (e.g., navigation with exploration). Finally, in LCHO, the swarm aggregates, displaying a high density (i.e., low coverage) and maintaining a high order. This structure is suitable for e.g., transporting tasks or navigating through narrow spaces. In all three structures, the swarm needs to maintain connectivity (i.e., remain in a single cluster).

We consider two models of collective motion: the metric model, where

each individual interacts within a defined radius, and the long-range model, where short-interactions occur topologically and long-range interactions randomly. By analyzing the influence of the different model parameters on the emergence of the above-mentioned structures, our results highlight the limitations of the metric model and the long-range model to generate these structures. We propose an extension to the metric model that enables the swarm to display the three structures and to switch between them as a response to environmental stimuli. We test our extended metric model (EMM) in a simulated dynamic environment, where different stimuli appear over time and are perceived by few individuals. Upon receiving information about a stimulus, individuals adapt model parameters in order to generate the corresponding structure. We show how the three target structures—i.e., HCLO, HCHO, and LCHO—have properly emerged, and how EMM scales with the system size. The remainder of this paper is organized as follows. In Section 2.2, we provide a brief overview of the literature on collective motion models. In Section 2.3, we describe the different models we are using in addition to the EMM. In Section 2.4, we present and discuss the key results. Finally, we conclude this paperin Section 2.5.

2.2 Related Work

A large number of models have been suggested to generate collective motion and display group cohesion. Models inspired from both biology [15–17] and physics [18, 19] suggested simple rules of interaction among individuals to be sufficient to induce collective motion. Such rules capture the three forces of attraction, repulsion and alignment [15]. The vast majority of these models adopt a definition of physical distance, i.e., metric distance. These models make the assumption that individuals align and attract each other, and that such interaction decays with increasing distance between individuals. Examples of widely-used models from the metric class include the Vicsek model [2], which exploits the individuals' alignment by considering the neighbor's velocity as the model input. The Couzin model [20] is widely used in theoretical biology, and has been extended to achieve collective motion in group robotics. We use the Couzin model as our metric model in this paper.

Special cases of the metric model were suggested and became widely used such as the topological model [21, 22], in which each individual interacts with a fixed number of the neighbors within their interaction range—e.g., the m-nearest neighbors. The topological model was suggested based on experimental findings which revealed that birds in a flock tend to interact with a fixed number of nearest neighbors [22]. Another special-case of the metric model is the visual model [23]. This model was proposed to account of the interaction limitations imposed by sensory constraints. In this model, individuals interact only with the neighbors that are visually ob-

servable (based on ray-casting). Recently, in [24] long-range models were suggested as a solution to avoid group dispersion when moving into open spaces. Long-range model is an extension of the Vicsek model which introduces long-range alignment interactions between the individuals. The authors also show that a small number of these interactions is enough for the group to generate long-lasting ordered states while moving in open spaces.

Finally, several studies have addressed spatial structures in natural organisms. In [25], the authors studied the two structures of shoals and schools in Zebrafish. shoals refer to any group of fish that is highly aggregated, whilst, schools are shoals that display a high level of order (alignment) when moving together. Nevertheless, the focus of this work was to investigate model parameters else than group density that may impact the emergence of these two structures. The authors in both [20, 26] studied how the individual heterogeneity (e.g., behavioral state, age) may affect the individual's spatial position in the group leading to specific structures to emerge. In [27], the authors highlight the importance of the system size and the number of influential neighbors, as well as the weight assigned to the three zones of attraction, alignment and repulsion on the emergence of different structures. Despite interesting results, the work doesn't suggest how the model can be used to generate the different spatial structures. Moreover, the obtained structures were defined in terms of dispersion rather than structures of a single cluster.

2.3 Model

Let $\mathcal{R}$ denote the set of robotic agents of size N. The state of system at time t is defined by the position $p_i(t) \in \mathbb{R}^2$ and orientation $\theta_i(t) \in [0, 2\pi)$ of each individual $i \in \mathcal{R}$. At time t, an individual computes its desired orientation $\hat{\theta}(t) \in [-\pi, \pi)$ based on information gathered from neighbors. The desired velocity $\hat{v}(t)$ of each individual is always set to the maximum linear velocity $v_{\max}$.

2.3.1 Motion Kinematics

Let w denote the angular velocity. The discrete-time motion kinematics are then defined by the difference in orientation $\phi(t) = \theta(t) - \hat{\theta}(t)$. The required time to turn $\tau(t)$ is then defined as $\tau(t) = \min(\frac{|\phi(t)|}{w}, \Delta t)$. Consequently, the turn $\varphi(t)$ is given by $\varphi(t) = \text{sgn}(\phi(t))\tau(t)w$, with time interval Δt and sgn as the sign function. The orientation θ and velocity v of the next time step $t' = t + \Delta t$ are respectively computed as $\theta(t') = \theta(t) - \varphi(t) + \mathcal{N}(0, \sigma)$ and $v(t') = \min\Big(\hat{v}(t), v_{\max}\big(\Delta t - \tau(t)\big)\Big)$. with angular Gaussian noise σ. The position at time t' is then given by

$$p(t') = p(t) + v(t') \begin{bmatrix} \cos\theta(t') \\ \sin\theta(t') \end{bmatrix}$$

2.3.2 Metric and Long-Range Model

Each individual $i \in \mathcal{R}$ updates its direction of motion based on three non-overlapping behavioral zones around i, in which its neighbors are located. Each behavioral zone corresponds to a distinct interaction; (i) repulsion from others, to establish a minimum inter-individual distance, inside the circular zone with radius ZoR, (ii) alignment of orientation with others, to all move in the same direction, inside the zone with width ZoO, and (iii) attraction to others, to remain as one cohesive group, inside the zone with width ZoA.

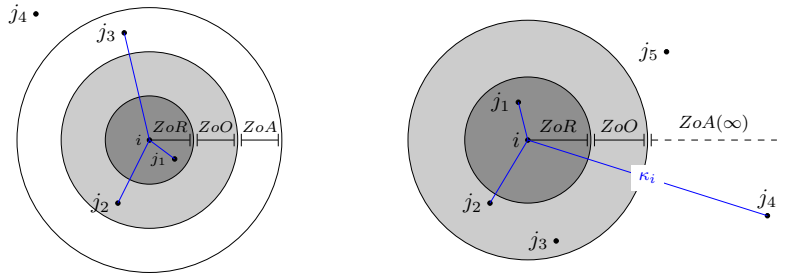

Figure 2.1: Schematic representation of the (**left**) metric and (**right**) long-range models for an arbitrary individual i. Other individuals of the swarm are denoted as $j \in \mathcal{R} \setminus \{i\}$, and those selected as neighbors are connected by a blue line from i. The metric model defines the set of neighbors $\mathcal{N}_i$ as every other individual positioned in one of the three behavioral zones surrounding i. In this example, $\mathcal{N}_i = \{j_1, j_2, j_3\}$. In the long-range model, short-range interactions are considered by selecting m nearest neighbors, while long-range interactions are defined by randomly selecting κ_i individuals from the remaining individuals of the swarm. The number of long-range neighbors κ_i is sampled from a Poisson distribution with average κ as parameter of the model. In this example, $\mathcal{N}_i = \{j_1, j_2, j_4\}$, with $m = 2$ and $\kappa_i = 1$.

The metric and long-range models differ from each other by neighbor selection. For the metric model, the neighbors $\mathcal{N}_i$ of individual i consists of all individuals within the interaction-radius ($ZoR + ZoO + ZoA$) around p_i (Fig. 2.1 left) [20], whereas for the long-range model $\mathcal{N}_i$ is the union of the set of m nearest neighbors $\mathcal{M}_i$ and κ_i randomly selected neighbors of $\mathcal{N}_i \setminus \mathcal{M}_i$ (Fig. 2.1 right) [24].

We define r_{ij} as the relative position of individual j from i, i.e. $r_{ij}(t) = p_j(t) - p_i(t)$. Let $\mathcal{N}_i^r$, $\mathcal{N}_i^o$, and $\mathcal{N}_i^a$ then denote the distinct subsets of neighbors by separating $\mathcal{N}_i$ based on the repulsion, orientation, and attraction zones respectively. More specifically, $\mathcal{N}_i^r = \{j \mid \|r_{ij}\| \leq ZoR \wedge j \in \mathcal{N}_i\}$, $\mathcal{N}_i^o = \{j \mid ZoR < \|r_{ij}\| \leq ZoR + ZoO \wedge j \in \mathcal{N}_i\}$, and

$\mathcal{N}_i^a = \{j \mid ZoR + ZoO < \|r_{ij}\| \leq ZoR + ZoO + ZoA \wedge j \in \mathcal{N}_i\}$. It follows for both models that $\mathcal{N}_i = \mathcal{N}_i^r \cup \mathcal{N}_i^o \cup \mathcal{N}_i^a$, as we set $ZoA = \infty$ for the long-range model. The desired direction vector $\hat{q}_i(t)$ at time t of individual i is then computed as follows:

$$\hat{q}_i(t) = -\alpha_r \sum_{j \in \mathcal{N}_i^r} \frac{r_{ij}}{\|r_{ij}\|} + \alpha_o \sum_{j \in \mathcal{N}_i^o} \frac{q_j(t)}{\|q_j(t)\|} + \alpha_a \sum_{j \in \mathcal{N}_i^a} \frac{r_{ij}}{\|r_{ij}\|} \tag{2.1}$$

and consequently the desired direction $\hat{\theta}_i(t) = \arg \hat{q}_i(t)$, with weights $\alpha_r \geq 0$, $\alpha_o \geq 0$ and $\alpha_a \geq 0$ of repulsion, orientation, and attraction respectively.

2.3.3 System Measures

To quantify whether all individuals move in approximately the same direction, we measure the amount of order Ψ defined as $\Psi = \lim_{t \to T} \frac{1}{N} \| \sum_{i=1}^{N} q_i(t) \|$, where T is the time at which the order has converged. Consequently, $\Psi = 1$ indicates an ordered system where all individuals follow the same direction, while $\Psi = 0$ corresponds to a disordered system where individuals have no alignment in orientation.

Furthermore, we define the relative coverage Ω as the ratio of the area of the convex hull at time of convergence $t = T$ and the initial area at $t = t_0$. More specifically, let $A(t)$ denote the area of the convex hull of $\mathcal{R}$ at time t, then the relative coverage is defined as $\Omega = \lim_{t \to T} \frac{A(t)}{A(t_0)}$. Thus, $\Omega > 1$ indicates that the individuals have spread out since the initial time t_0, whereas $\Omega < 1$ indicates that the inter-individual distance between individuals has decreased.

In order to observe fragmentation of the group, the number of clusters is measured where individuals i and j are part of the same cluster if their relative distance is lower than the interaction-radius RoI. In the experiments with the long-range model, we use a value of $ZoA = \frac{1}{2}(\frac{N}{\rho})^{\frac{1}{2}}$ to compute the interaction-radius, with $\rho = 0.01$ as initial density.

2.3.4 Extended Metric Model (EMM)

Our proposed model, the extended metric model (EMM), relies on adapting the impact of the behavioral orientation zone while following the same neighbor selection approach as the metric model. In order to obtain a relative coverage $\Omega > 1$ with low order $\Psi \approx 0$, we maintain the width of the orientation zone ZoO, but set $\alpha_o = 0$ in Equation (2.1)—i.e., deactivating the orientation zone. Consequently, individuals are able to spread out until attraction interactions ensure that they remain cohesive. Transitioning to high relative coverage $\Omega > 1$ and high order $\Psi \approx 1$ is then accomplished by resetting $\alpha_o > 0$ (i.e., activating the orientation zone). To obtain a low relative coverage $\Omega < 1$ with high order $\Psi \approx 1$, the width ZoO is decreased while keeping the zone activated ($\alpha_o > 0$).

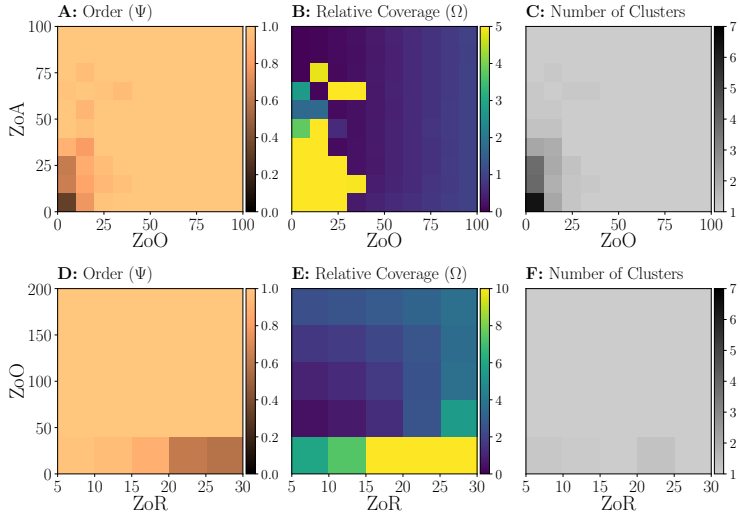

Figure 2.2: System measures of the metric model; order (A,D), relative coverage (B,E), and number of clusters (C,F), with varying combinations of ZoO-ZoA (top row) and ZoR-ZoO (bottom row).

2.4 Results and Discussion

We simulate a robot swarm in 2D open space environment. At the beginning of a simulation, robots are placed within a confined box of the size $(\frac{N}{\rho})^{\frac{1}{2}}$ with initial density $\rho = 0.01$. This is to facilitate immediate interactions between robots. Within this box, both the robots' positions and moving directions are initially uniformly distributed. Figures (2.2,2.3,2.4,2.5,2.7b) are results from simulation runs with 10 different seeds, while Figure 2.7 was obtained with 30 seeds. All experiments are run with $w = \frac{\pi}{2}$, $v_{\max} = 2$, $\sigma = 0.05$, $\Delta t = 1$, $\alpha_r = 100$, $\alpha_o = 50$, and $\alpha_a = 1$ based on preliminary experiments to obtain a system of a single cluster. Unless varied, the system size is $N = 100$.

We start with the metric model, looking at the swarm order Ψ, group relative coverage Ω, and the number of clusters. The emergence of the target structures (i.e., HCLO, HCHO, and LCHO) is investigated using a combination of these system measures. We enable the width of the orientation zone (ZoO) and the attraction zone (ZoA) to vary over the range of $[0-100]$ in Figures 2.2A,B,C, while keeping the width of the repulsion zone constant ($ZoR = 1$). Structures that arise while the swarm is preserved in a single connected cluster (light-gray color in Figure 2.2), have a low relative coverage and a high group order, which corresponds to the target structure (LCHO). We note that the swarm splitting in numerous clusters fits with the structure of high coverage and low order (HCLO) (left-bottom corner).

Finally, the high-coverage, high-order (HCHO) structure is completely absent.

The width of the repulsion zone (ZoR) and the orientation zone (ZoO) are then varied throughout a range of $[5-30]$ and $[0-200]$, respectively, while the width of the attraction zone remains constant ($ZoA = 50$). The LCHO structure is formed when high order corresponds with low relative coverage, as seen in Figures 2.2D,E,F. However, given a medium level of order with a possibility of more than one cluster, the right-bottom corner shows a likelihood of a high coverage, low order (HCLO) to emerge. Finally, the high-coverage, high-order (HCHO) structure is again absent.

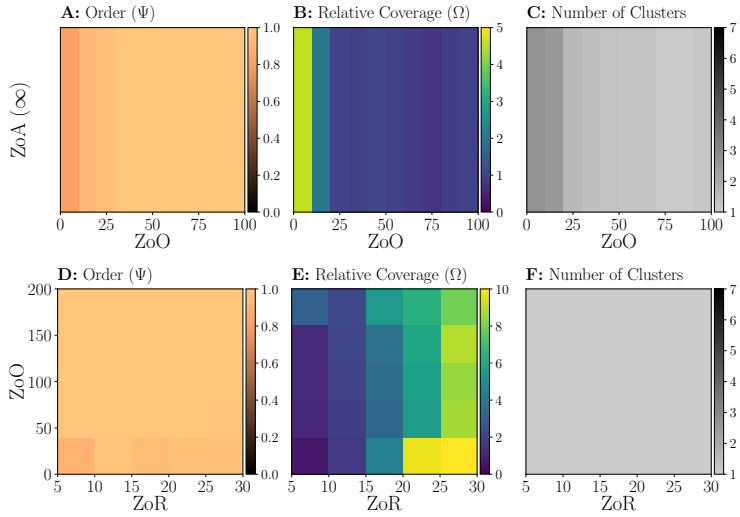

Figure 2.3: System measures of the long-range model with $\kappa = 0.05$; order (A,D), relative coverage (B,E), and number of clusters (C,F), with varying ZoO (top row) and combinations of ZoR-ZoO (bottom row).

Next we perform the same analysis of the system measures (order Ψ, relative coverage Ω, and number of clusters) using the long-range model. The results are shown in Figure 2.3 with the average long-range connectivity set to $\kappa = 0.05$. In Figures 2.3A,B,C we vary the width of the orientation zone (ZoO) over the range of $[0-100]$. Our results show that the swarm can move in a single cluster, while maintaining a low relative coverage, and a high swarm order. This aligns with the structure of low coverage, high order (LCHO). The other two structures of HCLO and HCHO are fully absent. In Figures 2.3D,E,F show the system measure while varying the width of the repulsion zone (ZoR) over the range of $[5-30]$. There we can notice an evidence of high relative coverage with high order in the right-bottom corner (while maintaining a single cluster). This corresponds to the structure of HCHO. The long-range model shows similarly to the

metric model the ability to generate low coverage, high order structures (LCHO). The structure of high coverage, low order (HCLO) is missing.

As the average long-range connectivity κ increases, the long-range model's ability to create a high coverage, high order (HCHO) decreases, as demonstrated in Figure 2.4. The low coverage, high order (LCHO) structure becomes the only one that the long-range model can generate. Hence, the emergence of the high coverage, high order (HCHO) is κ-dependent for the long-range model.

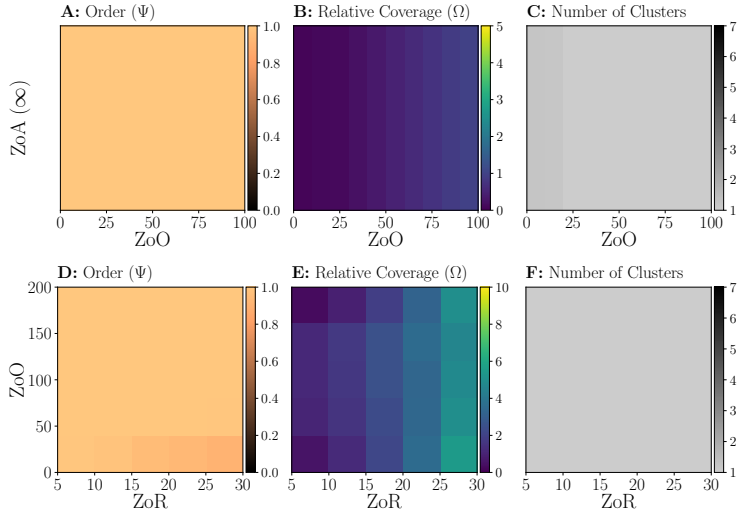

Figure 2.4: System measures of the long-range model with $\kappa = 0.9$; order (A,D), relative coverage (B,E), and number of clusters (C,F), with varying ZoO (top row) and combinations of ZoR-ZoO (bottom row).

So far, we have demonstrated that neither the metric nor the long-range model is suitable for generating a HCLO structure, under any considered parameter settings. However, the HCLO structure is a significant addition to the previously observed LCHO and HCHO structures, particularly in the context of underwater exploration and environmental sensing. In natural systems, such as fish shoals, HCLO-like configurations are observed during foraging behavior, where individuals explore independently while maintaining loose cohesion. From a robotics perspective, HCLO offers a compelling strategy for tasks that require maximal spatial exploration with diverse directional sampling. This is especially relevant when onboard sensors are not omnidirectional, such as acoustic sensors, sonar arrays, or directional cameras, which require physical reorientation to scan different parts of the environment. Even when robots possess omnidirectional sensing, low alignment increases the likelihood that agents will disperse toward different areas, reducing redundancy in gathered data, and

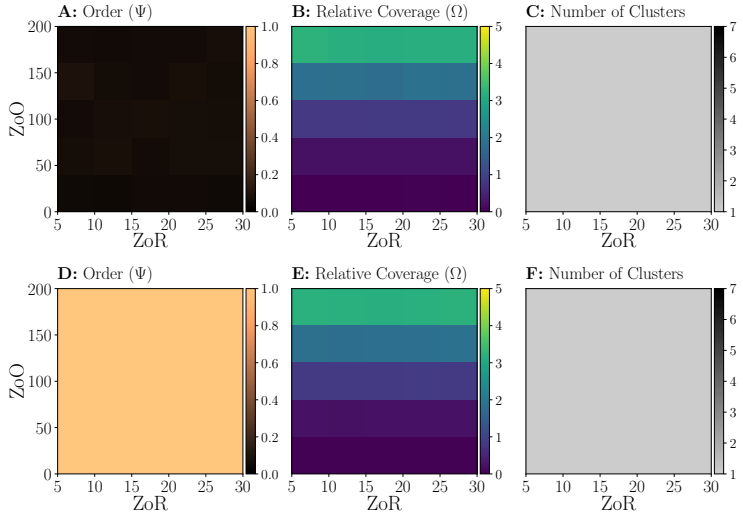

Figure 2.5: System measures of the proposed model (EMM); order (A,D), relative coverage (B,E), and number of clusters (C,F), with deactivation (top row) and activation (bottom row) of the orientation zone while varying ZoR-ZoO.

increasing overall spatial coverage over time.

In the following we show the system measures resulting from applying the extended metric model (EMM). Figure 2.5 shows the swarm order Ψ, relative coverage Ω, and the number of clusters. The width of the repulsion zone (ZoR) and the orientation zone (ZoO) are varied throughout a range of $[5-30]$ and $[0-200]$, respectively, while the width of the attraction zone remains constant ($ZoA = 50$). These results demonstrate how the EMM can display the three target structures, while maintaining the group moving in a single cluster for all structures (Figure 2.5 right column). Figures 2.5A,B show the system measures when deactivating the orientation zone. Hence the values at the y-axis define the distance at which the attraction zone starts. These two figures show the possibility to generate the HCLO structure through expanding the swarm coverage while pushing the attraction zone away by increasing the width of the deactivated orientation zone. Figures 2.5D,E are obtained after activating the orientation zone. They show the ability of the EMM model to display, both, the HCHO and the LCHO structures. The HCHO (LCHO) structure is achieved by increasing (decreasing) the width of the activated orientation zone. Both results in Figures 2.5A,B and in Figures 2.5D,E show that the emerging structure (i.e., the change in the relative converge Ω, and the order Ψ is independent of the width of the repulsion zone (ZoR).

Figure 2.5 suggests that while the EMM model can generate a continuum of structures with varying relative coverage, the degree of alignment

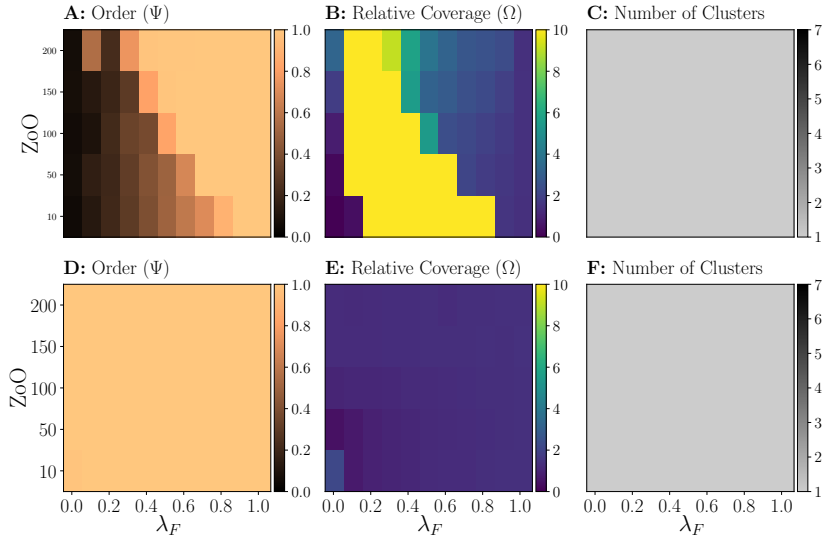

Figure 2.6: Hybrid swarm with λ_F denoting the fraction of individuals following the metric model, while the rest follow the proposed EMM model. A total of $N = 100$ individuals were simulated. Panels show system measures: order (A,D), relative coverage (B,E), and number of clusters (C,F), under deactivation (top row) and activation (bottom row) of the orientation zone, as λ_F and ZoO vary.

remains polarized. The swarm typically exhibits either low order (top row; $\Psi = 0$) or high order (bottom row; $\Psi = 1$), with no intermediate degrees of alignment. We hypothesize that this result comes from the binary parameterization of the alignment force coefficient α_O, which is set to 0 (deactivated) or 50 (fully active) in Figure 2.5. Intermediate values of Ψ could potentially emerge from tuning α_O across a continuous range. From a practical perspective, such medium-alignment structures may offer a useful balance between directional persistence and exploratory coverage. For instance, a group moving with moderate alignment would avoid redundant revisits to the same locations (as in low-order dispersal) while still diversifying their spatial footprint more than in highly aligned motion.

To better understand the transition between the traditional metric model and the proposed EMM, we additionally explore a hybrid swarm configuration in which a fraction λ_F of individuals follows the metric model, while the remaining agents operate under the EMM (see Figure 2.6). This intermediate setup helps reveal how even partial adoption of the EMM dynamics can impact collective behavior. The results show that as λ_F decreases (i.e., as more agents follow the EMM), the system progressively shifts towards displaying the full range of target structures. In particular, under activation of the orientation zone (bottom row), the hy-

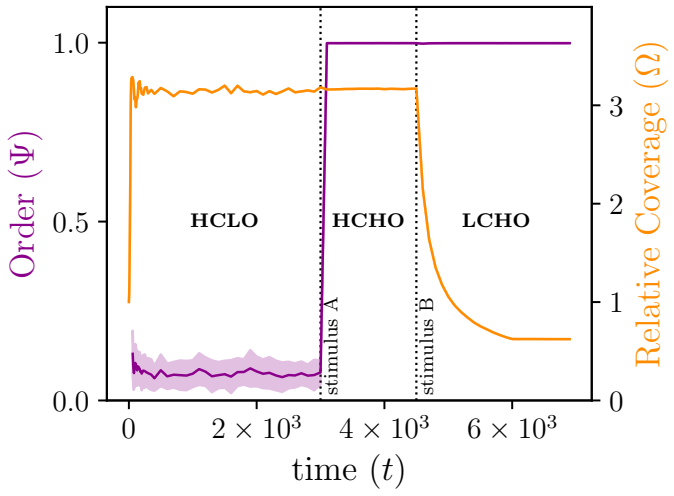

Figure 2.7: For the proposed model (EMM); the order (purple) and the relative coverage (orange) over time t, as information is introduced and propagated in the swarm.

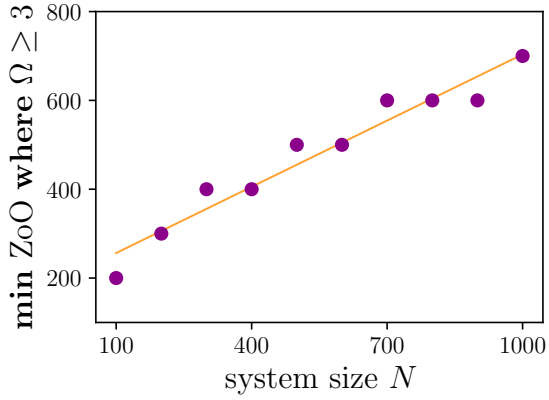

Figure 2.8: For the proposed model (EMM); the minimum required width of orientation zone ZoO to obtain a relative coverage $\Omega > 3$ for different system sizes N.

brid swarm is capable of producing both high-order and high-coverage (HCHO) configurations, which are absent in the pure metric model. These findings suggest that the EMM model exerts a strong organizing influence on the collective, even when only partially adopted by the swarm.

Next, we simulate a swarm of robots using the EMM model to perform the following sequential set of tasks: (i) explore the environment looking

for a particular stimulus (**stimulus A**) that define the direction they need to move into. (ii) Navigate in the direction of stimulus A until a **stimulus B** appears. (iii) As a response to stimulus B (e.g., a narrow path) the swarm needs to shrink in coverage while still navigating to its target. Figure 2.7 shows the system measures recorded over 7×10^3 simulated time steps. The swarm order Ψ starts low as the robots are initialized with random directions. Following the EMM model, every robot deactivates its orientation zone, whose width is set to 200, aiming for the HCLO structure. In a few time steps, the relative coverage increases to $\Omega = 3$, while the system maintains a low order. At time step 3×10^3, stimulus A (which triggers the swarm to navigate into one direction) is introduced and perceived by a single robot, who spreads the message to its neighbors. As the message spreads, robots start activating their orientation zone (see Figure 2.5D,E). This enables the swarm to converge to the HCHO structure after introducing stimulus A, as shown in Figure 2.7. Thanks to activating the orientation zone, fluctuations in both system measures disappear. Finally, at time step 4.5×10^3 stimulus B is introduced and perceived by a single robot, who spread it further. The system converges to the LCHO structure when informed robots reduce the width of their orientation zone. (In this paper, we reduce ZoO to 50 based on findings where we varied ZoO in $[5 - 200]$).

Finally, we show how the width of the orientation zone required to generate the HCLO, and HCHO scales with the system size. Note that the width of the other two zones is irrelevant to the behavior of the EMM (as it only modifies ZoO). Figure 2.8 shows that the minimum width of the orientation zone, that is needed to create the HCLO and the HCHO structures, scales linearly with the system size.

2.5 Conclusion

We studied in this paper the emergence of three target structures in robots swarms: HCLO, HCHO, and LCHO. We introduced these three structures as fundamental for a large set of robot tasks such as exploration, navigation, and moving through narrow paths. We define our structures using two system measures; the swarm order that expresses the degree of alignment, and the swarm relative coverage that expresses how the area occupied by the swarm is expanded or compressed in comparison to its value at initial time step.

We initially looked into the ability of the widely-used metric model to generate the three structures and we showed that across a wide range of parameter values only the structure of low coverage, high order (LCHO) has emerged. Similarly, we investigate the recently proposed long-range model, which shows the ability to generate only the low coverage, high order structure (LCHO) for all the configurations where κ is high. When κ is low and the width of the repulsion zone is large, the model displayed the HCHO structure.

We proposed an extension of the metric model (EMM) that modifies both the activation state of the orientation zone and the width of this zone to dynamically generate the three target structures. After showing the parameter range over which the EMM displays each of the three structures, we use a dynamic environment scenario to test the EMM by introducing stimuli over time that require the swarm to change its structure. EMM succeeds in enabling the swarm to form the necessary structure based on the displayed stimulus. We finally show how our model parameter scales linearly with the system size.

Bibliography

[1] S. Van Havermaet, P. Simoens, and Y. Khaluf. An adaptive metric model for collective motion structures in dynamic environments. In *International conference on swarm intelligence*, pages 257–265. Springer, 2022.

[2] T. Vicsek and A. Zafeiris. Collective motion. *Physics reports*, 517(3-4):71–140, 2012.

[3] I. Rausch, Y. Khaluf, and P. Simoens. Scale-free features in collective robot foraging. *Applied Sciences*, 9(13):2667, 2019.

[4] A. Font Llenas, M. S. Talamali, X. Xu, J. A. Marshall, and A. Reina. Quality-sensitive foraging by a robot swarm through virtual pheromone trails. In *International conference on swarm intelligence*, pages 135–149. Springer, 2018.

[5] J. Nauta, S. Van Havermaet, P. Simoens, and Y. Khaluf. Enhanced foraging in robot swarms using collective lévy walks. In *ECAI 2020*, pages 171–178. IOS Press, 2020.

[6] J. Nauta, P. Simoens, and Y. Khaluf. Group size and resource fractality drive multimodal search strategies: A quantitative analysis on group foraging. *Physica A: Statistical Mechanics and its Applications*, 590:126702, 2022.

[7] I. Rausch, P. Simoens, and Y. Khaluf. Adaptive foraging in dynamic environments using scale-free interaction networks. *Frontiers in Robotics and AI*, 7:86, 2020.

[8] M. Kegeleirs, D. Garzón Ramos, and M. Birattari. Random walk exploration for swarm mapping. In *Annual conference towards autonomous robotic systems*, pages 211–222. Springer, 2019.

[9] G. Valentini, D. Brambilla, H. Hamann, and M. Dorigo. Collective perception of environmental features in a robot swarm. In *International Conference on Swarm Intelligence*, pages 65–76. Springer, 2016.

[10] Y. Khaluf, M. Allwright, I. Rausch, P. Simoens, and M. Dorigo. Construction task allocation through the collective perception of a dynamic environment. In *International Conference on Swarm Intelligence*, pages 82–95. Springer, 2020.

[11] A. Cavagna, L. Del Castello, I. Giardina, T. Grigera, A. Jelic, S. Melillo, T. Mora, L. Parisi, E. Silvestri, M. Viale, et al. Flocking and turning: a new model for self-organized collective motion. *Journal of Statistical Physics*, 158(3):601–627, 2015.

[12] V. Dossetti and F. J. Sevilla. Emergence of collective motion in a model of interacting brownian particles. *Physical review letters*, 115(5):058301, 2015.

[13] A. Cavagna, L. Del Castello, S. Dey, I. Giardina, S. Melillo, L. Parisi, and M. Viale. Short-range interactions versus long-range correlations in bird flocks. *Physical Review E*, 92(1): 012705, 2015.

[14] D. H. Kelley and N. T. Ouellette. Emergent dynamics of laboratory insect swarms. *Scientific reports*, 3(1):1–7, 2013.

[15] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34, 1987.

[16] A. Huth and C. Wissel. The simulation of the movement of fish schools. *Journal of theoretical biology*, 156(3):365–385, 1992.

[17] H. Kunz and C. K. Hemelrijk. Artificial fish schools: collective effects of school size, body size, and body form. *Artificial life*, 9(3):237–253, 2003.
[18] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Physical review letters*, 75(6):1226, 1995.
[19] G. Grégoire and H. Chaté. Onset of collective and cohesive motion. *Physical review letters*, 92(2):025702, 2004.
[20] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.
[21] M. Camperi, A. Cavagna, I. Giardina, G. Parisi, and E. Silvestri. Spatially balanced topological interaction grants optimal cohesion in flocking models. *Interface focus*, 2(6): 715–725, 2012.
[22] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, et al. 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.
[23] W. Poel, C. Winklmayr, and P. Romanczuk. Spatial structure and information transfer in visual networks. *Frontiers in Physics*, page 623, 2021.
[24] M. Zumaya, H. Larralde, and M. Aldana. Delay in the dispersal of flocks moving in unbounded space using long-range interactions. *Scientific reports*, 8(1):1–9, 2018.
[25] N. Miller and R. Gerlai. From schooling to shoaling: patterns of collective motion in zebrafish (danio rerio). *PloS one*, 7(11):e48865, 2012.
[26] W. L. Romey. Individual differences make a difference in the trajectories of simulated schools of fish. *Ecological Modelling*, 92(1):65–77, 1996.
[27] V. Mirabet, P. Auger, and C. Lett. Spatial structures in simulations of animal grouping. *Ecological modelling*, 201(3-4):468–476, 2007.

The work presented in this chapter has been published in:
Royal Society Open Science (Van Havermaet et al. 2023)

3

Steering herds away from dangers in dynamic environments

Abstract
In this chapter, we focus on a different task of nature conservation, namely protection of fish schools[1]. Although fish are an important part of our natural environment, they are increasingly faced with dangers such as illegal fishing, pollution, and invasive species. Robots capable of guiding fish away from these dangers could provide a viable solution. In the literature, guiding a group of autonomous individuals in a desired direction is known as the shepherding problem. So far, only single-robot or centralized multi-robot solutions have been proposed. The former is unable to observe dangers at any place surrounding the fish, and the latter does not generalize to unconstrained environments. Therefore, we propose a swarm robotic shepherding algorithm in which robots maintain a caging formation around the fish school to detect nearby dangers. When a threat is detected, part of the swarm positions itself to repel the fish toward a safer area. The algorithm is evaluated across different collective motion models and tested in two dynamic scenarios: (i) avoiding dangerous patches that appear over time and (ii) staying within a safe circular enclosure. Simulation results show that the robots successfully shepherd the fish when the group remains cohesive and a sufficient number of robots are deployed.

[1]In the main text of this chapter, we use the term "herd" to align with the terminology commonly used in shepherding literature, as this work was originally written for an international journal paper. However, the parameter settings of the "herd" in our study correspond specifically to fish behavior.

3.1 Introduction

Agents capable of guiding a group of other autonomous agents are essential to a wide range of applications [2]. Robotic agents can be used to herd animals such as sheep [3, 4], cattle [5], and ducks [6], to enable crowd control [7], to keep birds away from aircraft [8, 9], and to bring people to safety [10]. In the literature, this type of guidance is studied under the name of shepherding, as it is inspired by dogs herding sheep to a desired location [11]. Sheepdogs are particularly skilled at shepherding, with a single dog capable of herding more than 80 sheep [12]. These dogs have learned to exploit the sheep's collective behavior of aggregating and together escaping from a threat [13]. This behavior is widely accepted as an example of the selfish herd theory [14]. Supporting evidence for this theory has been found for other animal species. For instance, fish move together in schools to reduce predation risk [15].

The individual mechanisms that animals apply to establish and maintain aggregated formations have been studied broadly in research. Various models of collective motion have been proposed in both theoretical [16–19] and empirical [20, 21] studies over the past decades. A recurring feature in these animal models is that the individual behavior results from a few simple rules that take as input the motion of a relatively small set of neighbors [22, 23]. Reynolds [24] proposed one of the earliest models for the flocking of birds based on three distinct rules: (i) to avoid collision with close neighbors, (ii) to move in the same direction and at the same speed as others, and (iii) to remain as a cohesive group. These interactions have been proposed as the founding blocks for the underlying behavioral mechanisms in fish schools [25], mammal herds [26], pedestrian crowds [27], and other vertebrates [28]. Most notably, Couzin et al. [29] showed how the information from only a few informed individuals can propagate to the entire collective system. When some individuals gain information about the location of a danger and consequently change their directions towards a safer region, these changes propagate in the entire group. Hence, the actions of a single sheepdog can influence the collective motion of the entire herd of sheep, even though only a small part of the herd directly observes the sheepdog.

Robotic agents, visually styled to trigger an aversive response of the herd, can rely on the same shepherding mechanisms as natural perceived threats, like dogs [30]. This corresponds to the robot computing the optimal motion control vector based on force vectors representing the interactions among herd members, and the interaction between the herd and a shepherd. In the classic shepherding problem, the shepherds are tasked with guiding the herd to a certain goal location, that is automatically known to the shepherds as prior information. In this chapter, however, we study the problem setting where this assumption is relaxed. The shepherds must actively determine the goal location based on local infor-

mation and communication. Such shepherds could be beneficial in protecting animal herds from unforeseen dangers in dynamic environments, by guiding them to a safer location. To mimic real-life use cases, we thus consider that the robotic shepherds can only obtain and communicate information in a local radius, and they have no prior information about the dangers. Once a danger has been detected, the shepherds are tasked with guiding the herd away from this danger. To do so, they must first reach a consensus on the direction in which to guide the herd, which requires a shared directional frame among them.

We propose a solution to the previously described problem where the herd is consistently surrounded by multiple shepherds. We also refer to this formation as *caging* the herd, as the shepherds are evenly spread out across the contour of the herd. As such, the shepherds can detect any nearby danger approaching the herd. Additionally, with a sufficient number of shepherds, the task of steering the herd can be allocated to the shepherds already present in the appropriate region based on the location of the detected danger. Similar to other shepherding research works, we model a member of the herd to move in the opposite direction of a shepherd when the relative distance is lower than a certain threshold. The shepherds positioned between the herd and a danger will therefore move close enough to the herd, which should trigger the herd to change direction. When the herd is not approaching a danger, the shepherds remain far away enough, which allows the herd to continue their natural behavior (e.g. foraging).

Long et al. [2] stated several challenges in their literature review of robotic shepherding. We believe our approach advances the state-of-the-art of shepherding with respect to three of these challenges. Firstly, the models they have reviewed are not easily transferable to dynamic environments. In our work, the proposed algorithm is demonstrated in two dynamic environments: (i) dangerous patches appear nearby the herd at a certain probability, and (ii) a circular safe zone containing the herd decreases in size over time. Secondly, a shepherd should be able to dynamically adapt the distance threshold where the herd tends to move away from the shepherd, while being limited in energy consumption and computation time. Therefore, we design the artificial intelligence of the shepherds through a set of control rules, where the threshold is explicitly incorporated in the computations. Thirdly, Long et al. argue that practical applications of shepherding require robustness with regard to the failure of robotic agents. Hence, we propose a decentralized, multi-agent shepherding algorithm where each shepherd gathers information from local radius-based observation and is able to communicate information to nearby neighbors. Contrary to the majority of the literature so far [31], no central unit providing commands or global information of the herd is available to the shepherds in our work.

We apply our approach specifically to the shepherding of fish, motivated by the biomimetic robotics state-of-the-art of fish-like robots [32] and

the potential environmental impact. Although fish are an important part of our natural environment [33], they are increasingly faced with dangers such as illegal fishing [34], pollution [35], and invasive species [36]. The efficiency of the proposed algorithm is demonstrated through simulation, where four different models of collective motion are used to simulate the fish behavior. One of these models was proposed by Couzin et al. [25] and has been shown by multiple empirical studies to capture the key features of fish behavior such as nearest-neighbor distance, polarization, group speed, and turning rate [21, 37–41]. More specifically, we simulate the herd as guppies who live in shallow waters, which applies to our proposed algorithm designed for two-dimensional environments [42].

The rest of this chapter is organized as follows. Section 3.2 briefly discusses the most prominent research works on the topic of shepherding and the studies that have inspired our proposed algorithm. In section 3.3, we formulate the three main parts of the problem scenario: (i) the four considered models of collective motion to simulate the herd, (ii) the definition of caging, and (iii) the descriptions of the shepherding tasks. Next, in section 3.4, an algorithm for establishing and maintaining a caging formation is proposed, and afterwards one for shepherding (while caging) the herd in the presence of danger. Results of the respective main parts of the problem scenario are discussed in section 3.5. Finally, we conclude the chapter in section 3.6.

3.2 Related Work

3.2.1 Shepherding

The Robot Sheepdog Project [43] was one of the first research projects to develop a robot capable of solving the classic shepherding problem, based on a computation of force vectors that represent the inter-individual rules proposed by Reynolds [24]. A single robot used a ceiling-mounted camera to track a flock of ducks and manipulate their movement. In order to direct the flock in the right direction, the robot positions itself on the opposite side of the flock to the goal. When the robot is correctly aligned with respect to the goal and the flock, it advances towards the flock and hence the flock moves towards the goal. However, the robot and the ducks were placed in a circular enclosure where the goal location was always placed at the edge of the enclosing circle. The robot shepherded the ducks to move along the edge until the goal location was reached.

Some follow-up works have adapted the aforementioned algorithm in order to shepherd in unconstrained environments. Strömborn et al. developed a single-agent shepherding algorithm based on force vectors, where the shepherd switches between collecting dispersed herd members and steering the cohesive herd [44]. This results in a side-to-side motion of the shepherding agent behind the herd, which mimics the behavior

of real sheepdogs. Miki and Nakamura developed a similar algorithm, where the shepherd adaptively switches between collecting and steering the herd [45], but included a notion of cooperation between multiple shepherds who avoid overlapping. Their experiments showed that two shepherds are more efficient in guiding the herd than only one. Other studies corroborated the finding that single-agent solutions are limited to smaller herd sizes [46], and thus multiple shepherds can control large herds more efficiently than a single shepherd [47].

However, most shepherding control approaches assume that the shepherds have global knowledge of the positions of every individual in the environment [2]. Applying robotic shepherding to any environment means that the robots can only rely on local information gathered by sensors with limited range [48–50]. Tsunado et al. proposed an algorithm where a single shepherd constantly aims to repel the furthest herd member from the goal location towards it, using only information collected via a simulated local camera [31]. This algorithm was shown to be successful in simulations, while the algorithm proposed by Strömborn et al. [44] mostly failed to guide the herd to the goal when only local information is available to the shepherd. The challenge of moving a herd to a goal with multiple shepherds, using only local information, was tackled by Lee and Kim [51]. Shepherds coordinate to aggregate one cohesive herd, where some steer wandering members towards the main herd, while others focus on keeping the main herd at its current position. Only when the herd is cohesive enough, each shepherd repels the closest herd member towards the goal location.

A simpler but effective algorithm was presented by Miki et al. [45], where the shepherd moves in a circular motion behind a member of the herd (relative to the guidance direction) and then comes closer to repulse the herd member. Each herd member followed the traditional Reynolds flocking algorithm in their work, where stochastic behavior is considered. Experiments were conducted with one shepherd steering a herd of 25 individuals, and two shepherds cooperating to steer 30 individuals. This steering algorithm is similar to the one we present in this chapter, where the shepherding robots have to cooperate and only obtain partial knowledge of the environment through local sensing.

As discussed beforehand, in the problem setting that we address, the shepherds are required to maintain a *caging* formation while guiding the herd away from dangers. Varava et al. [52] refer to the problem of steering a herd while maintaining a caging formation as *herding by caging*. This work was the first, and to our knowledge remains the only, study to explicitly apply the concept of caging to the context of shepherding. They proposed a centralized RRT-based (rapidly-exploring random tree) algorithm using computational topology techniques to verify the caging formations. However, their approach assumes global information and an initial correctly caged formation. Furthermore, they did not consider the natural

motion of the herd in the absence of robots, and their simulations suggest that small increases in the number of robots result in a significant decrease in the algorithm's performance.

Our approach for shepherds to establish and maintain a caging formation is inspired by the computational modeling study of wolf-pack hunting strategies presented in Muro et al. [53]. For a stationary single prey, the wolves become arranged in a stable configuration of a regular polygon, which we refer to as a caging formation. While the authors do not provide any formulas, they state that wolves hunt following only two simple decentralized rules: (i) the wolf moves towards the prey until a safe distance threshold is reached, and (2) when close enough to the prey, the wolf moves away from the other wolves that are close to the safe distance to the prey. To successfully follow these rules, the wolves do not need to rely on direct communication, nor is there a role of leader needed in the group. Hence, these rules are applicable to our problem scenario, where the shepherds operate in a decentralized system.

3.2.2 Formation control

The contributions of this chapter can also be situated within the broader context of multi-robot formation control. In particular, the caging approach used to steer the herd aligns with objectives found in the formation control literature, especially those involving decentralized shape controllers [54, 55]. For example, Fink et al. [56] proposed a decentralized method in which a team of agents surrounds and transports an object toward a designated goal. Their cooperative manipulation strategy is based exclusively on the local detection of each agent, combined with a shared understanding of the global task. Similarly, our method establishes a caging formation and steers the herd using only the local information available to each agent. Local communication is used solely to reach a consensus on the optimal escape direction when multiple robots detect nearby dangers. For this purpose, the agents require a shared directional reference frame. However, a key distinction lies in the nature of the caged entity and the resulting formation dynamics. Whereas prior approaches target passive objects, our scenario involves multiple active agents exhibiting both intrinsic (self-propelled) and collective (mutually influential) motion. Consequently, the proposed caging algorithm enables the robot team to maintain a dynamic, reactive formation that continuously adapts to the unpredictable movements of a living group.

A related distinction arises when comparing our work to formation–containment control approaches, in which a group of follower agents is enclosed within a convex hull formed by a set of leader agents [57, 58]. In these approaches, follower agents do not exhibit autonomous behaviors such as wandering or flocking. Their objective is to converge to, and remain within, the convex hull maintained by the leader agents, who

track a time-varying formation. In this chapter, the roles are reversed: the robots (serving as "leader agents") aim to contain the herd (the "followers") within their self-organized convex hull. As the herd operates according to a biologically grounded model of collective motion, the robots must continuously adapt to its dynamics in real time, rather than enforcing a predefined formation.

3.3 Problem Formulation

Let $\mathcal{A}$ and $\mathcal{H}$ denote the respective sets of robotic agents and herd members. The state of the discrete-time system at time t is defined by the position $p_i(t) \in \mathbb{R}^2$ and orientation $\theta_i(t) \in [0, 2\pi)$ of each individual $i \in \mathcal{A} \cup \mathcal{H}$. At every time step, an agent first rotates with angular velocity w to its desired orientation $\hat{\theta}$ that is computed based on local information. The agent stops rotating once $\hat{\theta}$ is reached, or when the time interval has passed. In the remainder of the time interval, the agent moves straight forward at linear velocity v. At every time step, Gaussian noise σ is added to the orientation of each individual.

3.3.1 Collective motion models

Each herd member updates its direction of motion based on three concentric non-overlapping zones (see Figure 3.1), containing distinct subsets of neighbors. Each zone corresponds to a type of interaction; (i) repulsion from others inside the disk with radius z_R, to establish a minimum inter-individual distance, (ii) alignment of orientation with others inside the annulus with width z_O, to all move in the same direction, and (iii) attraction to others inside the annulus with width z_A, to remain as one cohesive group.

In this chapter, we consider four different models of collective motion that have been empirically studied in fish schools [59]: metric, topological, visual reconstruction and long-range (see Figure 3.2). These models differ

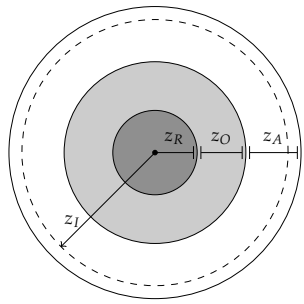

Figure 3.1: Schematic representation of the behavioral zonal model

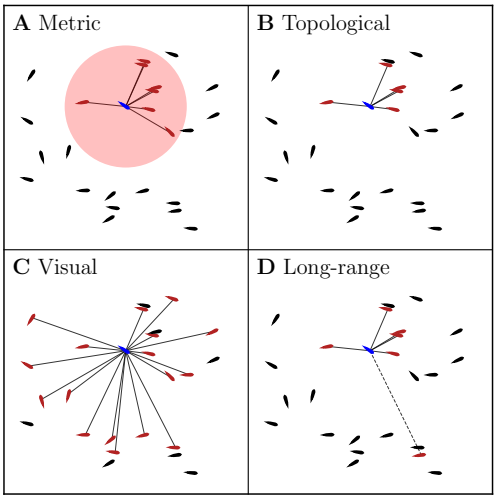

Figure 3.2: The set of neighbors (red) are illustrated in regards to a focal individual (blue) for the different models of collective motion. (**A**) Metric: all individuals within a certain radius. (**B**) Topological: only the k-nearest neighbors ($k = 5$ in this example). (**C**) Visual: all individuals that are visually observable, i.e. not obstructed by any other neighbor. (**D**) Long-range: in addition to topologically selecting neighbors (solid line with $k = 5$), there are λ_i randomly selected neighbors (dashed line with $\lambda_i = 1$).

from one another in the way that neighbors $\mathcal{N}_i$ of each individual i are selected. Using the metric model, the set of neighbors consists of all other individuals within the interaction-radius $z = z_R + z_O + z_A$. In the topological model, $\mathcal{N}_i$ only contains the k nearest neighbors. Visual reconstruction only selects other individuals that can be visually perceived as neighbors, which means no other individual obstructs the focal individual i from observing the peripheries of a neighbor. Finally, in the long-range model, short-range interactions are considered by selecting k nearest neighbors (i.e. topologically), while long-range interactions are defined by randomly selecting λ_i neighbors from the remaining members of the herd. The number of long-range neighbors λ_i is sampled from a Poisson distribution with average λ as parameter of the model.

To model the aversive behavior of a herd individual i interacting with a robotic agent, an additional disk with radius z_I surrounding i is introduced. A herd individual wants to move in the opposite direction (i.e. repulsion) of any robot present in this zone. Later in this chapter, we will introduce two different scenarios for shepherding, where $\mathcal{D}$ is a set containing (i) the centers of dangerous patches or (ii) the center of a safe circular enclosure. The herd interacts with these dangers likewise to robotic agents.

Let $\mathcal{N}_i^R$, $\mathcal{N}_i^O$ and $\mathcal{N}_i^A$ then denote the distinct subsets of neighbors by separating $\mathcal{N}_i$ based on the repulsion, orientation, and attraction zones respectively. Thus, each neighbor is only assigned to one of the subsets; $\mathcal{N}_i^R \cap \mathcal{N}_i^O \cap \mathcal{N}_i^A = \emptyset$. Let $\mathcal{A}_i$ and $\mathcal{D}_i$ denote the respective subsets of robotic agents and dangerous patches located in a radius of z_I around p_i. In the scenario of the enclosure, $\mathcal{D}_i$ contains the center of the safe disk if and only if the individual i is located outside of the disk.

Let us define the relative position of individual j from i as $r_{ij}(t) = p_j(t) - p_i(t)$. Furthermore, let q_i be the motion vector of an individual i, which is computed as follows:

$$\begin{aligned} q_i(t) = &- \alpha_R \sum_{j \in \mathcal{N}_i^R} \frac{r_{ij}}{\|r_{ij}\|} + \alpha_O \sum_{j \in \mathcal{N}_i^O} \frac{q_j}{\|q_j\|} \\ &+ \alpha_A \sum_{j \in \mathcal{N}_i^A} \frac{r_{ij}}{\|r_{ij}\|} - \alpha_I \sum_{j \in \mathcal{A}_i \cup \mathcal{D}_i} \frac{r_{ij}}{\|r_{ij}\|}, \end{aligned}$$

with weights $\alpha_R \geq 0$, $\alpha_O \geq 0$, $\alpha_A \geq 0$ and $\alpha_I \geq 0$ of repulsion, orientation, attraction, and animal-robot interactions respectively.

3.3.2 Caging

A caging formation can be constructed by the robotic agents based on the repulsive animal-robot threshold z_I. When two robotic agents are at a distance lower than $2z_I$ of each other, they exert a combined repulsive force on the herd which prevents them from intersecting the path between those agents. In other words, caging is equivalent to a closed chain formation where the distance between consecutive agents satisfies the upper bound of $2z_I$.

In order to measure whether the agents established an appropriate caging formation, we see if a polygon can be constructed from the edges between agents where the length is shorter than $2z_I$. Following [52], we then define the herd to be successfully caged, if and only if, the polygon is closed and the entire herd is located in the interior of the polygon.

3.3.3 Shepherding

The agents' objective is to prevent any herd member from entering dangerous areas. When a danger is detected to be approaching the herd, the agents are tasked to preemptively steer the herd away. In the absence of dangers, the agents should remain at a minimum distance of z_I from the herd to avoid any unnecessary stress induced on the herd. The agents have no prior knowledge about the positions and the movements of the dangers, which can only be locally observed.

In this chapter, we consider two different shepherding tasks based on common scenarios in which animal herds would benefit from safety interventions. In the first task, a stationary dangerous patch appears in front

of the herd at a certain probability. This task resembles several use cases of different dangers such as pollution (e.g. fish near oil leaks) and poachers. In the second task, the robots are expected to ensure that the herd remains in a given safe zone. This task is similar to deploying virtual fences. The use of mobile robots also allows for the safe zone to be dynamically changed, which is useful for caretaking of animals. For example, it allows to navigate the herd between nests and optimize grazing patterns.

We model the task environment by a time-variant potential function f that reaches a local minimum in the mean direction opposite of all dangers (assuming that each danger is equally important to avoid). Consequently, the herd should move in the direction of the potential gradient ∇f. In the first task, agent i computes the gradient based on the relative positions of the observable dangerous patches $\mathcal{D}_i$ as follows: $\nabla f(p_i) = -\sum_{d \in \mathcal{D}_i} \frac{p_d - p_i}{\|p_d - p_i\|}$. In case the agent does not detect any dangers, the gradient is undefined. In the second task, the potential gradient is defined as $\nabla f(p_i) = \frac{p_e - p_i}{\|p_e - p_i\|}$ with p_e the center of a circular safe zone. We make the simplifying assumption that each robot has access to the position of the safe zone center. This can be interpreted as a minimal form of shared task information that could be obtained via environmental cues. This assumption allows us to model directional guidance toward the interior of the zone in a tractable way. In case the agent is positioned within the safe zone and does not detect the boundary, the gradient is undefined. As we propose a decentralized multi-agent solution to this shepherding problem, the potential function is a way of representing locally observed information that will be communicated between agents.

3.4 Algorithm

3.4.1 Caging

We describe the algorithm from the perspective of an individual agent $a_i \in \mathcal{A}$. Let $\mathcal{A}_n$ and $\mathcal{H}_n$ respectively represent the neighboring subsets of agents and herd members, which are located within the detection distance d_d from a_i. In case the agent is unable to detect any member of the herd, an arbitrary search method is applied (e.g. a random walk). We opt for an adaptation of the herd motion model presented in subsection 3.3.1 to ensure neighboring agents remain within communication range of each other. Agents share whether they have detected herd members with nearby agents. This communication is minimal and can be implemented as a one-bit signal (e.g., via a light cue or beacon) indicating the presence of a nearby herd member. Consequently, agents that are unable to directly observe the herd will be attracted toward neighboring agents who have signaled a positive detection.

When members of the herd are directly detected, the agent attempts to reside at a given distance R^* from the closest herdable $h^* \in \mathcal{H}_n$. The agents

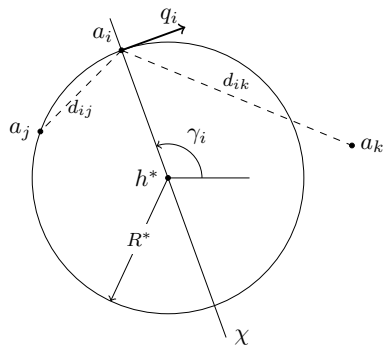

Figure 3.3: Illustration of the proposed caging method. Agent a_i follows tangential motion along the circle defined the position of the closest herdable h^* as the center and R^* as the radius. From the two closest neighboring agents a_j and a_k of opposite sides of the axis χ, the agent a_i moves towards the neighbor which has the largest relative distance. Here, the motion vector $\vec{q}_i$ of a_i indicates clockwise rotation towards a_k as $d_{ik} > d_{ij}$.

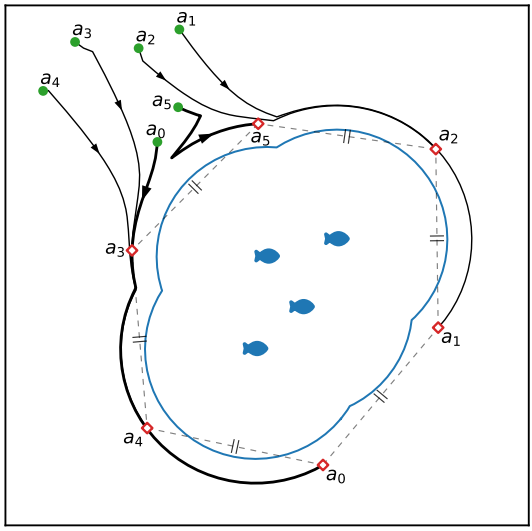

Figure 3.4: Trajectories (black arrows) of agents $\mathcal{A} = \{a_i\}_{i=0\ldots5}$ caging a stationary herd of four (blue fish). The agents are initially located (green dots) in a box of density ρ_a near the herd, from which they establish and maintain an equidistant caged formation in steady-state (red diamonds). At all times, the agents remain at a distance R^* based on the herd's repulsive zone (blue contour) of radius z_I. The highlighted (thicker arrows) trajectories of a_0 and a_5 respectively show how agents could either continuously move by the same direction or change directions due to the movement of neighbors.

should be close enough to observe, follow, and quickly interact with the herd when needed. However, they should not be closer than z_I, as this would otherwise induce unwanted stress on the herd. While approaching to and residing at the circular boundary of h^*, the agent a_i attempts to remain equidistant from the two closest neighboring agents a_j and a_k from opposite sides of the axis χ defined by the bearing $\gamma_i \in [-\pi, \pi)$ of a_i from h^*. As shown in Figure 3.3, the neighboring agents can be positioned anywhere in their respective half-plane.

Every agent attempts to maintain a distance of R^* from the closest herd member, while moving to position themselves at equal distance from the two neighboring agents. In other words, the agents moves along the boundary of the union of circles defined by the positions of every herd member as the center points and R^* always as the radius. Figure 3.4 shows how the agents follow this boundary, which can be partially seen by the black trajectories (For illustrative purposes, the herd does not move). To provide full flexibility, an agent should be able to eventually re-encounter any point on the boundary after moving along the boundary in the same direction for enough time. This means that the union of circles is a connected set.

Assumption 3.1. *The union of the circles, where the center points are the positions of every herd member $h \in \mathcal{H}$ and the radii are R^*, is a connected set.*

In order to approach and rotate along a circular path of radius R^*, we apply the following method to compute the agent's desired orientation $\hat{\theta}$ as proposed by [60]:

$$\hat{\theta} = \gamma_i + \vartheta \left(\frac{\pi}{2} + \arctan\left(\kappa(d^* - R^*)\right) \right) \tag{3.1}$$

where $\vartheta \in \{-1, 1\}$ determines the direction; i.e. clockwise ($\vartheta = -1$) or counterclockwise ($\vartheta = 1$), $\kappa > 0$ influences the rate of transition between moving towards and moving along h^*, and d^* is the distance from a_i to h^*.

To obtain the cage formation, the value of ϑ is computed based on the relative distances to the neighboring agents. We define the subsets of neighboring agents $\mathcal{A}_j \subseteq \mathcal{A}_n$ and $\mathcal{A}_k \subseteq \mathcal{A}_n$ divided by the axis χ as follows:

$$\mathcal{A}_j = \{j \mid a_j \in \mathcal{A}_n \wedge 0 < \operatorname{sgn}(\gamma_i) \cdot (\gamma_i - \gamma_j) \leq \pi\}$$
$$\mathcal{A}_k = \mathcal{A}_n - \mathcal{A}_j$$

with γ_j as the bearing of neighboring agent a_j from h^*. The shortest distances to each neighbor subset, $d_{ij} = g_i(\mathcal{A}_j)$ and $d_{ik} = g_i(\mathcal{A}_k)$ are computed by:

$$g_i(U) = \begin{cases} \min_{u \in U} \|p_u(t) - p_i(t)\| & \text{if } U \neq \emptyset \\ 2z_I & \text{otherwise} \end{cases},$$

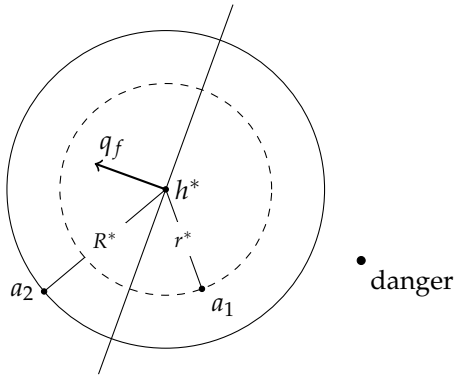

Figure 3.5: Illustration of the proposed shepherding method. Agent a_1 is located behind the closest herdable h^*, in perspective of the mean orientation θ_f (illustrated by the corresponding vector q_f with perpendicular axis). Therefore, a_1 positions itself at a distance of r^* from h^*. Agent a_2, positioned in front of h^*, remains at a distance of R^*.

where the agent assumes a distance of $2z_I$ in case the neighbor subset is empty. The direction ϑ is then given by

$$\vartheta = \begin{cases} \text{sgn}(\gamma_i) & \text{if } d_{ij} < d_{ik} \\ -\text{sgn}(\gamma_i) & \text{otherwise} \end{cases}.$$

The orientation $\hat{\theta}$ of the motion vector to move along the circular path of h^* is computed with Equation 3.1 based on the direction ϑ. We define the magnitude of this motion vector as $\eta = \frac{1}{2}|d_{ij} - d_{ik}|$, so that two agents moving towards each other will eventually reach a stable solution. This motion vector $\langle \hat{\theta}; \eta \rangle$ is then added with the predicted motion vector of the closest member of the herd, such that the cage formation is at the appropriate relative distance of the herd. Each member of the herd is predicted to move at mean velocity $\bar{v}_h$ in its current orientation. Thus, the predicted motion vector of the nearest herd member is $\langle \theta_{h^*}(t); \bar{v}_{h^*}(t) \rangle$. Note that $\eta = 0$ when the agent has positioned itself at equal distance from both neighbors, which results in the agent aligning its orientation and velocity with the herd. However, when the agent is not equidistant to a_j and a_k, we ensure that the maximum velocity is not exceeded as the magnitude η is upper bounded by $-\bar{v}_{h^*}(t)\cos\left(\theta_{h^*}(t) - \hat{\theta}(t)\right) + \sqrt{\bar{v}^2_{h^*}(t)(\cos^2(\theta_{h^*}(t) - \hat{\theta}(t)) - 1) + v^2_{\max}}$, with $v_{\max}$ as the maximum linear velocity of the robotic agents (See section 3.A for a derivation of the bound).

3.4.2 Shepherding

The idea behind the proposed algorithm is based on each individual agent adaptively changing between two caging formations, where one causes a repulsive force on the herd and the other does not.

Each agent a_i observes the gradient of the local potential $\nabla f(p_i)$ as described in the problem formulation (see subsection 3.3.3). A message containing a unique agent identifier, version and the local gradient is communicated to all nearby neighbors within the communication range d_c. Whenever a new local gradient is observed, the agent updates its own version and sends a new message. The received messages are then filtered so that only messages with the newest version of each agent remain. All filtered messages are then forwarded to their nearby neighbors. The local gradients of the filtered messages are combined as $q_f = \sum_j \frac{\nabla f(p_j)}{\|\nabla f(p_j)\|}$, with a_j as an agent of which the gradient (with most recent version) is defined and received by a_i. In case any a_j exist (including a_i), we attempt to steer the herd in the direction of q_f. In this step, we have assumed that the agents share a directional frame. This allows the agents to reach consensus in the mean direction when combining multiple locally observed potential gradients.

Assumption 3.2. *All agents have a shared directional frame.*

Based on the aversive behavior described in subsection 3.3.1, a robotic agent is capable of triggering a herd individual to move in the direction opposite along the axis between the herd individual and this agent, when the relative distance is lower than z_I. If an agent is positioned behind all nearby herd members, in the direction of q_f, then it performs the caging algorithm to remain at the circular boundary of the closest herdable at a distance of r^*, with $r^* < z_I < R^*$. The distance r^* can be dynamically changed, as long as the upper bound is satisfied. In our experiments, we use a fixed value. In the other case, positioned in front of the nearby herd, the agent executes the caging algorithm with a distance of R^*. Figure 3.5 shows how agents remain at different distances from the herd, based on their position relative to the direction of q_f. More specifically, agent a_i steers by a caging with a radius of r^* if the smallest angle γ_i^h between the agent and each individual of the nearby herd $h \in \mathcal{H}_n$ is greater than $\frac{\pi}{2}$:

$$\arctan\left(\sin\left(\theta_f - \gamma_i^h\right), \cos\left(\theta_f - \gamma_i^h\right)\right) > \frac{\pi}{2},$$

with θ_f as the orientation of the vector q_f.

If no messages of gradients are received, the agents remain in a caging formation at the distance R^*.

3.5 Results

First in subsection 3.5.1, we study the behavior of a simulated herd, when there are no robotic agents nearby, following the four different models of

collective motion. We vary the most relevant parameters in regards to cohesiveness of the group [17], the widths z_O and z_A of the orientation and attraction zones respectively, in order to examine in which parameter range Assumption 3.1 holds. Afterwards in subsection 3.5.2, we run an experiment with agents deployed in the environment that follow the proposed caging algorithm. We consider the same models of collective motion and varying parameter values (z_O and z_A) of the herd, which allows us to verify if caging indeed only fails when Assumption 3.1 does not hold. Finally in subsection 3.5.3, we run experiments of the two aforementioned shepherding tasks: (i) agents steering the herd away from randomly appearing dangerous patches, and (ii) agents keeping the herd inside a safe circular zone. Results were obtained from simulation runs with 10 seeds. All experiments are run with the following parameters unless explicitly stated otherwise; $N_H = 100$, $z_I = 19$, $w_h = w_a = \frac{\pi}{2}$, $v^a_{\max} = 4$, $v^h_{\max} = 2$, $d_d = d_c = 3z_I$, $\sigma_a = \sigma_h = 0.05$, $\alpha_R = 100$, $\alpha_O = 50$, $\alpha_A = 1$ and $\alpha_I \in [500, 2500]$.

3.5.1 Collective motion models

The herd is simulated, in the absence of any robotic agents, under the different models of collective motion; (A) metric, (B) topological, (C) visual, and (D) long-range. We simulate the herd in 2D open space environment. To facilitate immediate interaction between the N_H individuals at the beginning of a simulation, they are placed within a square region of the size $(\frac{N_H}{\rho_h})^{\frac{1}{2}}$ with initial density $\rho_h = 0.01$. Each individual's position and moving direction are initially uniformly distributed.

We aim to study which parameter values lead to fragmentation of the herd, since this would prohibit agents from successfully caging the herd. Thus, we say that the herd is fragmented when Assumption 3.1 does not hold. As previously described, a graph can be constructed where the vertices are the positions of every herd member, and there is only an edge between two vertices if the relative distance is lower than or equal to R^*. Based on this graph, the degree of fragmentation is measured as the number of groups N_G where each distinct group is connected. We vary the respective widths z_O and z_A of the orientation and attraction zones over $[0, 100]$. For the parametric models (topological and long-range), we chose values that produce most similar results of fragmentation to the other two models. This allows us to make a fair comparison between models. We set $k = 50$ in the topological model, since lower values of k lead to only higher probabilities of fragmentation, as shown by [17]. We verify this by running simulations for $k \in \{10, 25, 50\}$ (see section 3.C). This led them to proposing the long-range model, for which we use their proposed values of $k = 7$ and $\lambda = 0.1$. As visual reconstruction, a ray-casting algorithm is used in this chapter where the individuals are simulated guppy fish based on physical measurements.

In addition to the number of groups, we examine the relative area cov-

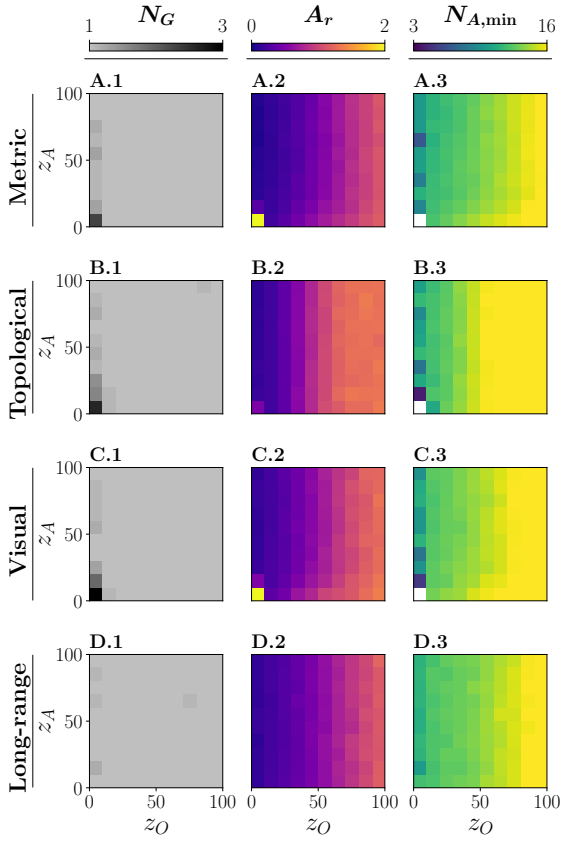

Figure 3.6: Quantitative herd measurements ($N_H = 100$) obtained in simulation without robotic agents, following different models of collective motion: **(A.1-3)** metric, **(B.1-3)** topological with $k = 50$, **(C.1-3)** visual, and **(D.1-3)** long-range with $k = 7, \lambda = 0.1$. For each model, the widths of orientation (z_O)and attraction (z_A) zones are varied. The number of groups N_G **(A-D.1)** shows that group fragmentation only occurs with for smaller values of z_O, independent of the width of zone of attraction z_A, for all models (Although, the long-range model causes significantly less fragmentation). The relative area coverage A_r **(A-D.2)** increases with z_O, mostly independent of z_A. When the herd remains as one cohesive group, the theoretical minimum number of agents $N_{A,\min}$ needed to form a caging pattern **(A-D.3)** increases with z_O in correlation with the relative coverage A_r for each model respectively, since more agents are needed to cage a larger area. On the other hand, $N_{A,\min}$ is high when fragmentation occurs.

erage A_r which is computed as the ratio of the areas of the convex hulls of the herd at convergence time T over initial time t_0. Related to this, the minimum number of robots needed to form a caging pattern $N_{A,\min}$ is computationally estimated by placing agents on the boundary of the union of circles, described in Assumption 3.1, at a maximum distance of $2z_I$ from one another. Based on these results, we can test our hypothesis that the proposed caging algorithm fails if and only if Assumption 3.1 does not hold or, there was an insufficient number of robotic agents deployed.

Figure 3.6 shows the aforementioned quantitative measurements for the four considered models, and varying the widths z_O and z_A of the orientation and attraction zones respectively. Figure 3.6A-C.1 shows that the herd becomes fragmented ($N_G > 1$) for lower values of the orientation width ($z_O < 10$), using a metric, topological, or visual model. In this range of z_O, the herd fragments into the most groups when z_A is lowest. As opposed to the other models, Figure 3.6D.1 shows that fragmentation is unlikely to occur with the long-range model, independent of z_O and z_A. Furthermore, we see in Figure 3.6A-D.2 that the relative area A_r is also dependent on the width of the orientation zone z_O, while mostly independent of z_A. The herd becomes denser than its initial distribution ($A_r < 1$) for the lowest values of z_O. Increasing z_O leads to increasing the relative area coverage until A_r is approximately equal to 1 or slightly higher. Note how in the topological model A_r eventually stabilizes once z_O is approximately higher than 60. Evidently, when the orientation zone is so large that each individual aligns its orientation with all its neighbors, the herd will maintain its initial area. On the contrary, when z_O is small enough so that neighbors also appear in the attraction zone, the herd will become more compact. As the relative area A_r increases with z_O, so does $N_{A,\min}$ (Figure 3.6A-D.3) as more robotic agents are needed to form a caging pattern where consecutive robots should remain at a distance lower than $2z_I$. Note that the theoretical minimum of number of agents is not computed when the herd is fragmented ($N_G > 1$). In the figure, a white square means no possible value as all seeds contain fragmentation.

3.5.2 Caging

In this second experiment, we placed N_A robotic agents following the proposed caging algorithm in the same environment as the herd. The swarm of N_A agents is placed in a square region based on initial density $\rho_a = 0.01$, at a random but detectable position from the school (i.e. within distance d_d). Figure 3.7 shows the ratio of successfully caging a herd with different number of agents (1) $N_A = 10$, (2) $N_A = 20$, and (3) $N_A = 30$, and for the different collective motion models of the herd.

Figure 3.7A-D.1 shows that, in the range of $z_O > 10$, deploying $N_A = 10$ robotic agents is inadequate to successfully cage a herd of size $N_H = 100$. This was expected as $N_A = 10 < N_{A,\min}$ (See the first column inside the

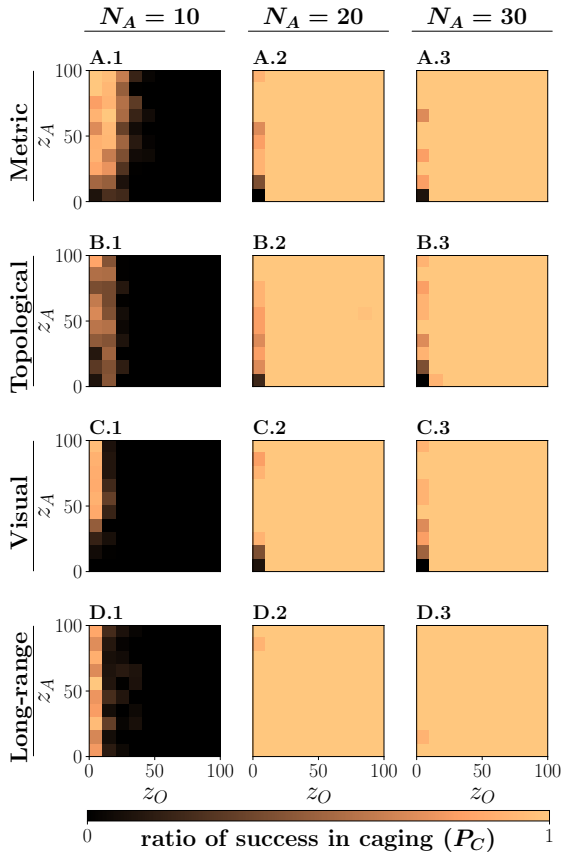

Figure 3.7: The ratio of robots successfully caging a herd of size $N_H = 100$, following different models of collective motion: **(A.1-3)** metric, **(B.1-3)** topological with $k = 50$, **(C.1-3)** visual, and **(D.1-3)** long-range with $k = 7, \lambda = 0.1$. Additionally, different sizes of the robot swarm N_A are studied: **(A-D.1)** $N_A = 10$, **(A-D.2)** $N_A = 20$ and **(A-D.3)** $N_A = 30$. For each combination of motion model and swarm size, the widths of orientation (z_O)and attraction (z_A) zones are varied. Although the herd remains cohesive in the range of $z_O > 10$, caging with $N_A = 10$ robots is entirely unsuccessful in this range since more agents are required ($N_A = 10 < N_{A,\text{min}}$). On the other hand, deploying $N_A = 20 > N_{A,\text{min}}$ robots guarantees successful caging in every seed. The same results are found for $N_A = 30$, which shows that adding redundant robots has no negative effects on the collective performance. In the range of $z_O \leq 10$, the theoretical lower bound $N_{A,\text{min}}$ is satisfied for all considered N_A, but fragmentation of the herd may occur, and thus caging is not always successful (Except for the long-range model where significantly less fragmentation occurs).

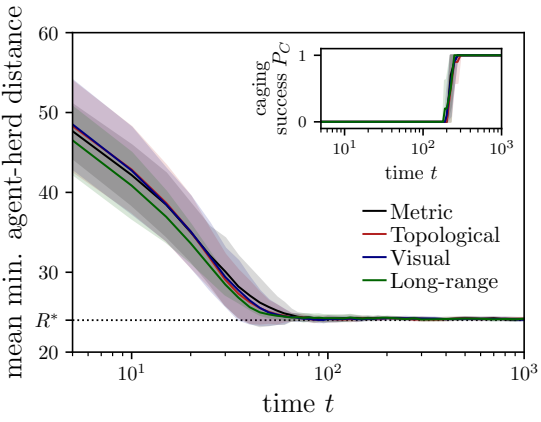

Figure 3.8: The average minimum distance between the robotic agents and the herd, and ratio of success in caging P_C in the inset, over time t from 0 tot 10^3 in log scale with $N_H = 100, N_A = 20, z_O = 50$, and $z_A = 50$. The mean is drawn as a solid line and the standard deviation as shaded area. The desired convergence value R^* of the minimum agent-herd distance is indicated on the y-axis and drawn as a dashed horizontal line. For each plot, measurements of four different models are shown: (black) metric, (red) topological, (blue) visual, and (green) long-range.

color-grids of Figures 3.6A-D.3). However, when a sufficient number of robots is deployed ($N_A = 20 > N_{A,\min}$), caging is successful in every seed with $z_O > 10$ (Figure 3.7A-D.2). Figure 3.7A-D.3 shows that adding a redundant number of robots to the task does not decrease the ratio of success. This is an important quality of the algorithm, as the theoretical minimum number of agents to be deployed is usually unknown a priori. For any $N_A \in \{20, 30\}$, the theoretical lower bound $N_{A,\min}$ is satisfied when fragmentation of the herd does not occur in the range of $z_O \leq 10$. However, the herd is not guaranteed to remain as one cohesive group in this range of z_O, in which case the robots will fail to cage properly. In this experiment, the agents consistently try to remain at a distance of $R^* > z_I$ from the herd, which means that caging does not prohibit the herd from fragmenting. Small adaptations to the proposed caging algorithm could prevent fragmentation, for instance the agents should maintain the maximum relative distance of $2z_I$ between agents instead of following the herd's movement. As an exception, the robot swarm is still successful in this range (see Figure 3.7D.2-3) as fragmentation of the herd is unlikely to occur (Figure 3.6D.1). In all cases, we find that failure in caging the entire herd occurs if and only if the herd is fragmented, or there is an inadequate number of robots deployed based on the theoretical minimum.

In order to observe the time needed for the robot swarm to converge to a stationary caging formation, we measure the average of minimum rela-

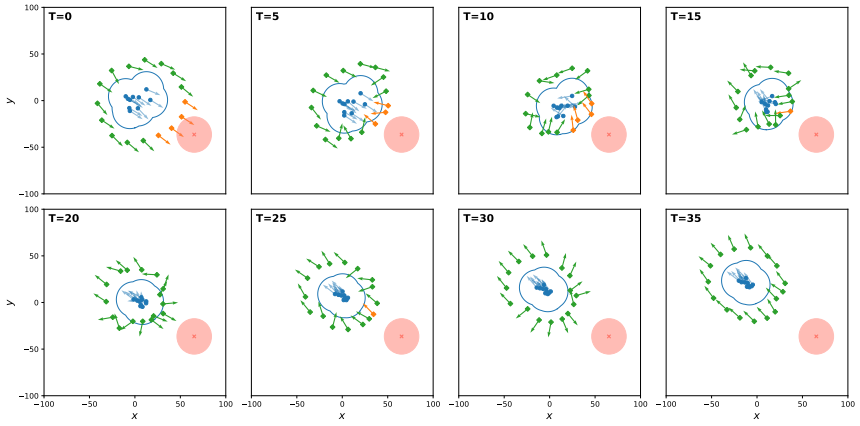

Figure 3.9: Visualization of robotic agents (green or orange diamonds) shepherding the herd (blue dots) away from a dangerous patch (red circle). When a robot has detected a danger, it is colored orange and green otherwise. The orientation of each individual is represented by an arrow. The blue contour is the union of each herd member's circular zone with radius z_I. In other words, only agents that are positioned in the blue contour generate a repulsive force onto the herd. At $t = 0$, agents detect the danger. Next ($t = 5$ to 20), the agents positioned between the herd and the danger now execute the caging algorithm with distance $r^* < z_I$, while the other agents remain at $R^* > z_I$. This causes the herd to re-orientate and move away from the danger. Once the danger is enough far away, all agents occupy a distance of R^* from the herd.

tive distances between each robotic agent and their respective closest individual of the herd $\|p_a(t) - p_{h^*}(t)\|$. As described in subsection 3.4.1, the agents should approach a distance of R^*, which is lower bounded by z_I to account for prediction errors of the herd's movement. Figure 3.8 shows that the robots are able to converge to R^* over time for every collective motion model of the herd. More specifically, the average minimum distance between agents and herd reaches R^* at approximately $t = 10^2$. After this time, the robots spread out around the herd and eventually reach a successful caging formation ($P_C = 1$).

3.5.3 Shepherding

To examine the proposed shepherding algorithm, we simulate two different tasks: (i) dangerous patches, and (ii) safe enclosure as described in subsection 3.3.3. We study the performance of our shepherding algorithm for the four models of collective motion that have been presented in previous sections; metric, topological, visual and long-range. Results in this chapter are shown for a herd size $N_H = 100$.

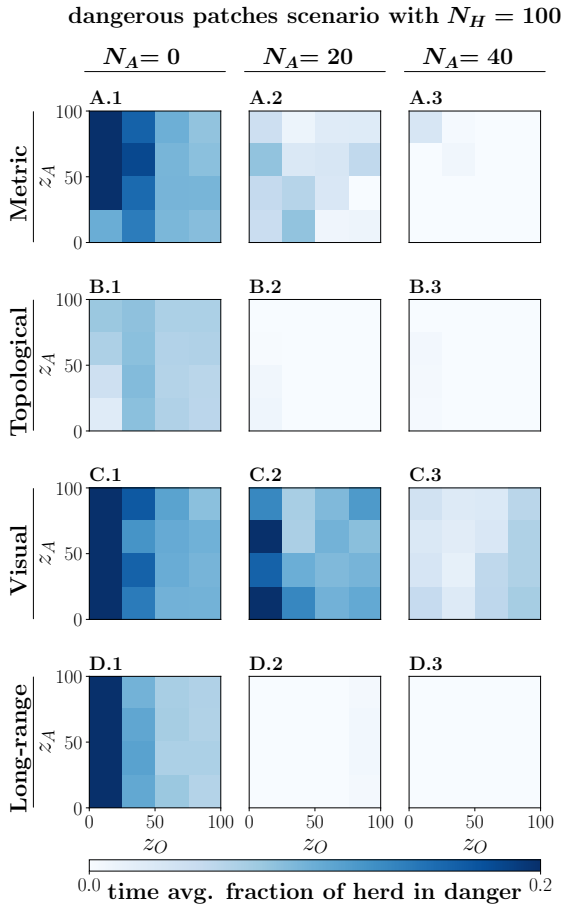

Figure 3.10: Fraction of the herd ($N_H = 100$) located in dangerous patches, averaged over 2000 time steps, and following different models of collective motion: **(A.1-3)** metric, **(B.1-3)** topological with $k = 50$, **(C.1-3)** visual, and **(D.1-3)** long-range with $k = 7, \lambda = 0.1$. The left column **(A-D.1)** $N_A = 0$ shows the results where no robots are present and thus no shepherding takes place. In the other columns, a shepherding robot swarm of sizes **(A-D.2)** $N_A = 20$ and **(A-D.3)** $N_A = 40$ are studied. For each combination of motion model and swarm size, the widths of orientation (z_O)and attraction (z_A) zones are varied. Without shepherding, a certain fraction (non-zero) of the herd is on average in danger for every model and values of z_O and z_A. The topological model results in the lowest endangered fraction, followed by the long-range model, while the metric and visual both perform the worst. When modeled as topological or long-range, deploying 20 shepherding robots successfully ensures the safety of the entire herd. Adding more robots ($N_A = 40$) shows to only potentially improve performance; the metric-modeled herd is successfully shepherded for most values of z_O and z_A, however the visual-modeled remains endangered (although with low probability).

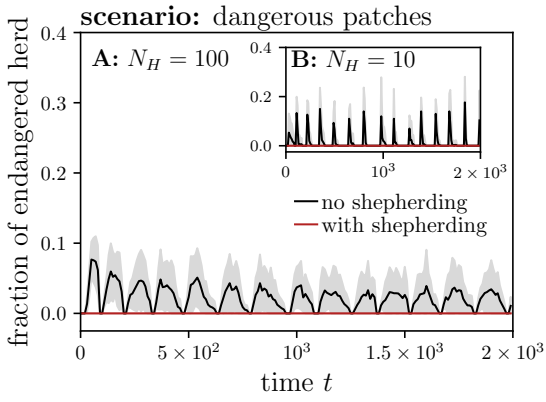

Figure 3.11: Fraction of herd located in dangerous patches over time t, deploying shepherding robots (red) or no robots (black). Results for sizes of the herd and robot swarm respectively: **(A)** $N_H = 100$ with $N_A = 40$ **(B)** $N_H = 10$ with $N_A = 20$. The herd follows a metric model with $z_O = 50$ and $z_A = 50$. The mean is drawn as a solid line and the standard deviation as shaded area.

3.5.3.1 Dangerous patches

Figure 3.9 visualizes the first scenario over time, where a dangerous patch is placed in front of the herd. At $t = 0$, the robots are in a caging formation, remaining at a distance $R^* > z_I$ from the herd. At this time, some robots detect the danger and locally communicate this to their neighbors. The appropriate subset of robots now apply a cage formation at a distance $r^* < z_I$, in order to push the herd in the opposite direction of the danger. In the following time steps, the herd changes movement and other robots also redirect their orientation. The herd and robots move away together from the danger. Once the danger is far enough away (the danger is out of the detection distance d_d at $t = 20$), each robot transitions back to caging at the distance R^*, which is fully established again at $t = 35$.

In order to observe whether the robotic agents are capable of repeatedly shepherding the herd in different directions, dangerous patches are placed in front of the herd after a certain time interval. Each time interval is newly sampled from a Gaussian distribution with mean 150 and standard deviation 25. An individual of the herd only detects the danger once it is located inside the patch, which is similar to animals encountering oil leaks and plastic pollution. When the herd is dense, multiple individuals may detect the danger at the same time. In some way this may be necessary, as the mean direction of the herd is hardly changed by the action of a single individual. Figure 3.10 shows the fraction of the herd that is located outside of the safe enclosure, on average over time. Without shepherding, the herd is endangered for any model of its collective motion. With 20 shepherding

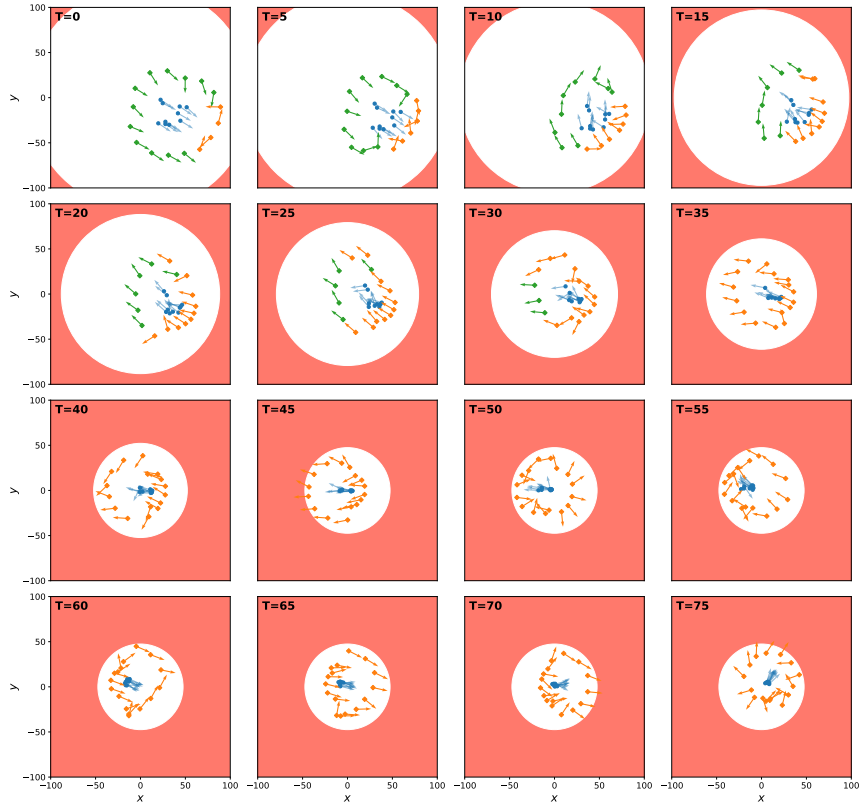

Figure 3.12: Visualization of robotic agents (green or orange diamonds) shepherding the herd (blue dots) to remain in the safe enclosure (white area). When a robot has detected a danger, it is colored orange and green otherwise. The orientation of each individual is represented by an arrow. The blue contour is the union of each herd member's circular zone with radius z_I. In other words, only agents that are positioned in the blue contour generate a repulsive force onto the herd. Agents adapt and repeatedly switch in caging with a distance of $r^* < z_I$ or $R^* > z_I$, depending on how close the herd is to the boundary of the safe enclosure.

robots, the herd stays out of any danger when modeled as topological or long-range for any values of z_O and z_A. For the metric model, the mean fraction is significantly decreased by deploying robots and eventually becomes zero with a robot swarm of size $N_A = 40$ for most values of z_O and z_A. Although the performance of shepherding a herd that follows a visual model improves by adding more robots ($N_A = 20$ to 40), a certain fraction of the herd remains endangered. We argue that the models likely differ in speed of information diffusion, which influences the number of individuals that continues to move towards the danger. For the visual model,

deploying additional robots may not improve information diffusion when those additional robots are positioned in an area occluded by others. We find that this issue only arises for larger herd sizes, as section 3.D shows that for any collective motion model, a herd of size $N_H = 10$ remains safe when shepherded by robots following the proposed algorithm.

Figure 3.11 shows the fraction of the herd that is positioned in a dangerous patch at time t. We see that for both $N_H = 100$ and $N_H = 10$, a fraction of the herd is endangered without the aid of shepherding robots, while the herd remains safe at all time in the presence of robots following our proposed algorithm.

3.5.3.2 Safe enclosure

Figure 3.12 visualizes the second scenario over time, where the safe zone (white) shrinks up until a certain area. At $t = 0$, the robots are in a caging formation, remaining at a distance R^* from the herd. Some robots have detected the danger zone (red) for the first time and begin re-orientating. At $t = 5$, the robots push the herd away from the boundary of the safe zone in a certain direction, until the boundary is reached again ($t = 50$). At this time, the agents redirect the herd towards another direction. This pattern continues as the robots attempt to keep the herd in the center of the safe zone.

The fraction of the herd that is located outside of the safe enclosure, on average over time, is presented in Figure 3.13. Without shepherding, the herd is endangered when modeled as metric, topological and visual. Following the long-range model, however, only a small fraction of the herd is in danger. With 20 shepherding robots, the herd stays out of any danger when modeled as topological or visual for any $z_O > 25$. The long-range herd are continuously shepherded in the safe enclosure, both with 20 and 40 robots, for any values of z_O and z_A. The metric model, however, performs the worst, as a fraction of the herd moves out of the safe enclosure when z_O is large.

Figure 3.14 shows the fraction of the herd that is positioned out of the safe enclosure at time t. We see again that for both $N_H = 100$ and $N_H = 10$, a fraction of the herd is endangered when no robots are present, while the herd remains safe at all time with the aid of shepherding robots. Most notably, there is a peak in the fraction of endangered herdables at the beginning of the scenario. This is most likely due to the herd being slower in detecting and reacting to the rapidly shrinking of the safe enclosure.

3.6 Conclusion

In this chapter, we proposed an algorithm on the individual level, which enables a swarm of robotic agents to shepherd an autonomous herd (i.e. a group following collective motion) away from dangers before the latter succumbs to these dangers. The robots rely on caging the herd at all times,

so that any danger surrounding the herd can be immediately detected. We assume that herd individuals move away from robots that are closer than a certain distance. Based on this assumption, robots are capable of steering the herd in a desired direction by placing themselves close enough to the herd in the appropriate relative position. We studied the performance of the proposed caging algorithm in simulation for four different models of

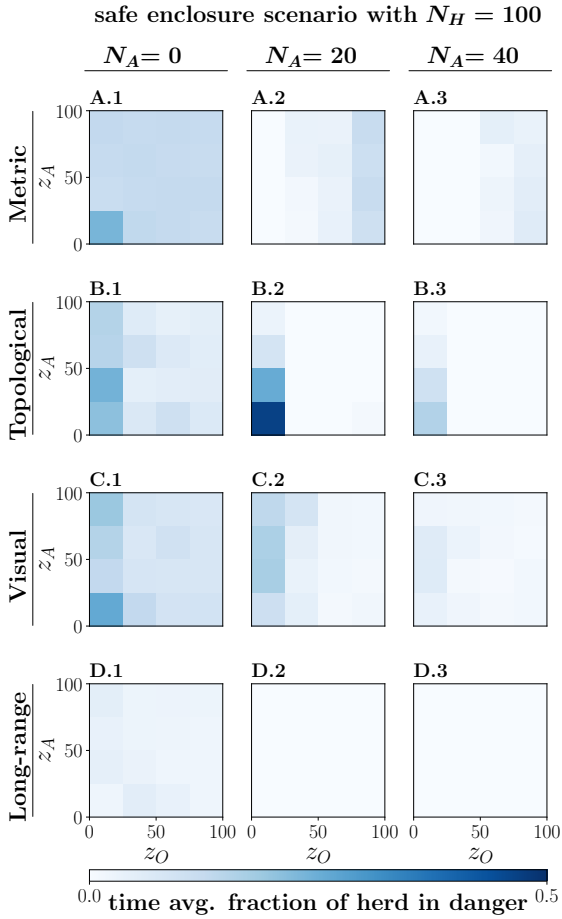

Figure 3.13: Fraction of the herd ($N_H = 100$) located outside of the safe enclosure, averaged over 2000 time steps, and following different models of collective motion: **(A.1-3)** metric, **(B.1-3)** topological with $k = 50$, **(C.1-3)** visual, and **(D.1-3)** long-range with $k = 7, \lambda = 0.1$. The left column **(A-D.1)** $N_A = 0$ shows the results where no robots are present and thus no shepherding takes place. In the other columns, a shepherding robot swarm of sizes **(A-D.2)** $N_A = 20$ and **(A-D.3)** $N_A = 40$ are studied. For each combination of motion model and swarm size, the widths of orientation (z_O)and attraction (z_A) zones are varied.

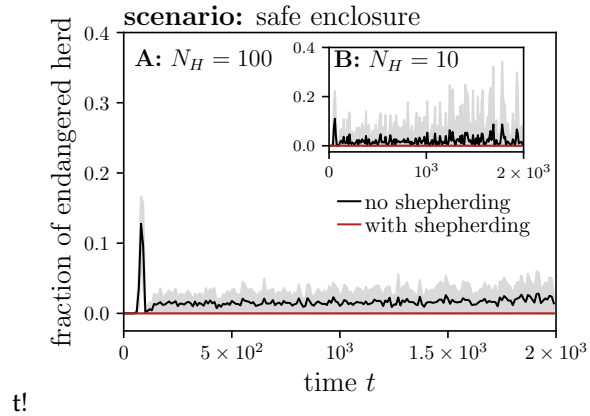

t!

Figure 3.14: Fraction of herd located outside of the safe enclosure over time t, deploying shepherding robots (red) or no robots (black). Results for sizes of the herd and robot swarm respectively: **(A)** $N_H = 100$ with $N_A = 40$ **(B)** $N_H = 10$ with $N_A = 20$. The herd follows a metric model with $z_O = 50$ and $z_A = 50$. The mean is drawn as a solid line and the standard deviation as shaded area.

collective motion to simulate the herd: (i) metric, (ii) topological, (iii) visual, and (iv) long-range. We found that when the herd remains cohesive (i.e. information can be passed between every pair of individuals in the herd), the robots are always capable of caging successfully if an appropriate number of robots is deployed. The required number of robots can be computed from Monte Carlo simulations if the collective motion model of the herd is known. Otherwise, it is best to overestimate the number as the performance of the algorithm does not decrease by deploying redundant robots. For parameters where the herd remains cohesive following a metric model, we examined the performance of robots shepherding by caging. We defined two different shepherding task scenarios; a first where dangerous patches appear with a probability over time, and a second where the herd should remain in safe circular area. Simulation results for herds of sizes 10 and 100 both show that shepherding prevents any individual of the herd to encounter a danger. Without shepherding, a fraction of the herd is endangered. To the best of our knowledge, equipping robots with the proposed algorithm results in the first decentralized multi-robot system only using local observations and communication capable of shepherding a herd by caging. Future work includes simulating the motion of the herd by trained models of real animal trajectories, and eventually real-life experiments.

Bibliography

[1] S. Van Havermaet, P. Simoens, T. Landgraf, and Y. Khaluf. Steering herds away from dangers in dynamic environments. *Royal Society Open Science*, 10(5):230015, 2023.

[2] N. K. Long, K. Sammut, D. Sgarioto, M. Garratt, and H. A. Abbass. A comprehensive review of shepherding as a bio-inspired swarm-robotics guidance approach. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 4(4):523–537, 2020.

[3] B. Bat-Erdene and O.-E. Mandakh. Shepherding algorithm of multi-mobile robot system. In *2017 First IEEE International Conference on Robotic Computing (IRC)*, pages 358–361. IEEE, 2017.

[4] M. Evered, P. Burling, M. Trotter, et al. An investigation of predator response in robotic herding of sheep. *International Proceedings of Chemical, Biological and Environmental Engineering*, 63:49–54, 2014.

[5] Z. Butler, P. Corke, R. Peterson, and D. Rus. From robots to animals: Virtual fences for controlling cattle. *The International Journal of Robotics Research*, 25(5-6):485–508, 2006.

[6] R. Vaughan, N. Sumpter, J. Henderson, A. Frost, and S. Cameron. Experiments in automatic flock control. *Robotics and autonomous systems*, 31(1-2):109–117, 2000.

[7] J.-M. Lien and E. Pratt. Interactive planning for shepherd motion. In *AAAI Spring Symposium: Agents that Learn from Human Teachers*, pages 95–102, 2009.

[8] A. A. Paranjape, S.-J. Chung, K. Kim, and D. H. Shim. Robotic herding of a flock of birds using an unmanned aerial vehicle. *IEEE Transactions on Robotics*, 34(4):901–915, 2018.

[9] S. Gade, A. A. Paranjape, and S.-J. Chung. Herding a flock of birds approaching an airport using an unmanned aerial vehicle. In *AIAA Guidance, Navigation, and Control Conference*, page 1540, 2015.

[10] A. Garrell, A. Sanfeliu, and F. Moreno-Noguer. Discrete time motion model for guiding people in urban areas using multiple robots. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 486–491. IEEE, 2009.

[11] J.-M. Lien, O. B. Bayazit, R. T. Sowell, S. Rodriguez, and N. M. Amato. Shepherding behaviors. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 4, pages 4159–4164. IEEE, 2004.

[12] L. Coppinger, R. Coppinger, et al. Dogs for herding and guarding livestock. *Livestock handling and transport*, 13:235–253, 1993.

[13] A. J. King, A. M. Wilson, S. D. Wilshin, J. Lowe, H. Haddadi, S. Hailes, and A. J. Morton. Selfish-herd behaviour of sheep under threat. *Current Biology*, 22(14):R561–R562, 2012.

[14] W. D. Hamilton. Geometry for the selfish herd. *Journal of theoretical Biology*, 31(2):295–311, 1971.

[15] L. J. Morrell, G. D. Ruxton, and R. James. Spatial positioning in the selfish herd. *Behavioral Ecology*, 22(1):16–22, 2011.

[16] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, et al. 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.

[17] M. Zumaya, H. Larralde, and M. Aldana. Delay in the dispersal of flocks moving in unbounded space using long-range interactions. *Scientific reports*, 8(1):1–9, 2018.

[18] A. E. King and M. S. Turner. Non-local interactions in collective motion. *Royal Society open science*, 8(3):201536, 2021.

[19] S. Hubbard, P. Babak, S. T. Sigurdsson, and K. G. Magnússon. A model of the formation of fish schools and migrations of fish. *Ecological Modelling*, 174(4):359–374, 2004.

[20] Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, and I. D. Couzin. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725, 2011.

[21] D. J. Hoare, I. D. Couzin, J.-G. Godin, and J. Krause. Context-dependent group size choice in fish. *Animal Behaviour*, 67(1):155–164, 2004.

[22] I. Giardina. Collective behavior in animal groups: theoretical models and empirical studies. *HFSP journal*, 2(4):205–219, 2008.

[23] J. Nauta, S. Van Havermaet, P. Simoens, and Y. Khaluf. Enhanced foraging in robot swarms using collective lévy walks. In *ECAI 2020*, pages 171–178. IOS Press, 2020.

[24] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34, 1987.

[25] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.

[26] S. Gueron, S. A. Levin, and D. I. Rubenstein. The dynamics of herds: from individuals to aggregations. *Journal of Theoretical Biology*, 182(1):85–98, 1996.

[27] J. Porzycki, J. Was, L. Hedayatifar, F. Hassanibesheli, and K. Kułakowski. Velocity correlations and spatial dependencies between neighbors in a unidirectional flow of pedestrians. *Physical review E*, 96(2):022307, 2017.

[28] I. D. Couzin, J. Krause, et al. Self-organization and collective behavior in vertebrates. *Advances in the Study of Behavior*, 32(1):10–1016, 2003.

[29] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.

[30] F. Mondada, A. Martinoli, N. Correll, A. Gribovskiy, J. I. Halloy, R. Siegwart, and J.-L. Deneubourg. A general methodology for the control of mixed natural-artificial societies. Technical report, Pan Stanford Publishing Singapore, 2013.

[31] Y. Tsunoda, Y. Sueoka, Y. Sato, and K. Osuka. Analysis of local-camera-based shepherding navigation. *Advanced Robotics*, 32(23):1217–1228, 2018.

[32] F. Berlinger, M. Gauci, and R. Nagpal. Implicit coordination for 3d underwater collective behaviors in a fish-inspired robot swarm. *Science Robotics*, 6(50), 2021.

[33] C. M. Holmlund and M. Hammer. Ecosystem services generated by fish populations. *Ecological economics*, 29(2):253–268, 1999.

[34] B. Le Gallic and A. Cox. An economic analysis of illegal, unreported and unregulated (iuu) fishing: Key drivers and possible solutions. *Marine Policy*, 30(6):689–695, 2006.

[35] L. Jacquin, Q. Petitjean, J. Côte, P. Laffaille, and S. Jean. Effects of pollution on fish behavior, personality, and cognition: some research perspectives. *Frontiers in Ecology and Evolution*, 8:86, 2020.

[36] F. Essl, B. Lenzner, S. Bacher, S. Bailey, C. Capinha, C. Daehler, S. Dullinger, P. Genovesi, C. Hui, P. E. Hulme, et al. Drivers of future alien species impacts: An expert-based assessment. *Global Change Biology*, 26(9):4880–4893, 2020.

[37] S. V. Viscido, J. K. Parrish, and D. Grünbaum. Individual behavior and emergent properties of fish schools: a comparison of observation and theory. *Marine Ecology Progress Series*, 273:239–249, 2004.

[38] A. Huth and C. Wissel. The simulation of fish schools in comparison with experimental data. *Ecological modelling*, 75:135–146, 1994.

[39] J. K. Parrish and L. Edelstein-Keshet. Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science*, 284(5411):99–101, 1999.

[40] E. Hensor, I. D. Couzin, R. James, and J. Krause. Modelling density-dependent fish shoal distributions in the laboratory and field. *Oikos*, 110(2):344–352, 2005.

[41] J. Gautrais, C. Jost, M. Soria, A. Campo, S. Motsch, R. Fournier, S. Blanco, and G. Theraulaz. Analyzing fish movement as a persistent turning walker. *Journal of mathematical biology*, 58(3):429–445, 2009.

[42] T. Landgraf, G. H. Gebhardt, D. Bierbach, P. Romanczuk, L. Musiolek, V. V. Hafner, and J. Krause. Animal-in-the-loop: Using interactive robotic conspecifics to study social behavior in animal groups. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:487–507, 2021.

[43] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot sheepdog project achieves automatic flock control. In *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour*, volume 489, page 493, 1998.

[44] D. Strömbom, R. P. Mann, A. M. Wilson, S. Hailes, A. J. Morton, D. J. Sumpter, and A. J. King. Solving the shepherding problem: heuristics for herding autonomous, interacting agents. *Journal of the royal society interface*, 11(100):20140719, 2014.

[45] T. Miki and T. Nakamura. An effective simple shepherding algorithm suitable for implementation to a multi-mmobile robot system. In *First International Conference on Innovative Computing, Information and Control-Volume I (ICICIC'06)*, volume 3, pages 161–165. IEEE, 2006.

[46] J.-M. Lien, S. Rodriguez, J.-P. Malric, and N. M. Amato. Shepherding behaviors with multiple shepherds. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3402–3407. IEEE, 2005.

[47] A. Pierson and M. Schwager. Bio-inspired non-cooperative multi-robot herding. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1843–1849. IEEE, 2015.

[48] Y. Khaluf, I. Rausch, and P. Simoens. The impact of interaction models on the coherence of collective decision-making: a case study with simulated locusts. In *International Conference on Swarm Intelligence*, pages 252–263. Springer, 2018.

[49] J. Hu, A. E. Turgut, T. Krajnik, B. Lennox, and F. Arvin. Occlusion-based coordination protocol design for autonomous robotic shepherding tasks. *IEEE Transactions on Cognitive and Developmental Systems*, 14(1):126–135, 2020.

[50] A. Özdemir, M. Gauci, and R. Groß. Shepherding with robots that do not compute. In *Artificial Life Conference Proceedings*, pages 332–339. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info . . . , 2017.

[51] W. Lee and D. Kim. Autonomous shepherding behaviors of multiple target steering robots. *Sensors*, 17(12):2729, 2017.

[52] A. Varava, K. Hang, D. Kragic, and F. T. Pokorny. Herding by caging: a topological approach towards guiding moving agents via mobile robots. In *Robotics: Science and Systems*, pages 696–700, 2017.

[53] C. Muro, R. Escobedo, L. Spector, and R. Coppinger. Wolf-pack (canis lupus) hunting strategies emerge from simple rules in computational simulations. *Behavioural processes*, 88(3):192–197, 2011.

[54] E. Tuci, M. H. Alkilabi, and O. Akanyeti. Cooperative object transport in multi-robot systems: A review of the state-of-the-art. *Frontiers in Robotics and AI*, 5:59, 2018.

[55] P. Culbertson and M. Schwager. Decentralized adaptive control for collaborative manipulation. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 278–285. IEEE, 2018.

[56] J. Fink, M. A. Hsieh, and V. Kumar. Multi-robot manipulation via caging in environments with obstacles. In *2008 IEEE International Conference on Robotics and Automation*, pages 1471–1476. IEEE, 2008.

[57] W. Jiang, G. Wen, Z. Peng, T. Huang, and A. Rahmani. Fully distributed formation-containment control of heterogeneous linear multiagent systems. *IEEE Transactions on Automatic Control*, 64(9):3889–3896, 2018.

[58] M. Porfiri, D. G. Roberson, and D. J. Stilwell. Tracking and formation control of multiple autonomous agents: A two-level consensus approach. *Automatica*, 43(8):1318–1328, 2007.

[59] A. Strandburg-Peshkin, C. R. Twomey, N. W. Bode, A. B. Kao, Y. Katz, C. C. Ioannou, S. B. Rosenthal, C. J. Torney, H. S. Wu, S. A. Levin, et al. Visual sensory networks and effective information transfer in animal groups. *Current Biology*, 23(17):R709–R711, 2013.

[60] D. R. Nelson, D. B. Barber, T. W. McLain, and R. W. Beard. Vector field path following for miniature air vehicles. *IEEE Transactions on Robotics*, 23(3):519–529, 2007.

Appendices

3.A Bound on the motion vector in caging to not exceed the maximum linear velocity

The motion vector q of an agent following the proposed caging algorithm described in subsection 3.4.1 is the summation of vectors $\langle \hat{\theta}; \eta \rangle$ and $\langle \theta_{h^*}; \bar{v}_{h^*} \rangle$. We now define an upper bound on η such that $\|q\| \leq v_{\max}$. This allows for the robot to move at maximum velocity, while maximizing its contribution to caging (i.e. by setting η to its upper bound). The value of η indicates how much velocity can be spend on establishing a caging pattern, while replicating the movement of the nearest herd member h^*. Hence, it is required that the maximum linear velocity of the robot $v_{\max}$ is higher than the herd's.

Let $q_{h,x} = \bar{v}_{h^*} \cos(\theta_{h^*})$ and $q_{h,y} = \bar{v}_{h^*} \sin(\theta_{h^*})$. It follows that:

$$q = \langle q_{h,x} + \eta \cos(\hat{\theta}); q_{h,y} + \eta \sin(\hat{\theta}) \rangle$$

Thus, a bound on the velocity is defined as:

$$\begin{aligned} v_{\max}^2 &\geq (q_{h,x} + \eta \cos(\hat{\theta}))^2 + (q_{h,y} + \eta \sin(\hat{\theta}))^2 \\ v_{\max}^2 &\geq \eta^2 + 2\eta(q_{h,x} \cos(\hat{\theta}) + q_{h,y} \sin(\hat{\theta})) + q_{h,x}^2 + q_{h,y}^2 \\ 0 &\geq \eta^2 + 2\eta(q_{h,x} \cos(\hat{\theta}) + q_{h,y} \sin(\hat{\theta})) + (\bar{v}_{h^*}^2 - v_{\max}^2) \end{aligned}$$

We can rewrite $q_{h,x} \cos(\hat{\theta}) + q_{h,y} \sin(\hat{\theta})$ as follows:

$$\begin{aligned} &= \bar{v}_{h^*} \cos(\theta_{h^*}) \cos(\hat{\theta}) + \bar{v}_{h^*} \sin(\theta_{h^*}) \sin(\hat{\theta}) \\ &= \bar{v}_{h^*} \cos(\theta_{h^*} - \hat{\theta}) \end{aligned}$$

Substituting this in the upper bound:

$$\begin{aligned} 0 &\geq \eta^2 + 2\eta \bar{v}_{h^*} \cos(\theta_{h^*} - \hat{\theta}) + (\bar{v}_{h^*}^2 - v_{\max}^2) \\ \eta &\leq -\bar{v}_{h^*} \cos(\theta_{h^*} - \hat{\theta}) \pm \sqrt{\bar{v}_{h^*}^2 \cos^2(\theta_{h^*} - \hat{\theta}) - (\bar{v}_{h^*}^2 - v_{\max}^2)} \\ \eta &\leq -\bar{v}_{h^*} \cos(\theta_{h^*} - \hat{\theta}) \pm \sqrt{\bar{v}_{h^*}^2 (\cos^2(\theta_{h^*} - \hat{\theta}) - 1) + v_{\max}^2} \end{aligned}$$

After root analysis, to ensure that η is positive, we find that:

$$\eta \leq -\bar{v}_{h^*} \cos(\theta_{h^*} - \hat{\theta}) + \sqrt{\bar{v}_{h^*}^2 (\cos^2(\theta_{h^*} - \hat{\theta}) - 1) + v_{\max}^2}$$

3.B Simulation parameters

Table 3.B.1 summarizes all parameters used in the simulation experiments, where topological and long-range are respectively abbreviated as top and

symbol	parameter	value	association
$w_{h,\max}/w_{a,\max}$	max. angular velocity	$\frac{\pi}{3}/\frac{\pi}{2}\ \frac{\text{rad}}{\text{s}}$	herd/robot
$v_{h,\max}/v_{a,\max}$	max. linear velocity	$2/4\ \frac{\text{cm}}{\text{s}}$	herd/robot
σ_h/σ_a	Gaussian orientation noise	0.05/0.05 rad	herd/robot
z_R	radius of repulsion zone	1 cm	herd
z_O	radius of orientation zone	$[10, 100]$ cm	herd
z_A	radius of attraction zone	$[10, 100]$ cm	herd
k	nearest neighbors	50 (top), 7 (l-r)	herd
λ	average of Poisson distribution (l-r)	0.1	herd
z_I	radius of aversive zone	19 cm	herd
α_R	weight of repulsion interaction	100	herd
α_O	weight of repulsion interaction	50	herd
α_A	weight of repulsion interaction	1	herd
α_I	weight of repulsion interaction	$[500, 2500]$	herd
d_d	robot-detection radius	57 cm	robot(cage)
R^*	distance robot-herd	24 cm	robot(cage)
κ	transition rate	0.2	robot(cage)
r^*	distance robot-herd	17 cm	robot (shepherd)
d_c	communication radius	57 cm	robot (shepherd)
d_z	danger-detection radius	40 cm	robot (shepherd)
N_H	number of herd members	$[10, 100]$	herd
N_A	number of robotic agents	$[0, 10, 20, 30, 40]$	robot

Table 3.B.1: Simulation parameters

l-r. The parameters of the herd are inspired by live Trinidadian guppies, who live in shallow waters, which applies to our problem setting of a two-dimensional environment. The robot parameters are derived from two-wheeled motion robots used in the real-world experiments of Chapter 5, in order to simulate physically plausible behavior and ensure consistency with the capabilities of the actual robotic platform.

3.C Quantitative herd measurements of the topological model for different values of k nearest neighbors

Figure 3.C.1 shows that a topological model with $k = 10$ leads to group fragmentation for any widths of the zones. When k is set to 25, the herd remains cohesive for most values of z_O and z_A. Finally, for $k = 50$, the topological model produces most similar results to that of the other three models (see Figure 3.6).

3.D Shepherding a herd of 10 individuals

In Figure 3.D.1, the average fraction of endangered herd over time is shown for all four herd models. Without shepherding, the herd is on average endangered for every model and any values of the zone widths z_O and z_A. Shepherding with 20 robots ensures that the herd remains entirely out

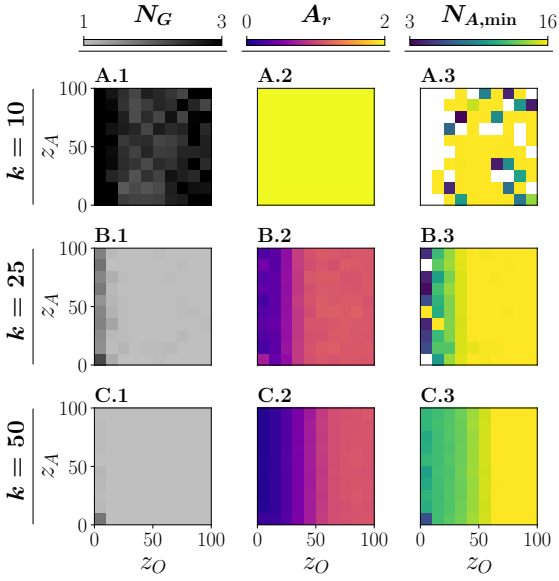

Supplementary Figure 3.C.1: Quantitative herd measurements ($N_H = 100$) obtained in simulation without robotic agents, following the topological model with different values of k nearest neighbors: **(A.1-3)** 10, **(B.1-3)** 25, **(C.1-3)** 50. For each model configuration, the widths of orientation (z_O)and attraction (z_A) zones are varied. The left column **(A-C.1)** shows the number of groups N_G, the middle column **(A-C.2)** shows the relative area coverage A_r, and the right column **(A-C.3)** shows the theoretical minimum number of agents $N_{A,\min}$ needed to form a caging pattern.

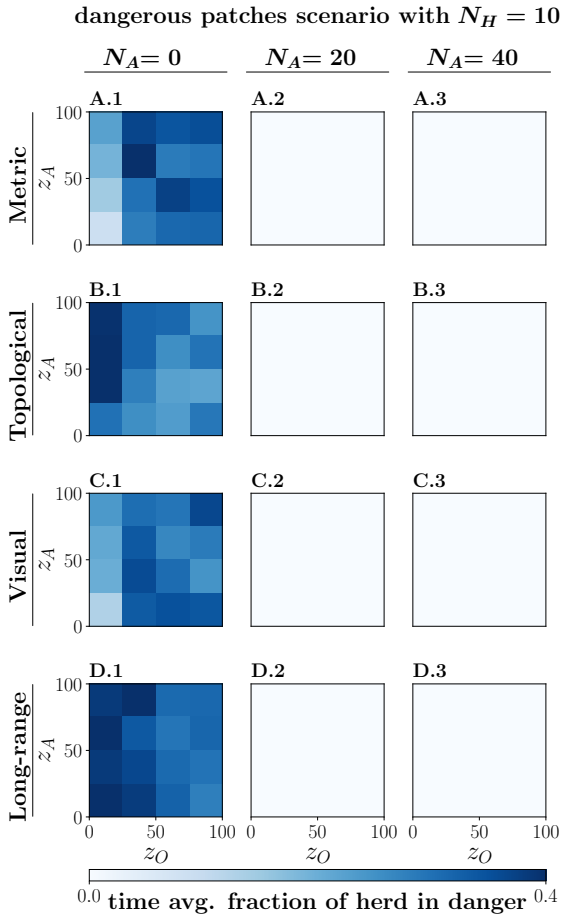

Supplementary Figure 3.D.1: Fraction of the herd ($N_H = 10$) located in dangerous patches, averaged over 2000 time steps, and following different models of collective motion: **(A.1-3)** metric, **(B.1-3)** topological with $k = 50$, **(C.1-3)** visual, and **(D.1-3)** long-range with $k = 7, \lambda = 0.1$. The left column **(A-D.1)** $N_A = 0$ shows the results where no robots are present and thus no shepherding takes place. In the other columns, a shepherding robot swarm of sizes **(A-D.2)** $N_A = 20$ and **(A-D.3)** $N_A = 40$ are studied. For each combination of motion model and swarm size, the widths of orientation (z_O)and attraction (z_A) zones are varied.

of danger, for any model configuration. Deploying redundant robots is shown to not decrease performance, as the herd remains out of danger with $N_A = 40$ shepherding robots.

The work presented in this chapter has been published in:
Scientific Reports (Van Havermaet et al. 2024)

4

Reactive shepherding along a dynamic path

Abstract
In the previous chapter, we focused on protecting fish [1] by guiding them away from dangers without imposing additional trajectory constraints. However, effective conservation also requires directing fish toward safe breeding and feeding grounds. Simply steering fish away from threats is insufficient, as they must reach specific destinations to sustain their populations. To address this, we propose a robotic swarm capable of guiding fish along predefined paths that avoid dangers, obstacles, or undesirable terrain. By confining fish within the boundaries of these paths, we can ensure their safety while directing them to target locations. Although this approach imposes stricter movement constraints, it eliminates the need for consensus decision-making among robots, as required in the method presented in Chapter 3. Therefore, we adapt this method by introducing a position-dependent potential field that encodes path information, allowing robots to respond to shared directional cues without inter-robot communication. Simulation results reveal that, especially in sections with sharp turns and narrow boundary spacing, robots struggle to maintain fish within the boundaries. Sensitivity to these path characteristics is influenced by the robot-to-fish ratio, with larger ratios reducing sensitivity. Additionally, performance is sensitive to fish velocity, as higher speeds make maintaining a caging formation more challenging.

[1] In the main text of this chapter, we use the term "group" to generalize across animal species. However, the parameter settings of the "group" in our study correspond specifically to fish behavior.

4.1 Introduction

In our rapidly changing environment, both anthropogenic and natural threats are placing immense pressure on animal populations, with fish being a poignant example. Illegal fishing [2], habitat destruction [3], and pollution [4] are just a few examples of contemporary challenges. Using robots to guide and protect these vulnerable populations emerges as promising avenue for intervention. Robots could be used to assist in guiding fish away from hazardous zones [5], directing them along secure trajectories, and monitoring environmental parameters to ensure the well-being of the population [6].

The term shepherding was originally coined to describe the process by which dogs guide and direct sheep towards predetermined destinations. The field of robotics has adopted this terminology to characterize tasks involving robotic agents herding various animals, including sheep [7–9], cattle [10], and ducks [11], enabling crowd control [12], averting bird-aircraft collisions [13, 14], or safely evacuating people [15]. Robotic agents, designed with distinctive visual cues or equipped with actuators eliciting aversive responses from animal groups, can employ analogous shepherding mechanisms as naturally perceived threats, akin to dogs herding sheep [16]. The collective behavior displayed by these shepherded animal groups encompasses both aggregation and evasion dynamics in response to potential threats [17].

In this study, we are moving beyond the conventional shepherding task of guiding animals towards predetermined destinations. Our algorithm is designed to guide an animal collective along a secure path, as opposed to merely steering them toward a fixed goal location. The path is demarcated by two boundary lines. This area is constructed by enveloping a series of line segments with a margin. The margin width varies as the path unfolds, simulating more realistic scenarios where the safe zone exhibits variability. We assume that the herded individuals lack the capacity to discern the pathway autonomously, indicating that their motion remains unaffected by the state and structure of the path. This assumption reflects real-world applications where the environment may harbor hidden dangers for the group that only become apparent after damage has occurred to the animals, such as in the case of an oil leak. Alternatively, the area might need to be avoided for ecological or economic reasons, such as preventing the group from grazing on newly seeded fields.

The Robot Sheepdog Project, led by Vaughan et al. [18], was a pioneering initiative in the realm of shepherding problem solutions. They devised a novel approach using force vectors inspired by Reynolds' inter-individual rules for flock behavior [19]. The single robot strategically positioned itself on the opposing side of the flock with respect to the goal and moved towards them. However, it is crucial to note that the project's shepherding method had a limitation, as it confined the ducks to move

along the periphery of an enclosed environment until reaching the goal. Some follow-up works have built upon the aforementioned algorithm to extend its applicability to unconstrained environments, where individuals naturally disperse [20–22]. In the work by Miki and Nakamura [23], a shepherd dynamically alternates between guiding the cohesive group and gathering dispersed individuals. As a consequence, the shepherding agent displays the behavior of real sheepdogs, moving side-to-side behind the group. Their experiments further demonstrated that employing two shepherds is more efficient in guiding the group than having only one. This aligns with other studies suggesting that single-agent solutions are constrained in handling larger group sizes [24]. Therefore, multiple shepherds prove to be more effective in controlling larger groups [25]. However, the majority of shepherding control approaches presume that the shepherds possess global knowledge of the positions of every individual in the environment [26, 27]. Addressing the challenge of guiding a group to a goal with multiple shepherds, utilizing only local information, was first tackled by Lee and Kim [28]. In their approach, shepherds coordinate to aggregate one cohesive group. Some shepherds steer wandering members towards the main group, while others concentrate on maintaining the position of the main group.

Similar to Lee and Kim, the algorithm we propose operates within a distributed fashion, without any central control unit. Our work extends beyond theirs as our robots utilize local observations to collectively guide a group along a safe path, as opposed to guiding it toward a goal location without trajectory constraints. Although the robots make decisions based on local information, they access a shared spatial representation of the path, encoded as a potential vector field. This field reflects directional cues that could, in practice, be derived from environmental stimuli. While the system operates in a decentralized manner without explicit inter-agent coordination, it assumes that each robot can obtain consistent task-related information at its current position and time. Rather than relying on real-time sensory input, the robots respond to this predefined field, which serves as a simplified representation of environmental guidance. This abstraction enables a tractable implementation of local decision-making and provides a foundation for future extensions where guidance would emerge from real-time, local sensing.

Decentralized multi-robot systems offer advantages over centralized ones by eliminating a single point of failure. Moreover, our algorithm is designed to operate without prior knowledge of the environment. Each robot determines its next action based on (i) the positioning of other robots and fish in proximity, and (ii) locally observed pathway information. This approach ensures that the entire group is guided towards the goal location while remaining within the safe path. The proposed algorithm is an extension of our previous work on cage formation [5], and now simultaneously employs a strategy to maintain a cage around the group and nudge the

group in the desired direction along the path.

Steering a group while preserving a caging formation was first tackled by Varava et al. [29]. Their proposed solution relies on a centralized RRT-based (rapidly-exploring random tree) algorithm, incorporating computational topology techniques to validate caging formations. However, their approach assumes access to global information and starts from an already formed caged formation. Our approach is inspired by the behavioral rules of wolf-pack hunting strategies [30], where the hunting wolves strategically position themselves to evenly spread out across the contour of the group. To replicate this behavior, our robots self-organize through the implementation of two simple decentralized rules:: (i) a robot moves toward the group until a specified distance threshold is reached, and (ii) when in close proximity to the group, the robot moves away from other nearby robots. By implementing these results, the robots effectively construct a cage around the animal group. The cage serves the purpose of preventing any potential escapes, and additionally, this strategy allows the robots to gather new information about the surrounding area before the group interacts with it. This information can be crucial for the timely redirection of the group away from undesirable areas. Once the cage has been formed, the robots strategically nudge the group in the desired direction along the path while maintaining the caging formation. Drawing from previous shepherding research, a fish is modeled to move in the opposite direction of a robot when the relative distance is below a certain threshold. Consequently, each robot decides between approaching to push the fish or staying further away to create free space for the fish to move freely. While guiding the group, robots might temporarily move beyond the boundaries of the path, but it is essential for the group to always remain within these boundaries. This temporary deviation can occur due to various reasons, including delays in the robots' response caused by variations in the speed at which they detect the boundaries and initiate maneuvers. Additionally, it can result from encountering narrow or sharp turns in the path.

We investigate the impact of various parameters on the efficacy of our shepherding algorithm. These parameters encompass the characteristics of the path (i.e. the space between the boundaries and the degree of a turn), the size of the system which consists of the number of robots and fish, and the maximum velocities of the robots and fish. Additionally, we investigate the correlation between the success of the caging phase and the overall success of the shepherding process. This multifaceted examination aims to provide insights into the complex dynamics at play and to shed light on the algorithm's performance under different conditions and configurations. We found that the success of the caging is dependent on the magnitude of the fish's maximum velocity. Additionally, we observe that the success of the shepherding process may decline when navigating sharp turns along the path, especially in regions where the spacing between the boundaries narrows. This decline is contingent upon the system size, as

well as the space between the boundaries.

Our findings offer valuable insights that can inform the analysis and refinement of complex shepherding algorithms, especially in scenarios involving hazardous and undesirable areas, across a spectrum of applications. These include biological [31], environmental [32], and robotics applications [26]. Furthermore, our research may find utility in the study of animal and insect collectives, where one species may need to guide, hunt, or gather individuals of another, providing a bridge between the natural systems and innovative technological solutions.

4.2 Related Work

Our approach bears conceptual similarity to formation tracking problems [33], in which agents are controlled to maintain a desired spatial configuration while following a trajectory. Classical methods often use potential functions and gradient-based laws to assign and track predefined trajectories for each agent. In contrast, our robotic agents form a flexible caging formation around an animal group, responding to local interactions and a shared environmental potential field. Rather than maintaining a rigid formation, the agents form a deformable caging structure that adapts to both the shape of the path and the motion of the animal group. The agents guide the group along the path, rather than following it themselves. This shifts the control objective from self-localization to external guidance of other autonomous entities.

As discussed in Chapter 3, a further distinction lies in the nature of the entity being shepherded. Traditional formation- or containment-control frameworks often assume passive or compliant followers that do not exhibit intrinsic or emergent behaviors [34, 35]. In contrast, our agents interact with a collectively moving animal group, governed by inter-individual dynamics (i.e., repulsion, alignment, and attraction). This results in bidirectional influence, where the agents adapt their formation to the group's behavior while simultaneously steering the group within environmental constraints.

4.3 Results and Discussion

In the experiments, N_R robots are tasked with guiding an animal group of size N_G along a safe path, which is defined as the area between two boundary lines. Starting from a series of line segments that evolve monotonically along positive X and Y direction, the path boundaries are constructed by applying an orthogonal margin at both sides of the line segments. This margin width is sampled from a truncated mixture of K Gaussian distribution components, with the minimum margin width M_{min} as a parameter and the maximum derived as $M_{\text{max}} = 3M_{\text{min}}$ (see Methods). An exam-

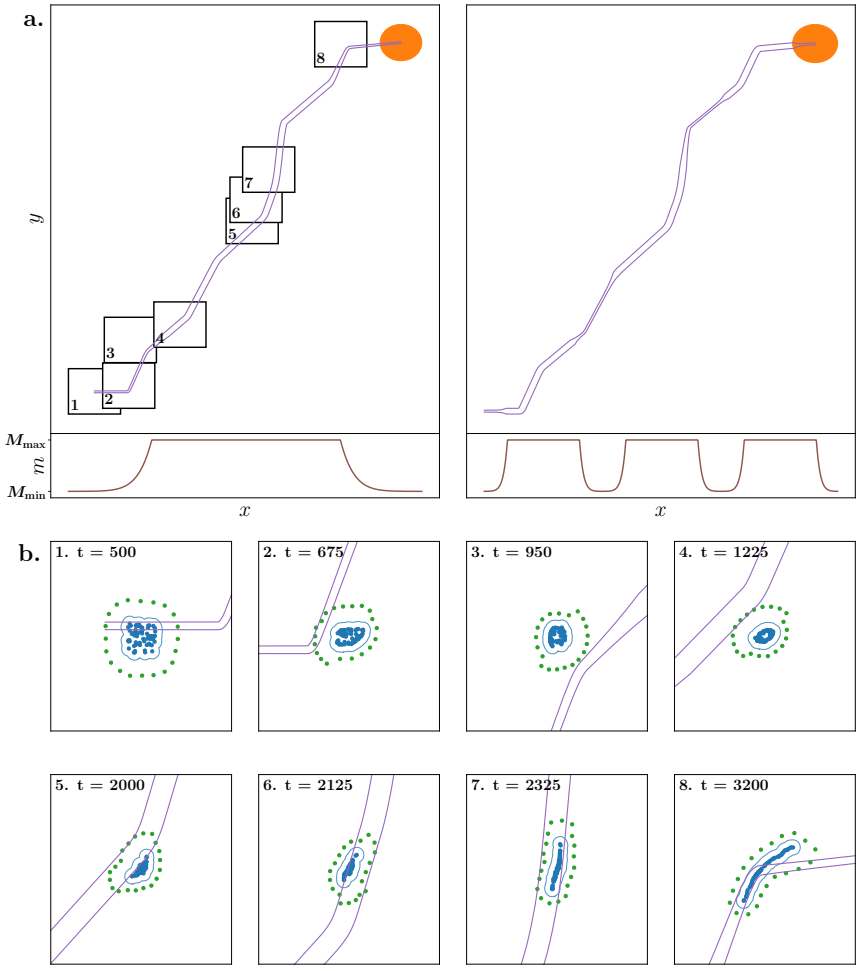

Figure 4.1: (a) Visualization of two different paths (first row), based on the same sequence of linesegments. The respective margin width of each path is shown below (second row). The margin of the left path is generated based on a single Gaussian, while the right margin is generated based on a 3-component Gaussian mixture. The goal location is illustrated by a full orange circle. **(b)** Visualization of the robots (green dots) shepherding the herd (blue dots) along a path with varying margins (purple lines) at different time steps. The location of the time step figures in (a) are respectively shown on the left path. The contour of the herd is marked by a blue line. During the initial time steps ($t < 500$), the herd is positioned at the start of the path and the robots construct the caging formation. Then, the robots try to move the herd inside the margin along the local direction of the path. After some time, the robots are able to maintain the herd inside the margin, until a steep turn occurs at a narrow point of the path ($t = 3200$). Generated using $N_R = 20$, $N_G = 60$, and $M_{\text{min}} = 10$.

ple of a randomly generated path is shown in Figure 4.1a for two different Gaussian mixtures (left $K = 1$, and right $K = 3$).

The purpose of the shepherding algorithm is to advance the animal group along the path while ensuring the group remains within the boundaries of the safe path. Figure 4.1b depicts an experimental instance wherein the robots strive to cage and shepherd the group inside the boundaries at different points in time. When a robot enters a defined radius around an animal, the animal responds by moving away from the robot in the opposite direction. Simultaneously, the animal group exhibits collective behavior, aligning their orientation and maintaining cohesion. Through the utilization of these interaction behaviors, the robots adeptly manipulate the movement of the group. The performance of the shepherding algorithm is defined as the percentage of group animals that is within the safe path, denoted by ρ. The performance of caging is measured by the percentage of the group that is not inside the cage, denoted by ψ. In case the robots fail to construct a closed cage formation, we set $\psi = 100$, i.e. the whole group is outside the cage. The performances ρ and ψ are both measured at any x-coordinate of the path.

We conducted simulation experiments to assess the impact of the considered control parameters on the shepherding and caging performances of the proposed algorithm. The control parameters are different minimum margin intervals of $[M_{\min}, 3M_{\min}]$ with $M_{\min} \in \{10, 20, 50\}$, system sizes which consist of the number of robots and animals $(N_R, N_G) \in \{(20, 30), (20, 60), (40, 60)\}$, and animal maximum velocities $v_G \in \{2, 6, 10\}$ with robot maximum velocities $v_R \in \{\frac{3}{2}v_G, 2v_G, 3v_G\}$. Each experiment with the same configuration is iterated 30 times, each time employing different seeds. A single simulation run ends when either the time limit of 5000 time steps has been surpassed or the group has reached the goal location at the end of the path. In what follows, we study the relationship between the performance and (i) the challenging characteristics of the path, (ii) the system size, and (iii) the robot/animal velocities.

Shepherding performance declines at path narrowing and sharp turns

Figure 4.2 displays the percentage of the group within the safe path (ρ) and the percentage of the group not caged (ψ), as a function of the mean x-coordinate of the group. Each column showcases outcomes for distinct system sizes, while each row varies in the minimum margin width $M_{\min}$ (A: 10, B: 20, C: 50) used to generate the margins. All paths were constructed from the same series of line segments but differ in margin widths on both sides of these segments. A visual representation of each path is provided as an inset in the first column of the corresponding row (A-C.1).

For paths with the smallest considered margin width ($M_{\min} = 10$), varying patterns emerge based on the different system sizes. In the in-

stance of a system with a 20/30 robot-animal composition (A.1), optimal shepherding performance ($\rho = 100$) is achieved; however, this is intermittently interrupted by several decline points in ρ. Observing these points on the path, we note that declines tend to occur where the margin width is minimal or a sharp turn occurs. Despite the sudden decline in ρ at these points, the group remains fully caged ($\psi = 0$), allowing these systems to recover and eventually reach optimal performance again. Increasing the number of animals to a system of size 20/60 (A.2), introduces notable ρ variance for every x-coordinate, with mean ρ never reaching optimum. Notably, this system is the only one among the considered systems where

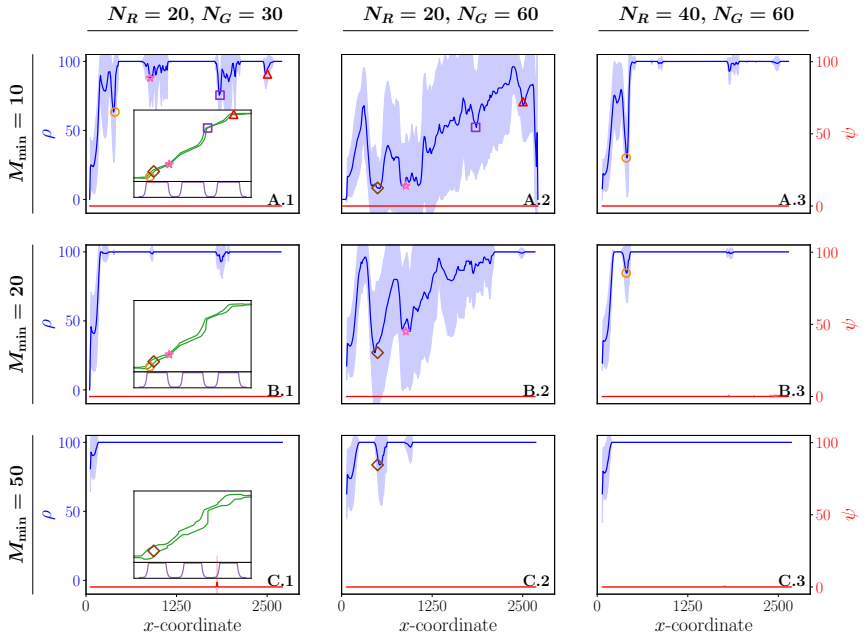

Figure 4.2: The percentage of the group within the safe path (denoted as ρ, colored in blue) and the percentage of the group that is not caged (denoted as ψ, colored in red) in function of the mean x-coordinate of the group. Each row corresponds to simulation results for different values of the minimum margin width $M_{\min}$ of the path: 10 (**A.1-3**), 20 (**B.1-3**), 50 (**C.1-3**). Each column considers different sizes of the system (i.e. the number of robots N_R and the size of the group N_G): $N_R = 20, N_G = 30$ (**A-C.1**), $N_R = 20, N_G = 60$ (**A-C.2**), $N_R = 40, N_G = 60$ (**A-C.3**). The first column shows the corresponding path of its row as an inset. Significant declines in performance are marked by different colors and symbols: a brown diamond, an orange circle, a red triangle, a purple square and a pink star. The location on the path where these declines happen is correspondingly marked in the insets. The mean (solid line) and variance (shaded area) of ρ and ψ results from averaging over 30 stochastically independent simulations for each configuration, with $K = 3$.

not all animals reach the desired goal, as shown in Supplementary Figure S1.1a. This observation is also reflected in Figure 4.2 (A.2), where ρ is not optimal at the maximum x-coordinate. Once again, substantial declines in shepherding performance can be detected, occurring mostly (all except the one indicated by an open brown diamond) at the same points in the path as observed for the 20/30 system size. Doubling the number of robots to a system of 40/60 robot-animals (A.3), the system becomes impervious to the narrow areas in the middle of the path. The system attains a stable optimal state, wherein it consistently maintains its maximum performance ($\rho = 100$ and $\psi = 0$) over a substantial distance along the path. Prior to reaching this stable optimal state, the system appears to necessitate some initial distance to stabilize, notably interrupted in the first turn of the path, as indicated by the open orange circle. Comparing 20/30 (A.1) and 40/60 (A.3) systems with a constant 2/3 ratio reveals distinct performances. The former consistently faces disruptions, while the latter achieves a stable optimal state. This suggests a non-linear relationship between system size and ρ, which we investigate in the following subsection.

Across all system sizes considered, shepherding performance improves when increasing the minimum margin width to $M_{\text{min}} = 20$, resulting in a margin in the interval $[20, 60]$. Specifically, declines in ρ disappear for the 20/30 system size (B.1) and the system reaches a stable optimal state. For a system of size 20/60 (B.2), the number of declines decreases, and the system eventually reaches optimal performance at the end of the path. Although the decline indicated by an open orange circle persists in the system of size 40/60 (B.3), it is substantially weaker compared to the one observed for $M_{\text{min}} = 10$ (A.3). Further increasing M_{min} to 50, all considered system sizes achieve a stable optimal state that persists until the end of the path after an initial stabilization period. Supplementary Figure S2.1 shows that for paths with a longer narrow segment at the start of the path (by generating margins with $K = 1$), in comparison to the paths of Figure 4.2 ($K = 3$), the systems exhibit a significantly later stabilization point. The delay is attributed to the robots facing challenges in aligning the group within the path at the beginning, causing a delay in achieving the desired level of stability.

To explore the impact of path characteristics on declines in ρ, we conducted experiments with 20 robots and 30 animals on five diverse paths, as depicted in Figure 4.3a. These paths feature distinct series of line segments but share the same margins, generated with $M_{\text{min}} = 10$ and $K = 3$. The selection of $K = 3$ allows us to observe multiple instances of narrowing within a single path, while variations in the line segments introduce different degrees of turns. Notably, when margin narrowing occurs along a straight path segment (e.g., open red triangles in B, C, D, and E), the performance decline is relatively small. Similarly, a sharp turn in a wider segment of the margin leads to a minor decline (e.g., the open orange circle in C). The most significant declines occur when a sharp turn happens

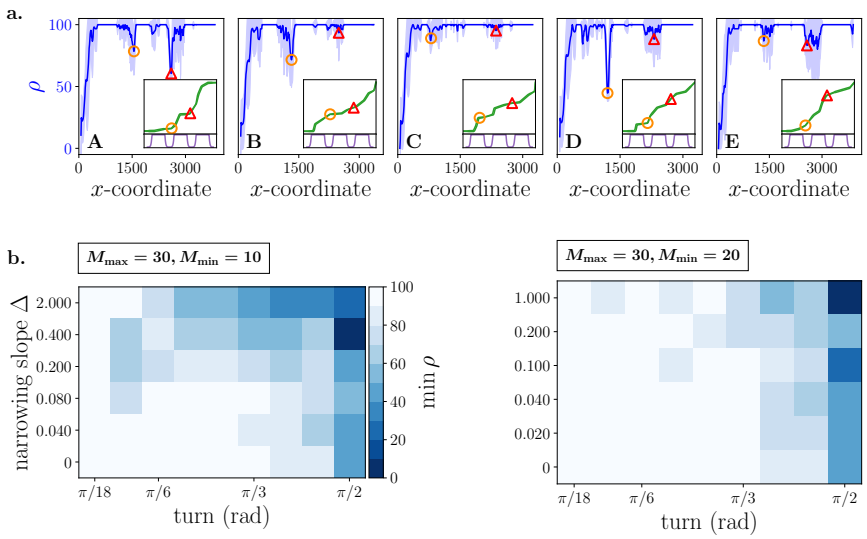

Figure 4.3: (a) The percentage of the group within the safe path (ρ) as a function of the mean x-coordinate of the group across five randomly generated paths. Each path is individually displayed as an inset alongside the corresponding results, featuring a visualization of the margin. Notable declines of ρ are denoted by distinct colors and symbols, specifically an open red triangle and an open orange circle. These dips are located along the path and are similarly marked in the insets. The mean (solid line) and variance (shaded area) of ρ results from averaging over 30 stochastically independent simulations for a configuration with $N_R = 20$, $N_G = 30$, $M_{\text{min}} = 10$ and $K = 3$. **(b)** The minimum percentage of the group inside the safe path ($\min \rho$), measured within the local range of the turn. Each path consists of only one turn, where the margin width decreases linearly from $M_{\text{max}} = 30$ to M_{min} with a (narrowing) slope Δ. The left and right panels respectively show the results for $M_{\text{min}} = 10$ and $M_{\text{min}} = 20$. The values of $\min \rho$ represent the mean over 5 seeds for every parameter configuration of the turn, minimum margin width, and narrowing slope. Other parameters $N_R = 20$ and $N_G = 30$ are constant.

within a narrow path segment (e.g., the open red triangle in A). Moreover, it appears that the greater the turn, the more pronounced the decline (as observed in the comparison of the open orange circles between B and D). To more precisely investigate what constitutes a *sharp* turn and a *narrow* segment, we conducted additional experiments with $N_R = 20$ and $N_G = 30$, involving paths consisting of only one turn. In these experiments, the margin width decreases linearly from $M_{\text{max}} = 30$ to M_{min} with a (narrowing) slope Δ. In Figure 4.3b, the minimum percentage of the group inside the path ($\min \rho$), within the local range of the turn, is depicted for every parameter configuration. A turn of $\frac{\pi}{2}$ radians is evident to cause a substantial decline in ρ for any narrowing slope, even when the margin width does not

decrease ($\Delta = 0$). Notably, when the margin reduces three times in width ($M_{\max} = 30, M_{\min} = 10$), performance experiences a more significant decline with a faster reduction in margin width (i.e., an increase in the slope Δ) and a larger turn. Conversely, when the margin only reduces 1.5 times in width ($M_{\max} = 30, M_{\min} = 20$), performance is more influenced by the turn than the narrowing slope.

The system size plays a crucial role in reaching and maintaining a stable optimal state

To explore in more detail the relationship between the size of the system (N_R, N_G) and the percentage of the group inside the safe path (ρ), we define τ as the continuous percentage of the total path length measured on the x-axis. Figure 4.4 presents the probabilities $\Pr(\rho \geq \rho_{\min} \cap \tau \geq \tau_{\min})$ across various values of $\rho_{\min} \in [5, 10, ..., 100]$ and $\tau_{\min} \in [5, 10, ..., 100]$, considering the same configurations of system sizes and minimum margin widths as previously explored.

We observe that there is no assurance that the robots can maintain the entire group within the safe path for the entire path length, as $\Pr(\rho = 100 \cap \tau = 100) < 1$ for each configuration. Interestingly, only configuration (C.1) exhibits a non-zero probability. This observation aligns with the findings in Figure 4.2, where, even with the largest margin widths ($M_{\min} = 50$), the system requires some initial distance to reach a stable optimal state and thus $\rho < 100$ at the beginning of the path. We identify a dual factor contributing to this initial stabilization requirement. Firstly, the group must adapt its initial shape to fit the margin width. As illustrated in Figure 4.1b, at $t = 500$ (1), the contour around the group exhibits a relatively large circular shape. To ensure the entire group remains inside the safe path, the robots must subsequently coerce the group into a more rectangular shape, achieving this feat at time step $t = 2325$ (7). Secondly, the robots are required to guide the group to align with the current direction of the path. When a part of the group is influenced to alter its direction, this information must propagate across the entire group, a process that does not occur instantaneously. Figure 4.4 demonstrates that the majority of configurations complete the stabilization process since they maintain a stable optimal state for most of the path. However, systems with 20 robots and 60 animals encounter challenges in achieving a stable state, particularly for minimum margin widths of 10 (A.2) and 20 (B.2). Even though successfully caging 60 animals with only 20 robots is achievable, the system faces difficulties in consistently keeping the entire group within the safe path.

We find this to be a consequence of the limited time available to the robots for executing precise maneuvers required to redirect the large group of 60 animals. The detection of the path orientation change is a local process, and by the time the robots react, the narrow margin of 10 or 20 has

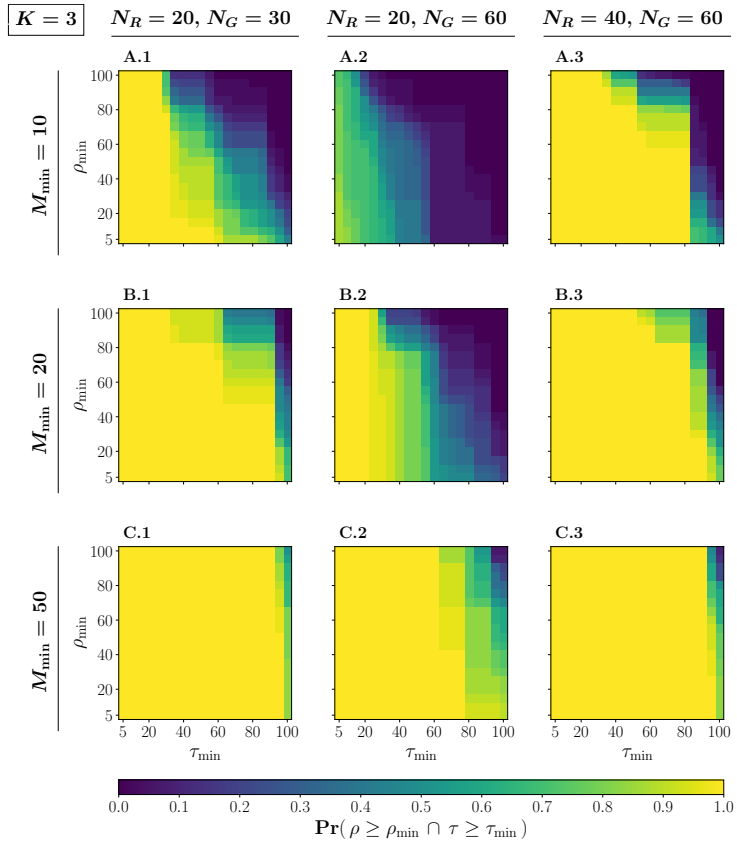

Figure 4.4: Probabilities $\Pr(\rho \geq \rho_{\min} \cap \tau \geq \tau_{\min})$ of the system maintaining at least a specified percentage of the group inside the safe path ($\rho_{\min}$) for at least a specified percentage of the total path length, measured on the x-axis, continuously ($\tau_{\min}$). Each row corresponds to simulation results for different values of the minimum margin width $M_{\min}$ of the path: 10 (**A.1-3**), 20 (**B.1-3**), 50 (**C.1-3**). Each column considers different sizes of the system (i.e. the number of robots N_R and the size of the group N_G): $N_R = 20, N_G = 30$ (**A-C.1**), $N_R = 20, N_G = 60$ (**A-C.2**), $N_R = 40, N_G = 60$ (**A-C.3**). Probabilities are obtained by averaging over 30 stochastically independent simulations with $K = 3$.

already been breached. Consequently, a sequence of transitions from one side of the path to the other occurs, which we also refer to as crossing behavior. Figure 4.5a visually depicts this behavior, highlighting the intricate interplay of the aforementioned dual factors. Over time, as the spatial arrangement of the group gradually conforms to the shape of the path, the robots can more rapidly realign the group towards the path, eventually achieving stability within the path boundaries. We systematically explore

nuanced aspects through a more detailed analysis (see Methods for the mathematical definitions of the following metrics). Group-path alignment is quantified as the average alignment between each animal's direction and the desired direction along or towards the path, with a measurement of 1 indicating full alignment. Additionally, we measure group compression to assess how compact the spatial arrangement is. A circular formation corresponds to compression equal to 1, while a stretched-out group (adhering to the path geometry) results in compression as zero. Figure 4.5b illustrates how the group undergoes rapid oscillations in aligning with the current path direction. The severity of these oscillations gradually mitigates over

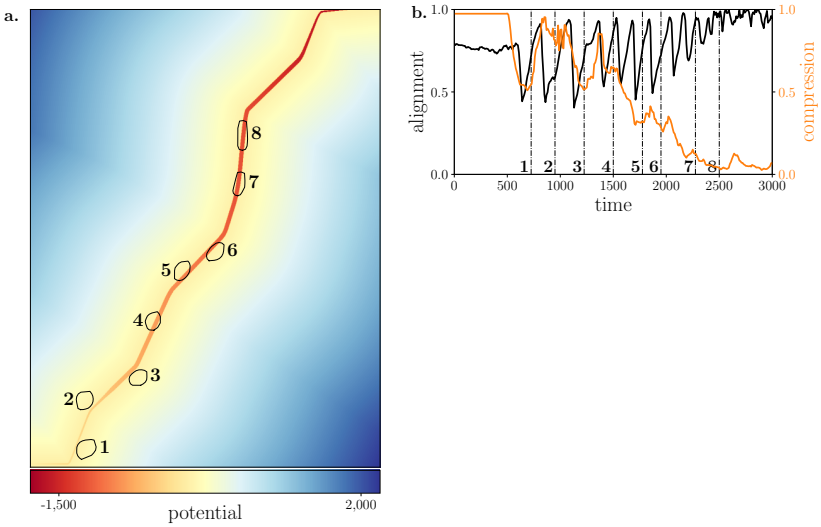

Figure 4.5: (a) The positions of the robots and the group are drawn by the convex hull of all individuals (depicted by the black contour) at different time points (chronologically labeled). The movement illustrates the crossing behavior, where the group sequentially transitions from one side of the path to the other. The colormap illustrates the artificial potential field of the environment, with the global minimum located at the end of the path. Outside the path, the potential diminishes toward the boundaries. Within the path boundaries, the potential decreases along the x-axis. These positions were extracted from an experiment with configuration parameters $N_R = 20$, $N_G = 60$, $M_{\text{min}} = 10$ and $K = 1$. **(b)** The group-path alignment (black) as the average alignment between each animal's direction and the desired direction along or towards the path, with a measurement of 1 indicating full alignment. The group compression (orange) measures how compact the spatial arrangement is. A circular formation corresponds to compression equal to 1, while a stretched-out group (adhering to the path geometry) results in compression as 0. The analysis was conducted on the same experiment as presented in (a). Both metrics are plotted over time, with different time points (annotated at the corresponding vertical lines) aligning with those in (a).

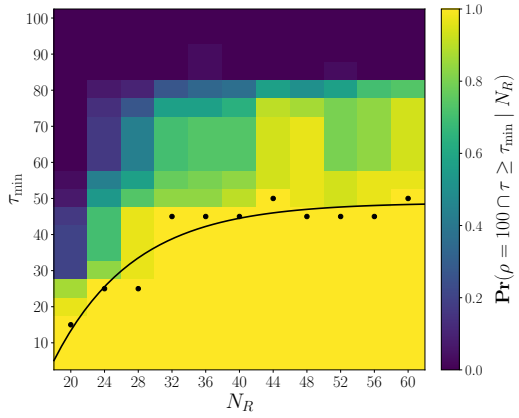

Figure 4.6: Probabilities $\Pr(\rho = 100 \cap \tau \geq \tau_{\min} \mid N_R)$ of the system maintaining the entire group inside the safe path ($\rho = 100$) for a percentage of the total path length measured on the x-axis ($\tau_{\min}$) in function of the number of robots (N_R). Solid black line serves as a guide to the eye, where the probability is greater than 0.975 for the highest value of $\tau_{\min}$, i.e. $\tau'_{\min}$. The curve results from non-linear regression to the exponential function $\tau'_{\min} = a \exp(-b \cdot N_R) + c$ with $a = -289.5, b = 0.1050$, $c = 48.86$ (residual standard error is 4.488). Probabilities are obtained by averaging over 30 stochastically independent simulations for each value of N_R, with $N_G = 60$, $M_{\min} = 20$, $K = 3$.

time, until complete alignment with the path is achieved. Simultaneously, as the alignment process undergoes damping, the spatial arrangement of the group eventually conforms to the shape of the path, leading to a decrease in compression.

While Figure 4.5 illustrates an experiment instance where the system of (B.2) eventually attains a stable state after the crossing behavior, this outcome is not consistent across all instances. Moreover, the duration of the crossing behavior varies inconsistently among systems with configurations of (A.2) or (B.2). If we revisit Figure 4.4 (B.2), it becomes apparent that for a given value of $\tau_{\min}$ (e.g., 50), the probability diminishes as $\rho_{\min}$ increases but consistently remains above zero. This suggests that, in some experiments, the entire group remains within the safe path, while in others, only 20% of the group remains within the path for half the length of the path. This phenomenon accounts for the high variance observed in configurations (A.2) and (B.2) in Figure 4.2. In comparing configurations (A-B.2) with (A-B.3), it is evident that deploying additional robots enhances the reliability of reaching and maintaining a stable state. However, upon comparing the outcomes between the 20/30 and 40/60 systems, the relationship between the system size and the duration of the stable state does not appear linear.

Figure 4.6 offers a comprehensive analysis of the influence of deploying

additional robots on the effectiveness of the proposed algorithm. Specifically, it presents the probabilities of a system with $N_G = 60$ maintaining a stable optimal state (with $\rho = 100$) within a safe path with $M_{\min} = 20$ covering continuously at least $\tau_{\min}$ percent of the path length. Additionally, we approximate the relationship between the maximum length of the stable optimal state ($\tau'_{\min}$) and the number of robots (N_R), for which the probability exceeds 0.975, using a non-linear regression to the exponential function $\tau'_{\min} = a\exp(-b \cdot N_R) + c$ with $a = -289.5, b = 0.1050$, and $c = 48.86$. The residual standard error is 4.488 (on 8 degrees of freedom) after optimization over 6 iterations (with achieved convergence tolerance $8.502 \cdot 10^{-6}$). The results indicate that increasing the number of deployed robots exponentially improves the length of the stable state, up to a certain threshold ($\tau_{\min} = 50$). However, the system fails to reach a stable state for the full length of the path, suggesting that deploying redundant robots ($N_R > 32$) does not alleviate the initial stabilization process or mitigate the decline in ρ at the first sharp turn (see Figure 4.2 B.3). Thus, the deployment of redundant robots neither enhances nor diminishes the length of the stable optimal state. In addition, Supplementary Figure S4.1 shows that the system demonstrates fault tolerance when redundant robots are deployed.

Shepherding success relies on successful caging, which is dependent on the robot-animal maximum velocities

One of the crucial parameters in this study is the maximum velocities of both the animals and the robots. In the presented algorithm, the robot must maintain its desired distance from its closest animal while remaining equidistant from its two neighboring robots. The robot predicts the movement of the fish but might have to correct prediction errors. Additionally, the robot might need to navigate from the back to the front of the group to establish and maintain the caging formation, especially in the case of malfunctioning robots. These scenarios require the robot to be able to move at a higher velocity (v_R) than the fish (v_G) in our system. In Figure 4.7, the relationship between the percentage of the group inside the path (ρ) and the percentage that is not caged (ψ) is depicted as a function of the x-coordinate. This analysis considers varying maximum velocities for the animal, denoted as $v_G \in [2, 6, 10]$, and different robot-animal velocity ratios, expressed as $\frac{v_R}{v_G} \in [\frac{3}{2}, 2, 3]$. The results exhibit a notable similarity across all robot-animal velocity ratios, suggesting that further increases in the maximum robot velocity beyond $\frac{3}{2}v_G$ do not yield significant improvements in maintaining the group within the path.

When the animals move at a maximum velocity of 2 (A-C.1), successful caging is achieved throughout the entire trajectory, leading to effective maintenance of the group within the path. The pronounced declines in ρ observed in these systems can be attributed to the specific characteristics of

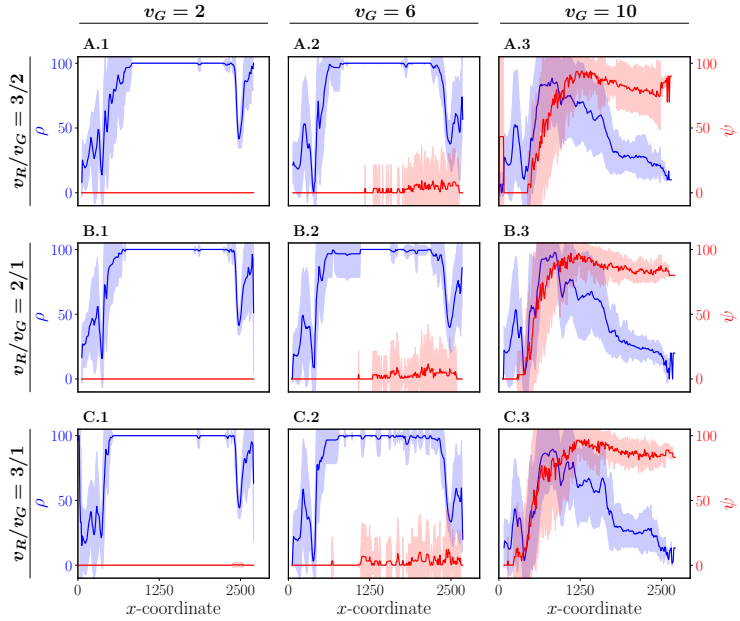

Figure 4.7: The percentage of the group within the safe path (denoted as ρ, colored in blue) and the percentage of the group that is not caged (denoted as ψ, colored in red) in function of the mean x-coordinate of the group. Each column considers different maximum velocities of the animals (denoted as v_G): $v_G = 2$ (**A-C.1**), $v_G = 6$ (**A-C.2**), $v_G = 10$ (**A-C.3**). Each row corresponds to results for different robot-animal maximum velocity ratios $\frac{v_R}{v_G} = \frac{3}{2}$ (**A.1-3**), $\frac{v_R}{v_G} = \frac{2}{1}$ (**B.1-3**), $\frac{v_R}{v_G} = \frac{3}{1}$ (**C.1-3**). The mean (solid line) and variance (shaded area) results from averaging over 30 stochastically independent simulations for each configuration, with $N_R = 20$, $N_G = 30$, $M_{\text{min}} = 10$, and $K = 1$.

the path, as illustrated in Supplementary Figure S2.1 (A.1), and previously discussed in detail. As the group is able to move at a higher velocity of $v_G = 6$ (A-C.2), the caging formation becomes less consistently optimal after the group has traveled halfway along the path. Nevertheless, ρ remains similar to the results obtained at $v_G = 2$. However, when the maximum animal velocity is significantly increased to $v_G = 10$ (A-C.3), ψ starts to increase rapidly from the early stages of the path, reaching approximately 90%. This indicates that the vast majority of the group is no longer caged. Importantly, as ψ increases, there is a corresponding decrease in ρ, signifying that the effectiveness of the caging formation is crucial to the algorithm. When this formation fails, it allows the group to venture beyond the path boundaries. Consequently, only a small percentage of the group reaches the desired goal as shown in Supplementary Figure S1.1b, either due to random factors or because a minority of the group was steered by an in-

complete caging formation. Increasing the maximum velocity of the robots does not enhance the success of caging, suggesting that other parameters may need adjustment. In the algorithm, the robots utilize two different distance thresholds to position themselves relative to a nearby animal, which are parameters of the proposed algorithm. In this paper, these thresholds are computed based on the maximum velocities of the animals, thus influencing the success of caging.

4.4 Conclusion

Our findings offer insights into the complex dynamics of robotic agents controlling the movement of an autonomous group. We introduce a novel collective shepherding algorithm that enables a team of robots to self-organize around a group of animals (i.e. caging) and guide the group along a predefined path, contributing to safeguarding them from anthropogenic threats, including pollution and illegal exploitation. Our results have far-reaching implications for the conservation of fish and other populations, aligning with broader environmental and ecological benefits on a global scale. In the proposed algorithm, robots operate with limited knowledge about the path, relying solely on local perception and possibly communication with their direct neighbors. Our shepherding algorithm leverages the principle of robot-animal repulsion as the primary method of interaction to direct the group. As a result, the robots must collaboratively determine the distribution of forces at each of their respective positions within the cage to achieve successful steering. It is worth highlighting that our algorithm is a pioneering example of a collective shepherding approach that relies exclusively on local information to guide animals using robots. This innovative approach holds great promise for advancing the field of robotic shepherding control.

Shepherding is successfully achieved through the robots forming a caging formation around the group, and the application of repulsive forces to steer the group along the safe path, which is defined as the area between two boundary lines. The objective is for the animal group to reach the end of the path while staying within the boundaries. The effectiveness of our shepherding algorithm is profoundly influenced by a set of critical parameters. These parameters include the structural characteristics of the path, the number of robots and animals, and their maximum velocities. We observed that the robots face the greatest challenges in maintaining the group within the safe path encountering a sharp turn, especially in instances where the path narrows. Nevertheless, we noted that the severity of these disruptions tends to diminish as the space between the boundaries increases. Furthermore, hybrid systems with a $2/3$ robot-animal size ratio demonstrate significant performance advantages over a $1/3$ ratio. The results illustrate that deploying additional robots leads to an exponential improvement in performance. Finally, we found that when the animals

move at relatively high velocity, robots do not achieve successful caging and consequently are unable to maintain the group within the safe path.

While we have demonstrated the proposed algorithm to work in simulations, there are limitations to applying it in real-life settings. First and foremost, the robots need equipment to sense the direction of the path. If only specific robots can sense this direction (e.g., when the path is illuminated, and only the front of the robot swarm can capture this information), then a communication system is also required to share this information. Additionally, the robots require sensors to measure distances and angles from nearby animals and other robots. Lastly, the robots should be capable of moving at least as fast as the maximum velocity of the fish. Apart from following the group, they need to reposition themselves from the back to the front of the group if necessary.

4.5 Methods

4.5.1 Simulation kinematics

The simulations were implemented using the Python programming language. Let $\mathcal{R}$ and $\mathcal{G}$ denote the respective sets of the robot swarm and the animal group. The state of the discrete-time system at time t is defined by the position $p_i(t) \in \mathbb{R}^2$ and orientation $\theta_i(t) \in [0, 2\pi)$ of each individual $i \in \mathcal{R} \cup \mathcal{G}$. At every time step, the group and robots compute their new vector of motion q_i respectively based on the model of collective motion or the proposed shepherding algorithm, as detailed in the following two sections. To consider physical constraints of body mass, the individual changes the motion vector with the highest priority to the average direction that is opposite to the position of all other individuals within a radius of 1 cm. To avoid collisions, the length of a motion vector is capped when it would cause the individual to be near any other individual of 1 cm distance. After computing its new motion vector, an individual first rotates with angular velocity w to its desired orientation $\hat{\theta}_i = \arctan(q_i)$ that is computed based on local information. The individual stops rotating once $\arctan(q_i)$ is reached, or when the time interval has passed. The individual then moves straight forward at linear velocity v, until the remainder of the time interval has passed or the individual has moved for a distance of $\|q_i\|$. At every time step, Gaussian noise σ is added to the orientation of each individual.

4.5.2 Collective animal motion model

We simulate the collective motion of the group as proposed by Couzin et al. [36]. Each animal i updates its direction of motion based on the relative positions of neighbors within a radius z. The neighbors are divided into three distinct subsets based on their positioning in one of the concentric zones surrounding i. Each zone corresponds to a different type of

interaction: (i) repulsion from others inside the disk with radius z_R, to establish a minimum inter-individual distance, (ii) alignment of orientation with others inside the annulus with width z_O, to all move approximately in the same direction, and (iii) attraction to others inside the annulus with width z_A, to remain as one cohesive group. Following the approach of Van Havermaet et al. [5], we extend this model by adding a zone to model the interaction between the animals and the robots : (iv) aversion from robots inside the disk with radius z_I, to move away from unknown mobile objects (e.g. robots) that could be predators.

Let $\mathcal{G}_i^R$, $\mathcal{G}_i^O$ and $\mathcal{G}_i^A$ denote the neighbor subsets by separating other animals of the group based on the repulsion, orientation, and attraction zones respectively. Thus, each neighbor is only assigned to one of the subsets; $\mathcal{G}_i^R \cap \mathcal{G}_i^O \cap \mathcal{G}_i^A = \emptyset$. Let $\mathcal{R}_i$ denote the subset of robots located in a radius of z_I around p_i. Furthermore, the relative position of individual j from i is defined as $r_{ij}(t) = p_j(t) - p_i(t)$. The motion vector q_i of a animal i is then computed as follows:

$$q_i(t) = -\alpha_R \sum_{j \in \mathcal{G}_i^R} \frac{r_{ij}}{\|r_{ij}\|} + \alpha_O \sum_{j \in \mathcal{G}_i^O} \frac{q_j}{\|q_j\|} + \alpha_A \sum_{j \in \mathcal{G}_i^A} \frac{r_{ij}}{\|r_{ij}\|} - \alpha_I \sum_{j \in \mathcal{R}_i} \frac{r_{ij}}{\|r_{ij}\|}, \tag{4.1}$$

with weights $\alpha_R \geq 0$, $\alpha_O \geq 0$, $\alpha_A \geq 0$ and $\alpha_I \geq 0$ of repulsion, orientation, attraction, and aversive animal-robot interactions respectively.

4.5.3 Safe path generation

A safe path is defined as the area between two boundaries, formed by adding equal upper and lower margins to a sequence of line segments. The sequence of segments is generated based on a given total length $L = 5000$ and normal distribution of segment length $\mathcal{N}(500, 100)$. The subsequent angles between two segments are sampled from a uniform distribution between $[0, \frac{\pi}{2}]$. As such, the path always progresses in the positive direction of the x-axis, enabling the assessment of whether the group follows the intended path direction. From the segments, we generate a discrete sequence of points $P = (X, Y)$ representing the curve at the middle of the path. The margin width $m(x)$ is generated by a clipped mixture of K Gaussian distribution components. The minimum margin width M_{min} is a parameter in the experiments, while the maximum is 3 times the minimum value. The real top and bottom boundaries are defined by sequences of points $M^{t,r}$ and $M^{b,r}$ respectively. To ensure that the herd remains between the real boundaries, even under the presence of noise and the uncertainty of the progression of the path, we generate safety margin widths that are smaller than the real margin widths. Let $M^{t,s}$ and $M^{b,s}$ be the respective sequences of points representing the safety top and bottom margin boundaries. For each point i in P, we compute the approximate angle between the curve and the x-axis as $\alpha = \arctan\left(\frac{Y_{i+1} - Yi}{X_{i+1} - Xi}\right)$. The top and bottom boundaries

are then respectively given by the following equations:

$$M_i^t = \left(X_i + m' \cos\left(\alpha + \frac{\pi}{2}\right), Y_i + m' \sin\left(\alpha + \frac{\pi}{2}\right)\right)$$
$$M_i^b = \left(X_i + m' \cos\left(\alpha - \frac{\pi}{2}\right), Y_i + m' \sin\left(\alpha - \frac{\pi}{2}\right)\right)$$

where $m' = m(X_i)$ for $M^{t,r}$ and $M^{b,r}$, and $m' = 0.6\, m(X_i)$ for $M^{t,s}$ and $M^{b,s}$. The safety margin width is thus 60% of the real margin width. Note that the M sequences may produce loops so that the curves represented by the sequences have multiple y-values for x-values. We remove these loops in a post-processing by assigning only the minimum or maximum y-values for x-values of the top or bottom boundaries respectively.

4.5.4 Robotic shepherding algorithm

The robots are tasked to steer the group, while maintaining a caging formation. The number of robots N_R required to construct a caging formation depends on the size of the group N_G (see Supplementary for a derivation of the lower bound on N_R given N_G).

Assumption 4.1. *The upper bound on the size of the group, N_G, is known a priori.*

A caging formation can be constructed by the robots based on the repulsive animal-robot threshold z_I. When two robots are at a distance lower than $2z_I$ of each other, they exert a combined repulsive force on the group which prevents them from intersecting the path between those robots. In our past work [5], we proposed a decentralized algorithm for the robots to steer the group away from dynamic dangers while maintaining a caging formation. In what follows, we provide a brief description of the algorithm.

In case the robot is unable to detect any animal, the robots follow a search strategy. In this paper, we use an adaptation of the aforementioned collective animal motion model. Additionally, robots communicate their search status (i.e. if they have detected an animal) to others. Consequently, robots who are unable to directly observe any animal, will become attracted to the neighboring robot senders.

When part of the group is detected, the robot attempts to reside at a given distance of R^* from the closest animal. Simultaneously, the robot moves to position itself equidistant from its two consecutive neighboring robots. During this movement, the robot may find a closer animal and will then orbit around this new animal. In other words, the robot moves along the circular perimeter of the group. To allow continuous motion along the perimeter, we hold the following assumption:

Assumption 4.2. *The union of the circles, where the center points are the positions of every animal and the radii are R^*, is a connected set.*

The robot i computes the desired orientation $\hat{\theta}_i$, to approach and rotate along a circular path of radius R^*, as proposed by Nelson et al. [37]:

$$\hat{\theta}_i = \gamma_i + \phi_i \left(\frac{\pi}{2} + \arctan\left(\kappa(d_i - r_i)\right) \right) \tag{4.2}$$

where γ_i is the angular position of robot i to the closest animal, $\phi_i \in \{-1, 1\}$ determines the direction; i.e. clockwise ($\phi_i = -1$) or counterclockwise ($\phi_i = 1$), $\kappa > 0$ influences the transition rate, d_i is the distance from the robot to the animal, and r_i is the desired distance. To cage, we set $r_i = R^*$ Let ℓ_{i+1} and ℓ_{i-1} be the distances of the two closest neighboring robots from opposite sides of the axis defined by γ_i. We set $\phi_i = \{-1, 1\}$ to move towards the neighbor with the highest distance. The motion vector $q_i^c = \langle \eta_i; \hat{\theta}_i \rangle$ is then combined with the predicted motion vector of the animal q_i^g, which is defined by the mean group velocity and the current animal orientation. The resulting motion vector q_i is the summation of these two vectors.

In this paper, we adjusted the algorithm for the group to move within the boundaries of a safe path, instead of moving away from dangers. We model the information about the path in an abstract way, by using an artificial potential field. Figure 4.5 exemplifies an instance of a generated potential field. Let $f : \mathbb{R}^2 \to \mathbb{R}$ denote the function representing the potential at a given position of the environment. If (x, y) is located outside of the path, the potential is equal to the minimum distance from the path. Otherwise, the potential is negatively proportionate to the distance traveled along the pathway (i.e. there is a direct negative relationship between x and $f(x, y)$). As such, if the robot continuously moves according to the gradient ∇f, it will reach the goal location with the minimum potential value. For a position $p = (x, y)$, the one-dimensional piecewise linear interpolants of x to $M^{t,s}$ and $M^{b,s}$ are respectively computed as y_b and y_t. The potential function f is then defined as:

$$f(x, y) = \begin{cases} f_{\min} \frac{x - x_{\min}}{x_{\max} - x_{\min}}, & \text{if } y_b \leq y \leq y_t \\ \inf\{d(p, i) : i \in P\}, & \text{otherwise} \end{cases}$$

with $f_{\min} = -2000$ as the minimum potential value. The respective minimum and maximum x-values of P are denoted as $x_{\min}$ and $x_{\max}$. The function d computes the Euclidean distance between two points.

The gradient ∇f can also be explicitly expressed. If p is inside the safety path, then the gradient is equal to the derivative of P at p', where p' is the closest point of p to P. Otherwise, the gradient is equal to the direction to p', i.e. $\nabla f = p' - p$. As we propose a decentralized multi-robot solution to this shepherding problem, the potential function is a way of representing locally observed information. We assume that robots can localize themselves within a shared spatial reference frame, such that they can access the potential field consistently at their own position. However, they do not

share their positions, orientations, or gradient estimates with one another. The system therefore relies on shared spatial access to a global field, but not on a shared global orientation or centralized coordination. For real-life applications, this artificial potential field can be implemented to represent various environmental stimuli, such as natural light or planned landmarks.

In order to steer the group in the direction of ∇f, we utilize the motion model of an animal as given by Equation 4.1, where an animal moves in the opposite direction of a robot when their relative distance is less than z_I. The robots position themselves relative to the gradient of the potential field, $\nabla f(x, y)$. When a robot is located behind all nearby animals, it moves closer to the animals along the direction of ∇f, thereby nudging the group forward. This is done by adjusting the robot's desired distance to the animal to r^* in Equation 4.2. Conversely, when the robot is positioned in front of the group (relative to ∇f), it reverts to the caging behavior by setting $r_i = R^*$.

4.5.5 Analysis

To measure the percentage of the group inside the path (ρ), we take for every individual position $p = (x, y)$, the one-dimensional piecewise linear interpolants of x to $M^{t,r}$ and $M^{b,r}$ are respectively computed as y_b and y_t. If $y_b \leq y \leq y_t$, then the individual is considered to be inside the path. To measure the percentage of the group that is caged ($1 - \psi$), a caging formation can be constructed by the robotic agents based on the repulsive animal-robot threshold z_I. When two robots are at a distance lower than $2z_I$ of each other, they exert a combined repulsive force on the group which prevents them from intersecting the path between those agents. In order to measure whether the robots established an appropriate caging formation, we see if a polygon can be constructed from the edges between robots where the length is shorter than $2z_I$. Following Varava et al. [29], we then define an animal to be caged, if and only if, the polygon is not closed and the animal is located in the interior of the polygon. To approximate the computationally intensive method of constructing a polygon from all edges (with a large number of points), the robots are first clustered. A robot is in the same cluster as its two nearest neighbors, with a relative distance lower than $2z_I$, from opposite sides of the axis defined by the bearing of this robot from the nearest animal. If a robot does not have two neighbors, it is not in a cluster. Afterward, the convex and concave hulls are generated for each cluster. If any of these shapes are closed (i.e. each segment of the shape has a length shorter than $2z_I$), the number of animals inside the hull is measured.

Quantifying group-path alignment involves considering the positions and directions of individual animals. Let $p_a = (x_a, y_a)$ and q_a represent the position and direction of animal a respectively. The alignment, defined as the average alignment between each animal's direction and the de-

sired direction along or towards the path, is expressed as $\frac{1}{2N_G}\sum_{a=1}^{N_G}|\hat{q}_a+\hat{v}|$. Here, $\hat{q}_a$ is the unit direction vector, and $\hat{v}$ is the unit vector of the gradient $\nabla f(x_a, y_a)$ based on the potential field function f (refer to the previous subsection). To assess the compression of the animal group, a convex hull is constructed for all animal positions. From this convex hull, the minimum bounding rectangle (also known as the bounding box) is generated. The compression is then determined by the ratio of the smaller (width) to the larger (length) dimensions of this rectangle.

4.5.6 Simulation parameters

The simulation parameters are based on our previous work [5] and are shown in Supplementary Table S2.1. The parameters of the animals are inspired by live Trinidadian guppies, who live in shallow waters, which applies to our problem setting of a two-dimensional environment. The parameters of the robots are based on the work of Landgraf et al. [38] designing fish-like robots.

Bibliography

[1] S. Van Havermaet, Y. Khaluf, and P. Simoens. Reactive shepherding along a dynamic path. *Scientific Reports*, 14(1):14915, 2024.

[2] B. Le Gallic and A. Cox. An economic analysis of illegal, unreported and unregulated (iuu) fishing: Key drivers and possible solutions. *Marine Policy*, 30(6):689–695, 2006.

[3] C. G. Scanes. Human activity and habitat loss: destruction, fragmentation, and degradation. In *Animals and human society*, pages 451–482. Elsevier, 2018.

[4] L. Jacquin, Q. Petitjean, J. Côte, P. Laffaille, and S. Jean. Effects of pollution on fish behavior, personality, and cognition: some research perspectives. *Frontiers in Ecology and Evolution*, 8:86, 2020.

[5] S. Van Havermaet, P. Simoens, T. Landgraf, and Y. Khaluf. Steering herds away from dangers in dynamic environments. *Royal Society Open Science*, 10(5):230015, 2023.

[6] Y.-S. Ryuh, G.-H. Yang, J. Liu, and H. Hu. A school of robotic fish for mariculture monitoring in the sea coast. *Journal of Bionic Engineering*, 12(1):37–46, 2015.

[7] B. Bat-Erdene and O.-E. Mandakh. Shepherding algorithm of multi-mobile robot system. In *2017 First IEEE International Conference on Robotic Computing (IRC)*, pages 358–361. IEEE, 2017.

[8] M. Evered, P. Burling, M. Trotter, et al. An investigation of predator response in robotic herding of sheep. *International Proceedings of Chemical, Biological and Environmental Engineering*, 63:49–54, 2014.

[9] K. J. Yaxley, K. F. Joiner, and H. Abbass. Drone approach parameters leading to lower stress sheep flocking and movement: Sky shepherding. *Scientific reports*, 11(1):7803, 2021.

[10] Z. Butler, P. Corke, R. Peterson, and D. Rus. From robots to animals: Virtual fences for controlling cattle. *The International Journal of Robotics Research*, 25(5-6):485–508, 2006.

[11] R. Vaughan, N. Sumpter, J. Henderson, A. Frost, and S. Cameron. Experiments in automatic flock control. *Robotics and autonomous systems*, 31(1-2):109–117, 2000.

[12] J.-M. Lien and E. Pratt. Interactive planning for shepherd motion. In *AAAI Spring Symposium: Agents that Learn from Human Teachers*, pages 95–102, 2009.

[13] A. A. Paranjape, S.-J. Chung, K. Kim, and D. H. Shim. Robotic herding of a flock of birds using an unmanned aerial vehicle. *IEEE Transactions on Robotics*, 34(4):901–915, 2018.

[14] S. Gade, A. A. Paranjape, and S.-J. Chung. Herding a flock of birds approaching an airport using an unmanned aerial vehicle. In *AIAA Guidance, Navigation, and Control Conference*, page 1540, 2015.

[15] A. Garrell, A. Sanfeliu, and F. Moreno-Noguer. Discrete time motion model for guiding people in urban areas using multiple robots. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 486–491. IEEE, 2009.

[16] F. Mondada, A. Martinoli, N. Correll, A. Gribovskiy, J. I. Halloy, R. Siegwart, and J.-L. Deneubourg. A general methodology for the control of mixed natural-artificial societies. Technical report, Pan Stanford Publishing Singapore, 2013.

[17] A. J. King, A. M. Wilson, S. D. Wilshin, J. Lowe, H. Haddadi, S. Hailes, and A. J. Morton. Selfish-herd behaviour of sheep under threat. *Current Biology*, 22(14):R561–R562, 2012.

[18] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot sheepdog project achieves automatic flock control. In *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour*, volume 489, page 493, 1998.

[19] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34, 1987.

[20] D. Strömbom, R. P. Mann, A. M. Wilson, S. Hailes, A. J. Morton, D. J. Sumpter, and A. J. King. Solving the shepherding problem: heuristics for herding autonomous, interacting agents. *Journal of the royal society interface*, 11(100):20140719, 2014.

[21] P. Nalepka, R. W. Kallen, A. Chemero, E. Saltzman, and M. J. Richardson. Herd those sheep: Emergent multiagent coordination and behavioral-mode switching. *Psychological science*, 28(5):630–650, 2017.

[22] A. Ranganathan, A. Heyde, A. Gupta, and L. Mahadevan. Optimal shepherding and transport of a flock. *arXiv preprint arXiv:2211.04352*, 2022.

[23] T. Miki and T. Nakamura. An effective simple shepherding algorithm suitable for implementation to a multi-mmobile robot system. In *First International Conference on Innovative Computing, Information and Control-Volume I (ICICIC'06)*, volume 3, pages 161–165. IEEE, 2006.

[24] J.-M. Lien, S. Rodriguez, J.-P. Malric, and N. M. Amato. Shepherding behaviors with multiple shepherds. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3402–3407. IEEE, 2005.

[25] A. Pierson and M. Schwager. Bio-inspired non-cooperative multi-robot herding. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1843–1849. IEEE, 2015.

[26] N. K. Long, K. Sammut, D. Sgarioto, M. Garratt, and H. A. Abbass. A comprehensive review of shepherding as a bio-inspired swarm-robotics guidance approach. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 4(4):523–537, 2020.

[27] Y. Tsunoda, Y. Sueoka, Y. Sato, and K. Osuka. Analysis of local-camera-based shepherding navigation. *Advanced Robotics*, 32(23):1217–1228, 2018.

[28] W. Lee and D. Kim. Autonomous shepherding behaviors of multiple target steering robots. *Sensors*, 17(12):2729, 2017.

[29] A. Varava, K. Hang, D. Kragic, and F. T. Pokorny. Herding by caging: a topological approach towards guiding moving agents via mobile robots. In *Robotics: Science and Systems*, pages 696–700, 2017.

[30] C. Muro, R. Escobedo, L. Spector, and R. Coppinger. Wolf-pack (canis lupus) hunting strategies emerge from simple rules in computational simulations. *Behavioural processes*, 88(3):192–197, 2011.

[31] B. Webb. What does robotics offer animal behaviour? *Animal behaviour*, 60(5):545–558, 2000.

[32] D. Romano, E. Donati, G. Benelli, and C. Stefanini. A review on animal–robot interaction: from bio-hybrid organisms to mixed societies. *Biological cybernetics*, 113:201–225, 2019.

[33] K.-K. Oh, M.-C. Park, and H.-S. Ahn. A survey of multi-agent formation control. *Automatica*, 53:424–440, 2015.

[34] W. Jiang, G. Wen, Z. Peng, T. Huang, and A. Rahmani. Fully distributed formation-containment control of heterogeneous linear multiagent systems. *IEEE Transactions on Automatic Control*, 64(9):3889–3896, 2018.

[35] M. Porfiri, D. G. Roberson, and D. J. Stilwell. Tracking and formation control of multiple autonomous agents: A two-level consensus approach. *Automatica*, 43(8):1318–1328, 2007.

[36] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.

[37] D. R. Nelson, D. B. Barber, T. W. McLain, and R. W. Beard. Vector field path following for miniature air vehicles. *IEEE Transactions on Robotics*, 23(3):519–529, 2007.

[38] T. Landgraf, D. Bierbach, H. Nguyen, N. Muggelberg, P. Romanczuk, and J. Krause. Robofish: increased acceptance of interactive robotic fish with realistic eyes and natural motion patterns by live trinidadian guppies. *Bioinspiration & biomimetics*, 11(1):015001, 2016.

Appendices

4.A Percentage of the group reaching the desired goal

Figure S4.A.1 shows the percentages of the group reaching the desired goal for varying parameter configurations. A comparison of Figure 2 and Figure S4.A.1a indicates that maintaining the group on the safe path near the end ensures reaching the desired goal. From Figure S4.A.1a, we observe that 20 robots cannot guarantee that all 60 animals reach the desired goal under the considered configuration parameters. In Figure S4.A.1b, we observe that the percentage of animals reaching the desired goal decreases with increasing maximum velocity v_G.

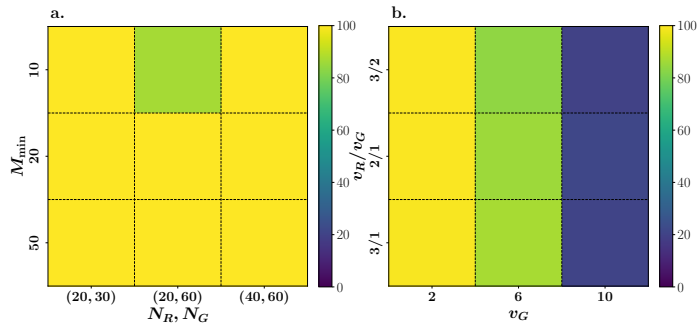

Supplementary Figure 4.A.1: Percentages of the group reaching the desired goal under varying configuration parameters. **(a)** and **(b)** respectively correspond to the simulations presented in the manuscript in Figure 2 and 7. Results are averages over 30 stochastically independent simulations. Configuration of **(a)** includes $v_R = 4, v_G = 2$ and $K = 3$. Configuration of **(b)** includes $N_R = 20, N_G = 30, M_{\min} = 10$ and $K = 1$.

4.B Performance for $K = 1$

Supplementary Figure S4.B.1 displays the percentage of the group within the safe path (ρ) and the percentage of the group not caged (ψ), as a function of the mean x-coordinate of the group. Each column showcases outcomes for distinct system sizes (i.e. the number of robots N_R, and the number of animals in group N_G), while each row varies in the minimum margin width $M_{\min}$ (A: 10, B: 20, C: 50) used to generate the margins. Although the resulting paths consist of the same series of line segments, they consequently differ in margin widths on both sides of these segments. A visual representation of each path is provided as an inset in the first column of the corresponding row (A-C.1).

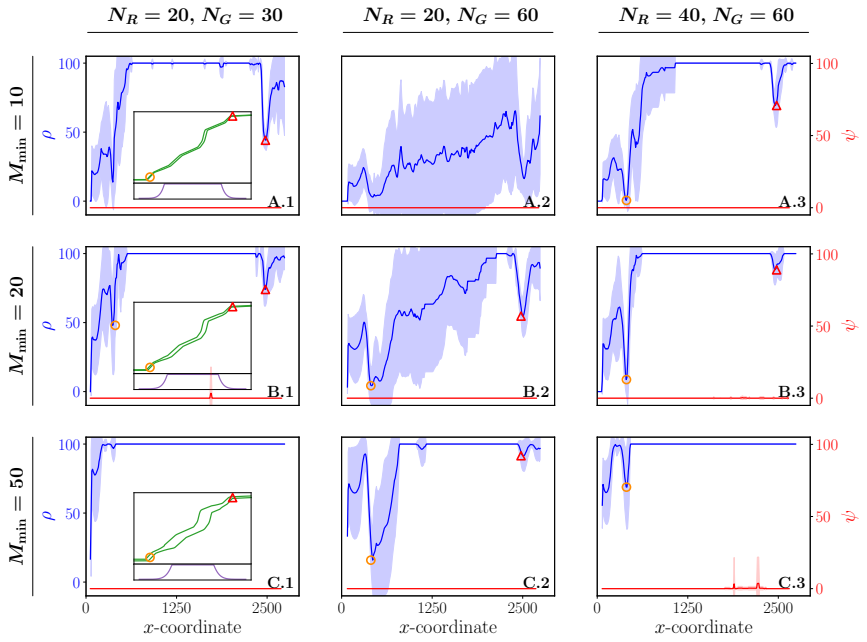

Supplementary Figure 4.B.1: The percentage of the group within the safe path (denoted as ρ, colored in blue) and the percentage of the group that is not caged (denoted as ψ, colored in red) in the function of the mean x-coordinate of the group. Each row corresponds to simulation results for different values of the minimum margin width M_{min} of the path: 10 (**A.1-3**), 20 (**B.1-3**), 50 (**C.1-3**). Each column considers different sizes of the system (i.e. the number of robots N_R and the size of the group N_G): $N_R = 20, N_G = 30$ (**A-C.1**), $N_R = 20, N_G = 60$ (**A-C.2**), $N_R = 40, N_G = 60$ (**A-C.3**). The first column shows the corresponding path of its row as an inset. Significant declines in performance are marked by different colors and symbols: an open red triangle, and an open orange circle. The location on the path where these declines happen is correspondingly marked in the insets. The mean (solid line) and variance (shaded area) of ρ and ψ results from averaging over 30 stochastically independent simulations for each configuration, with $K = 1$.

For all considered values of $M_{\min}$, the systems with a size ratio of 2/3 (i.e. sizes 20/30 and 40/60) consistently reach a stable optimal state, wherein the entire group remains within the safe path ($\rho = 100$) continuously for a substantial length of the path. Systems with a 1/3 size ratio only reach a stable state when the minimum margin is at its largest (C.2). To attain this stable optimal state, these systems require some initial distance, which diminishes as the minimum margin $M_{\min}$ increases. While observing the initial stabilization process, we note a significant disruption for most configurations at a specific point along the path where the margin width is minimal and a sharp turn occurs. Supplementary Figure S4.B.1 indicates this position by an open orange circle. Despite the sudden decline in ρ at this point, the group remains fully caged ($\psi = 0$), enabling these systems to recover and eventually reach a stable optimal state. However, this state is eventually disrupted by another point along the path, marked by an open red triangle in Supplementary Figure S4.B.1. Again, this location is characterized by a minimal margin width and a sharp turn. We observe that these disruptions were mitigated to a lesser extent with an increase in the minimum margin width. Indeed, for $M_{\min} = 50$, the stable optimal state of a system with size 20/30 remains unaffected by the aforementioned scenarios.

Finally, when comparing the initial distance to stabilize with the group contained within the safe path between $K = 1$ (Supplementary Figure S4.B.1) and $K = 3$ (Main Figure 2), we notice that the path with margin width generated by $K = 3$ Gaussian distributions exhibits a significantly earlier stabilization point. This can be attributed to the distinctive characteristics of the single Gaussian component ($K = 1$), which tends to produce a longer narrow segment at the start and end of the path, given that the total length of the path remains constant for both $K = 1$ and $K = 3$. As a consequence, the robots struggle with aligning the group within the path at the beginning, causing a delay in achieving the desired level of stability. For this reason, the declines for $K = 3$ at the first and last turn of the path, respectively indicated by an open orange circle and red triangle, are less substantial in comparison to $K = 1$.

4.C Worst-case lower bound on the number of robots

We estimate the minimum number of robots N_R needed to cage a group of size N_G in the worst-case scenario, where caging is defined as follows:

Definition 4.1. *In a caging formation, each robot maintains a distance of R^* from their respective closest animal, while positioning themselves at equal distance from their two neighboring robots. These two neighbors are positioned on opposite sides of the axis defined as the line intersecting the positions of the focal robot and their closest animal. The distance between consecutive robots should be at most $2R^*$.*

From a geometrical perspective, each animal can be considered to construct a circle with its position as the center and a radius of R^*. Assuming the union of these N_G circles forms a connected set, the animal group can then be represented by a single shape.

Assumption 4.3. *The union of the circles, where the center points are the positions of every animal and the radii are equal to R^*, is a connected set.*

Under these conditions, positioning N_R robots in a caging formation surrounding the group corresponds to uniformly distributing N_R points along the perimeter P of this shape. We then derive an estimate of the absolute minimum N_R points needed for a spatial arrangement of the group that maximizes P. In such an arrangement, the circles touch (Assumption 4.3) but do not overlap more than necessary. An example of this is N_G circles in a linear arrangement, where the distance between consecutive circle center points is $2R^* - \varepsilon$, with ε being a small value that ensures minimal overlap. We will derive the following equations based on this spatial arrangement.

The perimeter P can then be approximated as the sum of all individual circles minus the parts hidden by the overlaps. To calculate the overlap, we find the intersection points (x_e, y_e) of consecutive circles, which can be approximated for small ε as follows:

$$x_e \approx R^* - \frac{\varepsilon}{2}$$
$$y_e \approx \pm\sqrt{R^*\varepsilon}$$

The length of the shortest arc ℓ_e between these two points can then be computed as $\ell_e = R^*\theta_e$ where θ_e is the angle subtended by the chord at the center of one circle, given by:

$$\theta_e = 2\cos^{-1}\left(1 - \frac{\varepsilon}{2R^*}\right)$$
$$\theta_e \approx 2\sqrt{\frac{\varepsilon}{R^*}}$$

Thus, the length that should be removed from the perimeter of each circle for each overlap is approximately $\ell_e \approx 2\sqrt{R^*\epsilon}$. The first and last circles will have one overlap, while each other circle has two distinct overlaps. Therefore, the perimeter P of the shape defined by N_G circles can be approximated as:

$$P \approx 2\pi R^* N_G - 4(N_G - 1)\sqrt{R^*\varepsilon}.$$

Next, we estimate how many points N_R are required to be equidistantly distributed on this perimeter with maximum distance $2R^*$ between them. For the robots, we lower the maximum distance to $z = 2(R^* - e)$ with $e \ll R^*$, to account for the physical size of the robot and the their stochastic actuators and sensors. The perimeter P is an accumulation of (curved) arc

lengths of every circle. In order to measure a ratio between P and z, we must translate z to arc length. Let z be a chord length, then the central angle θ_z subtended by the chord of length z is approximated by Taylor series expansion as

$$\begin{aligned}\theta_z &= 2\sin^{-1}\left(1-\frac{e}{R^*}\right)\\ &\approx \pi - 2\sqrt{\frac{2e}{R^*}}.\end{aligned}$$

The arc length corresponding to this angle is then $\ell_z = \pi R^* - 2\sqrt{2R^*e}$. The lower bound on N_R is given by $\frac{P}{\ell_z}$. Thus, to cage N_G animals behaving under Assumption 4.3, at least N_R robots are required, as given by the following inequality:

$$N_R \geq \frac{2\pi R^* N_G - 4(N_G-1)\sqrt{R^*\varepsilon}}{\pi R^* - 2\sqrt{2R^*e}}. \tag{4.3}$$

We note that the lower bound (4.3) of N_R is derived when the group is exactly positioned to maximize the group perimeter. However, in our simulations, we observe that the spatial arrangement of the group gradually conforms to the shape of the path, eventually becoming highly elongated (see Figure 5b in the main manuscript, where $M_{\min} = 10$). This spatial arrangement resembles the one used to derive the lower bound. Nevertheless, our simulations indicate that the robots can maintain the cage with fewer robots than the calculated lower bound when ε is small (i.e. the group is not tightly connected). This suggests that the dynamics of the animal motion model and the robot control laws used in our simulations result in a more tightly connected group.

4.D Fault tolerance

Supplementary Figure S4.D.1 presents the percentage of the group within the safe path (ρ), the percentage of the group not caged (ψ), and the number of active robots (N_R^{act}) as functions of time t. At each time step t_F, one robot is uniformly randomly selected to malfunction from the set of robots positioned at l_F. To ensure a minimally functioning system, malfunctions do not occur when only 3 are currently functioning ($N_R^{\text{act}} = 3$), or when there is no active robot positioned at l_F. The system's performance of 20 robots and 30 animals is observed through two variables: (i) Δt_F, representing the time between consecutive malfunctions, and (ii) l_F, the coordinate relative to the group's frame of reference, with $l_F \in [BACK, FRONT, LEFT]$. The direction $RIGHT$ is omitted due to symmetry with $LEFT$.

With a fault interval of $\Delta t_F = 200$ (A.1-3), the robots consistently maintain a cage formation, ensuring all animals of the group remain inside the

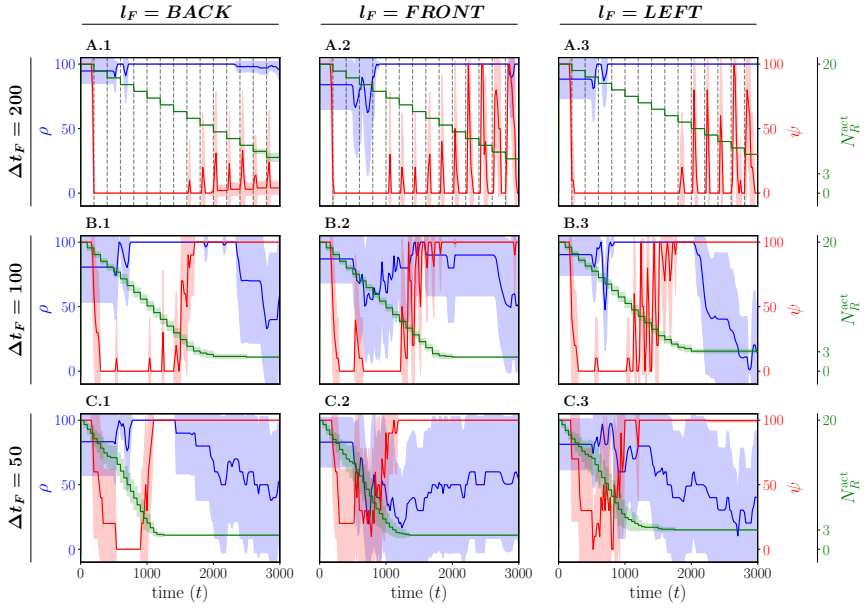

Supplementary Figure 4.D.1: The percentage of the group within the safe path (denoted as ρ, colored in blue), the percentage of the group that is not caged (denoted as ψ, colored in red), and the number of active robots (denoted as N_R^{act}, colored in green) in the function of time t. At each time step t_F, a robot is randomly selected to malfunction from the set of robots positioned at l_F. Each row corresponds to simulation results for different values of the fault interval Δt_F which determines the duration between consecutive faults: 200 (**A.1-3**), 100 (**B.1-3**), 50 (**C.1-3**). Each column considers different coordinates relative to the group's frame of reference l_F: *BACK* (**A-C.1**), *FRONT* (**A-C.2**), *LEFT* (**A-C.3**). The time steps t_F at which a robot malfunctions are indicated by vertical dashed lines. While they are displayed for configurations in the first row, they are omitted in subsequent rows for better visibility. The mean (solid line) and variance (shaded area) of ρ and ψ results from averaging over 10 stochastically independent simulations for each configuration, with $N_R = 20, N_G = 30, M_{\text{min}} = 50$, and $K = 1$.

cage ($\psi = 0$) even as robots periodically malfunction, until less than a critical number of robots remain operational. However, upon comparing failures across different group-centric coordinates (l_F), it becomes evident that this ability to keep the entire group within the cage diminishes much faster when robots malfunction in front of the group compared to malfunctions occurring on the left side or at the back. As the number of operational robots decreases, a certain percentage of animals is left uncaged each time a robot malfunctions. Nevertheless, the robots continuously manage to re-establish their caging formation within the $\Delta t_F = 200$ duration before another robot becomes inactive. This recovery period appears to lengthen as the number of functioning robots decreases, and correspondingly, the percentage of the group that remains uncaged during recovery seems to increase as the number of active robots decreases. It is noteworthy that despite temporary difficulties in keeping the group caged, the group consistently remains fully within the safe path ($\rho = 100$). This resilience is facilitated by the sufficiently wide margin width ($M_{\min} = 50$), which allows for temporary issues without compromising the group's safety within the designated path. s Δt_F is reduced to 100 (B.1-3) and 50 (C.1-3), it becomes apparent that the robots reach a point where they are unable to restore the cage formation when the number of functioning robots is approximately less than 10. Consequently, the group remains entirely uncaged, allowing the animals to move freely. This lack of containment may lead to the group deviating from the safe path.

4.E Varying robot velocity for high-speed animals in straight-line paths

In the manuscript, we have shown that the robots significantly fail to cage the vast majority of the group when the maximum animal velocity is set to $v_G = 10$, independent of the maximum robot velocity v_R. Consequently, the group moves freely and a significant part of the group leaves the safe path. In order to investigate whether the path characteristics play a role in the ability to maintain a caging formation, we have run additional simulations where the path has no turns (i.e. the path represents a horizontal line from the starting position to the desired goal). Supplementary Figure S4.E.1 demonstrates how the percentage of the group within the safe path (ρ) and the percentage of the group not caged (ψ) as functions of the x-coordinate of the group, for various values of maximum robot velocity v_R. In comparison to the results presented in the manuscript for dynamic paths, it seems that the significantly low performance for high-speed animals is not related to path characteristics. We therefore hypothesize that the failure in caging high-speed animals may be due to the prediction error of the animal direction that is used in the robotic caging algorithm. These directional prediction errors cause larger positional errors when velocity

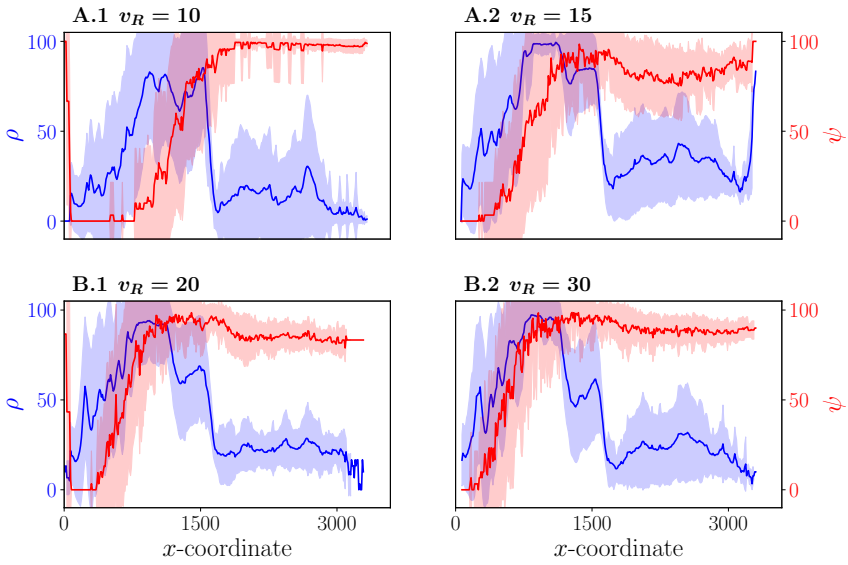

Supplementary Figure 4.E.1: The percentage of the group within the safe path (denoted as ρ, colored in blue), and the percentage of the group that is not caged (denoted as ψ, colored in red), in the function of the x-coordinate of the group. The robots are tasked with shepherding a high-speed animal group along a safe path without turns. The sub figures show results of various maximum velocities of the robot v_R. The mean (solid line) and variance (shaded area) of ρ and ψ results from averaging over 10 stochastically independent simulations for each configuration, with $N_R = 20, N_G = 30, M_{\text{min}} = 10$, and $v_G = 10$.

symbol	parameter	value	association
N_G	size of the group	$\{30, 60\}$	animal
N_R	size of the robot swarm	$[20, ..., 40, ..., 60]$	robot
M_{min}	minimum margin width	$\{10, 20, 50\}$	path
K	number of Gaussian distribution components	$\{1, 3\}$	path
v_G & v_R	max. linear velocity	2 & 4 $\frac{\text{cm}}{\text{s}}$	animal & robot
w_G & w_R	max. angular velocity	π & π $\frac{\text{rad}}{\text{s}}$	animal &robot
σ_G & σ_R	Gaussian orientation noise	0.05 & 0.05 rad	animal & robot
z_R	radius of repulsion zone	1 cm	animal
z_O	radius of orientation zone	50 cm	animal
z_A	radius of attraction zone	30 cm	animal
z_I	radius of aversive zone	40 cm	animal
α_R	weight of repulsion interaction	100	animal
α_O	weight of orientation interaction	50	animal
α_A	weight of attraction interaction	1	animal
α_I	weight of aversion interaction	100	animal
d_d	detection radius	50 cm	robot
κ	transition rate	0.2	robot
R^*	relative distance robot-animal to observe	$z_I + v_G$ cm	robot
r^*	relative distance robot-animal to steer	$z_I - v_G$ cm	robot

Table 4.F.1: Simulation parameters. For each parameter, the table lists: (i) the symbol used for reference in the main manuscript (or refers to Van Havermaet et al. [5] if not present in this manuscript), (ii) the name of the parameter, (iii) the value(s) employed in the experiments, and (iv) the association indicating whether it relates to the animal, robot, or path.

of the animals is increased. A first priority of future work is to address this issue.

4.F Simulation parameters

Supplementary Table S4.F.1 provides a comprehensive overview of the simulation parameters governing the dynamics of the robotic and animal entities within the two-dimensional environment. The parameters are categorized based on their association with animals, robots, and the path, each contributing to the simulation's realism and relevance. The parameters of the animal group are inspired by live Trinidadian guppies, who live in shallow waters, which applies to our problem setting of a two-dimensional environment. The robot parameters are derived from two-wheeled motion robots used in the real-world experiments of Chapter 5, in order to simulate physically plausible behavior and ensure consistency with the capabilities of the actual robotic platform.

The work presented in this chapter is in review for publication in:
Scientific Reports (Van Havermaet et al. 2025 [1])

5

The influence of a fish-like robot on the avoidance behavior of fish

Abstract
In the two previous chapters, we proposed algorithms for robots to protect fish by directing them away from danger or guiding them along safe trajectories. These methods relied on influencing fish movement through avoidance behaviors elicited by robot-fish interactions. Traditionally, robots designed to resemble natural predators have been used to provoke avoidance responses. However, fish also exhibit avoidance behaviors in response to conspecifics in various contexts. In a mixed society of fish and robots, we argue that conspecific-like robots offer greater flexibility by integrating with the group and acting only when needed. Yet, how fish respond to conspecific-like robots compared to traditional avoidance models remains largely unexplored, particularly regarding repeated exposure, approach speed, and proximity. To investigate this, we programmed a conspecific-like robot to repeatedly approach live fish at varying speeds in a free-swimming setup. Our findings indicate that repeated exposure to the robot increased the likelihood of avoidance responses, suggesting behavioral adaptation over time. In some instances, the fish exhibited classic anti-predator behaviors, such as freezing and escape maneuvers, indicating that the robot may be perceived as a threat. Notably, escape initiation was influenced by both the robot's speed and proximity, challenging conventional distance-based avoidance models. Avoidance speed also varied dynamically with these factors. These results provide insights into fish-robot interactions and inform the design of robotic systems capable of predicting and influencing fish behavior.

5.1 Introduction

Recent advancements in robotics have enabled the creation of biomimetic robots capable of interacting with live animals in a naturalistic manner [2, 3]. These robots mimic biological forms and behaviors, providing a gateway to the establishment of mixed animal-robot societies where robots and live animals mutually influence each other's behavior [4–6]. Such systems offer significant advantages for studying animal behavior, as they allow for precise control of robots to test specific hypotheses about social interactions in real-world experiments with standardized, reproducible designs [7–10]. Beyond research, these animal-robot systems have applications ranging from livestock management [11, 12] to environmental protection [13, 14]. For instance, robots could help steer fish away from ecological threats they may not naturally recognize, such as pollution [15], by leveraging their avoidance behavior [16, 17].

Theoretical [18] and empirical [19, 20] studies have identified three primary interaction forces governing conspecific fish behavior: attraction, alignment, and avoidance. Previous research has focused primarily on how conspecific-like robots elicit attraction [21] and alignment [22]. These robots mimic the appearance of a member of the same species. However, live animals do not always exhibit attraction toward their conspecifics. Depending on the social context, they may instead display avoidance behavior. For example, in mammals, encounters between conspecifics with conflicting goals can become distinctly agonistic, ranging from mutual avoidance to fighting [23]. While avoidance behavior often incurs the cost of lost social benefits, it can be advantageous in certain contexts. For instance, guppies avoid infected conspecifics to mitigate the risk of disease transmission [24]. In the context of mating, male guppies may engage in chasing behavior, where one male flees and ceases courting a female after being rapidly approached by another male [25]. Understanding when attraction is elicited and when avoidance occurs in response to conspecifics is crucial for a more comprehensive analysis of these social forces in animal behavior. Moreover, this knowledge is essential for advancing the design of more adaptable and effective robots that can dynamically respond to these interactions in complex environments. This study aims to provide insights that will inform future improvements in modeling and controller design, particularly in the elicitation of avoidance behavior.

In this study, we present findings from an experiment in which live Trinidadian guppies (*Poecilia reticulata*) interacted with a mobile, conspecific-like robot in the same tank, without physical separation. The objective was to assess the guppies' behavioral responses to the approaching robot. Specifically, we sought to examine various factors that influence avoidance behavior, such as robot speed, proximity, and repeated exposure. Additionally, we aimed to investigate whether key anti-predator behaviors [26–28], such as freezing (where the fish remains motionless as

a defense mechanism) and escape maneuvers (a sharp turn followed by rapid acceleration to evade the robot), manifest during these interactions. Finally, we evaluated whether the observed responses align with predictions from established interaction force models between conspecifics. To explore this, the robot was programmed to approach a fish from an initial distance of 40 cm, continuing until it reached close proximity. Each fish underwent multiple consecutive trials for at least 10 minutes. In each trial, the robot's approach speed was set to a fixed value, ranging from 15 to 30 cm/s. These findings provide valuable insights for improving prediction models of fish behavior, which could potentially be used to guide fish away from threats by influencing their movement in response to robot approaches [16, 17].

The growing body of research in animal-robot interactions has particularly emphasized the integration of robots with fish species [29, 30]. Much of this work has explored how biomimetic robots can be accepted by fish as conspecifics. For instance, recent findings suggest that biomimetic robots capable of adjusting their behavior in response to a guppy's actions can more effectively elicit following behavior and reduce avoidance reactions, further improving integration within fish groups [31]. In contrast, studies on the avoidance behavior of fish are primarily conducted in the context of anxiety-related and anti-predator scenarios. In response to predation threat, guppies increase anti-predator behaviours including escape movements, freezing, predator inspections and thigmotaxis (i.e. area avoidance) [32]. Escape responses typically involve rapid acceleration and a directional change to distance the fish from the threat [28]. Success in avoiding predators depends on factors such as timing, reaction distance, locomotor performance, and escape direction. Studies have shown that prey do not escape at the furthest distance a threat is perceived, but at a distance determined by the relative costs and benefits of escaping, which is influenced by factors like predator approach speed [33]. Recent findings suggest that the avoidance-initiation distance is also influenced by repeated exposure to the threat [34]. Although the robot in our study mimics a conspecific, we hypothesize that its repeated approaches may lead the guppy to perceive it as a threat, thereby increasing the likelihood of avoidance behavior with successive trials.

Researchers have also employed interactive robots to simulate anxiety-related [35] and anti-predator [36–39] scenarios to study the behavioral responses of fish. Most studies involved exposing live zebrafish to a 3D-printed predator-like replica of a red tiger. Spinello et al. [39] found that the robotic replica triggered fear-related responses, evidenced by geotaxis, although this effect diminished over time as the robot's interactivity increased. Surprisingly, the zebrafish's visual exposure to the replica did not result in a significant increase in the average distance from the robot. Ladu et al. [36] demonstrated a robotic replica of a predator induced significant fear responses in zebrafish, evidenced by increased thrashing behavior and

spatial avoidance, compared to the control and computer-animated conditions. These studies consistently involve physical separation when investigating avoidance behavior. Removing such barriers could allow for a finer analysis of how relative distance and robot speed impact avoidance behavior, as well as the spatial aspects of the fish's responses.

To the best of our knowledge, few studies have examined avoidance behavior in live fish interacting with conspecific-like robots in free-swimming contexts. Kruusmaa et al. [40] found that collective avoidance responses in mackerel schools were influenced by both the size and speed of a robotic replica. However, their study only examined the probability of avoidance for two distinct speed levels, rather than analyzing fish speed in function of a continuous range of robot speeds. Examining a wider range of speeds is important, as speed regulation plays a critical role in avoidance behavior during interactions between conspecifics [41]. This could help developers design robots that elicit specific fish speeds, which may be crucial for guiding fish with higher precision. Pino et al. [42] focused on conceptual models to investigate stress, fear, and anxiety in zebrafish caused by bio-inspired robotic fish with varying components and designs. Their freely swimming robots enabled the analysis of fish speed relative to robot-fish distance, suggesting that a decrease in distance results in increased fish velocity, indicative of avoidance behavior. Nevertheless, their findings relied on descriptive results and did not account for individual differences through statistical analyses. Building on these studies, we hypothesize that both robot speed and proximity interact to influence fish avoidance speed, with their combined effects shaping the fish's response to the robot's approach.

Although these studies provide valuable insights, they do not address how fish actions, such as turns and accelerations, vary with the relative distance and speed of an approaching robotic conspecific. These actions are critical indicators of anti-predator behavior, such as escape maneuvers, which could reveal whether conspecific-like robots have the potential to serve multiple roles, acting as a threat or conspecific depending on the context. This flexibility could be particularly valuable in the dynamic environments these animals inhabit. Furthermore, analyzing how avoidance speed varies as a function of both robot speed and robot-fish distance could also help robot developers better understand which actions to program in order to elicit specific responses from fish. Typically, avoidance behavior between conspecifics in fish is modeled solely by distance, where fish move opposite to the average position of neighbors within a defined radius [18–20]. More recent studies have incorporated decay functions that gradually decrease the strength of the avoidance force as the distance increases [43, 44]. However, we hypothesize that fish may adjust their avoidance responses based on the interaction between robot speed and robot-fish distance. We therefore seek to determine whether the guppies in our experimental setup consistently respond within a certain range, aligning

with traditional model assumptions, or if their responses are more dynamic.

We used the RoboFish system, originally introduced by Landgraf et al. [21], which enables an interactive guppy replica to engage with live guppies within the same tank. The 3D-printed replica is attached to a magnetic base below, aligned with a two-wheeled robot carrying a magnet on top that moves beneath the tank floor. This robot design has been shown to be accepted as a conspecific by guppies in multiple studies [21, 22, 31, 45, 46]. We integrated high-definition video tracking with a closed-loop feedback system to steer the robot toward the fish in real time. Given that guppies naturally inhabit shallow waters, all analyses were conducted in two dimensions, focusing on horizontal movement within the tank. Fish behaviors were classified into avoidance, attraction, and slow motion to examine how repeated exposure to the robot's approach influences the onset of avoidance behavior. We also investigated the occurrence of anti-predator or stress behaviors, such as freezing and thigmotaxis, to assess whether the conspecific-like robot was perceived as a threat. Additionally, a fixed-effects model was used to examine how the robot's approach speed and relative distance influenced the fish's avoidance speed. Finally, we analyzed the execution of avoidance behavior, focusing on how relative speed and proximity to the robot influence the fish's decision-making process for avoidance. We also discuss whether these actions align with traditional models of avoidance behavior between fish.

5.2 Material and methods

5.2.1 Study organism and maintenance

For this study, we selected a fish species previously shown to exhibit a high level of acceptance of the robotic model as a conspecific [21]. We used wild-type guppies (Poecilia reticulata), which have been bred in the laboratory for multiple generations. These guppies originally descended from individuals captured in the Arima River, Trinidad, in 2010. The test fish were sourced from large, randomly outbred single-species populations housed at the animal care facilities of the Faculty of Life Sciences, Humboldt University of Berlin. To prevent inbreeding, these stocks are regularly supplemented with wild-caught fish brought back from fieldwork conducted in Trinidad and Tobago. A natural light cycle of 12 hours of light and 12 hours of darkness was provided, with the water temperature consistently maintained at 26°C. The fish were fed twice daily ad libitum with commercially available flake food (TetraMin™). The experiments were conducted in the afternoon, three hours after the initial feeding. Afterwards, all fish were given a second feeding. In the experiments, only adult female guppies were used to avoid any potential influence of sex-specific and life stage-specific differences in behavior, and to ensure that the fish were large

Figure 5.1: Figure reproduced from Maxeiner et al. [31] with permission, showcasing the RoboFish system. Left panel; a 3D-printed fish replica is attached to a magnetic base, which aligns with the robot. Hence the replica can be moved directly by the robot at constant height. Center panel; a wheeled differential-drive robot moves on a transparent platform below the test tank, carrying a neodymium magnet. Right panel; the tank is a quadratic (1 m × 1 m) with a triangular start box used as shelter for the live fish at the beginning of each experiment.

enough for the tracking device to accurately locate them. Adult fish were identified as those with a body length exceeding 15 mm and displaying external signs of maturity, such as a gravid spot.

5.2.2 Experimental apparatus

The RoboFish system (see Figure 5.1), originally introduced by Landgraf et al. [21], enables an interactive fish replica to engage with live guppies within the same environment. The experimental setup consists of a glass tank measuring 120 cm × 120 cm, filled with 7 cm of aged tap water, which simulates the shallow conditions typically found in the guppies' natural habitats, such as rivers and streams. These low water levels are intended to encourage natural behavior in the fish. An experimental area of 100 cm × 100 cm is defined in the center of the tank by four plastic walls, while the surrounding space houses heating elements and a pump to maintain a consistent water temperature of 25°C and to aerate the water. The tank is elevated 1.40 m above the ground, supported by an aluminum rack.

Beneath the tank floor, a two-wheeled differential drive robot operates on a transparent plastic surface, carrying an upward-facing neodymium magnet that aligns with a corresponding magnetic base inside the tank, to which a 3D-printed fish replica is attached. The replica's movement is directly controlled by the robot's motion. The robot is equipped with three red-light LEDs—two on the right side and one on the left—that are visible from below through the transparent pane, aiding in the estimation of the robot's current position and orientation. A Basler acA1300-200um camera (1280 px × 1024 px) is positioned on the floor to track the robot's location, while a second camera, the Basler acA2040-90uc (2040 px × 2040 px), is mounted 1.5 m above the tank to monitor both the live fish and the replica.

To minimize exposure to external disturbances, the entire system is en-

closed within an opaque canvas. The tank is illuminated from above by artificial LED lights that mimic the daylight spectrum. The system operates on a personal computer (i7-6800K, 64GB RAM, GTX1060) that runs custom software for robot control. This software tracks the robot via the bottom camera's feed and controls it through a WiFi connection. Additionally, a second program records video from the top camera, detects and tracks all agents within the tank, and transmits positional data to the robot control software. At each time step (25 Hz), the robot control software updates the positions and orientations of both the fish and the robot in an internal data structure. Behavior modules can access this data to calculate target positions for the robot based on the current or previously observed states. Once the active behavior determines a new target position, the robot moves toward it. If a significant turn is required, the robot first rotates before advancing forward. If only a slight adjustment is needed, the robot moves toward the target while subtly turning by adjusting the relative speeds of its two wheels.

A 3D-printed triangular retainer, with a side length of 19 cm, served as a shelter box to house the fish prior to the experiment. This shelter featured a cylindrical section with a 10 cm diameter, from which the fish could access the experimental area through a 3 cm × 2.5 cm door. Apart from the retainer, the environment was symmetrical and uniform in appearance.

5.2.3 Experimental procedure

In the experiment, we aim to assess the impact of the robotic fish's approach behavior on the avoidance responses of live fish. To initiate the experiment, this study utilizes an approach similar to that of Maxeiner et al. [31]. For each trial, an adult female guppy was randomly selected from its holding tank and gently introduced into the shelter box. The front door of the shelter was then opened, and the robot remained stationary until the fish exited the shelter. If the test fish did not leave the shelter within two minutes, the shelter box was removed. Once the fish left the shelter, the robot commenced a circular milling movement in the center of the tank, with a circumference of 10 cm, for one minute. This milling behavior served as part of the acclimatization process.

Trials began immediately after the milling behavior was completed. The robot followed a consistent protocol for each trial. Initially, the robot remained stationary, waiting for the fish to be 40 cm distance away. If the fish did not reach this distance within 5 seconds, the robot relocated to maintain the desired separation. The 40 cm distance was chosen to maximize the space between the robot and the fish, given the constraints of the tank. Even when the fish was in the center, the robot could keep this separation while staying far enough from the tank walls. When the robot and the fish were sufficiently far apart, the robot began to approach the fish until one of two conditions was met. The first condition ensured that the dis-

tance between the robot and the fish was less than 2 cm, allowing the robot to approach closely to potentially tigger avoidance behavior. The second condition involved the follow metric, a measure introduced by Maxeiner et al. [31] to quantify the fish's avoidance response over a specified time window. It is defined as the time average of the dot product between the fish's velocity vector and the direction vector from the fish to the robot. Negative values indicate movement away from the robot, while larger values reflect faster movements. If the distance was less than 8 cm and the follow metric fell below -1 cm/s, the approach was halted, as it indicates that the fish is beginning to move away. Once either condition was satisfied, the trial concluded.

The robot was set to move directly towards the continuously updated tracked position of the fish. For each trial, the robot was set to approach the fish at a certain speed. We used ten speed configurations, ranging from 15 to 30 cm/s in 1.5 cm/s increments. At the start of each trial, a configuration was randomly chosen from the remaining set. Once a configuration was selected, it was removed from the set. After all configurations were used, the set was reset. To prevent the robot from colliding with the tank walls, we implemented a wall avoidance feature. Since the robot may approach the fish at relatively high speeds in some trials, a significant buffer was necessary. Whenever the robot came within 10 cm of the walls, it relocated to the nearest suitable position that maintained an appropriate distance from the tank's edges.

Each live fish was tested continuously for at least 10 minutes before being returned to the holding tank. This method provided observations of the same fish to the robot approaching from various angles. Additionally, repeated measures allowed for the assessment of any changes in the fish's behavior over time as it adapted to the repeated scenario.

5.2.4 Data processing

A total of 276 trials were conducted with $n = 5$ subjects. Tracking data were manually and automatically checked for errors, processed to correct for missing frames, and converted to centimeters. For the missing frames, the positions were linearly interpolated. Given that our analysis involves the distribution of velocities, interpolation is not performed when there are 25 or more consecutive frames (i.e., one second or longer) without tracking data. Such frames were excluded from the analysis. Additionally, frames from the initialization phase (milling process) were also excluded. After processing, a total of 82,002 frames (approximately 54 minutes) of data were retained across all subjects. A detailed breakdown of the dataset at each step of the data processing pipeline can be found in Supplementary Table S5.A.1.

The tracking data provided two-dimensional positions $\vec{p} = (x, y)$ at a frame rate of 25 Hz for both the robot and the fish. These positions were

smoothed using a running average with Gaussian weights over a time window of 5 frames. This time window effectively captures short, spurious changes in motion [45]. From the smoothed positions, the instantaneous velocity vector $\vec{q}$ with its norm as speed v, and acceleration vector $\vec{u}$ with its norm as acceleration a, were computed for both the robot and the fish, while the distance d_{FR} between them was also calculated. The relative angle between the robot's position and the fish's heading direction was defined as $\theta_{RF} = \arctan(\vec{p}_{RF}) - \theta_F$ where $\vec{p}_{RF} = \vec{p}_R - \vec{p}_F$ is the relative position of the robot with respect to the fish, and $\theta_F = \arctan(\vec{q}_F)$ is fish's heading angle. Furthermore, the robot's approach speed was computed as $v^R_{app} = \vec{q}_R \cdot \hat{p}_{FR}$ and the fish's avoidance speed as $v^F_{avoid} = \vec{q}_F \cdot \hat{p}_{FR}$, with $\hat{p}_{FR} = \frac{\vec{p}_F - \vec{p}_R}{d}$ as the normalized relative position of the fish with respect to the robot. These two quantities, v^R_{app} and v^F_{avoid}, served as the primary variables in our analysis. Although it was possible to normalize the approach speed by the robot's speed and separately use the approach cosine, significant errors in the computed angle (and thus the cosine) were observed when the robot moved very slowly or remained almost stationary. These errors occurred because, at low speeds, even minor tracking inaccuracies disproportionately affected the angle calculation. By incorporating the robot's speed into the approach speed calculation, these inaccuracies were suppressed. The same principle applied to the fish's avoidance speed.

5.2.5 Trial selection

In a substantial portion of the trials, the fish remained in close proximity to the tank wall for the entire duration, which hindered the robot's ability to exhibit its intended behavior of continuously approaching the fish until an encounter occurred. In such instances, the robot executed wall-avoidance maneuvers and repeatedly attempted to approach the fish until one of the predefined halting conditions was met. A subset of trials was automatically selected based on the criterion that wall proximity did not significantly influence the robot's behavior. This subset was later manually reviewed and confirmed, resulting in a final total of 88 trials out of the 276 original trials. This approach allows for a comprehensive evaluation, with the subset providing insights into fish behavior under more consistent robot performance.

5.2.6 Data analysis

The analyses in this study draw on two datasets: (i) the full dataset comprising 82,002 observations collected across all 276 trials, used for general behavioral analysis, and broad statistical modeling, and (ii) a subset of 88 manually selected trials, comprising 15,168 observations, used for in-depth phase-based behavioral analysis and extreme action classification.

Classification of fish behavior

We classified fish behavior to identify different responses to the robot, as well as to determine the occurrence of anti-predator behaviors. When the fish moved at speeds less than 2 cm/s, it was classified as slow motion, in accordance with prior research on Trinidadian guppies [31]. A specific subbehavior of slow motion, referred to as freezing, has been identified as a characteristic stress or anti-predator response in fish, particularly in Trinidadian guppies [47–49]. We initially applied the methodology described by Houslay et al. [47], which defined freezing behavior of guppies as a continuous speed below 4 cm/s for at least 2.5 seconds. However, upon direct observation, we found that this approach resulted in a high number of false positives. Consequently, we adopted a more stringent definition from Ladu et al. [36], which was based on zebrafish, where freezing behavior is defined as the fish moving less than 2 cm in 2 continuous seconds. When the speed exceeded 2 cm/s, the behaviors were further classified based on avoidance speed (v^F_{avoid}): if $v^F_{avoid} > 0$, the behavior was identified as avoidance, whereas $v^F_{avoid} \leq 0$ indicated attraction behavior. The open field arena away from the tank walls, is assumed to be perceived as riskier by guppies [47]. Thigmotaxis has been defined as the tendency to stay close to physical boundaries, a behavior suggested to help avoid exposure to potential threats [26]. A fish was considered to exhibit thigmotaxis when it was located within 10 cm of the nearest boundary ($d_{FW} < 10$). This threshold was determined based on the robot's programmed wall-avoidance behavior at this distance.

Statistical modeling of avoidance speed

We used a fixed-effects (FE) model to investigate how robot behavior affected the fish's avoidance speed v^F_{avoid} across the entire experiment (full dataset). The predictor variables included the distance d_{FR} between the robot and the fish, the robot's approach speed v^R_{app}, and the distance d_{FW} between the fish and the nearest tank wall. Taking response latency into account, we shifted the avoidance speed forward by $\Delta = 0.32s$ (8 frames), a value determined through preliminary analysis of the time gap between the initiation of an accelerative motion and the resulting speed increase. We replicated the model for shifts of 4, 12, and 16 frames, all of which yielded results consistent with the 8-frame shift (see Supplementary Table S5.B.1). To account for individual heterogeneity, an individual-specific fixed effect was included, with the fish's ID serving as a fixed effect. This approach allows the model to estimate a distinct intercept for each subject, which remains constant over time. In our study, this captured characteristics such as the fish's personality. For example, Houslay et al. [47] found that Trinidadian guppies exhibited time-invariant behavioral differences (e.g., evasive behavior versus cautious exploration) when exposed to predator models. Their findings suggest that such behaviors

in stressful situations are driven by stable personality traits rather than temporary states solely influenced by the environment. The inclusion of a time-specific fixed effect in a two-way FE model did not alter the results (see Supplementary Table S5.C.1); therefore, the results presented in the manuscript focus on one-way individual FE models. To run the FE models, we utilized the `plm` package in `R`. We computed the models' coefficients using HC1 robust standard errors.

Phase-based behavioral analysis

Phase-based behavioral analyses were conducted on the subset of 88 manually selected trials. Trials were divided into three phases based on d_{FR}: (i) pre-encounter, (ii) at encounter, and (iii) post-encounter. The encounter phase was defined as the period from 0.4 seconds before to 0.4 seconds after the moment of minimal distance. To investigate how fish adapt their behavior during a trial, we computed the marginal and transition probabilities of fish behaviors across the pre-encounter, encounter, and post-encounter phases. For this analysis, we focused on manually selected trials where the robot's behavior was consistent. In these trials, the robot reached a stable maximum approach speed (v_{app}^{R}) during the pre-encounter phase, encounters the fish, and then remained stationary ($v_{app}^{R} = 0$) in the post-encounter phase. Following the post-encounter phase, the robot moved to a new location to reset the scenario, but this phase was not included in the analysis, as the robot's influence on the fish was likely to vary during this time.

Classification of extreme fish actions

We classified fish turns and accelerations occurring within 40 cm of the robot by comparing them to a baseline dataset. This baseline (12,692 observations) comprised instances when the fish and robot were more than 40 cm apart, during which the robot was not programmed to actively approach, thereby minimizing interaction. This baseline data was used to define the reference distributions for each subject. For turns, baseline distributions were modeled as normal distributions, parameterized by subject-specific means (μ_i) and standard deviations (σ_i), where i denotes the subject identifier. Extreme turns for each subject were then defined as those exceeding the threshold $\mu_i + 2\sigma_i$, calculated separately for each individual to account for variability in baseline behavior. In contrast, acceleration baseline data did not follow a normal distribution. Instead, we classified fish accelerations as extreme if they exceeded the 95th percentile of their respective baseline distributions. Both extreme action classifications were applied to the subset of 88 trials to provide insights under consistent robot performance.

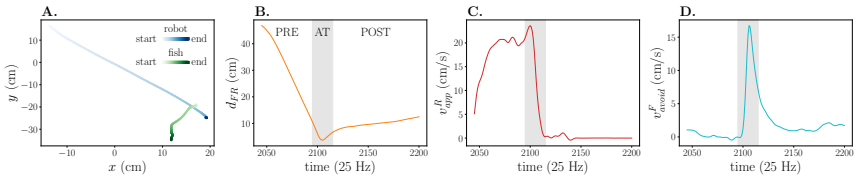

Figure 5.2: (A) The trajectories of both the robot (blue) and the fish (green) are shown in 2D space, with the progression over time indicated by a color gradient from light (start) to dark (end). Measurements over time include **(B)** the distance between the robot and the fish d_{FR}, **(C)** the robot's relative approach speed v^R_{app}, and **(D)** the fish's avoidance speed v^F_{avoid}. The different phases of the trial are labeled in (B), where the (at) encounter phase is highlighted by a gray overlay, also visible in (C) and (D). Time is recorded in frames at a frequency of 25 Hz.

5.3 Results

In this study, we investigate the interaction dynamics between a robot and a fish, focusing on spatial and behavioral responses during their encounter. Figure 5.2 presents the trajectories of both the robot and fish (A), along with the measured spatial and dynamic quantities (B-D) from a single trial. The trial is divided into three phases based on the fish-robot distance (d_{FR}): (i) pre-encounter, (ii) at encounter, and (iii) post-encounter. During the encounter phase, the robot decelerates its approach speed (see Figure 5.2C, where v^R_{app} decreases to zero) as it meets the halting conditions, while the fish accelerates its avoidance speed (v^F_{avoid} increases to a local maximum), demonstrating avoidance behavior as $v^F_{avoid} > 0$. This acceleration further indicates escape behavior due to the fish's rapid motion in response to the robot.

The robot's approach behavior induces avoidance in fish throughout the experiment

Figure 5.3A illustrates the percentage of time, averaged across all fish for the duration of their respective experiments, during which various behaviors were observed. These behaviors include attraction, slow motion, freezing, avoidance, and thigmotaxis. The first four are categorized as dynamic behaviors, while the last is considered a spatial behavior. Among the dynamic behaviors, the fish primarily engage in avoidance behavior, moving away from the robot. This is an expected response to the robot moving into the fish's space. The second most frequent behavior observed is slow motion, with a small percentage of this behavior identified as freezing. Attraction to the robot is also observed, but in much smaller quantities compared to avoidance. This pattern is consistent in most subjects, with the exception of one individual (Fig 5.3B, subject 4), which is the only one to

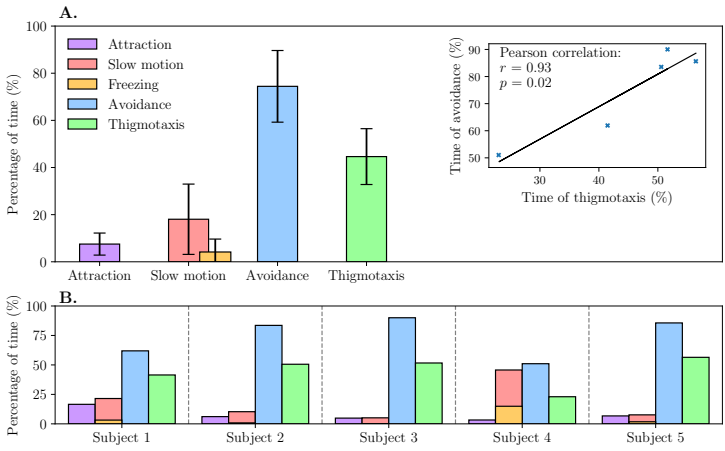

Figure 5.3: (A) The percentage of time the fish spent in various behaviors throughout the experiment: attraction (purple), slow motion (red), freezing (orange), avoidance (blue), and thigmotaxis (green). The bars represent the mean across all subjects, with error bars showing the standard deviation, based on $N = 82,002$ observations. The inset shows the correlation between the percentage of time spent in thigmotaxis and avoidance across subjects. The Pearson correlation coefficient is $r = 0.93$ with a significance of $p = 0.02$ and a 95% confidence interval of $[0.28, 1]$. The black line represents a linear fit, included as a visual aid to highlight the correlation trend. **(B)** The percentage of time each subject (1-5) spent in the different behaviors: attraction, slow motion, freezing, avoidance, and thigmotaxis. Each bar represents the proportion of time allocated to each behavior for the respective subject. Note that the freezing bar overlays the slow motion bar, and in some subjects, the minimal occurrence of freezing makes the bar less visible.

show a substantial amount of freezing, resulting in a reduced percentage of avoidance compared to the others. The fish exhibit notable thigmotaxis, spending approximately half of the experiment near the tank walls. The inset of Figure 5.3A illustrates, for each subject, the percentage of time the fish displays avoidance behavior as a function of the time spent exhibiting thigmotaxis. A strong linear correlation is observed (Pearson's $r = 0.93$, $p < 0.05$), as demonstrated by the linear fit applied to these data points. This correlation suggests that fish exhibiting more frequent avoidance behavior also spend more time near the walls. One possibility is that thigmotaxis serves as a spatial strategy, with fish seeking the safety of the walls to avoid the robot's movements. Alternatively, thigmotaxis could be a natural consequence of the fish traveling farther as they attempt to avoid the robot, eventually reaching the walls where they remain because further movement to escape is no longer possible.

Repeated approaches by the robot increase avoidance behavior and reduce attraction behavior

As the fish were repeatedly exposed to trials, their tendency to avoid the robot increased, even before the robot fully approached during the pre-encounter phase. Figure 5.4A shows the cumulative proportion of fish exhibiting avoidance behavior during this phase, indicating that the elevated level of avoidance in this phase is attributed to repeated trial exposure. A generalized linear mixed model (GLMM) with a logit link function

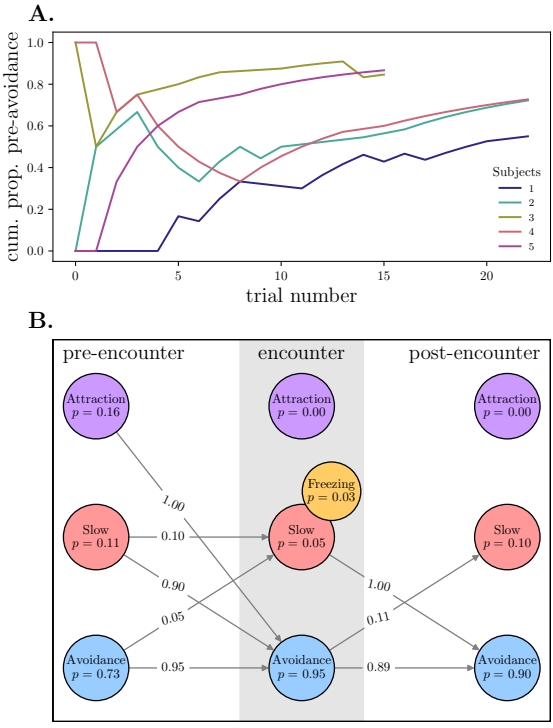

Figure 5.4: (A) The cumulative proportion of the fish displaying avoidance behavior before the robot's approach (i.e. pre-encounter phase), across trials for each subject. **(B)** Transition probability diagram of attraction (purple), slow motion (red), and avoidance (blue) behaviors during the pre-encounter, encounter, and post-encounter phases. The encounter phase is highlighted by a gray overlay. The arrows show the probability of transitioning from one behavior to another, and the associated marginal probabilities p indicate the probability of being in a specific behavior during each phase. Freezing (orange), shown as a sub-behavior of slow motion in the encounter phase, is also represented with its corresponding probability. Probabilities were computed based on $N = 88$ trials. Transitions with zero probability are not shown.

($\chi^2(3) = 27.01$, $p < 0.001$) reveals that the likelihood of avoidance behavior significantly increases with each successive trial (Estimate $= 0.206$, $p < 0.001$). The model includes trial number and robot speed as fixed effects, and subject identity as a random effect. Robot speed did not significantly affect the likelihood of avoidance behavior in pre-encounter phase.

Figure 5.4B shows that, before the robot has encountered the fish (pre-encounter phase), most fish avoid the robot ($p = 0.73$), while the remainder are either moving towards the robot ($p = 0.16$) or moving at a slow pace ($p = 0.11$). Upon encounter, fish that were moving toward the robot transitioned to avoidance behavior in all instances. In most cases where the fish were moving slowly, they also switched to avoidance, with only a small fraction exhibiting freezing behavior ($p = 0.03$). After the robot becomes stationary in the post-encounter phase, the remaining slow or immobile fish also switch to avoidance. Only a minority (11%) of the fish that avoided the robot during the encounter remain stationary or nearly immobile in the post-encounter phase. Notably, none of the fish exhibit attraction behavior during or after encountering the robot, highlighting their aversive response following the robot's approach.

Fish moves away faster as the robot approaches faster and their relative distance decreases

We conducted a fixed-effects model to examine how the distance to the robot (d_{FR}), the robot's approach speed (v^R_{app}), and the distance to the nearest wall (d_{FW}) influence the fish's avoidance speed (v^F_{avoid}). The model is significant ($p < 0.001$), with coefficient estimates and their statistical significance shown in Table T.5.1. Both the distance and the approach speed significantly affect the fish's avoidance speed ($p < 0.001$). Specifically, a negative coefficient for d_{FR} suggests that as the distance to the robot decreases, the fish's avoidance speed increases. Conversely, the positive coefficient for v^R_{app} indicates that the fish retreats faster when the robot approaches at higher speeds. No significant effect of d_{FW} was found on v^F_{avoid}.

response at $t + \Delta$	predictor at t	coeff.	p-value	test statistics
fish avoidance speed v^F_{avoid}	robot-fish distance d_{FR}	-0.029	$p < 0.001$	$R^2 = 0.0386$
	robot approach speed v^R_{app}	0.066	$p < 0.001$	$Re = 0.199$
	wall-fish distance d_{FW}	-0.019	$p > 0.05$	$F_{(3,81954)} = 1095.45$
				$p < 0.001$

Table T.5.1: Results from the fixed-effects model analyzing the predictors of fish avoidance speed (v^F_{avoid}) at time $t + \Delta$ with $\Delta = 8$. The predictors include robot-fish distance (d_{FR}), robot approach speed (v^R_{app}), and wall-fish distance (d_{FW}). The coefficients, p-values, and test statistics are reported. The model, based on $N = 82,002$ input observations, is significant with $p < 0.001$. The measured repeatability ($Re = 0.199$) indicates that individual-level effects (i.e., variation due to differences between individual fish) explain only about 19.9% of the total variance in the data.

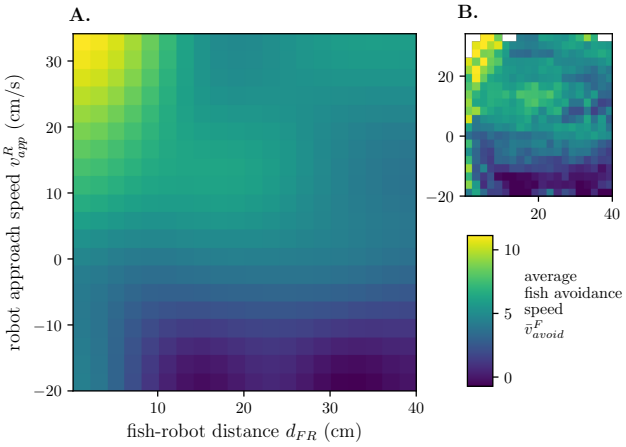

Figure 5.5: Smoothed **(A)** and unsmoothed **(B)** heatmaps showing the guppies' avoidance speed v^F_{avoid} at time $t + \Delta$, in function of the fish-robot distance (x-axis) and robot's approach speed (y-axis) at time t, with $\Delta = 8$. The heatmaps were each generated using $N = 69,265$ observations after applying range constraints to the x- and y-axes. Both axes are divided into 20 bins. To smooth the heatmap, missing data is first replaced with neighborhood averages, followed by Gaussian smoothing with $\sigma = 2$. The color scale indicates areas of high (yellow) and low (blue) speed. Missing data is represented in white in (B).

Figure 5.5A and B display the smoothed and unsmoothed heatmaps of v^F_{avoid}, averaged per bin, as a function of d_{FR} and v^R_{app}. The figure supports the results of the fixed-effects model, showing that v^F_{avoid} increases as d_{FR} decreases and v^R_{app} increases. Notably, even when the robot is further away (10 to 40 cm), the fish continues to move away from the robot ($v^F_{avoid} > 0$) if it is approaching ($v^R_{app} > 0$). Conversely, when the robot is not approaching ($v^R_{app} < 0$), the fish remains still ($v^F_{avoid} \approx 0$), possibly indicating a sense of safety in this scenario.

Reactive turns and accelerations as key mechanisms of fish avoidance from the robot

Avoidance behavior in fish can be measured by their speed v^F and their directional movement relative to the robot, as indicated by the relative angle θ_{RF}. The actions of the fish, such as accelerations and turns, influence these measures. Within the broader context of avoidance behavior, escape movements are characterized by extreme turns and accelerations that exceed baseline behavior, representing more pronounced responses to the robot's approach. Figure 5.6A and C show that the frequency of these extreme actions increases as the distance to the robot decreases, consistent

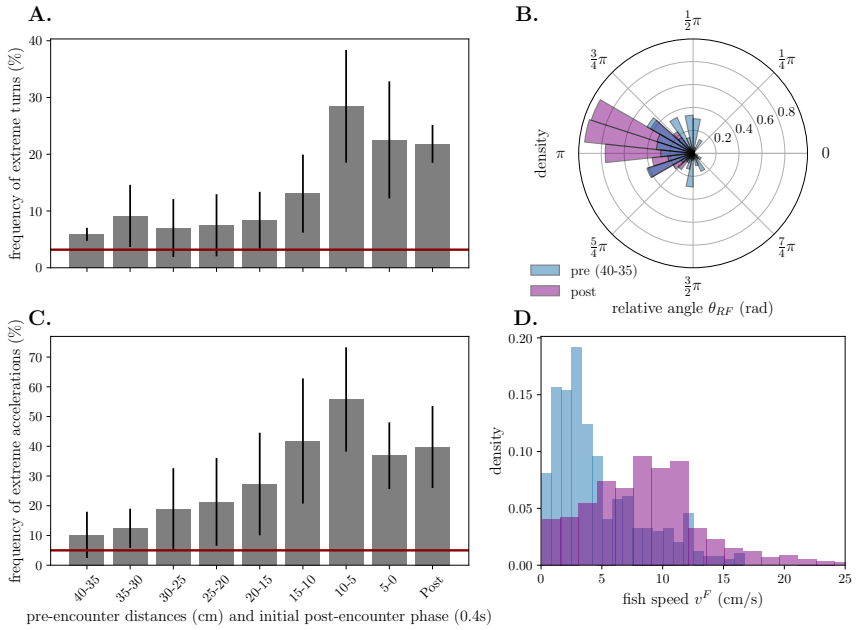

Figure 5.6: Frequencies of extreme turns **(A)** and accelerations **(C)** across fish-robot distances, calculated per distance bin during the pre-encounter phase and the first 0.4 seconds of the post-encounter phase. Extreme turns were defined as those exceeding the subject-specific threshold of $\mu_i + \sigma_i$, where μ_i and σ_i respectively represent the mean and standard deviation of the baseline turn distribution. Extreme accelerations were classified as those exceeding the 95th percentile of the baseline acceleration distribution for each subject. The bars represent the mean across all $N = 88$ trials and subjects, with error bars showing the standard deviation. Red horizontal lines indicate baseline frequencies for turns and accelerations, defined when the fish and robot are at least 40 cm apart. Density distributions of **(B)** the relative angle between the robot's position and fish's heading direction (θ_{RF}) and **(D)** the fish speed, sampled from the initial periods of the pre-encounter (blue) and post-encounter (purple) phases. When $\theta_{RF} = \pi$, the fish is moving directly away from the robot. The encounter with the robot increases both the fish's opposing movement and speed.

with the findings in Table T.5.1. The fish perform extreme turns and accelerations, classified based on deviations from the baseline behavior, most frequently within the 5–10 cm range from the robot.

A GLMM analysis confirms that the likelihood of turning actions increases significantly as the distance between the fish and robot decreases and as the fish-robot speed ratio ($\frac{v^F}{v^R}$) decreases (see Supplementary Table S5.D.1). This suggests that fish are most likely to perform extreme turns when moving significantly slower than the robot. In contrast, the relative

angle (θ_{RF}) does not have a statistically significant effect. One possible explanation is that the fish are likely to turn irrespective of where the robot is approaching from. This indicates that the primary purpose of the turn is to avoid immediate collision by moving out of the robot's path, rather than to align their movement directly opposite to the robot. Further supporting this explanation, the fish's turns do not appear to significantly affect their relative directional movement away from the robot. This is indicated by a separate FE analysis using the cosine of v^F_{avoid} (1 indicates direct retreat; -1 the opposite) as the response variable, which shows no significant difference before and after turns ($p > 0.1$). This interpretation is also reinforced by the sudden increase in the frequency of extreme turns at close proximity, compared to the more gradual rise in acceleration frequencies across distance bins (Fig. 5.6A, C). Additionally, pre-turn and post-turn speeds were measured as the maximum speed over a duration of 0.16 seconds, both before and after an extreme turn. An FE analysis reveals that post-turn speeds are significantly higher than pre-turn speeds (Estimate = 2.34, SE = 0.81, $t = 2.90$, $p < 0.01$), indicating that fish tend to accelerate following an extreme turn.

For the majority of trials, the fish do not perform extreme turns (Fig. 5.6A), suggesting that they often move at speeds comparable to or greater than the robot during its approach. The fish rely on accelerations, with the likelihood of these actions increasing as the robot approaches closer. Simultaneously, the fish make smaller orientation adjustments to optimize their directional movement, gradually aligning their heading to move more directly opposite to the robot (i.e., $\theta_{RF} \to \pi$). Figure 5.6B illustrates a high-density clustering of θ_{RF} values between $\frac{10}{12}\pi$ and π during the post-encounter phase. This indicates that the fish predominantly end up moving in the opposite direction of the robot after the encounter, with a slight preference for veering to the right relative to the robot's trajectory (binomial test; $p < 0.01$). During the early pre-encounter phase, the distribution of θ_{RF} is more dispersed but remains concentrated between $\frac{1}{2}\pi$ and $\frac{3}{2}\pi$, indicating that the robot is generally positioned behind the fish to some extent. The fish also exhibit increased speeds following the encounter (Fig. 5.6D), driven by their accelerative responses. An FE model confirms that post-encounter speeds are significantly higher than those observed in the early pre-encounter phase (Estimate = 3.80, SE = 0.75, $t = 5.06$, $p < 0.001$).

5.4 Discussion

Our findings reveal several key insights into the avoidance behavior of guppies in response to a conspecific-like robot. We demonstrate that avoidance responses are influenced by robot speed and proximity with faster approach speeds and closer distances eliciting stronger avoidance behaviors. The fish became more avoidant of the robot due to repeated expo-

sure to the approaches. Additionally, guppies displayed characteristic anti-predator responses, including freezing and escape maneuvers, suggesting the robotic conspecific was perceived as a threat. Interestingly, avoidance behavior was dynamic, as the fish employed turns and accelerations that depended on relative speed and distance, diverging from traditional models which assume fixed responses based only on distance.

Between conspecifics, fish tend to move away from each other at close distances [19, 41]. We found that repeated approaches of the robot increased the likelihood of the fish moving away before close proximity (Figure 5.4A), which also led to a reduced attraction to the robot. Notably, freezing was observed in only one of the five subjects. Conrad et al. [50] have discussed how individual personality traits, such as boldness or shyness, as well as body size [51], which is related to energy reserves and metabolic demands, can influence fish responses to perceived threats. Their concept of behavioral syndromes suggests that consistent behavioral patterns, including threat responses, are linked to individual variability among fish. This bidirectional relationship may explain why subject 4 showed a preference for freezing over avoidance as a defense strategy. Additionally, all instances where the fish initially froze upon encountering the robot eventually transitioned to movement after the encounter. Eilam [52] explains that the transition from freezing to movement, once the threat subsides, is a key component of dynamic defense strategies. This supports the notion that the robot's approach may be perceived as a threat, as freezing is recognized as an adaptive defense strategy.

Further analysis of the fish's behavioral patterns revealed a positive correlation between avoidance behavior and thigmotaxis. These findings are consistent with previous studies identifying thigmotaxis as a key indicator of avoidance behavior. Animals exhibiting thigmotaxic behavior tend to avoid the center of an arena, staying or moving close to the walls. This evolutionarily conserved behavior is observed across various species, including fish [53–55]. Thigmotaxis is widely recognized as a measure of anxiety, as it is reduced by anxiolytic drugs and heightened by anxiogenic agents [56]. Additionally, research has shown that guppies display thigmotaxic behavior in response to stressors, such as threats [26, 57].

The fish's avoidance speed was found to increase significantly with higher robot speed and decreased relative distance. These findings align with previous studies on fish-robot interactions in free-swimming contexts (without barriers), which suggest that higher robot speeds amplify avoidance responses [40], while reduced relative distance similarly intensifies these behaviors [42]. The fish consistently evaded the robot across most of the approach speed and distance parameters; however, avoidance behavior was absent when the robot was moving away at high speeds (-10 to -20 cm/s) and was more than 10 cm from the fish. Furthermore, proximity to the tank walls did not significantly affect avoidance speed, suggesting that the fish's avoidance behavior in terms of speed was not influenced by the

presence of the tank boundaries.

While most models treat the avoidance-initiation distance (the distance at which others can approach before the focal fish moves away) as a fixed radius or decay function, recent research suggests that individual factors, such as body size and prior exposure to threats, also influence this distance [34]. Our findings support this, as we observed that repeated exposure to the approaching robot increased the likelihood of the fish exhibiting avoidance behavior earlier, during the pre-encounter phase. Furthermore, our results indicated that the avoidance-initation distance is influenced by both the fish's speed and the speed of the approaching robot. When the fish moved slower than the robot (likely arising from within-subject dynamics; Supplementary Section S5.D), they were likely to execute a rapid escape maneuver when the robot approaches closer than 10 cm. These maneuvers, such as the C-start response, involve an extreme turn (characterized by sharp deviations from baseline turning behavior), which reorients the fish's trajectory away from the robot's path, followed by a powerful thrust (high acceleration) [58]. This strategy maximizes the distance from the robot with minimal energy expenditure. We observed a bimodal distribution in the angles of these escape trajectories (Supplementary Figure S5.D.1). Although we did not explicitly determine whether these escape trajectories are C-start responses, the observed bimodal distribution has previously been associated with C-starts [59].

However, these rapid escape maneuvers do not occur in all instances. In many cases, the fish were already moving away from the robot at significant speeds. Under such circumstances, they relied primarily on acceleration, initiating avoidance behaviors at distances even greater than 10 cm (Figures 5.5 and 5.6C). After the robot's approach, the fish tended to move almost directly opposite to its trajectory, aligning with traditional models of avoidance behavior [18–20]. Notably, a slight bias to the right was observed in their trajectories, consistent with findings reported by Domenici et al. [59]. Additionally, our findings revealed substantial variability in fish speeds post-encounter, both between and within subjects (Supplementary Section S5.E). These results emphasize the need for models of fish avoidance behavior to account for the high degree of speed variability, adding a layer of complexity to traditional avoidance frameworks that often assume uniform or constant behavioral responses.

This study is not without limitations. While we collected a substantial dataset (82,002 observations) to investigate fish avoidance behavior, the sample size of subjects was relatively small (five fish). As a result, generalizing these behavioral patterns to the broader population should be approached with caution. Further research with a larger sample size is necessary to better understand the individual traits that may contribute to the behavioral variation observed in our findings, particularly as one of the five subjects exhibited a distinctly different behavior profile. Moreover, although we identified the conditions under which a robot can elicit vary-

ing degrees of avoidance in individual fish, future research should explore how these interactions unfold in a school of fish, where collective dynamics could offer new insights.

Our experiment provides evidence that a mobile conspecific-like robot can induce avoidance behavior on live fish. Our findings generally indicate that the fish perceived this robot as a threat to be avoided, and even displayed anti-predator behavior, such as freezing and escape maneuvers [60]. Conspecific-like robots have primarily been used to attract, lead, or swim alongside fish. However, this study opens the door to investigating whether such robots could adapt their role dynamically, alternating between invoking different behaviors such as attraction after avoidance. This adaptability could greatly expand the potential applications of robots in influencing fish behavior. A potential extension of this idea are heterogeneous robot swarms, where some robots take on the role of attractive leaders while others act as repulsive aggressors. This approach could enhance control performance, as a single control strategy may be limited by the heterogeneity observed in animal collectives [61]. Furthermore, we have highlighted dynamic and variable factors in this study that differ from the traditional avoidance models, which would be of interest in to robot developers when simulating fish behavior. Understanding how fish-like robots influence a range of behaviors, including avoidance, is essential for advancing the development of mixed animal-robot societies [4, 6, 9]. Such systems hold the potential not only for enhancing our knowledge of animal behavior [62–64] but also for practical applications [13], such as guiding fish away from environmental hazards [16] or promoting more harmonious coexistence between animals and robots [65].

Ethics note

Experiments reported in this study were carried out in accordance with the recommendations of 'Guidelines for the treatment of animals in behavioural research and teaching' (published in Animal Behavior 1997) and comply with current German law approved by LaGeSo Berlin (G0117/16 to D B).

Bibliography

[1] S. Van Havermaet, A. Gerken, D. Mazrekaj, D. Bierbach, P. Simoens, T. Landgraf, and Y. Khaluf. The influence of a fish-like robot on the avoidance behavior of fish. Manuscript under review at Scientific Reports, 2025.

[2] J. Krause, A. F. Winfield, and J.-L. Deneubourg. Interactive robots in experimental biology. *Trends in ecology & evolution*, 26(7):369–375, 2011.

[3] E. Datteri. Interactive biorobotics. *Synthese*, 198(8):7577–7595, 2021.

[4] J. Halloy, F. Mondada, S. Kernbach, and T. Schmickl. Towards bio-hybrid systems made of social animals and robots. In *Biomimetic and Biohybrid Systems: Second International Conference, Living Machines 2013, London, UK, July 29–August 2, 2013. Proceedings 2*, pages 384–386. Springer, 2013.

[5] D. Romano, E. Donati, G. Benelli, and C. Stefanini. A review on animal–robot interaction: from bio-hybrid organisms to mixed societies. *Biological cybernetics*, 113:201–225, 2019.

[6] D. Romano, M. Porfiri, P. Zahadat, and T. Schmickl. Animal–robot interaction—an emerging field at the intersection of biology and robotics. *Bioinspiration & Biomimetics*, 19(2):020201, 2024.

[7] B. A. Klein, J. Stein, and R. C. Taylor. Robots in the service of animal behavior. *Communicative & integrative biology*, 5(5):466–472, 2012.

[8] B. Webb. What does robotics offer animal behaviour? *Animal behaviour*, 60(5):545–558, 2000.

[9] T. Landgraf, G. H. Gebhardt, D. Bierbach, P. Romanczuk, L. Musiolek, V. V. Hafner, and J. Krause. Animal-in-the-loop: Using interactive robotic conspecifics to study social behavior in animal groups. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:487–507, 2021.

[10] N. Horsevad, H. L. Kwa, and R. Bouffanais. Beyond bio-inspired robotics: how multi-robot systems can support research on collective animal behavior. *Frontiers in Robotics and AI*, 9:865414, 2022.

[11] M. Kondoyanni, D. Loukatos, C. Maraveas, C. Drosos, and K. G. Arvanitis. Bio-inspired robots and structures toward fostering the modernization of agriculture. *Biomimetics*, 7 (2):69, 2022.

[12] C. Cheng, J. Fu, H. Su, and L. Ren. Recent advancements in agriculture robots: Benefits and challenges. *Machines*, 11(1):48, 2023.

[13] M. Chellapurath, P. C. Khandelwal, and A. K. Schulz. Bioinspired robots can foster nature conservation. *Frontiers in Robotics and AI*, 10:1145798, 2023.

[14] K. Arts, R. Van der Wal, and W. M. Adams. Digital technology and the conservation of nature. *Ambio*, 44:661–673, 2015.

[15] L. Jacquin, Q. Petitjean, J. Côte, P. Laffaille, and S. Jean. Effects of pollution on fish behavior, personality, and cognition: some research perspectives. *Frontiers in Ecology and Evolution*, 8:86, 2020.

[16] S. Van Havermaet, P. Simoens, T. Landgraf, and Y. Khaluf. Steering herds away from dangers in dynamic environments. *Royal Society Open Science*, 10(5):230015, 2023.

[17] S. Van Havermaet, Y. Khaluf, and P. Simoens. Reactive shepherding along a dynamic path. *Scientific Reports*, 14(1):14915, 2024.

[18] S. Hubbard, P. Babak, S. T. Sigurdsson, and K. G. Magnússon. A model of the formation of fish schools and migrations of fish. *Ecological Modelling*, 174(4):359–374, 2004.

[19] Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, and I. D. Couzin. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725, 2011.

[20] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.

[21] T. Landgraf, D. Bierbach, H. Nguyen, N. Muggelberg, P. Romanczuk, and J. Krause. Robofish: increased acceptance of interactive robotic fish with realistic eyes and natural motion patterns by live trinidadian guppies. *Bioinspiration & biomimetics*, 11(1):015001, 2016.

[22] D. Bierbach, L. Gómez-Nava, F. A. Francisco, J. Lukas, L. Musiolek, V. V. Hafner, T. Landgraf, P. Romanczuk, and J. Krause. Live fish learn to anticipate the movement of a fish-like robot. *Bioinspiration & Biomimetics*, 17(6):065007, 2022.

[23] M. Lukas and T. R. de Jong. Conspecific interactions in adult laboratory rodents: friends or foes? *Social behavior from rodents to humans: neural foundations and clinical implications*, pages 3–24, 2017.

[24] J. F. Stephenson, S. E. Perkins, and J. Cable. Transmission risk predicts avoidance of infected conspecifics in trinidadian guppies. *Journal of Animal Ecology*, 87(6):1525–1533, 2018.

[25] P. Guevara-Fiore, J. Stapley, and P. J. Watt. Mating effort and female receptivity: how do male guppies decide when to invest in sex? *Behavioral Ecology and Sociobiology*, 64: 1665–1672, 2010.

[26] T. M. Houslay, R. L. Earley, S. J. White, W. Lammers, A. J. Grimmer, L. M. Travers, E. L. Johnson, A. J. Young, and A. Wilson. Genetic integration of behavioural and endocrine components of the stress response. *Elife*, 11:e67126, 2022.

[27] J. L. Kelley, J. P. Evans, I. W. Ramnarine, and A. E. Magurran. Back to school: can antipredator behaviour in guppies be enhanced through social learning? *Animal Behaviour*, 65(4):655–662, 2003.

[28] P. Domenici, J. M. Blagburn, and J. P. Bacon. Animal escapology i: theoretical issues and emerging trends in escape trajectories. *Journal of Experimental Biology*, 214(15):2463–2473, 2011.

[29] A. Raj and A. Thakur. Fish-inspired robots: design, sensing, actuation, and autonomy—a review of research. *Bioinspiration & biomimetics*, 11(3):031001, 2016.

[30] J. J. Faria, J. R. Dyer, R. O. Clément, I. D. Couzin, N. Holt, A. J. Ward, D. Waters, and J. Krause. A novel method for investigating the collective behaviour of fish: introducing 'robofish'. *Behavioral Ecology and Sociobiology*, 64:1211–1218, 2010.

[31] M. Maxeiner, M. Hocke, H. J. Moenck, G. H. Gebhardt, N. Weimar, L. Musiolek, J. Krause, D. Bierbach, and T. Landgraf. Social competence improves the performance of biomimetic robots leading live fish. *Bioinspiration & Biomimetics*, 18(4):045001, 2023.

[32] J. A. Fox, M. W. Toure, A. Heckley, R. Fan, S. M. Reader, and R. D. Barrett. Insights into adaptive behavioural plasticity from the guppy model system. *Proceedings of the Royal Society B*, 291(2018):20232625, 2024.

[33] P. Domenici. Context-dependent variability in the components of fish escape response: integrating locomotor performance and behavior. *Journal of Experimental Zoology Part A: Ecological Genetics and Physiology*, 313(2):59–79, 2010.

[34] D. A. Feary, A. M. Fowler, and D. J. Booth. Predator-avoidance behaviour of target and non-target temperate reef fishes is lower in areas protected from fishing. *Marine Biology*, 171(3):66, 2024.

[35] V. Cianca, T. Bartolini, M. Porfiri, and S. Macrì. A robotics-based behavioral paradigm to measure anxiety-related responses in zebrafish. *PloS one*, 8(7):e69661, 2013.

[36] F. Ladu, T. Bartolini, S. G. Panitz, F. Chiarotti, S. Butail, S. Macrì, and M. Porfiri. Live predators, robots, and computer-animated images elicit differential avoidance responses in zebrafish. *Zebrafish*, 12(3):205–214, 2015.

[37] G. Cord-Cruz, T. Ruberto, D. Neri, and M. Porfiri. Zebrafish response to live predator and biologically-inspired robot in a circular arena. In *Bioinspiration, Biomimetics, and Bioreplication 2017*, volume 10162, pages 120–125. SPIE, 2017.

[38] R. El Khoury, R. B. Ventura, G. Cord-Cruz, T. Ruberto, and M. Porfiri. Interactive experiments in a robotics-based platform to simulate zebrafish response to a predator. In *Bioinspiration, Biomimetics, and Bioreplication VIII*, volume 10593, pages 134–140. SPIE, 2018.

[39] C. Spinello, Y. Yang, S. Macrì, and M. Porfiri. Zebrafish adjust their behavior in response to an interactive robotic predator. *Frontiers in Robotics and AI*, 6:38, 2019.

[40] M. Kruusmaa, G. Rieucau, J. C. C. Montoya, R. Markna, and N. O. Handegard. Collective responses of a large mackerel school depend on the size and speed of a robotic fish but not on tail motion. *Bioinspiration & biomimetics*, 11(5):056020, 2016.

[41] J. E. Herbert-Read, A. Perna, R. P. Mann, T. M. Schaerf, D. J. Sumpter, and A. J. Ward. Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, 108(46):18726–18731, 2011.

[42] A. Pino, R. Vidal, E. Tormos, J. M. Cerdà-Reverter, R. Marín Prades, and P. J. Sanz. Towards fish welfare in the presence of robots: Zebrafish case. *Journal of Marine Science and Engineering*, 12(6):932, 2024.

[43] D. S. Calovi, A. Litchinko, V. Lecheval, U. Lopez, A. Pérez Escudero, H. Chaté, C. Sire, and G. Theraulaz. Disentangling and modeling interactions in fish with burst-and-coast swimming reveal distinct alignment and attraction behaviors. *PLoS computational biology*, 14(1):e1005933, 2018.

[44] A. D. Hartono, L. T. H. Nguyen, and T. V. Ta. A stochastic differential equation model for predator-avoidance fish schooling. *Mathematical Biosciences*, 367:109112, 2024.

[45] J. W. Jolles, N. Weimar, T. Landgraf, P. Romanczuk, J. Krause, and D. Bierbach. Group-level patterns emerge from individual speed as revealed by an extremely social robotic fish. *Biology letters*, 16(9):20200436, 2020.

[46] D. Bierbach, H. J. Mönck, J. Lukas, M. Habedank, P. Romanczuk, T. Landgraf, and J. Krause. Guppies prefer to follow large (robot) leaders irrespective of own size. *Frontiers in Bioengineering and Biotechnology*, 8:441, 2020.

[47] T. M. Houslay, M. Vierbuchen, A. J. Grimmer, A. J. Young, and A. J. Wilson. Testing the stability of behavioural coping style across stress contexts in the trinidadian guppy. *Functional Ecology*, 32(2):424–438, 2018.

[48] A. J. Brusseau, L. E. Feyten, A. L. Crane, I. W. Ramnarine, M. C. Ferrari, and G. E. Brown. Antipredator decisions of male trinidadian guppies (poecilia reticulata) depend on social cues from females. *Current Zoology*, page zoae040, 2024.

[49] A. L. Crane, L. E. Feyten, I. W. Ramnarine, and G. E. Brown. Temporally variable predation risk and fear retention in trinidadian guppies. *Behavioral Ecology*, 31(4):1084–1090, 2020.

[50] J. L. Conrad, K. L. Weinersmith, T. Brodin, J. Saltz, and A. Sih. Behavioural syndromes in fishes: a review with implications for ecology and fisheries management. *Journal of fish biology*, 78(2):395–435, 2011.

[51] C. Brown, F. Jones, and V. Braithwaite. In situ examination of boldness–shyness traits in the tropical poeciliid, brachyraphis episcopi. *Animal Behaviour*, 70(5):1003–1009, 2005.

[52] D. Eilam. Die hard: a blend of freezing and fleeing as a dynamic defense—implications for the control of defensive behavior. *Neuroscience & Biobehavioral Reviews*, 29(8):1181–1191, 2005.

[53] S. Sharma, S. Coombs, P. Patton, and T. B. De Perera. The function of wall-following behaviors in the mexican blind cavefish and a sighted relative, the mexican tetra (astyanax). *Journal of Comparative Physiology A*, 195:225–240, 2009.

[54] R. Blaser, L. Chadwick, and G. McGinnis. Behavioral measures of anxiety in zebrafish (danio rerio). *Behavioural brain research*, 208(1):56–62, 2010.

[55] S. Schnörr, P. Steenbergen, M. Richardson, and D. Champagne. Measuring thigmotaxis in larval zebrafish. *Behavioural brain research*, 228(2):367–374, 2012.

[56] L. Prut and C. Belzung. The open field as a paradigm to measure the effects of drugs on anxiety-like behaviors: a review. *European journal of pharmacology*, 463(1-3):3–33, 2003.

[57] G. E. Brown and J.-G. J. Godin. Chemical alarm signals in wild trinidadian guppies (poecilia reticulata). *Canadian Journal of Zoology*, 77(4):562–570, 1999.

[58] A. N. Peterson, A. P. Soto, and M. J. McHenry. Pursuit and evasion strategies in the predator–prey interactions of fishes. *Integrative and comparative biology*, 61(2):668–680, 2021.

[59] P. Domenici and R. W. Blake. The kinematics and performance of fish fast-start swimming. *Journal of Experimental Biology*, 200(8):1165–1178, 1997.

[60] R. Gerlai. Antipredatory behavior of zebrafish: adaptive function and a tool for translational research. *Evolutionary Psychology*, 11(3):147470491301100308, 2013.

[61] J. W. Jolles, A. J. King, and S. S. Killen. The role of individual heterogeneity in collective animal behaviour. *Trends in ecology & evolution*, 35(3):278–291, 2020.

[62] G. De Schutter, G. Theraulaz, and J.-L. Deneubourg. Animal–robots collective intelligence. *annals of mathematics and artificial intelligence*, 31:223–238, 2001.

[63] J. Abdai and A. Miklosi. *An Introduction to Ethorobotics: Robotics and the Study of Animal Behaviour*. Taylor & Francis, 2024.

[64] S. Mitri, S. Wischmann, D. Floreano, and L. Keller. Using robots to understand social behaviour. *Biological Reviews*, 88(1):31–39, 2013.

[65] F. Bonnet, R. Mills, M. Szopek, S. Schönwetter-Fuchs, J. Halloy, S. Bogdan, L. Correia, F. Mondada, and T. Schmickl. Robots mediating interactions between animals for interspecies collective behaviors. *Science Robotics*, 4(28):eaau7897, 2019.

[66] J. M. Butler and K. P. Maruska. The mechanosensory lateral line is used to assess opponents and mediate aggressive behaviors during territorial interactions in an african cichlid fish. *Journal of Experimental Biology*, 218(20):3284–3294, 2015.

[67] E. Scott, D. E. Edgley, A. Smith, D. A. Joyce, M. J. Genner, C. C. Ioannou, and S. Hauert. Lateral line morphology, sensory perception and collective behaviour in african cichlid fish. *Royal Society Open Science*, 10(1):221478, 2023.

Appendices

5.A Size of dataset after preprocessing

subject	original	selection	exclusion	corrections
0	16040	14500 (-1540)	12605 (-1895)	231 (1.83%)
1	19668	17530 (-2138)	16575 (-955)	307 (1.85%)
2	18629	16700 (-1929)	15455 (-1245)	256 (1.66%)
3	23045	20500 (-2545)	18961 (-1539)	149 (0.79%)
4	24547	20300 (-4247)	18406 (-1894)	351 (1.91%)
total	101929	89530 (-12399)	82002 (-7528)	1294 (1.57%)

Table 5.A.1: The number of frames per subject and in total, after each step of data processing. Original: the initial frame count for each video. Selection: the removal of frames associated with the initialization phase (milling process) and the post-experiment period (the delay between the end of the experiment and video termination). Exclusion: the number of frames discarded when 25 or more consecutive frames have no tracking data. Corrections: the number of frames where positions were interpolated, expressed as a percentage of the remaining frames after exclusion.

Table 5.A.1 shows the remaining number of frames at each step of the the data preprocessing pipeline. Frames were systematically removed based on several criteria to ensure the integrity of the analysis. First, frames corresponding to the initialization phase (milling process) and the post-experiment period were excluded, followed by the removal of frames where 25 or more consecutive frames lacked tracking data. By applying this exclusion criterion, we kept the number of frames requiring interpolation—referred to as corrections—relatively low. Corrections were applied to less than 2% of the remaining frames after exclusion. This was an important consideration, as interpolation, while necessary for handling small gaps in data, could alter key distributions (e.g. speed) if applied too extensively.

5.B Fixed-effects models of various response latencies

The fixed-effects (FE) models investigate how robot behavior affected the fish's avoidance speed v^F_{avoid}. The predictor variables included the distance d_{FR} between the robot and the fish, the robot's approach speed v^R_{app}, and the distance d_{FW} between the fish and the nearest tank wall. In the manuscript, we shifted the avoidance speed forward by $\Delta = 0.32s$ (8 frames), to take response latency into account. Table 5.B.1 shows that other values of $\Delta \in 4, 12, 16$ yield results consistent with the 8-frame shift.

response at $t+\Delta$	predictor at t	coeff.	signif. code	statistics
v^F_{avoid} ($\Delta = 4$)	fish-robot distance d_{RF}	-0.0235	$**$	$R^2 = 0.0252$
	robot approach speed v^R_{app}	0.0521	$***$	$Re = 0.1985$
	fish-wall distance d_{WF}	-0.0195		$F_{(3,81974)} = 705.48$
				$p < 0.001$
v^F_{avoid} ($\Delta = 8$)	fish-robot distance d_{RF}	-0.0294	$***$	$R^2 = 0.0386$
	robot approach v^R_{app}	0.0657	$***$	$Re = 0.1986$
	fish-wall distance d_{WF}	-0.0185		$F_{(3,81954)} = 1095.45$
				$p < 0.001$
v^F_{avoid} ($\Delta = 12$)	fish-robot distance d_{RF}	-0.0330	$***$	$R^2 = 0.0502$
	robot approach v^R_{app}	0.0762	$***$	$Re = 0.1988$
	fish-wall distance d_{WF}	-0.0177		$F_{(3,81934)} = 1442.04$
				$p < 0.001$
v^F_{avoid} ($\Delta = 16$)	fish-robot distance d_{RF}	-0.0351	$***$	$R^2 = 0.0591$
	robot approach v^R_{app}	0.0837	$***$	$Re = 0.1990$
	fish-wall distance d_{WF}	-0.0170		$F_{(3,81914)} = 1714.43$
				$p < 0.001$

Table 5.B.1: Results from 4 fixed effects models, each examining the relationship between fish avoidance speed (v^F_{avoid}) at future time points ($t + \Delta$) and predictor variables at time t, including fish-robot distance (d_{RF}), robot approach speed (v^R_{app}), and fish-wall distance (d_{WF}). Each model uses a different shift Δ to model the response latency of the fish to the robot's behavior. All models are significant ($p < 0.001$) and yield consistent results. Significance codes of the predictor terms: $p < 0.001$ ($***$); $p < 0.01$ ($**$); $p < 0.05$ ($*$); $p < 0.1$ (.); $p \geq 0.01$ ( )

predictor at $\bar{t}$	coeff.	p-value	test statistics
robot-fish distance d_{FR}	-0.033	$p < 0.001$	$R^2 = 0.0383$
robot approach speed v^R_{app}	0.057	$p < 0.001$	$Re = 0.278$
wall-fish distance d_{FW}	-0.018	$p > 0.05$	$F_{(3,12202)} = 162$
			$p < 0.001$

Table 5.C.1: Results from the two-way fixed-effects model analyzing the predictors of fish avoidance speed (v^F_{avoid}) at aggregated time $\bar{t} + \Delta$ with $\Delta = 2$, with individual- and time-specific fixed effects. The predictors include robot-fish distance (d_{FR}), robot approach speed (v^R_{app}), and wall-fish distance (d_{FW}). The coefficients, p-values, and test statistics are reported. The model is significant with $p < 0.001$.

5.C Two-way fixed-effects model on aggregated dataset

To include the time-specific fixed effect, the dataset was aggregated by subject and time. The frame number was first converted to seconds, using a sampling rate of 25 Hz. Afterwards, the data was aggregated and averaged within 5-second time windows. This process resulted in a dataset with one observation per 5-second interval per subject, allowing us to run a two-way fixed-effects model. The results presented in Table 5.C.1 align with

predictor	coeff.	SE	p-value	test statistics
relative angle θ_{RF} (fixed)	-0.2624	0.168	$p > 0.05$	pseudo $R^2 = 0.146$
speed ratio $\frac{v^F}{v^R}$ (fixed)	-5.3713	1.076	$p < 0.001$	$Re = 0.6879$
distance d_{RF} (fixed)	-0.1025	0.014	$p < 0.001$	$\chi^2(4) = 99.42$
subject identity (random)	-0.0642	0.095	$p > 0.05$	$p < 0.001$

Table 5.D.1: Results from the generalized linear mixed model (GLMM) with a logit link function, analyzing the occurence of an extreme turn as the binary response variable. The fixed effects are the relative angle θ_{RF}, the fish-robot speed ratio $\frac{v^F}{v^R}$, and the fish-robot distance d_{RF}. The subject identity is included as a random effect. Coefficients, standard errors (SE), and p-values are reported for the predictors. Test statistics of the model are also reported (right column). The model is statistically significant (p ¡ 0.001) based on the likelihood ratio test $\chi^2(2) = 99.42$. The repeatability estimate (Re = 0.688) indicates that 68.8% of the total variance in turning behavior is explained by individual-level differences between subjects.

the findings from the one-way (individual) fixed-effects model presented in the manuscript, demonstrating consistency across both approaches.

5.D Logistic regression analysis of extreme turns

To investigate the factors influencing the likelihood of a fish executing an extreme turn (defined as a turn exceeding its individual baseline), a Generalized Linear Mixed Model (GLMM) with a logit link function was applied. The binary response variable indicated whether an extreme turn occurred (1) or not (0) within each distance bin during the pre-encounter phase. Predictor variables included the relative angle between the fish and the robot (θ_{RF}), the fish-robot speed ratio ($\frac{v^F}{v^R}$), the distance between the fish and the robot (d_{RF}), and the subject identity as a random effect to account for inter-individual variability.

Table 5.D.1 presents the model results, showing statistical significance ($p < 0.001$) based on the likelihood ratio test ($\chi^2(2) = 99.42$). Among the predictors, the fish-robot speed ratio ($\frac{v^F}{v^R}$) exhibits a significant negative effect ($\beta = -5.371$, $SE = 1.076$, $p < 0.001$). This suggests that as the speed ratio increases, the likelihood of an extreme turn decreases, likely because a faster-moving fish is already distancing itself from the robot, making turning unnecessary. Similarly, the distance between the fish and the robot (d_{RF}) shows a significant negative effect ($\beta = -0.102$, $SE = 0.014$, $p < 0.001$), indicating that the likelihood of turning increases as the robot approaches closer proximity. This behavior aligns with the expectation that closer encounters with the robot heighten the need for evasive maneuvers. In contrast, the relative angle (θ_{RF}) does not have a statistically significant effect ($\beta = -0.262$, $SE = 0.168$, $p > 0.05$). This may be due to the limited variability in θ_{RF} across conditions, with most observations concentrated

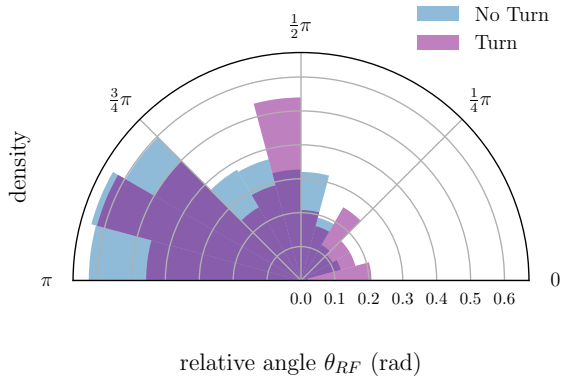

Supplementary Figure 5.D.1: Density distribution of the relative angle θ_{RF} for the instances of extreme turn (purple) and no turn (blue), sampled from all trials and subjects.

between $\frac{3}{4}\pi$ and π (Fig 5.D.1). Alternatively, it is possible that when fish are nearly immobile, turning behavior is less dependent on the relative angle since any directional movement suffices for evasion. Lastly, the subject identity random effect does not indicate significant differences in the likelihood of turning across individuals ($\beta = -0.064$, $SE = 0.095$, $p > 0.05$). The high repeatability ($Re = 0.688$) suggests that the variability appears to be more influenced by within-subject dynamics rather than subject-specific characteristics.

5.E Fish movement patterns in the post-encounter phase

Figure 5.E.1 illustrates the movement patterns of the fish during the post-encounter phase, after the robot finished approaching. The fish predominantly moved in the opposite direction of the robot, with a slight preference for veering to the left from the robot's perspective (highest density of θ_{RF} observed at between $\frac{10}{12}\pi$ and π, Fig 5.E.1A). The highest density of maximum speeds lies in the range of 10 to 15 cm/s. Analysis using a FE model showed that post-encounter speeds were significantly higher compared to the early pre-encounter phase (Estimate = 3.80, SE = 0.75, $t = -5.06$, $p < 0.001$), supporting the claim that fish exhibit high-speed avoidance behavior. Figure 5.E.1B shows the density distribution of fish speeds one second after the robot ceases its approach, fitted with a two-component Gaussian model. The first component represents relative slow

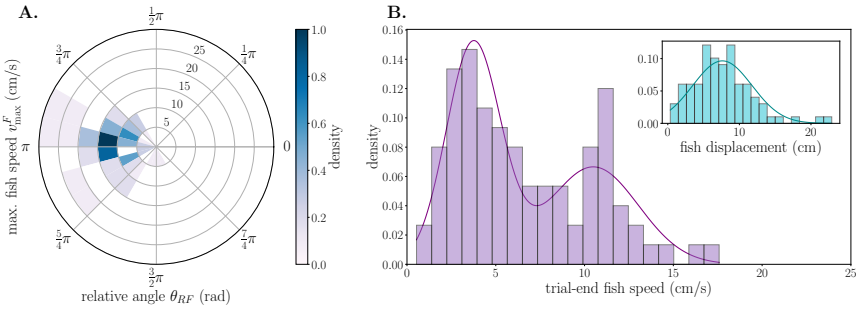

Supplementary Figure 5.E.1: (A) Density of the fish's maximum speed (v^F_{max}) and the relative angle (θ_{RF}) in the initial period of the post-encounter phase (0.4s). The radial gridlines represent relative angles in radians, with π indicating movement directly away from the robot. The concentric circles indicate levels of maximum speed with values increasing outward from the origin. Color intensity represents density, with darker colors (blue) indicating higher densities and lighter colors (purple, white) indicating lower densities. **(B)** Density distribution of the of the fish's speed 1s into the post-encounter phase. The inset shows the density distribution of fish's straight-line displacement from its position at the start of the post-encounter phase to its position 1s later. The solid lines represent Gaussian fits, with two components in the main plot and one component in the inset.

movements, with speeds around 3-4 cm/s, while the second component reflects higher-speed movements, centered around 11-12 cm/s. This suggests that while the majority of movements occur at a slow pace, a significant proportion of instances involve higher-speed motion. A linear mixed model (LMM) of these speeds, incorporating subject identifier as a random effect ($\beta = 7.019$, $SE = 1.345$, $p < 0.001$), indicates substantial variability both between subjects ($\sigma^2 = 8.547$) and within subjects ($\sigma^2 = 8.491$). The inset depicts fish displacement ($\mu = 7.62$, $\sigma = 4.14$), calculated as the change in robot-fish distance after 1 second.

5.F Per-subject analysis of fish displacement in the post-encounter phase

Figure 5.F.1 illustrates the relationship between mean fish speed and post-encounter displacement across individual fish. Notably, individuals with lower fish speeds also show lower displacements during the post-encounter phase. The displacement is calculated as the difference between the fish-robot distance at the start of the post-encounter phase and the distance 2 seconds afterwards. The Spearman test reveals a positive correlation ($r_s = 0.90$, $p < 0.05$, 95% CI $= [0.11, 1]$), indicating that less active fish tend to move shorter distances away from the robot. As the inset illustrates, the relationship between mean speed and displacement mirrors

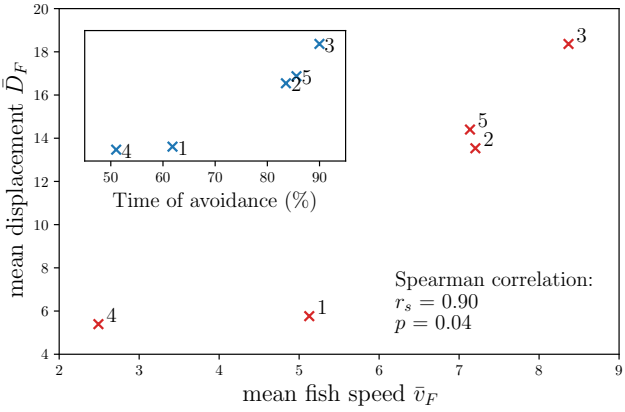

Supplementary Figure 5.F.1: For each individual fish, the mean displacement $\bar{D}_F$ in the first 2 seconds of the post-encounter phase is plotted against the mean speed $\bar{v}_F$. Each data is labeled with its respective subject identifier. The Spearman correlation coefficient is $r_s = 0.90$ with significance of $p = 0.04$ and a 95% confidence interval of $[0.11, 1]$. The inset shows the relationship between mean displacement and the time of avoidance.

the trend observed between the time of avoidance behavior and displacement. Subjects exhibiting less avoidance behavior during the experiment also tend to show shorter displacements.

5.G Relative angle between the robot's position and the fish's heading direction

Figure 5.G.1 shows the density distribution of the relative angle θ_{RF} using data collected across the entire experiment for all subjects, restricted to instances where the fish was moving at a speed greater than 2 cm/s. The results show that the vast majority of the density lies between $\frac{1}{2}\pi$ and $\frac{3}{2}\pi$, indicating that the robot was positioned behind the fish for most of the experiment. Notably, the regions with the highest density (0.4) are centered around π.

5.H Fish avoidance acceleration in function of the robot's relative position

Figure 5.H.1A and B show, respectively, the smoothed (Gaussian filter with $\sigma = 2$) and unsmoothed heatmaps of the guppies' avoidance acceleration (a^F_{avoid}) as a function of the relative position of the robot. The color scale visually represents areas of high (yellow) and low (blue) values, with miss-

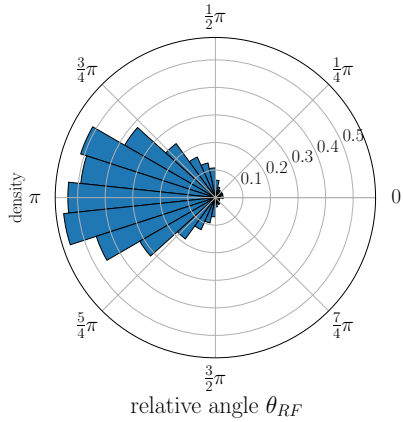

Supplementary Figure 5.G.1: Density distribution of the relative angle θ_{RF}, derived from data collected across the entire experiment for all subjects, where the fish is moving at a speed greater than 2cm/s.

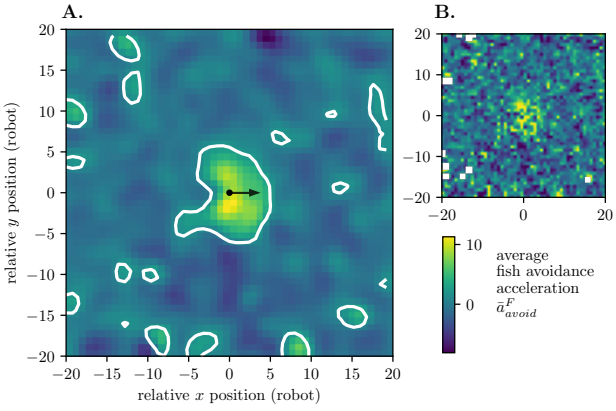

Supplementary Figure 5.H.1: Smoothed **(A)** and unsmoothed **(B)** heatmaps showing the guppies' avoidance acceleration a^F_{avoid} as a function of the robot's relative position, with the fish positioned at the origin, facing east (indicated by a black arrow). To smooth the heatmap, missing data is first replaced with neighborhood averages, followed by Gaussian smoothing with $\sigma = 1$. The color scale indicates areas of high (yellow) and low (blue) acceleration. Contour lines, with 3 levels, are added for visual clarity in (A). Missing data is represented in white in (B).

ing data shown in white. Contour lines in Figure A are provided for visual clarity.

The fish show the highest avoidance accelerations when the robot is located close the fish, particularly when positioned beside them (as indicated by the contour lines). Thus, the fish is not merely avoiding the robot upon encounter, rather, this suggests the fish is performing an escape maneuver. The fish perceives the robot as most threatening when it is present at these directions, prompting a rapid acceleration away from the robot. Additionally, there appears to be a clear pattern of reduced acceleration at greater distances, indicating that the fish only exhibit significant bursts of movement when the robot is in close proximity.

Escape movements, characterized by rapid, high-acceleration avoidance behavior, were most prominent when the robot was in close proximity. The distinct acceleration patterns observed laterally suggest heightened sensitivity in these regions, likely due to enhanced sensory detection and threat perception. This observation is consistent with previous studies [66, 67] showing that fish rely heavily on lateral-line systems to detect nearby objects or potential threats. Our findings suggest that fish are more likely to prioritize escape movements over freezing behaviors, especially when the robot approaches from the side at close range.

6 Conclusion

This chapter concludes the dissertation by critically reviewing the key research questions in light of the main findings, highlighting their contributions, limitations, and the open questions that remain. The chapter then broadens in scope to outline promising directions for future research beyond those research questions.

6.1 Revisiting the research questions

Throughout this dissertation, we explored how to develop robots that can assist in nature conservation tasks involving animal collectives, using fish as a case study. Swarm intelligence-based algorithms have been proposed for exploration and protection tasks, alongside experimental insights into how robots can influence fish behavior. In this section, we revisit the four main research questions posed (see Section 1.5), where each is critically reviewed by summarizing key findings, discussing the answers obtained, and highlighting limitations as well as areas for further research.

Research question 1. *How can robot teams explore and navigate like fish in dynamic environments?*

In Chapter 2, we proposed a novel collective motion model that adaptively generates different collective structures essential for exploration and navigation in dynamic environments. Three key structures were identified, each characterized by the group's spatial coverage and degree of alignment (order), corresponding to different functional roles in both nature conservation efforts and the ecological dynamics of fish exploration. We found

that conventional collective motion models, particularly behavioral zonal models, were unable to generate all three structures by varying the zone width parameters. We hypothesized that this limitation arises because these models typically assume that the weight factors of interaction forces are strictly positive, meaning that individuals are always influenced by all three interaction forces of repulsion, alignment, and attraction. To address this limitation, we proposed an adaptation of a metric-based behavioral zonal model and demonstrated that by allowing the alignment weight factor to be non-negative (i.e., including zero), all three structures could be generated solely through adjustments to the alignment force parameters. To evaluate its adaptability in dynamic environments, we demonstrated that a robot swarm following the proposed model can dynamically transition between these structures in response to stimuli that emerge over time and are detected by only a few individuals, enabling various exploration and navigation tasks as conditions change.

One of the three collective structures observed is characterized by low spatial coverage with high order, which may be particularly advantageous for maneuvering through narrow spaces. However, in Chapter 2, we did not consider environmental constraints such as obstacles. A swarm following the traditional zonal model may inherently adjust its spatial coverage to fit narrow spaces when such constraints are incorporated as an additional interaction type in the motion model [1, 2]. In fact, we demonstrated this concept in Chapter 4, where a fish school compressed its spatial structure under the external influence of fish-robot interactions to navigate a constrained trajectory. Nevertheless, fish schools may adopt this collective structure even in open spaces as a strategy to reduce predation risk [3]. In a future where robots integrate into fish schools, understanding how this structure naturally emerges from the fundamental three interaction forces remains valuable for both biological insight and effective swarm design.

The proposed model relies on the adaptation of the alignment rules in the behavioral zonal model, which has been extensively used to model the collective motion of fish [4–6]. As discussed in Section 1.3.1, the presence of alignment forces in the collective motion of fish has been a topic of debate. Some studies have found little evidence supporting alignment forces, suggesting either that no alignment zone exists [7, 8] or that it functions as a neutral zone [9]. However, these studies emphasize that this does not imply that alignment rules are never used by fish. Instead, they hypothesize that fish adapt their alignment rules in response to environmental stimuli or transmitted information (e.g., in the context of predation threats). Given our findings that changes in alignment rules lead to the emergence of different structures in terms of spatial coverage and order, an interesting direction for future research is to investigate the contexts that trigger these adaptations and how the resulting structural changes benefit fish in these different situations.

A comprehensive understanding of the various forms of collective mo-

tion is essential for establishing fish-robot mixed societies, where robotic swarms not only replicate fish movement patterns but actively integrate into schools, swimming alongside them. A potentially interesting direction for future research is the extent to which robots can influence the collective structure of these mixed societies. Several studies have already examined how specific robot characteristics affect collective properties. For instance, zebrafish group cohesion was found to decrease as the speed of a robotic fish increased [10], while in mosquitofish, cohesion was not influenced by the robot's visual features [11]. For robots to effectively shape collective structures, fish must adapt to the changes in the robots' alignment rules. Even if fish recognize and respond to these behavioral changes, successful emergence of collective structures would still depend on effective information propagation throughout the group [12–14]. This transfer of information is influenced by critical factors such as the ratio of informed robots to fish [15] and the spatial positioning of these robots within the group [16]. For example, individuals at the front of a fish school tend to exert a stronger influence on directional movement [17]. Consequently, future research may benefit from exploring the role of robots as leaders [18, 19] within these mixed societies.

Research question 2. *How can robot teams effectively protect fish schools from dynamic dangers?*

Focusing on another aspect of nature conservation, we explored how robot swarms could help protect fish schools from unrecognized threats, such as oil spills or invasive species. Specifically, we considered a scenario where dangers emerge dynamically, and robots, without prior knowledge of these threats, must reactively guide the fish away from the locally detected dangers. To address this, we proposed a swarm intelligence-based algorithm in which robots guide fish away from danger while maintaining a cage formation around the school. This structure allows the swarm to detect threats from any direction while ensuring that all fish are influenced to change direction when necessary. The approach presented in Chapter 3 is based on two simple decentralized rules: (i) each robot maintains a target distance from its nearest fish, and (ii) each robot positions itself equidistantly between two neighboring robots. This method depends on a certain level of school cohesion, to allow robots to move along the spatial contour of the school. Maintaining a cage formation also requires a minimum spacing between adjacent robots, meaning its success depends on both the number of robots deployed and the spatial coverage of the school. To assess this, we measured school cohesion, spatial coverage, and the minimum computationally determined number of robots required for different collective motion models of fish schooling. Our findings demonstrated that the proposed algorithm enables robots to effectively enclose the school, provided that cohesion constraints and swarm size requirements are met.

In a caging formation, robots steer fish by switching between two states based on their position relative to the desired direction. Robots behind the fish move closer, applying a repulsive force, while those ahead maintain distance, minimizing repulsion to facilitate movement. We demonstrated the effectiveness of this approach in two dynamic environments: (i) avoiding a stationary danger that appears stochastically, and (ii) remaining within a circular safe zone with a dynamically changing radius. The first scenario represents real-world threats such as pollution (e.g., fish near oil spills) or poaching, while the second mirrors virtual fencing applications, enabling controlled habitat management and animal care. We demonstrated that the robot swarm effectively ensured all fish avoided the dangers and remained within the safe zone. The proposed algorithm advances the state-of-the-art in shepherding by addressing three key challenges identified in the literature [20]: flexibility to diverse and dynamic environments, robustness against robotic agent failures, and adaptability to variations in animal responses to robot behavior. These features stem from the inherent advantages of swarm intelligence-based systems.

In the proposed algorithm, local communication is used only for the purpose of reaching consensus in scenarios where multiple dangers are present at distinct locations. However, transitioning this algorithm to real-world application requires a careful evaluation of its performance under the constraints and challenges of underwater communication. Specifically, underwater communication technologies are characterized by limited bandwidth and range, frequent disconnections, and sensitivity to environmental conditions such as temperature and salinity, all of which affect reliability [21]. Although our simulations did not explicitly constrain communication bandwidth, the messages used are relatively small, ranging from 16 to 24 bytes. The algorithm operates using an event-driven and asynchronous message-passing mechanism, which is generally compatible with the high-latency characteristics of underwater networks, such as acoustic communication. Although we did not simulate varying degrees of communication reliability, we examined fault tolerance in Chapter 4 by analyzing performance under repeated individual robot failures. Nonetheless, to transition from simulation to real-world deployment, future research should determine the operational limits of these algorithms depending on the available technology of bioinspired robots, including the depth limitations, communication ranges, and environmental conditions under which they remain effective. Recent research is also exploring novel underwater communication methods, such as electrocommunication systems inspired by weakly electric fish [22]. These systems have been shown to allow effective communication in a range of underwater environments, including still water, flowing water, water with obstacles and natural water conditions. This highlights their potential as alternatives to traditional communication techniques.

Our approach depended on a certain level of school cohesion. How-

ever, it should also be capable of protecting less cohesive groups. In practice, schools may initially exhibit lower cohesion, as observed in certain parameter ranges across all four collective motion models in Chapter 3, or may become less cohesive over time depending on the context, as previously discussed in relation to Research Question 1. Maintaining group cohesion is a common challenge in shepherding research [23, 24]. Traditional sheep herding models suggest that dogs alternate between collecting dispersed sheep and driving them when they are aggregated [25]. However, unlike fish, sheep exhibit less agile movement, and these models often overlook the animals' intrinsic motion. Lee and Kim [26] addressed the collection and steering of moving animals and, although they proposed a decentralized arc formation for steering, they still relied on a centralized approach to gather scattered groups. A possible decentralized solution could involve robots maintaining a chain formation to link subgroups, ensuring continuous information flow about each subgroup's location. However, fully decentralized collection remains an open question.

A potential approach to counteract the loss of school cohesion, and the resulting increase in spatial coverage, is for robots to regulate area coverage through the caging formation. As theoretically suggested [27] and empirically observed in predator-prey dynamics [28–30], animals are expected to remain within the cage, which could help maintain a desirable level of cohesion. However, in natural cases, prey often exhibit a freezing response as an acute stress response in such scenarios, and if a predator moves closer, the group can become highly disordered, reacting with rapid, uncoordinated movement [29]. This raises the question of whether animals, such as fish, can still be effectively guided when cohesion is externally imposed by robotic shepherds.

Research question 3. *How can robot teams effectively guide fish along a dynamic path?*

Continuing our study into how robot swarms can protect fish schools in dynamic environments, this research question introduces a more constrained problem setting than the previous one. In Chapter 3 (Research Question 2), multiple escape directions are often available for guiding fish away from local dangers, requiring consensus decision-making among the robot swarm to determine the best course of action. In Chapter 4 (Research Question 3), we focused on guiding fish along a dynamic trajectory without prior knowledge of the path. This scenario has real-world applications in the context of harmful environmental threats or ecologically sensitive areas. For instance, fish may need to be steered away from an oil spill that spreads unpredictably [31], or directed to safer waters during habitat restoration efforts [32]. Similarly, certain areas may need to be avoided for ecological reasons, such as preventing invasive species from entering fragile habitats or protected spawning grounds where their presence could

disrupt the ecosystem [33–35]. When local path information, such as direction and space between the boundaries, is represented using an artificial potential field, we found that robots could guide fish along a path using the algorithm proposed in Chapter 3 without the need for consensus decision-making. For real-life applications, this artificial potential field can be implemented to represent various environmental stimuli, such as water quality gradients indicating pollution levels or temperature variations affecting fish migration. Additionally, it could model physical barriers like dam openings or conservation zones.

The effectiveness of the proposed algorithm is measured by the percentage of the school that remains within the path boundaries as a function of the distance traveled. Optimal performance is achieved when the entire school stays within these boundaries from start to finish. Since path information is only available locally in real-time, the robot swarm must adaptively react, making optimal performance not always attainable. To assess the factors influencing performance, we examined three key aspects: the structural characteristics of the path (i.e., boundary spacing and turn sharpness), the system size (number of robots and fish), and the maximum velocities of both robots and fish. Our findings indicated that maintaining the group within the safe path was most challenging at sharp turns, especially in narrow sections. However, as boundary spacing increased, the severity of these disruptions decreased. Additionally, hybrid systems with a 2:3 robot-to-animal ratio significantly outperformed those with a 1:3 ratio, with further analysis showing that increasing the number of robots led to exponential performance improvements. Finally, when fish moved at relatively high velocities, the robots were unable to sustain the caging formation, resulting in failure to keep the group within the safe path.

The problem setting we considered is complex, as it involves a time-varying environment without prior information. Although prior data can quickly become obsolete in dynamic settings, some real-world scenarios involve trajectories that change gradually or infrequently. However, obtaining prior information remains challenging, as fish may cover vast and difficult-to-map areas, despite advances in underwater localization and mapping [36]. A potential solution to navigate challenging path structures is to deploy scout robots ahead of the school. These robots could relay trajectory information to the guiding swarm, effectively acting as copilots, similar to those in rally racing. In particular, a recent study used control barrier functions to shepherd in unknown and cluttered environments with limited trajectories, but relied on the assumption of centralized environment estimation and trajectory planning [37]. As such, whether fully decentralized strategies can optimally shepherd in time-varying, complex environments remains an open question.

The speed of the agents being shepherded is a critical [38] yet often overlooked [20] factor in shepherding problems. In our algorithm, robots maintain a set distance from the nearest fish to form a cage, factoring in the

speed of the fish to minimize prediction errors. Maintaining the appropriate distance is crucial, as unintended proximity can alter fish movement and lead to incorrect guidance. However, at higher fish speeds, maintaining cage formation becomes increasingly challenging with the same number of robots. In addition to challenges in maintaining formation, the robots themselves may also influence fish speed [10, 39]. A deeper understanding of how animal responses, including speed, are shaped by robot properties is essential for advancing shepherding toward real-world applications.

Research question 4. *How do the avoidance responses of fish to a conspecific-like robot differ from conventional models of animal-conspecific and animal-shepherd interactions?*

In Chapters 3 and 4 (Research Questions 2 and 3), we modeled fish-robot avoidance interactions using the conventional approach from shepherding literature. In these models, fish move at a constant speed in the opposite direction of a robot when the robot enters a predefined radius (see Subsection 1.3.2). This approach is also used to model avoidance behavior among live fish in schools but with significantly smaller predefined radii (see Subsection 1.3.1). However, the specific aspects of robot behavior that influence fish avoidance responses and how these responses deviate from conventional models remain largely unexplored. To investigate this, we programmed a robot to repeatedly approach a guppy at a fixed speed. The robot was designed to visually resemble a conspecific, aligning with our broader goal of advancing research in animal-robot mixed societies, where robots integrate into animal collectives (see Subsection 1.1.1). Each fish was subjected to multiple consecutive trials during an experiment, with the robot sampling a different speed within the medium to high range for each trial.

Our findings showed that repeated exposure to the robot's approach increased avoidance behavior in fish. Guppies also exhibited typical anti-predator responses, such as freezing and escape maneuvers, indicating that the robotic conspecific was perceived as a threat. We found that avoidance speed was influenced by the interaction between the robot's approach speed and its distance to the fish, with faster approaches and closer proximity eliciting higher avoidance speeds. Avoidance behavior was dynamic, as fish adjusted their turns and accelerations based on relative speed and distance, diverging from traditional models that assume fixed responses based solely on distance. Specifically, avoidance-initiation distance depended both on the speed of the fish and the speed of the robot approach. As the fish moved more slowly relative to the robot, they were more likely to initiate an escape maneuver (rapid acceleration and sharp directional changes [40]) at a shorter distance. Escape trajectories exhibited a bimodal distribution, clustering around 180 (opposite direction) and 90 (perpendicular) degrees. This pattern deviates from conventional avoidance models,

which typically assume that fish always move in the opposite direction from the other individual. Otherwise, fish tended to avoid the robot at greater distances, relying only on acceleration to swim in the opposite direction rather than executing an escape maneuver.

Recent research of predator-prey dynamics suggest that repeated exposure to a predator increases the avoidance-initiation distance [41]. Our findings align with this, as repeated robot approaches increased early avoidance, despite its conspecific-like appearance. This suggests that avoidance behavior adapts over time, though it remains unclear whether fish begin perceiving the robot as a threat or simply become more aversive. Since avoidance behavior also occurs among guppies in various contexts, such as mating [42] and disease transmission risk [43], further research is needed to clarify its underlying mechanisms. More specifically, a comparative study examining fish behavior in response to robots resembling either a conspecific or a predator could provide deeper insights.

Nevertheless, we found some evidence that the robot was perceived as a threat, as one subject exhibited freezing—a characteristic acute stress response in which the fish becomes motionless—before resuming movement, consistent with dynamic defense strategies [44]. Individual differences may also play a role, as personality traits and body size influence threat responses [45, 46], potentially explaining why some fish favored active avoidance over freezing. Additionally, fish exhibited escape maneuvers, another typical response to threats [47]. As mentioned before, we observed a bimodal distribution in escape trajectory angles, previously linked to C-start responses. However, as C-start responses are defined at the kinematic level [48], which we did not analyze, additional analysis is needed to confirm this and strengthen the hypothesis that the robot was perceived as a threat. Supporting this, we found a positive correlation between avoidance behavior and thigmotaxis, i.e. the tendency to stay near tank walls. While thigmotaxic behavior in guppies has been linked to stress responses, such as threats [49, 50], it is also observed in non-stressed natural shoals [51], warranting further investigation.

Finally, although we collected a large dataset to analyze fish avoidance behavior, the sample size of subjects was relatively small, limiting the generalization of our findings. A larger sample is needed to better understand individual behavioral variation, particularly as one subject exhibited a distinctly different response. Additionally, while we identified conditions influencing avoidance at the individual level, future research should examine these interactions within a school, where collective dynamics may reveal new insights.

6.2 Future perspectives

Below, we outline future research directions within the broader scope of this dissertation that we believe to be worthwhile of further investigation.

Other animal species. While this thesis focused on fish as a case study, we want to emphasize the broader need for the preservation of other animal species (see Subsection 1.2.1). The algorithms proposed in this dissertation (Chapters 2-4) may be applicable to other species, as they are based on a general framework of attraction, alignment, and repulsion forces (see Subsection 1.3.1). Although our simulation parameters were tailored to fish behavior, this framework has been used to model collective motion in birds, mammals, and even humans [52–55]. Thus, we believe that the proposed algorithms may be applicable, or at least serve as inspiration, for use cases of other animal species. Alternatively, research on shepherding in other species could inform future work on fish. Exciting developments are emerging, such as robotic falcons used to guide birds away from hazardous areas for their protection [56]. In particular, this robotic falcon has been shown to elicit similar collective avoidance responses in different bird species [57], raising the question of whether a robotic fish predator could similarly generalize avoidance behavior across different fish species. Studies on sheep offer further insight into group behavior in response to a shepherd. While individual sheep flee, large herds exhibit selfish herd behavior, whereas small groups transition unpredictably between the two [58]. As school cohesion in fish is also influenced by group size [59], this raises the possibility that comparable dynamics occur in fish schools and should be explored further.

Combining AI methods. In this thesis, we proposed swarm intelligence-based solutions using heuristic and model-based approaches that rely only on local observation and communication. As a result, the proposed swarm robotic systems benefit from key advantages such as flexibility, scalability, and robustness, which are essential features for nature conservation tasks (see Subsection 1.2.3). However, while these solutions provide a strong foundation, additional advancements are needed to fully bridge the gap to real-world applications. As discussed in Chapter 5, fish avoidance behavior in response to robots is influenced by multiple factors and evolves with repeated interactions. This underscores the need to integrate additional AI methods in future research. For instance, adaptive learning methods could improve real-time prediction models of avoidance behavior, allowing robots to adjust their strategies dynamically. Comparisons between AI approaches would also be valuable. In the same task of maintaining fish within a circular zone (introduced in Chapter 4), a shepherding solution has been proposed based on ordinary and partial differential equations [60], which resulted in an emergent caging formation. The authors

highlighted the similarity with our approach, despite not being explicitly designed for it. Recently, reinforcement learning (RL) has been increasingly explored for shepherding problems, offering improved execution time optimization compared to heuristic model-based rules while demonstrating scalability [61]. As RL continues to advance, hybrid approaches that integrate swarm intelligence with RL could present promising future solutions [62, 63].

Heterogeneity of the robot on the individual and collective level. In our proposed solutions for protecting fish schools (Chapters 3 and 4), we did not specify whether the robot resembled a conspecific or a predator, only assuming that it could trigger an avoidance response in fish. Traditionally, shepherding literature has employed robots that visually mimic a predator [56] or lack a specific visual design [64]. However, as discussed in relation to Research Question 4 (see Subsection 1.5), live conspecifics also exhibit avoidance behavior among themselves. Furthermore, studies using conspecific-like robots have demonstrated their ability to elicit avoidance responses in fish [65, 66]. Our findings further support this, as the experimental evidence in Chapter 5 suggests that the conspecific-like robot may even be perceived as a threat. This raises the question of whether a conspecific-like or predator-like robot is more effective for guiding fish. Building on this, an intriguing research direction is whether conspecific-like robots could dynamically adapt their roles, switching between evoking avoidance (perceived as a threat) and attraction (perceived as a leader) as needed. At the collective level, this concept could be extended to heterogeneous robot swarms, where some robots act as informed agents swimming alongside the fish (e.g., monitoring their behavior), while others serve as repulsive agents to guide them. This dual-role approach may improve control performance, particularly given the behavioral heterogeneity in animal collectives, which could limit the effectiveness of a single control strategy. However, guiding fish may not even require avoidance behavior. A conspecific-like robot assuming a leadership role could strongly influence the movement of its followers [17]. Hence, we find that investigating which interaction forces (avoidance or attraction) are most effective for controlling fish movement would be another interesting direction for future research.

Moving beyond simulation. The algorithms proposed in this dissertation have so far been demonstrated and analyzed only in simulation. While this represents a necessary first step, the ultimate goal is to enable their deployment in real-world scenarios. Among the necessary steps towards that goal ¡refine this beginning to avoid redundancy with previous sentence¿ is the specification of the robot technology to be used, and to consider the constraints, specifically in underwater environments, that these robots will

have to deal with. As discussed in Subsection 1.2.3, underwater environments pose unique challenges for sensing, actuator and communication technology. Nevertheless, research is showing promising avenues. In our algorithms (Chapters 2-4), the robots rely on information of the relative positions and orientations of nearby fish or other robots. Researches have designed artificial lateral lines, which are distributed sensing mechanisms inspired by how fish sense each others position in close range proximity [67–69]. Unlike cameras or optical sensors, lateral line-inspired systems are even effective in dark, murky, or cluttered environments. Moreover, they allow for fast and accurate control response. In the problem setting of Chapter 3, robots also have to be able to sense dangers. Depending on their nature, different sensors would have to be employed to the robot depending on the nature of the danger. For example, fluorometers to detect oil spills [70] or multispectral optical sensors for plastic debris [71]. Moreover, the specific choice of animal species plays a significant role in determining the system's design and conservation relevance. In this thesis, guppies were chosen due to their accessibility, prior use in fish-robot interaction studies, and availability of empirical data for simulation parameters. However, future research should also consider species that are more critically endangered or ecologically important. Adapting the proposed algorithms to species-specific behavioral dynamics and ecological requirements will be a key step toward real-world deployment.

Ecological impact. As we move toward a future in which robots can successfully integrate into animal collectives and influence their behavior, it is essential to consider the potential ecological risks of such interventions. While robots are increasingly employed to study and manipulate animal behavior [72–75], the long-term effects on animals' natural instincts remain insufficiently understood. If animals begin to rely on artificial or predictable cues from robotic agents, there is a risk that their natural responses may become dulled or maladaptive in real-world contexts. Furthermore, embedding robots into animal groups to mediate or control collective behaviors could potentially lead to unintended and irreversible changes in ecosystem functioning. To avoid compromising the complexity and adaptability that characterize healthy ecosystems, it is crucial to assess both the operational and environmental impacts of such technologies. Studies exploring how animals respond to prolonged exposure to robots are therefore an important and necessary direction for future research.

Towards animal-robot mixed societies for nature conservation. Achieving fully integrated animal-robot mixed societies requires further research across several key areas. While this thesis focused on developing robotic intelligence in such systems, research gaps remain in other areas [76]. First, enhancing locomotion capabilities is essential, as bioinspired designs mim-

icking fish can improve maneuverability and efficiency in underwater environments [77]. Second, addressing energy constraints is crucial; recent advancements in harvesting ocean thermal energy for autonomous underwater vehicles represent a step forward in renewable energy solutions for such robots [78]. Third, incorporating biodegradable materials in robot construction can reduce environmental impact and support conservation goals. For example, researchers have developed a robotic body composed primarily of silk hydrogel with embedded fibers to mimic the structure of natural fish skin [79]. Finally, understanding the visual and behavioral cues that influence animal perception is key. For instance, studies show that factors such as size [80] and coloration [81] significantly affect whether fish are attracted to robotic agents. Advancing these areas will be critical to successfully integrating robots into natural ecosystems for conservation efforts.

Final remarks

We may still be far from robots actively integrating with animals in natural environments. However, as the title of this thesis suggests, I hope to have taken a step toward that future. More importantly, I hoped to have raised awareness of the ongoing challenges in nature, particularly those affecting animals, and to highlight how swarms of bioinspired robots could serve as a tool to address some of these issues.

With this in mind, I would like to draw attention to the quote at the end of this chapter. By developing swarm intelligence inspired by animal behaviors, we have demonstrated that the strength of the robot is nature. Hopefully, the robot will also become the strength of nature.

"The strength of the pack is the wolf, and the strength of the wolf is the pack."

Rudyard Kipling, The Jungle Book (1894)

Bibliography

[1] H. Serna and W. T. Góźdź. The influence of obstacles on the collective motion of self-propelled objects. *Physica A: Statistical Mechanics and its Applications*, 625:129042, 2023.

[2] H. Zhao, H. Liu, Y.-W. Leung, and X. Chu. Self-adaptive collective motion of swarm robots. *IEEE Transactions on Automation Science and Engineering*, 15(4):1533–1545, 2018.

[3] C. K. Hemelrijk and H. Kunz. Density distribution and size sorting in fish schools: an individual-based model. *Behavioral Ecology*, 16(1):178–187, 2005.

[4] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.

[5] U. Lopez, J. Gautrais, I. D. Couzin, and G. Theraulaz. From behavioural analyses to models of collective motion in fish schools. *Interface focus*, 2(6):693–707, 2012.

[6] A. Huth and C. Wissel. The simulation of fish schools in comparison with experimental data. *Ecological modelling*, 75:135–146, 1994.

[7] Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, and I. D. Couzin. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725, 2011.

[8] J. E. Herbert-Read, A. Perna, R. P. Mann, T. M. Schaerf, D. J. Sumpter, and A. J. Ward. Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, 108(46):18726–18731, 2011.

[9] J. H. Tien, S. A. Levin, and D. I. Rubenstein. Dynamics of fish shoals: identifying key decision rules. *Evolutionary Ecology Research*, 6(4):555–565, 2004.

[10] S. Butail, T. Bartolini, and M. Porfiri. Collective response of zebrafish shoals to a free-swimming robotic fish. *PLoS One*, 8(10):e76123, 2013.

[11] G. Polverino and M. Porfiri. Mosquitofish (gambusia affinis) responds differentially to a robotic fish of varying swimming depth and aspect ratio. *Behavioural brain research*, 250: 133–138, 2013.

[12] A. Brabazon, W. Cui, and M. O'Neill. Information propagation in a social network: The case of a fish schooling algorithm. *Propagation Phenomena in Real World Networks*, pages 27–51, 2015.

[13] A. Strandburg-Peshkin, C. R. Twomey, N. W. Bode, A. B. Kao, Y. Katz, C. C. Ioannou, S. B. Rosenthal, C. J. Torney, H. S. Wu, S. A. Levin, et al. Visual sensory networks and effective information transfer in animal groups. *Current Biology*, 23(17):R709–R711, 2013.

[14] G. Rieucau, A. J. Holmin, J. C. Castillo, I. D. Couzin, and N. O. Handegard. School level structural and dynamic adjustments to risk promote information transfer and collective evasion in herring. *Animal Behaviour*, 117:69–78, 2016.

[15] A. J. Ward, D. J. Sumpter, I. D. Couzin, P. J. Hart, and J. Krause. Quorum decision-making facilitates information transfer in fish shoals. *Proceedings of the National Academy of Sciences*, 105(19):6948–6953, 2008.

[16] C. C. Ioannou, I. D. Couzin, R. James, D. P. Croft, and J. Krause. Social organisation and information transfer in schooling fish. *Fish cognition and behavior*, 2:217–239, 2011.

[17] J. Krause, D. Hoare, S. Krause, C. Hemelrijk, and D. Rubenstein. Leadership in fish shoals. *Fish and Fisheries*, 1(1):82–89, 2000.

[18] M. Maxeiner, M. Hocke, H. J. Moenck, G. H. Gebhardt, N. Weimar, L. Musiolek, J. Krause, D. Bierbach, and T. Landgraf. Social competence improves the performance of biomimetic robots leading live fish. *Bioinspiration & Biomimetics*, 18(4):045001, 2023.

[19] C. Wang, X. Chen, G. Xie, and M. Cao. Emergence of leadership in a robotic fish group under diverging individual personality traits. *Royal Society open science*, 4(5):161015, 2017.

[20] N. K. Long, K. Sammut, D. Sgarioto, M. Garratt, and H. A. Abbass. A comprehensive review of shepherding as a bio-inspired swarm-robotics guidance approach. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 4(4):523–537, 2020.

[21] Y. Zhang, S. Wang, M. K. Heinrich, X. Wang, and M. Dorigo. 3d hybrid formation control of an underwater robot swarm: Switching topologies, unmeasurable velocities, and system constraints. *ISA transactions*, 136:345–360, 2023.

[22] W. Wang, J. Liu, G. Xie, L. Wen, and J. Zhang. A bio-inspired electrocommunication system for small underwater robots. *Bioinspiration & biomimetics*, 12(3):036002, 2017.

[23] J.-M. Lien, O. B. Bayazit, R. T. Sowell, S. Rodriguez, and N. M. Amato. Shepherding behaviors. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 4, pages 4159–4164. IEEE, 2004.

[24] J.-M. Lien, S. Rodriguez, J.-P. Malric, and N. M. Amato. Shepherding behaviors with multiple shepherds. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3402–3407. IEEE, 2005.

[25] D. Strömbom, R. P. Mann, A. M. Wilson, S. Hailes, A. J. Morton, D. J. Sumpter, and A. J. King. Solving the shepherding problem: heuristics for herding autonomous, interacting agents. *Journal of the royal society interface*, 11(100):20140719, 2014.

[26] W. Lee and D. Kim. Autonomous shepherding behaviors of multiple target steering robots. *Sensors*, 17(12):2729, 2017.

[27] A. Varava, K. Hang, D. Kragic, and F. T. Pokorny. Herding by caging: a topological approach towards guiding moving agents via mobile robots. In *Robotics: Science and Systems*, pages 696–700, 2017.

[28] C. Muro, R. Escobedo, L. Spector, and R. Coppinger. Wolf-pack (canis lupus) hunting strategies emerge from simple rules in computational simulations. *Behavioural processes*, 88(3):192–197, 2011.

[29] R. J. Schmitt and S. W. Strand. Cooperative foraging by yellowtail, seriola lalandei (carangidae), on two species of fish prey. *Copeia*, 1982(3):714–717, 1982.

[30] R. Escobedo, C. Muro, L. Spector, and R. Coppinger. Group size, individual role differentiation and effectiveness of cooperation in a homogeneous group of hunters. *Journal of the Royal Society Interface*, 11(95):20140204, 2014.

[31] J. Michel and M. Fingas. Oil spills: Causes, consequences, prevention, and countermeasures. In *Fossil fuels: current status and future directions*, pages 159–201. World Scientific, 2016.

[32] J. Geist and S. J. Hawkins. Habitat recovery and restoration in aquatic ecosystems: current progress and future challenges. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 26(5):942–962, 2016.

[33] C. C. Koenig, F. C. Coleman, C. B. Grimes, G. R. Fitzhugh, K. M. Scanlon, C. T. Gledhill, and M. Grace. Protection of fish spawning habitat for the conservation of warm-temperate reef-fish fisheries of shelf-edge reefs of florida. *Bulletin of Marine Science*, 66 (3):593–616, 2000.

[34] J. S. Weis. Invasion and predation in aquatic ecosystems. *Current Zoology*, 57(5):613–624, 2011.

[35] A. R. Cooper, D. M. Infante, J. R. O'Hanley, H. Yu, T. M. Neeson, and K. J. Brumm. Prioritizing native migratory fish passage restoration while limiting the spread of invasive species: a case study in the upper mississippi river. *Science of the total environment*, 791: 148317, 2021.

[36] S. Wang, L. Chen, H. Hu, Z. Xue, and W. Pan. Underwater localization and environment mapping using wireless robots. *Wireless personal communications*, 70:1147–1170, 2013.

[37] M. Hamandi, F. Khorrami, and A. Tzes. Robotic shepherding in cluttered and unknown environments using control barrier functions. *arXiv preprint arXiv:2407.15701*, 2024.

[38] G. Y. Dosieah, A. Özdemir, M. Gauci, and R. Groß. Moving mixtures of active and passive elements with robots that do not compute. In *International Conference on Swarm Intelligence*, pages 183–195. Springer, 2022.

[39] P. P. Klamser, L. Gómez-Nava, T. Landgraf, J. W. Jolles, D. Bierbach, and P. Romanczuk. Impact of variable speed on collective movement of animal groups. *Frontiers in Physics*, 9:715996, 2021.

[40] P. Domenici, J. M. Blagburn, and J. P. Bacon. Animal escapology i: theoretical issues and emerging trends in escape trajectories. *Journal of Experimental Biology*, 214(15):2463–2473, 2011.

[41] D. A. Feary, A. M. Fowler, and D. J. Booth. Predator-avoidance behaviour of target and non-target temperate reef fishes is lower in areas protected from fishing. *Marine Biology*, 171(3):66, 2024.

[42] P. Guevara-Fiore, J. Stapley, and P. J. Watt. Mating effort and female receptivity: how do male guppies decide when to invest in sex? *Behavioral Ecology and Sociobiology*, 64: 1665–1672, 2010.

[43] J. F. Stephenson, S. E. Perkins, and J. Cable. Transmission risk predicts avoidance of infected conspecifics in trinidadian guppies. *Journal of Animal Ecology*, 87(6):1525–1533, 2018.

[44] D. Eilam. Die hard: a blend of freezing and fleeing as a dynamic defense—implications for the control of defensive behavior. *Neuroscience & Biobehavioral Reviews*, 29(8):1181–1191, 2005.

[45] J. L. Conrad, K. L. Weinersmith, T. Brodin, J. Saltz, and A. Sih. Behavioural syndromes in fishes: a review with implications for ecology and fisheries management. *Journal of fish biology*, 78(2):395–435, 2011.

[46] C. Brown, F. Jones, and V. Braithwaite. In situ examination of boldness–shyness traits in the tropical poeciliid, brachyraphis episcopi. *Animal Behaviour*, 70(5):1003–1009, 2005.

[47] A. N. Peterson, A. P. Soto, and M. J. McHenry. Pursuit and evasion strategies in the predator–prey interactions of fishes. *Integrative and comparative biology*, 61(2):668–680, 2021.

[48] P. Domenici and R. W. Blake. The kinematics and performance of fish fast-start swimming. *Journal of Experimental Biology*, 200(8):1165–1178, 1997.

[49] G. E. Brown and J.-G. J. Godin. Chemical alarm signals in wild trinidadian guppies (poecilia reticulata). *Canadian Journal of Zoology*, 77(4):562–570, 1999.

[50] T. M. Houslay, R. L. Earley, S. J. White, W. Lammers, A. J. Grimmer, L. M. Travers, E. L. Johnson, A. J. Young, and A. Wilson. Genetic integration of behavioural and endocrine components of the stress response. *Elife*, 11:e67126, 2022.

[51] T. K. Kleinhappel, T. W. Pike, and O. H. Burman. Stress-induced changes in group behaviour. *Scientific Reports*, 9(1):17200, 2019.

[52] A. Deutsch, G. Theraulaz, and T. Vicsek. Collective motion in biological systems, 2012.

[53] A. Okubo. Dynamical aspects of animal grouping: swarms, schools, flocks, and herds. *Advances in biophysics*, 22:1–94, 1986.

[54] T. Vicsek and A. Zafeiris. Collective motion. *Physics reports*, 517(3-4):71–140, 2012.

[55] W. H. Warren. Collective motion in human crowds. *Current directions in psychological science*, 27(4):232–240, 2018.

[56] R. F. Storms, C. Carere, R. Musters, H. Van Gasteren, S. Verhulst, and C. K. Hemelrijk. Deterrence of birds with an artificial predator, the robotfalcon. *Journal of the Royal Society Interface*, 19(195):20220497, 2022.

[57] R. F. Storms, C. Carere, R. Musters, R. Hulst, S. Verhulst, and C. K. Hemelrijk. A robotic falcon induces similar collective escape responses in different bird species. *Journal of the Royal Society Interface*, 21(214):20230737, 2024.

[58] T. Chakrabortty and S. Bhamla. Controlling noisy herds. *ArXiv*, pages arXiv–2406, 2024.

[59] D. S. Shelton, B. C. Price, K. M. Ocasio, and E. P. Martins. Density and group size influence shoal cohesion, but not coordination in zebrafish (danio rerio). *Journal of Comparative Psychology*, 129(1):72, 2015.

[60] B. Di Lorenzo, G. C. Maffettone, and M. di Bernardo. A continuification-based control solution for large-scale shepherding. *arXiv preprint arXiv:2411.04791*, 2024.

[61] I. Napolitano, A. Lama, F. De Lellis, and M. di Bernardo. Emergent cooperative strategies for multi-agent shepherding via reinforcement learning. *arXiv preprint arXiv:2411.05454*, 2024.

[62] M.-A. Blais and M. A. Akhloufi. Reinforcement learning for swarm robotics: An overview of applications, algorithms and simulators. *Cognitive Robotics*, 3:226–256, 2023.

[63] M. Hüttenrauch, A. Šošić, and G. Neumann. Deep reinforcement learning for swarm systems. *Journal of Machine Learning Research*, 20(54):1–31, 2019.

[64] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot sheepdog project achieves automatic flock control. In *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour*, volume 489, page 493, 1998.

[65] M. Kruusmaa, G. Rieucau, J. C. C. Montoya, R. Markna, and N. O. Handegard. Collective responses of a large mackerel school depend on the size and speed of a robotic fish but not on tail motion. *Bioinspiration & biomimetics*, 11(5):056020, 2016.

[66] A. Pino, R. Vidal, E. Tormos, J. M. Cerdà-Reverter, R. Marín Prades, and P. J. Sanz. Towards fish welfare in the presence of robots: Zebrafish case. *Journal of Marine Science and Engineering*, 12(6):932, 2024.

[67] G. Liu, A. Wang, X. Wang, and P. Liu. A review of artificial lateral line in sensor fabrication and bionic applications for robot fish. *Applied bionics and biomechanics*, 2016(1): 4732703, 2016.

[68] X. Zheng, C. Wang, R. Fan, and G. Xie. Artificial lateral line based local sensing between two adjacent robotic fish. *Bioinspiration & biomimetics*, 13(1):016002, 2017.

[69] L. DeVries, F. D. Lagor, H. Lei, X. Tan, and D. A. Paley. Distributed flow estimation and closed-loop control of an underwater vehicle with a multi-modal artificial lateral line. *Bioinspiration & biomimetics*, 10(2):025002, 2015.

[70] A. Vasilijevic, N. Stilinovic, D. Nad, F. Mandic, N. Miskovic, and Z. Vukic. Auv based mobile fluorometers: System for underwater oil-spill detection and quantification. In *2015 IEEE Sensors Applications Symposium (SAS)*, pages 1–6. IEEE, 2015.

[71] B. Basu, S. Sannigrahi, A. Sarkar Basu, and F. Pilla. Development of novel classification algorithms for detection of floating plastic debris in coastal waterbodies using multispectral sentinel-2 remote sensing imagery. *Remote Sensing*, 13(8):1598, 2021.

[72] N. Afzal, M. ur Rehman, L. Seneviratne, and I. Hussain. The convergence of ai and animal-inspired robots for ecological conservation. *Ecological Informatics*, page 102950, 2024.

[73] G. Polverino, M. Karakaya, C. Spinello, V. R. Soman, and M. Porfiri. Behavioural and life-history responses of mosquitofish to biologically inspired and interactive robotic predators. *Journal of the Royal Society Interface*, 16(158):20190359, 2019.

[74] G. Polverino and M. Porfiri. Controlling invasive species with biologically-inspired robots. In *ALIFE 2021: The 2021 Conference on Artificial Life*. MIT Press, 2021.

[75] F. Bonnet, R. Mills, M. Szopek, S. Schönwetter-Fuchs, J. Halloy, S. Bogdan, L. Correia, F. Mondada, and T. Schmickl. Robots mediating interactions between animals for interspecies collective behaviors. *Science Robotics*, 4(28):eaau7897, 2019.

[76] M. Chellapurath, P. C. Khandelwal, and A. K. Schulz. Bioinspired robots can foster nature conservation. *Frontiers in Robotics and AI*, 10:1145798, 2023.

[77] X. Jian and T. Zou. A review of locomotion, control, and implementation of robot fish. *Journal of Intelligent & Robotic Systems*, 106(2):37, 2022.

[78] I. R. Mohamed, A. S. Shehata, W. M. El-Maghlany, and M. A. Kotb. Experiment study on harvesting ocean thermal energy using phase change material for autonomous underwater vehicle powering. In *AIP Conference Proceedings*, volume 2769. AIP Publishing, 2023.

[79] C. M. Donatelli, S. A. Bradner, J. Mathews, E. Sanders, C. Culligan, D. Kaplan, and E. D. Tytell. Prototype of a fish inspired swimming silk robot. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 60–65. IEEE, 2018.

[80] T. Bartolini, V. Mwaffo, A. Showler, S. Macrì, S. Butail, and M. Porfiri. Zebrafish response to 3d printed shoals of conspecifics: the effect of body size. *Bioinspiration & biomimetics*, 11(2):026003, 2016.

[81] D. Romano and C. Stefanini. Any colour you like: fish interacting with bioinspired robots unravel mechanisms promoting mixed phenotype aggregations. *Bioinspiration & Biomimetics*, 17(4):045004, 2022.